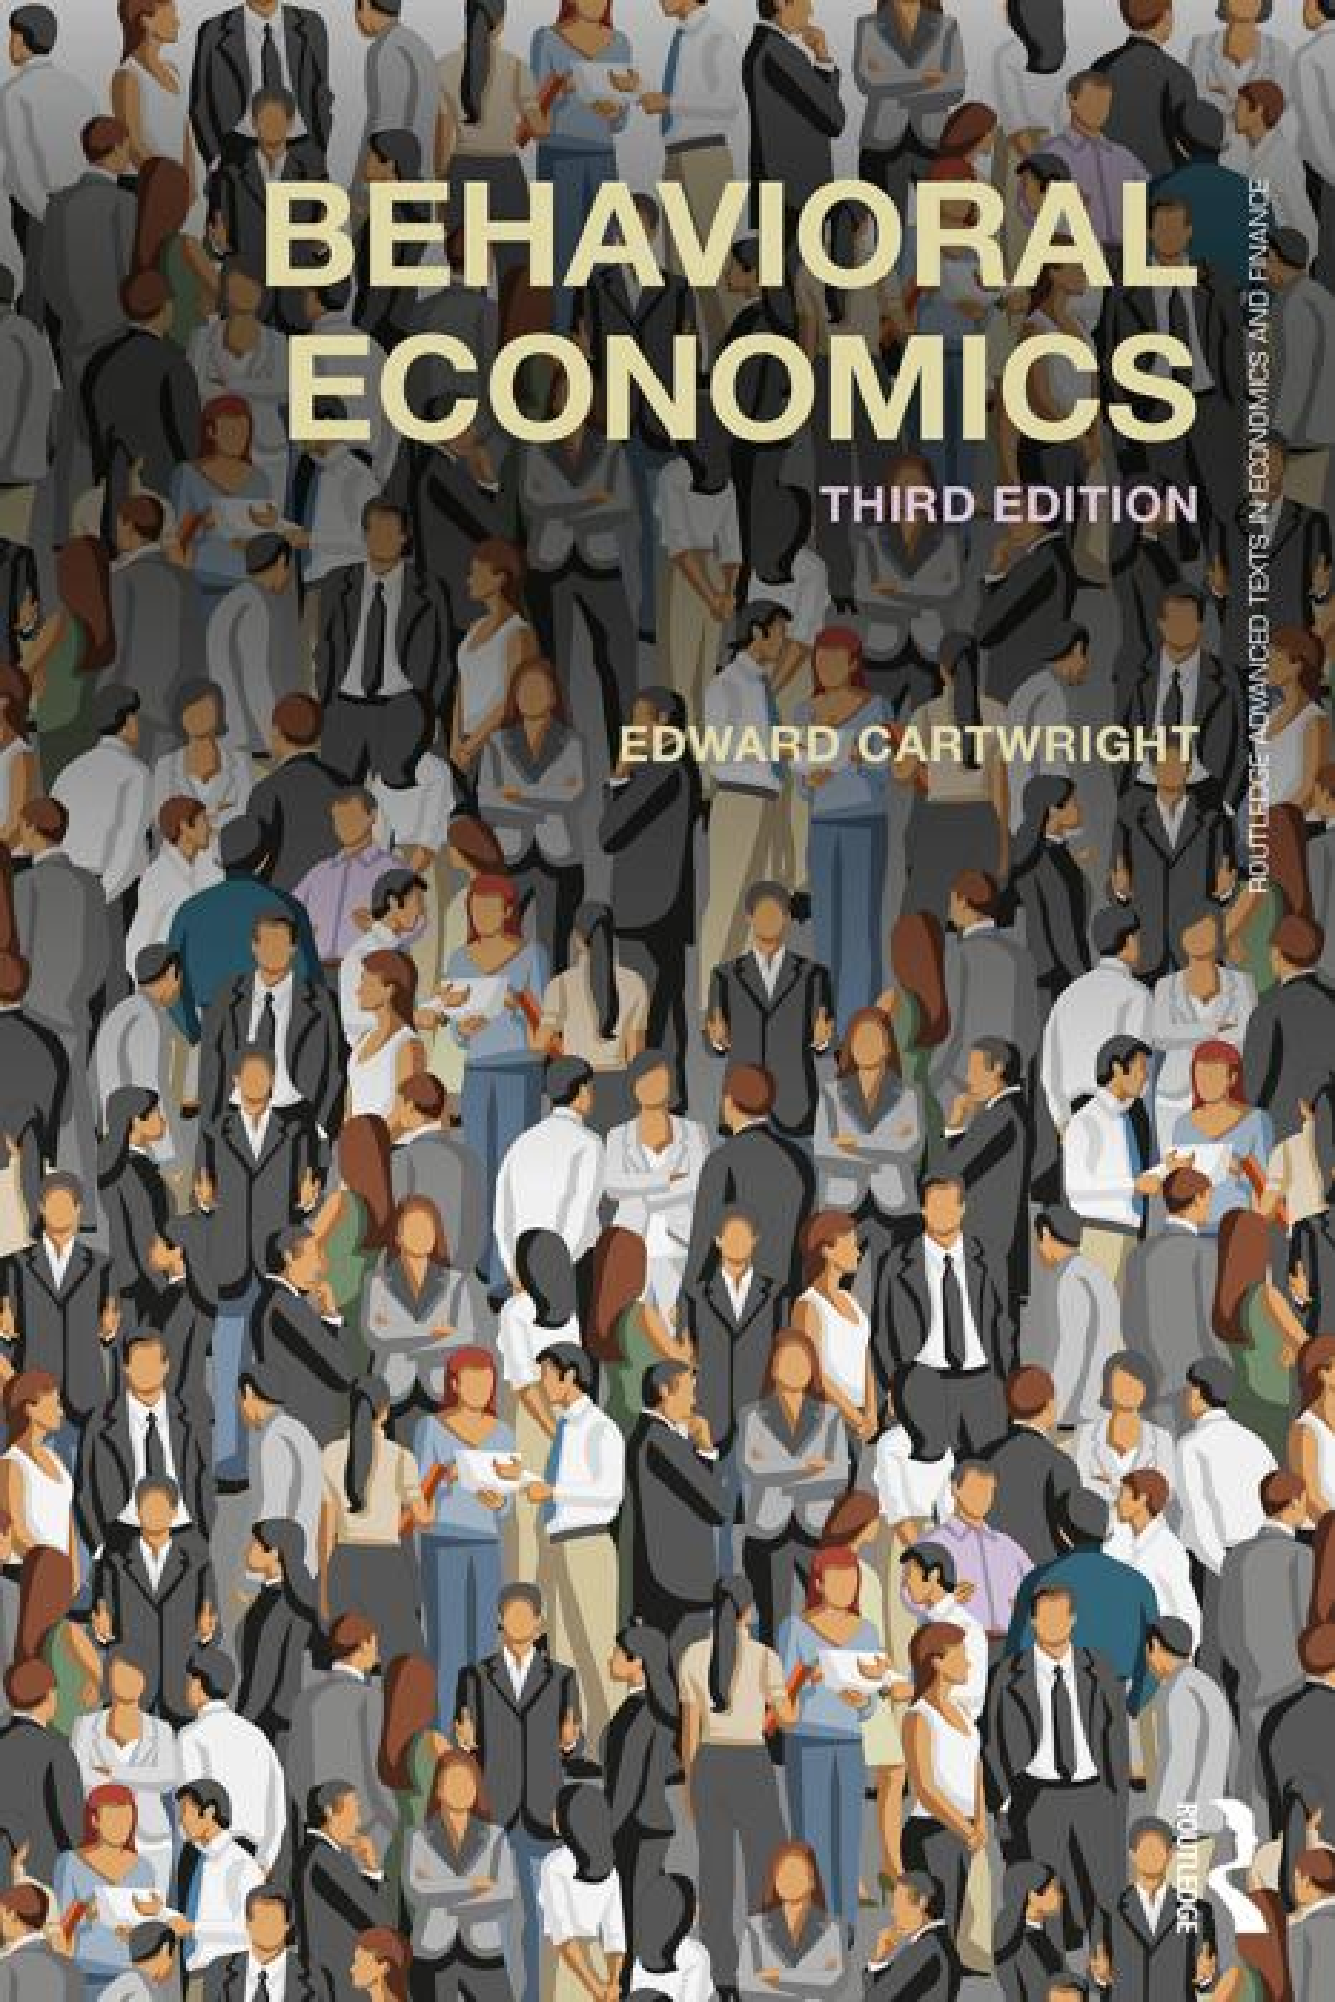
BEHAVIORAL ECONOMICS
THIRD EDITION
EDWARD CARTWRIGHT
ROUTLEDGE ADVANCED TEXTS IN ECONOMICS AND FINANCE
ROUTLEDGE

Behavioral Economics

Over the last few decades behavioral economics has revolutionized the discipline. It has done so by putting the human back into economics, by recognizing that people sometimes make mistakes, care about others and are generally not as cold and calculating as economists have traditionally assumed. The results have been exciting and fascinating, and have fundamentally changed the way we look at economic behavior.

This textbook introduces all the key results and insights of behavioral economics to a student audience. Ideas such as mental accounting, prospect theory, present bias, inequality aversion and learning are explained in detail. These ideas are also applied in diverse settings such as auctions, stock market crashes, charitable donations and health care, to show why behavioral economics is crucial to understanding the world around us. Consideration is also given to what makes people happy, and how we can potentially nudge people to be happier.

This new edition contains expanded and updated coverage of contract theory, bargaining in the family, time and risk, and stochastic reference points, among other topics, to ensure that readers are kept up to speed with this fast-paced field. The companion website is also updated with a range of new questions and worked examples. This book remains the ideal introduction to behavioral economics for advanced undergraduate and graduate students.

Edward Cartwright is a Reader in Economics at the University of Kent, Canterbury, UK.

Routledge Advanced Texts in Economics and Finance

12. Culture and Economics
    On Values, Economics and International Business
    Eelke de Jong
13. Modern Public Economics, Second Edition
    Raghbendra Jha
14. Introduction to Estimating Economic Models
    Atsushi Maki
15. Advanced Econometric Theory
    John Chipman
16. Behavioral Economics
    Edward Cartwright
17. Essentials of Advanced Macroeconomic Theory
    Ola Olsson
18. Behavioral Economics and Finance
    Michelle Baddeley
19. Applied Health Economics, Second Edition
    Andrew M. Jones, Nigel Rice, Teresa Bago d'Uva and Silvia Balia
20. Real Estate Economics
    A Point to Point Handbook
    Nicholas G. Pirounakis
21. Finance in Asia
    Institutions, Regulation and Policy
    Qiao Liu, Paul Lejot and Douglas Arner
22. Behavioral Economics, Second Edition
    Edward Cartwright
23. Understanding Financial Risk Management
    Angelo Corelli
24. Empirical Development Economics
    Måns Söderbom and Francis Teal with Markus Eberhardt, Simon Quinn and Andrew Zeitlin
25. Strategic Entrepreneurial Finance
    From Value Creation to Realization
    Darek Klonowski
26. Computational Economics
    A Concise Introduction
    Oscar Afonso and Paulo Vasconcelos
27. Regional Economics, Second Edition
    Roberta Capello
28. Game Theory and Exercises
    Gisèle Umbhauer
29. Innovation and Technology
    Business and Economics Approaches
    Nikos Vernardakis
30. Behavioral Economics, Third Edition
    Edward Cartwright

Behavioral Economics

Third Edition

Edward Cartwright

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

Third edition published 2018
by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

and by Routledge
711 Third Avenue, New York, NY 10017

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2018 Edward Cartwright

The right of Edward Cartwright to be identified as author of this work has been asserted by him in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

First edition published by Routledge 2011
Second edition published by Routledge 2014

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data
Names: Cartwright, Edward, author.
Title: Behavioral economics / Edward Cartwright.
Description: 3rd Edition. | New York : Routledge, 2018. |
Revised edition of the author's Behavioral economics, 2014. |
Includes bibliographical references.
Identifiers: LCCN 2017037171 | ISBN 9781138097117 (hardback) |
ISBN 9781138097124 (pbk.) | ISBN 9781315105079 (ebook)
Subjects: LCSH: Economics–Psychological aspects.
Classification: LCC HB74.P8 C37 2018 | DDC 330.01/9–dc23
LC record available at https://lccn.loc.gov/2017037171

ISBN: 978-1-138-09711-7 (hbk)
ISBN: 978-1-138-09712-4 (pbk)
ISBN: 978-1-315-10507-9 (ebk)

Typeset in Gill Sans
by Out of House Publishing

Visit the companion website: www.routledge.com/cw/cartwright

Dedicated to Anna, Darren,
Sofie, Kerry and Robert



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Contents

Detailed contents ix
List of figures xvi
List of tables xxiv
List of research methods boxes xxviii
Preface xxx

Part I Introduction 1
1 An introduction to behavioral economics 3

Part II Economic behavior 29
2 Simple heuristics for complex choices 31
3 Choice with risk 97
4 Choosing when to act 162
5 Learning from new information 206
6 Interacting with others 250
7 Social preferences 321

Part III Origins of behavior 371
8 Evolution and culture 373
9 Neuroeconomics 415

Part IV Welfare and policy 463
10 Happiness and utility 465
11 Policy and behavior 504

Bibliography 530
Index 550



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Detailed contents

List of figures xvi
List of tables xxiv
List of research methods boxes xxviii
Preface xxx

Part I Introduction 1
1 An introduction to behavioral economics 3
1.1 The history and controversies of behavioral economics 4
1.1.1 Behavioral economics is reborn 6
1.1.2 Behavioral economics and policy 9
1.1.3 The different faces of behavioral economics 11
1.1.4 Debate and controversy 13
1.1.5 Too far or not far enough 16
1.2 Some background on behavioral economics methods 18
1.2.1 Some background on experiments 18
1.2.2 Some background on theory 22
1.3 How to use this book 24
1.3.1 Chapter previews 25
1.3.2 Behavioral finance 26
1.4 Further reading 26
1.5 Online material 27
1.6 Review questions 27

Part II Economic behavior 29
2 Simple heuristics for complex choices 31
2.1 Utility and search 31
2.1.1 How to search 33
2.1.2 Choice arbitrariness 38
2.2 Mental accounting and framing 45
2.2.1 Reference-dependent utility 46
2.2.2 The endowment effect 47
2.2.3 Willingness to pay or accept 49
2.2.4 Transaction utility 51
2.2.5 Narrow framing 52
2.2.6 Hedonic editing 55
2.2.7 Choice bracketing 55
2.3 The role of emotions 57
2.3.1 Aversion to lying 58
2.3.2 Deception 59

2.3.3 Honesty and framing 63
2.4 Summary 66
2.5 Demand, supply and markets 67
2.5.1 Double-auction markets in the lab 69
2.5.2 Posted offer markets and market power 72
2.5.3 The law of one price 73
2.6 Labor supply and reference dependence 75
2.6.1 A target income and target wage 76
2.6.2 Taxicab drivers 77
2.6.3 Female labor supply 79
2.7 The housing market 81
2.7.1 Reluctance to sell 82
2.7.2 What are buyers willing to pay? 83
2.8 The behavioral life cycle hypothesis 84
2.8.1 Fungibility and mental accounting 86
2.9 Saving for the future 91
2.9.1 Let's diversify 92
2.9.2 Let's not diversify 93
2.10 Further reading 95
2.11 Online material 96
2.12 Review questions 96
3 Choice with risk 97
3.1 Expected utility 97
3.1.1 The Allais paradox 102
3.1.2 Risk aversion 103
3.1.3 Risk-loving for losses 106
3.1.4 When expected utility will work 107
3.2 Independent and fanning out 108
3.2.1 Disappointment 110
3.2.2 Rank-dependent expected utility 112
3.3 Reference dependence and prospect theory 115
3.3.1 Reference-dependent utility 118
3.3.2 The reference point and expectations 119
3.3.3 Combined gambles 122
3.3.4 Stochastic reference point 123
3.4 Preference reversals 126
3.4.1 Procedural invariance 130
3.4.2 Regret theory 131
3.4.3 Prospect theory and preference reversals 133
3.4.4 Why preference reversals matter 135
3.5 Summary 135
3.6 Financial trading 136
3.6.1 The equity premium puzzle 136
3.6.2 The disposition effect 139
3.6.3 The ostrich effect 142
3.7 Insurance 144
3.8 Tax evasion 146
3.8.1 Standard model of tax evasion 148

3.8.2 Behavioral theories of tax evasion 150
3.8.3 Taxes and reference points 152
3.8.4 Tax evasion in the laboratory 152
3.9 Legal settlements 155
3.9.1 Fourfold pattern of risk attitudes 156
3.9.2 Frivolous litigation 158
3.10 Further reading 160
3.11 Online materials 160
3.12 Review questions 160
4 Choosing when to act 162
4.1 Exponential discounting 162
4.1.1 The discount factor 164
4.1.2 The utility of sequences 168
4.2 Hyperbolic discounting 170
4.2.1 Quasi-hyperbolic discounting 172
4.2.2 The consequences of time inconsistency 174
4.2.3 Temptation and self-control 176
4.3 Loss aversion and sequences 179
4.3.1 Reference dependence 179
4.3.2 Preferences for sequences 182
4.4 Time and risk 184
4.5 Summary 189
4.6 Borrowing and saving 190
4.6.1 Saving equals growth or growth equals saving? 191
4.6.2 Why save when you have debts? 192
4.7 Exploiting time inconsistency 193
4.7.1 Time inconsistency and consumer behavior 193
4.7.2 Firm pricing 194
4.7.3 Choosing the correct calling plan 197
4.8 Environmental economics 198
4.8.1 Inter-generational discount factor 199
4.8.2 Reducing CO_2 emissions 201
4.9 Further reading 203
4.10 Online materials 204
4.11 Review questions 204
5 Learning from new information 206
5.1 Bayesian updating and choice with uncertainty 206
5.1.1 Models of choice with uncertainty 208
5.1.2 The Ellsberg paradox 210
5.2 Two cognitive biases 212
5.2.1 Confirmatory bias 212
5.2.2 A model of confirmatory bias 214
5.2.3 Law of small numbers 216
5.2.4 A model of the law of small numbers 218
5.2.5 Generating random sequences 220
5.2.6 Do biases matter? 221
5.3 Learning from others 222
5.3.1 To conform or not 222

5.3.2 Cascade experiments 223
5.3.3 What happened to conformity? 226
5.3.4 Signaling games 228
5.4 Summary 231
5.5 Health care 231
5.5.1 Patients 231
5.5.2 Practitioners 234
5.6 Bubble and bust 236
5.6.1 Bubbles in the lab 236
5.6.2 Experience and bubbles 240
5.6.3 Explaining bubbles 244
5.7 Voting in elections 245
5.8 Further reading 247
5.9 Online materials 248
5.10 Review questions 248
6 Interacting with others 250
6.1 The beauty contest 250
6.1.1 Strategy and Nash equilibrium 251
6.1.2 Choice in a beauty contest 253
6.1.3 Learning in a beauty contest 255
6.2 Playing for the first time 256
6.2.1 Level-k thinking 256
6.2.2 Sophisticated beliefs 258
6.2.3 Focal points 260
6.2.4 Equilibrium refinement 265
6.2.5 Nash equilibrium with mistakes 269
6.3 Learning from experience 272
6.3.1 Reinforcement learning 273
6.3.2 Belief-based learning 274
6.3.3 Experience-weighted learning 278
6.3.4 Learning and prediction 279
6.4 Teams make decisions 283
6.4.1 Teams and the beauty contest 285
6.4.2 The sophistication of teams 288
6.4.3 Are teams smarter? 289
6.5 Summary 290
6.6 Auctions 290
6.6.1 Revenue equivalence 291
6.6.2 Winner's curse 293
6.7 Learning to coordinate 296
6.7.1 Weakest link games 297
6.7.2 Threshold public good games 303
6.7.3 Coordinating on networks 305
6.8 Monetary policy by committee 309
6.9 Industrial organization 310
6.9.1 Limit pricing 311
6.9.2 Market entry 314

6.9.3 Quantity leadership 316
6.10 Further reading 319
6.11 Online materials 319
6.12 Review questions 319
7 Social preferences 321
7.1 The experimental evidence for social preferences 321
7.1.1 The nice side of social preferences 322
7.1.2 The nasty side of social preferences 327
7.1.3 Reciprocity 330
7.1.4 Fairness and competition 334
7.1.5 The terminology of reciprocity 335
7.1.6 Social preferences and teams 335
7.2 Inequality aversion 338
7.2.1 Inequality aversion with incomplete information 338
7.2.2 Inequality aversion with complete information 341
7.2.3 An evaluation of inequality aversion models 345
7.3 Intentions and social norms 346
7.3.1 A model of fairness based on intentions 347
7.3.2 What is fair? 349
7.4 Summary 353
7.5 Giving to charity 354
7.5.1 Crowding out 354
7.5.2 Who is watching? 356
7.5.3 Why do people give? 358
7.6 Price and wage rigidity 358
7.6.1 A model of worker reciprocity 360
7.6.2 Wage stickiness in the lab 362
7.6.3 How long to forget a wage change? 363
7.6.4 Firm pricing 363
7.7 Contract theory 365
7.7.1 Contracts for loss-averse workers 366
7.7.2 Exploitation of an overconfident worker 368
7.8 Further reading 369
7.9 Online materials 369
7.10 Review questions 369

Part III Origins of behavior 371
8 Evolution and culture 373
8.1 Evolution and economic behavior 373
8.1.1 Looking for food and finding a utility function 374
8.1.2 Choosing when to have children 376
8.1.3 Aggregate risk 377
8.1.4 Competing with others 380
8.2 Culture and multi-level selection 385
8.2.1 Cross-cultural comparisons 385
8.2.2 Group selection 389
8.2.3 Gene–culture coevolution 392

8.2.4 Reciprocity in children and chimpanzees 393
8.3 Summary 396
8.4 The gender gap 397
8.4.1 Attitudes to risk 398
8.4.2 Attitudes to competition 399
8.4.3 Social preferences 401
8.4.4 Why are men and women different? 402
8.5 The economics of family 403
8.6 Development economics 407
8.6.1 The education production function 407
8.6.2 Microfinance 410
8.7 Further reading 413
8.8 Online materials 413
8.9 Review questions 414
9 Neuroeconomics 415
9.1 An introduction to the brain 415
9.1.1 An economist's map of the brain 418
9.1.2 Brain processes 421
9.1.3 Executive control systems 422
9.2 Valuing rewards and learning 424
9.2.1 Reward evaluation 424
9.2.2 Learning about rewards 427
9.2.3 Risk and uncertainty 430
9.2.4 Different types of reward 432
9.3 Making decisions 433
9.3.1 Choice and strategy 434
9.3.2 Framing effects 437
9.3.3 Strategic behavior 438
9.3.4 Fairness and norms 440
9.3.5 Punishment and inequality aversion 443
9.3.6 Present bias and a brain in conflict 445
9.3.7 Multiple-self models 449
9.4 Summary 450
9.5 Addiction 451
9.5.1 A model of rational addiction 452
9.5.2 Biases and addiction 454
9.5.3 Cues and addiction 457
9.5.4 Addiction and neuroscience 459
9.6 Further reading 460
9.7 Online materials 461
9.8 Review questions 461

PART IV Welfare and policy **463**
10 Happiness and utility 465
10.1 What makes us happy? 465
10.1.1 Happiness is relative 468
10.1.2 Adaption and habituation 471

10.2 Do we know what makes us happy? 473
10.2.1 Remembered utility 474
10.2.2 Projection bias 477
10.3 Choice and commitment 480
10.3.1 Does present bias matter? 481
10.3.2 Pre-commitment 483
10.3.3 Do people like having choice? 484
10.4 Summary 486
10.5 Health and happiness 487
10.5.1 Measuring the value of treatment 487
10.5.2 Improving the remembered utility of treatment 490
10.6 Saving and retirement 492
10.6.1 Projection bias in saving 492
10.6.2 Investor autonomy 494
10.7 Welfare trade-offs 495
10.7.1 The inflation–unemployment trade-off 495
10.7.2 Tax saliency 497
10.8 Further reading 501
10.9 Online materials 502
10.10 Review questions 502
11 Policy and behavior 504
11.1 Designing good institutions 504
11.1.1 The tragedy of the commons 505
11.1.2 Matching markets 508
11.1.3 Spectrum auctions 512
11.1.4 Behavioral economics and institution design 516
11.2 Nudge and behavior change 517
11.2.1 Savings accounts 518
11.2.2 A default to save 519
11.2.3 Nudge 521
11.2.4 Nudge and behavior change 524
11.2.5 Consumer protection, health and the environment 525
11.3 Summary 528
11.4 Further reading 528
11.5 Online materials 529
11.6 Review questions 529

Bibliography 530
Index 550

Figures

1.1	An overview of the project	21
2.1	How well search heuristics fit the observed search and choices in a search experiment	36
2.2	The decision that subjects made when the choice is conflicting or non-conflicting, compared to when there is no choice	39
2.3	Whether subjects prefer an extra coupon to $10 depends on whether they were exposed to a background stage when coupons cost $15 or $5	40
2.4	The proportion choosing a mid-level Sony radio-cassette player when the alternative was an Emerson player, a top-quality Sony player or both the Emerson and Sony	41
2.5	The proportion who chose spot A and B to book and to cancel	42
2.6	The amount that subjects were willing to pay for a box of chocolates was higher for those with a higher social security number	43
2.7	A representation of different brain processes	45
2.8	Whether subjects chose chocolate depended on the endowment	48
2.9	The valuation put on ten chocolates	49
2.10	How subjects valued a bottle of wine that cost $20 and is now worth $75 depending on whether they drank it, gave it as a gift or broke it	56
2.11	How subjects felt about a task depending on how long they did it after getting paid	57
2.12	A matrix for the adding to 10 task	58
2.13	Number of matrices that subjects reported to have solved	59
2.14	The proportion of subjects who chose option B or lied that option B will earn more money	60
2.15	Proportion of senders who were liars and sophisticated truth tellers distinguished into individuals and groups	62
2.16	The observed distribution of earnings from ten coin tosses compared to what we would expect with pure chance comparing those with no priming or priming about banking	65
2.17	The demand and supply curves before and after a shift in demand	68
2.18	Convergence of market price and quantity to equilibrium in four double-auction markets	70
2.19	The equilibrium price and average price traded in a market that lasted 15 periods	70
2.20	The deviation from equilibrium price and quantity in experimental double-auction and posted offer markets	72
2.21	The prices traded in a negotiated price market	74
2.22	The law of one price does not hold in many markets	75

2.23	Two possible utility functions	77
2.24	If there is a reference level of earning $160, the utility function becomes S-shaped	77
2.25	The probability that a driver stops his shift compared to when he has worked nine hours and earned $150 to $175	78
2.26	The participation rate of women aged 16 and above in the United States and United Kingdom	80
2.27	The change in probability of a woman being in employment depends on whether her sister-in-law works and her sister's husband earns more or not than her husband	81
2.28	The average selling price of residential property and sales volume in England and Wales	82
2.29	The estimated value of two properties depending on the reported list price	84
2.30	An illustration of consumption smoothing	85
2.31	An inverse U-shaped consumption and income pattern derived from UK couples with a husband born between 1936 and 1943	85
2.32	Consumption patterns over time derived from UK data	86
2.33	The amount of the $2,400 that subjects thought they would spend in the next year	87
2.34	The change in spending in the three months after receiving the rebate compared to the three months before receiving the rebate	88
2.35	The marginal propensity to consume the windfall gain compared to normal income of Israeli recipients of German restitution payments	89
2.36	The marginal propensity to consume clothing out of normal income and child benefit	90
2.37	The contribution rate of people enrolled in a large defined contribution plan	91
2.38	The proportion of funds that were based on equities and the proportion of savings that savers invested in equity	92
2.39	The proportion of investments in each fund of a company's contribution savings plan	93
2.40	The number of funds that savers invested in and were on offer in a large sample of defined contribution plans	94
2.41	The proportion of new participants whose behavior was consistent with the conditional l/n heuristic	95
3.1	Three possible utility functions: one is concave, which would imply Alan is risk-averse, and one convex, which would imply Alan is risk-loving	99
3.2	What can we infer about the utility function of someone who prefers prospect I to prospect K for increasing levels of wealth?	105
3.3	A probability triangle diagram of the Allais paradox	108
3.4	Indifference curves in a probability triangle diagram	109
3.5	The Allais paradox	110
3.6	The Allais paradox can be explained by the fanning out of indifference curves	111
3.7	Indifference curves no longer fan out if disappointment is very strong	113
3.8	An illustration of probability weighting	114
3.9	Indifference curves in the Allais paradox with rank-dependent utility	115
3.10	A prospect theory utility function for money	118

3.11 The proportion of subjects appearing risk-averse, risk-neutral or risk-loving if the reference point was a sure thing or gamble 126
3.12 The proportion of subjects who price the $ bet higher, choose the $ bet and exhibit a preference reversal 128
3.13 Whether the rate of preference reversal differs if there are no, truth-revealing or intermediate incentives 129
3.14 The expected regret and rejoice from choosing prospect AA rather than AB 132
3.15 The equity premium puzzle 137
3.16 The returns in money and utility from investing in two hypothetical assets 138
3.17 The proportion of gains realized (PGR) minus the proportion of losses realized (PLR) in simulated investments 141
3.18 Why prospect theory does not cause the disposition effect 141
3.19 The change in the number of investors looking up their investments depending on whether the stock market was up, down or neutral over the previous week 143
3.20 The estimated size of the shadow economy between 1999 and 2007 as a percentage of GDP 147
3.21 Expected utility (rescaled) for different amounts of income reported and for different values of the risk aversion parameter r^r_u 149
3.22 Expected utility (rescaled) for different amounts of income reported depending on the reference level 151
3.23 The proportion of subjects who filed a tax return (tax form filing rate) and the proportion of income that is declared (reporting compliance rate), by treatment 153
3.24 The amount of money the defense was willing to offer and the prosecution was willing to accept for settling out of court 158
4.1 Estimated discount factors for four types of choices 165
4.2 Choice of when to visit an aunt and friends 168
4.3 Hyperbolic discounting compared to exponential discounting 171
4.4 The difference between time-consistent and present-biased preferences 173
4.5 The amount of money allocated to the earlier payment, for different levels of interest rate, comparing certain payment (1, 1) with a 50 percent chance of payment (0.5, 0.5) 187
4.6 The amount of money allocated to the earlier payment, for different levels of interest rate, comparing four different levels of risk 188
4.7 Average cost per visit to the gym for the first six months of monthly membership or year of annual membership 194
4.8 The timeline of the model of gym pricing 195
4.9 Willingness to pay for a more fuel-efficient car and the estimated price taking into account predicted fuel savings 203
5.1 The willingness to sell a bet on each of the boxes 210
5.2 A cascade experiment 224
5.3 The proportion of subjects who chose to see what others had done and then copied or not the majority 228
5.4 The game tree of a signaling game 229
5.5 Preferences of respondents in surveys of patients in 1996, 2000, 2004 and 2008 232

5.6	Cesarean rates in different counties of the United States for normal-weight and low-weight babies	235
5.7	Price of the S&P 500 index compared to the estimated fundamental value	237
5.8	A description of an experimental asset market that can give rise to bubbles	238
5.9	A bubble in an experimental asset market	238
5.10	The second time that the group takes part in an experimental asset market the bubble is much reduced	241
5.11	The deviation of price from fundamental value when subjects have already experienced two markets	243
5.12	The buy–sell imbalance for different types of investors depending on whether the stock appeared in the news or not	245
5.13	The proportion of outcomes that were correctly predicted by looking at photos of the two candidates	246
6.1	Choices the first time that subjects play a p-beauty contest	254
6.2	The mean number chosen when playing a p-beauty contest four consecutive times, for each of the experimental sessions	255
6.3	Classifying subjects by type after they had played 16 different games	257
6.4	The proportion of people at each cognitive level in a cognitive hierarchy	260
6.5	A sequential battle of the sexes game	266
6.6	A sequential battle of the sexes game with an outside option	267
6.7	The QRE for the game in Table 6.7	270
6.8	The estimated distribution of λ across subjects for each of the three games	271
6.9	A model of belief-based learning	276
6.10	Median estimates of ω in a belief-based model of learning	277
6.11	The adjusted R^2 for the three different learning models in six games	279
6.12	Assessing the goodness of fit of various equilibria relative to the Nash equilibrium	283
6.13	Within FlyMe and YouFly the board of directors play a game	284
6.14	The mean number chosen when playing a 0.67-beauty contest four consecutive times, comparing individuals versus individuals and teams versus teams	285
6.15	The proportion of individuals and teams of each type of level-k thinking	286
6.16	The average number of times (out of four rounds) that different types of decision-maker were closest to the optimal number	287
6.17	The proportion choosing the Nash equilibrium strategy depending on whether a dominant strategy existed or not and whether the game was solvable in one, two, three or no steps of iterated deletion of dominated strategies	288
6.18	Percentage of subjects of each type in four auction experiments	295
6.19	Categories of bidder	296
6.20	An example of coordination failure in a weakest link game	298
6.21	The proportion of subjects choosing 1 and 7 in weakest link games	298
6.22	How well various models fit data from the weakest link game	300
6.23	Average effort in a weakest link game comparing teams and individuals	300
6.24	The minimum effort in the turnaround game	301
6.25	Average effort in the turnaround game	302
6.26	Cumulative contributions in a typical threshold public good experiment	304

6.27	Success at providing a threshold public good depends on the institution	304
6.28	How nine subjects were allocated to groups within a network	305
6.29	A subject is more likely to choose high effort in a group that has a higher minimum effort	306
6.30	Three different network structures for five people	307
6.31	The score of monetary policy-makers trying to keep unemployment and inflation on target	310
6.32	Timeline in a limit price entry game	311
6.33	The proportion of strategic play in a limit pricing game	313
6.34	The predicted number of entrants in a market entry game compared to the average number of entrants in four groups	314
6.35	The behavior of subjects in a market entry game	315
6.36	Profits in a quantity-setting game and classification of outcomes	316
6.37	The output of firm 1 and firm 2 combined in Cournot and Stackelberg markets compared to the equilibrium	317
6.38	The best response and observed response of followers to the output of a Stackelberg leader	318
7.1	The amount that proposers gave in dictator experiments	323
7.2	Investments and returns in a trust game with and without social history	325
7.3	Wage and effort levels in a gift exchange experiment comparing one shot (OS) and repeated game (RG)	326
7.4	The proportion of offers rejected in an ultimatum game	327
7.5	The share that proposers gave in ultimatum experiments when the amount of money given to the proposer ranged from $5 to the equivalent of $1,500	329
7.6	Punishment in a dictator game with third-party punishment compared to in the ultimatum game	330
7.7	Contributions in linear public good games with and without punishment and with stranger and partner matching	331
7.8	The amount of punishment depends on how much the subject contributed relative to the average contribution in the group	332
7.9	Responders' impact on proposers in the moonlighting game	333
7.10	The share that proposers offer in ultimatum games with competition	334
7.11	The share offered by proposers in the dictator game	336
7.12	The share proposed in the ultimatum game, invested in the trust game and returned in the trust game comparing individuals and teams of three	337
7.13	The amount shared by dictators in dictator games with take and earnings	352
7.14	A simple model of worker reciprocity	361
7.15	The average wage and reservation wage in experimental labor markets	362
8.1	Accuracy of children's ability to predict the giving of fellow students in a dictator game, above that of guessing randomly	384
8.2	Ultimatum game offers across different cultures	386
8.3	Punishment in a public good game in different locations	387
8.4	The proportion of respondents who said 'Most people can be trusted'	389
8.5	An example of the haystack model	391
8.6	Experimental setup for an ultimatum game with chimpanzees	394
8.7	The proportion of 8 and 2 offers by game and the proportion of 8 and 2 offers rejected	395

8.8 The share given in the dictator game and offered in the ultimatum game by school-aged children 396
8.9 Running times of boys and girls 399
8.10 The proportion who chose competition conditional on how well they had performed 400
8.11 Food production and consumption by age in males and females, and in chimpanzees 404
8.12 The proportion of subjects who cooperate in public good games when playing with their partner or a stranger 406
8.13 The difference in average test score comparing schools involved in the project and those not, in a particular year 409
8.14 The probability of the loan being repaid given the observed levels of honest reporting 412
9.1 A schematic representation of modularity 416
9.2 Some areas of the brain seen from the medial surface of the brain (i.e. by slicing the brain down the middle) 418
9.3 Terminology to describe location in the brain 420
9.4 Activity in the nucleus accumbens on seeing the circle, around two seconds, and knowing how much is won, around eight seconds 425
9.5 Activity in the ventromedial prefrontal cortex on seeing the circle, around two seconds, and knowing how much is won, around eight seconds 426
9.6 The response in the nucleus accumbens and ventromedial prefrontal cortex to anticipating, and then winning, or not winning, money 428
9.7 The change in activation in two areas of the VMPFC following a reward during a learning stage, or loss in a task reversal stage 429
9.8 The change in activation in two areas of the VMPFC following a reward or loss of varying magnitudes 430
9.9 Activity in the NAcc (left ventral striatum) increases with expected value and risk 431
9.10 Activation in the NAcc (left striatum) and VMPFC (left orbitofrontal cortex) depends on whether the choice is with risk or uncertainty 432
9.11 The fit between values and activity in three areas of the brain 433
9.12 Activation in the dorsal and ventral striatum during a learning task on seeing the options or cue and getting the reward 434
9.13 Areas of the lateral prefrontal cortex and parietal cortex are more activated in those with a preference for ambiguity and risk 435
9.14 Different choices correlate with increased activation in different areas of the brain 436
9.15 Increased connectivity is observed between the dorsomedial prefrontal cortex and the insula and parietal cortex when someone makes a choice inconsistent with their normal strategy 437
9.16 Activation in the amygdala and anterior cingulated cortex depending on the frame and choice 438
9.17 Activation in the VMPFC, anterior cingulated cortex and medial prefrontal cortex (MPFC) depending on level of reasoning and whether the opponents are humans or computers 439
9.18 Timeline of activation in the (right anterior) insula depending on whether the offer is fair or unfair 440

9.19 Activation in the (right) insula and (right) dorsolateral prefrontal cortex for offers that are subsequently rejected or accepted 441
9.20 Acceptance rate and perceptions of an unfair offer in the ultimatum game 442
9.21 Activation in the striatum depends on the type of punishment 443
9.22 The effect of inequality on the valuation of subsequent transfers and activation in the NAcc and VMPFC 444
9.23 Activation in the NAcc and DLPFC depending on whether the reward is immediate or delayed 446
9.24 Areas associated with immediate reward (beta) are more activated when the subject chooses the early option, and areas associated with delayed reward (delta) are more activated when choosing the later option 446
9.25 The proportion of subjects choosing chocolate cake is higher when the subject has to remember a seven-digit number and is shown the snacks for real 447
9.26 Accuracy on Raven's matrices for shoppers in New Jersey and farmers in India depending on whether they are facing a hard or easy financial scenario 448
9.27 The proportion of people surveyed, by age group, who want to quit smoking 455
9.28 What smokers who are planning to quit in the next 12 months expect will happen after 12 months 455
10.1 The U-index compared to life satisfaction 468
10.2 The change in reported life satisfaction following a life-changing event 469
10.3 Life satisfaction and income 470
10.4 Life satisfaction and GDP per capita in different countries in 2006 470
10.5 Life satisfaction and income over time 471
10.6 Changes in happiness relative to the baseline in, before and after life-changing events 472
10.7 Estimated effect on life satisfaction of being employed versus unemployed 473
10.8 The instant utility over time of Ian on a holiday in Italy 474
10.9 Instant utility while hand is immersed in cold water for 60 or 90 seconds 475
10.10 How the 90-second trial is rated compared to the 60-second trial 476
10.11 Students' expectations before, experience during and recollections after a spring vacation 477
10.12 Predicted or recalled quality of life compared to actual before and after a renal transplant 478
10.13 The endowment effect 479
10.14 The proportion who chose an unhealthy snack for one week's time depends on current hunger 479
10.15 More chose an unhealthy snack than predicted they would, but fewer chose the unhealthy snack if they said they would choose a healthy snack 480
10.16 How many unplanned items shoppers bought depended on whether or not they ate a muffin and whether or not they had a shopping list to hand 481
10.17 Performance in a proofreading task 484
10.18 The proportion who chose chocolate rather than money depending on whether there was a limited or extensive choice, and whether they got to taste a chocolate they chose or the experimenter chose 485
10.19 How opinions and behavior changed if subjects had a limited choice or extensive choice of chocolates 486

10.20 Changes in quality of life following two possible treatments, A and B, or no treatment 488
10.21 Changes in quality of life following two possible treatments, taking into account adaptation 490
10.22 Measures of pain and rate of return for a repeat colonoscopy depending on whether the procedure was prolonged or not 491
10.23 The average rating of portfolios at UCLA and SwedishAmerican 495
10.24 The inflation–unemployment trade-off of those on the political left is higher than of those on the political right 497
10.25 Change in demand if sales tax is displayed 498
10.26 The consequences of ignoring the sales tax 500
11.1 The average net profit as a percentage of the maximum over 20 rounds of a common resource problem 507
11.2 Matching algorithms compared 512
11.3 The radio wave portion of the electromagnetic spectrum and how it is allocated in the United Kingdom 514
11.4 The revenue in euros per capita from the 3G spectrum auctions in different European countries in date order 515
11.5 Changes in saving behavior after one year 519
11.6 Participation rate in a 401(k) retirement savings plan 520
11.7 Contribution rate and asset allocation in the new and window groups 521

Tables

2.1	Cereals for sale, and their characteristics	32
2.2	The utility of each cereal if initial wealth is $100	32
2.3	Five search heuristics with a very brief statement of how they can help with search	34
2.4	Examples of how Anna's choice may be influenced by the three psychological effects	41
2.5	The proportion of subjects who answered 'No', 'Yes', 'Yes' to purchasing a good when the alternative was a related purchase, gift or unrelated purchase	54
2.6	Four different types of behavior possible in cheap-talk sender–receiver game	62
2.7	The number of sellers and buyers before and after a shift in demand	67
2.8	The market is most efficient in equilibrium, as evidenced by the high total profit	68
2.9	The market efficiency of five experimental markets, comparing human participants to zero-intelligence traders	71
3.1	The possible consequences of Alan buying insurance	98
3.2	Alan's expected utility depends on his utility function	99
3.3	The expected value and risk of different choices	100
3.4	The Allais paradox	102
3.5	The near-certainty effect	102
3.6	Prospects that can lead to extreme risk aversion	103
3.7	The amount a person would need to gain on a 50:50 bet to take the bet, if risk-averse for small gambles	105
3.8	Prospects with gains and losses	107
3.9	Framing of losses versus gains	107
3.10	The Allais paradox can be explained by disappointment	112
3.11	Rank-dependent utility and the Allais paradox	115
3.12	How disappointment theory, rank-dependent expected utility and prospect theory differ from expected utility	119
3.13	How the utility of Alan will depend on his expectations	120
3.14	Risk aversion for small gambles is less if there is existing risk	122
3.15	A set of choices for which the reference point is a certain amount	125
3.16	A set of choices for which the reference point is a prospect	125
3.17	Prospects that can lead to a preference reversal	128
3.18	Four possible explanations for a preference reversal and the number of choices by subjects fitting each explanation	131
3.19	Choice of deductible on home insurance	144

3.20	Deductible choice in Israeli car insurance	146
3.21	Estimated tax evasion by profession	148
3.22	A summary of the outcome if Alan reports all, half or none of his income	149
3.23	The fourfold pattern of risk attitudes in litigation	156
3.24	Outcomes of going to court or settling out of court	157
3.25	A summary of outcomes of litigation	159
4.1	The day-by-day utility for Maria of doing her homework on different days	163
4.2	The inter-temporal utility of each plan for three different discount factors	163
4.3	Choice of when to eat at a restaurant	169
4.4	Maria deciding when to visit an aunt and friends	169
4.5	Quasi-hyperbolic discounting in the homework example	174
4.6	Payoffs from Maria going to watch a movie	175
4.7	Inter-temporal utility in the movie example with time-consistent and present-biased preferences	175
4.8	A summary of the two examples: choice and utility	175
4.9	The utility of Maria on Saturday	178
4.10	The homework example with reference-dependent preferences	181
4.11	Preference over sequences of outcomes with the anticipated and recollected utility	183
4.12	Preferences over sequences in the homework example	183
4.13	Discounted expected utility if taking the money today or tomorrow	185
4.14	Two sequences of consumption compared in a model of habit formation	192
4.15	The consequences of optimal pricing in the gym example with present-biased preferences	196
4.16	Choices of calling plan, estimates of usage and the proportion of customers making the wrong choice	198
4.17	The long-run consequences of four climate change policies	199
4.18	The discount factor for which each policy is optimal, and the inter-temporal utility of each policy for different discount factors	201
5.1	John updates his beliefs after a test drive	208
5.2	The proportion of subjects fitting each model of preferences and the average willingness to sell	211
5.3	Evidence of confirmatory bias in attitudes to capital punishment	213
5.4	Evidence of confirmatory bias in assessing the ability of a child	213
5.5	The probability of an incorrect belief after reading ten or 100 reviews	215
5.6	Payouts to winners in state lotteries	217
5.7	Evidence (or lack of it) for the hot hand fallacy	217
5.8	Beliefs when there is the law of small numbers	219
5.9	Probability of calling out a sequence of blacks and reds in a deck of cards with five black and five red cards	220
5.10	Data from the 1982 Wimbledon men's singles final	221
5.11	An information cascade	222
5.12	Did subjects play the empirical optimal strategy in cascade experiments?	225
5.13	Play in a signaling game	230
5.14	Approximate number of deaths in the United States and estimated number of deaths from different causes	233
5.15	The highest and lowest rates of performing a treatment in different areas of the United States	234

5.16	The relative absolute deviation and relative deviation of prices from fundamental value in eight asset market treatments	240
5.17	The three different treatments in an experiment with experienced subjects	242
6.1	Proportion of subjects of each type in four p-beauty contests	255
6.2	Level-k thinking in the p-beauty contest if type L-0s choose randomly, compared to choosing the Nash equilibrium	257
6.3	Level-k thinking in the 0.5-beauty contest compared to a model of cognitive hierarchy	262
6.4	The answers subjects gave, and the proportion giving each answer, to two questions	262
6.5	A matrix game	265
6.6	The Nash equilibria in the sequential battle of the sexes	267
6.7	A matrix game	269
6.8	Emma using reinforcement learning in the 0.5-beauty contest	273
6.9	Emma using belief-based learning in the 0.5-beauty contest with $\rho = 0.5$	275
6.10	EWA parameter estimates	279
6.11	A matrix game to illustrate payoff sampling	280
6.12	Transforming the payoffs to take into account loss aversion	281
6.13	Impulses after taking into account loss aversion	281
6.14	EWA parameter estimates comparing individuals and teams	287
6.15	Bidding in the Magic game	292
6.16	Bidding strategies of a type L-k comparing what happens with two different assumptions about type L-0s	294
6.17	Payoffs in a weakest link game	297
6.18	The efficient equilibrium probability of high effort and payoff compared to observed proportion choosing high effort	308
6.19	The efficient equilibrium probability of high effort and observed proportion choosing high effort with incomplete information	308
6.20	Possible equilibria in the limit price entry game	312
7.1	A brief overview of the games I will look at	322
7.2	Some possible payoffs in a four-person linear public good game	331
7.3	The ERC model and linear public good game	341
7.4	The linear public good game and Fehr–Schmidt model from the perspective of Anna	343
7.5	Payoffs, predictions and choices in distribution games	345
7.6	The proportion who offered and the proportion who would reject an $8, $2 split in ultimatum games with restricted choice	346
7.7	Offers in a dictator game with varying amounts of moral anonymity	357
7.8	The proportion who contribute, the average contribution conditional on giving something and the average contribution to Poas National Park	359
7.9	The percentage of employers who thought that a 10 percent wage cut or increase would change worker effort, and change it by more than 10 percent	360
7.10	The probability of low, medium or high output depending on the effort of Miguel	366
8.1	The four examples involving aggregate risk	378
8.2	The two examples with interdependence	380

8.3	The payoffs if people are cooperative or non-cooperative	381
8.4	Payoffs in the hawk–dove game	382
8.5	The three recurring stages of the haystack model	390
8.6	Gender equality data	397
8.7	Predicted proportion of portfolio invested in risky assets	398
8.8	The percentage of $1 offers accepted from a model fitted to the experimental data	401
8.9	Minimal acceptable offers	402
9.1	Parts of the central nervous system and some of the main functions that go on in that part	419
9.2	Parts of the cerebrum and some of the main functions that go on in that part	421
9.3	Four different brain processes with an example	421
9.4	Main components of executive control systems	423
9.5	Two areas of the brain that connect with the VTA and are important in economic valuation and decision-making	424
9.6	The utility of several different addiction plans	453
9.7	The consequences of underestimating the addictiveness of cigarettes, and present bias in the example	456
9.8	The effect of the discount factor on Rob's desire to smoke during dinner	458
9.9	The discount factor matters if there are two cues	459
10.1	The day reconstruction method	467
10.2	Ian's utility depends on when he finishes the dissertation	482
10.3	Emma's utility depends on when she buys the book	483
10.4	Examples of QALY measures using the EQ-5D system	489
10.5	The consequences of projection bias in a simple example	493
10.6	Changes in reported life satisfaction if there is unemployment or inflation	496
10.7	The consequences of a 5 and 10 percent tax	501
11.1	Some of the possible payoff combinations in a common pool resource game involving Anna and seven others	507
11.2	An example of possible preferences of junior doctors looking for work and hospitals looking for junior doctors	510
11.3	The deferred acceptance algorithm allocating doctors to hospitals for the example	510
11.4	Different methods used in the United Kingdom	511
11.5	Bidding over the rounds in the UK 3G spectrum auction	516
11.6	Principles for good choice architecture, spelling 'NUDGES'	523
11.7	An overview of MINDSPACE	525
11.8	The leading risks for lost years of healthy life	527

Research methods boxes

2.1	The infinite regress problem	35
2.2	How long should an experiment take?	37
2.3	Can you replicate it?	43
2.4	Adversarial collaboration	50
2.5	Eye tracking and lie detection	61
2.6	Videoing team decision-making	63
2.7	Intelligent institutions or people?	71
3.1	Measuring risk aversion	101
3.2	Money or something else	103
3.3	Risk aversion and incentives in the lab	106
3.4	The development of an idea	116
3.5	TV game shows	121
3.6	Behavioral economics' longest saga	127
3.7	Do incentives matter?	129
3.8	Don't always trust intuition	142
3.9	External validity, something is better than nothing	154
4.1	Empirical versus lab evidence	167
4.2	Axioms and behavior	177
4.3	Delayed payments in the lab	186
4.4	Inter-generational lab experiments	200
5.1	Empirical optima versus Nash equilibrium	224
5.2	Meta-data analysis	226
5.3	Deception in experiments	227
5.4	What if subjects do not understand instructions?	239
5.5	Experienced trader versus experienced group	241
5.6	Why do experiments?	243
6.1	Seeing how subjects choose	258
6.2	Algebra has limits?	263
6.3	Equilibrium refinements and selection	268
6.4	Estimating learning models	276
6.5	Field experiments	293
6.6	Money versus real effort	299
6.7	Artefactual field experiments	302
7.1	Single- and double-blind: anonymity versus reality?	323
7.2	Finding out what you actually want to know: the strategy method	328
7.3	Does stake size matter?	329
7.4	Letting subjects chat	336
7.5	Psychological game theory	351

8.1	Evolutionary game theory	381
8.2	Experimenting with children	383
8.3	Statistics and lies: more, less, or just different?	386
8.4	Incentives over time	400
8.5	Experiments in the household	406
8.6	Experimenting with policy	409
9.1	Monitoring brain activity	416
9.2	The terminology of neuroanatomy	419
9.3	Sample sizes and confidence intervals	426
10.1	The law of large numbers: can we trust happiness data?	466
10.2	An alternative measure of happiness	466
10.3	Double-blind experiments	491
11.1	Communication	508

Preface

The first edition of this book came about from teaching a course called 'Strategy and games' for first-year undergraduates. The idea of the course was to illustrate through classroom experiments some of the basic ideas of economics and game theory. In teaching that course my interest and knowledge of behavioral economics grew a lot; as did my frustration at the limited resources available to students wanting to learn more about this fascinating area of economics. Eventually, I decided it was time to write a book.

My basic objective when writing the book was to convey both the excitement and the importance of behavioral economics. I wanted to explain the basic principles, ideas and results of behavioral economics and show how fascinating they are. I also wanted to illustrate through applications why behavioral economics is fundamentally important in understanding the world around us. Finally, I wanted a book that was accessible to a general reader, not just those trained in economics and comfortable with algebra.

This third edition builds on the previous two. Behavioral economics is an evolving subject, and so there is always new material to add and updates to take note of. There are, for example, new sections on stochastic reference points and the relationship between risk and time preferences. New applications are also considered, including behavioral contract theory and bargaining in the family.

This book would not have been possible without those who have contributed to the literature on behavioral economics and given me such great material to work with. Particular thanks are due to those who made my job so easy by writing papers over many years that, when put together as a whole, gave very clear insights into economic behavior. I would also like to extend a sincere thanks to very many people who have given me useful feedback on the first two editions of the book. It is important to make clear, though, that the opinions expressed in this book are mine, and not necessarily those of the researchers whose work I refer to.

A lot of thanks is due to my wife, Anna, who has helped so much in writing this and the earlier editions of the book. On a personal level, she has been patient and supportive while I wrote them. On a practical level, she is my greatest critic, never happy until everything is explained as fully and clearly as possible; she is also a great source of new ideas and new ways of thinking about old ideas. I should also thank all the students and colleagues connected with the Strategy and Games course, for teaching me so much. Finally, I want to give a big thanks to Myrna Wooders, who, as well as being a fantastic person to know, has long been, and continues to be, a great inspiration and mentor.

Part I

Introduction



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Chapter 1

An introduction to behavioral economics

> However beautiful the strategy, you should occasionally look at the results.
>
> Sir Winston Churchill

It seems natural to start the book by asking: what is behavioral economics? That's a tough question, but I will give you three answers. You might think that this is two too many, but, if we put all three together, we get a definition of behavioral economics that I am happy with. So, here's the first answer.

> Behavioral economics is about understanding economic behavior and its consequences. It's about understanding why someone buys a hotdog, goes to work, saves for retirement, gives to charity, gets a qualification, sells an old car, gambles on a horse race, cannot quit smoking, etc. It's also about understanding whether people make good or bad choices, and could be helped to make better choices.

I like this definition, because I want to know the answer to all those questions. The problem is that we end up with far too broad a definition to be very useful. So, we need something a bit different.

For that something different, we can use the fact that economists already have a fantastic model with which to try and understand economic behavior and its consequences. You can find this **standard economic model** in any economics textbook, and it gives us a beautiful story of how things *might* work. We can, for example, derive demand and supply curves, solve for equilibrium prices and quantities, and design mechanisms or policy to increase efficiency and social welfare.

The potential problem with this standard economic model is that it is based on some strong assumptions. For example, it assumes people can be approximated by a ***Homo economicus*** who is rational, calculating and selfish, has unlimited computational capacity and never makes systematic mistakes. It also assumes **market institutions** work, and so, for instance, prices should converge, as if by magic, to equilibrium. There is no obvious reason why humans can be well approximated by *Homo economicus*, or why market institutions should work. This leads to my second definition.

> Behavioral economics is about testing the standard economic model on humans, seeing when it works and when it does not, and asking whether it can be tweaked, or given an overhaul, to better fit what we observe.

This is a more practical and specific definition, but it does leave open the question of how we are going to test and tweak the model. After all, surely economists have been testing the model all along? Well, unfortunately, we now know the standard economic model is flexible enough that it can be manipulated to explain just about any anomaly we ever observe. We somehow need to step out of this trap. This leads to my third and final definition.

> Behavioral economics is about applying insights from laboratory experiments, psychology and other social sciences in economics.

We will, therefore, be guided by the evidence on human behavior and psychology, both in testing and in tweaking or overhauling the standard economic model. In short, we will replace *Homo economicus* with *Homo sapiens*.

If you combine all three definitions, I think we can strike a nice balance. Behavioral economics is about working constructively with the standard economic model to get a better understanding of economic behavior. The objective is definitely not to criticize the standard economic model, or to accentuate the negatives. Testing the standard model is a means to an end, and that end is to understand economic behavior as best we can.

Behavioral economics has really come of age in the last 50 years or so, and so a lot of progress has been made. Various things could have happened as we started to put the standard economic model to the test. The model could have worked perfectly; that would have been fantastic news for economics, but not so exciting for the future of behavioral economics. At the other extreme, the model could have proved useless; that would be bad news all round (except for those who like to poke fun at economics). What has actually happened is an exciting mix in which the standard economic model sometimes seems to work very well, sometimes to work very badly, but most of the time is not far off, and with a bit of tweaking gets a lot better.

This is exciting because it means that behavioral economics can teach us a lot. It can tell us when the standard economic model does a good job and when not, and it can tell us how to change the model to get a better fit with reality. That is progress, and that is why behavioral economics is revolutionizing economics and our ability to understand economic behavior. In this book I hope to talk you through some of the main results and insights of behavioral economics in order that you can see for yourself how exciting and informative it is.

This chapter will set the scene a little by introducing some basic ideas and concepts, before previewing the rest of the book.

1.1 The history and controversies of behavioral economics

Behavioral economics has an interesting and checkered history. I will not delve too deeply into that history here, but an overview is useful because it will allow me to expand on the motivations behind behavioral economics. I will also explain some of the controversies and debates that surround it. These controversies and debates still rumble on, and so it is helpful to have some understanding of them before you see what behavioral economics has to offer.

It is difficult to say when behavioral economics began, but I am going to credit Adam Smith with being its founder. Any student of economics should be familiar with Adam Smith's book *An Inquiry into the Causes of the Wealth of Nations*, first published in 1776. In that book Smith famously explained the invisible hand of the market. Less well known to most economists is a book that Smith first published in 1759, called *The Theory of Moral Sentiments*. It was actually in this book that the invisible hand first made an appearance. More interesting, for our purpose, is how Smith explains that people are not motivated solely by self-interest, but also feel a natural sympathy with others, and have a natural sense of virtue.

In short, in *The Theory of Moral Sentiments*, Smith talks about many things that in the last 30 years or so have become major issues in behavioral economics. For example, a theme through many parts of his book is the importance of reward and punishment (something we will pick up on in Chapter 7). Another theme is the influence of custom and fashion (which we will pick up on in Chapter 5). In my mind, Adam Smith comfortably does enough to be the father of behavioral economics. The intriguing question is why we had to wait over 200 years for the ideas that he talked about to be taken seriously by economists.

That turns out to be slightly the wrong question, because psychology did go hand in hand with economics for a long time after Adam Smith. Early economists gave much weight to emotions, impulses, stimulus, morals, and the like. For example, the law of diminishing marginal utility, one of the most fundamental principles of the standard economic model, was based on psychological ideas; economists would often appeal to the work of psychologists in supporting and developing their work. At the beginning of the twentieth century, however, economics turned away from psychology, and behavioral economics, if we can call it that, disappeared for over half a century.

This shift was initiated by Vilfredo Pareto. In a letter dated 1897 Pareto writes: 'Pure political economy has therefore a great interest in relying as little as possible on the domain of psychology.' This is presumably why, when he published a paper in 1900 outlining a new approach to the theory of choice, he claims as one of its main achievements that 'every psychological analysis is eliminated' (these quotes taken from Bruni and Sugden 2007). Why would we want to rid economics of psychology, and how is it possible to do so?

It is easier to start with the latter question. Psychology can be taken out of economics by focusing on choice rather than desire. Instead of trying to work out why people do things, we can make inferences based solely on what they do. To quote again from Pareto: 'I am not interested in the reason why man is indifferent between [one thing and another]: I notice the pure and naked fact.' This approach makes a lot of sense, because it allowed Pareto, and subsequent economists, to abstract away from difficult psychological questions and develop a mathematical theory of rational choice. If people are rational then they will reveal their desires through their choices, and so we need focus only on choice.

Asking, as Pareto did, what happens if people are rational is a good, logical thing to do, because it provides a natural benchmark to work with. The same could be said of asking, as Smith did in discussing the invisible hand, what happens if people are selfish. Assuming for mathematical convenience that people are rational and selfish clearly does not mean, however, that people actually are rational or selfish. Smith clearly recognized this, and Pareto did so too by distinguishing when rational choice should or should not be expected to approximate what people do.

The problem is that these caveats can easily be forgotten in the beauty or simplicity of the argument. Smith's conjecture that someone, 'by pursuing his own interest, frequently promotes that of the society more effectually than when he intends to promote it' is truly intriguing. Similarly, Pareto's new theory of choice seemed widely applicable. In the face of such appeal and convenience it became easy to overlook the fact that people are neither rational nor selfish; *Homo economicus* became king, and economics became very distant from psychology.

Interestingly, *Homo economicus* became ever cleverer as the twentieth century progressed, Pareto's caveats long since forgotten. The assumption of rational expectations, for instance, meant that *Homo economicus* knew far more about the workings of the economy than any economist. As behavioral economist Richard Thaler, writing while behavioral economics was still in its relative infancy (1990: 203), says:

> The problem seems to be that while economists have gotten increasingly sophisticated and clever, consumers have remained decidedly human. This leaves open the question of whose behavior we are trying to model. Along these lines, at an NBER [National Bureau of Economic Research] conference a couple of years ago I explained the difference between my models and Robert Barro's by saying that he assumes the agents in his model are as smart as he is, while I portray people as being as dumb as I am.

In my mind, to assume people are like rational and selfish *Homo economicus* is the most natural, objective place to begin thinking about modeling economic behavior. Indeed, in many of the chapters that follow I will start by asking what a selfish, rational person would do. The crucial point, though, is that it is the start point and not the end point. It is the best way to start thinking about modeling economic behavior but not necessarily the best way to model economic behavior. A crucial distinction!

The standard economic model does, though, march ever onwards, and behavioral economics has so far done relatively little to quell the tide. Many economists continue to derive ever more complicated and elaborate results based on the standard economic model. Whether these results are any use in understanding what happens on planet Earth, I am not so sure. That's why we need behavioral economics, and that's why the time would inevitably come for behavioral economics to be reborn.

1.1.1 Behavioral economics is reborn

From the 1960s onwards psychology gradually made a return to economics. I would suggest four distinct elements to its comeback. Let's look at each in turn.

The first element I would call the 'you cannot be serious' attack, and give the main credit to Herbert Simon. Simon seriously questioned the sense of approximating people by *Homo economicus*. For example, in a paper published in 1955, he solves for how a rational person should behave before stating: 'My first empirical proposition is that there is a complete lack of evidence that, in actual human choice situations of any complexity, these computations can be, or are in fact, performed.' Instead, Simon suggested looking at the information and computational capacities that

humans possess, and using this as the starting point for economic models. Recognizing the limitations faced by humans led to the term 'bounded rationality'.

Simon won the Nobel Prize in Economics in 1978 for his 'pioneering research into the decision-making process within economic organizations'. His calls to replace *Homo economicus* with something more human-like fell largely on deaf ears, however. Symptomatic of this was Simon leaving the Graduate School of Industrial Administration at Carnegie Mellon University for the psychology department. In his autobiography (Simon 1991) he writes: 'My economist friends have long since given up on me, consigning me to psychology or some other distant wasteland.'

One thing notably lacking in much of what Simon wrote was proof that *Homo economicus* is not a good approximation of how people behave. He may have thought this was obvious (many do), but the lack of any formal proof made it easy for economists to ignore his work. The same could not be said of the second element I want to talk about. I will call this the 'your assumptions are wrong' attack, and give the main credit to Daniel Kahneman and Amos Tversky. The approach here is one of demonstrating that people really are very different from *Homo economicus*. One way to make the point is to get people, including economists, to answer simple questions such as this (which I took from Thaler and Sunstein 2008):

> A bat and ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost?

If you said ten cents then you are like most people, but you need to think again, because the answer is five. Here's another famous example from Tversky and Kahneman (1983):

> Bill is 34 years old. He is intelligent, but unimaginative, compulsive, and generally lifeless. In school, he was strong in mathematics but weak in social studies and humanities.
>
> Rank the following eight statements from most probable to least probable:
>
> Bill is a physician who plays poker for a hobby.
> Bill is an architect.
> Bill is an accountant.
> Bill plays jazz for a hobby.
> Bill surfs for a hobby.
> Bill is a reporter.
> Bill is an accountant who plays jazz for a hobby.
> Bill climbs mountains for a hobby.

If you are like most then you said it is more likely Bill is an accountant than Bill is an accountant who plays jazz for a hobby than Bill plays jazz for a hobby. Now, how can it be more likely that Bill is an accountant who plays jazz than that Bill plays jazz? It cannot.

If most people make mistakes such as this, how can we possibly expect them to do the complex calculations that *Homo economicus* routinely does? With this example, and many others, very clear evidence was provided that people are not like *Homo economicus*, or, at least, not like the version assumed in the standard economic model.

Different assumptions, therefore, seemed sensible. To quote from an earlier paper published by Tversky and Kahneman (1981: 453):

> The definition of rationality has been much debated, but there is general agreement that rational choices should satisfy some elementary requirements of consistency and coherence. In this article we describe decision problems in which people systematically violate the requirements of consistency and coherence, and we trace these violations to the psychological principles that govern the perception of decision problems and the evaluation of options.

Daniel Kahneman won the Nobel Prize in Economics in 2002 for 'having integrated insights from psychological research into economic science, especially concerning human judgment and decision-making under uncertainty' (Amos Tversky was, unfortunately, no longer alive). Even so, I am not so sure their work, and that of others in a similar vein, was ultimately that crucial in the comeback of psychology in economics. That's because the attack was still too easy to dodge for economists confident in the standard economic model. After all, was it not obvious that people are not like *Homo economicus*? The real issue is whether models in which people are approximated by *Homo economicus* make good predictions. The early work of Kahneman, Tversky and others had less to say on this issue.

To illustrate the point, we can get to the third element, which I will call the 'markets work' revelation, and give the main credit to Vernon Smith. Starting in 1955 Smith performed a series of experiments to see whether basic predictions of the standard economic model about markets would prove correct. We will look at these experiments in more detail in the next chapter, but basically the predictions proved good. A stunning result! Maybe, therefore, it does not matter if people are not like *Homo economicus*; the standard economic model can still work.

These initial experiments led to a continuing line of research on market institutions that, to my mind, provides the most important results to have come out of behavioral economics. In 2002 Vernon Smith won the Nobel Prize in Economics 'for having established laboratory experiments as a tool in empirical economic analysis, especially in the study of alternative market mechanisms'. I am still reluctant, though, to think that this work had too much of a hand in the comeback of psychology in economics. That's because, while moving beyond the standard economic model, it largely confirmed that the standard economic model was getting something right. Why, therefore, would we need behavioral economics?

While all three of the elements I have discussed so far were instrumental in the comeback of behavioral economics, you have probably guessed by now that it is to the last element that I would give most importance. This final element I will call the 'what equilibrium to choose?' problem, and give the main credit to Reinhard Selten. The problem became apparent with the rapid progress of game theory in the 1950s and 1960s. Game theory looks to capture behavior in strategic situations, and meant the demands on *Homo economicus* became ever more stringent. Not only should he or she be selfish, rational and cleverer than any economist, *Homo economicus* also needs to be telepathic in order to predict what others will do (and even that is not enough). Basically, in strategic situations, it usually becomes ambiguous what *Homo economicus* should do; it is ambiguous what the rational thing to do is.

The technical way to express this problem is to say that there are multiple equilibria. Somehow we need to try and say which of the equilibria 'makes more sense' or 'seems more likely to occur'. That's a bit like throwing darts at a dartboard while blindfolded. To have any chance of success it makes sense to question how people might think or reason in such strategic situations and observe what people do when they play games. In other words, it made sense to draw a little on psychology and to run controlled experiments.

In the late 1950s Selten began running experiments. These first experiments were primarily concerned with industrial economics and ultimately led to important theoretical ideas such as subgame perfection, something we will look at in Chapter 6. The main thing, though, was that experiments increasingly became seen as a useful way to learn more about economic behavior. To quote from an autobiographical sketch by Selten (1994):

> More and more I came to the conclusion that purely speculative approaches like that of our paper of 1962 are of limited value. The structure of boundedly rational economic behavior cannot be invented in the armchair, it must be explored experimentally.

Selten won the Nobel Prize in Economics in 1994, together with John Nash and John Harsanyi, 'for their pioneering analysis of equilibria in the theory of non-cooperative games'. More than anything else, I think that game theory was instrumental in the rebirth of behavioral economics. That's because it meant that the next logical step in developing the standard economic model was to draw from psychology and use experiments. The standard economic model had hit a dead end, and behavioral economics was needed to move it forward.

If we put together these four different elements in the rebirth of behavioral economics, it is no surprise that it has made a strong comeback. Behavioral economics has the potential to improve our understanding of economic behavior from so many different angles, and this eventually shone through. (Before moving on, it would be remiss of me not to mention that it did so through the work of many more than just the five people to whom I gave credit above.)

1.1.2 Behavioral economics and policy

If behavioral economics can improve our understanding of the economy then we should be able to put it to good use. Consequently, I will add a fifth element into the mix. This fifth element has little to do with the rebirth of behavioral economics but is proving instrumental in its rapid growth. I will call it the 'policies that work' problem. By its nature, economics is an applied subject; it should inform on how to alleviate poverty, avoid unemployment, regulate industry, and so on. Many, however, have become frustrated by the inability of economists to provide good answers to the important policy questions!

Increasingly, this problem is being traced back to an over-reliance on the standard economic model. The standard economic model suggests that intervention is needed only when markets fail because of things such as externalities, imperfect information or imperfect competition; if markets work, then people make the rational decision.

So, if people do not save for retirement, then they clearly want to end their life in poverty. If someone buys a mortgage they cannot afford then they knowingly gambled everything on house prices rising. Similarly, if someone becomes addicted to heroin then they chose to do so taking into account their financial constraints. To anyone other than economists these kinds of statements sound weird. They also sound weird to a behavioral economist.

Once we take into account the mistakes people make and the difficulties of coordinating on an equilibrium, the rationale for intervention becomes stronger. But it is important to realize that behavioral economics does not prescribe big government; rather, it prescribes clever government. I would distinguish two different elements to this.

One thing behavioral economics does is give fresh insight into what policies will work and what will not. For instance, the traditional approach to increase saving for retirement has been complex tax breaks; these are the kinds of things that appeal to *Homo economicus* but are ignored by *Homo sapiens*. A behavioral economics approach suggests things such as the save more tomorrow plan, which I will look at in Chapter 11; these are the kinds of things that appeal to *Homo sapiens* but are ignored by *Homo economicus*. It is partly for his groundbreaking work in this area that Richard Thaler was awarded the 2017 Nobel Prize in Economics 'for his contributions to behavioral economics'.

Another thing that behavioral economics does is make us question the purpose of policy. This is because the emphasis turns more towards measuring happiness than measuring rational choice. In fact, we shall see that choice can be influenced by all kinds of context effects, such as a '50 percent off' sticker or a TV advert on the consequences of evading tax; this makes the very notion of rational choice become somewhat meaningless. Suppose, for instance, Anna 'loves a bargain' and cannot resist buying a handbag that she has absolutely no need for, because the price is reduced 50 percent. We might call this a 'mistake', because Anna was only won over by the '50 percent off' sticker. But, if buying the handbag makes her happy, then was it a mistake? Is Anna going to be any happier if we point out to her that she has just been duped into buying a handbag she does not need? Will she be any happier if the government bans misleading price offers? Similarly, suppose Anna always pays her taxes because she overestimates the probability of being caught evading tax. Will she be any happier if we point out to her that she has been duped by the advertising of the tax authorities? Clever policy should make people happier, and behavioral economics gives us a better picture of what makes people happier.

Behavioral economics, therefore, offers a fresh approach to economic policy. Indeed, it is arguably in terms of policy prescription that behavioral economics differs most from the standard economic model. Behavioral economics has started to deliver on its promise and make a real difference, not only in economics but to people's lives.

This seems as good a point as any to also highlight that behavioral economics is not just about microeconomics and finance. It can easily seem that way given that the focus is on individual decision-making. Behavioral economics has a lot to offer across the spectrum of economics, however, including macroeconomics. This is forcefully argued by George Akerlof and Robert Shiller in their book *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Concepts such as wage and price stickiness, money illusion, fairness and social norms are fundamental to understanding how the aggregate economy works. Behavioral economics can have a big role, therefore, in informing macroeconomic policy. For instance, one thing

we have learnt is that being unemployed makes people particularly unhappy (see Chapter 10 for more). This suggests that reducing unemployment should take center stage in designing policy.

1.1.3 The different faces of behavioral economics

In suggesting, as I have, that there were four distinct elements in the rebirth of behavioral economics, it is not too surprising that behavioral economics is somewhat splintered into subtly different subfields. In this book I am going to use a very broad and encompassing notion of behavioral economics that tries to cut across any arbitrary divisions. Divisions do exist, however, and it is worth knowing something about these, and some terms used to describe them that you may come across.

I will start with the notion of **bounded rationality** that we have already come across in talking about Herbert Simon. The idea here is to recognize the constraints that people face, in terms of computational capacity, memory, information, time, and the like. We should not, for instance, assume that *Homo economicus* can do mathematical calculations that a human cannot do, or can remember more things than a human can. This sounds clear enough, but there are two quite different ways in which the idea has been put into practice.

Today the term 'bounded rationality' is commonly reserved for work in which the constraints people face are explicitly modeled. The approach is thus one of solving for what a rational person will do if they have, say, limited memory. For instance, what password should Anna use on her computer if she knows she might forget it? With this approach it is still assumed that people can be approximated by a selfish and rational *Homo economicus*; just one with a bit less memory and mathematical ability. Such an approach is prone to something called the infinite regress problem, which I will talk about in the next chapter, but it does give us an idea of how a person can optimally cope with their limitations, or bounded rationality.

This common usage of the term is somewhat removed from what Simon originally envisaged. More in keeping with his work is the idea of **simple heuristics that make us smart**. This approach starts with the idea that people use heuristics, or 'rules of thumb', to make decisions. For example, Anna might use the same password whenever she is asked to give a password; this way she is less likely to forget the password for her computer. Whether or not heuristics such as this are optimal is not really considered important. More important is to find what heuristics people do use, and to think about the consequences.

This approach most naturally fits with the work of psychologists, such as Daniel Kahneman and Amos Tversky. Heuristics tend to come with biases, however, and sometimes it's the bias that gets more of the headlines. For example, a reasonable heuristic in repetitive situations is 'Do what I did last time'. Anna, for instance, might always go grocery shopping in the same store. Sensible though this may be, it can lead to a **status quo bias** in which she fails to change her behavior 'often enough'. Maybe she keeps going to the same store even though a cheaper and better store opened nearer to her house. Such biases usually go by the name **cognitive biases**.

In exploring the connection between simple heuristics and cognitive bias it is useful to introduce the notion of **ecological rationality.** This recognizes that a heuristic can be smart or dumb depending on the environment in which it is used.

For example, Anna may use one heuristic when shopping in her local superstore and a completely different one when shopping at a bazaar in the Middle East. Each heuristic is matched to the context. Cognitive bias is most likely to result when a heuristic is used out of context or in an environment for which it was not designed. Haggling, for example, is unlikely to get you a great deal at the local supermarket.

Cognitive biases are simplest to see in well-designed experiments, and experiments have long been the most common form of research method in psychology. In economics the story is a bit different, with experiments basically unheard of until the early ones involving Vernon Smith and Reinhard Selten. Since then they have become progressively more common, but the proportion of academic economists who have run an experiment is still way below the corresponding proportion of psychologists. Where I work at the University of Kent, for example, I estimate the respective proportions at 15 percent and 100 percent; this is typical. Consequently, one still often sees the term **experimental economics** as referring to any economic research based on experiments.

As more and more economists use experimental methods, however, the term 'experimental economics' becomes ever more non-descriptive. Experiments are now being used in very different ways and in very different subject areas. Sometimes, for instance, as with the market experiments of Vernon Smith, the focus is on institutions or aggregate behavior, and there may be relatively little interest in individual behavior. Other times, as with the experiments of Reinhard Selten, there is a greater focus on individual behavior and testing theories of individual behavior. Experiments have also been run across the subject boundaries of economics from game theory to macroeconomics.

The contrasts between game theory and bounded rationality, on the one side, and experimental economics and research on cognitive biases, on the other, can be quite stark. As far as I am concerned, however, it is all behavioral economics. Good behavioral economics can involve theoretical research with no experiments. It can also involve experimental research with little or no theory. It seems natural, therefore, to split things according to subject of investigation and not method of investigation. That's how I will do things in this book: trying to mix theory and experiment. It's notable, however, that some big gaps in our knowledge still exist between theory and experiment; these are the kinds of gaps that will, hopefully, be filled by the next generation of behavioral economists.

Before I finish, a few more terms you may come across seem worth mentioning. I will start with 'economic psychology'. In principle, **economic psychology** is a subdiscipline of psychology that deals with the psychology of economic decision-making. For example, it might ask what cognitive processes were used in deciding to buy this book. The difference between behavioral economics and economic psychology is sometimes a little subtle, and I will come back to the difference shortly. But it seems fair to say that a lot of what has historically been called behavioral economics would probably be more appropriately called economic psychology. The trouble is that behavioral economics needs to draw heavily on economic psychology, and that's why the dividing line becomes a bit blurred. It's also why a lot of what I am going to talk about would probably be more appropriately called economic psychology.

The next term I want to mention is 'behavioral finance'. As the name suggests, **behavioral finance** is, in principle, a subset of behavioral economics that deals with

financial decision-making and financial behavior. I say 'in principle' because, in reality, behavioral finance has taken on something of a life of its own. As detailed more in section 1.3.2, I will cover key issues and results of behavioral finance. Indeed, many of the applications I will look at will be from behavioral finance.

Finally, I want to introduce 'neuroeconomics'. **Neuroeconomics** gets its own chapter in the book, and so probably needs little mention here. But I will mention that it is the latest development in behavioral economics, and brings neuroscience to bear on economics. It seems safe to say that neuroeconomics has contributed more to neuroscience than it has to economics, so far. For this reason it remains somewhat controversial, but only time will tell whether it can significantly advance economics. And behavioral economics is not short of controversy, as I now want to explain.

1.1.4 Debate and controversy

Yes, behavioral economics does have its controversy and disputes. Economists question whether we need behavioral economics, and whether we need economic experiments. Behavioral economists question the methods and techniques of others, such as the need for neuroeconomics. There is nothing unusual in academics disagreeing, but it can be useful to know something about the disagreements. Here are four interesting debates for you to think about, both now and while you read through the remainder of the book.

Let's start with the big question, one that I talked about a little in tracing the decline and rebirth of behavioral economics.

> Is it enough to assume people can be approximated by *Homo economicus*, or do we need psychologically grounded assumptions?

Common sense might suggest that we do need psychologically grounded assumptions. The **methodology of positive economics**, advocated by Milton Friedman, suggests things are not so obvious. The basis of this methodology is that a model and theory should be judged on its predictions and not its assumptions. Friedman (1953: 40–41) writes:

> One confusion that has been particularly rife and has done much damage is confusion about the role of 'assumptions' in economic analysis. A meaningful scientific hypothesis or theory typically asserts that certain forces are, and other forces are not, important in understanding a particular class of phenomena...
>
> Such a theory cannot be tested by comparing its 'assumptions' directly with 'reality.' Indeed, there is no meaningful way in which this can be done.
>
> Complete 'realism' is clearly unattainable, and the question whether a theory is realistic 'enough' can be settled only by seeing whether it yields predictions that are good enough for the purpose in hand or that are better than predictions from alternative theories.

That people are not like *Homo economicus* is not, therefore, reason to *assume* they are not like *Homo economicus*. We need to look at the predictions. To illustrate this distinction, suppose we see Anna go into the bookstore and are interested in whether

she will buy this book. We come up with two models of what she does; call them the 'cognitive model' and the 'choice model'. Suppose that the cognitive model better describes the thought processes that Anna will go through when deciding to buy the book. Suppose that the choice model better predicts whether or not Anna will buy the book. Which model is better?

You might say that this is a trick question, because something that better describes her thought processes should better predict what she does. Unfortunately, however, the need to use models that are tractable abstractions from a complex reality makes it common to face such a dilemma. So, you do need to make your choice!

To a psychologist, the cognitive model is surely better, because it gives a better description of what Anna was thinking; this is good economic psychology. To an economist, the choice model is surely better, because we get a more accurate economic prediction; this is good behavioral economics. It is not so obvious, therefore, that psychologically grounded assumptions are better for economics, and we will see examples in this book of models with psychological grounded assumptions doing a lot worse than the standard economic model.

Search though we might, therefore, for the 'perfect model' that combines the best elements of both the cognitive and choice model, we may not find it. Behavioral economics needs to prove itself by coming up with good economic predictions, and better predictions than that of the standard economic model.

This leads on to the second debate that I want to talk about.

> Should more emphasis be put on things that the standard economic model does well or badly?

An interesting example of this debate playing itself out came after the 'Behavioral Foundations of Economic Theory' conference, held in 1985. The conference organizers stressed 'a growing body of evidence – mainly of an experimental nature – that has documented systematic departures from the dictates of rational economic behavior' (Hogarth and Reder 1987). Vernon Smith (1991), who was present at the conference, preferred to stress 'a growing body of evidence that is consistent with the implications of rational models'! Elsewhere, Smith (1998: 105) writes that '[o]ur scientific advance is handicapped by our failure to pursue the exciting implications of the fact that things sometimes work better than we had a right to expect from our abstract interpretations of theory'.

It is clear that the argument carries on to this day. Some behavioral economists prefer to emphasize the good, and some the bad. Why no agreement? One reason is probably the distinction between assumptions and predictions. Behavioral economists from an economics background, such as Vernon Smith, have usually been the ones emphasizing what the standard economic model does well, and that's probably because they more naturally focus on the predictions, not the assumptions. Those coming from a psychology background seem more prone to emphasize what the standard economic does badly, probably because they more naturally focus on the limitations of the assumptions. If the assumptions of the standard economic model are pretty bad but the predictions are often pretty good, then we can see why there may be grounds for disagreement.

This is far from the end of the story, however, because people can also disagree on how to evaluate predictions. This stems partly from how far people are

prepared to push the standard model. As game theorist and experimental economist Ken Binmore (2008: F249) explains:

> At one end of the spectrum, there are conservative experimenters who defend traditional economic theory by looking at situations in which it predicts fairly well. At the other end of the spectrum, there are radical experimenters who seek to show that traditional economic theory does not work at all by looking at situations in which its predictions fail.

Even on relatively safe ground, however, there can be disagreement. To see why, consider this question.

> What should we conclude if the standard economic model predicts well only what experienced people do – i.e. people familiar with a task or decision?

This turns out to be a tough question to answer. The problem is that in many instances we find that the standard economic model gives a poor prediction of what happens the first time someone faces a particular situation, but a much better prediction the fourth, fifth, sixth time they face the same situation. This is the **discovered preference hypothesis**, that the standard economic model is a good predictor if people have had ample opportunity to learn from experience.

There is debate over the validity of the discovered preference hypothesis, and we shall see situations in which no amount of experience helps the standard economic model predict well. This gets us back to a 'half full' or 'half empty' debate, however. More interesting is to ask what we should conclude if the discovered preference hypothesis is correct. This will depend on whether people have ample opportunity to learn from experience in most of the things we are interested in. Arguably, they do not because there are many important things that a person does only once or a few times in their life, such as retiring, choosing a career, buying a house, and so on. We want to be able to predict what will happen in these situations too, so maybe the standard economic model is not so great after all? I shall conveniently avoid that question and finish with a much easier one.

> Should behavioral economics look to rewrite economics from a psychological perspective, or adapt the standard economic model to take account of psychological insight?

Some would probably like to start from scratch, and scrap the standard economic model. Maybe that is what you expected of behavioral economics. This is not, however, what behavioral economics is about. Behavioral economics is very much about working with the standard economic model, whether it is a good predictor or not. That's because, as mentioned earlier, it is the natural starting point. Daniel Kahneman (2003: 1449), for example, writes about his work with Amos Tversky: 'The rational-agent model was our starting point and the main source of our null hypotheses.' To quote from two other leading behavioral economists, Colin Camerer and George Loewenstein (2004: 3): 'At the core of behavioral economics is the conviction that increasing the realism of the psychological underpinnings of economic analysis will improve economics on its own terms.'

1.1.5 Too far or not far enough

Disagreement on the kinds of questions we have just discussed means that behavioral economics has no shortage of critics. Some criticize it for not being radical enough. Others criticize it for going too far! As you learn more about behavioral economics it is useful to be aware of this criticism and see how it relates back to the questions above.

Pretty much all criticism stems from the same basic concern, namely that behavioral economics is too much like the standard economic model. As we have seen, behavioral economics takes the standard economic model as its starting point. In practice, this means that behavioral models are almost always the standard economic model plus something new. That 'something new' might be a reference point, loss aversion, greater weight on the present than the future, a desire to earn no less than others, and so on. The basic assumptions of the standard economic model are retained, however.

This means that behavioral models retain many non-psychologically grounded assumptions. Put another way, behavioral economics is firmly based on the methodology of positive economics. Some suggest that a more radical approach is needed. For example, in a recent critique of behavioral economics Berg and Gigerenzer (2010) write:

> The unrealistic models now being defended [in behavioral economics] are endowed with additional parameters given psychological labels, resting on the claim that people behave as if they are solving a complicated constrained optimization problem with bounds on self-interest, willpower, or computational capacity explicitly modeled in the objective function or constraint set... To become more genuinely helpful in improving the predictive accuracy and descriptive realism of economic models, more attention to decision process will be required.

This brings us back to the debate on the methodology of positive economics: will more attention to the decision process result in improved predictive accuracy? I cannot answer that question. I will, however, use this as an opportunity to warn against some of the hyperbole that surrounds behavioral economics – and a lot of hyperbole does surround behavioral economics! Although behavioral economics is revolutionizing economics, it is not as radical a departure from the standard economic model as some would have you believe.

Let us turn now to those who criticize behavioral economics for going too far. If behavioral economics is the standard economic model plus something else, then it is no surprise that we can get a better fit. The danger is one of **over-fitting**. The basic critique goes something like this. The behavioral economist observes behavior that 'should not happen', according to the standard economic model, such as Anna borrowing at a high interest rate on a credit card while simultaneously saving for retirement at a low interest rate. We then come up with something that can 'explain' this anomaly, such as an impulsive desire to buy expensive handbags while out shopping. We then congratulate ourselves on a better fit.

There are three related problems here. (1) Can we really be sure there is an anomaly? Maybe Anna will learn over time to not borrow so much on her credit card. (2) If there is an anomaly, then maybe other things could explain it. How can we be

sure an impulsive desire to buy handbags is the real cause? (3) It is easy to come up with ad hoc explanations for observed anomalies. More difficult is to come up with **novel testable predictions** that are subsequently proved correct. Those who criticize behavioral economics for going too far would say that the extent of anomalies is overblown, the derived models are too ad hoc and there is not enough focus on novel testable predictions. For instance, in an article critical of the methods being used in behavioral economics, Binmore and Shaked (2010) write:

> [T]he scientific gold standard is prediction. It is perfectly acceptable to propose a theory that fits existing experimental data and then use the data to calibrate the parameters of the model. But, before using the theory in applied work, the vital next step is to state the proposed domain of application of the theory and to make specific predictions that can be tested with data that was used neither in formulating the theory nor in calibrating its parameters...
>
> [S]ome events are easier to predict than others. A good theory will successfully predict events that are difficult to predict. Choosing events that a theory is to predict in a manner that favors the theory – especially if done after the test experiment has been run – is unacceptable.

A further, related concern is the **external validity** of laboratory experiments. Behavioral economics draws heavily on lab experiments when looking for the existence of anomalies, comparing models and testing predictions. Given that lab experiments are typically very context-specific, however, the concern is that we cannot infer much from them about real behavior. In a review of lab experiments Levitt and List (2007) write:

> Perhaps the most fundamental question in experimental economics is whether the findings from the lab are likely to provide reliable inferences outside of the laboratory. [...] [W]e argue that lab experiments generally exhibit a special type of scrutiny, a context that places extreme emphasis on the process by which decisions and allocations are reached, and a particular selection mechanism for participants. In contrast, many real-world markets are typified by a different type of scrutiny, little focus on process, and very different forms of self-selection of participants.

The issue here is not whether we should do experiments; the broad consensus is that experiments are an important tool of economic analysis. The issue is the balance between experiment and theory. Experiments are inevitably going to be somewhat artificial. For instance, experiments done in the field, rather than in the lab, are arguably even more context-specific (a point I will return to in Chapter 8 when looking at development economics). The key thing to guard against, therefore, is becoming over-reliant on experiments. Some would argue that behavioral economics is too quick to draw conclusions from lab experiments. This can easily lead to ad hoc models and overblown claims of anomalies.

What to make of all this criticism? It should definitely be taken seriously. There is a danger that behavioral economics is not focused enough on decision processes, that it is too ad hoc and that it is over-reliant on experiments. Some 'findings' of behavioral economics will inevitably not stand the test of time. Taken as a whole, however, I think

it is clear that behavioral economics has delivered. In particular, it has come up with many novel testable predictions that have proved correct, not only in the lab but also in real applications. Indeed, over the last few decades behavioral economics has come under closer scrutiny than any other area of economics, and it has survived the test well. So, take note of the criticism but also be reassured that you can learn a lot from behavioral economics!

1.2 Some background on behavioral economics methods

We have already seen that behavioral economics involves both theory and experiment. In this section I want to briefly sketch a little more about the methods of behavioral economics and, in particular, give some background on economic experiments. Before doing so I want to make clear that my objective is not to explain how to run experiments, or to perform statistical tests on the experimental data. These warrant books on their own, and there are lots of good books out there (see the further reading at the end of the chapter). All I want to do is explain enough that you will be able to follow the rest of this book. That means, in particular, getting straight some terminology that I am going to use.

Rather than talk abstractly about experiments and theory I thought it would be more interesting to talk through some research that I did (with Federica Alberti and Anna Stepanova) while writing the first edition of the book. This will give me a chance to introduce all the concepts you need at this stage. Other concepts will be introduced as and when needed or relevant in 'research methods boxes' that you will find throughout the book.

I will use the term **study** for a particular piece of research. The objective of our study was to see whether the amount of money people are endowed with in a threshold public good game has any effect on their ability to coordinate. That might not make much sense at this point, but do not worry, because it should make more sense shortly. A study may be part of a more general **project**. Our ongoing project is to see how people can coordinate better in threshold public good games. The main part of this study involved running experiments, and so I will talk about that first.

1.2.1 Some background on experiments

To run an experiment we need people who are willing to take part. Those who do are given the title **subject** or participant. In order to recruit subjects we sent e-mails and placed adverts, and eventually got the 120 volunteers we needed. Most of these were students at the University of Kent. Each subject was asked to come to a particular **experimental session**, and we ran six sessions in all, with 20 people invited to come to each session. (We actually invited a few more, because the chance that 20 students will turn up on time in the right place is clearly quite small!) All the sessions took place in a computer lab at the university.

When the subjects arrived, after a brief introduction they were each allocated to a computer and signed a consent form. After this there was no talking until the

experiment ended, and a subject could see only their own computer screen. Things were, therefore, completely anonymous, and subjects were reassured this would be the case. Beside their computer each subject would find an instruction sheet, which read something like this:

> In this experiment you will make decisions, and earn an amount of money that depends on what you and others choose. The money will be given to you at the end of the experiment in an envelope. Only you will know how much money you earned.
>
> You have been organised into groups of five. Each group will consist of the same five people for the dluration of the experiment. The experiment will last for 25 rounds. In each period you will be required to make a decision, and your total earnings will depend on your decisions in all rounds.
>
> At the beginning of every round you, and all other members of your group, will receive 55 tokens. Each of you must decide, on your own, how many of the 55 tokens to allocate to a group account.
>
> If the total number of tokens allocated to the group account is 125 or more then you will each receive an additional 50 tokens.
>
> If the total number of tokens allocated to the group account is less than 125 then you will receive no additional tokens but will get back any tokens you allocated to the group account.
>
> So, at the end of the round:
>
> If the total number of tokens allocated to the group account ≥ 125, your earnings = initial 55 tokens – tokens allocated to group account + additional 50 tokens.
>
> If the total number of tokens allocated to the group account ← 125, your earnings = initial 55 tokens.
>
> At the end of the session, you will be asked to fill in a short questionnaire. You will be paid in cash the total amount that you earned for all rounds in the session plus £2. Each token will be worth 0.5 pence.

Once all the subjects had read the instructions the experiment would commence, with everything taking place on computer (using an experimental economics program called z-Tree). So, the first thing that would happen is each subject was asked to input how many tokens they wanted to allocate to the group account. To illustrate, in group one, subject one input 25, subject two input zero and subjects three to five input 50, 25 and 10. The total number allocated was, therefore, 110, which was short of the 125 target. The subjects were told this and told their earnings were 55. That was the end of round one.

Something similar happened in rounds two to 25. This did not take long, and so the experiment was over after around 30 to 45 minutes. Once the last round had finished we worked out how much money each subject had earned and paid them in cash. The average subject earned around £10, which is not bad for 30 minutes.

One thing I want to highlight is that, in an experiment, subjects are typically asked to do the same thing several times. In this case they were asked to play the same threshold public good game 25 times. There are various reasons for doing this, of which the main one is to see how people learn with experience. For instance, the group I have just talked about were short of the target in round one, and we might

wonder how they will respond. In fact, the total number of tokens allocated in round two went up to 161, well above the target. As you may have seen by now, the term **round** keeps track of how far subjects have progressed through the experiment. (The word 'period' can also be used, but I will try and avoid this because we will use 'period' for something slightly different.)

Hopefully, this gives you some idea of what a threshold public good game is (we are going to talk about them in more detail in Chapter 6). Hopefully, it also gives you some idea of what an experiment is, and what it's like to take part in one. But what were we hoping to get out of this experiment?

The primary thing of interest in a threshold public good game is whether group members give enough to the group account to reach the threshold of 125. If this all seems a little abstract, then imagine five housemates who want to buy a new TV that costs $125. Between them they need to find enough to buy the TV or they go without. Broadly speaking, it is in their interests to reach the threshold, but there are good reasons why they may fall short. From prior research we know groups can be quite unsuccessful at reaching the threshold.

The specific question that motivated our study was whether groups will be more successful at achieving the threshold if they have more tokens to start with. The instructions above correspond to the case when subjects got 55 tokens, which you can think of as each housemate having $55 of spending money. What if they have $30 or $70; will they be any more or less successful at reaching the threshold?

This is relatively easy to test with experiments, because we can just give some groups 55 tokens to start with, some 30, some 70, and record how often groups reach the threshold. That's basically what we did. The term **experimental treatment** is used to distinguish these different versions of the experiment. We had five treatments, but I will just focus on the three of them that I have already mentioned. These are summarized in Figure 1.1, and correspond to subjects having 55, 30 or 70 tokens.

It is common to have one treatment called the **baseline treatment**, which serves as a reference point for comparison. In our study the treatment in which subjects got 55 tokens was the baseline. We chose this as the baseline because it had already been used many times by other researchers. We thus have an expectation of what should happen in this treatment, and can more easily relate our results to the results of other researchers. This is invaluable in making good scientific progress, and so the baseline typically corresponds to a treatment that has already been done in previous studies.

One thing I want to make clear is the distinction between 'experimental session' and 'experimental treatment'. **Experimental session** refers to a particular instance in time and space when an experiment was run; for instance, we had a 2 p.m. session on Friday 29 January. **Experimental treatment** refers to a specific version of the experiment used; for instance, we had the baseline treatment with subjects getting 55 tokens. Session and treatment are quite different. You might have, as we did, one session with multiple treatments, or multiple sessions with the same treatment. In session four, for example, we had four groups, two of which were in the high treatment (so subjects in these groups got 70 tokens) and two were in the low treatment (so subjects got 30 tokens).

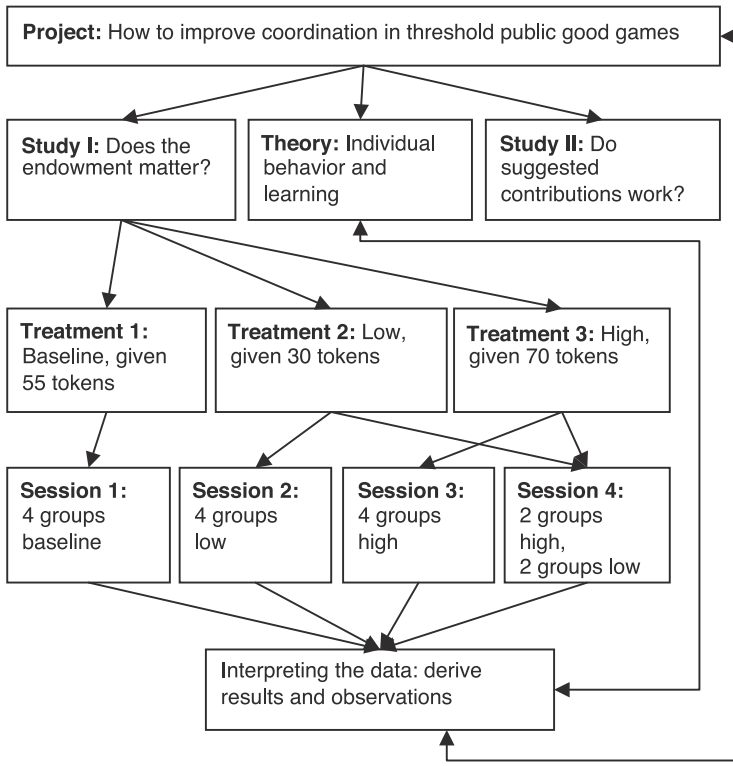

Figure 1.1 An overview of the project.

Another thing I want to mention is 'deception'. In psychology it is not unknown to see **deception** in experiments, meaning, basically, the experimenter lies to subjects about what they are doing. (Afterwards they are told that they were lied to.) This is a definite no-no in economics, in which we are not allowed to deceive subjects in any way. So, if the instructions say there will be 25 rounds, or say one token is worth 0.5p, then it really is so. That subjects cannot be deceived does not, however, mean that they have to be told everything. Indeed, a basic requirement of most experiments is that they be **single blind**, meaning subjects are not biased by knowing too much. For instance, we did not want the subjects in our experiment to know that the objective of our study was to vary the number of tokens given. If they knew this, it might have biased how they behaved. So, while subjects cannot be deceived, they can be given a minimum amount of information.

Before moving on to look at theory I should point out that the experiment I have described above is not untypical, but experiments do vary considerably. For example, some experiments are done with pen and paper rather than via computer. Some do not involve interaction between subjects in any way. Others allow subjects to communicate with each other, possibly via computer chat. Some use schoolchildren as subjects, or experienced financial traders. Some are done outside the lab in a supermarket, or at a sports show. There are many other possibilities, and I will talk about some of these as and when we get to them.

1.2.2 Some background on theory

As we have already noted, experiments are not much use without some theory behind them. So, I also want to give a little background information about basic theoretical concepts that we are going to need. Let's start by thinking what someone called Edward does when he is in the experiment described above. To keep things simple, imagine for now that there is only one round.

All Edward has to do is to decide how many tokens to allocate to the group account. He can choose any integer between 0 and 55. We call this the **set of actions** or **set of possible options**, and he needs to pick one action or choose one option from this set. An **action** thus describes what he does – for instance, 'Allocate 20 tokens to the group account'.

I now want to contrast action with strategy. You can think of a **strategy** as a description of how Edward decides what action to choose. To explain what this means, it's useful to talk briefly about a second study we did. In this study, one person in the group could suggest how much others should allocate to the group account, before others choose how much they will allocate to the group. If Edward is not the one giving the suggestion, then his set of actions is the same as before; he needs to choose an integer between 0 and 55. His strategy can, however, be conditional on what was suggested. For instance, his strategy might be 'Allocate as many tokens as was suggested'. With this strategy his action will be 'Allocate 20 tokens' if the suggestion is to allocate 20 tokens, but 'Allocate 25 tokens' if the suggestion was to allocate 25, and so on.

Actions and strategies are, therefore, different, and you might have noticed that actions are observable but strategies need not be. For instance, we observe how many tokens Edward gave but, generally speaking, will not know how many tokens he would have given if, say, the suggestion had been different.

The next thing to introduce is the payoff function. Edward's **payoff function** says what payoff he will get for any possible combination of actions he and everyone else uses. For instance, if everyone allocates 20 to the group account, his payoff is 55. There are various different notations for the payoff function. Given that it is typically more convenient to write payoffs in terms of strategies rather than actions, we can write $u_E(s_1, s_2, s_3, s_4, s_5)$ for Edward's payoff as a function of the strategies of the five group members, including him. We can also write $u_E(s_E, s_{-E})$ for his payoff as a function of his strategy and the strategy of others. The payoff function in the experiment we are looking at is

$$u_E(s_E, s_{-E}) = \begin{cases} 55 - s_E + 50 & \text{if } s_1 + s_2 + s_3 + s_4 + s_5 \geq 125 \\ 55 & \text{if } s_1 + s_2 + s_3 + s_4 + s_5 < 125 \end{cases}$$

For simplicity, it's not uncommon, as I have just done, to substitute action for strategy when working out the payoff function.

It is useful to be aware of a slightly troublesome thing about the interpretation of payoff. Easiest is to think of payoff as the amount of money earned. Given that subjects were paid money for every token they earned during the experiment, this is the approach I have taken above when writing down the payoff function; payoff equals money. The problem with this approach is that we typically want to equate payoff with utility, a subjective measure of happiness. In Chapter 10 we'll see that money and happiness are quite different things. Doubling the amount of money earned need not

lead to a doubling in the amount of increased happiness. So, why not just write down a payoff function that reflects utility? Measuring utility is difficult, while measuring money is simple. The norm, therefore, is to wish the problem away and equate utility with money. I shall do the same here. I will, at least, point out that this is done mainly for convenience rather than any good economic reason! Be reassured, though, that I shall return to this issue at various points during the book.

If we have a group of people and know the set of actions and strategies and payoff function of each person, then we have a **game**. Recall that subjects in an experiment typically play a game several times. What happens in any one round is called the **stage game** or **base game** or **constituent game**. So, the stage game in our experiment was the threshold public good game. The whole experiment – i.e. the 25 rounds – constitutes the **game**. The reason for making this distinction is that someone may want to make their action in round two conditional on what happened in round one, and so on. For instance, a strategy might be 'If we reached the threshold in round one, I will allocate the same number of tokens to the group account in round two'.

Normally, the distinction between stage game and game is not so important, but it can be useful if the stage game changes at some point during the experiment. For example, we shall see experiments in which subjects play one stage game for ten rounds and then a different stage game for a subsequent ten rounds.

I should also briefly mention that a game can be between **teams** of people as well as between individuals. In this case each team chooses an action and strategy, with the members of a team collectively deciding what they should do. We'll return to this distinction in Chapter 6.

The final concepts I want to talk about are possible outcomes of a game. Two particularly important concepts are Pareto efficiency and Nash equilibrium. Let's do Pareto efficiency first. We say that an outcome is **Pareto efficient** if no one could be made better off without making someone worse off. To see why this is interesting, imagine an outcome that is not Pareto efficient. Then we would be able to do something that made at least one person better off and no one worse off. This looks like a good thing to do. In the threshold public good game, if everyone were to allocate 25 tokens to the group account, everyone would get a payoff of 80 tokens. This is Pareto efficient. If everyone were to allocate zero tokens to the group account, they would get a payoff of 55. This is not Pareto efficient, because everyone is getting fewer, 55, than they could have – e.g. 80. Pareto efficiency is thus one way to formalize the idea of something that is desirable versus something that is not.

Contrast this with Nash equilibrium. A **Nash equilibrium** is a strategy for each person such that no one could change their strategy and increase their own payoff. Again, imagine this were not the case. Then someone can change their strategy and increase their payoff. We might expect them to do so. A formulation you might come across is that Nash equilibrium requires

$$u_i(s_i, s_{-i}) \geq u_i(s, s_{-i})$$

for any person called i, where s_i is the strategy they do choose and s is any other strategy they might have chosen. In the threshold public good game, everyone allocating 25 tokens to the group account is a Nash equilibrium. Everyone allocating 24 tokens to the group account is not a Nash equilibrium, because the total will fall short of the threshold by five, and, if one person were to give 29 rather than 24, their payoff would increase.

Interestingly, everyone allocating zero tokens to the group account is also a Nash equilibrium, because no one person on their own can do anything to reach the threshold. This illustrates that a Nash equilibrium need not be Pareto efficient. It's also true that a Pareto-efficient outcome need not be a Nash equilibrium (see if you can find an example in the threshold public good game). A Nash equilibrium and Pareto efficiency are, thus, two very different things.

I will finish this section by making a connection with some of the ideas I talked about in the last section. In many ways Nash equilibrium is the epitome of the standard economic model. It is pretty much always assumed in standard economic analysis that behavior can be captured by Nash equilibrium. But there can be many Nash equilibria. I have already shown you that there are at least two in the threshold public good game, and there are actually many more. The standard economic model has nothing to say about which of these is supposed to happen. That renders it pretty useless (for me at least). We need to know which equilibrium, if any, is most likely to happen, and that, as I have discussed above, inevitably leads us into behavioral economics. Experiments can help us to understand what happens when people play such a game. This can give us the clues to improve our theoretical ability to model economic behavior, and better understand how people reason and learn.

1.3 How to use this book

To put things in some context I should give a brief preview of the rest of the book. Before I do that, I want to talk a little about my approach in writing this book.

As I said in the preface, my objective in writing the book was to explain the basic principles, ideas and results of behavioral economics and also illustrate through applications how these can help us understand the world around us. This will be reflected in the make-up of each chapter. The first half of each chapter will be dedicated to explaining the basic principles, ideas and results, while the second half will give various applications. Hopefully, this will give a nice mix of theory and practice.

The approach I have taken throughout is to pick a few experiments and theoretical papers that I think best illustrate the more general lessons that have been learnt. I like this approach because it gives me a chance to look in a bit more detail at some studies and give a better flavor of what behavioral economics is all about. One downside is that it can give the impression that our knowledge hangs on the results of one or two experiments or theoretical examples (that you might spot a potential flaw in). Sometimes this impression would not be far wrong, but for the most part it is a long way off. Most of the ideas I will talk about in this book are robust. In terms of experiments, that primarily means the results have been replicated in many other similar studies. In terms of theoretical examples, this means the examples can be generalized to give the same conclusions. The further reading suggestions, at the end of each chapter, as well as the online resources, give you the chance to check this and explore things in a bit more depth.

To make the book as accessible as possible, I have tried to keep algebra and notation to the minimum. Algebra, though, is part and parcel of behavioral economics, so a few equations are unavoidable to get a full grasp of the main issues. I also do not like discriminating against those who do like a bit of algebra. So, sometimes you will see '[Extra]' written, which is a cue that the rest of the paragraph will contain some algebra but can be skipped without missing anything.

In addition, each chapter contains theoretical examples. Essentially, this is about illustrating a model by giving a specific example that you can work through. We shall look at Anna shopping for breakfast cereal, Alan buying car insurance, Maria deciding when to go to the movies, and so on. In order to focus on behavioral economics, rather than algebra, I will not give a blow-by-blow account of where the numbers in these examples come from. The more advanced and confident you are, the more I would encourage you to try working through the examples yourself. In case you get stuck the companion website contains a lot more detail on all the examples.

Finally, I have little time for the politically correct issue of whether to use 'he or she' or 'she or he'. So, each chapter will be based on one or two characters, and the language will reflect the sex of those characters. The fact that the characters have the same names as my wife and colleagues is pure coincidence.

1.3.1 Chapter previews

The book will be split into three further parts, of which the next part will be by far the biggest, looking at what behavior we observe and how we can model it. It begins, in Chapter 2, with an introduction to heuristics and context effects. The main message I want you to get from this chapter is that people can be influenced in a systematic way by the context or setting in which choices are made. This might seem obvious enough, but it is quite a big step away from the standard economic model.

In Chapter 3 I look at choice with risk, and in Chapter 4 at choice over time. The basic approach in both chapters will be the same: to start with a simple and standard way to model behavior, namely expected utility and exponential discounting, and then to propose various modifications to the model so that we can better fit the behavior we observe. This will bring us to some of the more important concepts in behavioral economics, such as prospect theory and time inconsistency.

A theme throughout Chapters 2 to 4 is the importance of reference dependence and context. In particular, we shall see that behavior often depends on whether outcomes are perceived as a loss or gain relative to some reference point. We shall also see that the reference point a person has in mind can change a lot depending on the context.

In Chapters 5 and 6 the focus switches to how people use information. In Chapter 5 the focus is on how people can learn from the information around them, including what they get from other people. In Chapter 6 the focus is on learning and behavior in interactive situations, when it's necessary to predict the behavior of others. I think the basic message that comes out of these two chapters is that people are clever and do learn, but also struggle to cope with more complex situations. We are going to see, for example, some worrying biases in how people learn from new information.

Chapter 7 continues the theme of Chapter 6 in looking at interactive situations, but has a very different focus. The focus here is on looking at whether people are selfish and, if not, how we can model this. We shall see that it is necessary to take account of people's preferences for fair and just outcomes, and there are various ways we can do this. One of the more important lessons from this chapter will be the heterogeneity of behavior. We'll see that some people give, some not, and what is considered fair or just can vary from one person to the next and one context to the next.

The next part of the book will consist of two chapters that look at the origins of behavior. That is, we shall try to answer questions such as why people are risk-averse, are loss-averse, value fairness, and so on. Chapter 8 focuses on the role of evolution and culture in shaping economic behavior. Chapter 9 looks at neuroeconomics and recent attempts to see what is going on in the brain as people make economic decisions. We'll see that an evolutionary and neuroscience perspective provides some fascinating insights into economic behavior, and really does tell us something about why people are risk-averse and the like.

In the final part of the book my focus will primarily be on whether behavioral economics can help us do something useful in improving the happiness of people. In Chapter 10 I look at what makes people happy and whether they know what makes them happy. We are going to see that people do not always do the thing that would make them happiest, and may be willing to pre-commit or delegate choice to others. It's clear, however, that people still desire some choice or self-determination.

Chapter 11 follows this up by looking at how behavioral economics can help inform policy. I look at both the design of institutions and policies aimed more at individual incentives. Making the world a better place is always easy on paper, but I am going to argue that behavioral economics does give us some nice tools to improve policy for the better.

1.3.2 Behavioral finance

As I have already mentioned, behavioral finance has taken on something of a life of its own alongside behavioral economics. Given that this is a textbook on behavioral economics, rather than behavioral finance, there will naturally not be a comprehensive treatment of behavioral finance. The key ideas of behavioral finance are, however, covered. For those interested in behavioral finance I shall briefly point out three key things to watch out for.

Of primary importance is the efficiency of markets. This is a theme in Chapter 2 (see sections 2.5 and 2.7). Particularly relevant for finance is the efficiency of markets in the presence of uncertainty. This is covered in Chapter 5 (see section 5.6), including a discussion of bubbles and busts. Chapter 3 (see section 3.6) focuses on some specific finance 'puzzles', such as the equity premium puzzle.

An often neglected area of finance is the interaction between traders. Two important concepts for behavioral finance are information cascades and beauty contests. Information cascades help us study how one trader may learn from another (see section 5.3). The beauty contest helps us study the incentives of traders (see sections 6.1 to 6.4).

The saving and investment decisions of individuals will be a recurring theme throughout the book. See in particular section 2.8 on the life cycle hypothesis, sections 2.9 and 10.6 on retirement saving and section 4.5 on borrowing.

1.4 Further reading

Many have given good introductions to what behavioral economics is all about: Mullainathan and Thaler (2000), Camerer (2003) and Camerer and Loewenstein (2004) are a small selection. The excellent review articles by Conlisk (1996), Rabin

(1998) and DellaVigna (2009) are also worth a mention at this point. For more on the methodology of experiments, the books by Cassar and Friedman (2004) and Bardsley and co-authors (2010) are recommended; see also Samuelson (2006) and Croson and Gächter (2010). I should also mention Fischbacher (2007) for the details about z-Tree. For more information on ecological rationality, see Smith (2003, 2007), and on the discovered preference hypothesis, see Loomes, Starmer and Sugden (2003). For more on the history of behavioral and experimental economics, see Roth (1995), Sent (2004), Bruni and Sugden (2007) and Svorenčík and Maas (2015).

1.5 Online material

On the companion website you will find access to a timeline history of behavioral economics, flashcards, quizzes and answers to the review questions.

1.6 Review questions

1.1 Why is the standard economic model a good thing, and why is it a bad thing, in trying to understand economic behavior?
1.2 Why do we need to run economic experiments?
1.3 Why does a heuristic usually come hand in hand with a cognitive bias? Should we emphasize how clever people are for having good heuristics, or how dumb they are for being biased?
1.4 Why does it make sense to mix up experimental treatments and sessions – i.e. to have multiple treatments in each session and multiple sessions for each treatment?
1.5 Is it good that experiments usually involve students as subjects?
1.6 What are the objectives of behavioral economics?
1.7 What are the objectives of studying the standard economic model?
1.8 What is the discovered preference hypothesis and why does it matter?
1.9 Why is there a danger that behavioral economics will over-fit experimental data?
1.10 What is the methodology of positive economics and why is it controversial?
1.11 Why should economics pay attention to ecological rationality?
1.12 What are the main differences between *Homo economicus* and *Homo sapiens*?
1.13 What is bounded rationality?



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Part II

Economic behavior



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Chapter 2

Simple heuristics for complex choices

> Rule no. 1 is never lose money. Rule no. 2 is never forget rule number one.
>
> Warren Buffet

Even the most innocuous of economic choices are in principle very complicated. For example, consider a shopper in a grocery store looking at rows of breakfast cereal and deciding which one to buy. Should she buy the cereal she usually buys? Should she try a new cereal the store has just introduced? Should she buy the cereal on special offer? Will the cereal she usually buys be on special offer next week? Will it be cheaper in another store? Should she be tempted by the cereal with the chance to win a holiday in the Caribbean?

Clearly, most of us do not spend much time considering all these issues. Indeed, most of us simply buy the cereal that we usually buy. That way we can make a quick decision that will probably keep us happy. This is an example of a heuristic. A **heuristic** is any 'rule of thumb' or simple rule of behavior by which a person solves a problem. The shopper can solve their problem of what cereal to buy with the heuristic 'Buy what I usually do'. Almost all the economic decisions we make are based on such heuristics; otherwise, life would get far too complicated. This makes it important for us to know how heuristics work and what their consequences can be.

Heuristics, and the biases they give rise to, are going to be a recurring theme throughout the book, and are the primary focus of this chapter. I will start by giving some examples, but the main thing I want to do is provide some structure for how to think about the consequences of heuristics. In doing this we shall come across some of the key concepts and ideas in behavioral economics.

2.1 Utility and search

I will start by taking up the story of a shopper called Anna in a grocery store deciding what breakfast cereal to buy. There is a large selection of potential choices, all with different characteristics, but we will narrow things down to the four listed in Table 2.1. How can she decide what to buy?

The standard way of thinking about this in economics is to assume that Anna has a **utility function** that says how much utility she gets from particular combinations of money and goods. In this context we would write $u(x, TQ, HQ)$ as the utility she gets from having money x and a cereal with taste quality TQ and health quality HQ.

Table 2.1 Cereals for sale, and their characteristics, with 1 = low, 2 = medium and 3 = high.

Product	Price	Taste quality	Health quality
Budget	$1	1	1
Nutty	$3	2	2
Honey	$4	3	2
Superior	$6	3	3

Table 2.2 The utility of each cereal if initial wealth is $100.

Choice	Wealth	Taste quality	Health quality	Utility
No cereal	$100	0	0	200
Budget	$99	1	1	202
Nutty	$97	2	2	203
Honey	$96	3	2	204
Superior	$94	3	3	203

One important point to note is that, in the standard description, utility should be a function of money wealth, and so when evaluating each choice we need to focus on how much money Anna will have after buying the cereal. To illustrate, Table 2.2 works through an example in which she initially has $100 and her utility function is

$$u(x, TQ, HQ) = 20\sqrt{x} + 2TQ + HQ.$$

She can buy no cereal, keep wealth $100 and have utility 200, or pay $1 for Budget, have wealth $99 and utility 202, and so on.

You might be wondering what it means to say that Anna's utility is 200, or 202. The interpretation we should give to utility is crucial but ultimately very tricky, and so, on the basis that we have plenty of other things to worry about at this stage, I am going to dodge that question for now. In Chapter 10 we will look at it in some detail. In the meantime you can think of utility as a general measure of happiness or satisfaction. So, more utility is better, and Anna wants to choose the cereal with highest utility. We are also not going to get far unless we think in terms of **cardinal utility**. What this means is that the differences in utility should mean something. So, we can say things such as: 'Honey is better than Budget by the same amount as Budget is better than no cereal.'

In the example, we can see that Honey offers the highest utility and so looks like the best choice. It does so because it offers the best trade-off of quality for price. Anna is willing to pay the extra $1 to $4 that Honey costs over other choices in order to improve the taste and health quality, but is not willing to pay a further $2 to get the highest quality. We have, therefore, a prediction of what Anna should buy: she should buy Honey.

This is fine, if Anna knows what maximizes her utility. Realistically, however, she probably will not know. Maybe she has never tried Budget or Superior, or maybe

she did try them once but has forgotten what they tasted like, her preferences have changed or the manufacturers have subsequently improved the quality.

This lack of knowledge is crucial, and means that it is not enough for us to say that Anna should do the thing that maximizes her utility. She does not know what that is, and so we need to delve a little deeper. Two things that become relevant when we do so are search and choice arbitrariness. I will look at each in turn.

2.1.1 How to search

If Anna does not know the quality of goods, or her utility function, then she can gather more information in order to become better informed; we call this search. A **search heuristic** specifies what Anna should do in order to become better informed. There are lots of possible search heuristics, and I shall look at five to give you some idea of what they, and heuristics in general, look like.

The most obvious search heuristic is 'Try everything'. For example, Anna could try a different cereal every week until she has tried them all, and then subsequently buy the one she liked the most. This means she will end up knowing a lot about cereals. Potentially, however, the process will be costly. To see why, suppose that in the first week she tries Honey and can tell that she likes it a lot. If she sticks to her heuristic then in subsequent weeks she will have to buy and try cereals that are not going to give her as much utility. She would have been better off just sticking with Honey (if you will excuse the pun).

In this case, search proves costly in terms of **forgone utility**. Search can also be costly in terms of time and money. A good search heuristic needs to trade off the benefits of acquiring more information with these costs. Characterizing optimal, or good, search algorithms has a long history in mathematics, computer science and economics. Optimal search algorithms are typically very complicated, however, so we need to think what search heuristics people could realistically use that come close to the optimum. Three such heuristics are satisficing, elimination by aspects and directed cognition.

The basic idea behind **satisficing** is that a person sets some **aspiration level** for what they are looking for, and continues to search until they find something above the aspiration level. For example, Anna may decide she wants something that tastes good and is reasonably healthy. This determines her aspiration level, and she will keep on trying cereals until she tries Honey or Superior. Satisficing relaxes the objective from finding the optimal choice to merely finding a choice that is good enough. This means a person may not end up with the best, but they will end up with something relatively good, while avoiding the costs of excessive search. If Anna, for instance, tries Superior before Honey, then she will never get to know how good Honey is, but she might also avoid knowing how bad Budget is.

How close satisficing comes to the optimum will depend on the aspiration level. This is where satisficing become slightly more tricky, because it is not trivial what the aspiration level should be, or how it should change. If, for example, Anna's aspiration is to find a cereal costing less than $7, then she may end up with Budget; the aspiration level looks too pessimistic. If her aspiration is to find a very tasty and healthy cereal for less than $5, then she will be disappointed because no such cereal exists; the aspiration level looks too optimistic. Anna will need, therefore, to set the aspiration

level appropriately and revise it as she goes along; it's just not clear how exactly she will do so.

A heuristic that partly addresses this problem is that of directed cognition. The idea behind **directed cognition** is that a person treats each chance to gather information as if it is the last such chance before they have to make a choice. Typically, it will not actually be their last chance. Directed cognition simplifies Anna's task, because she does not need to forward-plan. To illustrate, suppose she knows only the characteristics of Nutty. Using directed cognition, she should ask herself: 'Shall I try one alternative cereal, and, if so, which one?' This is a much simpler question than would be needed with forward-planning. With forward-planning there are many more permutations that need to be considered, such as: 'Shall I try Budget and, if that tastes nice, try Honey and, if not, try Superior, and then Honey?'

The final search heuristic I want to look at, for now, is that of elimination by aspects. The basic idea of **elimination by aspects** is to consider the aspects of possible choices one by one and sequentially eliminate choices that fall below some aspiration level. For example, if Anna's aspirations are to buy a medium taste and health quality cereal for under $5, then on the price aspect she would eliminate Superior, and on the quality aspects eliminate Budget. This would leave a choice between Honey and Nutty.

Elimination by aspects is different from the previous two heuristics in that it compares across aspects, such as price, rather than across choices, such as Honey or Nutty. Conceptually it is simpler to compare across aspects, because there is likely to be a simple ordering from best (e.g. least expensive) to worst (most expensive). The problem is that comparing across aspects presupposes that the person has information about all the possible choices. With aspects such as price and health quality this is plausible, because the prices and ingredients will be displayed on the box. With an aspect such as taste quality it is difficult to know the differences without trying them all. Elimination by aspects can take us only so far, therefore, in explaining search, but it does offer vital clues on how a person can choose what to try next and what to not try at all.

The three heuristics that we have discussed should be seen as complementary rather than alternatives, as summarized in Table 2.3. Satisficing, for instance, suggests how long Anna should search, but offers no clues as to what she should try. Directed cognition has less to say about how long to search, but can offer clues as what to try next. Together, therefore, they give us a picture of how people might search. In general, no one heuristic is ever going to be perfect, and so we can expect people to

Table 2.3 Five search heuristics with a very brief statement of how they can help with search.

Heuristic	What it does well	What it does not do so well
Try them all	Make the person well informed	Minimize the cost of search
Satisficing	Say when to stop search	Say what choice to try next
Directed cognition	Suggest what choice to try next	Give a forward-looking plan of search
Elimination by aspects	Say what choices not to try	Say when to stop searching
Search for *x* minutes	Give certainty how long search will last	React to success or failure in search

use different combinations of heuristics for different tasks. Indeed, Anna may have a heuristic to say what heuristic to use (see Research methods 2.1).

❖ RESEARCH METHODS 2.1 THE INFINITE REGRESS PROBLEM

We began by talking about Anna's problem of 'what cereal to choose' and suggested she should choose the cereal that maximized her utility. Having noted that she will probably not know what cereal maximizes her utility, we got the additional problem: 'how to find out what cereal maximizes utility'. This has to be a harder problem to solve than the original one, but we have not finished yet. That is because there are lots of ways to search for information about preferences, so we get a further additional problem: 'how to find out, how to find out what cereal maximizes utility'. I could go on: 'how to find out, how to find out, how to find out...', and so on.

More generally, finding the optimal answer to a problem involves finding the optimal way to find the optimal answer, and the optimal way to find the optimal way, and so on. This is known as the **infinite regress problem**. Clearly, if Anna, or anyone else, is going to make a decision, they need some way to break the loop. Heuristics, emotions and intuition (which we shall get to shortly) are one way to do this because they provide simple rules to make decisions. That answer is not entirely satisfactory, however, because Anna will need heuristics to choose what heuristic to use, and then we are back in the cycle again!

In short, the problem is not going to go away, and so we need to be aware of it when we try to model behavior. Behavioral economists have primarily focused on basic heuristics – i.e. heuristics designed to solve a specific problem. Some attention, but much less, has been given to heuristics for choosing heuristics, and we will get to them in later chapters. The problem of how people choose heuristics for choosing heuristics is ignored, but in Chapter 8 I will suggest that evolution might do that job.

To get a feel for how people search we can look at the results of a study by Gabaix and co-authors (2006). In the study, subjects were asked to choose between eight options with the payoff of each option the sum of ten numbers. One of the numbers was visible, but to find the value of each of the other nine numbers the subjects had to click on the screen. To relate this back to our earlier example, imagine that the price is visible, but to find the health quality Anna has to pick up the box and look up the ingredients. Software called Mouselab can track what information subjects choose to look up and gives us some idea how subjects searched (see Research methods 6.1).

There are lots of things we can look at in the data in terms of how much information subjects choose to look up, what sequence they choose to look it up in, and

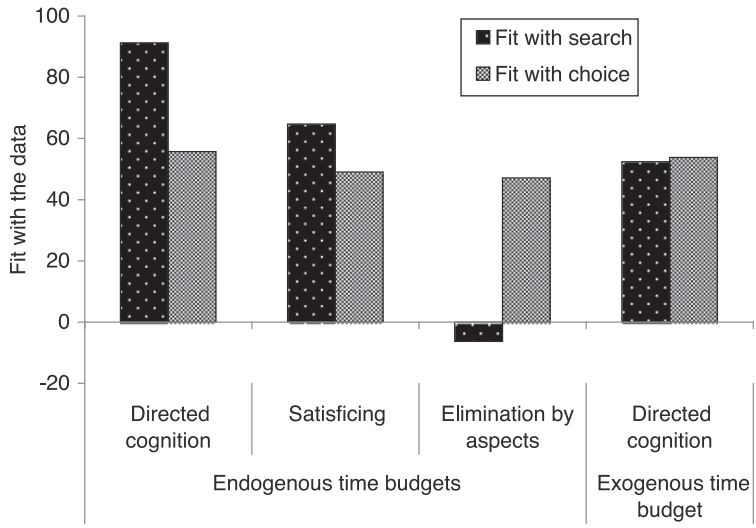

Figure 2.1 How well search heuristics fit the observed search and choices in a search experiment. The fit is measured relative to what could have been expected by chance. Directed cognition and satisficing give the best fit.
Source: Gabaix *et al.* (2006).

so on. One way to see if subjects were using a particular heuristic is to compare how much information subjects should have looked up about each choice, if they were using, say, directed cognition, and how much information they did look up. This measures how well the search heuristic fits observed search behavior. A second thing we can do is to see how well the search heuristic fits the final choice made. This measures how well the search heuristic fits observed choice behavior.

Figure 2.1 summarizes the results of the study. We see that directed cognition does a good job of fitting how subjects search, while elimination by aspects does less well. In terms of fitting choice, all algorithms do roughly as well as each other and the choice is what we would have expected around 50 percent of the time. A 50 percent success rate at predicting what subjects will choose is much better than we could have done by guessing (remember there were eight possible choices), but still not as high as we might like. It looks as if we are missing something.

One thing that becomes apparent is that subjects were not as selective in what information they looked for as the three heuristics suggested they should have been. In different ways they looked up both too much and too little. First, the too much: if one choice ranks poorly on the visible aspect, then it pays to not search for other information about that choice. For example, if Superior is too expensive, then there is little point in looking at the health quality. Subjects tended to look up more information than we might expect in such instances (see Research methods 2.2). Next, the too little: it does not matter how long the search has gone on; if the search has not yet uncovered something useful, it pays to carry on searching. Subjects seemed to stop searching if they had looked up a lot of information, however, even if that information had not been very informative.

❖ RESEARCH METHODS 2.2 HOW LONG SHOULD AN EXPERIMENT TAKE?

One slightly troublesome issue in running experiments is how much time to give subjects to make decisions.

On a practical level, this issue is important if we do not want some subjects to be delayed by others. In particular, it immediately becomes apparent when running an experiment that some people take a lot longer to make decisions than others! If the experiment can proceed only when each subject has made a decision, this can leave some subjects twiddling their thumbs for a long while. On the one hand, this may be good if it means some subjects spend a bit more time thinking about the experiment. On the other hand, it can lead to boredom and less thinking about the experiment.

On a theoretical level, it may be important how much time subjects have, if we are interested in the consequences of time pressure. Time is very relevant in search and so important to take into account. Gabaix and co-authors did this by comparing exogenous and endogenous time budgets. In an exogenous time budget treatment, subjects are given a fixed amount of time during which they can look at a particular set of eight choices. In an endogenous time budget treatment, subjects are given 25 minutes to look at as many sets of eight choices as they wish.

This distinction turns out to be potentially informative. In particular, subjects were more selective about what information they looked up with an endogenous than exogenous budget. So, the low success at predicting search with an exogenous time budget (see Figure 2.1) may reflect subjects having too much time. Optimally, they should have found out more information about the current best-looking choices, but maybe they had already made up their mind and so were filling in time.

Putting these observations together suggests a fifth search heuristic: a person decides how much time to spend on search and then searches for that long. This may mean that Anna would search too long, if she initially tried something she liked, or search too little, if she has yet to try anything she really likes. But it does mean search will last a definite length of time, and this may be a useful thing.

There are many more possible search heuristics and many more studies (see the suggested reading) that look at how people search. We will come back to search elsewhere in the book, however, notably Chapter 5, and so I will leave the issue for now. The main thing I want you to see at this stage is how people can use simple (or not so simple) heuristics to resolve complex search problems. Hopefully, this gives you some idea of how people can search effectively, and gives some feel for what heuristics in general look like.

Now that we have looked at some heuristics it is worth briefly expanding on their role in decision-making. The ideal heuristic is simple and effective: simple in the sense that it can be performed relatively easily, particularly when compared to what

would be required to make an optimal decision; effective in the sense that its use typically results in a choice close enough to the optimal decision that extra effort is not justified. The best heuristic strikes the right balance between simplicity and effectiveness.

It is important to recognize that the right balance will depend on the context, because in some instances simplicity may be all-important while in others it is vital to make as good as choice as possible. For example, Anna is probably not going to spend too much time thinking about breakfast cereals, but she might want to think a lot more about which university to attend or what house to buy. Moreover, some heuristics may be better suited to some environments than others. It is useful, therefore, to have a large set of complementary heuristics from which we can draw on to solve particular problems (although see Research methods 2.1). The term **adaptive toolbox** nicely captures the idea that people have such a set of heuristics at their disposal.

One issue the previous discussion leaves open is the size of trade-off between simplicity and effectiveness. Indeed, what does it mean to say that a choice is 'close enough' to the optimum? This is a matter of judgment rather than exact science. It is interesting, though, to reflect on the different perspectives that behavioral economists can take on this issue (see also section 1.1.4). Some emphasize that very simple heuristics can lead to optimal or very nearly optimal choices. So, there is not much of a trade-off between simplicity and effectiveness. This view is particularly associated with the idea of **fast-and-frugal heuristics**, as well as that of ecological rationality (see Research methods 2.7). Others take a different view in recognizing that simplicity is typically going to come with a cost. So, a simple heuristic may end up generating choices well away from the optimum. This is sometimes called the **heuristics-and-biases framework**, in which the 'biases' term recognizes the potential for suboptimal choice. In the next section we shall see some of the biases that can arise from search heuristics.

2.1.2 Choice arbitrariness

In the process of search, Anna is going to face some fairly arbitrary choices. That's because she does not yet know what maximizes her utility but still has to choose something. For example, she may have narrowed her choice to Nutty or Honey, and there is no real reason to try one ahead of the other. Which one to try first? Her choice will be arbitrary, and she might as well toss a coin to decide. Arbitrary does not, however, have to mean random. For instance, Anna might be attracted by the bright red packaging of Honey or the '50 percent off' sticker on Nutty, or she might choose Budget because she just saw it advertised on TV. In each of these cases choice is systematic. The crucial thing is that choice is influenced by factors that just happened to be like that and could have been different: this is **choice arbitrariness**. Let's look at some examples.

I will start by looking at the difference between conflicting and non-conflicting choices. We say a set of choices are **conflicting** if one choice is better on one aspect and a different choice better on some other aspect. For example, Budget is better on the price aspect but Superior is better on the health quality aspect, and so these are conflicting choices. A set of choices are **non-conflicting** if one choice is better

on all aspects. For example, if Superior were on sale for $0.50 then there would be a non-conflicting choice.

To illustrate the potential consequences of conflicting versus non-conflicting choice, consider this example from a study by Tversky and Shafir (1992). Subjects were asked to imagine that they want to buy a CD player, and walk past a shop with a one-day clearance sale. Some subjects were given the conflicting choices of a Sony player for $99 and a top-of-the-range Aiwa player for $169; this is a conflicting choice because the Sony is cheaper while the Aiwa is better quality. Some were given the non-conflicting choice of the Sony player for $99 or an inferior Aiwa player for $105; this is a non-conflicting choice because the Sony is better in terms of price and quality. Others were just given the option of the Sony player for $99. All subjects were asked whether they would buy one of the players or wait and learn more about the models. Figure 2.2 summarizes the choices made.

As we would expect, more people buy the Sony when the choice is non-conflicting than when it is conflicting. The more interesting comparison is that between a non-conflicting choice and no choice. Crucially, we see that more people choose the Sony when the choice is non-conflicting than when there is no choice at all. This latter observation violates the **regularity condition** of choice that an increase in the number of available options should not increase the share buying a particular option. It seems that the presence of an inferior option increased the likelihood of buying the Sony.

What we have just seen suggests that one alternative can look more or less desirable depending on what it is compared to. A slightly different possibility is that particular *aspects* of an alternative can look more or less desirable depending on what they are compared to. To illustrate, let's suppose Anna has narrowed her choice to either Nutty or Honey and there are only three cereals on display: Nutty, Honey

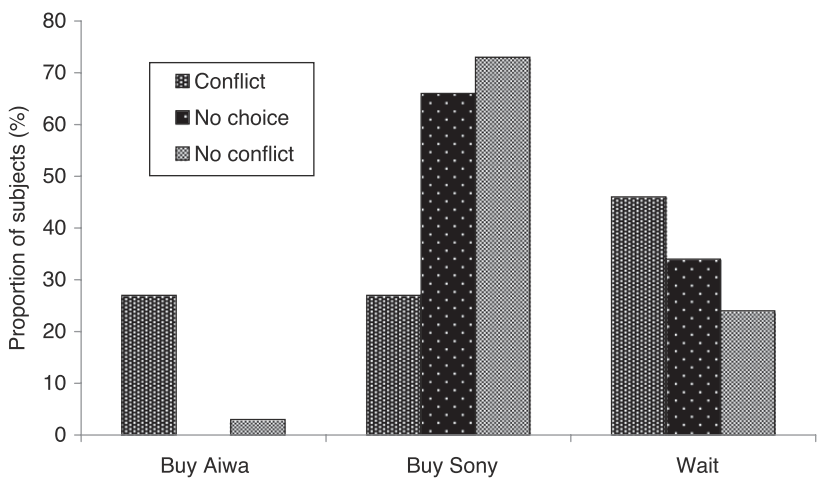

Figure 2.2 The decision that subjects made when the choice is conflicting or non-conflicting, compared to when there is no choice. Subjects were more likely to wait, the more conflicting the choice.
Source: Shafir, Simonson and Tversky (1993).

and one of Budget or Superior. Honey looks relatively cheap when contrasted with Superior but relatively expensive when contrasted with Budget. She might, therefore, be more likely to buy Honey when it is contrasted with Superior than with Budget. This would be an example of trade-off contrast. The **trade-off contrast hypothesis** is that a product with a desirable quality will appear cheaper if contrasted with a product for which that desirable quality is more expensive. In this example, the high taste and health quality of Honey appear cheaper if contrasted with Superior.

A study by Simonson and Tversky (1992) shows the relevance of trade-off contrast. In one part of the study, subjects were asked to choose between coupons and cash, with each coupon redeemable for books or CDs at local stores. In a background stage, subjects were exposed to choices in which a coupon cost either $15 or $5. After this they were all exposed to the same choices with a coupon costing $10. Figure 2.3 summarizes the results. The first choice offered to subjects was $47 and five coupons, or $37 and six coupons (the extra coupon costing $10). We see that those exposed to a background stage in which coupons cost $15 were more likely to choose the extra coupon, while those exposed to a background stage in which coupons cost $5 were not. This is consistent with the trade-off contrast hypothesis; the $10 coupon looks cheap, or expensive, depending on whether coupons previously cost $15 or $5.

Let us now go back to the scenario in which there are two or three cereals on display out of Budget, Nutty and Honey. Budget has the advantage of being cheap, Honey has the advantage of being tasty but Nutty strikes a good compromise. Maybe, therefore, Anna will buy Nutty because it's 'in the middle'. If true, this means she should be more likely to buy Nutty when all three cereals are on display rather than just two. This would be an example of **extremeness aversion with compromise**.

There is another, more subtle, alternative. It may be that the presence of Honey and Budget on display emphasizes the importance of quality and results in Anna 'going for the best'. In this case, she might go for the lowest price, and lowest quality, when

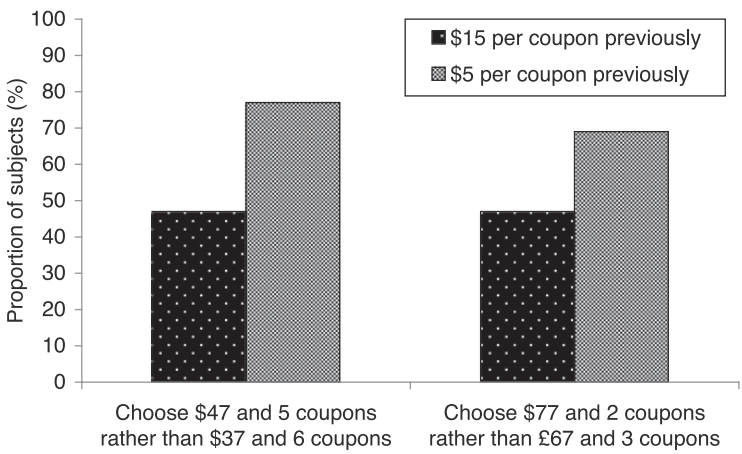

Figure 2.3 Whether subjects prefer an extra coupon to $10 depends on whether they were exposed to a background stage when coupons cost $15 or $5.
Source: Simonson and Tversky (1992).

Table 2.4 Examples of how Anna's choice may be influenced by the three psychological effects.

Psychological effect	Examples of what Anna may choose
Trade-off contrast	If choices are Budget, Nutty and Honey, buys Nutty If choices are Nutty, Honey and Superior, buys Honey
Extremeness aversion with compromise	If choices are Nutty and Honey, buys Honey If choices are Budget and Nutty, buys Budget If choices are Budget, Nutty and Honey, buys Nutty
Extremeness aversion with polarization	If choices are Nutty and Honey, buys Nutty If choices are Budget and Nutty, buys Budget If choices are Budget, Nutty and Honey, buys Honey

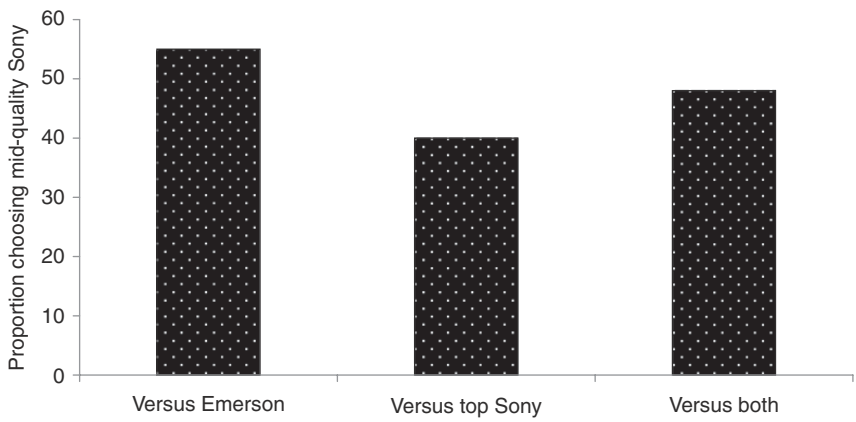

Figure 2.4 The proportion choosing a mid-level Sony radio-cassette player when the alternative was an Emerson player, a top-quality Sony player or both the Emerson and Sony. We observe extremeness aversion with polarization.
Source: Simonson and Tversky (1992).

comparing Nutty and Honey or Budget and Nutty. It is only when all three cereals are on display that she focuses on the differences in quality and chooses Honey. So, the presence of Budget on display increases the likelihood that Anna will buy Honey. This would be an example of **extremeness aversion with polarization**. Table 2.4 distinguishes these possible effects with examples.

To give an example of extremeness aversion we can return to the study of Simonson and Tversky, and a different scenario. Subjects were given information about three radio-cassette players: a mid-quality Emerson, mid-quality Sony and top-quality Sony. In Figure 2.4 we see the proportion that chose the mid-quality Sony. Adding the top-quality Sony decreases the proportion choosing the mid-quality Sony compared to the Emerson. Also, adding the Emerson increases the proportion choosing the mid-quality Sony compared to the top-quality Sony. This is an example of extremeness aversion with polarization. It is also a second example violating the regularity condition.

In the next example we push extremeness aversion a bit further. First, imagine that Anna asks herself: 'What cereal shall I buy?' She might say Superior, because it is more extreme, higher quality. Now imagine she asks herself: 'What cereal shall I not buy?' She might say Superior, because it is more extreme, more expensive. In both cases the fact that Superior is more extreme makes it more salient. That means it is the first thing she considers, and so might be the first thing she accepts and/or rejects. This can lead to a **choose–reject discrepancy**, when a person chooses the same thing whether asked what they want or what they do not want.

Here is an example from a study by Shafir, Simonson and Tversky (1993). Subjects were given two choices for a vacation. Spot A was relatively neutral: average weather, average beaches, average nightlife, etc. Spot B was more extreme: lots of sunshine, gorgeous beaches, very strong winds, no nightlife. Some subjects were asked which they would prefer to book. Other subjects were told they have a provisional booking at each location and now have to decide which to cancel. In Figure 2.5 we see that most subjects preferred spot B over A but were equally split on whether to cancel A or B. Overall, the proportion of choices for spot B exceeded 100 percent. This basically means that some people would both prefer spot B and want to cancel spot B!

To give a final example, suppose that on going into the grocery store Anna is stopped and asked to do a survey. One question they ask is whether she would be willing to pay $8 for a new cereal they are trialing. After that, Superior may look relatively cheap at $6. If the question had said $4, then maybe Superior would look expensive. This is an example of an **anchoring effect**, when a person's choice is influenced by some prior cue or anchor. This might look like a trade-off contrast. In trade-off contrast, however, two or more different products are compared, while, with the anchoring effect, a person's thoughts on a particular product are influenced by some prior event.

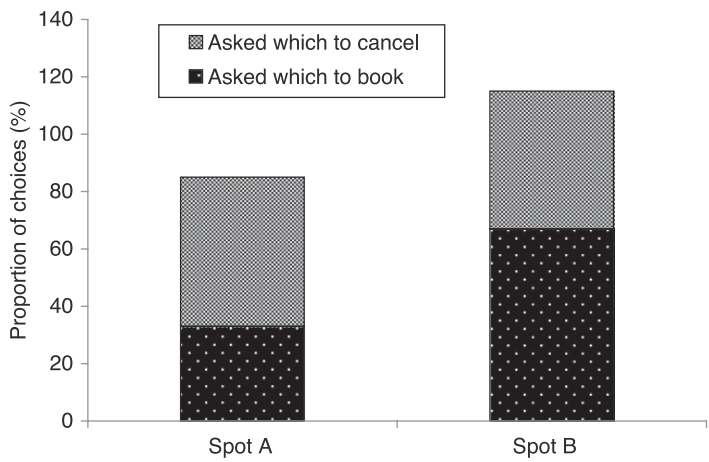

Figure 2.5 The proportion who chose spot A and B to book and to cancel. Both columns should add up to 100 percent because if, say, 30 percent prefer spot A, 70 percent should reject spot A.
Source: Shafir, Simonson and Tversky (1993).

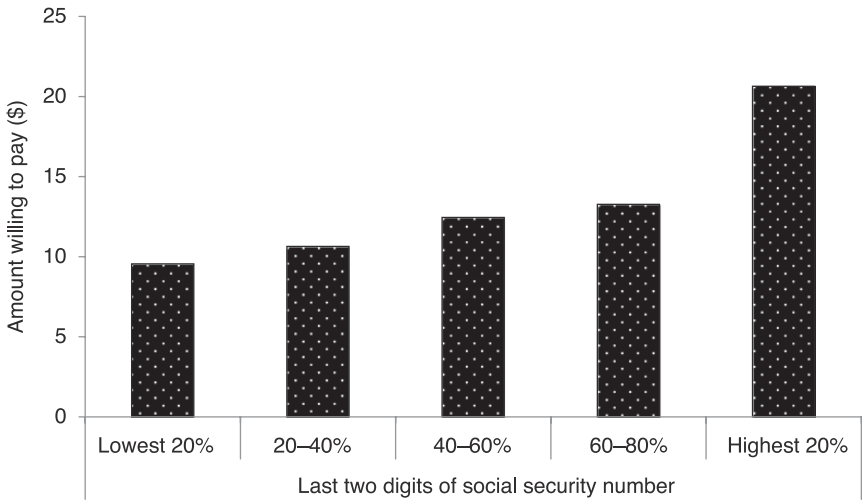

Figure 2.6 The amount that subjects were willing to pay for a box of chocolates was higher for those with a higher social security number.
Source: Ariely, Loewenstein and Prelec (2003).

To illustrate how the anchoring effect can happen, we shall look at part of a study by Ariely, Loewenstein and Prelec (2003). Subjects were first asked whether they would buy a box of Belgian chocolates, and some other items, for more than the last two digits of their social security number. For example, if the last two digits are 25 they were asked whether they would pay more than $25. They were then asked how much they would be willing to pay. The last two digits of a social security number are random, but did matter, as we see in Figure 2.6. On average, those asked whether they were willing to pay, say, $55 subsequently said they were willing to pay a higher price than those who were asked to pay, say, $15 (but see Research methods 2.3).

❖ RESEARCH METHODS 2.3 CAN YOU REPLICATE IT?

The study by Ariely, Loewenstein and Prelec, of which the box of chocolates experiment is just one small part, is one of the best known and most widely cited in behavioral economics. Why? In large part it is because an anchoring effect should not exist according to the standard economic model, and so the study raises new questions. The popularity of the study is also, inevitably, helped by there being some 'fun' and 'quirky' results. Who would have thought the last two digits of someone's social security number influences how much they are willing to pay for a box of chocolates?

Herein, however, we potentially have a problem. The profession and society like revolutionary, quirky results! To see why this is a problem, suppose

that 100 research teams did the exact same box of chocolates experiment and only one team found any effect. The team that found an effect can publish their results in a top academic journal and get the publicity. The 99 that found no effect will throw their results away, because the 'social security number does not influence willingness to pay for chocolates' result is hardly exciting.

Interestingly, the anchoring effect results are now so widely accepted that showing 'social security number does not influence willingness to pay for chocolates' has become revolutionary and quirky! A recent study by Fudenberg, Levine and Maniadis (2012) reports a series of experiments showing no evidence of any anchoring effect. In this case, though, we do not know if there were 99 research teams that did the same experiment and found evidence of an anchoring effect.

To be really confident of experimental results we need independent **replication** of findings and the publication of those findings. There are no incentives within the profession for this to happen, however. Running replication studies is a dead-end career path, and top journals will not publish them. The simple truth, therefore, is that some of the experimental results I will discuss in this book are likely to be flukes rather than representative of the truth! Do not be too disheartened, however, because knowledge generally goes on an upward path (see, e.g., Research methods 2.4 and 3.4). There are enough studies, for instance, for us to be confident that anchoring effects are a reality. What is less clear is how big the effects are.

The five psychological effects we have looked at are the consequences of choice heuristics, such as 'Pick the one in the middle' or 'Pick the most extreme'. They result in choice arbitrariness. Put another way, they cause **context effects**, a general name I shall give to any external factors, like the other choices on offer that influence choice. Recall that this all comes about because people are unlikely to know what maximizes their utility. We should therefore expect context effects in just about any economic choice a person ever makes. This means choice arbitrariness and context effects are important – important enough that I devote the next section to trying to understand them in more detail.

Before doing that I want to make one point clear. Some would have you believe that things such as trade-off contrast and extremeness aversion are evidence of people not being rational and not being like *Homo economicus*. This is not true. In a complicated world where there are lots of decisions to make it may be optimal to 'Pick the one in the middle' or 'Pick the most salient' or 'Be influenced by the other choices on offer'. We are not, therefore, going to spend time wondering whether such things are evidence for or against the standard economic model. Indeed, even the more bizarre effects, such as choice being influenced by the last two digits of a social security number, can be consistent with rational behavior. Such effects just mean that the heuristics that people use are not always well adapted to some of the things experimenters have us do in the lab!

2.2 Mental accounting and framing

What I want to do now is question why context effects exist: why is it that external factors can influence the choice someone makes? A good starting point is to focus on a subset of context effects called framing effects.

Framing effects are when essentially equivalent descriptions of the same thing lead to different choices. That some people would choose to go to spot B for their vacation but also choose to not go to spot B for their vacation is an example of a framing effect: what people choose is dependent on the way the question is asked, or framed. If we can explain why people choose different things when offered the same set of choices, it should be simple to see why they may choose different things when offered a slightly different set of choices.

To understand why we observe framing (and context) effects, I want you to look at the simple classification of cognitive processes summarized in Figure 2.7. (In Chapter 9 we will revisit this classification to see how accurate it is, and see that, while not perfect, it is more than good enough for our purposes.) The classification suggests that, when we initially see something, perception and intuition kick in automatically to give us impressions of what we are looking at. This process happens spontaneously and the person has no or very little control over it. Anna, for example, may make a choice purely based on intuition. More generally, however, particularly for the kinds of choices of interest in economics, she might use reasoning to think through options and make a more deliberate, informed judgment. In this case she will use reasoning processes.

Now, here's the key thing we learn from this classification: even if Anna does use reasoning, her initial perception and intuition will inevitably influence the starting point for her subsequent reasoning. Initial perceptions and intuition will, therefore, matter. How a choice is framed will likely affect perception and intuition, which will in turn affect reasoning. To see this in action, recall the example of trade-off contrast. If someone was led to think a coupon was worth $15, then being asked to pay $10 will be perceived as a good deal and maybe one worth taking. If they were led to think it was worth $5, then it looks like a bad deal.

That context and framing influence perception and intuition, which influences reasoning, is one of the most important ideas in behavioral economics. *Homo economicus* only ever uses reasoning. So recognizing the role played by perception and intuition gives us a completely different perspective on why people behave the way

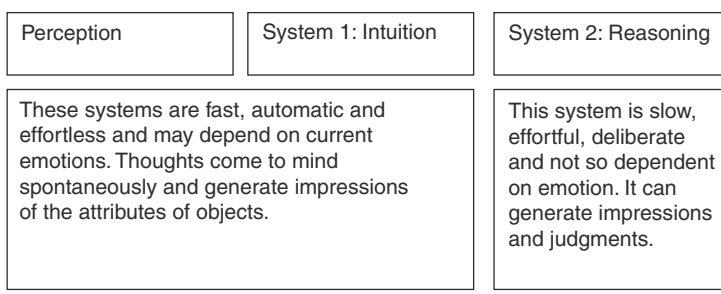

Figure 2.7 A representation of different brain processes.

they do compared with the standard economic model. Furthermore, this explanation of context and framing effects means that they are inevitable, and not just a quirk of some experiment. Every time a person makes a choice, that choice has to be framed in a particular way, and how it is framed will likely affect perception, intuition, reasoning and the choice made. If we want to understand economic behavior, we therefore need to understand framing effects. This, as I shall now explain, means we need to understand reference points.

2.2.1 Reference-dependent utility

Some attributes of an object are more **accessible** than others when we first perceive it. Such **natural assessments** of an object include size, distance, loudness, temperature, similarity and whether it is good or bad. Importantly, a natural assessment will usually be a relative rather than absolute one. It is far more natural for us to say what is bigger, longer, louder, hotter and better, without knowing the exact volume, length, temperature, etc. To be able to judge relative magnitude we need some standard of comparison, and this is called the **reference point** or **reference level**.

To illustrate why the distinction between absolute and relative can be important for economic behavior, consider this question, taken from Kahneman (2003):

> Two persons get their monthly report from a broker. Carol is told that her wealth went from \$4 million to \$3 million. Amanda is told that her wealth went from \$1 million to \$1.1 million. Who is happier?

As we have already discussed, the standard in economics is to measure utility using a utility function $u(x)$, where x is money and, the more money there is, the higher the utility. The answer, therefore, should be that Carol is happier because $u(\$3\ million) > u(\$1.1\ million)$. Things do not seem quite this simple, however, because many of us may be tempted to think that Amanda would be the happier. She might be the happier because her wealth is bigger than before, while Carol's is smaller than before.

This example suggests that we need to take seriously the distinction between relative and absolute measures. The simplest way to do this is to suppose some reference point about which outcomes are measured. So, we can use a utility function or **value function** $v(x - r)$ that measures the utility of outcome x relative to some reference point r. If the reference point is $r = 0$, then the value function will coincide with the original formulation, but if r is different – say, wealth last year – then we get something different. Setting r as wealth last year we see that Carol would get $v(-\$1\ million)$ and Amanda would get $v(\$0.1\ million)$. With this formulation it seems that Amanda should be happier, because $v(\$0.1\ million) > v(-\$1\ million)$.

Now Amanda is happier. Things still seem a little incomplete, however, because Carol does have more money than Amanda. A more general formulation is the **reference-dependent utility function**,

$$u^r(x) = \eta u(x) + v(x - r),$$

where η is some parameter and r is the reference point. Total utility is now a weighted sum of utility from total wealth and the utility from a relative gain or loss. With this

we can capture Carol being potentially happier because she has more wealth and potentially less happy because her wealth has decreased.

Judging things relative to a reference point raises the important distinction between gains and losses. Intuitively, there is something distinctly different about being above the reference point, a gain, or being below the reference point, a loss. A large amount of evidence, which we shall get to shortly, backs this up and points towards something called loss aversion. Someone is **loss-averse** if a loss causes a bigger fall in utility than a similar-sized gain causes an increase in utility. In terms of the value function, loss aversion implies that $v(-g) < -v(g)$ for any $g > 0$. Basically, losses are worse than gains are good.

We are going to see that reference dependence and loss aversion have big implications for economic behavior. Before we look at this in more detail I want to briefly preview why reference dependence is so important.

At first glance it might look as though we have just rewritten the utility function to include an extra parameter r called the reference point. If this were the case, then economists could carry on modeling utility and utility maximization the way they traditionally have done and not a great deal would need to change. This is not how it is! That's because the reference point a person uses when making a particular decision can be relatively arbitrary and will likely be determined by perception and intuition. In other words, the reference point can depend on the context, causing important framing and context effects. This is why reference dependence is more than just adding a parameter to the utility function. The ways in which the reference point can depend on context will be a recurring theme in the rest of this chapter and those to follow.

2.2.2 The endowment effect

One important consequence of reference dependence and loss aversion is the endowment effect. To explain what this is, we shall look at a study by List (2004). The study was conducted on the floor of a sports card show in the United States. Subjects were recruited throughout the day and completed a short survey before being asked to choose whether they would like to take a coffee mug or chocolate bar as reward for doing the survey. Both the mug and chocolate bar retailed at around $6. The key thing about the experiment was to vary the perception of ownership. Some people were physically given a mug and told they could swap for a chocolate bar, some were given the chocolate bar and told they could swap for a mug, others were either given both or neither and asked to choose whether they wanted the mug or chocolate bar.

Figure 2.8 details the proportion of subjects who chose the chocolate. Look first at the non-dealers, who were people attending the show. Here we see the same as has been observed in many other experiments: those given the chocolate bar do not want to swap for a mug, and those given a mug do not want to swap for a chocolate bar! It is highly unlikely that this is just an accident of giving chocolate to those who like chocolate and mugs to those who like coffee. Instead, it seems that people are biased by physically having the good in their hand – a framing effect. This is clear from the 'both' and 'neither' cases, when approximately 50 percent choose the chocolate bar.

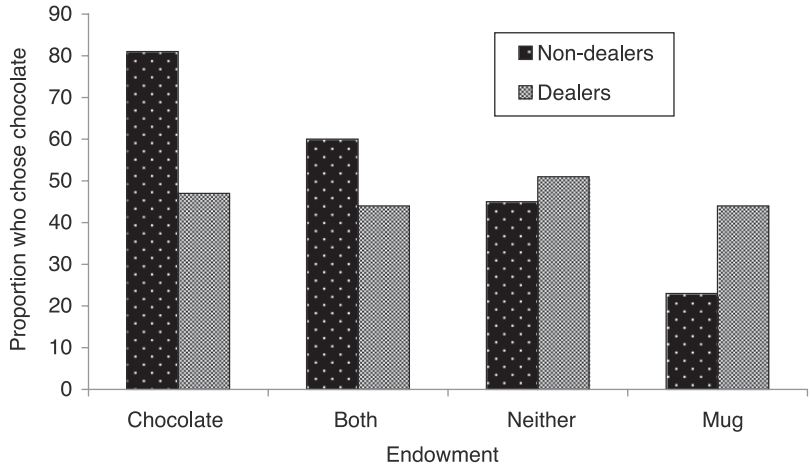

Figure 2.8 Whether subjects chose chocolate depended on the endowment. Those given chocolate were more likely to choose chocolate and those given a mug to choose the mug.
Source: List (2004).

That people value more highly goods that they have some ownership over is called the **endowment effect**. If the endowment effect shows up in people who randomly possess for a few seconds a mug they probably do not want, then economists need to take it seriously, because ownership and the trading of goods are at the heart of economics.

One way to explain the endowment effect is through reference dependence and loss aversion. If Anna has the coffee mug in her hand, then she may feel that she already owns that coffee mug. To give it away would be a loss. By contrast, if she does not have the chocolate bar in her hand then she feels no sense of ownership over it. To get the chocolate would be a gain. We have suggested that losses are weighted more heavily than gains. If so, Anna would be reluctant to swap the mug in her hand for a chocolate bar, even if she would have marginally preferred a chocolate bar.

Being overly possessive of a mug does not seem particularly useful. In some situations, however, the endowment effect may be beneficial, as we shall see when we talk about housing and saving. In situations when it is not beneficial, we should expect more experienced people to have learnt to avoid the effect. This is what we generally do observe. To illustrate, look at what dealers, people with booths at the show, did when offered a mug. Dealers had more market experience in trading (sports cards) than non-dealers, and showed no signs of being biased.

Interestingly, however, people may learn to redirect the endowment effect, rather than suppress it. This follows from an appreciation of opportunity cost. To someone familiar with the concept of opportunity cost, to keep the coffee mug in your hand is a loss, because of the value that can be got from swapping the mug. The distinction between what is a loss and what is a gain can thus be blurred. In this way the endowment effect is lessened and we get our first suggestion that the reference point depends on the person and context.

2.2.3 Willingness to pay or accept

One place the endowment effect shows up is contingent valuation. **Contingent valuation** is an often used technique to get people to reveal how much they value things such as health, safety or the environment. Two basic methods of contingent valuation are willingness to pay (WTP) and willingness to accept (WTA). In the first, **willingness to pay**, people are asked how much money they would be willing to pay to get an extra unit of some good; that is, they are asked how much they would pay to consume $c + 1$ rather than c units of the good. In the second, **willingness to accept**, people are asked how much money they would be willing to accept for having one unit less of the good; so, to get a like-for-like comparison, they are asked how much they would accept to consume c rather than $c + 1$ units. In both cases we get a valuation on the $c + 1$st unit of the good.

Given what we know of loss aversion and the endowment effect, we should not be surprised that WTA and WTP can lead to different valuations. WTP makes the person think of losing money and they may not be willing to pay much – i.e. lose a lot of money – to get an extra unit of the good. WTA, by contrast, makes the person think of losing the good, and they may want a lot of money to compensate for getting less of the good. This is consistent with what we observe. The valuation of a good based on willingness to accept typically exceeds one based on willingness to pay, and this difference can be large. This is a **violation of procedural invariance**, in that the answer we get to a question depends on the procedure we use to find the answer, a specific form of framing effect.

To illustrate, Figure 2.9 provides some data from a study by Bateman and co-authors (2005), in which subjects were asked to put a contingent valuation on ten chocolates. For now, just focus on the WTA and WTP. The WTA valuation is around 60 percent more than the WTP valuation. This difference is a problem for anyone who

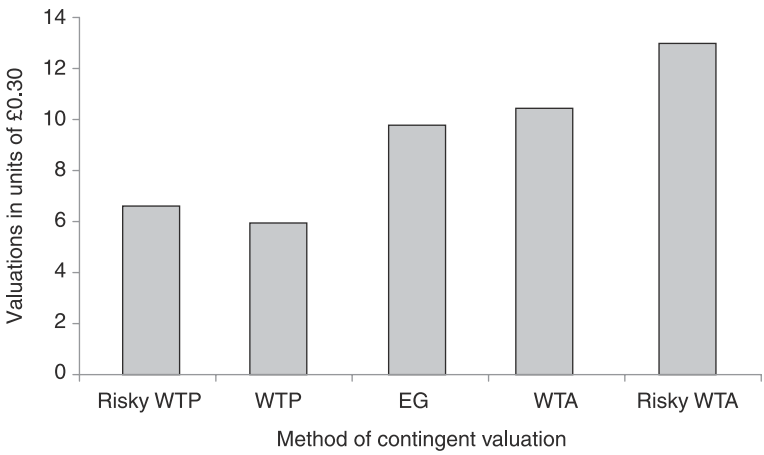

Figure 2.9 The valuation put on ten chocolates. The WTA exceeds the WTP but not the EG.
Source: Bateman *et al.* (2005).

wants to obtain a reliable estimate of the value that people put on goods. Should they use the WTP or WTA measure?

It also raises some fundamental questions about how people value goods. That's because when talking about the endowment effect it is natural to have in mind physical ownership of a good. So, we think of the subject owning a mug or owning a box of chocolates, if only for a few seconds. The WTA and WTP measures, however, get subjects to think about potentially owning a good. This seems enough to trigger an endowment effect and shift of the reference point.

One place where this becomes relevant is in thinking through the **'no loss in buying' hypothesis**. The idea behind this hypothesis is that, when people are deciding whether to buy a good, they mentally deduct money from their current income and then decide whether to gain the good or (re)gain the money they would spend buying the good. For example, if Anna usually buys Honey cereal for $4, then she feels no loss in spending $4 buying Honey because she expected to pay $4. We can rephrase this in terms of the reference point. Is the reference point current income, and so buying Honey involves a loss of $4, or is it current income minus the amount needed to buy Honey, meaning that not buying the good involves regaining $4?

For me, it is more natural to think of the latter. For example, I might say that 'Anna's reference point for Honey is $4'. What I really mean is that Anna expects to spend $4 on Honey, and so she has mentally deducted that from her wealth to give a reference point of wealth minus $4. The study by Bateman and co-authors, however (see Research methods 2.4), put the 'no loss in buying' hypothesis to the test and finds little support for it. What to make of this result?

The 'no loss in buying' hypothesis essentially assumes that people take account of the opportunity cost of money, as we suggested experienced traders might do to avoid the endowment effect. This suggests that experience is important. If people are inexperienced in a particular task, which is likely when finding the WTP and WTA of goods that the person rarely if ever buys, then buying may be interpreted as a loss. By contrast, if people are more experienced in a task, such as Anna buying Honey every week, then buying may not be interpreted as a loss, unless the price is more than they expected.

The interesting thing this illustrates is that it can be quite ambiguous what a person thinks they own or do not own, or what is counted as a gain or a loss. For instance, when Anna walks into the grocery store, does she own $4, or the box of Honey that she expects to buy for $4? That depends on how she thinks. That means it's somewhat arbitrary and subject to perception, intuition and framing. We will look next at why this matters.

❖ RESEARCH METHODS 2.4 ADVERSARIAL COLLABORATION

Observed behavior can often be explained by two or more competing theories, with no way to distinguish between the theories using existing data. That means that we need a more nuanced experiment to try and evaluate one theorem over another. Typically, such experiments are done by researchers who historically have supported one side of the debate. Typically, they find data that

supports their side of the debate. This may be for pure reasons, such as confirmatory bias (which we shall explain in Chapter 5), or for less pure reasons, such as wanting to further their career. Either way, it does not often seem the most constructive way to do things.

An alternative is adversarial collaboration. This is when researchers on both sides of a debate get together and jointly design and evaluate an experiment to test respective theories. In the study we have been looking at, a team, loosely based at the University of East Anglia, joined with Daniel Kahneman to explore the 'no loss in buying' hypothesis. Given that both sides are given equal chances to influence the design of the experiment and look at and evaluate the data, the conclusions should be free of systematic bias. This seems, therefore, a much more constructive way of doing things.

For completeness I should briefly explain what the researchers did. A third valuation method is **equivalent gain** (EG), in which a person is asked the smallest amount of money that would be equivalent to one more unit of the good. Like WTA, this question asks about consuming c rather than $c + 1$ units of the good. The difference is the endowment point. In WTA, the person should feel as though they are endowed with $c + 1$ units and are being asked to give up a unit. With EG, the person should feel as though they are endowed with c units and have a choice between more money or an extra unit of the good.

All agree that WTA will exceed EG. The more interesting question is whether EG exceeds WTP. If giving up money in exchange for a good is counted as a loss, then EG should exceed WTP; if it is considered a forgone gain, then EG and WTP should be equal. This seems to provide a simple test. What happens, however, if WTA is a 'lot bigger' than EG and EG only 'slightly bigger' than WTP? This is the kind of thing open to different interpretations. So, valuation methods called risky WTP and risky WTA are also used. For reasons we will skip, this gives a test of the 'no loss in buying' hypothesis:

if $EG > WTP$ and $WTP = riskyWTP$ evidence that exchange is a loss of money
if $EG = WTP$ and $WTP > riskyWTP$ evidence that exchange is a forgone gain of money

Other checks that things are working well include that risky WTA equals the WTA. The data, summarized in Figure 2.9, suggests that exchange is a loss of money.

2.2.4 Transaction utility

If reference dependence fundamentally affects how much we value goods, as loss aversion and the endowment effect suggest it does, then it should influence our decisions to buy or not. To illustrate, suppose that Anna is on holiday and wants to buy Nutty, so goes into the local grocery store. How much should she be willing to pay? There are two things worth considering. First is the maximum she is willing to pay for a box, ignoring any possible psychological effects. Let's suppose that she likes Nutty a lot, so

this is $8. Next is the reference point, and, given that she normally pays $3, we can start with a reference point of r = $3. Remember that what this actually means is that 'the reference point is wealth minus $3'.

We can then distinguish between acquisition and transaction utility. **Acquisition utility** measures the net gain from buying something for less than it is valued, so in this case is $v(8 - p)$, where p is the price paid. **Transaction utility** measures the net gain or loss from buying something for less or more than expected, so is $v(r - p)$. The total value from buying the box of cereal is therefore $v(8 - p) + v(3 - p)$. If p = $2 then she will feel a double gain: she buys something she likes and pays less than she expected to! If = $4 then she will probably buy but will feel a loss of paying more than she expected. If = $6 then she will probably not buy because the loss from 'paying too much' exceeds the gain from buying the cereal she likes.

Transaction utility is the key thing here, in that it captures how Anna may feel good about buying a bargain, or bad about being charged too much. This can influence whether or not she buys Nutty. But this in itself is not so fundamental. What is fundamental is that the reference point she has will likely depend on perception and intuition. Whether she perceives Nutty as a bargain or a rip-off will depend on the context. An example from Thaler (2008a) illustrates:

> You are lying on the beach on a hot day. All you have to drink is water. For the last hour you have been thinking about how much you would enjoy a nice cold bottle of your favorite brand of beer. A companion gets up to go make a phone call and offers to bring back a beer from the only nearby place where beer is sold, a fancy resort hotel. He says that the beer might be expensive and so asks how much you are willing to pay for the beer... What price do you tell him?
>
> Now, substitute 'a small, run-down grocery store' instead of 'a fancy resort hotel'.

People surveyed by Thaler said they would pay on average $2.65 in the fancy hotel and $1.50 in the store. The difference probably reflected concerns for fairness, which we shall look at in Chapter 7, but the key point is that there was a difference. The reference point, and thus the perception of what is a bargain or a rip-off, depended on the context. Indeed, the reference point by which transaction utility is judged can seemingly change quite easily. We are back to framing effects and choice arbitrariness. Someone might buy a beer from a fancy resort hotel but not from a small, run-down grocery store, even if it's the same beer at the same price.

The notion of transaction utility may seem intuitive, but we do hit a snag. Anna might have a reference price for Nutty, or beer, but it is unlikely she has a reference price for everything she might buy. It is also unlikely she knows for sure what things she is going to buy. It is even more unlikely she keeps track of each gain or loss of transaction utility when buying 100 or more different products in a grocery store. We have some work to do, therefore, in order to apply the notion of transaction utility.

2.2.5 Narrow framing

It could be that Anna enters the grocery store with a reference point for various things she might buy, such as 'I will probably spend $3 on Nutty'. More likely, however, is that Anna has an expectation of how much she should spend in the grocery store,

or how much she should spend a week on food. The important reference point will thus be something like '$200 a week on food'. If Nutty is cheaper than $3, then great, but that will soon be counteracted if apples are more expensive than usual. It is the final total at the checkout that matters most.

What's happening here is that things are being grouped together, or not, so that Anna can keep track of what is going on. This way she knows whether she is spending more or less than she wanted. She can keep track of spending on individual items, such as a box of Honey, and keep track of the total amount spent in the store.

The process of grouping things together not only brings some things together but also separates those that end up in different groups. This can cause **narrow framing**, wherein a choice or outcome is seen in isolation rather than being integrated with other things. For instance, spending on food items might be seen as separate from spending on leisure. Here's another example.

Question one. Imagine that you have decided to see a play for which admission costs $10 per ticket. As you enter the theater you discover that you have lost a $10 bill. Would you pay $10 to watch the play?

Question two. Imagine that you have bought the ticket to see a play for which admission costs $10. As you enter the theater, you discover you have lost the ticket and there is no way to recover it. Would you pay $10 for another ticket?

Tversky and Kahneman (1981) report that 88 percent of subjects asked question one said they would pay, but only 46 percent asked question two said they would pay. In both cases the loss is $10, so where the loss comes from clearly matters.

To capture such things we can introduce mental accounting. **Mental accounting** is the process of coding, categorizing and evaluating choices and outcomes. The primary component of mental accounting is to put any spending or income into separate accounts for specific purposes. For example, Anna might have a 'go to the theater account', a 'loose change account' and a 'food account'. Outcomes will be perceived and experienced relative to the particular account that is brought to mind. Losing a $10 bill brings to mind the loose change account, whereas losing the theater ticket brings to mind the theater account. Losing the $10 bill has no implications for the theater account, therefore, and so should not influence Anna's choice to go to the theater; she expected to pay $10 and so should still pay $10 to watch, and she just might look to save $10 somewhere else to bring the loose change account back below the reference point. By contrast, losing the ticket does matter for the theater account; she expected to watch for no extra cost but now has to pay an additional $10. This makes the cost to go to the theater feel like $20, and that may put her off going.

A pertinent question is how narrowly defined accounts are. For example, is there a 'theater account', an 'entertainment account' or just a general 'current spending account'? To get some idea, we can look at a study by Heath and Soll (1996). In the study, subjects were asked how their spending would likely be affected by a related purchase, a gift or an unrelated purchase. For instance, subjects were asked whether they would purchase a $25 theater ticket if they had ((1), related purchase) already bought a $50 sports ticket, ((2), gift) been given for free a sports ticket worth $50 or ((3), unrelated purchase) heard of a flu epidemic and had to spend $50 on a flu

Table 2.5 The proportion of subjects (%) who answered 'No', 'Yes', 'Yes' to purchasing a good when the alternative was a related purchase, gift or unrelated purchase. The numbers in bold indicate that the related purchase was not considered in the same mental account as the good being purchased.

The 'related' purchase	The good being purchased		
	Theater ticket (entertainment)	Chinese takeaway (food)	Sweatshirt (clothes)
Sports ticket	38	**12**	**12**
Boat tour	41	**16**	**6**
Party snacks	21	12	**6**
Salmon	21	24	**0**
Dinner out	32	32	**6**
Wine	15	18	**3**
Jeans	**0**	**0**	12
Watch	15	**3**	9
Costume	38	**9**	9

Source: Heath and Soll (1996).

inoculation. Of particular interest are people who answer 'No', 'Yes' and 'Yes'. That they answer 'Yes' to (2) and (3) means they would be willing to go to both the theater and sports event and can afford to do so even if they have already spent $50. That they answer 'No' to (1) is evidence of narrow framing.

Table 2.5 summarizes what proportion of subjects answered 'No', 'Yes', 'Yes' for various combinations of the good being purchased and the related purchase. For instance, 38 percent did so for our example question of buying a theater ticket after having purchased a sports ticket. We see that a high proportion of subjects said 'No', 'Yes', 'Yes', and so there is evidence of narrow framing. By comparing goods we can also see how narrowly defined mental budgets appeared to be. The combinations in bold were considered by subjects to not be related purchases. For example, buying a sweatshirt was not considered related to buying a sports ticket. In these cases the proportion of 'No', 'Yes', 'Yes' answers goes down. This suggests that mental accounts can be quite narrowly defined. Spending on a sweatshirt does not come out of the same account as spending on a sports ticket, boat tour, party snacks, etc.

More generally, the evidence suggests that people use mental accounts to keep track of income and spending, and that gains or losses are coded relative to an account reference point. One interesting thing about this is that we now have two dimensions to the arbitrariness of the reference point. There is the question of what reference point will be brought to mind. Will Anna think of the 'Nutty is usually $3' reference point or the 'I should spend $200 on food' reference point? Will the person on the beach think of the '$2 for beer is expensive' reference point or the 'I have $8 loose change' reference point? There is then the question of what level the reference point should be: a person might pay $2.50 from a fancy resort hotel but only $1.50 from a small, run-down grocery store. This gives lots of room for perception and intuition to matter. No wonder we get framing effects and choice arbitrariness. But can we go even further?

2.2.6 Hedonic editing

So far we have primarily been thinking of a person responding to external factors beyond their control that influence perceptions, intuition, the reference point and possibly behavior. Mental accounting, however, suggests that people may have some control of such factors. Anna can potentially choose what mental accounts she has, the reference point for each account and how she codes gains and losses.

To illustrate, suppose that Anna buys one box of Nutty for herself and a box of Superior for her son. She goes to the shop and, fortunately, finds that both are in a half-price sale. This means she gains $1.50 on the box of Nutty and $3 on the box of Superior. This could be coded as two separate gains or as one big gain. If the outcomes are **segregated**, they are valued separately to give $v(\$1.50) + v(\$3)$. If they are **integrated**, they are combined to give $v(\$4.50)$. We are going to show in the next chapter that the value function is probably concave, which means that $v(\$1.50) + v(\$3) > v(\$4.50)$. Thus, it would be better for Anna to segregate the gains in her mind.

Suppose that a month later, the sale long over, Anna goes to the store and finds that the prices of Nutty and Superior have increased by 50 percent. This means she needs to spend an extra $1.50 on Nutty and $3.00 on Superior. Segregating these losses gives $v(-\$1.50) + v(-\$3)$, while integrating the losses gives $v(-\$4.50)$. We shall show that the value function is probably convex in losses, which means that $v(-\$1.50) + v(-\$3) < v(-\$4.50)$. This means that Anna would do best to integrate losses.

So, if Anna has a choice how to code gains and losses, then: multiple gains should be segregated; multiple losses should be combined; and gains should be used to cancel small losses. Does such integration and segregation make sense? Intuitively, it does seem to, as an example by Thaler (2008a) illustrates:

> Mr A was given tickets to some lotteries. He won $50 in one lottery and $25 in the other. Mr B was given a ticket for another lottery and won $75. Who was happier?

Most people say that A will be happier. As Thaler suggests, 'Moral: don't wrap all the Christmas presents in one box!' We call this process of categorizing gains or losses as **hedonic editing**. It suggests that people can partly manipulate how they perceive or interpret things in order to make themselves feel better. Another way to do this is choice bracketing.

2.2.7 Choice bracketing

The focus in this final section is on the coding of losses and gains over time. Often there is a delay between paying for a good and consuming it, or consuming it and paying. For example, Anna might buy some frozen food she does not expect to eat for a month or two, but also might put everything on credit, meaning that she does not need to pay for a month or two. Such things mean that losses and gains will appear in accounts at different points in time. People can choose how narrowly or broadly they want to group together events spread over time, an example of **choice bracketing**.

I will look first at the implications of a delay between paying and consumption, using a study by Shafir and Thaler (2006) to illustrate. They first asked subscribers to a wine newsletter how they would think about giving away as a gift, drinking or dropping

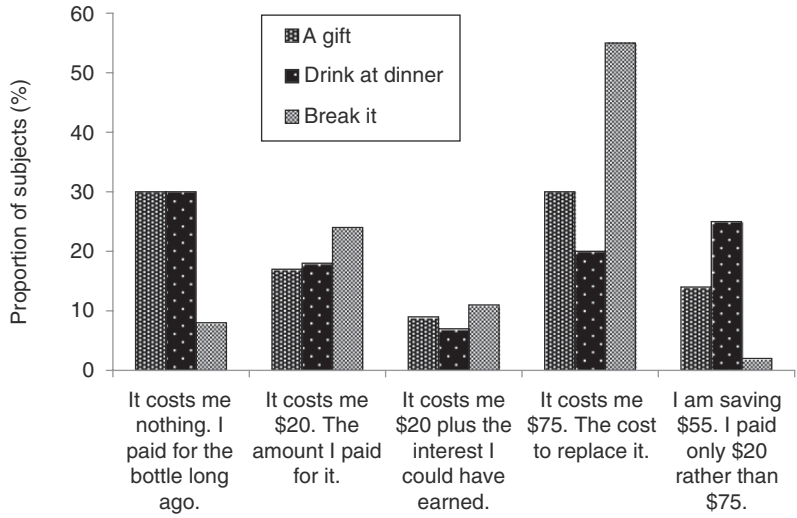

Figure 2.10 How subjects valued a bottle of wine that cost $20 and is now worth $75 depending on whether they drank it, gave it as a gift or broke it.
Source: Shafir and Thaler (2006).

and breaking a bottle of wine that cost $20 and now sells for $75. Figure 2.10 summarizes the results. Overall, we see a fairly mixed response. Many, however, thought that giving away or drinking the bottle of wine costs them nothing, or even saves them money. This suggests the cost must be accounted for when the bottle is bought. Further questioning, however, suggested that, when buying the wine, most framed the purchase as an investment rather than a cost. This motivates the title of Shafir and Thaler's paper, 'Invest now, drink later, spend never'. It is as if delayed consumption allows people to frame the initial purchase as an investment that has depreciated by the time the consumption takes place. This is a convenient way to never feel any loss! Only if the bottle is broken does the loss become apparent.

We can next look at what happens when there is a delay between consumption and paying, with a study by Siemens (2007). In the study, subjects were paid $5 for doing a 30-minute survey. All subjects were given the $5 immediately. What differed was when they did the study. Some did it immediately, some three days later, some a week later and some two weeks later. Having finished the survey, subjects were asked what they thought about the task. Figure 2.11 summarizes the results. We see a fairly clear and consistent drop-off in perceived satisfaction, fairness and willingness to do the task again, the more the delay between getting the $5 and doing the survey.

Both these studies give a consistent picture in which the costs and benefits, or losses and gains, diminish over time. So, the cost of buying the wine is forgotten by the time it is drunk, and the gain of $5 is forgotten by the time the task has to be done.

One interesting thing to take from this is that we now have a third dimension to the arbitrariness of the reference point, namely how far ahead or back a person looks. For instance, Anna's entertainment account could be '$50 a week' or '$200 a month'. The former option means she will probably go to only one event a week. The latter option allows her to attend a few events one week if she makes up for that with

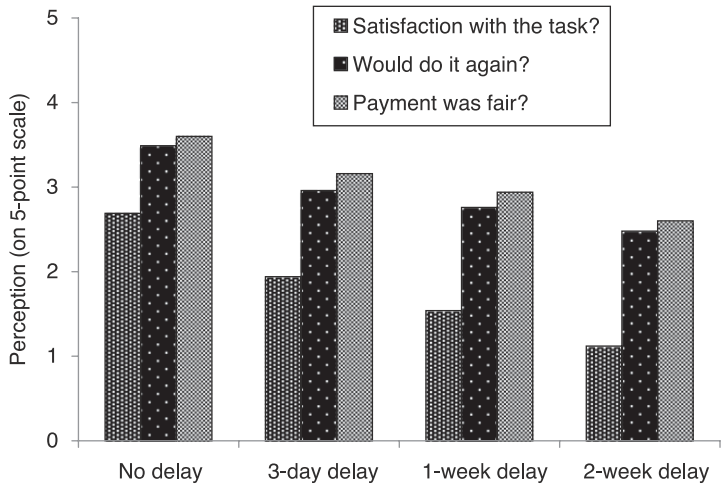

Figure 2.11 How subjects felt about a task depending on how long they did it after getting paid. Answers measured on a five-point scale from unpleasant, definitely would not, extremely unfair (1) to pleasant, definitely would, extremely fair payment (5).
Source: Siemens (2007).

subsequent weeks of abstinence. What time frame comes to mind at any one time will likely depend on her perception and intuition of the situation.

Hopefully, by now you have got the idea that reference points matter and are somewhat arbitrary. They are likely to depend on the context and the way that someone perceives a particular situation. This is a very important idea to keep in mind as you read through the rest of the book.

2.3 The role of emotions

All the heuristics mentioned so far have been relatively practical ones, such as 'Do the same as last time', 'Spend no more than $200 a week on food' and 'Pick the one in the middle'. It is important to recognize that heuristics can also have a much more emotive and moral character. One particularly interesting example is the heuristic 'Do not lie'.

Clearly, the decision whether or not to lie has various moral, religious and philosophical undertones. This does not stop it being a crucial component in many economic interactions, however. The saleswoman can lie about the quality of her product, the job candidate can lie about her qualifications, the buyer can lie about her ability to pay, and so on. Whether or not the person lies, and whether or not others trust her to not lie, are fundamental to how such interactions will play out. Heuristics such as 'Do not lie' and 'Trust others to not lie' are, therefore, easily as relevant to economics as heuristics such as 'Do the same as last time'.

Indeed, emotions are a vital part of what makes us different from *Homo economicus*. We get excited, angry, sad, upset and happy, while *Homo economicus* just keeps on

maximizing objective functions. Emotions inevitably influence choice, and so behavioral economics needs a way to take account of that. What I want to do here is illustrate how emotions naturally fit within the system of mental accounting and framing developed in the previous section. In order to do that, I shall focus on the 'Do not lie' heuristic.

2.3.1 Aversion to lying

Homo economicus would lie whenever he or she can benefit from doing so. *Homo sapiens*, however, does not like lying. To illustrate the point I'll look at an experimental study by Mazar, Amir and Ariely (2008). Subjects were given a test sheet containing 20 matrices, such as the one in Figure 2.12, and asked to find, for each matrix, the two numbers that add to ten. In the matrix shown, for instance, the answer is 4.81 and 5.19. Subjects had four minutes to solve as many matrices as possible and were asked to record the number solved on a separate answer sheet. They were subsequently paid either $0.50 or $2 for each matrix solved.

Three treatments were compared. In the 'answers will be checked' treatment subjects had to hand in both the test sheet and answer sheet. The experimenter would then verify how many matrices the subject solved. In the 'chance to lie' treatment subjects had to hand in only the answer sheet. The experimenter, therefore, had no opportunity to check how many matrices the subject really had solved. Finally, in the 'honor code' treatment subjects had the chance to lie but were asked to sign below a statement that read 'I understand that this short survey falls under MIT's [Massachusetts Institute of Technology's] honor system'. This setup captures three basic ways in which employers can monitor workers. They can closely monitor worker output, let the worker self-report or let the worker self-report with reminders of moral obligations.

Letting a worker self-report may seem bizarre. What is to stop, for instance, someone in the 'chance to lie' treatment saying they solved all 20 matrices? They walk away with the maximum payout, and the employer (or experimenter, in this case) is short-changed. Self-reporting is very common in the workplace, however, particularly in skilled professions. A professor, for example, does most of her work behind the closed, or half-closed, door of her office. She might be carefully marking students' work, or she might be rushing through the marking with one eye on the football that is streaming live on the internet. Monitoring performance is very difficult, and that

1.69	1.82	2.91
4.67	4.81	3.05
5.82	5.06	4.28
6.36	5.19	4.57

Figure 2.12 A matrix for the adding to 10 task. Subjects had to find the two numbers that add up to 10.
Source: Mazar, Amir and Ariely (2008).

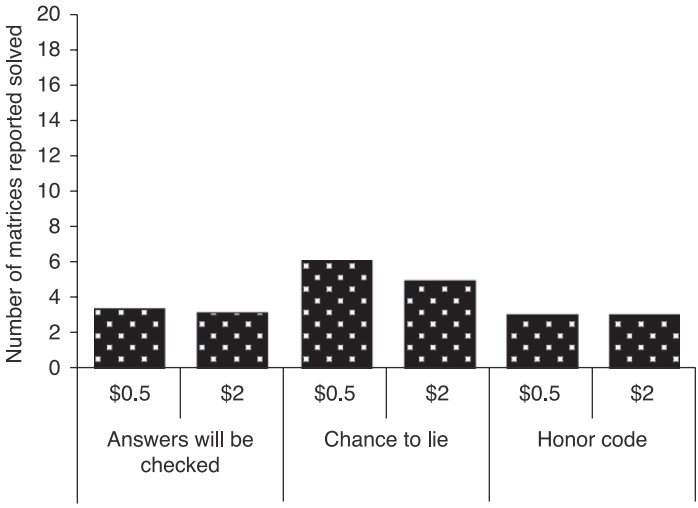

Figure 2.13 Number of matrices that subjects reported to have solved. In the 'chance to lie' treatment, subjects 'solved' more correctly. But cheating was well below that possible. Subjects could have reported 20 solved matrices.
Source: Mazar, Amir and Ariely (2008).

means we rely on the honesty of the professor. Figure 2.13 shows why that may not be such a bizarre thing to do.

Figure 2.13 summarizes the number of matrices that subjects reported to have solved. Subjects in the 'chance to lie' treatment did on average 'solve' more than those in the 'answers will be checked' treatment. This suggests that some cheating went on. The numbers are way below the maximum of 20, however. So, subjects did not cheat anywhere near as much as they could have done. Also, getting subjects to sign up to the honor code was enough to eliminate any cheating. And increasing the rewards to cheating from $0.5 to $2 decreased the amount of cheating. It seems as though subjects did use a 'Do not lie' heuristic.

2.3.2 Deception

In the discussion so far we have focused on a person who could lie for financial gain and has to choose whether or not to do so. For example, the worker can misreport output and, for sure, earn a higher payout. In reality, things are rarely so simple. A more realistic exchange would be one in which the worker reports output and the employer decides whether or not to believe that report. Or one in which the saleswoman reports on the quality of her product and the buyer decides whether or not to trust her. Such interrelationships bring us to the issue of **deception.** To illustrate the ideas we can look at a so-called **cheap-talk sender–receiver game**.

The game involves two options, called A and B, and two people, called the sender and the receiver. The sender gets to see the payoffs from each option. She learns that option A would give her $5 and the receiver $15, and that option B would give her $15 and the receiver $5. Clearly, option A looks better for the receiver and option B

looks better for her. Having seen these payoffs she then has to send a message to the receiver. She can either send the message 'Option A will earn you more money than option B' or 'Option B will earn you more money than option A'. The first message is honest and the second is dishonest. The receiver, without knowing anything about payoffs, must choose option A or B.

To put the game in context, you can think of a saleswoman trying to sell a student a calculator. Model A is the better model but not very profitable for the seller. Model B is not so good, but a lot more profitable for the seller. The saleswoman must choose whether to be honest and say model A is better. The student must decide whether to believe the saleswoman's claims. The saleswoman has no material incentive to be honest. So, why should the student believe anything she says? Well, if people have an aversion to lying, then it may make sense to trust her word.

A study by Gneezy (2005) had subjects play this cheap-talk game and others in which the payoffs from options A and B differed. Subjects also played a dictator version of the game in which the sender chose between options A and B. The dictator game is important for comparison purposes, because people, as we'll see in Chapter 7, often voluntarily give money to others. Figure 2.14 illustrates the main results of the study.

To make sense of the results, let us start by looking at the games in which option A yields \$5 and \$15 and option B yields \$15 and \$5. In the dictator version of the game around 90 percent of senders chose option B. This is not surprising, given that the sender chooses to have \$15 rather than \$5. In the cheap-talk version of the game only around 50 percent of senders sent the 'B will earn you more money' message. The drop from 90 percent to 50 percent could be an aversion to lying. Senders did not want to lie to receivers about which option was best (see Research methods 2.5).

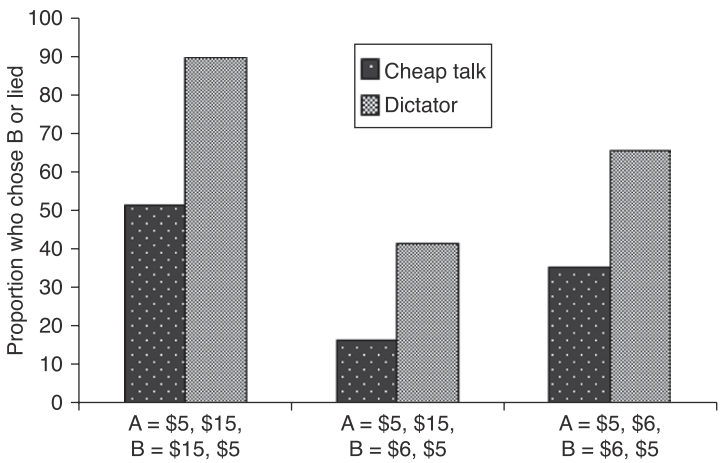

Figure 2.14 The proportion of subjects who chose option B or lied that option B will earn more money. Subjects showed an aversion to lying, as shown by a lower proportion in the cheap-talk game than the dictator game.
Source: Gneezy (2005).

❖ RESEARCH METHODS 2.5 EYE TRACKING AND LIE DETECTION

In a study by Wang, Spezio and Camerer (2010) subjects played a sender–receiver game not too unlike the one already described. The main differences were that there were five options instead of two and the receiver had some information about payoffs. These differences mean that strategy is more of an issue. I don't, however, want to focus on this. Instead, I want to focus on the fact that the eye movements of senders were monitored as they made their decision.

This was possible with an eye-tracking device that can tell where subjects are looking (or on what their eyes are fixated). For example, we know that senders spent an average 3.99 seconds looking at what their own payoff would be for each option, and only 2.72 seconds looking at what the payoff of the receiver would be. The device can also tell how dilated the pupil is. This is fascinating, because pupils are known to dilate during stress or while doing cognitively difficult tasks. As expected, the pupils of senders become more dilated the more they attempt to deceive the receiver.

So, what do we gain from eye tracking? Without eye tracking we observe only the choice of the sender. We then have to try and infer why the sender made the choice they did and how they felt about it. That can be a difficult, if not impossible, task. Eye tracking gives us a lot more information to go on. By analyzing the things they chose to look at we can potentially piece together the strategy they were using (see also Research methods 6.1). By measuring pupil dilation we get some limited insight on how they felt about the decision (see also Chapter 9). To prove the point, the authors of the study showed that information on eye movements and dilation was enough to relatively accurately predict the truth. A lie detector test!

For the two other payoff combinations, senders were less likely to choose option B. This is understandable given the change in payoffs and a willingness to give money to others. We still, though, see a clear drop in the proportion of senders who lied by saying that option B was better. Again, this suggests some aversion to lying.

Interpreting behavior is a little bit tricky in sender–receiver games, however, given their strategic element. For instance, suppose the saleswoman expects the student will use the heuristic 'Never trust a saleswoman'. Then the saleswoman may recommend model A as the better in the hope that the student's mistrust will cause her to buy model B. In this case, the saleswoman attempts to deceive the student by telling the truth! Such strategic concerns are more the topic of later chapters (see, in particular, Research methods 7.4), but I want to give you some idea of how prevalent such 'false truth telling' can be. To do that, I turn to a study by Sutter (2009).

The key thing about this study was that senders were asked which option they expected receivers to choose. This allows us to pick out the four types of behavior summarized in Table 2.6. We are particularly interested in the sophisticated truth

Table 2.6 Four different types of behavior possible in cheap-talk sender–receiver game.

Behavior of sender	Recommends to receiver	Expects receiver to choose
Benevolent truth teller	A	A
Sophisticated truth teller	A	B
Liar	B	B
Benevolent liar	B	A

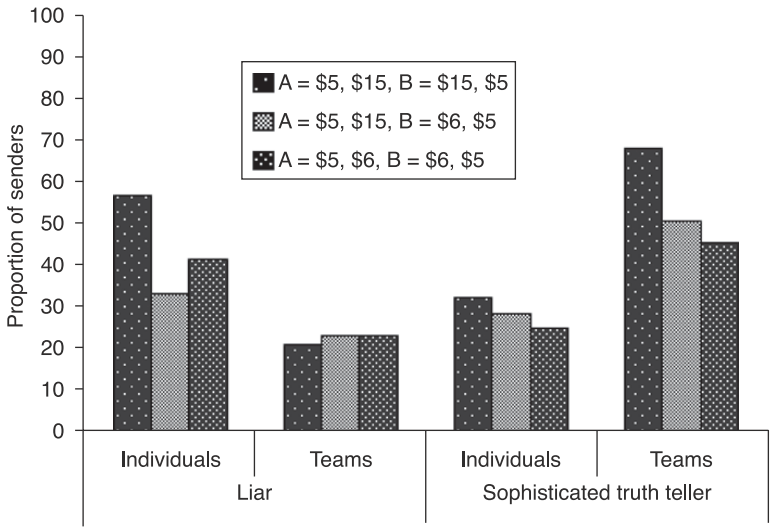

Figure 2.15 Proportion of senders who were liars and sophisticated truth tellers distinguished into individuals and groups. A large proportion of individuals were sophisticated truth tellers. An even larger proportion of groups were sophisticated truth tellers.
Source: Sutter (2009).

teller. This is someone who 'tells the truth' expecting the receiver to not trust her. Figure 2.15 details the proportion of senders that could be classified as liars and sophisticated truth tellers (hardly anyone was a benevolent liar). The study contrasted the situation in which the sender was an individual with the situation in which the sender was a team of two or three that made a joint decision.

Focus first on individuals. Previously, in Figure 2.14 in particular, we picked out only the liars. You can see that the proportion of liars is similar in this study by Sutter to that of Gneezy. We can now also see, however, that around 20 to 30 percent of subjects were sophisticated truth tellers. In teams, the proportion is much higher (see Research methods 2.6). Such evidence shows that we need to take seriously the idea that people can deceive by telling the truth. Once we take account of sophisticated truth tellers, the proportion of senders who do not attempt to deceive (benevolent truth tellers) falls to around 30 percent. This is still a large proportion, but it is clear that people feel less aversion to deceiving than to lying. An aversion to lying does not necessarily, therefore, translate into always telling the truth.

This qualification should not, though, detract from the basic message that senders reveal much more information than *Homo economicus* would. It is also clear that receivers are a lot more trusting of the information sent than *Homo economicus* would be. For example, in both the studies by Gneezy and Sutter, over 70 percent of receivers followed the advice of the sender. So, the heuristic 'Trust others to not lie' seems to fit much better than 'Never trust a saleswoman'.

❖ RESEARCH METHODS 2.6 VIDEOING TEAM DECISION-MAKING

There are various reasons why a sender might look like a sophisticated truth teller but not want to deceive. For example, she might think that the receiver is going to choose B no matter what she says. Or, she might be averse enough to lying that she says A is better but doesn't want the receiver to choose B. In neither case is her objective in choosing A the deliberate deception of the receiver. It would be nice, therefore, to know what is going through the mind of sophisticated truth tellers when they send message A. Are they really trying to deceive the receiver?

This is impossible to tell from the choices that senders make. So, in order to gain some insight on this question, the study by Sutter looked at decision-making in teams. The teams were videoed reaching their decision. By analyzing the discussion between team members it is possible to get a much clearer picture of how they arrived at their decision. Of the 37 sender teams that were classified as sophisticated truth tellers, 34 explicitly discussed the argument that sending option A would increase their payoff because the receiver would not believe them and choose option B. This strongly suggests that most sophisticated truth tellers did want to deceive.

Here we see a clear advantage of using teams in the laboratory. By eavesdropping on the conversation between team members we can get a clearer picture of each subject's thought processes. This can help us understand why a particular decision was made. The drawback is that teams may behave differently from individuals. This is clear from Figure 2.15, where we see that teams were more likely to be sophisticated truth tellers than individuals. It does not seem unreasonable, however, to argue that the deception observed in teams is a good indication that many sophisticated truth-telling individuals also set out to deceive. In Chapter 6 we shall look at team decision-making in more detail.

2.3.3 Honesty and framing

The general picture we get, from many different studies, is that 'Do not lie', 'Do not deceive' and 'Trust others to not lie' are important heuristics in economic decision-making. This is not to say that people do not lie and deceive. They clearly do. The collapse of Enron and doping in professional sport are two powerful examples of

people's willingness to lie for financial gain. The point is: people lie and deceive a lot less than they could do.

So, why are people so honest? A leading explanation is that honesty helps a person maintain a **positive self-concept**. The person wants to feel that she is a good person, and by not lying she can help maintain a positive opinion of herself. If this is true, then we should be able to observe some interesting framing effects. We have already seen one example: the honor code reminds a person of the importance of honesty and so brings the 'Do not lie' heuristic to mind. On the flip side may be possibilities to lie without harming self-concept. For example, if a person perceives that 'everyone is lying', then it may excuse her lying. People may also be more willing to lie in order to avoid a loss than to achieve a gain. And they may be even more willing to lie in order to avoid a loss of something they feel they have earned. We'll come back to some of these issues when we look at tax avoidance in Chapter 3.

The differences we observed between deception and lying can also arise from a framing effect. In a sender–receiver game strategic interaction is foremost. This more easily allows a person to maintain positive self-concept by passing on the blame to others. Should, for instance, a sophisticated truth teller feel any guilt from recommending A? After all, she told the truth. The receiver will lose out only if she mistrusts what she is told. Similarly, should a liar feel any guilt from recommending B? In this case the receiver loses out because of naïvely trusting what she is told.

Such reasoning can explain why people appear more willing to deceive than lie. Given that lying is inherently an interpersonal process, however, passing the blame is always possible. For example, the professor who self-reports output can say that the university is naïve to believe she will report honestly! The difference is that some contexts, and frames, make it easier for a person to lie and maintain self-concept by passing the blame than others.

For an interesting illustration of the way in which context can matter I will briefly look at a study by Cohn, Fehr and Maréchal (2014) on dishonesty in banking. Their study involved an experiment in which subjects had to self-report the outcome of ten coin tosses. For each coin toss a subject could earn $0 or $20 depending on whether they said the coin had come down 'heads' or 'tails'. Note that there was no way to verify the actual outcome of the coin toss, and so a subject could easily lie and walk away with $200. The important twist to the experiment was that the subjects were employees of a large, international bank and half of them had been primed to think about banking. For instance, those primed to think about banking were asked questions such as 'What is your function at this bank?', while others were asked questions such as 'How many hours per week do you watch television on average?'.

Figure 2.16 summarizes the results by comparing reported outcomes with those that would be expected by pure chance. We can see that the bankers who were not primed to think about banking seem honest; the distribution of their reported outcomes were what we would expect from the outcome of ten coin tosses. By contrast, the bankers who were primed to think about banking were 'remarkably successful' at tossing coins and earned considerably more than would be expected by pure chance. Indeed, 10 percent walked away with the maximum $200. These results would suggest that bankers are not dishonest in themselves but may behave dishonestly when working as bankers. Context matters, and potentially we need to worry about a culture of dishonesty within banking.

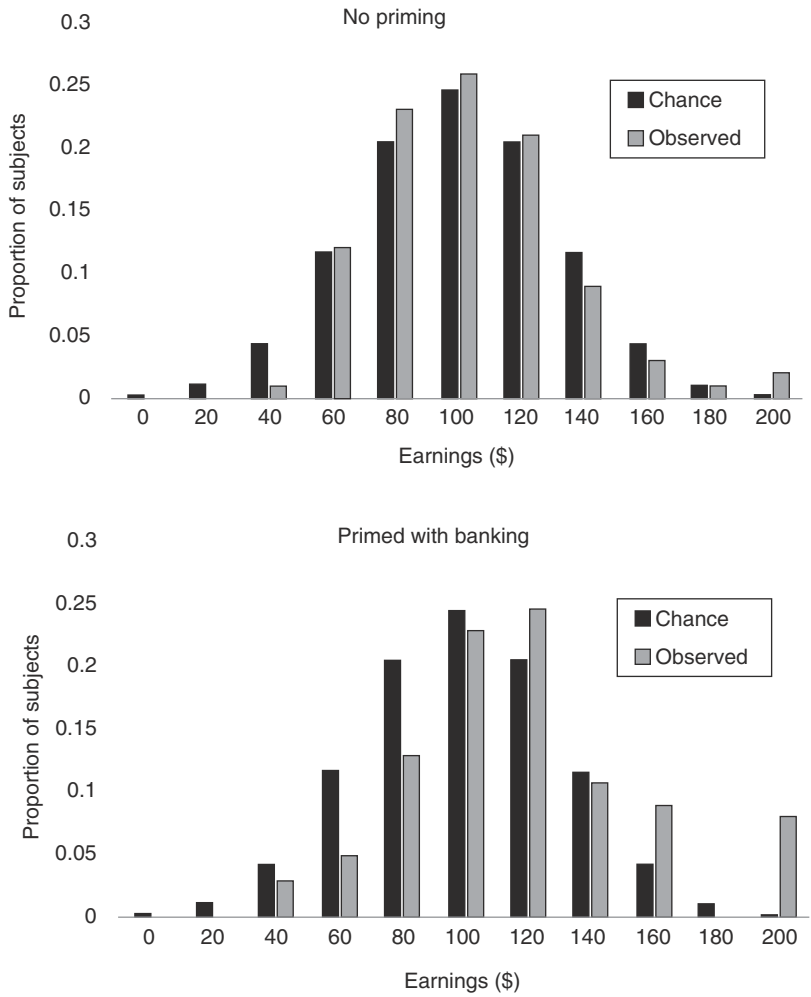

Figure 2.16 The observed distribution of earnings from ten coin tosses compared to what we would expect with pure chance comparing those with no priming or priming about banking.
Source: Cohn, Fehr and Maréchal (2014).

In bringing this section to a close I want to highlight the key thing you should take away: heuristics, perception and intuition can have a moral and emotional dimension to them. A person feels guilt from lying. She feels pride from being honest. She feels pleasure when someone tells the truth. She feels angry when she is deceived. Not all heuristics, therefore, are about simple, hard-headed cost–benefit calculation. Instead, emotions often drive perception and intuition, meaning that context effects can be caused by the mood that the context generates.

We can already see the consequences of emotions in things such as loss aversion, the endowment effect and transaction utility. Paying $2.50 for a beer from a fancy

resort hotel is fine, but being asked to pay $2.50 in a run-down grocery store makes a person angry. That anger means she would rather stay thirsty than buy the beer.

In the remainder of the book we shall see more and more examples of the role that emotions play in choice. For instance, emotions are particularly important when analyzing interpersonal choice. A person can lie only if they have someone to lie to; they feel guilty for having deceived the other person. In later chapters we shall see heuristics such as 'Punish those that harm me' and 'Reward those that help me'. Such heuristics can be understood only if we take account of the emotions that harm and help arouse in the person.

2.4 Summary

We have already covered a lot of important ideas, so it's high time that I put things into context with a short summary.

I started with the observation that saying someone should maximize her utility does not get us far if, as seems reasonable, she does not know what maximizes her utility. This lack of knowledge means she needs to search for more information about what she likes and dislikes. In order to solve this problem, and others, people have an adaptive toolbox of heuristics that they can draw upon.

During searches choices will be somewhat arbitrary until a person becomes better informed. Indeed, choices can be influenced by all sorts of external factors, including the way things are framed. We called these context effects.

I then argued that people focus more on relative rather than absolute judgments. In terms of economic behavior this means there will be a focus on gains or losses relative to some reference point. People seem particularly averse to losses, resulting in things such as the endowment effect.

We then saw that the reference point can be a very subjective thing. It is essentially what someone expected or thought would or should happen, and the person is free to think different things. She is also free to change the viewpoint from, say, what will happen this week to what will happen this month.

One step up from this is a person using 'creative accounting' to manipulate the reference point or viewpoint to increase her utility. This we called hedonic editing. Most of the time, however, the reference point will probably be influenced by subconscious processes such as intuition and perception. So, context influences perception and intuition, which influences the reference point, which influences choice. Framing matters and the reference point matters!

The final thing we looked at is the importance of emotions. Heuristics can have a moral and emotive character to them, such as the heuristic 'Do not lie'. Also, perception and intuition will often be influenced by the emotions that a context generates.

That there are context effects, and that emotions and framing matter, might seem obvious, particularly if you have studied some psychology, but it does mean a fairly fundamental departure from the standard economic model. Specifically, there are two new and crucial things we need to add to the model: (1) that utility depends on outcomes relative to some reference point; and (2) that the reference point and choice will depend on the context in which the decision is made in somewhat arbitrary ways. The first thing is relatively easy to add to the standard economic model, the second less so, but we will see more on that in the next chapter. What I want to do now is

look through some applications in which heuristics, search and reference dependence have a role to play.

2.5 Demand, supply and markets

Through much of this chapter I have used the example of Anna in the grocery store choosing a breakfast cereal. This is illustrative of a **market** where potential buyers of a good, such as Anna, interact with people who are collectively willing to provide that good (the owners and workers of the store, farm that supplies the wheat for the cereal, etc.). Markets are at the heart of economics, so they seem an apt place to begin looking at applications of what we have learnt.

It's easiest to start with a laboratory market. The standard way to create a market in the experimental lab begins by randomly allocating subjects to the role of buyer or seller. Each buyer is given a card that has written on it the value to her of buying one unit of a fictitious good that will be traded in the experiment. Each seller is given a card that has written on it the value to her of selling one unit of the good. Different subjects are given different values, and only they know their value. Buyers and sellers are then able to interact and do a deal. The profit of a buyer is her value minus the price she paid. The profit of a seller is the price she sold for minus her value. So, in principle, buyers want to buy as cheap as possible and sellers to sell as dear as possible.

If we know the values of all sellers and buyers we can derive a demand and supply curve. To illustrate, suppose that the values given to buyers and sellers are as in Table 2.7. For example, one buyer is given a value of $2.30, meaning that, if she bought for $2, she would make a profit of $0.30. To derive the **supply curve** we have to look at all possible prices and ask how many sellers would sell at each price. If the price is less than $1.50 then no one should sell. If the price is between $1.50 and $2 then there is one person who should sell, namely the seller with value $1.50. If the price is between $2 and $2.50 then there are two people who should sell, namely the sellers with value $1.50 and $2. Carrying this on, we obtain the supply curve in Figure 2.17. To derive the **demand curve** we have to look at all possible prices and ask how many buyers would buy at each price. If the price is more than $3.10 then no one should buy. If the price is between $2.80 and $3.10 then there is one person who should buy, namely the buyer with value $3.10. Carrying this on, we obtain the demand curve.

Demand and supply curves are one of the most basic ideas in economics, and these curves are no different from those you may be familiar with. For instance, the demand curve must slope down and the supply curve slope up. (You might be used

Table 2.7 The number of sellers and buyers before and after a shift in demand.

Type	Distribution of values
Sellers	There are five sellers, with values $1.50, $2, $2.50, $2.50 and $3
Buyers	There are five buyers, with values $1.80, $2.30, $2.60, $2.80 and $3.10
Buyers after change	There are five buyers, with values $1.50, $1.70, $2.20, $2.20 and $3

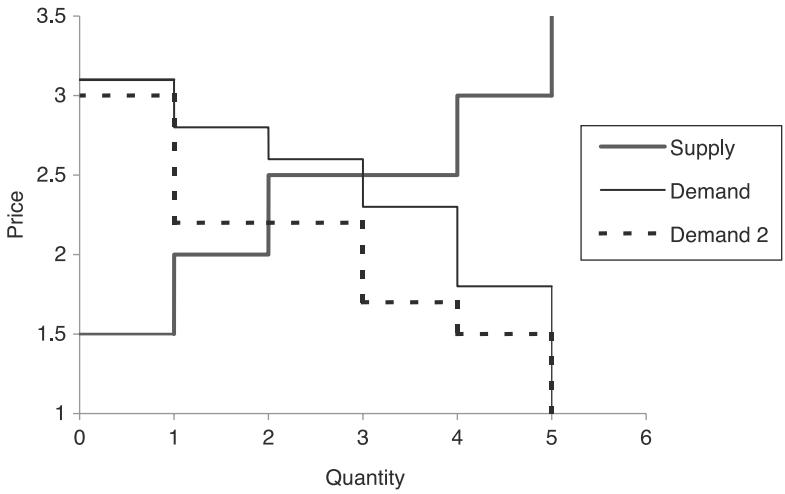

Figure 2.17 The demand and supply curves before and after a shift in demand.

Table 2.8 The market is most efficient in equilibrium, as evidenced by the high total profit.

	Trades			Profit		
	Buyer value	**Seller value**	**Price**	**Buyer**	**Seller**	**Total**
Equilibrium	$3.10	$1.50	$2.50	$0.60	$1.00	$1.60
	$2.80	$2.00	$2.50	$0.30	$0.50	$0.80
	$2.60	$2.50	$2.50	$0.10	$0.00	$0.10
						$2.50
Maximum trades	$1.80	$1.50	$1.65	$0.15	$0.15	$0.30
	$2.30	$2.00	$2.15	$0.15	$0.15	$0.30
	$2.60	$2.50	$2.55	$0.05	$0.05	$0.10
	$2.80	$2.50	$2.65	$0.15	$0.15	$0.30
	$3.10	$3.00	$3.05	$0.05	$0.05	$0.10
						$1.10
Maximum seller profit	$3.10	$1.50	$2.79	$0.31	$1.29	$1.60
	$2.80	$2.00	$2.79	$0.01	$0.79	$0.80
						$2.40

to smooth curves rather than the stepped ones drawn here, but if you have enough buyers and sellers the curve will start to look smooth.) The **market equilibrium price** and **quantity** are where demand equals supply. In the example, that is where price is $2.50 and quantity is three. In practice, this means we should expect there to be three trades, each at a price of $2.50. If we change the demand curve, we see that the equilibrium price drops to $2.20 and the equilibrium quantity to two. Thus, a shift in demand and supply should change the market price in a predictable way.

One remarkable thing about market equilibrium is that total profit is maximized at the equilibrium. Table 2.8 illustrates how this works by contrasting the equilibrium

with two other possible sets of trades. In equilibrium there are three trades at $2.50 and total profit is $2.50. In the second scenario all ten people trade and make a profit, but total profit is only $1.10. In the third scenario we can think of sellers colluding so as to set a price that maximizes total seller profit. There are lots of other possible combinations of trades we should check, but none will give as high a profit as the equilibrium. This is an illustration of the celebrated **first fundamental theorem of welfare economics**, that any market equilibrium is Pareto efficient. Notice, however, that the equilibrium need not be 'fair'. For instance, one seller makes $1 while several sellers would make $0.

This stylized example illustrates why demand and supply analysis is so fundamentally important in economics: it is a tool to predict what may happen in markets, and should tell us what the most efficient outcome is.

One thing we have not done is specify how people interact in order to make trades. Three of the more common ways we observe are the following.

1. A **negotiated price institution**, in which buyers and sellers are free to talk to each other and try to do a deal. This is similar to bartering in a local market.
2. A **double-auction institution**, in which buyers submit bids of what they are willing to pay and sellers submit asks of what they are willing to sell for. All bids and asks are displayed on a screen for all to see. A buyer can buy by accepting the lowest ask price. A seller can sell by accepting the highest bid price. This institution is used in most financial, commodity and currency markets.
3. A **posted offer institution**, in which each seller displays a 'take it or leave it' price. Buyers can go to any seller and agree to trade at the displayed price. This institution is what we are familiar with when we go to the grocery store or most shops.

That is enough background information. I now want to get on and look at some data on how markets work. I will start with markets in the experimental lab and work progressively towards real markets.

2.5.1 Double-auction markets in the lab

All three of the institutions above have been used in the laboratory, but the double auction is most common, so I will start by looking at that.

In one of the earliest economic experimental studies, Vernon Smith (1962) ran ten experimental double-auction markets. Each market was repeated up to six periods and had different demand and supply curves. Figure 2.18 summarizes the data from four of the sessions. For each market the equilibrium price is represented by a circle. We can then see the observed quantity and average price over the periods.

Market 6 is the easiest to explain, because there we see the average price increase over the periods from 5.29 in period 1 to 7.17, 9.06 and finally 10.9 in period 4. The quantity traded is 12 in every period. The equilibrium price was 10.75 and quantity was 12, and so by period 4 the market was trading at equilibrium. In the other three sessions I have plotted there is not much to see, and the same could be said of all the sessions run in the study; but that's the beautiful thing! The average price and average quantity traded were remarkably close to the market equilibrium in all periods.

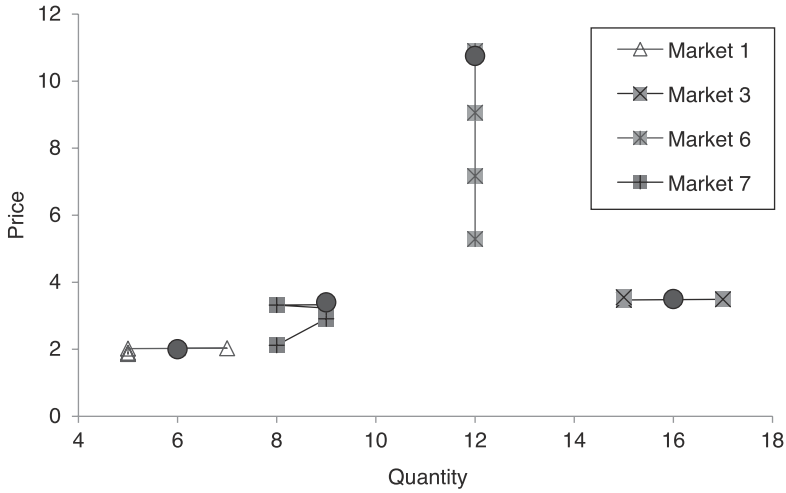

Figure 2.18 Convergence of market price and quantity to equilibrium in four double-auction markets. The circles show the equilibrium.
Source: Smith (1962).

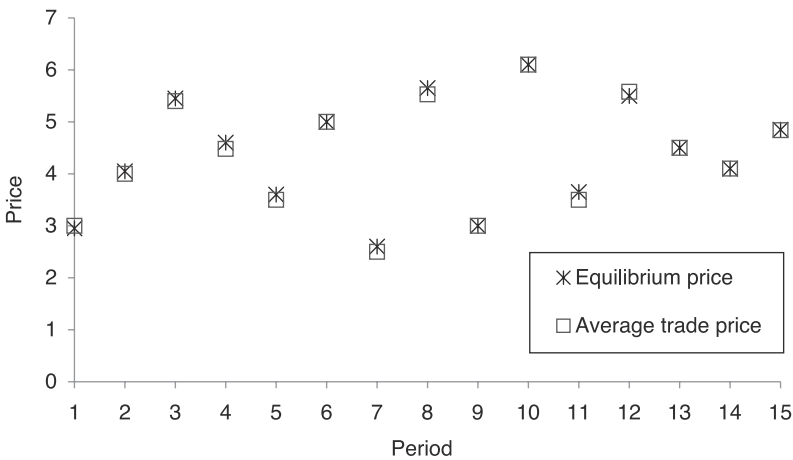

Figure 2.19 The equilibrium price and average price traded in a market that lasted 15 periods.
Source: Smith (2002).

The closeness between market and observed prices also holds when demand and supply are changed from period to period. This is illustrated in Figure 2.19, where the equilibrium and average price are plotted for a market that lasted 15 periods, with the equilibrium price changing each period. The match between equilibrium and observed average prices is nearly perfect!

Basically, the average price and quantity observed in experimental double-auction markets almost always get close to the market equilibrium with repeated trading, and

typically start close. The price and quantity also change as predicted by changes in demand and supply.

This is a fantastic result. It suggests that demand and supply are reliable predictors of what will happen in a market, which is good news for economists. It also suggests that markets are a great way to allocate resources, which is good news for everyone. Indeed, this result goes further than the first fundamental theorem of welfare economics, because the theorem is agnostic about how prices will be determined, but the double auction gives us an institution to determine prices and achieve efficiency.

The efficiency of double-auction markets is arguably the most important lesson so far from behavioral economics. The reason that double-auction markets work so well is that subjects need use only simple heuristics. Recall that the only thing a subject knows at the start is her own value to buy or sell. So she simply needs to buy for less or sell for more than her value to make a profit, and this is something very familiar to most of us (see Research methods 2.7).

❖ RESEARCH METHODS 2.7 INTELLIGENT INSTITUTIONS OR PEOPLE?

There are, broadly, two reasons why double-auction markets work: subjects are very clever, or double-auction markets are very clever. It is probably a bit of both. To illustrate why, we can look at a study by Gode and Sunder (1993). They compared what happens in double-auction markets if choices are made by 'zero-intelligence' computer programs rather than human subjects.

They use two versions of zero intelligence. In the first, each trader randomly generated bids or offers, and so there was no sense to behavior at all. This really does look like a zero-intelligence trader. In the second, each trader randomly generated bids or offers, but would never do a deal that would lose them money. So a seller never sold for less than value, and a buyer never bought for more than value. I will call this a 'Don't sell at a loss' trader. Table 2.9 gives a snapshot of the results by comparing the efficiency of markets with humans to those with zero-intelligence traders. The main result is that 'Don't sell at a loss' seems enough to get high efficiency.

Table 2.9 The market efficiency of five experimental markets, comparing human participants to zero-intelligence traders.

Participants	Market efficiency (%) in market				
	1	2	3	4	5
Zero intelligence	90.0	90.0	76.7	48.8	86.0
Don't sell at a loss	99.9	99.2	99.0	98.2	97.1
Human	99.7	99.1	100.0	99.1	90.2

Source: Gode and Sunder (1993).

How should we interpret this result? It would appear to demonstrate the relevance of heuristics, because the simple 'Don't sell at a loss' heuristic is enough to achieve market efficiency. We need to be a little careful, however, to distinguish institutions from individuals. This is because an individual trader may want to use something more subtle than this heuristic, and there was good evidence that human subjects did do that. The main lesson, therefore, is that some institutions can work well even if people use only the most basic of heuristics. This does not imply that people do use only the most basic of heuristics, but it does mean that institutions can be clever, and potentially make up for a lack of subtlety in people's strategy or behavior.

Having given the good news, it is time for some of the bad. Markets are not always so efficient, for a variety of reasons. In Chapter 5 we shall see how efficiency can decrease if buyers and sellers have some uncertainty about the value of the good. Here I shall look at a host of other reasons why efficiency may be less, starting with a brief look at posted offer markets.

2.5.2 Posted offer markets and market power

To give an illustration of what happens in posted offer markets I will look at a study by Ketcham, Smith and Williams (1984). Figure 2.20 summarizes the differences between prices and quantities observed in double-auction and posted offer markets averaged over the 12 experimental markets of the study. As expected, in the double auction, price and quantity converged over the periods towards the equilibrium. The same

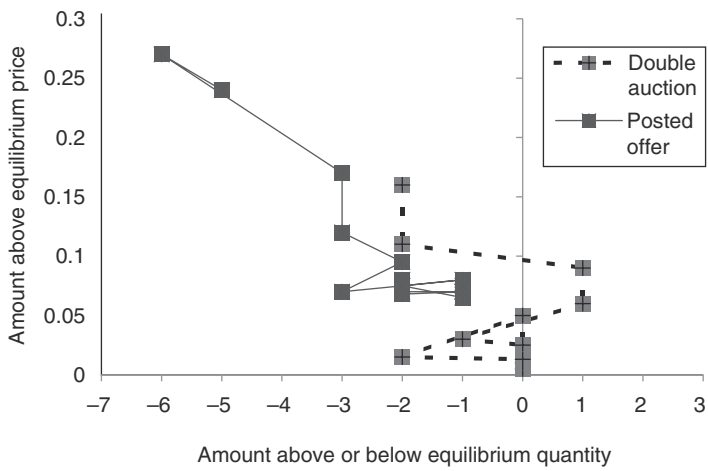

Figure 2.20 The deviation from equilibrium price and quantity in experimental double-auction and posted offer markets.
Source: Smith (2002).

cannot be said in the posted offer market. In this case prices stayed consistently above the equilibrium and quantity below the equilibrium, and there is no evidence that this would change if the market were to be repeated more times.

It is not too much of a surprise that we do not always see convergence to the market equilibrium in posted offer markets. This is because suppliers have an incentive to try and keep the price high. To illustrate, we can look back at the third set of possible trades in Table 2.8. In this case the price is pushed above the equilibrium and, while total profit is less than at equilibrium, the profit of sellers is higher than at equilibrium, at $2.08 compared to $1.50. Things are not that simple, because sellers do not share the profit and so have an incentive to compete amongst each other for their share of the profit. The basic incentive to push prices higher exists, however, and a posted offer market gives sellers more possibility to do so.

Sellers can keep prices high because they have the power to set prices. They can signal to each other from one period to the next through the prices they offer and potentially collude on higher prices. Prices are, therefore, higher because sellers have market power. More generally, market power can move the price away from the market equilibrium. This is a point I'll return to in Chapter 6, when we look at industrial organization.

2.5.3 The law of one price

The final institution to look at is negotiated price. In looking at this type of market it is pertinent to mention the law of one price. The **law of one price** is that all items should be sold at the same price in market equilibrium. For instance, in Figure 2.17 we predict that all items are sold at $2.50. In a double auction this is what we typically do see, and this is not surprising, given that all buyers and sellers automatically have access to the same information. If we consider institutions in which buyers and/or sellers have to search for information, then things change.

We can illustrate by looking at a negotiated price market from one of the first ever recorded economic experiments, by Chamberlin (1948). Recall that, in this institution, buyers and sellers move around the room, interact personally and try to do a deal on a one-to-one basis. This makes it possible that people on one side of the room could be agreeing a different price from people on the other side of the room. Figure 2.21 summarizes what happened in the market. The equilibrium price was 5.7, but we can see that the prices traded vary considerably from as low as 4 to as high as 7.2. The law of one price does not hold in this market. Interestingly, however, the average price comes out at 5.3, which is not that far from the equilibrium.

This large variation of individual prices but an average price close to the equilibrium is what we typically observe in this type of market. So, demand and supply are still reliable predictors of what will happen in the market but are not so reliable a predictor of what will happen in a specific transaction. Why?

The person who sold at $4 and the person who bought at $7.20 could have come together and agreed a price of $5.60. This would mean both make an extra profit of $1.60 and would mean the traded price would have been a lot closer to the equilibrium. Why did they not do that? Presumably they never met to talk about this possibility. Both could have searched a bit longer, or harder, for a better deal. The

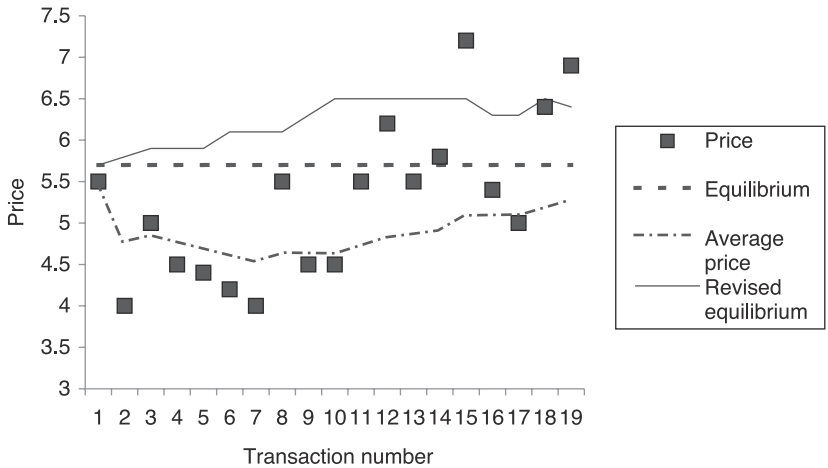

Figure 2.21 The prices traded in a negotiated price market. The prices vary a lot but the average ends up close to equilibrium. As trades are being done and buyers and sellers leave the market, the demand and supply curve changes, so we should concentrate more on the revised equilibrium price than equilibrium price.
Source: Chamberlin (1948).

problem is that such search is costly, not only in time and effort but also in potentially missing out on a good deal. For example, selling at $4 might look like a good deal to the seller; remember, they have no idea what the market demand is, so should they risk losing that deal to search for a better one? Potentially not, and so search costs and reference points are a good reason why we should not expect to see the law of one price.

Search costs are clearly relevant in many markets. You might have an idea what the price of a textbook is, but you cannot know for sure the price of different sellers without a little bit of effort. We should not be surprised, therefore, to see variation in prices for an identical good that is sold at more than one location. This is the case irrespective of the institution used to sell the good. To illustrate the point, Pratt, Wise and Zeckhauser (1979) randomly opened the Yellow Pages of the Boston telephone directory to select 39 products, and then rang all the advertised sellers for each product to see how much they charged. Figure 2.22 illustrates the large variability in prices across sellers. The maximum price was often twice that of the minimum.

The data in Figure 2.22 is arranged in order of increasing price. Carnations are cheapest, with a price of $0.33, and a boat the most expensive, with a price of $602.87. One slightly curious thing is that the relative variability in price does not decline for more expensive items. To better understand why this is curious, consider this question posed by Tversky and Kahneman (1981):

> You are about to purchase a jacket for $125 and a calculator for $15. The salesman tells you that the calculator is $5 cheaper at another branch of the store 20 minutes' drive away. Would you make the trip to the other store?
>
> Now imagine the jacket costs $15 and the calculator $125.

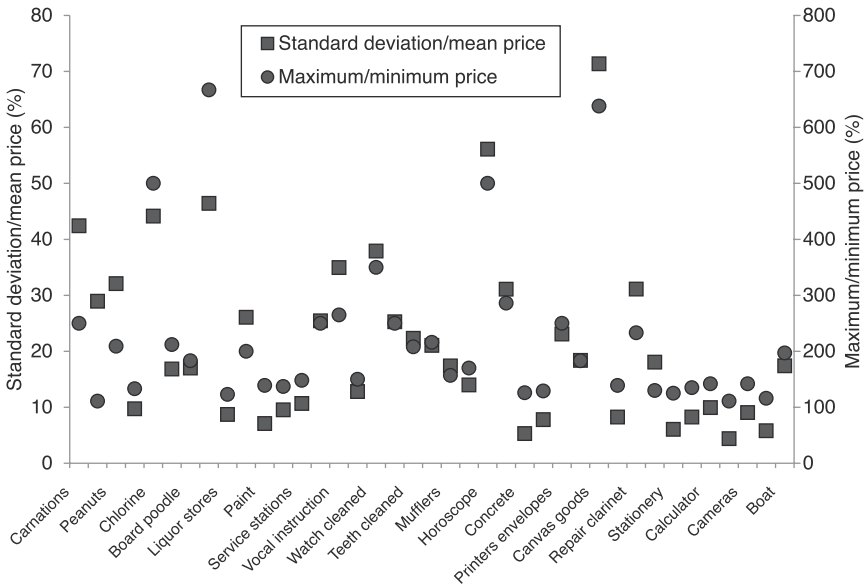

Figure 2.22 The law of one price does not hold in many markets. The standard deviation in price is typically more than 10 to 20 percent of the mean price. The maximum price is typically twice or more the minimum price available.
Source: Pratt, Wise and Zeckhauser (1979).

Tversky and Kahneman report how 68 percent of those they surveyed would make the trip to save $5 on an item costing $15, but only 29 percent would make the trip to save $5 on something costing $125. Either way, the person will save $5, and so it is slightly strange that it matters whether the saving is on a cheap or expensive item. Put another way, however, the person will save 33 percent in the one case and 4 percent in the other, so, if we have a heuristic that says 4 percent is not a good saving but 33 percent is, then things become easier to explain.

Now let's return to search costs. There is no reason to suppose the cost of search depends on the price of the item being sold. So, if someone puts the same effort into searching for a $1 saving on all goods, we should see less variability in the relative price of expensive items. We do not see this. If someone puts the same effort into searching for a 1 percent saving on all goods, then the variability in relative prices should not depend on the price of the good. That is what we do seem to observe.

This focus on relative rather than absolute savings is the first indication we have seen in this section that framing and perceptions can influence a person's willingness to trade. What I want to do now is look at two specific markets in which we might expect a bit more evidence that framing and perceptions matter.

2.6 Labor supply and reference dependence

The story told in the previous section is largely one in which it is obvious if someone has made a profit, and so gains and losses are easy to calculate. The buyer wants to

sell for less than she values the good, and the seller wants to sell for more than the cost of the good. In some markets that story is all we need, but in others it misses something, because it may not be obvious what counts as a gain or a loss, or selling at a profit. In other words, it may not be clear what the relevant reference point is. This can have implications for how the market works. I shall illustrate by looking at labor supply in this section and the housing market in the next section.

2.6.1 A target income and target wage

In the labor market it is natural to think of people deciding how many hours they want to work at the going wage. For example, if Anna is an accountant, then she needs to decide how many clients to take on, whether to work full-time or part-time, and so on. Anna may have a reference level of hours that she wants to work, or a reference level of income that she wants to earn. This reference level could be determined by a variety of things, but what has previously happened, or what friends and peers are doing, are two likely candidates. Before we look at some data it is useful to have a model of labor supply to see how the reference level may matter.

The standard way of modeling labor supply is to think of the utility function $u(x, H)$, which depends on wealth x and the number of hours spent working H. To measure the benefit of working, we need the hourly wage rate w. Anna trades wealth for leisure time, and if we solve for the optimal amount of hours she should work we get a condition such as

$$-\frac{du}{dH} = w\frac{du}{dx}. \tag{2.1}$$

This means that the disutility from an extra hour of work should be offset by the utility she gets from the money earned by working an extra hour.

Figure 2.23 gives two examples of how Anna's utility may change as she works more or less. The marginal disutility of working is assumed to increase the more she works in both examples. What differs is that in one she has a reference level for working eight hours a day. This causes a big change in the disutility of working at eight hours. Anna is relatively happy to continue work if she has worked less than eight hours, but is relatively unhappy to continue work if she has already worked more than eight hours. Consequently, equation (2.1) is likely to be satisfied when she works eight hours and labor supply is likely to be insensitive to changes in the wage.

The story when there is a target level of income is relatively similar. Figure 2.24 illustrates what may happen if Anna has $10,000 in wealth and wants to increase her income by $160. Extra income is worth somewhat more when she has earned less than the target than when she is above the target. Consequently, amount earned will be relatively insensitive to changes in the disutility of working. The number of hours worked would change, however, if the wage rate changes. A higher wage would mean Anna can work fewer hours and still reach her target, while a lower wage would mean Anna needs to work more hours to reach the target.

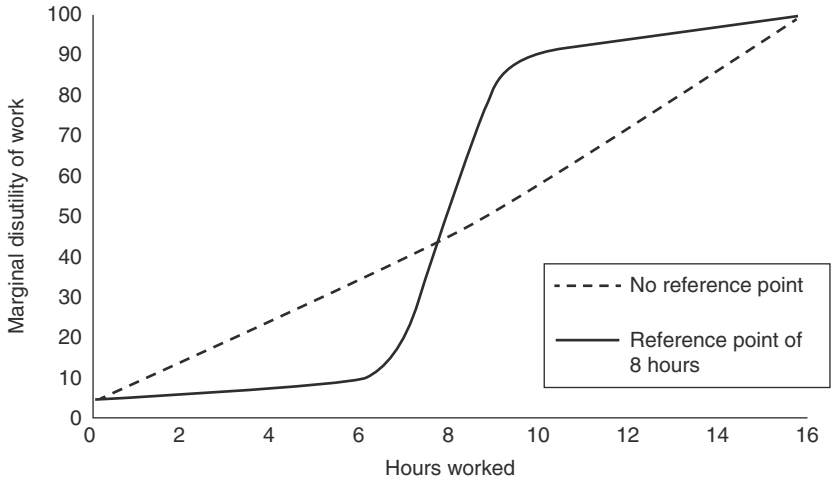

Figure 2.23 Two possible utility functions. If there is a reference level of eight hours' work, the utility function becomes S-shaped and the number of hours worked will not change much as the wage changes.

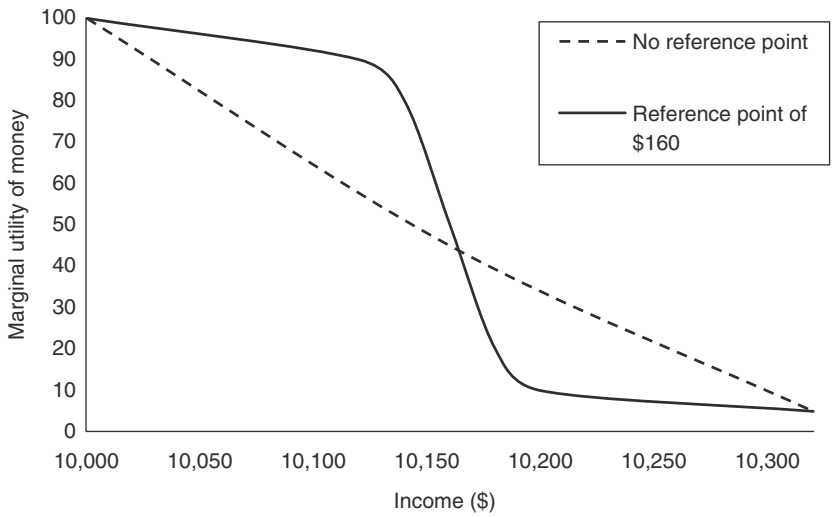

Figure 2.24 If there is a reference level of earning $160, the utility function becomes S-shaped. An increase in wages will cause a decrease in hours worked because the target of $160 can be achieved more easily.

2.6.2 Taxicab drivers

In practice we might expect that Anna has both a target income and a target number of hours she wants to work. This comes out in the first set of data I want to look at. The people who arguably have most control over how many hours they work are self-employed workers. An example that has attracted some interest among

economists is that of taxicab drivers. Many taxicab drivers lease a cab from a taxi-leasing company and keep any money they make from fares or tips, net of fuel costs. By choosing when to start and end the shift, and how many breaks to have, the driver has complete control over how long he works and some control over his income. In New York City taxicab drivers are required to fill out trip sheets that record details on all the trips done during a shift, including times, locations and fares. The trip sheets are a great source of data to look at labor supply. To illustrate, I will look at a study by Farber (2005, 2008) of 21 New York City taxicab drivers working in 1999 and 2000.

What we can focus on is the probability that a driver ends his shift at a particular point in time. Factors that might influence that decision are the hours already worked, income already earned and things such as the weather and time of day. Demand for taxicabs varies a lot across days and within the course of a day, meaning that the wage rate w will fluctuate a lot day to day and during the day. We saw that, if drivers focus on a reference level of hours that they want to work, fluctuations in the wage should make little difference to the number of hours they work. By contrast, if drivers focus on a target level of income, changes in the wage should affect hours worked.

To see what does happen, we can estimate the probability that a driver will stop after each trip. Figure 2.25 summarizes what we get from estimating this probability. It shows the change in the probability that a driver ends his shift, compared to if he has worked nine hours and earned $150–$175. We see that drivers are far more likely to stop after having worked ten or more hours. By contrast, income earned has very little effect.

Figure 2.25 suggests that drivers primarily stop after having worked a certain number of hours, and not after reaching a target level of income. Clearly, this is

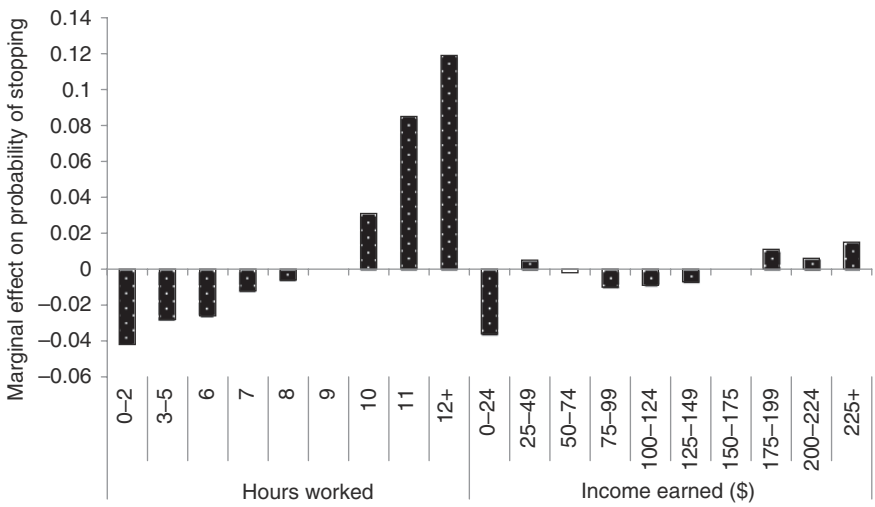

Figure 2.25 The probability that a driver stops his shift compared to when he has worked nine hours and earned $150 to $175. Drivers were significantly more likely to stop if they had worked ten or more hours. The decision to stop was not influenced much by income already earned.
Source: Farber (2005).

consistent with a reference level of hours worked (but see Review question 2.10). It does not, however, completely rule out there being a subsidiary target level of income. To check this, we can plug a target income into the equation that we estimate and see whether we get a better fit with the data. Doing this suggests that drivers do indeed have a target level of income. Somewhat strangely, however, we also find that the target fluctuates a lot from day to day. In particular, the mean target was $196 but the standard error $91. So, a driver might have a target income of $120 one day, $250 the next, and so on. It is not so easy to interpret what this means.

[Extra] The probabilities in Figure 2.25 were estimated by writing the probability that driver *i* on day *j* will stop after trip *t* by

$$P_{ijt} = \Phi(X_{ijt}\,\beta),$$

where Φ is the cumulative normal distribution and $X_{ijt}\,\beta$ is a term that includes things such as the number of hours worked and the time of day. To check whether drivers have a target income, we can slightly change the equation estimated to

$$P_{ijt} = \begin{cases} \Phi\left(X_{ijt}\beta + \delta\right) & \text{if } y_{ijt} \geq r_{ij} \\ \Phi\left(X_{ijt}\beta\right) & \text{if } y_{ijt} < r_{ij}, \end{cases}$$

where r_{ij} is a reference level of income for driver *i* on day *j* and y_{ijt} is current income. If current income is above the reference point, $y_{ijt} \geq r_{ij}$, then, the larger is δ, the higher the probability the driver will stop. Term *S* therefore measures the potential importance of a target level of income. We can assume that the reference point is determined by

$$r_{ij} = R_i + \varepsilon_{ij},$$

where R_i is the average reference level of driver *i*, which may fluctuate day to day. On estimating δ, R_i and β it turns out that δ is significantly greater than zero, and so there does appear to be a target level of income.

2.6.3 Female labor supply

The overall picture we get from the taxicab data is that drivers might have a target level of income, but the number of hours they have worked seems a more dominant factor. I now want to look at some data that gives a slightly different perspective. Most workers do not have as much flexibility as taxi drivers in how many hours a day they work. Often choices are restricted to a more discrete choice of whether to work full-time or part-time or to work at all. But this still raises some interesting issues. For instance, Figure 2.26 plots the growth in labor force participation of women in the United States and United Kingdom since the Second World War. Clearly, there is a significant rise in the number of women choosing to work. Why?

If we look at equation (2.1), then either the wage increased, the utility of money increased or the disutility of working decreased. Changes in wages and the disutility of work can account for a lot of the increase in participation. For example, lower

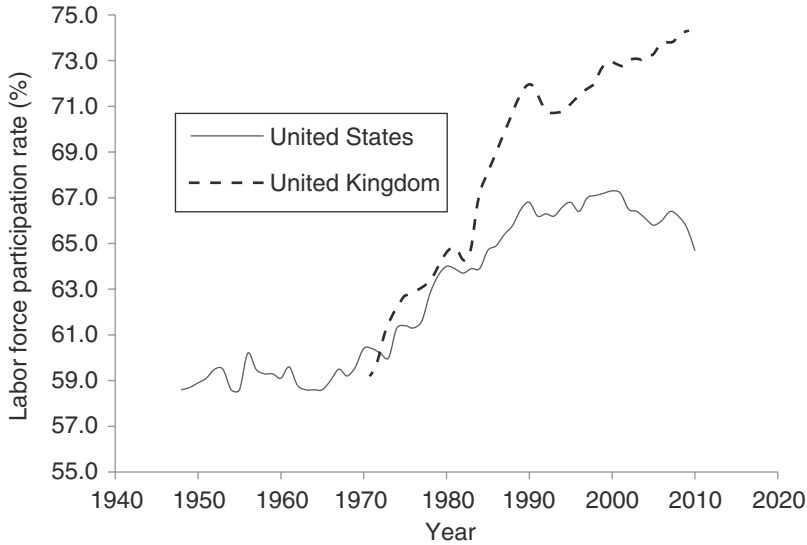

Figure 2.26 The participation rate of women aged 16 and above in the United States and United Kingdom.
Sources: Bureau of Labor Statistics and Office for National Statistics.

fertility and the growth of the service sector are two reasons why the disutility of work could have decreased for women. I want to focus more, however, on the possibility that the utility of money has changed. Why would it have? If people want to earn a similar amount to friends and family, then the fact that some women enter employment (because of the higher wages and decreased disutility of work) could spur other women to enter employment so as to not fall behind.

Neumark and Postlewaite (1998) explore this possibility with data on a sample of US women who become of employable age in the 1980s. They questioned whether a woman is more likely to work if her sister-in-law works, or if her sister's husband earns more than her husband.

[Extra] To see how this is done, let P_{it} denote the probability that woman i is employed at time t. To see whether a woman is more likely to work if her sister-in-law also works we can fit equation

$$P_{it} = \Lambda(X_{it}\beta + \gamma SL),$$

where $X_{it}\beta$ captures things such as education, Λ is the logistic cumulative probability distribution and SL is one or zero depending on whether the sister-in-law works. The bigger is γ, the more likely a woman will work if her sister-in-law works. To see whether a woman is more likely to work if her sister's husband earns more than her husband, we can fit

$$P_{it} = \Lambda(X_{it}\beta + \gamma HM),$$

where HM is one or zero depending on whether the sister's husband earns more.

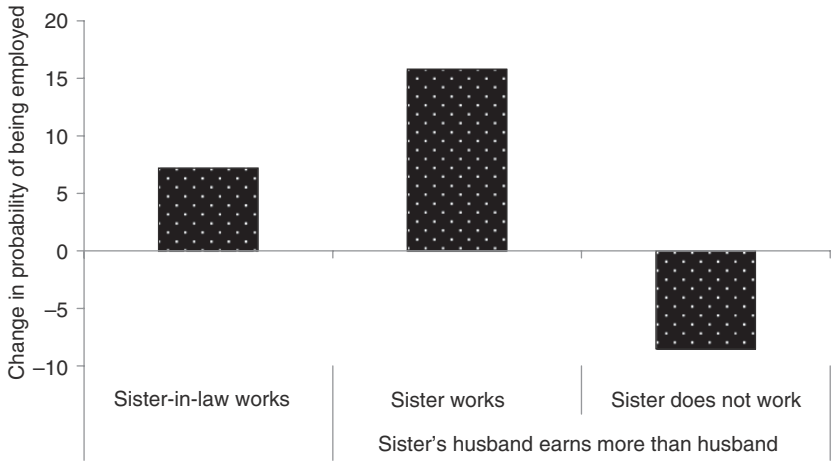

Figure 2.27 The change in probability of a woman being in employment depends on whether her sister-in-law works and her sister's husband earns more or not than her husband.
Source: Neumark and Postlewaite (1998).

Figure 2.27 summarizes the main results. We see that a woman's choice to work was significantly affected by what her sister or sister-in-law was doing. She was more likely to work if her sister-in-law worked, and a lot more likely to work if her sister worked and her sister's husband earned more than her husband! This is consistent with a person, or family, being influenced by a reference level of income that depends on what others are earning.

We have now seen evidence that workers can have a reference level of hours they want to work and a reference or target level of income they want to earn. This is important in understanding individual labor supply but also important in understanding how the labor market works. That's because the reference point will partly determine what is counted as a gain or loss and so also partly determine the equilibrium price and quantity. We see this in the large increase in the number of women working and in the price of a taxicab in the middle of the night.

2.7 The housing market

The second market I want to look at is the housing market. Buying property is one of the bigger decisions most people will make in their life and can be thought of as both a consumption decision and an investment decision. People want somewhere nice to live, but are also aware that the value of a house can increase or decrease, and the more it increases the better. To give an illustration of how much markets can move over time, Figure 2.28 plots the average property price and number of sales by month in England between January 1995 and December 2016. The price clearly increases over this time, but not without periods of decline. Knowing when to sell or buy, and what price to ask or bid, are financially important and complicated decisions. But what is an appropriate reference point to say what is a gain or loss?

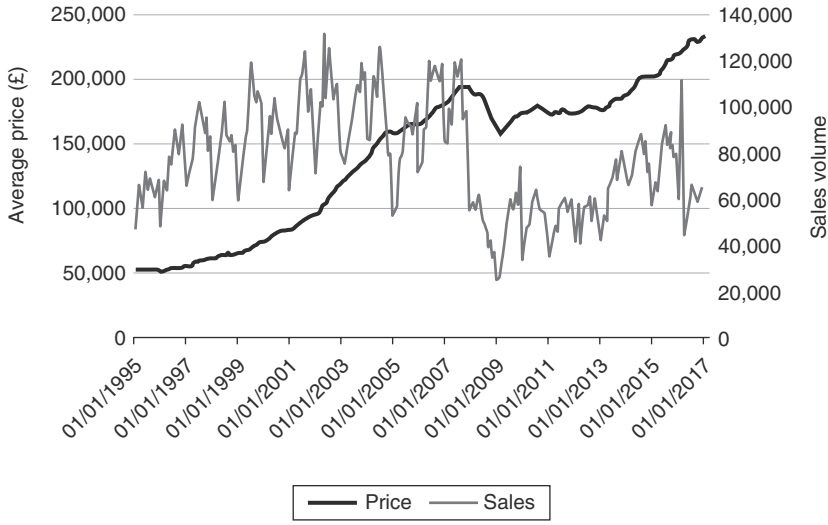

Figure 2.28 The average selling price of residential property and sales volume in England and Wales.
Source: Land Registry.

2.7.1 Reluctance to sell

Let's look first at what price Anna might want to sell her apartment. She could use two natural reference prices: the price that other similar properties are selling for, or the price she originally paid for her apartment. For instance, selling for less than she originally paid might be interpreted as a loss, while selling for more is interpreted as a gain. To see the implications of this, suppose that Anna bought her apartment for $80,000 while John bought an almost identical apartment for $100,000. Both are looking to sell, and the price of other similar apartments is $90,000. For Anna, selling at $90,000 feels like a gain, so she might be happy to sell at this price, or maybe even slightly less. For John, selling at $90,000 feels like a loss. Loss aversion might mean he tries selling his apartment at something like $95,000, or even decides to not sell his apartment at all.

A study by Genesove and Mayer (2001) puts this prediction to the test with data from downtown Boston in the 1990s. They use data from over 5,000 apartments put up for sale between 1990 and 1997. Most importantly, they have data on the price someone bought an apartment for and the price they first listed it when trying to sell. They also have estimates of market value and any outstanding loan at sale. This latter number is important in order to rule out people being reluctant to sell below the price they paid because they have to repay the loan used for purchase.

What Genesove and Mayer find is that someone, like John, who is expected to make a loss asks for a higher price and has to wait longer to sell, but does sell at a higher price. It seems, therefore, that sellers may be averse to selling a house for less than they paid for it. This helps us better understand the housing market. In particular, one puzzling thing we observe in housing data, and that can be seen in Figure 2.28, is that sales volume falls during a slump in prices. If prices fall, we might expect demand

to pick up, not fall. This becomes easier to explain once we account for loss aversion. If prices fall, and people are unwilling to sell at a loss, then we can predict less property being put up for sale and prices being above what they should be. This clearly suggests lower sales.

What is less clear is whether loss aversion benefits a potential seller. To put the magnitude of loss aversion into context, we can plug the numbers estimated by Genesove and Mayer into the Anna and John example. Anna would have asked for $90,000 and sold for this after around 60 days. John would have tried to recoup 25 to 35 percent of his 'loss' and so would have asked for around $93,000 and actually recouped 3 to 18 percent of the loss by selling for, say, $91,000. He would have had to wait around 6 percent longer, but this equates only to an extra four days. It seems as though John does not do so badly!

It is not as simple as that, however, because people who sell also, typically, buy. A slump in the market may be a good opportunity to sell, even if at a loss, and upsize to a more expensive property. Someone who does not want to sell at a loss could miss this opportunity. Or perhaps they need to do a bit of hedonic editing and alter their reference point to focus on the gain they can make on the property they will buy, rather than the loss on the property they will sell.

2.7.2 What are buyers willing to pay?

This brings us on to the buyers. What price should someone be willing to pay for property? The price of similar property may be the most natural reference point in this case. A nice way to see whether it is is to follow people who relocate from one area of the country to another. To illustrate, suppose Anna currently lives in an area where a good apartment costs $90,000 and she is going to relocate due to work. If, where she is moving to, apartments cost around $70,000, they will look cheap, but, if they average $120,000, they will look expensive. Reference dependence suggests that she might be happy to buy something for less than $90,000 but be averse to spending anything more than $90,000.

A study by Simonsohn and Loewenstein (2006) puts this prediction to the test by looking at data on people in the United States relocating between 1983 and 1993. They questioned whether the average rent or house price in the location a person moved from significantly affects the price they pay where they move to. They find that it does matter. A 10 percent difference in the median rent or property price at the location moved from causes an increase of around 2 percent in the rent paid or purchase price at destination. You might say that this could be because of preferences or something else. They also find, however, that, when people move for a second time, this time in the same area, any differences disappear. It seems, therefore, as if the reference point readjusts as people live in the area for some time. Reference dependence can, therefore, really matter.

To finish this section I want to show that anchoring can matter as well. The point was nicely illustrated in a study by Northcraft and Neale (1987). They had subjects view two different properties for sale in Tucson, Arizona. All subjects were given a detailed ten-page brochure full of facts about the property. The only thing that differed between subjects was the reported list price. The properties were actually listed for $74,900 and $134,900, but subjects were given prices around 5 and 12 percent either

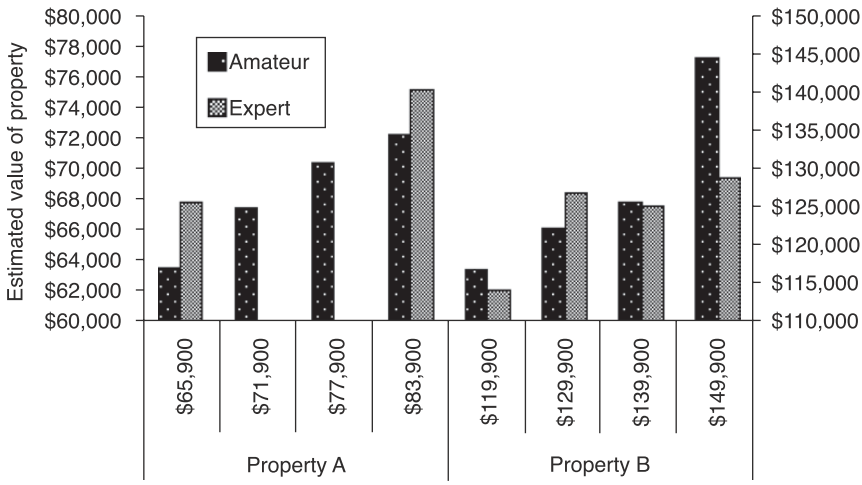

Figure 2.29 The estimated value of two properties depending on the reported list price. If the list price was higher, subjects estimated the value as higher.
Source: Northcraft and Neale (1987).

side of this. For example, some subjects were told the $74,900 property was listed for $65,900 and some that it was listed at $83,900. After viewing the property subjects were asked what value they would put on the property.

Figure 2.29 summarizes the results. An increase in the list price led to a significant increase in the value subjects put on the property. Keep in mind that subjects could view the property and had a host of information at their fingertips. The list price still affected judgment. This is a great example of anchoring and of a framing effect. The study also allows comparison of amateurs at real estate (i.e. students) with professional real estate agents. The amateur subjects were more biased than the experts, but not by much. A more interesting difference showed up when subjects were asked to describe how they came to an estimated value. The experts were less likely to admit that the list price affected judgment than the amateur subjects. This illustrates how blissfully ignorant we often are of framing effects!

2.8 The behavioral life cycle hypothesis

Having spent some time looking at markets I now want to move on to another of the more important and controversial ideas in economics, the life cycle hypothesis. The **life cycle hypothesis** comes in various different guises but the basic idea is relatively simple: a person should take a long-term view about their future income and smooth consumption over their lifetime to maximize expected utility. Figure 2.30 gives a stylized illustration of the model. We see that income goes through large variation over the life cycle, as Anna earns little when a student, has increased income as she works and gets promoted and then earns little income again when she retires. Presumably she will want her standard of living and consumption to be more stable than this (something we shall discuss in Chapter 4). Consequently, she may smooth

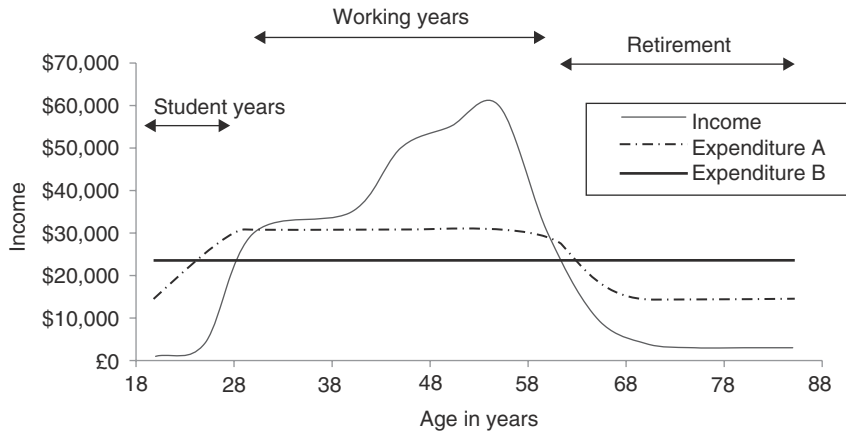

Figure 2.30 An illustration of consumption smoothing. Income varies a lot with age, but consumption can be completely smoothed, as in schedule B, or partly smoothed, as in schedule A.

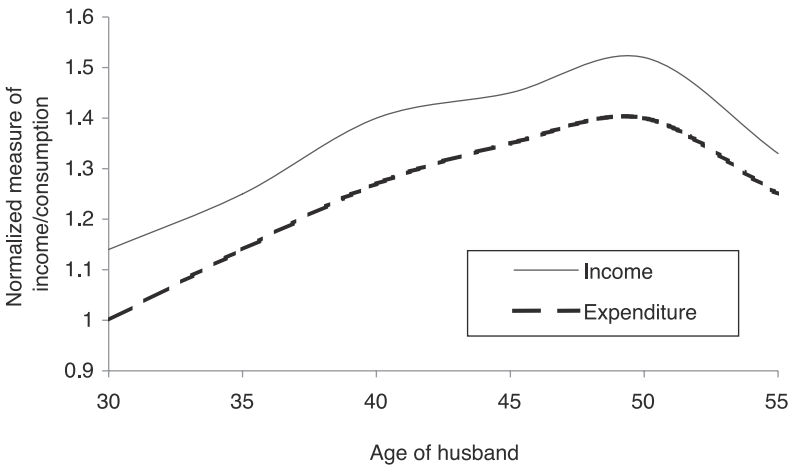

Figure 2.31 An inverse U-shaped consumption and income pattern derived from UK couples with a husband born between 1936 and 1943. Cyclical and growth effects are removed from the data.
Source: Browning and Crossley (2001).

consumption by borrowing in her student years, saving during her working years and using her savings during retirement. Smoothing consumption completely would give schedule B. More realistic is something like schedule A, in which consumption does change with age, because, for example, of having children.

Rather than rely on the stylized illustration in Figure 2.30 we can look at some actual data. When we look at consumption over the lifetime we typically find an inverted U shape like that in Figure 2.31. This inverted U shape has been observed in

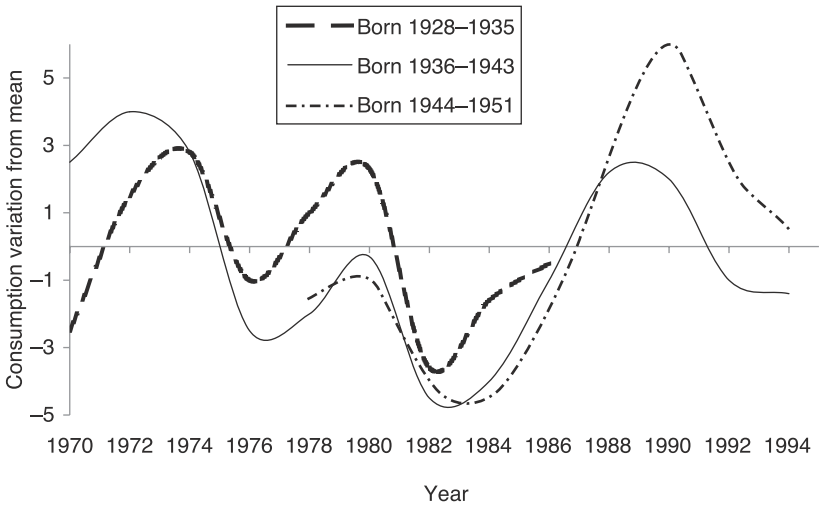

Figure 2.32 Consumption patterns over time derived from UK data.
Source: Browning and Crossley (2001).

different countries, and for all types of workers, education background, etc. An alternative view is obtained by looking at consumption over calendar time. Figure 2.32 shows what typically happens when we do this. It is very noticeable that consumption is cyclical and changes by as much as 11 percent from trough to peak. Unsurprisingly, these cycles coincide with the business cycle.

Figures 2.31 and 2.32 do not look very consistent with the life cycle hypothesis. Instead, it seems as though consumption tracks income with little sign of consumption smoothing. If we factor in the effects of children and a desire to be prudent, or to not want to spend uncertain and unearned future income, then we could argue that the data in Figures 2.31 and 2.32 is consistent with the life cycle hypothesis. But that essentially amounts to arguing that the optimal thing is to not smooth consumption! This seems an unsatisfactory conclusion.

A more realistic perspective is to acknowledge that people undeniably do smooth consumption over their lifetime by saving for retirement, borrowing when a student or borrowing to buy a first home. But they do not smooth consumption as much as we might expect. Instead, consumption seems to track current income more closely and be more sensitive to macroeconomic cycles than would seem optimal. We need to try and understand why this is the case. One potential contributory factor is mental accounting.

2.8.1 Fungibility and mental accounting

Mental accounting and narrow framing are very important in thinking about the life cycle hypothesis. That's because the life cycle hypothesis assumes **fungibility**, meaning that all money is treated the same, no matter where it came from. If people keep mental accounts, then money is likely to be segregated according to how it was obtained, and so would not be fungible.

To illustrate, consider these three scenarios.

You have been given a special bonus at work, meaning you will receive $200 a month over the next year.
You have been given a special bonus at work, meaning you will receive a lump sum payment of $2,400.
You have been told of a distant relative who has left you an after-tax inheritance of $2,400, but you will not receive the money for five years.

In all three scenarios your wealth will increase by $2,400. Fungibility requires that the use to which you put this extra money should not depend on where it came from. You should just think how you will spend the $2,400 over the rest of your lifetime. In reality, things seem different, as the results of a study by Shefrin and Thaler (1988) illustrate. Figure 2.33 shows a stark difference in how much of the extra $2,400 people said they would consume over the following year. If the money came from the regular payment, people planned to consume half of it in one year, but, if it came from a future inheritance, would spend none of it.

Broadly speaking, we can think that households keep in mind three basic mental accounts: a current income account, an asset account and a future income account. The **current income account** is for day-to-day spending. We would expect that the **marginal propensity to consume** from this account, or the proportion of the money spent, is relatively high, because the money is there to be spent. The **asset account** is for saving and investing for things such as retirement or a new house. We would expect a much lower marginal propensity to consume from this account. Indeed, its purpose may be to commit the household to save adequately. Finally, the **future income account** includes predictable future income such as from a pension, inheritance or expected salary increases. The marginal propensity to consume from

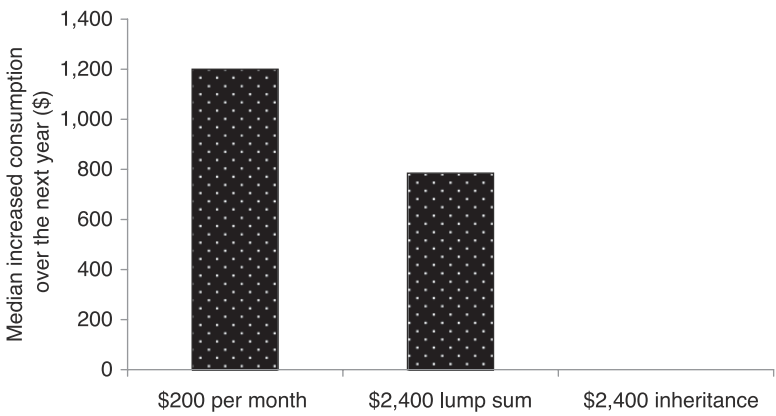

Figure 2.33 The amount of the $2,400 that subjects thought they would spend in the next year. The origin of the money significantly affects the amount that subjects expected to spend.
Source: Shefrin and Thaler (1988).

this account is likely to be very low if people are reluctant to spend income they do not have.

These predicted differences in the marginal propensity to consume are not consistent with fungibility. They are consistent with the data in Figure 2.33, however. A regular payment of $200 a month is more likely to be added to the current income account. A lump sum payment may be saved in the asset account. The inheritance is clearly put in the future income account. This can help explain the decreasing marginal propensity to consume.

Fortunately, we have a wealth of expenditure data with which to move beyond hypothetical surveys to real choices. I will look at three examples.

The first example concerns a tax rebate in the United States in 2001. Most US households over a ten-week period received a rebate of either $300 or $600. The rebate could have been predicted, because the policy was widely publicized and was part of a long-term change in tax policy. There was, therefore, a gap between a person knowing they would get $300 or $600 and actually getting the money. Fungibility requires that people should start spending the money as soon as they know they will get it, and not when they actually get it (unless they have insufficient money). Mental accounting would suggest that receipt of the money moves it from the future income account to either the current income or asset account, and so the money is treated differently once it is received.

Johnson, Parker and Souleles (2006) looked at expenditure data on what households did with the money, and paid particular attention to whether receiving the rebate made any difference to spending. Figure 2.34 shows the change in spending in the three months after the rebate arrived compared to the three months before the rebate. Different households received the money at different times, so there is no possibility that this difference is due to some seasonal effect. It looks as if receipt of

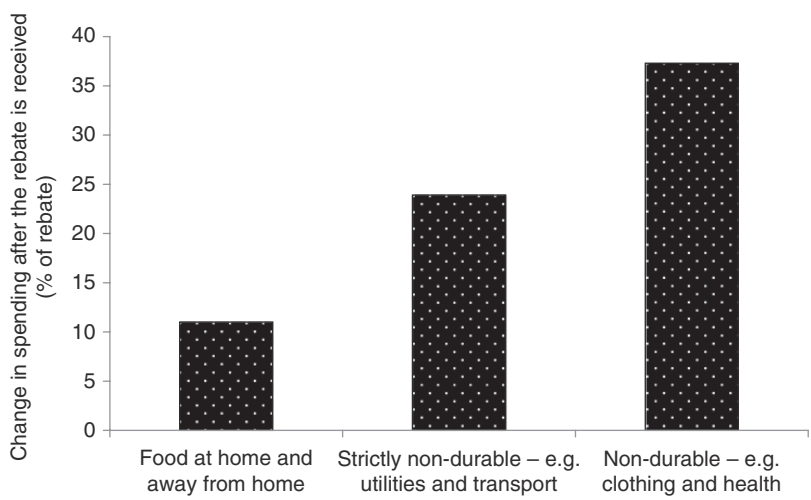

Figure 2.34 The change in spending in the three months after receiving the rebate compared to the three months before receiving the rebate.
Source: Johnson, Parker and Souleles (2006).

the money did make a difference to spending. This is consistent with mental accounting, not fungibility.

Another interesting thing about this rebate is that what households did with it seemed slightly at odds with what they planned to do with it. For example, Shapiro and Slemrod (2003) surveyed households about what they planned to do with the rebate. Only 22 percent of those surveyed expected to spend the rebate. Most expected to pay back debt. This suggests that households planned to put the money in the asset account. This is seemingly confirmed by Agarwal, Liu and Souleles (2007), who showed that households did initially 'save' the money by paying off debt, such as credit card debt. Soon afterwards, however, their spending rose.

We can pick up on this distinction between the current and asset account in the second example I look at. In this example the focus is on Israeli recipients of German restitution payments in the late 1950s. Families received a lump sum 'windfall payment' of varying amounts. Landsberger (1966) looked at what people did with the money and compared it to normal spending. The results are summarized in Figure 2.35. The difference between the marginal propensity to consume from normal income and the windfall is large. As with the tax rebate, this suggests a lack of fungibility. More interesting is how the size of the windfall seemed to matter in terms of how much a family spent of the money. A small payment was spent twice over! A large payment was primarily saved. This illustrates how the amount of money can matter in terms of what mental account it is put into. A small windfall is added to the current account and a big windfall added to the savings account.

The final example I want to look at is very different from the previous two. It demonstrates how accounts can be more narrowly defined than the three accounts I have mentioned so far. Often government benefits are given for a specific reason. For example, child benefit is given to families with children. Fungibility would suggest that

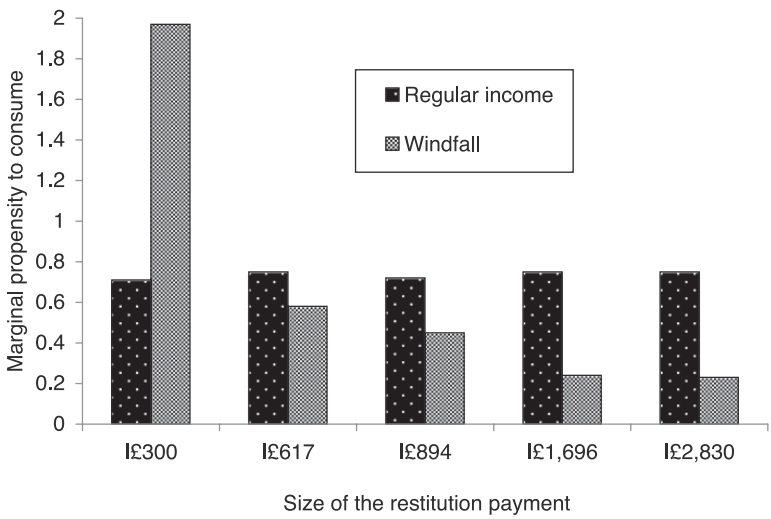

Figure 2.35 The marginal propensity to consume the windfall gain compared to normal income of Israeli recipients of German restitution payments.
Source: Landsberger (1966).

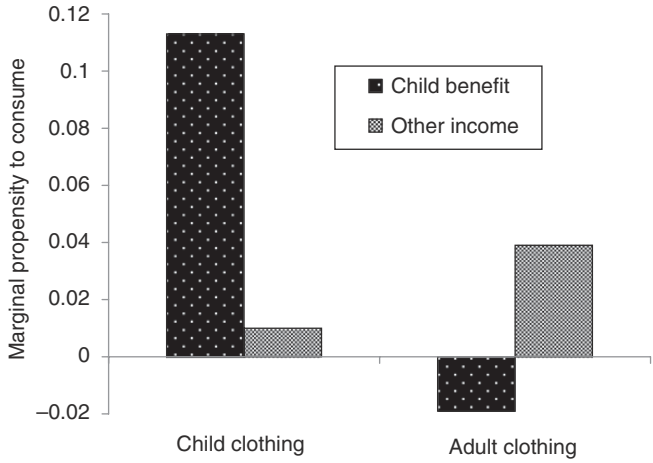

Figure 2.36 The marginal propensity to consume clothing out of normal income and child benefit.
Source: Kooreman (2000).

households should ignore the reason they got the money. Mental accounting suggests they may not. Kooreman (2000) looked at data from the Dutch child benefit system from 1978 to 1994.

Over these years various changes to the benefit system allow us to try and estimate whether child benefit is used to buy things for children. In the study, the focus was put on the amount spent on clothing in families with one child. Figure 2.36 summarizes the main finding. The marginal propensity to consume child clothing out of child benefit was much larger than that from other income. By contrast, the marginal propensity to consume adult clothing out of child benefit was much less than that from other income. It seems that child benefit was spent on the child.

All three of these examples show a lack of fungibility. Households treat money differently depending on where it comes from, how much it was and the purpose for which it was given. There are lots of other examples I could give to demonstrate the same basic conclusion (see Review question 2.12). But, given what we saw in section 2.2, it should come as no surprise that money is not fungible. If people keep mental accounts, then it is natural that money should not easily flow from one account to the next. The question we need to get back to is what implications this has for the life cycle hypothesis.

One thing we have seen is that people seem reluctant to spend future unearned income. Only when they got the money did households spend the tax rebate and restitution payment. We know, more generally, that people do spend unearned income. Student debt and families taking out a mortgage to buy a house are ample evidence of this. The point is, however, that people may spend less of their unearned income than the life cycle hypothesis might suggest they should. Given that most people get increasing amounts of money through their lifetime, this would mean spending tracks income more than would be predicted by the life cycle hypothesis. This is what we observe.

2.9 Saving for the future

One particularly important aspect of the life cycle hypothesis is saving. People need to save enough for retirement, to buy a house or new car, and so on. A survey of faculty and staff at the University of Southern California by Benartzi and Thaler (1999) finds that 58 percent of those surveyed spent less than one hour determining their contribution rate and investment decisions to a savings plan. This is presumably quite typical, but these are important, complicated decisions, and so it is slightly worrying that so little time is spent making them! Put another way, however, people must be using simple heuristics to help them make decisions so quickly. Maybe these heuristics help them make the right decisions.

To illustrate the heuristics we do observe, Figure 2.37 plots the distribution of contribution rates of those joining a large retirement savings plan in 2001 or 2002. The contribution rate is the proportion of income that people allocate to the savings plan. Three heuristics are apparent in this data. First is 'Pick a multiple of five', as seen in the notable spikes at 5, 10 and 15 percent. Second is 'Pick the maximum allowed'. In 2001 the maximum allowed in this contribution plan was 16 percent, but in 2002 the maximum was increased to 100 percent. We can see that this caused a big fall in the proportion choosing 16 percent. Finally, in most contribution plans, the employer will match the contribution of an employer up to some limit. The third heuristic is 'Pick the maximum to get a full employer match'. In this plan that maximum was 6 percent, explaining why so many chose this contribution rate.

With these three heuristics we can get a good feel for why the contribution rate is distributed as in Figure 2.37. Next I want to look in a bit more detail at one important kind of heuristic we have not yet seen.

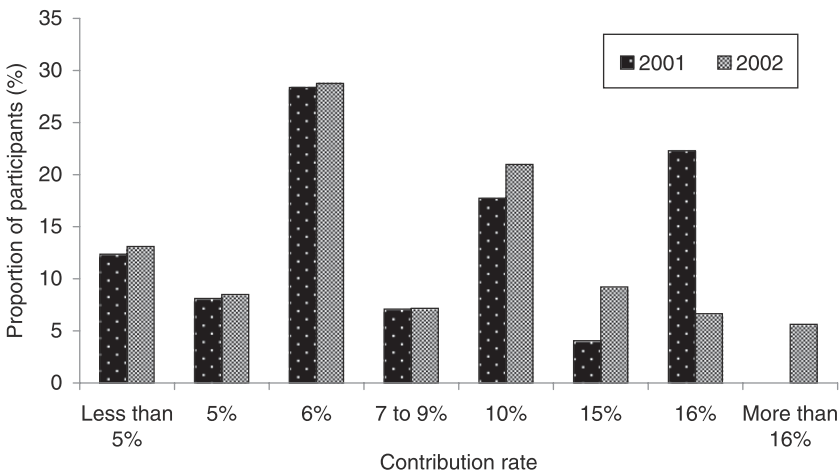

Figure 2.37 The contribution rate of people enrolled in a large defined contribution plan.
Source: Benartzi and Thaler (2007).

2.9.1 Let's diversify

As well as deciding how much to save, an investor also needs to decide where to save, or invest. The typical contribution retirement savings plan will give the saver a set of funds or investment opportunities that they can invest the money into. For example, Benartzi and Thaler (2001) looked at a database from Money Market Directories that covers 170 retirement savings plans with 1.56 million participants. The average number of investment options available in a plan was around six to seven. For example, a plan might allow participants to split investments between an equity, international equity, corporate bonds, balanced growth, government bonds and money market funds.

Deciding where to invest is a complex choice. One option is to use a simple **diversification strategy**, called the 1/n heuristic. The **1/n heuristic** is to split the amount to be saved or invested equally among the available funds. A simple heuristic, but one that leaves the saver exposed to any bias in the options on offer. For example, compare a plan that has three equity options and one money option to one with two equity options and two money options. In the first, someone using the 1/n heuristic would invest 75 percent in equity, and, in the second, they would invest 50 percent. That's a big difference.

Do we observe people using the 1/n heuristic? Benartzi and Thaler categorized the investment plans in their data as having a low, medium or high proportion of equity fund options. Figure 2.38 compares the proportion of equity fund options with the proportion invested in equities. We can immediately see that things are not as simple as the 1/n heuristic would suggest, because the proportion invested is more consistent across plans than the proportion of equity fund options offered. Even so,

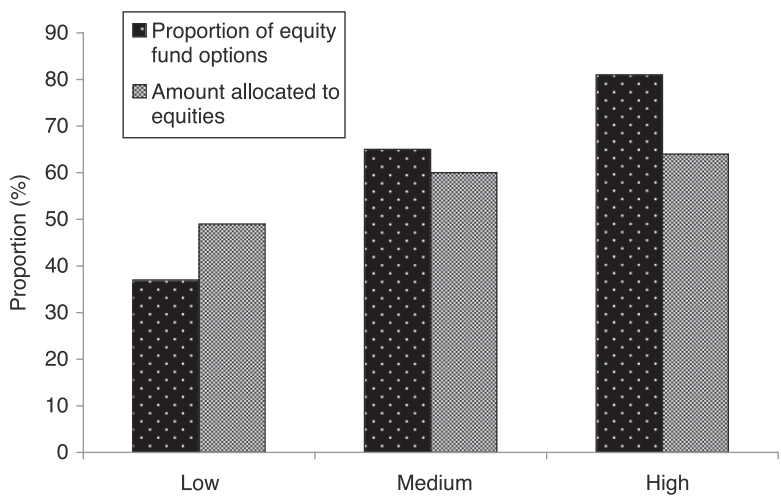

Figure 2.38 The proportion of funds that were based on equities and the proportion of savings that savers invested in equity. Those in plans with more equity funds invested more in equity.
Source: Benartzi and Thaler (2001).

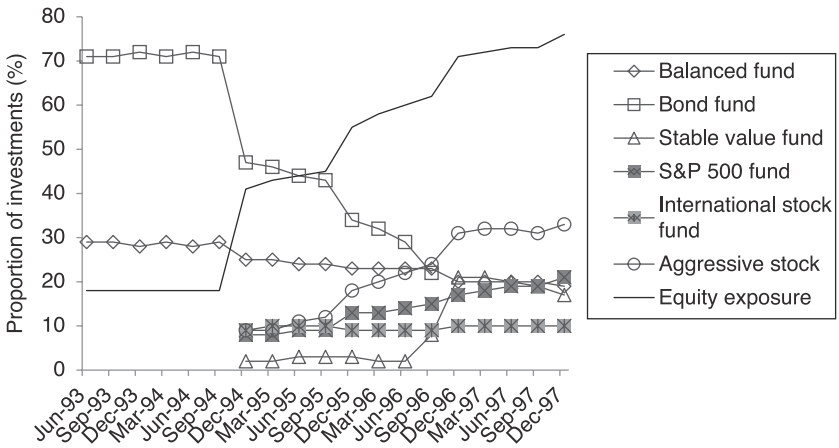

Figure 2.39 The proportion of investments in each fund of a company's contribution savings plan. In 1994 new funds were added and in 1996 the bond fund was withdrawn. The amount invested in equities increased a lot as a consequence. *Source:* Benartzi and Thaler (2001).

those in a plan with a larger number of equity fund options do invest significantly more in equities.

It is hard to justify such a systematic difference in investments without people being biased by the options they have to choose from. To illustrate this further, Benartzi and Thaler analyzed data from a specific company. The company is interesting because, at the end of 1994 and 1996, the fund options available to investors were changed. Figure 2.39 charts the change in investment patterns over time. In 1994 four new fund options were made available. In 1996 the bond fund was withdrawn. You can see that by the end of 1997 investments are quite evenly split among the remaining possible funds. The remarkable thing about this is that it means there was a huge shift in exposure to equity funds. In 1993 only 18 percent of investments were in equity but by 1997 that had risen to 76 percent. It is highly unlikely that this shift is because of a change in preferences or the optimal portfolio!

2.9.2 Let's not diversify

The 1/n heuristic makes most sense if n is not very large – i.e. the number of options is small. What if the number of options is large? Then we might get a **conditional 1/n heuristic** of picking a small number of funds and dividing investments equally amongst these. Huberman and Jiang (2006) looked at the investment decisions of over half a million participants of defined contribution plans in 2001. Figure 2.40 compares the number of funds participants chose to invest in versus the number of funds offered. The majority of participants were in plans with over ten funds to choose from. Despite this, the majority of participants chose to invest in only one to five of the possible funds. Investors were clearly selective, therefore, in where they invested, consistent with the first part of the conditional 1/n heuristic.

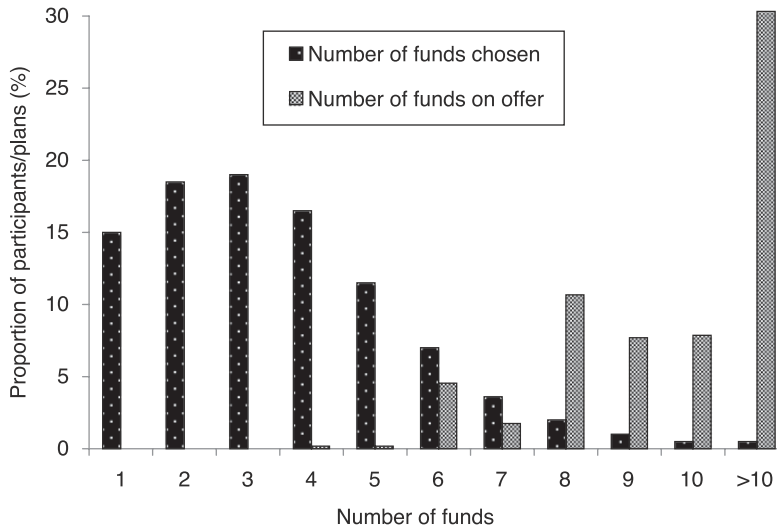

Figure 2.40 The number of funds that savers invested in and were on offer in a large sample of defined contribution plans. A typical plan offers ten or more funds. Most savers invest in one to five funds.
Source: Huberman and Jiang (2006).

To check the second part of the conditional 1/n heuristic we need to see how participants allocated their investments between the funds they did invest in. This is not simple to do, because the value of funds goes up or down over time, and so, even if an investor did split her investments equally across funds initially, it may not look this way a few years later. To compensate for this, Huberman and Jiang looked at new participants. Figure 2.41 summarizes for each choice of number of funds the proportion of new participants who invest consistent with the 1/n heuristic. We see that a high proportion of those investing in two, four, five and ten funds do seem to use the 1/n heuristic. Intuitively, this makes sense, because then 1/n is a round number. Figure 2.41 is misleading as to the total proportion who use the 1/n heuristic because we need to factor in the proportion of participants who chose one, two and three funds, etc. Doing this we find that over 70 percent of participants behaved consistently with a conditional 1/n heuristic.

Figure 2.38 suggested investors may be biased towards investing in equity if the proportion of equity funds offered is high. It was clear, however, that the differences in investments chosen were smaller than those of investments offered. Huberman and Jiang draw a similar conclusion, finding only a small positive correlation between the number of equity funds offered and the number chosen. One possibility is that, when the number of offered funds is small, as in Figure 2.39, investors are biased by the number of equity funds on offer, but, if the number of funds is large, investors are not so biased by the proportion of equity funds.

Whether this is true or not, the picture we get is one of investors using simple heuristics to make complex choices. What's not so clear is whether these heuristics are optimal for the investor. It's hard to believe that, by some happy coincidence, a

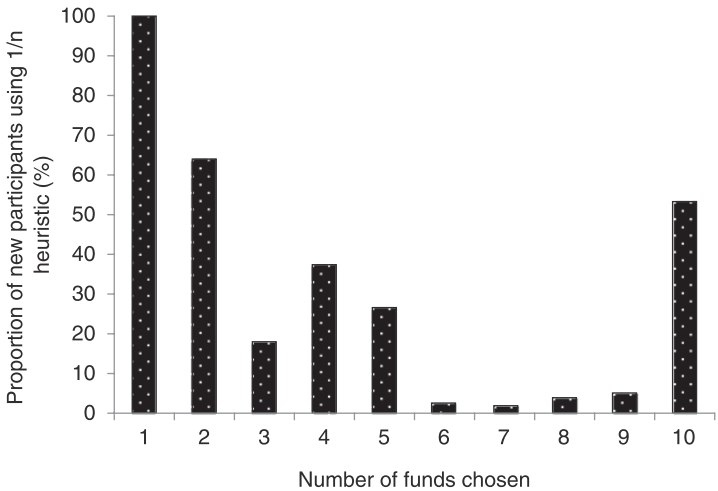

Figure 2.41 The proportion of new participants whose behavior was consistent with the conditional 1/n heuristic – i.e. split their investments equally between the funds they invested in.
Source: Huberman and Jiang (2006).

contribution rate of 10 percent, or a split of 75 to 25 percent between equities and money, is optimal. Maybe, however, a contribution rate of 10 percent is nearly as good as the true optimum of, say, 9.34 or 11.28 percent. Maybe it is not so bad that investors increase their exposure to equity because of a change in the funds on offer. These are important questions, and so we will come back to them in subsequent chapters.

2.10 Further reading

There is an overlap between this chapter and the next in terms of further reading for sections 2.1 and 2.2. That said, the articles by Tversky and Kahneman (1981, 1986), Kahneman and Tversky (1983), Rabin (1998), Kahneman (2003) and Thaler (2008a, 2008b) are particularly worth a look. The book by Kahneman (2011) is a must-read! To get a different perspective, the work on fast-and-frugal heuristics is summarized in the book by Gigerenzer, Todd and ABC Research Group (1999) and the review article by Hafenbrädl and co-authors (2016). One issue I covered relatively briefly is search. The paper by Schunk (2009) provides an interesting recent addition to the literature connecting search with loss aversion, while Caplin, Dean and Martin (2011) provide evidence on satisficing. For a recent paper exploring how context effects can be consistent with rationality, see Ok, Ortoleva and Riella (2014). The book by Ariely (2012) is a good place to start for more on lying and deception. See also Rosenbaum, Billinger and Stieglitz (2014). For more on markets, the articles by Vernon Smith (2002, 2003) are a good starting point. Farber (2015) provides a recent update on the evidence of New York taxicab drivers. For more on the life cycle hypothesis and savings, you could start with Shefrin and Thaler (1988), Thaler (1990), Browning and Crossley (2001) and Benartzi and Thaler (2007).

2.11 Online material

On the companion website you will find access to flashcards, quizzes and answers to the review questions.

2.12 Review questions

2.1 If you wanted to find a new mobile phone, outfit, car, house or job, how would you search for one?
2.2 Why might a company wanting to advertise a product be interested in the trade-off contrast hypothesis and extremeness aversion?
2.3 Is it possible that Anna could be influenced by both trade-off contrast and extremeness aversion simultaneously?
2.4 Why is time important in reference-dependent utility? Compare how happy Carol and Amanda will feel the day they get the report from their broker, the next day, after a month, etc.
2.5 List all the heuristics that have appeared in the chapter.
2.6 Argue that, if true, the 'no loss in buying' hypothesis means the WTP valuation is more reliable than the WTA valuation.
2.7 Why would we expect EG to be greater than WTA?
2.8 Suppose someone's willingness to pay for a good is $10 and their reference point is $20. If the good is $13, will they buy it? What does this tell us about sales and 'bargain buys'?
2.9 Write down a list of occasions when you could have gained from lying or deception. Why can context influence whether or not you lie?
2.10 How is it possible to detect lying in the lab when subjects are given tasks to do and self-report their success?
2.11 What effect do you think the internet has had on the validity of the law of one price?
2.12 What is ecological rationality and why is it important in understanding the efficiency of markets?
2.13 Does the fact that a taxi driver works nine hours a day mean he has a reference level of hours he wants to work?
2.14 Using search, contrast effects and reference dependence, suggest why a person who lives in an area where apartments cost around $200,000 may be reluctant to buy when she moves to an area where apartments cost around $300,000.
2.15 Give reasons why fungibility may not hold.
2.16 Suppose that Anna owns her own house and house prices go up. Should she spend more? Now, suppose she has money invested in the stock market and share prices go down. Should she spend less?
2.17 What are the advantages and disadvantages of using the 1/n heuristic to choose investments?

Chapter 3

Choice with risk

> The safest way to double your money is to fold it over once and put it in your pocket.
> Kin Hubbard

Economic choices are almost always made with some uncertainty as to what the outcome will be. A person buys groceries without knowing for sure how tasty they will be. He buys a new car without knowing how it will perform or how long it will last, a plane ticket without knowing whether the plane will be delayed, house insurance because he does not know whether his house will be burgled, and he invests in shares without knowing whether they will increase or decrease in value. In the last chapter we saw that uncertainty can lead to choice arbitrariness and all the consequences that entails. What I want to do in this chapter is look in more detail at some other important consequences of risk.

Before we get started there is one distinction I need to explain. We say that someone faces a **situation of risk** if they know what could happen and how likely it is. An example would be someone who bets $10 on the toss of a coin; they know that there is a 50:50 chance it could be heads or tails, and, if it's heads, they win $10 and, if it's tails, they lose $10. We say that someone faces a **situation of uncertainty** if they do not know some of the possible outcomes or how likely they are. An example would be someone booking a plane ticket, who is unlikely to know all the possible delays or problems that could happen to change their experience of the flight.

Most of the situations we face are ones of uncertainty. Even the toss of a coin could be biased in many different ways. It is more difficult to model situations of uncertainty than ones of risk, however, and without knowing the consequences of risk we cannot get very far thinking about uncertainty. It is traditional, therefore, to focus on situations of risk, and that is what I do in this chapter. That makes our task manageable, and, as we shall see, still gives us a lot to think about. But do not be too disappointed, because in Chapters 5 and 9 we will come back to look at choice with uncertainty.

3.1 Expected utility

To illustrate some of the issues, imagine someone called Alan deciding whether or not to insure his car. Table 3.1 summarizes the choice he has to make. This is a choice with risk, because he cannot know for sure whether the car will be stolen or will be

Table 3.1 The possible consequences of Alan buying insurance.

	Car is stolen	Car is in an accident	Stopped by police	No theft, accident or police
Probability	0.05	0.05	0.10	0.80
Final wealth if buys full insurance	\$61	\$61	\$61	\$61
Insurance against theft	\$64	\$14	\$64	\$64
No insurance	\$20	\$20	\$30	\$70

in an accident, or whether he will be asked to prove he has insurance. He begins with \$20 in money and his car is worth an additional \$50 (you can multiply these numbers by \$1,000 if you wish). Full insurance costs \$9, and guarantees his final wealth will be \$61 = 20 + 50 − 9. Insurance against theft costs only \$6 but does not insure against an accident. No insurance costs nothing but leaves him vulnerable to theft, an accident or being fined \$40 by the police for having no insurance. Should he insure his car?

More generally, we shall be interested in things called prospects. A **prospect** is a list of probabilities that things will happen, together with the monetary payoff if they do happen. In the insurance example there are three prospects. The 'full insurance prospect' gives probability one of having \$61. We can write this as (1, \$61). The 'insurance against theft prospect' gives probability 0.95 of \$64 and 0.05 of \$14. We can write this as (0.95, \$64; 0.05, \$14). Finally, the 'no insurance prospect' gives probability 0.1 of \$20, 0.1 of \$30 and 0.8 of \$70, which we write as (0.1, \$20; 0.1, \$30; 0.8, \$70). In general, we would summarize a **prospect** with the list $(p_1, x_1; p_2, x_2; \ldots; p_n, x_n)$, where p_i is the probability of getting monetary payoff x_i.

The reason we are interested in prospects is that this is what a person must choose. In the insurance example, Alan must decide whether to choose the 'full insurance', 'insurance against theft' or 'no insurance prospect'. We need, therefore, a model of how people choose prospects. The standard way to do this is to use expected utility.

To calculate expected utility we make use of a utility function u that translates money into utility. You can think of this as the same utility function we saw in the last chapter; it is just that, to make life easier for ourselves, we ignore things other than money. So, if Alan has x in money, his utility is $u(x)$. The **expected utility of a prospect** $A = (p_1, x_1; p_2, x_2; \ldots; p_n, x_n)$ is then found using the formula

$$U(\mathrm{A}) = \sum_{\mathrm{i}=1}^{\mathrm{n}} \mathrm{p_i u(x_i)}. \tag{3.1}$$

We can then predict that people will choose the prospect with the highest expected utility. This gives us a simple way to predict choice with risk.

To see how this works in the insurance example, I shall contrast the choices Alan would make for three different possible utility functions. Table 3.2 provides the numbers and Figure 3.1 plots the relevant functions. With a linear function we see that insurance against theft gives the highest expected utility. If the utility function is sufficiently concave, then full insurance gives the highest expected utility, and, if it is sufficiently convex, no insurance gives the highest expected utility.

Table 3.2 Alan's expected utility depends on his utility function.

Utility function	Expected utility if chose:			Predicted choice
	Full insurance	Insurance for theft	No insurance	
Linear: $u(x) = x$	61.0	61.5	61.0	Insurance for theft
Concave: $u(x) = 10\sqrt{x}$	78.1	77.9	76.9	Full insurance
Convex: $u(x) = x^2/50$	74.4	78.0	81.0	No insurance

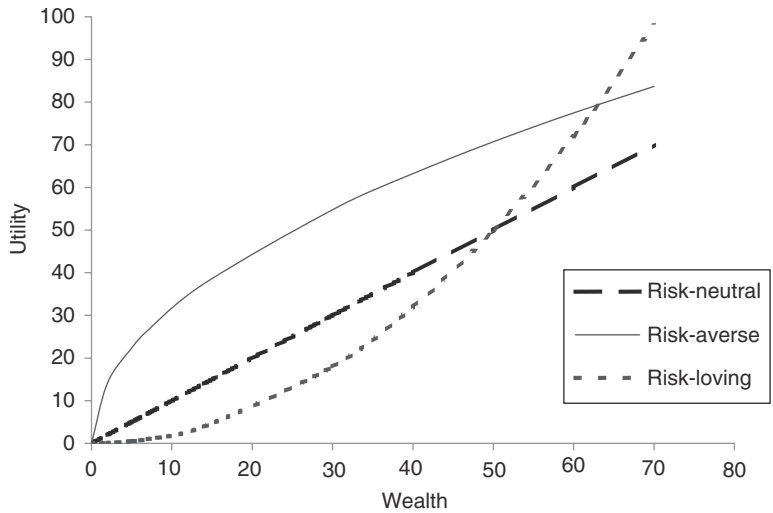

Figure 3.1 Three possible utility functions: one is concave, which would imply Alan is risk-averse, and one convex, which would imply Alan is risk-loving.

The example illustrates how the prospect that gives the highest expected utility will depend on the shape of the utility function. To explain why, I should first define the **expected value of a prospect A** as

$$V(A) = \sum_{i=1}^{n} p_i x_i. \tag{3.2}$$

The only difference between equations (3.1) and (3.2) is that we do not transform monetary payoffs into utility payoffs. This means that we obtain expected monetary payoff rather than expected utility payoff. In the example, insurance for theft has the highest expected value. It is not clear, however, that choosing the prospect with the highest expected value is best. This is because Alan will want to also take into account the riskiness of the prospect.

The relevant trade-offs are illustrated in Table 3.3. With full insurance, Alan knows his payoff will be \$61. There is no risk, and Alan might like this. Indeed, he might like it so much that he is willing to sacrifice \$0.50 of expected value to get it. In this case,

Table 3.3 The expected value and risk of different choices.

	Expected value	Minimum possible	Maximum possible	Probability wealth below expected value
Full insurance	\$61	\$61	\$61	0.00
Insurance against theft	\$61.5	\$14	\$64	0.05
No insurance	\$61	\$20	\$70	0.20

we would say that he is risk-averse. Someone is **risk-averse** if they prefer a certain amount of money to a prospect with the same expected value; they would rather avoid risk. Going to the other extreme, it may be that Alan likes to gamble and prefers to choose no insurance because that is where the risk is highest. In this case, we would say that he is risk-loving. Someone is **risk-loving** if they prefer a risky prospect to the expected value of the prospect for sure; they would prefer risk.

Whether or not Alan is risk-averse or risk-loving will depend in a simple way on the curvature of his utility function. If his utility function is concave, Alan loses relatively more if his wealth goes down than he gains if his wealth goes up, so he would rather not risk a loss for a gain, and he is risk-averse. If his utility function is convex, then he loses relatively less if his wealth goes down than he gains if his wealth goes up, so he would risk a loss for a gain and he is risk-loving. The curvature of the utility function tells us a lot, therefore.

There are two commonly used ways to measure this: one measures **absolute risk aversion** and the other **relative risk aversion**. It's not so important that you know exactly what these are, but for completeness the formulas are

$$\text{absolute risk aversion: } r_u^a = \frac{u''(x)}{u'(x)}; \text{ relative risk aversion: } r_u^r = -x\frac{u''(x)}{u'(x)},$$

where $u'(x)$ denotes the derivate and $u''(x)$ the second derivate of $u(x)$. If Alan is risk-neutral, then $u(x) = x$ and so $r_u^a, r_u^r = 0$. But, if he is risk-averse, $r_u^a, r_u^a > 0$, and the more risk-averse he is, the larger are r_u^a and r_u^r. Various methods can be used to measure risk aversion (see Research methods 3.1).

Economists like things simple, and so two utility functions are of particular note because they imply, respectively, **constant absolute risk aversion** (CARA) or **constant relative risk aversion** (CRRA). It's even less important that you know what these are, but here they are:

$$\text{CARA}: u(x) = -e^{-r_u^a x}; \quad \text{CRRA}: u(x) = \frac{x^{1-r_u^r}}{1-r_u^r}$$

The beauty of the CARA and CRRA functions is that we need to know only one number, the level of risk aversion, and then we can model choice with risk, because we know what the utility function is and can use expected utility to find what maximizes utility. No wonder that expected utility has become the standard way of modeling choice with risk: if we know Alan's level of risk aversion, then we can predict his choice over any set of risky prospects!

❖ RESEARCH METHODS 3.1 MEASURING RISK AVERSION

Given the risk inherent in many economic choices, it is important to know the level of risk aversion, even if our primary interest is not how a person reacts to risk. There are lots of methods that have been used to measure risk aversion, and we will see many of them as we go through the rest of the chapter. The basic idea behind most methods is the same: choices are observed and fed into either the CARA or CRRA utility function to estimate the parameter r_u^a or r_u^r. The differences between methods are the choices observed. Needless to say, the experimental laboratory is one setting that has been used. In this case, we look at prospects with relatively small amounts of money. Empirical settings in which we can measure risk aversion include asset and insurance markets. In this case, we can observe choices over prospects with both relatively small and large amounts of money (see section 3.7).

An intriguing and slightly different way to measure risk aversion was suggested by Chetty (2006). This method does not use choices over prospects but, instead, looks to directly estimate the curvature of the $u(x)$ function. This is, after all, what we need to calculate r_u^a or r_u^r, because the more curved is $u(x)$ the more risk-averse a person is. The curvature of $u(x)$ is estimated by looking at labor supply. To explain how this works: the more curved is $u(x)$ the more quickly a person becomes relatively satiated with money – that is, has enough money to not want more; if someone is satiated with money, then they will not want to risk losing money to gain more, they are risk-averse and they have no desire to increase their income, so can work less if the wage rate increases. By looking at how people react to changes in wages we can therefore estimate risk aversion.

The problem that economists face is that different studies have given very different estimates for the level of risk aversion, and we are not talking small differences. For example, Chetty estimates that r_u^r must be less than two, but we shall see estimates later in this chapter as high as 1,000. This suggests that there is no real measure of risk aversion, just measures that are appropriate for particular situations.

Expected utility does give us, therefore, a very simple and transparent way to model and understand how people make choices when there is risk. Maybe, however, it is a bit too simple. Is it possible to capture all the things that can influence choice, and that we are interested in predicting and modeling, in just one number? The only way to find out is to look at some data and see. To get us started on this I am going to give a series of examples that suggest that expected utility cannot explain everything we typically observe. These examples will serve as a precursor to us thinking more generally about how to move beyond expected utility. As we go through the examples, I would strongly encourage you to 'play along' and think what you would choose and why.

3.1.1 The Allais paradox

Consider the prospects in Table 3.4. When asked to choose between prospects A and B, most people choose prospect B. When asked to choose between prospects C and D, most people choose prospect C. This seems entirely sensible, but is not consistent with expected utility theory: the **Allais paradox**.

To see why this is a paradox, first observe that prospect C is obtained from prospect A, and prospect D is obtained from prospect B, by removing a 0.66 chance of winning \$2,400. Now, if Alan chooses B over A (and put $u(0) = 0$),

$$u(2{,}400) > 0.33\, u(2{,}500) + 0.66\, u(2{,}400),$$

which (removing the 0.66 chance of winning \$2,400) means that

$$0.34\, u(2{,}400) > 0.33\, u(2{,}500). \tag{3.3}$$

If, however, he chooses prospect C over D,

$$0.33\, u(2{,}500) > 0.34\, u(2{,}400). \tag{3.4}$$

Equations (3.3) and (3.4) cannot both be true.

So, what causes the Allais paradox? There is a sense in which removing the 0.66 chance of winning \$2,400 matters more for prospect B than A because changing a sure thing to a risk matters more than changing a risk to a risk. Generalizing from this, the evidence suggests that people favor outcomes that are perceived certain rather than probable or possible. This is called the **certainty effect**.

Interestingly, the certainty effect can be relevant even if outcomes are not theoretically certain. To illustrate, consider Table 3.5, with four more prospects. Most people choose prospect F over E and prospect G over H. This is, again, inconsistent with expected utility theory. It does seem consistent with a certainty effect because a

Table 3.4 The Allais paradox. Many people prefer B to A and C to D.

Prospect	Amount with probability of outcome		
A	\$2,500 with probability 0.33	\$2,400 probability 0.66	\$0 probability 0.01
B	\$2,400 for sure		
C	\$2,500 with probability 0.33		\$0 with probability 0.67
D	\$2,400 with probability 0.34		\$0 with probability 0.66

Table 3.5 The near-certainty effect. Many people choose F over E and G over H.

Prospect	Amount with probability of outcome	
E	\$6,000 with probability 0.45	\$0 otherwise
F	\$3,000 with probability 0.9	\$0 otherwise
G	\$6,000 with probability 0.001	\$0 otherwise
H	\$3,000 with probability 0.002	\$0 otherwise

Table 3.6 Prospects that can lead to extreme risk aversion.

Prospect	Amount with probability of outcome	
I	$0 for sure	
J	$100 loss with probability 0.5	$105 gain otherwise
K	$100 loss with probability 0.5	$125 gain otherwise
L	$100 loss with probability 0.5	$200 gain otherwise
M	$225 loss with probability 0.5	$375 gain otherwise
N	$600 loss with probability 0.5	$36 million gain otherwise

0.9 chance might look relatively certain compared to a 0.45 chance. What is perceived as certain can thus depend on the context.

3.1.2 Risk aversion

Next consider the prospects in Table 3.6. Anyone who prefers prospect I to prospects J to N is risk-averse. They prefer $0 for sure to a prospect with expected value greater than $0. The further we go down the list of prospects, the more risk-averse someone would have to be to prefer prospect I. Many people choose prospect I over prospect J. More generally, many people display risk aversion for relatively small gambles.

This, in itself, is consistent with expected utility. The problem comes when we try to infer what such a choice implies a person would do when choosing amongst prospects with bigger stakes.

❖ RESEARCH METHODS 3.2 MONEY OR SOMETHING ELSE

Experimental subjects are almost always given financial incentives. This means that at the end of the experiment they are given cash, and the knowledge that they will be given cash is supposed to be incentive enough to think about their choices. For economists, this is easy to justify, because money is a unit of exchange and easily substituted for other goods. Maybe, however, subjects would behave differently if the incentive were something other than money? We have already encountered an experiment that gave chocolates as well as money, but this does not really answer the question of interest here. What we need is an experiment in which subjects are provided with a monetary incentive in one treatment and a different incentive in another treatment.

Harrison, List and Towe (2007) did such a study on risk aversion. The study was done at the Central Florida Coin Show and involved subjects choosing amongst prospects that had either monetary or coin outcomes. In a monetary treatment the possible prize was money, and in a coin treatment it was a collectable coin. The coins that could be won had a retail value of $40, $125, $200 and $350. The money prizes were set to match these amounts.

The study showed very similar estimates of risk aversion whether subjects chose amongst prospects with a monetary or coin prize. The answer to our

question, at least on the basis of this study, would, therefore, be that subjects did not behave differently when the incentive was something other than money. There is one important caveat to this conclusion, however.

Two versions of the coin treatment were considered in the study. In one the grade of the coins that could be won was clear to subjects, and in the other it was not. The grading is important in determining the value of the coin. When the grade of coin is clear, someone who knows about coins will be able to give a relatively accurate value of the coin. When the grade is not clear, the valuation will be more uncertain. Risk attitudes were the same whether the prize was money or a coin of known grade. But, when the grade was not clear, subjects were more risk-averse. This suggests that people may be more risk-averse if there is background uncertainty about the value of possible outcomes. This also suggests that estimates of risk aversion based on monetary incentives may understate risk aversion for the many contexts in which people do face background uncertainty. Note that we could know this only by using non-monetary prizes. This is because there cannot be background uncertainty in a monetary prize; for instance, a subject in the United States has a pretty good idea what a $10 bill is worth.

To explain, suppose that Alan has a wealth of $100 and prefers prospect I to prospect K. This means that

$u(100) > 0.5\ u(0) + 0.5\ u(225)$.

If $u(0) = 0$ and $u(100) = 100$, it must be that $u(225) < 200$. Knowing that Alan prefers prospect I to prospect K thus tells us something about the utility he gets from having wealth of $225. Now, suppose that Alan has wealth of $225 and still prefers prospect I to prospect K. We can then infer something about the utility he gets from having wealth of $350; see below. Continuing in this way, we get a picture of what the utility function must be like if Alan continues to prefer prospect I to K for higher and higher levels of wealth. Figure 3.2 gives such a picture.

[Extra] I should explain how to get this picture. Keep in mind that our objective is to try and stop the utility function from bending too much, because this leads to too much risk aversion. So, let us say that $u(225) = 199$, meaning that Alan only just prefers prospect I to K when his wealth is $100. If Alan has wealth of $225 and would still prefer prospect I to prospect K,

$u(225) > 0.5\ u(125) + 0.5\ u(350)$.

We do not know $u(125)$. But we do know that $u(100) = 100$ and $u(225) = 199$. If we put a straight line between $u(100)$ and $u(225)$ then we infer that $u(125) > 119$. This means $u(350) \leq 279$. Suppose that $u(350) = 279$. Now imagine that Alan has wealth of $350 and would still prefer prospect I to prospect K. This means that

$u(350) > 0.5\ u(250) + 0.5\ u(475)$.

Again, we do not know $u(250)$, but we can infer that $u(250) > 215$, so $u(475) \leq 343$. We can go on and on doing this.

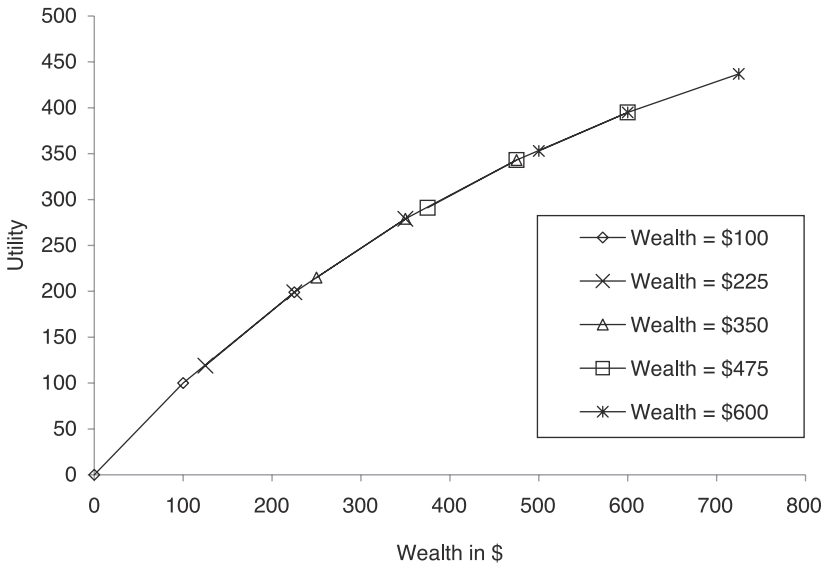

Figure 3.2 What can we infer about the utility function of someone who prefers prospect I to prospect K for increasing levels of wealth?

Table 3.7 The amount a person would need to gain on a 50:50 bet to take the bet, if risk-averse for small gambles.

Potential loss	**Potential gain needed to take the bet**	
	If prefers prospect I to J	**If prefers prospect I to K**
$400	$420	$1,250
$600	$730	$36,000,000,000
$800	$1,050	$90,000,000,000
$1,000	$1,570	$160,000,000,000
$20,000	$71,799,110	$540,000,000,000,000,000,000

Source: Rabin (2000).

The key point to notice in Figure 3.2 is how the utility function has already curved a lot. Going from $0 to $100 is worth 100 utility, but going from $600 to $700 is worth only 40 utility. We need the utility function to bend like this for low levels of wealth in order to understand why Alan would prefer prospect I. The problem is that, once we have started bending it, we have to carry on bending it. This relative bending of the utility function has implications for bigger gambles. To illustrate, if $u(225) = 199$, and if you carry on the reasoning above, you should find that $u(600) < 400$. Consequently, if Alan had a wealth level of $225, he would prefer prospect I to prospect M. This may well be the case, but we see that the difference between the relative gain and loss has increased quite a lot. If we were to look at realistic wealth levels of $50,000 to $300,000, then this effect would become very pronounced.

To demonstrate, Table 3.7 details the amount of money Alan would need as potential gain in order to take on a 50:50 gamble, if he had a wealth level of $290,000

and would prefer prospect I to J or K, for any level of wealth below $300,000. By this stage we see that the differences between relative gain and loss have become huge – implausibly huge. For example, it must be that anyone who would prefer prospect I to prospect K must prefer prospect I to prospect N. This is possible, but seems highly unlikely, suggesting that expected utility cannot explain risk aversion for small gambles (see Research methods 3.3).

❖ RESEARCH METHODS 3.3 RISK AVERSION AND INCENTIVES IN THE LAB

In many economic experiments we are interested in something other than attitudes to risk. The outcome may, however, be inherently risky. For example, in Chapters 6 and 7 we shall look at behavior in strategic contexts in which one person's payoff will depend on what others do. Our focus will be on things such as a person's ability to predict what another will do, or his desire to be fair. His payoff will depend on what another subject does, however, and so there is an element of risk. This complicates things. For example, we shall have to try and determine whether he gave something to another to be fair or because he is risk-averse. What we would like is that the subject's choice does not depend on his attitude to risk.

If subjects are risk-neutral then the problem goes away. If, however, subjects are risk-averse (or risk-loving) then things are not so simple. One potential solution is to frame choices in a slightly different way. Instead of asking subjects to make choices that might increase their payoff from, say, $5 to $10 we can offer them choices that increase the probability of winning, say, $100 from 0.05 to 0.10. If a subject maximizes expected utility then this approach should make attitudes to risk irrelevant. If, however, the subject maximizes expected utility then risk aversion should not matter anyway for the small sums of money in a typical experiment! It is not clear, therefore, how much such alternative framings really help.

Indeed, there seems to be no simple way to disentangle the consequences of risk and uncertainty from other things we might be interested in. This means that we need to be mindful of attitudes to risk in interpreting any experimental results.

3.1.3 Risk-loving for losses

The next set of prospects are in Table 3.8. Most people choose prospect P rather than O, and prospect Q rather than R. This suggests that people are risk-averse for gains, preferring the $3,000 for sure, but risk-loving for gambles, preferring to gamble on losing $0.

Such choices are theoretically possible while maximizing expected utility, but only by pushing things to the limit. Recall that utility should be measured according to the

Table 3.8 Prospects with gains and losses. Many prefer P to O and Q to R.

Prospect	Amount with probability of outcome	
O	$4,000 with probability 0.8	$0 otherwise
P	$3,000 for sure	
Q	$4,000 loss with probability 0.8	$0 otherwise
R	$3,000 loss for sure	

Table 3.9 Framing of losses versus gains. Many prefer T over S and U over V.

Prospect	In a preliminary stage given	Subsequently get amount with probability of outcome	
S	$1,000	$1,000 gain with probability 0.5	$0 otherwise
T	$1,000	$500 for sure	
U	$2,000	$1,000 loss with probability 0.5	$0 otherwise
V	$2,000	$500 loss for sure	

absolute level of wealth. So, we would need that, when asked to choose amongst these options, a person's current wealth just happened to be at a point where the utility function is concave for higher levels of wealth, and convex for lower levels of wealth. This is possible, but seems a bit too much of a coincidence. More plausible, particularly given what we saw in the last chapter, is that choices are made, in part, by measuring relative deviations from current wealth. This is called the **reflection effect**.

To illustrate the possible consequences of the reflection effect, consider the next set of prospects, given in Table 3.9. In principle, prospect S is equivalent to U and prospect T equivalent to V. Despite this, most people choose prospect T rather than prospect S and prospect U rather than prospect V. This can be explained by the reflection effect. In prospect T the $500 is perceived as a gain, while in prospect V the $500 is perceived as a loss.

3.1.4 When expected utility will work

Expected utility provides a brilliantly simple way to model choice with risk. We have already seen, however, that it can be a bit too simple to capture all that we observe. No model will work all the time, and so this is no great surprise, and it does not mean that we should ditch expected utility. If you want a model that can reliably capture how a person behaves for particular types of gamble, then expected utility may be good enough. If, though, you want a single model that can capture all the risky choices a person may make, whether it be big or small gambles, gambles with losses or with gains, then expected utility is probably not good enough. It is natural, therefore, to look for a model that can do a better job. A useful first step in doing this is to question when expected utility does provide a good model of choice. This can, hopefully, help pinpoint what expected utility misses.

The first thing we need for expected utility to work well is transitivity. A person's **preferences are transitive** if, for any three prospects X, Y and Z,

if $X \geq Y$ and $Y \geq Z$ then $X \geq Z$,

where the notation $X \geq Y$ means that the person prefers prospect X to prospect Y or is indifferent between prospects X and Y. In other words, there should be a natural ordering to preferences, with a best, X in this case, and worst, Z.

The second thing we need is independence. A person's **preferences satisfy independence** if, for any three prospects X, Y and Z,

if $X \geq Y$ then $(p, X; 1 - p, Z) \geq (p, Y; 1 - p, Z)$

for any number p between zero and one. In other words, if Alan prefers prospect X to Y, then he must also prefer a prospect that mixes X with some other prospect Z to one that mixes Y with the same prospect Z.

If preferences satisfy independence and transitivity (as well as two other properties of completeness and continuity) then preferences and choice can be modeled using expected utility. Expected utility would then be the best way to model choice. This is useful to know in trying to work out why expected utility may not work so well. In the following sections I shall take up in turn the possibility that preferences do not satisfy independence and that they do not satisfy transitivity. With this we shall see how to adapt expected utility to better account for choice with risk.

3.2 Independence and fanning out

I want to start this section by revisiting the Allais paradox. To help us understand the paradox, we can use a probability triangle diagram such as that in Figure 3.3. On the horizontal axis we measure the probability of having \$0, and on the vertical axis we

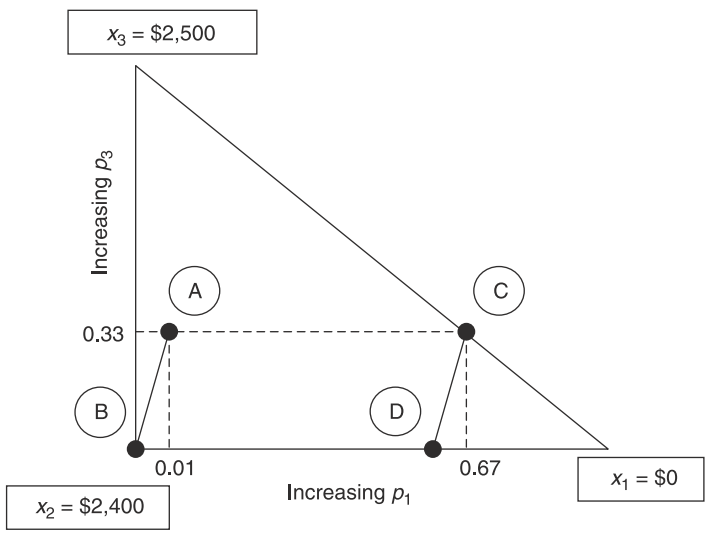

Figure 3.3 A probability triangle diagram of the Allais paradox. The horizontal axis is the probability of winning \$0 and the vertical axis the probability of winning \$2,500. The other possibility is to win \$2,400.

measure the probability of having \$2,500. Any point in the triangle is a prospect. This becomes clear by plotting in the triangle the four prospects from the Allais paradox. Prospect B guarantees \$2,400 for sure, so goes at the bottom left corner, where $p_1 = p_3 = 0$. Prospect A gives \$0 with probability 0.01 and \$2,500 with probability 0.33, so goes where $p_1 = 0.01$ and $p_3 = 0.33$. We can also plot on prospects C and D. Drawing this diagram makes clear that prospects C and D are similar to A and B but 'moved across' by adding a 0.66 chance of getting \$0.

What we now want to start to do is represent preferences on the probability triangle diagram. While we are doing this I want to clarify two different notions of utility. First, there is **outcome utility**. This is the utility we talked about in the previous chapter and can capture with the function u. It tells us the utility that Alan will get from a particular outcome, such as winning \$0, \$2,400 or \$2,500. Next is the **utility of a prospect**. You can think of this as the utility that Alan gets from a prospect before he knows the outcome of that prospect. We capture this with the function U, and to understand the Allais paradox we need to know the utility of prospects A to D. Expected utility is one way of getting at this (see equation (3.1)), but we now want to consider alternatives.

If we know the utility that Alan would get from any possible prospect then we can start to draw indifference curves in the probability triangle diagram. An **indifference curve** should show prospects that give the same utility. So, all we need to do is fix a level of utility, and find the prospects – that is, points in the probability triangle diagram – that give this utility. Figure 3.4 illustrates what indifference curves may look like. Indifference curves will likely slope up because a higher probability of the worst outcome, \$0, needs to be offset by a higher probability of the best outcome, \$2,500. This also means that the steepness of the indifference curves measures attitude to risk. The more risk-averse someone is, the more the probability of getting \$2,500 has

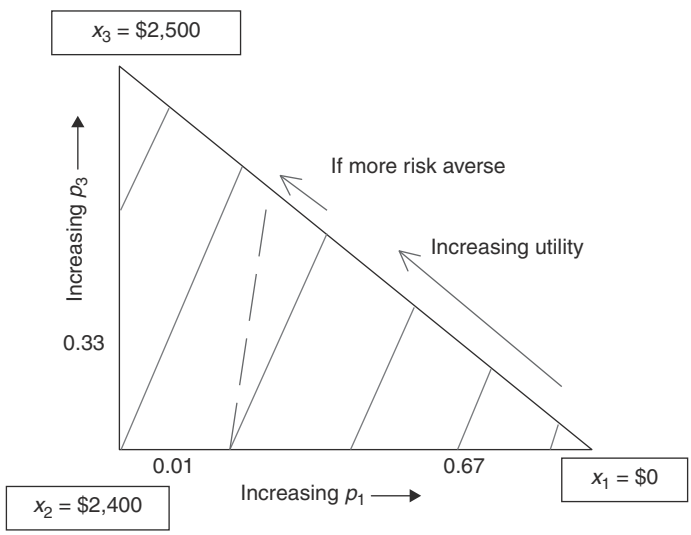

Figure 3.4 Indifference curves in a probability triangle diagram. The steeper the indifference curve, the greater the risk aversion.

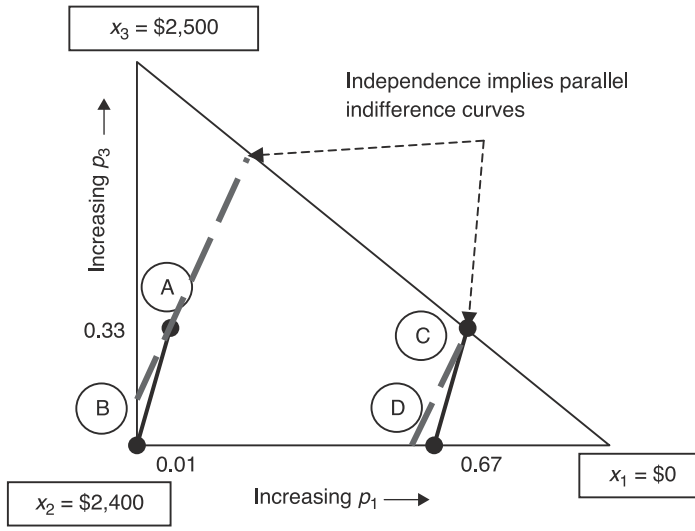

Figure 3.5 The Allais paradox. If preferences satisfy independence, the person should prefer either A and C or B and D.

to increase to compensate for any increase in the probability of getting $0. Thus, the steeper the indifference curve, the greater the risk aversion.

It turns out that, if preferences satisfy independence, the indifference curves must be straight parallel lines like those in Figure 3.4. Figure 3.5 illustrates the implications of this by combining parallel indifference curves with the four prospects. If the lines between prospects A and B, and C and D, are parallel, and the indifference curves are also parallel, then Alan should choose either A and C or B and D. As I have drawn it he would choose A and C. Recall, however, that most people chose B and C. This is the paradox.

From this we learn two things. First, the Allais paradox is probably caused by preferences not satisfying independence. Second, to explain the Allais paradox we will need indifference curves that are not parallel but instead **fan out** like those in Figure 3.6. The key thing about fanning out is that indifference curves get steeper the closer to the vertical axis, meaning that people are more risk-averse the closer they get to being guaranteed $2,400 or more. This is consistent with the certainty effect and can explain why someone prefers B over A, and C over D.

There are lots of ways that have been suggested for how to model the fanning out of indifference curves. Basically, if we want indifference curves to fan out then we can write down an equation for the utility of prospects that will give us what we want. More constructive is to develop models based on some intuition for why indifference curves might fan out. I shall look at two such models, one with disappointment and one with rank-dependent expected utility. In section 3.3 we shall look at other, more radical models that also cause fanning out.

3.2.1 Disappointment

One explanation for why Alan might dislike prospect A in the Allais paradox is that, if he chooses A and gets a payoff of zero, he will likely be very disappointed (you might

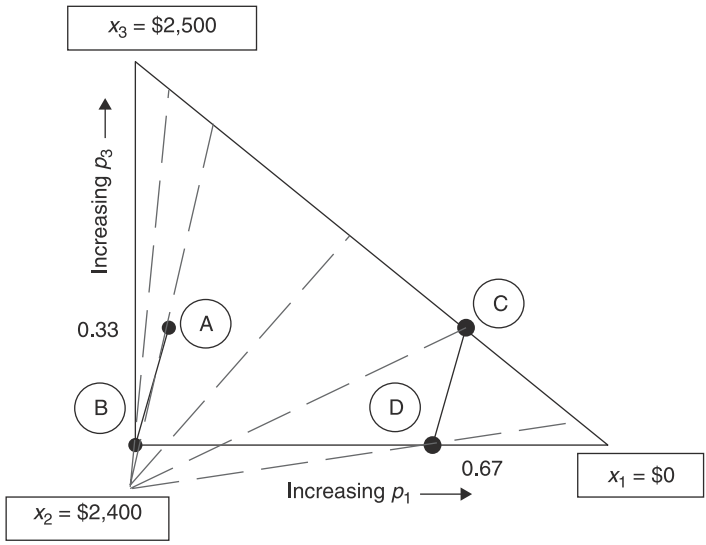

Figure 3.6 The Allais paradox can be explained by the fanning out of indifference curves. When indifference curves get steeper, the closer the person is to being guaranteed $2,400 or more.

want to substitute some stronger word here). Such disappointment can be avoided by choosing prospect B and getting a guaranteed high payoff. A way to capture such disappointment and its consequences was proposed by Loomes and Sugden (1986).

The basic idea is to measure outcome utility relative to some prior expectation of what utility will be. If outcome utility is below the expected level then Alan experiences **disappointment**. If the outcome utility is above the expected level then he experiences **elation**. Elation and disappointment are captured by some function, D. So, if the outcome is x his utility will be $u(x) + D(u(x) - \text{prior})$, the outcome utility plus or minus elation or disappointment. The utility of a prospect X is then written

$$U(X) = \sum_{i=1}^{n} p_i[u(x_i) + D(u(x_i) - prior)].$$

This is the same as expected utility, but taking into account disappointment and elation of being below or above the prior expectation. The prior expectation can be context-specific but a natural candidate is the expected utility of a prospect, calculated using equation (3.1). It is also natural to impose assumptions on D, including that $D(g) \geq 0$ if $g > 0$ and $D(l) \leq 0$ if $l < 0$, so elation is good and disappointment bad.

To demonstrate the implications of disappointment, we can work through the Allais paradox. Suppose that $u(x) = x$ and

$$D(g) = \begin{cases} \theta g^2 & \text{for any } g > 0 \\ -\theta g^2 & \text{for any } g < 0, \end{cases} \tag{3.5}$$

Table 3.10 The Allais paradox can be explained by disappointment. If $\theta = 0$, prospect A is preferred to B, and C to D. If $\theta = 0.0002$, prospect B is preferred to A, and C to D.

Prospect	Utility of the prospect		Value of D(.)			Net D(.)
	$\theta = 0$	$\theta = 0.0002$	If \$0	If \$2,400	If \$2,500	
A	2,409	2,398	−1,161	0	+2	−11
B	2,400	2,400	—	0	—	0
C	825	919	−136	—	+561	+94
D	816	899	−133	+502	—	+83

where $\theta \geq 0$ measures the importance of disappointment. Table 3.10 details the utility of each prospect for different values of θ. When $\theta = 0$ we have a standard expected utility model and Alan would choose prospect A over B, and C over D. If disappointment is given weight $\theta = 0.0002$, we see that prospect B becomes preferred to prospect A. Why? The expected utility of prospect A is \$2,409, and so if he gets \$0 his disappointment is $1{,}161 = 2{,}409^2\theta$, which is relatively large. For example, the elation he feels if he gets \$2,500 is only $1.66 = 91^2\theta$. These numbers must be weighted by the probability that they will happen, because he will be disappointed only with probability 0.01. Even so, we get a net effect of $-0.01 \times 2{,}409^2\theta - 0.66 \times 9^2\,\theta + 0.33 \times 91^2\,\theta = -11$. This is enough to just swing things in favor of prospect B.

The small probability of being very disappointed if he chooses prospect A can be enough for Alan to prefer prospect B. The Allais paradox can thus be caused by people avoiding potential disappointment. If we plot the indifference curves when $\theta = 0.0002$, we find that they do indeed fan out similarly to in Figure 3.6.

Disappointment provides a plausible explanation of the Allais paradox and fanning out, therefore. It can also point to other things that we might want to be aware of. For example, if we increase the weight given to disappointment, we can get indifference curves with weirder shapes. Figure 3.7 illustrates what happens when $\theta = 0.002$. There is still a preference for prospects B and C, but the indifference curves are no longer fanning out. In fact we obtain, counter-intuitively, that an increase in p_1, the probability of getting \$0, can increase the utility of a prospect! The explanation is that, when p_1 is relatively large, Alan is not so disappointed if he gets \$0, because he expected it. He might, therefore, prefer less chance of winning, because this means less disappointment if he does not win. Maybe not so counter-intuitive after all?

3.2.2 Rank-dependent expected utility

A model of disappointment deviates from expected utility by changing how utility is perceived. An alternative possibility is to change how probabilities are perceived. This is the approach taken by rank-dependent expected utility. I should probably warn you, before launching into the explanation, that the primary virtue of rank-dependent expected utility is that it works. It's not so intuitive why it works, so a little bit of faith is needed.

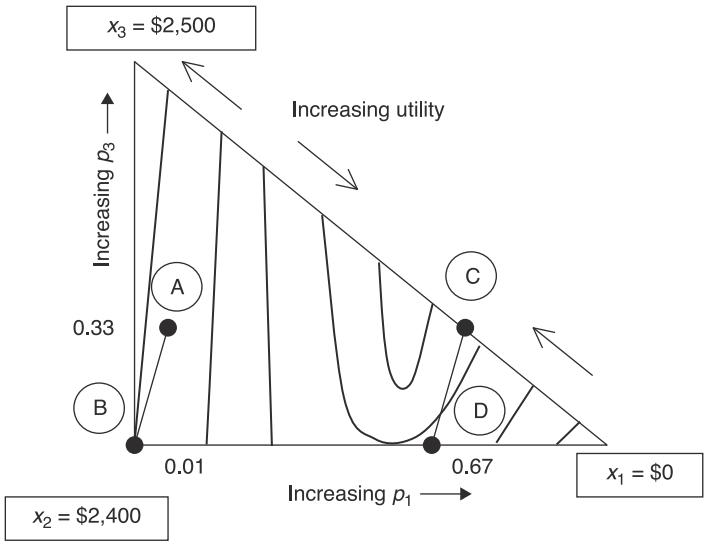

Figure 3.7 Indifference curves no longer fan out if disappointment is very strong.

One crucial component we need is a function π that **weights probabilities**. A commonly used function is

$$\pi(p) = \frac{p^{\gamma}}{(p^{\gamma} + (1-p)^{\gamma})^{\frac{1}{\gamma}}}, \tag{3.6}$$

where γ is some number between zero and one. The main reason for using this function is that it gives the nice S-shaped probability weights illustrated in Figure 3.8. You can see that if $\gamma = 1$ then $\pi(p) = p$, so there is essentially no weighting of probabilities. When γ is less than one there are two things that happen. First, there is an **overweighting of small probabilities**, in the sense that a small probability is given a relatively bigger weight. Second, there is an underweighting of large probabilities. We will come back later to why we would want these two things.

Now that we know how to weight probabilities, we can find the utility of a prospect. The first thing we need to do is rank the outcomes so that x_1 is the worst outcome and x_n is the best. For now we shall assume that $x_1 \geq 0$, and so the worst outcome is not a loss. The **rank-dependent expected utility** of a prospect X is then given by

$$U(X) = \sum_{i=1}^{n} w_i u(x_i),$$

where the decision weights are

$$w_i = \pi(p_i + \ldots + p_n) - \pi(p_{i+1} + \ldots + p_n)$$

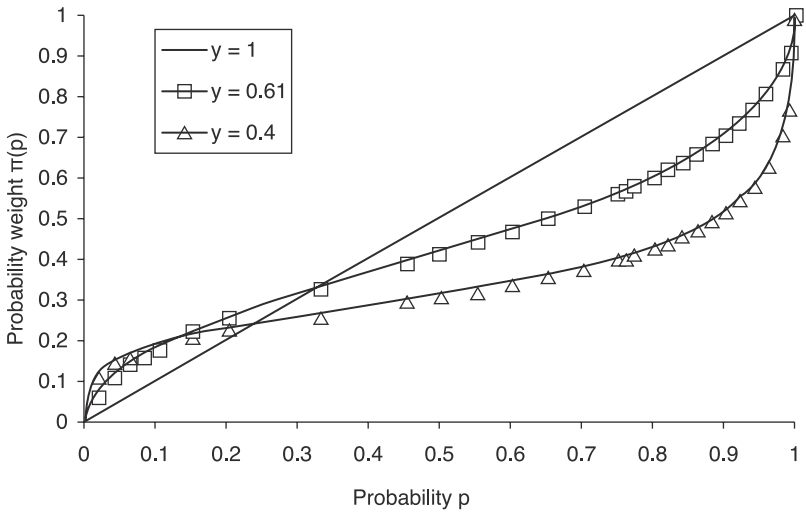

Figure 3.8 An illustration of probability weighting.

for all i. In interpretation, $\pi(p_i + \ldots + p_n)$ is the weighted probability of getting an outcome equal to or better than i, and $\pi(p_{i+1} + \ldots + p_n)$ is the weighted probability of an outcome better than i.

Rank-dependent expected utility is easy enough to work out if we know u and π. It's slightly harder to know why it works or why we would want to do it. I can try to provide some intuition, though. If $\pi(p) = p$ then $w_i = p_i$, and we have standard expected utility. More generally, decision weights allow us to capture pessimism and optimism. To illustrate, consider prospect X = (0.5, \$0; 0.5, \$100), in which Alan has a 50:50 chance of winning \$100. The decision weights will be $w_1 = 1 - \pi(0.5)$ and $w_2 = \pi(0.5)$. If $\pi(0.5) < 0.5$ then $w_1 > w_2$. This suggests that Alan is **pessimistic**, because he puts a higher decision weight on the worst outcome. If $\pi(0.5) > 0.5$ then $w_1 < w_2$. This suggests Alan is **optimistic**, because he puts a higher decision weight on the best outcome. Figure 3.8 illustrates that, with the weighting given by equation (3.6), we primarily observe pessimism with a little bit of optimism over small probabilities.

To illustrate rank-dependent expected utility, Table 3.11 works through the Allais paradox when $u(x) = x$. We can see that, for the most part, the weighting of probabilities has little effect. The important exception is prospect A, in which the 0.01 probability of the worst outcome is increased in decision weight to 0.088. This is pessimism and is enough to make prospect B be preferred to prospect A. The Allais paradox could thus be explained by pessimism and the overweighting of a small probability of getting \$0.

For completeness, we should have a look at the indifference curves we get in the probability triangle diagram. Figure 3.9 illustrates what they look like. They are not as exciting as with disappointment, but even here we see that the story is not as simple as fanning out. The indifference curves have an S shape.

My focus throughout this section has been on the Allais paradox, but, hopefully, you can see that there is enough intuition behind the models of disappointment and rank-dependent utility that we should be able to meaningfully apply them more

Table 3.11 Rank-dependent utility and the Allais paradox, when $u(x) = x$. If $\gamma = 1$, prospect A is preferred to B, and C to D. If $\gamma = 0.61$, prospect B is preferred to A, and C to D.

Prospect	If $\gamma = 1$				If $\gamma = 0.61$			
	$w1$	$w2$	$w3$	Utility	$w1$	$w2$	$w3$	Utility
A	0.01	0.66	0.33	2,409	0.088	0.577	0.334	2,221
B	0	1	0	2,400	0	1	0	2,400
C	0.67	0	0.33	825	0.666	0	0.334	836
D	0.66	0.34	0	816	0.661	0.339	0	815

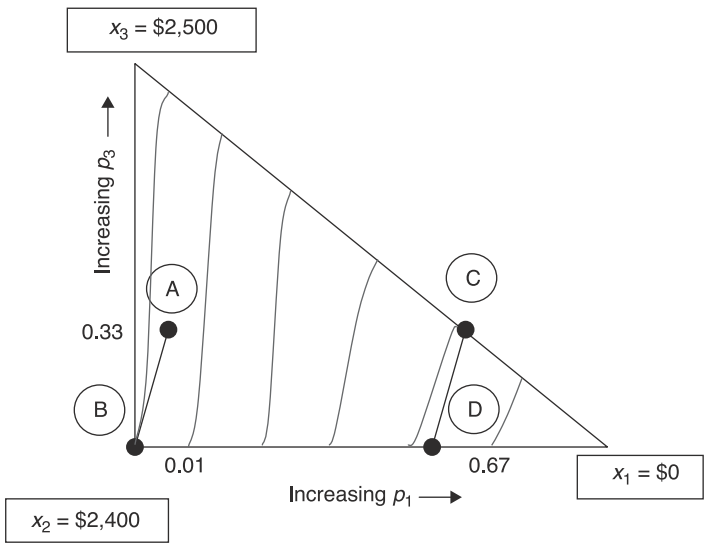

Figure 3.9 Indifference curves in the Allais paradox with rank-dependent utility.

generally. Indeed, given that both models capture something intuitively appealing, a next step is to combine them and try to account for things other than the Allais paradox. This is the direction we shall take in the next section.

3.3 Reference dependence and prospect theory

A theme through much of the previous chapter was the relevance of reference points and a potential asymmetry between losses and gains. I now want to show how we can extend that to choice with risk. Doing so brings us to **prospect theory** (see Research methods 3.4). Prospect theory can be as general or as specific as one wants it to be. Given that I have already talked about reference dependence a fair bit in the previous chapter, I am going to focus on a specific formulation of the theory here.

❖ RESEARCH METHODS 3.4 THE DEVELOPMENT OF AN IDEA

In 1979 Kahneman and Tversky published a paper entitled 'Prospect theory: An analysis of decision under risk'. This paper went on to become one of the most cited papers in economics, and prospect theory became arguably behavioral economics' most well-known idea.

The first half of the 1979 paper was a critique of expected utility. Kahneman and Tversky used hypothetical choice problems to demonstrate phenomena (such as the Allais paradox) that violated expected utility theory. The second part proposed a theory that could explain these phenomena and, therefore, provide a better account of decision-making with risk. The theory consisted of two parts: an editing phase, in which prospects would be organized and reformulated; and then an evaluation phase, when probabilities are weighted and the value of a prospect evaluated. The editing phase could account for things such as framing effects and the evaluation phase for things such as loss aversion.

The theory originally proposed by Kahneman and Tversky in 1979 was hard to apply. It was formulated only for prospects with two outcomes, and essentially set out what properties a model of decision-making with risk should have, rather than proposing a specific model. In 1992 a second paper was published that proposed a parameterized model of prospect theory that could more easily be applied. This is commonly referred to as second-generation prospect theory. The key step was to make use of rank-dependent weighting of probabilities.

The second-generation model was simpler and focused on the evaluation phase, with the editing phase almost completely forgotten. This lack of an editing phase left it unclear what the reference point should be. Recent models address this by allowing for a changing and probabilistic reference point. This is commonly referred to as third-generation prospect theory, and brings the attention back a little on the editing phase. Taken together, this work pointed out the shortcomings of standard economics while also, crucially, proposing something else that might work instead. That is progress, and prospect theory has rightly become a celebrated idea.

Unsurprisingly, one key ingredient of prospect theory is a reference point about which outcomes are judged. We then talk of gains or losses relative to the reference point. If x is final wealth and r is the reference level of wealth then $x \geq r$ is a gain and $x < r$ is a loss. Outcomes are judged relative to the reference point using utility function or value function:

$$v(x) = \begin{cases} (x-r)^{\alpha} & \text{if } x \geq r \\ -\lambda(r-x)^{\beta} & \text{if } x < r, \end{cases} \tag{3.7}$$

where α, β and λ are parameters. In interpretation, this value function is exactly the same as that which I talked about in the previous chapter. The only difference is that I have now given a specific functional form for what it might look like.

The thing we do need to do (which we did not do in the last chapter) is think about risk. To do this we shall use rank-dependent expected utility. One slight twist is that, when weighting probabilities, we should distinguish gains from losses. So, when we rank outcomes from worst x_1 to best x_n, we also need to distinguish outcomes that are losses – say, x_1 to x_L – and gains – x_{L+1} to x_n. Gains are then given decision weight,

$$w_i = \pi_g (p_i + \ldots + p_n) - \pi_g (p_{i+1} + \ldots + p_n),$$

as before, and losses are given decision weight,

$$w_i = \pi_l (p_1 + \ldots + p_i) - \pi_l (p_1 + \ldots + p_{i-1}),$$

where

$$\pi^g(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{\frac{1}{\gamma}}}, \quad \pi^l(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{\frac{1}{\delta}}}$$

are the probability weights for gains and losses. This formulation might look a little messy, but it is just a natural extension of what we did beforehand to capture the asymmetry whereby gains are good and losses are bad.

With equations (3.6) and (3.7), and this weighting of probabilities, we can work out easily enough the utility of any prospect, if we know $r, \alpha, \beta, \lambda, \gamma$ and δ. The overriding motivation for this formulation is that it is relatively simple but has the potential to explain what we observe. For instance, if $\lambda > 1$ then we capture loss aversion, because losses are given more weight than gains. If $1 > \alpha, \beta > 0$, then we capture risk aversion over gains and risk loving over losses. Finally, if $\gamma, \delta < 1$, then we can capture the certainty effect.

But, to see whether this formulation really does work, we need to go and estimate the parameters $\lambda, \alpha, \beta, \gamma$ and δ. Tversky and Kahneman (1992) did this by recruiting 25 graduate students from Berkeley and Stanford and having them choose amongst a series of prospects. Before giving the headline numbers I want to point out that these estimates are based on the median responses and we should not ignore the possibility that different subjects behave very differently. Now for the headline numbers: they were $\lambda = 2.25$, $\alpha = \beta = 0.88$, $\gamma = 0.61$ and $\delta = 0.69$. Thankfully, these numbers are consistent with what is needed to capture loss aversion and the like.

Figure 3.10 plots the utility function for parameters $\lambda = 2.25$ and $\alpha, \beta = 0.88$. The asymmetry between losses and gains is apparent, and comes from losses being given over twice the weight of gains. You might also notice, however, that there is not much curvature to the utility function. It is marginally concave for gains and convex for losses, but this is quite hard to see. This is because the estimates of $\alpha, \beta = 0.88$ are close enough to one for the utility function to be nearly linear above and below the x axis. Indeed, to simplify things, it is not uncommon to use $\alpha, \beta = 1$.

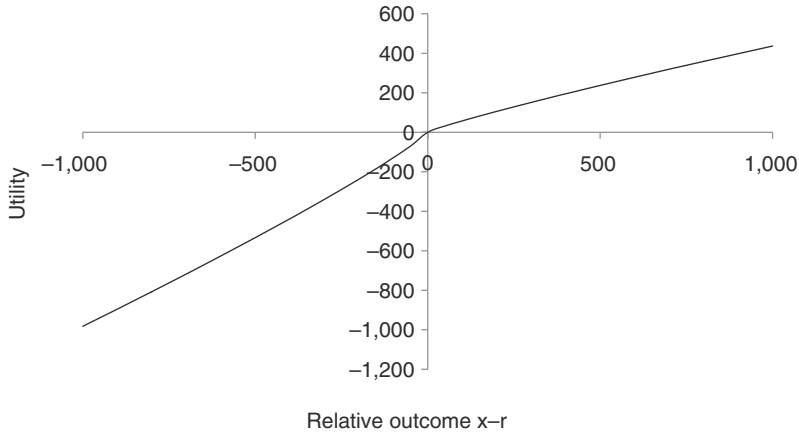

Figure 3.10 A prospect theory utility function for money with parameters $\lambda = 2.25$, $\alpha, \beta = 0.88$, $\gamma = 0.61$ and $\delta = 0.69$.

This lack of curvature might suggest that we are going to struggle to explain risk aversion over small gambles. The utility function in Figure 3.10, for example, is a lot flatter than that in Figure 3.2. Loss aversion can help us out, however. That's because the kink in the utility function at zero means that the utility function curves a lot when we compare negative and positive amounts. Loss aversion can, for example, explain why Alan would prefer prospect I in Table 3.6 to prospect J. This line of thinking does raise some questions (which I will come back to) about what the reference point should be and what should be coded as a gain and a loss. For now, though, let's focus on the good news: prospect theory has allowed us to capture everything we came across in sections 3.1.1 to 3.1.3.

3.3.1 Reference-dependent utility

To put prospect theory in context, we can write down a more general model that includes as special cases all the models we have seen so far. This extends the reference-dependent utility function that I introduced in the previous chapter. The **rank-dependent**, **reference-dependent expected utility** of a prospect X is

$$U(X) = \sum_{i=1}^{n} w_i[\eta u(x_i) + h(x_i, r)], \tag{3.8}$$

where h is some function that depends on a reference point r and η is some parameter. The three special cases we have looked at can be explained as follows.

If $\eta = 1$, $w_i = p_i$ and h is the same as the disappointment function D in equation (3.5) then we have a model of disappointment. In this case, there is no weighting of probabilities, but there is disappointment or elation if the utility of the outcome is above or below the prior expected level of utility.

Table 3.12 How disappointment theory, rank-dependent expected utility and prospect theory differ from expected utility.

Model	Differences from expected utility
Disappointment	There is disappointment or elation if the utility of the outcome is below or above the prior expectation
Rank-dependent utility	There is weighting of probabilities to take account of pessimism or optimism
Prospect theory	There is weighting of probabilities, and the utility of the outcome is determined solely by the gain or loss relative to the reference point

If $\eta = 1$ and $h(x, r) = 0$ for all x then we have rank-dependent expected utility. If $w_i = p_i$ then this reduces to expected utility. Another way to get at the same thing is to have $\eta = 0$, $r = 0$ and h the same as the utility function u. In this case, outcomes are not judged relative to a reference level, but there is weighting of probabilities.

Finally, if $\eta = 0$ and h is the same as the value function v in equation (3.7) then we have prospect theory. In this case, utility is determined by the loss or gain relative to the reference point, and there is weighting of probabilities.

The reference-dependent utility function (3.8) gives us, therefore, a framework to compare models. We see that they are all different, and all point to different ways that we might want to move beyond expected utility. This is summarized in Table 3.12. Clearly, when we combine the three models we also obtain a very general model to think about choice with risk. The model is probably a bit too general to be very practical, however. So, in practice, it is necessary to start thinking about narrowing things down a little by asking what features are most important, or most important in particular contexts.

That is an exercise I leave for you to think about, because what I want to do now is focus on two things that we looked at in some detail in the last chapter but are worth briefly revisiting here. First, equation (3.8) is going to depend crucially on reference point r, and so we need to think about that. Second, attitudes to risk will depend on how potential risks are bracketed or perceived relative to existing risk.

3.3.2 The reference point and expectations

We have seen so far, in this chapter, three candidates for the reference point: they are the current level of wealth (prospect theory), zero wealth (expected and rank-dependent expected utility) and the expected utility of the prospect (disappointment theory). In the last chapter we looked at other things that might influence the reference point, such as the 'no loss in buying' hypothesis, and said that these things primarily reflected expectations. For instance, if Anna expected to pay \$3 for a box of Nutty cereal, then paying \$3 for a box was not counted as a loss. In choice with risk, expectations of what may happen can become even more important.

To illustrate the implications of expectations, let's return to the insurance example and consider the choice between full insurance and insurance against theft. Recall that the insurance against theft prospect is AT = (0.05, \$14; 0.95, \$64), full

Table 3.13 How the utility of Alan will depend on his expectations. The expected utility from gains and losses measures the expected utility from gains and losses relative to the reference point.

Framing	Reference point	Expected utility from gains and losses		Reference-dependent utility	
		AT	FI	AT	FI
Unexpected choice, immediate outcome	\$70	–19.1	–20.3	58.8	57.9
Unexpected choice, delayed outcome, chooses AT	\$61.5	–3.0	—	74.9	—
Unexpected choice, delayed outcome, chooses FI	\$61	—	0.0	—	77.9
Expected choice of AT	\$61.5	–3.0	–1.1	74.9	77.0
Expected choice of FI	\$61	–2.4	0.0	75.4	77.9

insurance is prospect $FI = (1, \$61)$ and Alan's wealth is \$70. Suppose that Alan has the reference-dependent utility function

$$u^r(x) = \begin{cases} 10\sqrt{x} + (x - r) & \text{if } x \geq r \\ 10\sqrt{x} - 2.25(r - x) & \text{if } x < r, \end{cases} \tag{3.9}$$

where r is the reference point. I shall now contrast three scenarios worked through in Table 3.13.

First, suppose that it was a complete surprise to Alan that he was asked to make this choice, and once he has made the choice he will immediately find out the outcome. For example, the renewal letter arrives in the post when he is not expecting it, he needs to reply that day, and the next day he will find out whether his son is going to use the car (and his son has a history of being in accidents). This sounds a bit contrived but is the scenario that most closely fits what we have been analyzing so far. (It is also broadly representative of what happens in a typical lab experiment, in that a subject is thrust into a situation, makes a choice and immediately observes the outcome.) In this case the current level of wealth seems the most relevant reference point. Paying for insurance is thus seen as a loss, and Alan would prefer insurance against theft.

Next, suppose, more realistically, that Alan was still surprised to get the renewal letter but will have to pay for the insurance only after some delay and will not find out the outcome – i.e. whether the car is in an accident – for some time. In this case, it makes sense for the reference level to be the expected outcome of the prospect that Alan chooses, because he will have time to change his expectations of wealth before anything happens. His reference point will thus depend on what he chooses. In this case Alan will prefer full insurance, because by the time he pays for it he will not think of it as a loss. With insurance against theft, by contrast, he could still lose if the car is in an accident.

Finally, suppose that Alan was expecting the renewal letter. He presumably had some expectation of what he would choose and has already factored this into

his reference point. The reference point will, therefore, depend on whether he expected to choose insurance against theft or full insurance. In both cases he prefers full insurance. We see, however, that full insurance looks relatively better if he expected to buy full insurance, and insurance against theft looks relatively better if he expected to buy insurance against theft. This can lead to self-fulfilling expectations, whereby the thing that Alan expected to choose looks better because he expected to choose it.

Two general points are apparent in this example. First, whether or not Alan expected to face risk will influence how he deals with that risk. For instance, he is not willing to pay for full insurance when the renewal letter was a surprise, but is when the renewal letter was expected. Second, if risk was expected, the optimal choice can depend on what Alan expected to choose. Kőszegi and Rabin (2007) capture this with the idea of a **personal equilibrium**. The expected choice of full insurance is a personal equilibrium because it makes sense for Alan to choose full insurance if he expected to choose it. The expected choice of insurance against theft is not a personal equilibrium, because Alan would want to choose full insurance even if he expected to choose insurance against theft.

The picture we get is one in which expectations and context can matter a lot in determining the reference point, and hence the perception of gains and losses. This naturally builds on the things we talked about in the previous chapter, but the presence of risk adds even more things to the mix. There is even more ambiguity about how things will be perceived and what the reference point may be. This clearly raises some challenges in trying to model choice with risk. One small part of meeting this challenge requires asking how people might perceive interrelated gambles.

❖ RESEARCH METHODS 3.5 TV GAME SHOWS

The objective in most economic experiments is complete anonymity. The exact opposite is surely a TV game show, when contestants make choices in front of millions of viewers. Game shows provide an interesting alternative to an economic experiment, therefore, particularly as the stakes are typically a lot higher.

Post and co-authors (2008) are one group of several researchers to study the show *Deal or No Deal*. In the game, a contestant has to choose one of 26 briefcases. Each briefcase contains an amount of money from €0.01 to €5 million (in the Dutch edition). Having chosen a briefcase, the contestant is sequentially told what is in the other briefcases, and given a walk-away offer. If the contestant continues to turn down the offers they eventually get to see how much money is in the briefcase they first chose, and they can keep it! The walk-away offer depends on the expected value of the remaining suitcases.

The issue that Post and co-authors focus on is whether choice is path-dependent. Specifically, did the contestants' risk attitudes change as they began to see what was in the briefcases? They find that they did. Contestants

became less risk-averse when either high amounts (which is unlucky) or low amounts (which is lucky) were in the briefcases they did not choose. For example, one contestant called Frank was unlucky but carried on gambling, eventually rejecting a sure €6,000 in favor of a 50:50 gamble of €10 or €10,000. He won €10. A second contestant, Susanne, was lucky and rejected a sure €125,000 to eventually win €150,000. It is doubtful we would see such risk taking without the prior process of sequentially seeing what was in the other briefcases.

Post and his co-authors show that such behavior can be explained by prospect theory with a reference point based largely on prior expectations of the likely prize.

3.3.3 Combined gambles

Most people are constantly exposed to existing risks because they have money invested in the stock market, might lose their job, have already bought a lottery ticket, etc. This raises interesting issues about how a person should evaluate a new risk.

To get us started I will demonstrate that it may be in Alan's interest to take on a small gamble, as a form of diversification, if he is exposed to existing risk. Consider the prospects in Table 3.14. Suppose that Alan is already committed to prospect X, because, say, he has money invested in the stock market. He now needs to decide whether to take on prospect Y or Z. As you can see, the outcome of prospects may be revealed now or in one week's time. Suppose that his current wealth is $20,000 and his utility function is

$$u^r(x) = \begin{cases} x - 20{,}000 & \text{if } x \geq 20{,}000 \\ -2(20{,}000 - x) & \text{if } x < 20{,}000. \end{cases}$$

With these preferences Alan would not prefer prospect Y or Z to prospect W. More relevant, however, is to compare prospect Z with prospect X, because both will be

Table 3.14 Risk aversion for small gambles is less if there is existing risk. Alan is already exposed to prospect X, and so by choosing prospect Z would obtain prospect X + Z, which is preferred to prospect X.

Prospect	Probability and value	Timing	Expected utility
W	(1, $0)	Now	0
X	(0.5, –$10,000; 0.5, $30,000)	One week	5,000
Y	(0.5, –$500; 0.5, $550)	Now	–225
Z	(0.5, –$500; 0.5, $550)	One week	–225
X + Z	(0.25, –$10,500; 0.25, –$9,450; 0.25, $29,500; 0.25, $30,550)	One week	5,037.5

resolved in a week's time. If he chooses prospect Z then this should be combined with the existing prospect X to give a combined prospect X + Z. It turns out that prospect X + Z is preferred to prospect X.

The key point of the example is that Alan should choose prospect Z if he is already exposed to prospect X. That's because any new risk should be considered in the light of existing risk. This holds more generally than this simple example. People should be less risk-averse over prospects for small amounts of money if they are already exposed to risk.

Another thing the example illustrates is how the delay before knowing the outcome of a prospect should influence choice, if there are existing risks. Alan would do well to choose prospect W rather than Y, but also choose prospect Z. Or should he? If there is no escape from being told the outcome of prospect Y immediately then, yes, he should choose prospect W rather than Y. Suppose, however, he could somehow arrange that he will not be told the outcome of prospect Y until one week's time. Then he should do this. This brings to the fore the question of how often a person should choose to observe the outcome of prospects. If Alan is loss-averse it may be beneficial for him to delay or combine finding the outcomes of existing prospects so as to avoid too much 'bad news'.

These issues of how to combine prospects are the same as those we looked at in Chapter 2 in terms of mental accounting, hedonic editing, narrow framing, and the like.

A pertinent question is whether people do take existing risks, and delay, into account. Narrow framing would suggest they do not, but we also know there is a question of how narrowly people like to bracket things together. I will come back to some evidence on this later when looking at investing in the stock market. In the meantime here are two things for you to think about. Let's start with the puzzle of low stock market participation. Historically, though things have changed now, relatively few people invested in the stock market. If people are risk-averse this may seem reasonable. If we take into account all the other risks that people were exposed to, however, such as becoming unemployed, being in an accident, etc., then it becomes much harder to explain such low participation. This suggests narrow framing. On the other hand, it is not unknown for people to delay opening a letter that may bring bad news until they can also open a letter that will, it is hoped, have good news!

3.3.4 Stochastic reference point

One important issue that I have largely dodged until now is what should happen if the reference point is itself a gamble. Consider again, by way of illustration, the case in which Alan typically purchases insurance against theft. This means that his reference point is the prospect (0.05, $14; 0.95, $64). In section 3.3.2 and Table 3.13 I simply assumed that Alan acts as though his reference point is the expected value of this prospect – i.e. $61.5. But it does not have to be like that. In particular, Alan might behave differently with the reference point of insurance against theft from how he would do if his reference point was a sure $61.5. If, as we have just discussed, everyone faces risk as part of everyday life then it is natural to think that our reference point should be viewed as a gamble. This is an issue, therefore, that we have to take seriously.

An approach advocated by Kőszegi and Rabin (2006, 2007) is that of **stochastic reference points**. With this approach, if Alan's reference point is insurance against theft then, with probability 0.95, he will think in terms of a reference point of \$64 and, with probability 0.05, he will think in terms of a reference point of \$14. To illustrate how this works, let us look at the expected utility of Alan from choosing full insurance if his reference point is insurance against theft. I will use the same reference-dependent utility function as previously; see equation (3.9). If Alan chooses full insurance then he is guaranteed to have final wealth \$61, and so his expected utility is

$$10\sqrt{61} - 0.95 \times 2.25(64 - 61) + 0.05(61 - 14) = 74.04.$$

This calculation recognizes that, with probability 0.95, the reference point is \$64 and so he experiences a loss of \$3, while, with probability 0.05, the reference point is \$14 and so he experiences a gain of \$47.

Note that when working out the expected utility of insurance against theft we also have to take into account that the reference point is stochastic. Specifically, if Alan takes out insurance against theft then, with probability 0.95, he will end up with \$64. If his reference point is \$64 then he experiences no loss or gain. If his reference point is \$14 then he experiences a gain of \$50. Similarly, with probability 0.05 Alan will end up with \$14. In this case he experiences a loss of \$50 if his reference point is \$64 and no gain or loss if the reference point is \$14. Given that Alan's reference point is \$64 with probability 0.95 and \$14 with probability 0.05, his expected utility is therefore

$$0.95(10\sqrt{64} + 0.05(64 - 14)) + 0.05(10\sqrt{14} - 0.95 \times 2.25(64 - 14)) = 74.9.$$

You can see that Alan gets a higher expected utility from choosing insurance against theft, 74.9, than he does from choosing full insurance, 74.04. So, if his reference point is insurance against theft then Alan should choose insurance against theft. The more observant of you might have noticed that this is different from the result we obtained when we fixed the reference point as the expected value of insurance against theft! In that case, see Table 3.13, we found that full insurance gave a higher expected utility, 77.0, than insurance against theft, 74.9.

This example illustrates the need to think carefully about how people behave when the reference point is a gamble. If Alan thinks in terms of the expected value of insurance against theft then he chooses full insurance. But if he thinks in terms of a stochastic reference point then he chooses insurance against theft. Clearly, therefore, we need some insight into how Alan is likely to think. Before we do that it is worth briefly questioning why the stochastic reference point makes a difference.

If Alan takes as his reference point the expected value of insurance against theft then he ends up being risk-averse. This is because full insurance guarantees that he will not experience the loss of \$47.5 that would occur if his car was stolen, and he ended up with \$14 rather than the reference level of \$61.5. By contrast, if Alan uses a stochastic reference point he ends up being risk-neutral. This is primarily because gains and losses tend to cancel out. For instance, if he takes out full insurance and his

Table 3.15 A set of choices for which the reference point is a certain amount.

Choice	Sure payoff	Risky prospect
1	$10 for sure	$30 with probability 0.3 and $0 otherwise
2	$10 for sure	$30 with probability 0.35 and $0 otherwise
3	$10 for sure	$30 with probability 0.4 and $0 otherwise
4	$10 for sure	$30 with probability 0.45 and $0 otherwise
5	$10 for sure	$30 with probability 0.5 and $0 otherwise

Table 3.16 A set of choices for which the reference point is a prospect.

Choice	Sure payoff	Risky prospect
1	$10.50 for sure	$30 with probability 0.5 and $0 otherwise
2	$13.50 for sure	$30 with probability 0.5 and $0 otherwise
3	$16.50 for sure	$30 with probability 0.5 and $0 otherwise
4	$19.50 for sure	$30 with probability 0.5 and $0 otherwise
5	$21.50 for sure	$30 with probability 0.5 and $0 otherwise

reference level is $64 then he experiences a loss of $3. Given that his reference level will likely be $64, this loss weighs heavily on his decision and counteracts the potential loss he would experience if his car is stolen.

The general point to take from this example is that a stochastic reference point tends to make people appear more risk-neutral. It is worth pointing out that this is related to the idea of the previous section that a person should take on risk, and so appear more risk-neutral, if they already have existing risk. With this in mind we can now turn to some experimental evidence. I will look at a study by Sprenger (2015).

Consider, first, the set of choices in Table 3.15. With choice one you can either have $10 for sure or a gamble with expected value $9. A risk-averse person, looking at this gamble in isolation, would choose the sure $10. As you work down the list things turn in favor of the risky prospect. For instance, with choice five you can have $10 or a gamble with expected value $15. This increase in expected value might be enough to tempt a risk-averse person to gamble. Indeed, it can be expected that most people would choose the sure payoff in choice one and then switch to the risky prospect by choice five. The point at which they switch gives a measure of how risk-averse or risk-loving they are.

Consider, next, the set of choices in Table 3.16. With choice one you can either have $10.50 for sure or a gamble with expected value $15. A risk-averse person might be tempted by the gamble. As you work down the list things turn in favor of the sure payoff. For instance, by choice five only the most risk-loving would turn down the sure $21.50. So, we would expect that most people would choose the risky prospect in choice one and then switch to the sure payoff by choice five. Again, the point at which they switch gives a measure of how risk-averse or risk-loving they are.

Our main focus is on the reference point. If Alan is presented with the set of choices in Table 3.15 then it seems natural his reference point is a sure $10.

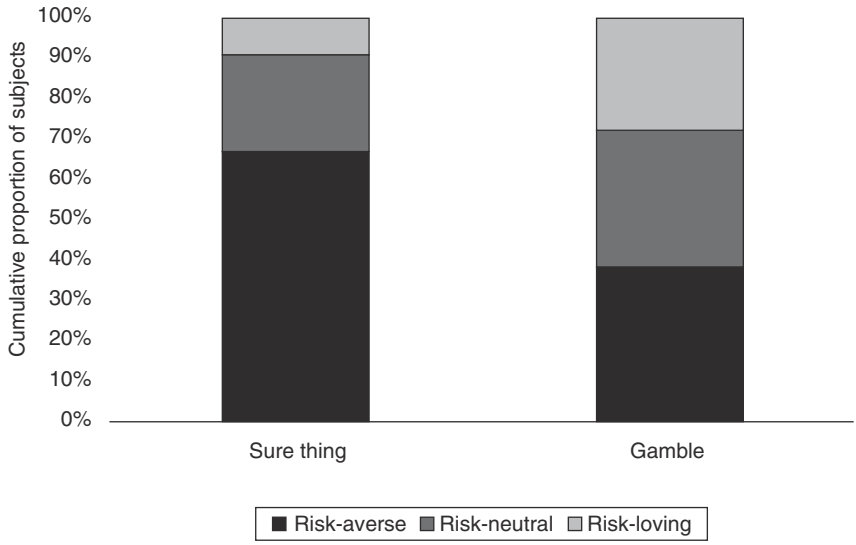

Figure 3.11 The proportion of subjects appearing risk-averse, risk-neutral or risk-loving if the reference point was a sure thing or gamble.
Source: Sprenger (2015).

If so, then we do not need to worry about stochastic reference points and we would expect him to appear risk-averse. Contrast that with the set of choices in Table 3.16. Crucially, here the risky prospect looks like the natural reference point, because that is held fixed across all choices. This allows us to distinguish between competing theories of how people behave when the reference point is a prospect. Specifically, if Alan thinks in terms of expected value then he should again appear risk-averse. But, if he thinks in terms of a stochastic reference point, he should appear risk-neutral.

In the study by Sprenger, some subjects were given choices like those in Table 3.15 and some given choices like those in Table 3.16. Figure 3.11 summarizes the main findings. You can see that when the reference point is a sure thing most subjects appeared, as we would expect, risk-averse. But when the reference point was a gamble there was no tendency towards risk aversion. These results would suggest that people do think in terms of a stochastic reference point. They also reaffirm the point that risk aversion seems to be driven to a large extent by loss aversion. In measuring and interpreting risk aversion it is, therefore, important to ask what the relevant reference point is and how this may influence choice. That is a point we will come back to in the next chapter when looking at time preference.

3.4 Preference reversals

Recall that in section 3.1.4 I said that expected utility should be appropriate if preferences satisfy independence and transitivity. Up until now our focus has primarily been on independence. For instance, preferences consistent with reference-dependent

expected utility need not satisfy independence but must satisfy transitivity. What I want to do now is change focus a little by questioning whether preferences satisfy transitivity. In order to do this I shall revisit the disparity between willingness to accept and willingness to pay, which we looked at in the previous chapter.

❖ RESEARCH METHODS 3.6 BEHAVIORAL ECONOMICS' LONGEST SAGA

Preference reversals have created a lot of debate over the years. Things got started with work by Lichtenstein and Slovic published in the late 1960s and early 1970s. Most widely cited is a paper they published in 1971, entitled 'Reversals of preferences between bids and choices in gambling decisions'. Since then debate has continued unabated about whether preference reversals are real, what causes them, whether they disappear with experience, and so forth.

One interesting aspect of the debate has been the different methods used to think about the issue. Psychologists have largely focused on the issue of procedural invariance. Lots of different ways to model and explain procedural invariance have been suggested. Some economists have focused on whether preference reversals are real, questioning whether the experimental methods used are reliable enough. Others have pointed out that preference reversals could be obtained by a violation of the independence axiom, while others, as we shall see, say that they can be obtained by a violation of the transitivity axiom. A more recent literature has looked at what happens if the willingness to accept is elicited from markets, and subjects have chance to learn over time.

This mix of approaches looks fantastic, but can end up seeming a little antagonistic. At times it seems as though the debate is one in which some argue that preference reversals are not real and others respond by showing they are. On the plus side, this has given us a lot of evidence to suggest that preference reversals really exist. It can detract slightly from some of the more subtle underlying issues, however, which tend to get muddled together or forgotten.

Look at the prospects in Table 3.17. Prospects AA and AB have the same expected value but get there in different ways: the **P bet** offers a high probability of winning a low payoff and the **$ bet** a low probability of winning a high payoff. Suppose that we elicit the willingness to accept for each of the prospects. You can refer back to the last chapter for a definition of willingness to accept, but, basically, we shall ask Alan the minimum amount of money he would accept rather than have prospect AA, or prospect AB, and shall interpret this as the price Alan puts on the prospect. Subsequently we can ask Alan what he would rather choose: prospect AA or prospect AB.

Figure 3.12 summarizes the results of a study by Tversky, Slovic and Kahneman (1990) in which subjects were asked to do this. Many put a higher price on the $

Table 3.17 Prospects that can lead to a preference reversal. Many value the $ bet higher than the P bet but choose the P bet over the $ bet.

Prospect	Amount with probability of outcome
AA, the '$ bet'	$100 gain with probability 0.08 and $0 otherwise
AB, the 'P bet'	$10 with probability 0.8 and $0 otherwise
AC, sure thing	$8 with probability one

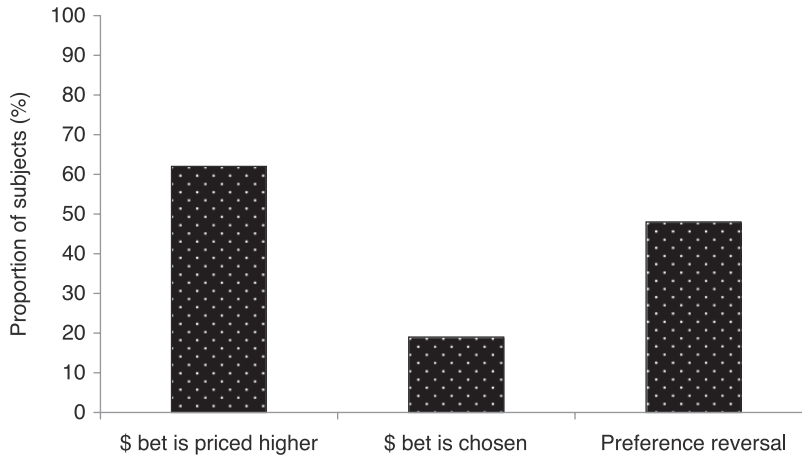

Figure 3.12 The proportion of subjects who price the $ bet higher, choose the $ bet and exhibit a preference reversal.
Source: Tversky, Slovic and Kahneman (1990).

bet than the P bet, but then chose the P bet over the $ bet. This is an example of a **preference reversal**. Another example is to value the P bet more than the $ bet yet choose the $ bet over the P bet. This is rarer.

Preference reversals are puzzling. To see why, suppose that Alan prices the $ bet at $9 and the P bet at $7 and also says he would choose the P bet over the $ bet. You could sell him the $ bet for $9, then offer to swap that for the P bet, and then buy the P bet off him for $7. We are back where we started, except you are $2 richer and he is $2 poorer. It looks like a good deal for you!

Because they are so puzzling, many have questioned whether preference reversals really exist. The evidence, however, suggests that they definitely do exist (see Research methods 3.6 and 3.7). The more important question, therefore, is why they exist. It might be because preferences are non-transitive. To see why, consider the prospect AC in Table 3.17. Given that Alan prices the $ bet at $9 and P bet at $7 it would appear that $AA \geq AC$ and $AC \geq AB$. But we know that $AB \geq AA$. These preferences appear non-transitive. We saw enough in the last chapter about framing effects and procedural invariance, however, to suggest that this might also be the cause. Maybe Alan perceives prospects differently depending on how they are presented to him. We need to think about both framing and non-transitivity.

❖ RESEARCH METHODS 3.7 DO INCENTIVES MATTER?

A perennial issue in experimental economics is how to incentivize subjects to make the choices that really capture their preferences. For example, when asking a subject to state a willingness to accept a $ bet, how can we be sure he tells us what he would really want?

In psychology it is common to ask subjects hypothetical questions, and so a subject is asked about his willingness to accept, knowing that his answer is ultimately not going to matter. Economists are generally skeptical about such an approach and prefer monetary incentives. So, having asked the subject his willingness to accept, a real game is played. For example, we could generate a random number between 0 and 10. If the random number is above his willingness to accept he gets to keep that amount of money (e.g. if his minimum willingness to accept is $8 and the number comes up 9, he gets $9). If the random number is below his willingness to accept, the $ bet is played out and he either wins $100 or gets $0. Knowing that this game will happen should get a subject to reveal his true willingness to accept.

A study by Berg, Dickhaut and Rietz (2010) looked back over the many studies on preference reversals and tried to pick out whether it matters what the incentives are for subjects. In some studies choices are hypothetical and in others monetary incentives are provided that should get subjects to reveal their true preferences. In an intermediate category subjects are given monetary incentives, but the procedure used may not get subjects to reveal their true preferences. Figure 3.13 summarizes some of their data. It turns out that

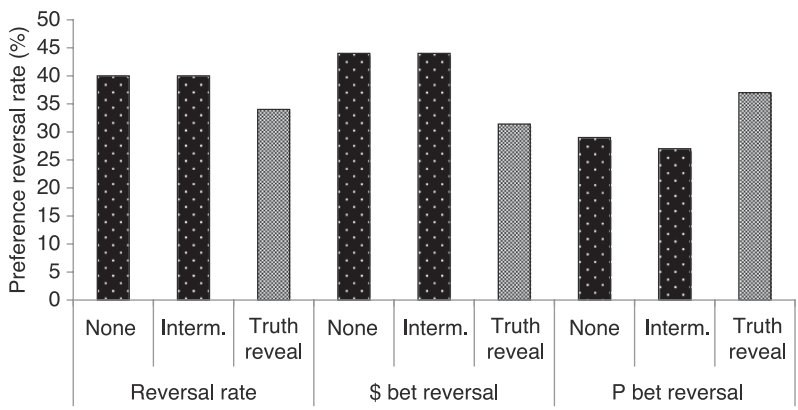

Figure 3.13 Whether the rate of preference reversal differs if there are no, truth-revealing or intermediate incentives. The overall reversal rate does not significantly change, but the rate of $ and P reversals becomes more similar. The $ and P reversal rates show, respectively, the proportion of those subjects pricing the $ or P bet higher that chose the P or $ bet.
Source: Berg, Dickhaut and Rietz (2010).

the overall preference reversal rate does not change much if there are truth-revealing incentives: it does go down, but not by a significant amount. This might suggest that incentives do not matter much.

If we look a little deeper, however, we do see that incentives start to matter. This is most apparent when comparing the $ bet reversal rate (the percentage of those subjects who priced the $ bet higher who subsequently chose the P bet) and the P bet reversal rate (the percentage of those subjects who priced the P bet higher who subsequently chose the $ bet). With non-truth-revealing preferences there is a systematic tendency for more $ bet reversals than P bet reversals. With truth-revealing preferences this disappears. This is important, because a systematic tendency for reversals of one kind and not another suggests something we can model and explain. If reversals are non-systematic they could be just due to mistakes or errors.

This study is not going to close the debate on preference reversals. We now have lots of evidence that there are systematic tendencies in preference reversals even with truth-revealing preferences. The study is also not going to close the debate on whether incentives matter (and we shall return to this issue a few times in the rest of the book). The study does illustrate the care needed in running and interpreting experimental results, however. The incentives that subjects are given do matter, but possibly in quite subtle ways that are not immediately apparent.

3.4.1 Procedural invariance

When describing prospect theory we focused on something Kahneman and Tversky called the **evaluation phase** and ignored the editing phase (see Research methods 3.4). The editing phase is important, though. In the **editing phase** Alan will organize and reformulate prospects to make them simpler to understand. This can consist of various operations. For example, in a coding phase he classifies outcomes as gains or losses. In a simplification phase he may round probabilities and outcomes up or down. It is basically a combination of mental accounting and the use of simplifying heuristics.

This editing and reformulating of prospects means that the way in which prospects are presented to Alan can affect his interpretation of them. We have already seen this in Table 3.9, where a different framing of the same prospect meant that outcomes were coded differently as gains or losses. A similar type of thing could explain preference reversals, because there is a big difference between asking Alan to price a prospect and asking him to choose between the prospect and another.

To pursue this possibility further, Tversky, Slovic and Kahneman also included the sure-thing prospect AC (see Table 3.17) and asked subjects whether they would prefer prospect AA or AC, and prospect AB or AC. To see why this is useful, suppose that Alan chose prospect AB over AA while pricing prospect AA more than $8 and prospect AB less than $8. This is a preference reversal. Table 3.18 summarizes four different ways to get such a reversal. For instance, if Alan chooses prospect AA over AC and AC over AB then it looks as though his preferences are intransitive. If, by contrast, he chooses prospect AC over both AB and AA then it looks as though his preferences are transitive but he just put too high a price on the $ bet, prospect AA.

Table 3.18 Four possible explanations for a preference reversal and the number of choices by subjects fitting each explanation. The sample is restricted to those who chose AB > AA, and price AA > $8 > price AB.

Explanation	Choices	Times observed (%)
Intransitivity	*AA* > *AC* and *AC* > *AB*	10.0
Overpricing of $ bet, AA	*AC* > *AA* and *AC* > *AB*	65.5
Underpricing of P bet, AB	*AC* > *AA* and *AB* > *AC*	6.1
Both overpricing of $ bet and underpricing of P bet	*AC* > *AA* and *AB* > *AC*	18.4

Source: Tversky, Slovic and Kahneman (1990).

I have also put in Table 3.18 the percentage of choices by subjects in each category. We see that most preference reversals in this study could be explained by an overpricing of the $ bet. Only 10 percent of preference reversals suggested intransitive preferences. It is plausible, therefore, that preference reversals are caused, in large part, by procedural invariance. To back this up it would be nice to have a good explanation for why there is procedural invariance. Two context effects that seem relevant are scale compatibility and the prominence effect.

Scale compatibility is when the importance of one aspect of an option (e.g. the probability or money that can be earned on a prospect) is enhanced by compatibility with the decision a person is asked to make. So, when Alan is asked to put a price on the $ bet, his attention is drawn to the relatively large payoff, causing him to overprice it. The **prominence effect** is when one aspect of an option becomes more prominent depending on the task. For example, the probabilities in a prospect may become more prominent when Alan is comparing prospects. So, when asked to choose between the P bet and the $ bet, he may focus on the higher probability of winning with the P bet.

Procedural invariance clearly seems relevant in explaining why we observe preference reversals. But have we done enough yet to completely rule out intransitive preferences?

3.4.2 Regret theory

In the study by Tversky, Slovic and Kahneman, around 10 percent of choices were suggestive of intransitive preferences. This might seem a relatively small percentage. If we think in a different way, however, we can get a much bigger percentage. For example, Loomes, Starmer and Sugden (1991) asked subjects to choose amongst several prospects like those in Table 3.17, and find that only 36 percent of subjects always made choices consistent with transitive preferences. Most subjects, therefore, made choices suggestive of intransitive preferences, some of the time.

You may be skeptical of this approach. Is it not a case of asking subjects to choose amongst so many prospects that they eventually make a mistake that we then interpret as intransitive preferences? It is probably not. The main evidence that it is not is the fact that subjects were far more likely to choose *AA* > *AC*, *AC* > *AB* and *AB* > *AA*. It was more rare to observe *AA* > *AB*, *AB* > *AC* and *AC* > *AA* or any other form of intransitivity. If subjects were making random mistakes we should not see such a

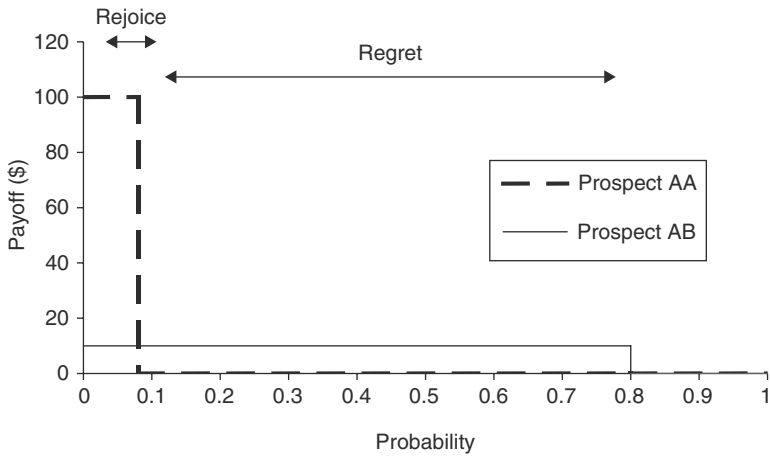

Figure 3.14 The expected regret and rejoice from choosing prospect AA rather than AB.

systematic tendency for one set of choices. It seems, therefore, that preferences can be intransitive.

None of the models of choice that we have considered so far in this chapter can lead to intransitive preferences (though I will come back to prospect theory later), so we need something new. One possibility is **regret theory**. Regret theory is most easy to apply when comparing two prospects, so, let's compare prospects AA and AB from Table 3.17.

The key to regret theory is a function $R(x, y)$ that assigns a number to any pair of payoffs x and y. In interpretation, $R(x, y)$ is the **regret** or **rejoice** that Alan experiences from getting payoff x if he would have got payoff y by choosing the other prospect. One slight complication is that, by choosing, say, prospect AA, Alan will, generally speaking, not know what his payoff would have been if he had chosen AB. We can get around this by thinking of distinct states of the world and saying what the outcome would be in each state of the world, ordering outcomes according to payoffs, as in Figure 3.14. In practical terms, this means that, if Alan chooses prospect AA and gets 100, the best outcome, then he assumes that if he had chosen prospect AB then he would have got 10, the best outcome. So, there is a 0.08 chance he will rejoice $R(100,10)$, a 0.72 chance he will regret $R(0,10)$ and a 0.2 chance he will have no regret $R(0,0)$. We then say that prospect AA is preferred to AB if $0.08\ R(100,10) + 0.72\ R(0,10) + 0.2\ R(0,0) > 0$. That is, prospect AA is preferred if the expected rejoice exceeds the expected regret.

Generalizing from this example, given any two prospects X and Y, we can distinguish possible states of the world (in the example there were three). We then say that

$$X > Y \text{ if and only if } \sum_i p_i R(x_i, y_i) > 0,$$

where i are possible states of the world and x_i and y_i are the outcomes of the respective prospects. Three assumptions are natural for the function R: (1) $R(x, y) = -R(y, x)$,

implying that regret is the opposite of rejoice; this also implies $R(0,0) = 0$; (2) $R(x, z) > R(y, z)$ if $x > y > z > 0$, implying that regret and rejoice are increasing in the difference in outcomes; and (3) $R(x, z) > R(x, y) + R(y, z)$ if $x > y > z > 0$, implying that a large difference in outcomes is regretted or causes rejoice more than two smaller differences combined.

Of primary interest for us is to see whether and why regret can lead to non-transitive preferences. The short story is that it can, and in particular preferences $AA > AC$, $AC > AB$ and $AB > AA$ are consistent with regret theory, while choices $AA > AB$, $AB > AC$ and $AC > AA$ are not. Remember, it was the first type of choices that is most often observed, so this is a really nice result.

[Extra] If you want the slightly longer story, then here it is. Applying regret theory, we get

$AA > AB$ if and only if $0.08R(100,10) + 0.72R(0,10) > 0$.

$AC > AA$ if and only if $0.08R(8,100) + 0.92R(8,0) > 0$.

$AB > AC$ if and only if $0.8R(10,8) + 0.2R(0,8) > 0$.

Now, let's consider whether it is possible that $AA > AB$, $AB > AC$ and $AC > AA$. Using symmetry, we know that $0.92\ R(8,0) + 0.2\ R(0,8) = 0.72\ R(8,0)$. It will also be convenient to write $0.8\ R(10,8) = 0.08\ R(10,8) + 0.72\ R(10,8)$. Doing this, we require the following to hold:

$$0.08[R(100,10) + R(8,100) + R(10,8)] + 0.72[R(0,10) + R(8,0) + R(10,8)] > 0.$$

We can rewrite this as

$$0.08[R(100,10) + R(10,8) - R(100,8)] + 0.72[R(8,0) + R(10,8) - R(10,0)] > 0.$$

If we now apply condition (3), we see that this is not possible (use $x = 100$, $y = 10$, $z = 8$, and $x = 10$, $y = 8$, $z = 0$). So, we cannot get preferences $AA > AB$, $AB > AC$ and $AC > AA$. In showing this, however, we have shown that we can get preferences $AB > AA$, $AA > AC$ and $AC > AB$!

Regret theory can explain why we observe non-transitive choices and can explain why we observe some choice cycles and not others. But what's the intuition behind this? Risk aversion could easily explain why prospect AB is preferred to AA and prospect AC is preferred to AB. What we need to do is explain why prospect AA is preferred to AC. That Alan would regret not having a gamble on prospect AA is one possible explanation. Essentially, he would regret missing out on $92. So, prospect AA is preferred to AC because of **regret aversion**. Regret theory does, therefore, give a plausible explanation for why preferences may be non-transitive.

3.4.3 Prospect theory and preference reversals

For another plausible explanation of preference reversals I am going to revisit prospect theory. To make this work we are going to have to make use of stochastic reference points.

Suppose that Alan is asked the minimum price for which he will give up the $ bet, prospect AA. Given that he is being asked to sell prospect AA, it makes sense, recalling the endowment effect, to think of that as his reference point. Applying the notion of a stochastic reference point, this would mean that, with probability 0.08, he thinks in terms of a reference point of $100 and, with probability 0.92, he thinks in terms of a reference point of $0. So, if he accepts some amount of money, say, $9, then, with probability 0.08, he will be down $91 on his reference point and, with probability 0.92, he will be up $9.

The minimum price he is willing to sell the $ bet, prospect AA, will then be where the expected gain equals the expected loss. Putting this into the standard prospect theory model with parameters $\lambda = 2.25$, $\alpha = 0.5$, $\gamma = 0.61$ and $\delta = 0.69$ gives a selling price of $16.7. If we do the same for the P bet, prospect AB, we get a selling price of $9.7. So, the $ bet is priced above the P bet.

[Extra] To see how we can get these numbers, let p_{AA} denote the price Alan puts on the $ bet. Setting the gains equal to the losses gives

$$\pi^g(0.92) \times p_{AA}^{\alpha} = \lambda \pi^l(0.08) \times (100 - p_{AA})^{\alpha}.$$

Using parameters $\lambda = 2.25$, $\alpha = 0.5$, $\gamma = 0.61$ and $\delta = 0.69$ implies that

$$0.74 \times p_{AA}^{0.5} = 2.25 \times 0.15 \times (100 - p_{AA})^{0.5}.$$

Squaring both sides gives

$$0.55 p_{AA} = 0.11(100 - p_{AA})$$

and $p_{AA} = \$16.7$ works. So, he would be willing to sell the $ bet for $16.7. Using the same logic for the P bet, prospect AB, we get that he would sell for p_{AB} where

$$\pi^g(0.2) \times p_{AB}^{\alpha} = \lambda \pi^l(0.8) \times (10 - p_{AB})^{\alpha}.$$

Working through the algebra gives $p_{AB} = \$9.7$.

Finally, we can ask whether he should choose prospect AA or prospect AB. In this case, the relevant reference point would appear to be his current wealth. So, he chooses the P bet if

$$\pi^g(0.08) \times 100^{0.5} < \pi^g(0.8) \times 10^{0.5}.$$

If you work this out, it comes to $1.5 < 2.1$, which is clearly true. So he would choose the P bet.

To recap, Alan priced the $ bet above the P bet and yet would choose the P bet over the $ bet. We have a preference reversal! Furthermore, we have done so because of an overpricing of the $ bet, consistent with the data in Table 3.18 (although $p_{AB} > \$8$). Basically, Alan would be reluctant to sell the $ bet both because of the overweighting of the small probability of winning and the large potential loss from selling a prospect that could have yielded him a high payoff.

Thus, prospect theory provides another plausible explanation for preference reversals. Is this explanation procedural invariance or intransitive preferences? I would say it is a bit of both. Framing is clearly relevant here, because the explanation relies on Alan perceiving that he does or does not have ownership of a prospect. The explanation also captures something of preferences, however; Alan puts a higher price on the $ bet because he does not want to regret missing out on a potentially big win.

3.4.4 Why preference reversals matter

We have seen that a combination of things seem to come together to make preference reversals a reality. But why is it important that we observe preference reversals? We saw in the last chapter that choice can depend on context effects for a host of different reasons; have we not just seen a further example of this? I want to finish this section by arguing that preference reversals go beyond being just another context effect.

To motivate the point, let's think back to the last chapter and Anna deciding what cereal to buy. According to Table 2.2, she would maximize her utility by choosing the Honey cereal. Maybe her preferences will change, but at that point in time the best thing she could do would be to buy Honey. The problem is, she may not know that this is the best thing to do, and because of context effects choose something different. For instance, if she is attracted to the bright red packaging of Superior she might choose that. This does not change the fact that Honey was what she should have chosen.

Now, let's think about Alan deciding what insurance to buy. Maybe we can say that full insurance is definitely the best thing for him to choose but, because of context effects, he might choose something else. This would be analogous to the story with Anna. If his preferences are intransitive, however, we get a quite different picture. Suppose, for example, that Alan prefers full insurance to no insurance, $FI \geq NI$, prefers no insurance to insurance against theft, $NI \geq AT$, and insurance against theft to full insurance, $AT \geq FI$. This does not seem so crazy.

With such preference it should come as no surprise that Alan's choice would depend on context effects. There is a deeper issue here, however, as to what is best for Alan. Whatever choice he makes we can find a choice that he seemingly prefers! There is no best thing he can do. This kind of thing plays havoc with the standard economic model. It also raises some interesting policy questions concerning what is in Alan's best interest. Intransitive preferences and preference reversals, therefore, matter and raise some fundamental questions. In choice with risk they might also be not so uncommon.

3.5 Summary

We started with the idea that the utility of a prospect is equal to the expected utility of the prospect. This is the standard way of modeling choice with risk in economics and is a simple, transparent way of doing so. It must be a good way to do so if preferences are transitive and satisfy independence.

We subsequently saw evidence that expected utility is a bit too simple to capture all that we observe. The Allais paradox, risk aversion for small gambles, risk-loving behavior over losses, and preference reversals are all hard to reconcile with expected utility. This motivated adapting how we model choice with risk.

One thing we considered was to take account of disappointment and elation when measuring the utility of a prospect. A second possibility was to weight probabilities to take account of pessimism and optimism. This is enough to explain the Allais paradox.

We subsequently looked at prospect theory and reference-dependent utility and showed that this can explain all the other deviations from expected utility that we looked at. We also saw how the presence of risk suggests additional reasons, to those we looked at in the previous chapter, why the reference point can depend on context effects. Expectations of what might happen are particularly important.

Finally we looked at preference reversals and the possibility that preferences are intransitive. Regret theory is one way to explain this.

There can be no denying that the things we have looked at in this chapter give us a much better understanding of choice with risk. The one minor problem is that we have lots of different ways of thinking that individually make sense and can explain a lot, but not everything. For example, rank-dependent expected utility can explain the Allais paradox but not preference reversals. Regret theory can explain preference reversals but not why someone would be risk-loving over losses. On this score, prospect theory looks like a clear winner, because it can explain all that we have seen. Arguably, however, this is only because we have the flexibility to change the reference point to something that works. The search still goes on, therefore, for a unified theory of choice with risk. Until we have such a theory it's important to think carefully which model or way of thinking is the one best adapted to the question you are interested in.

3.6 Financial trading

The first application I want to look at in this chapter is financial trading. Investing in stocks and shares is, clearly, an inherently risky business, and so the models and ideas that we have developed in this chapter should be able to tell us something useful about it. To illustrate that they can, I have picked out three interrelated things to talk about: the equity premium puzzle, disposition effect and ostrich effect. A particular reason for choosing these three things is that, in looking at them, not only do we learn something useful about investor behavior, but we also need to think more about the importance of time and risk, an issue we looked at only briefly in section 3.3.3.

3.6.1 The equity premium puzzle

The **equity premium puzzle** is that the real rate of return on stocks is much higher than the return on 'safer assets' such as Treasury bills. Figure 3.15 gives some data to illustrate the size of these differences. Such differences in annual return can multiply to huge differences over a few years. For example, $1 invested in stocks in the United States in 1926 would have been worth over $2,500 by 2000, while the

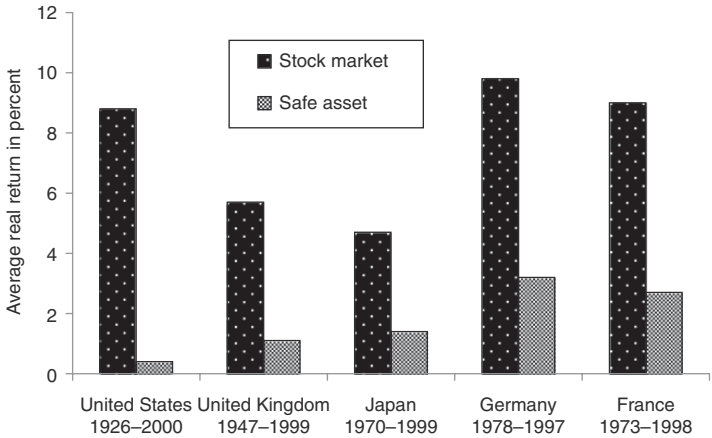

Figure 3.15 The equity premium puzzle. The average annual return on the stock of money invested exceeds that invested in a safe asset such as Treasury bills.
Source: Mehra and Prescott (2003).

same amount invested in Treasury bills would have been worth less than $20. Given such numbers, why would anyone invest in Treasury bills?

The obvious answer would be that investors are risk-averse and so invest only in stocks, the riskier of the two, if the average return is higher. The equity premium that we observe could not be explained solely by risk aversion, however, because to do so would require a coefficient of relative risk aversion of around 30, which is implausibly high. We need a different explanation. One explanation that does work is loss aversion and prospect theory. To make it work, however, we are going to have to think about a complicating factor: time.

The prices of stocks are changing minute by minute during a working day, and in principle an investor could change his investments at any time. There is, therefore, no definite moment in time when the outcome of investing in the stock market is known. Instead, an investor must choose how long to leave his investments before he evaluates how they are doing. We call this length of time the **evaluation period**. It seems sensible to assume that the investor gets utility when he evaluates his investments at the end of each evaluation period. For example, Alan might check how his investments are doing every Friday morning, and it's then that he perceives any loss or gain.

The evaluation period is going to be crucial. To explain why, it is easiest to work through an example. In Figure 3.16 I have plotted the value over time of $100 invested in two hypothetical assets. In any month the risky asset could fall or rise by $10 while the safe asset grows by a steady $1. If we just look at the value of the investment, in the top half of the figure, then the risky asset looks a clear winner. Things are not so simple, however, when we look at the investor's utility.

Let's look first at the case when the evaluation period is one month and Alan invested in the risky asset. In the first month the value of the asset went up $10. If we plug a gain of $10 into the prospect theory equation (3.7) with value $\alpha = 0.88$ then we get a utility of $10^{0.88} = 7.6$. The same thing happened at the end of month two.

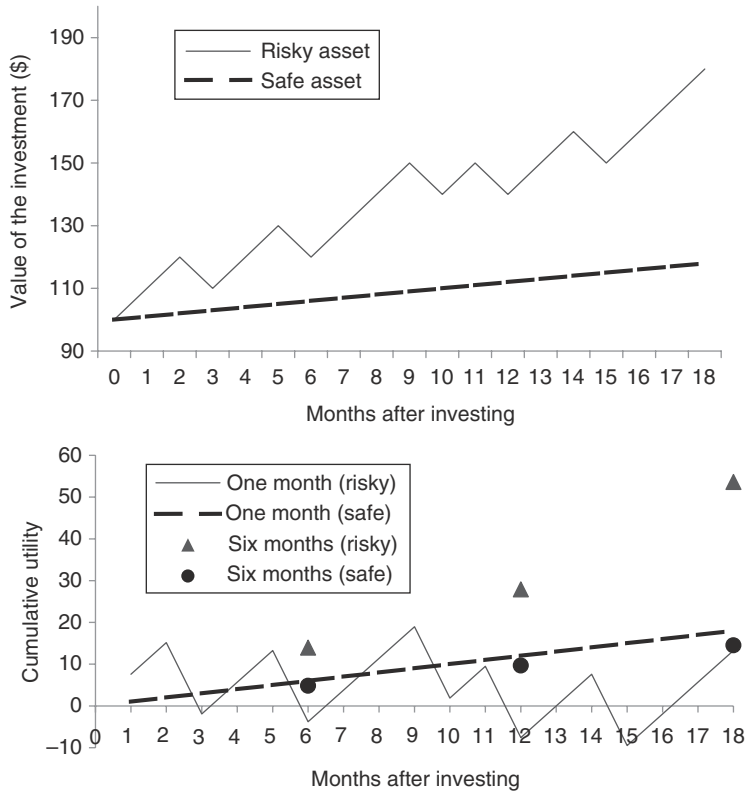

Figure 3.16 The returns in money and utility from investing in two hypothetical assets. If the evaluation period is one month, the investor gets more cumulative utility from the safe asset because he avoids the frequent falls in the risky asset. If the evaluation period is six months, the investor gets more cumulative utility from the risky asset.

What about at the end of month three? In the third month the asset fell in value by \$10. If we plug a loss of \$10 into equation (3.7) with values $\lambda = 2.25$ and $\beta = 0.88$, we get a utility of $-2.25 \times 10^{0.88} = -17.1$. This is bad, and bad enough to wipe out the positive utility of the previous two months. Given Alan's loss aversion, his cumulative utility after three months would have been higher if he had invested in the safe asset and avoided this loss. The same can be said after 18 months.

Now suppose the evaluation period is six months. At the end of every six months the value of the risky asset has gone up. Consequently, the cumulative utility from the risky asset works out to be a lot higher than that from the safe asset.

What we see in this example is how the volatility of a risky asset can make it undesirable for someone who is loss-averse, even if the overall returns are high. For every month that the investment goes down in value there need to be three months of gain to compensate Alan for the disutility of the bad month. This effect is lessened by lengthening the evaluation period. The longer the evaluation period, the less chance the investment will have decreased in value, and so the less chance that

Alan feels a loss. This means there should be an evaluation period when he is relatively indifferent between the risky and safe asset.

Benartzi and Thaler (1995) asked how long the evaluation period would have to be to explain the equity premium that we do observe. That is, how long the evaluation period would have to be to make the expected utility on safe assets the same as on stocks. They evaluated this by simulating the utility an investor would have got from investing in US stocks, bonds or Treasury bills for various time intervals between 1926 and 1990. As in our stylized example they find that, with an evaluation period of one month, estimated utility would be higher from investing in safe assets. With an evaluation period of 18 months, it would be higher from stocks. The point where the estimated utility from safe assets and stocks are equal is an evaluation period of 12 months.

The estimate of 12 months is remarkably plausible! It does seem intuitive that ordinary investors might use a one-year period to evaluate returns on investments. Managers of pension and investment funds might also be evaluated on their year-by-year performance. We have, therefore, an explanation for an equity premium. This explanation is based on people being both loss-averse and **myopic**, in the sense that the evaluation period is shorter than it could be.

3.6.2 The disposition effect

Our next puzzle is why investors prefer to sell stocks that are trading at a gain and hold on to stocks trading at a loss. To explain, suppose that Alan sells a stock on a particular day. We can ask whether there was a **realized gain** – the stocks were sold for more than was paid for them – or a **realized loss** – the stocks were sold for less than was paid for them. We can also look at all the other stocks in his portfolio that are not sold and ask whether there was a **paper gain** – the stocks could have been sold for more than was paid for them – or a **paper loss** – the stocks could have been sold only for less than was paid for them. Next we calculate two ratios:

$$\text{Proportion of gains realized} = \frac{\text{Number of realized gains}}{\text{Number of realized gains} + \text{Number of paper gains}}$$

$$\text{Proportion of losses realized} = \frac{\text{Number of realized losses}}{\text{Number of realized losses} + \text{Number of paper losses}}$$

The **disposition effect** is that the proportion of gains realized is significantly greater than the proportion of losses realized. In other words, investors prefer to sell stocks that will give them a realized gain than to sell stocks that will give them a realized loss. To give some numbers, we can look at a study by Odean (1998), who analyzed the trading activity from 1987 to 1993 of 10,000 households that held accounts at a large brokerage firm. He finds that the proportion of gains realized was 0.148 and the proportion of losses realized only 0.098.

You might be thinking that the disposition effect is just good investing. For example, if a stock has gone up in price, meaning realized gains, maybe it will subsequently fall in price, and so it is better to sell while the going is good. Unfortunately, the

evidence suggests otherwise. Odean finds, for instance, that stocks that were sold for a realized gain earned on average 3.4 percent more over the subsequent year than stocks that were not sold for a realized loss. Again, we need a different explanation. Given that prospect theory did so well at explaining the equity premium puzzle, maybe it can work again?

I shall stick with a very similar story to the one that we used to think about the equity premium puzzle. So, imagine assets that every month either go up or down in value. At the end of the year Alan will evaluate his portfolio. The utility he feels at that point depends on the gain or loss on his portfolio and is consistent with the prospect theory equation (3.7) with $\lambda = 2.25$ and $\alpha, \beta = 0.88$. Up until now everything is the same as when we looked at the equity premium puzzle. Here is the new thing: during the year Alan can check how his portfolio is doing and buy or sell assets if he wishes to. When he checks on his investments he does not feel any loss or gain; that comes only at the end of the evaluation period.

Now, suppose that Alan checks how his investments are doing six months into the year. If his investments have gone down during the first six months then, because of loss aversion, he will be very eager to regain those losses before the evaluation period ends. By contrast, if his investments have gone up during the first six months then he may be keen to guarantee those gains to avoid any loss. This looks like a possible explanation for the disposition effect, but we need to check that (see Research methods 3.8). Barberis and Xiong (2009) did that by working out the optimal strategy of an investor in such a situation, and then simulating asset prices to calculate the proportion of gains realized and losses realized that we would expect to see.

Figure 3.17 plots some of the data from the study. Recall that we are trying to explain why the proportion of gains realized is larger than the proportion of losses realized. Figure 3.17 suggests that prospect theory will actually give us the opposite! For the most part, we predict that the proportion of losses realized should exceed the proportion of gains realized.

To try and understand why we obtain this result we can work through an example illustrated in Figure 3.18. Suppose that Alan will evaluate his portfolio after 12 months and check it once after six months (the $T = 2$ case in Figure 3.16). He is interested in only one asset, which starts the year at price \$25 and every six months either goes up in value by \$5 or down by \$2.50. Alan buys two shares at the start of the year, investing \$50. We need to ask what he will do after six months if the price has gone up or down.

First, suppose that it has gone up. This means he is at point G in Figure 3.18. He could sell at this point and be guaranteed a gain by the end of the year, but why stop there? He could buy two more shares, still be guaranteed to not make a loss and have the chance to be at point 2G if the price goes up again. It turns out that, if he was willing to invest at the start of the year (when there was the chance of making a loss), then he must want to increase his investment at this stage (now that he is guaranteed not to make a loss).

Now, suppose the price has gone down after six months. This means he is at point L. He would like to regain his losses, but also wants to avoid further losses. This is possible by selling one share and keeping the other. If the price goes up then he will be back to zero by the end of the year, and if the price goes down a second time

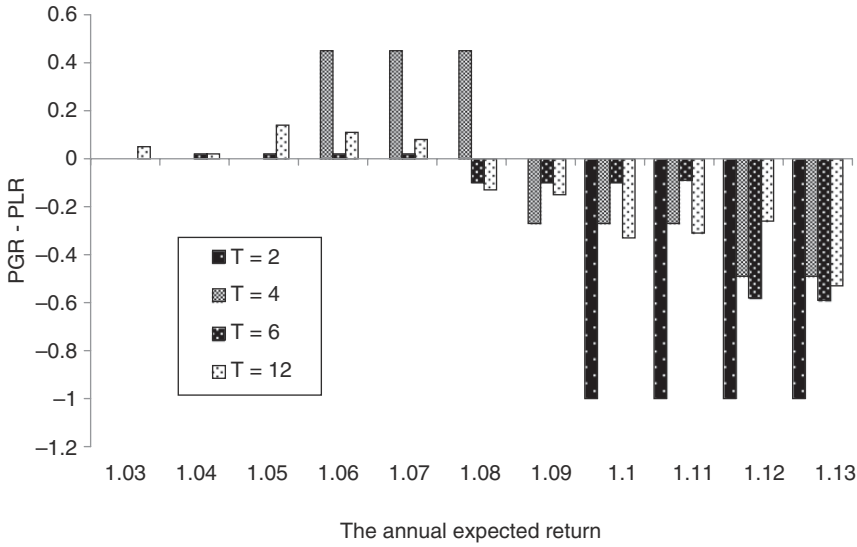

Figure 3.17 The proportion of gains realized (PGR) minus the proportion of losses realized (PLR) in simulated investments. T is the number of times that the investor looks up his portfolio during the evaluation period. The disposition effect is that the proportion of gains realized exceeds the proportion of losses realized. Prospect theory appears to give the opposite.
Source: Barberis and Xiong (2009).

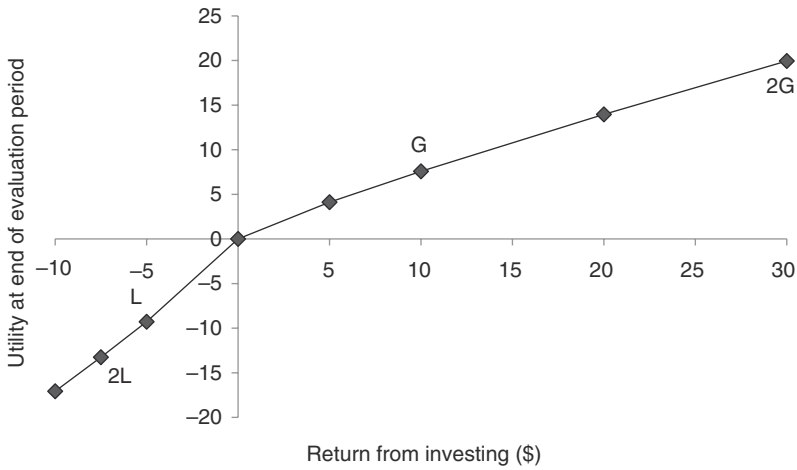

Figure 3.18 Why prospect theory does not cause the disposition effect. If the investor is at point G after six months, he should increase his investment and aim for 2G. If the investor is at point L after six months, he should decrease his investment to make sure he does no worse than 2L.

he will at least only be at point 2L with losses of $7.50. Again, this will be the optimal thing to do.

We see through this example that prospect theory fails to explain the disposition effect on both counts. If Alan had realized gains then he should increase his investment, and if he had realized losses he should decrease his investment. You might say this result is dependent on the asymmetry whereby the price of the asset increases by more than it decreases. This is what we do tend to observe about asset prices, however. Furthermore, if Alan was willing to invest at all, despite his loss and risk aversion, then there must have been this asymmetry, because he will invest only in stocks that are more likely to go up than down.

Prospect theory does not, therefore, seem a convincing explanation for the disposition effect. Indeed, it suggests we have even more of a puzzle to solve than we might have first imagined. Not all hope is lost, because we might argue that the investor will feel a loss or gain in utility every time he looks at the value of his portfolio and not just at the end of the evaluation period. Perhaps, however, it is better to accept reality. Prospect theory was able to explain the equity premium puzzle without any trouble, but raises more questions about the disposition effect than it solves. In Chapter 5 we shall come back and try to answer these questions.

3.6.3 The ostrich effect

In both the previous puzzles, time has played a pivotal role. First we looked at the evaluation period over which gains and losses are judged. Then we looked at how often an investor may check up on and trade within the evaluation period. The point I shall pick up on now is that the evaluation period and frequency of trade are not only pivotal but also typically chosen by the investor. This raises the issue of how often Alan should pay attention to his investments.

To get some insight into this issue, imagine that Alan periodically hears news about the performance of the stock market. This news gives him some idea what may be happening to his investments. It is only by looking up his investments, however, that he can definitively find out what has happened to them. What should he do after he has watched the news?

❖ RESEARCH METHODS 3.8 DON'T ALWAYS TRUST INTUITION

Ronald Reagan once said: 'One definition of an economist is somebody who sees something happen in practice and wonders if it will work in theory.' A nice quote, which sums up a lot of what economics is about. Many question whether this formalism is of any use, as is clear in a second quote from Alfred Knopf: 'An economist is a man who states the obvious in terms of the incomprehensible.' This study by Barberis and Xiong (2009) is one illustration of why economists believe it is good to check that things work in theory.

Prospect theory was an intuitive solution to the disposition effect. If a stock has risen in value it takes an investor to the concave risk-averse area of her utility function and if it falls in value it takes her to the convex risk-loving area of

her utility function. Consequently, she should take more risks on stocks that have fallen in value than gained in value. It was so intuitive an explanation that it became the most commonly mentioned explanation for the disposition effect. Barberis and Xiong took on the task of checking that the explanation does work in theory (to make the obvious incomprehensible). Their conclusion was that prospect theory is not so great an explanation for the disposition effect. Sometimes, therefore, intuition and common sense can lead us astray. That's why it's important to write down models and check that things really do work how they seem to.

If an investor feels the loss or gain of investing only when he evaluates his investments at the end of the year, then he should look up his investments, because he has got nothing to lose. Realistically, however, he is going to feel something when he checks on his investments, even if the biggest effect is at the end of the evaluation period. We can distinguish, therefore, two effects. First, there is an **impact effect**, whereby looking up his investments makes it more likely he will feel the gain or loss on his investments. Second, there is a **reference point updating effect**, whereby he is more likely to change the reference point by which he will evaluate his investments at the end of the evaluation period.

These two effects work in opposite directions. If the news is good then the investor will be tempted to look up his investments and feel happy about the gain that he is likely to observe. But he will also be reluctant to see the gain, increase his reference point and subsequently be more likely to experience a loss relative to this higher reference point. There are similar trade-offs if the news is bad. Intuitively, however, the reasons to look up the investments after good news seem a bit stronger than those

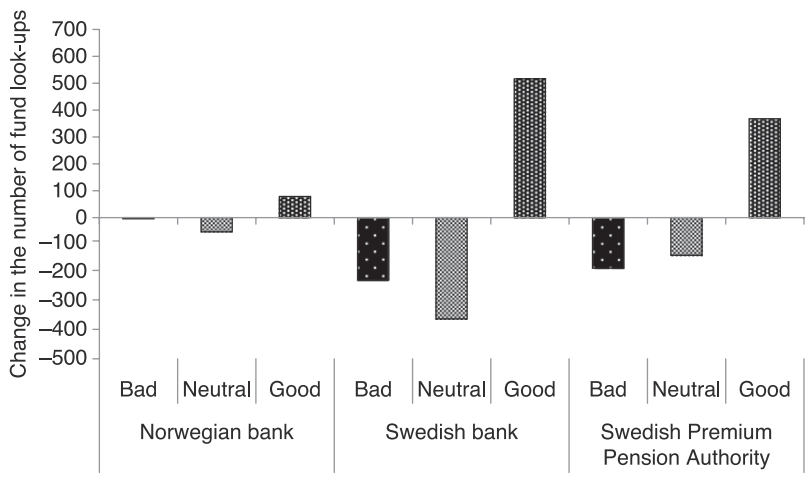

Figure 3.19 The change in the number of investors looking up their investments depending on whether the stock market was up, down or neutral over the previous week.
Source: Karlsson, Loewenstein and Seppi (2009).

to look up after bad news. This suggests the **ostrich effect**, whereby the investor is more likely to look up his investments following good news rather than bad.

A study by Karlsson, Loewenstein and Seppi (2009) looked at data from Norway and Sweden on how often investors do look up their investments. They compared the daily number of investors who look up their portfolio to the performance of the relevant stock market index in the previous week. Figure 3.19 summarizes the data and shows that there does indeed appear to be an ostrich effect. When the stock market index was going up (good news), investors were more likely to look up their investments than when it was flat or going down (neutral or bad news).

The consequences of an ostrich effect are interesting to hypothesize about. In discussing the equity premium puzzle, we suggested that investors may have too short an evaluation period. The ostrich effect partly alleviates this. In discussing the disposition effect we suggested that investors may not do the best thing when they look up a portfolio with gains and losses. The ostrich effect may be a way to avoid this. Might selective attention, therefore, be a good way of managing loss aversion and the like?

3.7 Insurance

I started this chapter with the example of Alan buying insurance, so it seems apt to look at some real data on insurance. One basic observation is that people seem to want insurance. Most people buy at least one of the following: house insurance, life insurance, car insurance, health insurance, travel insurance. In some instances this is because insurance is compulsory, but that can explain only so much. It seems, therefore, that people like insurance, and this clearly suggests that people are risk-averse for big gambles. That should be no surprise, given that we observe risk aversion for small gambles, and so it is not too hard to explain why people might want insurance.

What I will focus on, therefore, is a more subtle question of how much insurance people choose to have. One particularly interesting issue is that of the deductible. The **deductible** is the amount of any insurance claim that a person would have to pay themselves. So, for example, if Alan's car is involved in an accident with total damage

Table 3.19 Choice of deductible on home insurance. Few households chose the $1,000 deductible, but this is hard to explain using expected utility because it requires extreme levels of risk aversion.

Deductible	Proportion who chose it (%)	Average extra premium	Claim rate (%)	Expected saving with $1,000 deductible	Lower bound on risk aversion using CRRA
$1,000	17.0	$0	2.5	—	—
$500	47.6	$100	4.3	$80	1,719
$250	35.0	$159	4.9	$127	4,000
$100	0.3	$243	4.7	—	—

Source: Sydnor (2010).

of $2,000 and his deductible is $500, then Alan pays the first $500 and the insurance company the remaining $1,500.

Insurance companies typically give some choice of deductible, and a policy with a lower deductible will have a higher premium. In choosing the level of deductible, Alan faces a clear trade-off: by spending, say, $10 more in premium he will have to pay, say, $100 less if he makes a claim. The choice of deductible is, therefore, a nice thing to look at. It can give us a picture of how people choose insurance and also give us some general insights into choice with risk.

Sydnor (2010) looked at a sample of 50,000 standard policies from a large home insurance provider in the United States. Customers could choose a deductible of $1,000, $500, $250 or $100. Table 3.19 summarizes some of the data and shows that the majority chose the $500 or $250 deductible. The exact amount of premium would depend on the house value, location, etc., but on average people paid around $100 or $150 extra to lower the deductible.

There was evidence, which we cannot see in Table 3.19, that households were price-sensitive, in the sense that those buying the $250 deductible paid on average less extra premium than those who chose the $500 deductible would have had to do. Even so, the amount that households spent to lower the deductible looks high. We can say this with some confidence because we know the actual claim rate of the customers and so know how much they would have benefited from the lower deductible. The costs of the deductible do not seem to justify the benefits, and the estimated level of risk aversion to justify such choices is way above what is plausible.

How can we explain the preference for a lower deductible? One likely factor is an overestimation of the likelihood of making a claim. If the claim rate is 20 percent, then the estimates of risk aversion come down to around ten. This, in itself, is still not enough, however, because an expected claim rate of 20 percent would be way too pessimistic, and a coefficient of risk aversion of ten is still too high. A second likely factor is loss aversion.

Care is needed in applying loss aversion to insurance, however, because paying the premium might be interpreted as a loss. It can, therefore, become even harder to explain why someone would incur for sure the loss of, say $100, to lower the probability that they will subsequently lose $500. Indeed, this would directly contradict the evidence that people are risk-loving over losses. One way out of this puzzle is to apply the 'no loss in buying' hypothesis that we looked at in the last chapter. Recall that this hypothesis says giving up money to buy something is perceived as a forgone gain rather than a loss. So, spending an extra $100 on insurance is not perceived as a loss. Recall also the example in section 3.3.2.

If we accept the 'no loss in buying' hypothesis and apply the standard formulation of prospect theory with $\lambda = 2.25$, $\alpha, \beta = 0.88$ and $\delta = 0.69$, we find that a 4 percent claim rate would be weighted as 12 percent and a person would pay $100 for a $500 deductible. A combination of loss aversion and overestimating the probability of making a claim can, therefore, explain the choice of deductible that people made. Prospect theory again comes up trumps. It is crucial, though, to recognize that we need both loss aversion and an overweighting of small probabilities to makes sense of observed behavior. Either on its own is not enough.

Even with that, things may not be quite so simple. To give a first hint why, we can look at a study by Cohen and Einav (2007). They looked at the deductible choices of

Table 3.20 Deductible choice in Israeli car insurance. Monetary amounts were converted from the actual New Israel Shekels by dividing by 3.5, the average exchange rate over the period.

Deductible chosen	Average deductible	Proportion who chose it (%)	Extra premium	Claim rate (%)
Very high	$1,075	0.5	0	13.3
High	$745	0.6	$69	12.8
Regular	$415	81.1	$182	23.2
Low	$250	17.8	$237	30.1

Source: Cohen and Einav (2007).

over 100,000 individuals buying car insurance in Israel between 1994 and 1999. Again, people could choose from four deductible levels, called regular, low, high or very high. This time the premium and deductible were percentages of the regular premium and deductible. For example, someone opting for a low deductible paid 6 percent more than the regular premium and got 60 percent of the regular deductible. Table 3.20 gives some numbers to compare with those of Table 3.19.

The financial trade-offs that people faced were roughly similar to those of the Syndor study, with one notable exception: the claim rate is much higher. Unfortunately, it is difficult to calculate a precise estimate of risk aversion from the data. If we focus on the fact that the vast majority of people chose the regular deductible, and not the high or very high deductible, then we would obtain high estimates of risk aversion, similar to those in the study by Syndor, and too high to be plausible. Cohen and Einav noted, however, that customers may not have been informed about the high and very high deductible options when deciding what to do. In this case, we should focus on the fact that people chose the regular and not the low deductible. Doing this, we find low and plausible estimates of risk aversion of less than four.

We are simply unable to tell from this data what the attitudes to risk were. That, in itself, is informative, though, because it emphasizes that many things other than attitudes to risk will likely influence an insurance decision. It is surely not a coincidence that so many chose the regular option. Indeed, Syndor also found considerable inertia in his study, with many choosing the same deductible as they had in the past. To carry on this story we need to look at choice over time, which is the focus of the next chapter.

3.8 Tax evasion

Benjamin Franklin famously said: 'In this world nothing can be said to be certain, except death and taxes.' It might have been more accurate to say that nothing can be certain except death, taxes and people not paying their taxes. You typically don't have to wait long before the latest story of tax evasion hits the headlines. Usually, it is a politician, movie star, sports star or rich businessman caught underdeclaring income or moving money offshore. Tax evasion is not solely the domain of the rich, however. It goes a lot deeper than that. And that's why it is a big concern in many countries.

Tax evasion has two bad consequences. First, it is unfair that one person honestly pays their taxes while another illegally avoids doing so. Second, it creates uncertainty at the macroeconomic level, because a government has less security of income. This latter problem became all too clear in the Eurozone crisis of 2009 onwards, when large-scale tax evasion in some countries made it very difficult to know how a change in policy (e.g. an increase in tax rates) would feed through into government revenue. Eliminating tax evasion is, therefore, an important goal. In this section we shall look at how behavioral economics can inform as to the reasons why people evade taxes, and the things that might stop them doing so.

Before getting into the theory it is interesting to get a feel for how big a problem tax evasion is. Measuring tax evasion is inevitably difficult, given that evaders have a strong incentive to cover their tracks. Various indirect ways to measure tax evasion have been developed, however. Most reliable are measures that take account of a combination of different indicators of unreported activity. The kinds of indicators to look for are aggregate spending above reported income, debt out of line with reported income, a discrepancy between official employment statistics and the number of people seemingly in work and the use of cash rather than checks or bank cards.

Figure 3.20 provides estimates on the size of the shadow economy between 1999 and 2007 in a wide selection of countries. The shadow economy measures that part of the economy that goes unreported, and so is a good proxy for tax evasion. In Brazil, Nigeria and Russia the shadow economy is estimated at a whopping 40 percent of the economy. Even at the other end of the scale, in countries such as the United States and Japan, the shadow economy is still estimated at around one-tenth of the

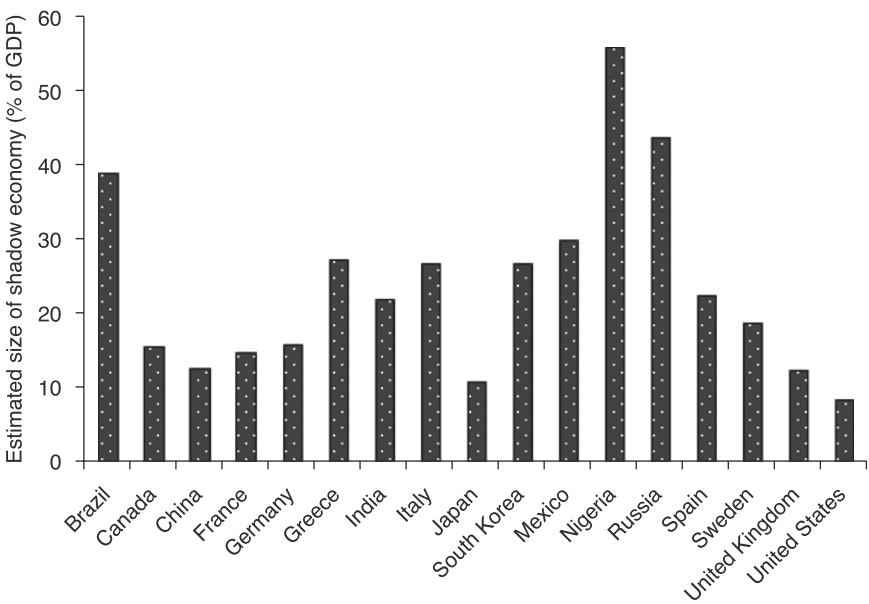

Figure 3.20 The estimated size of the shadow economy between 1999 and 2007 as a percentage of GDP.
Source: Buehn and Schneider (2012).

Table 3.21 Estimated tax evasion by profession. The income multiplier measures how much lenders are estimated to multiply reported income in order to obtain a more realistic estimate of true income.

Profession	Average tax evasion (€)	Income multiplier
Doctors and medicine	29,343	2.45
Engineering and science	28,625	2.40
Education	24,742	2.55
Accounting and financial services	24,572	2.22
Law	24,032	2.24
Lodging and restaurants	15,884	1.99
Construction	13,919	1.85
Agriculture	9,288	1.75
Retail	5,215	1.27

Source: Artavanis, Morse and Tsoutsoura (2015).

economy. Table 3.21 focuses on Greece and looks at the estimated tax evasion of self-employed workers by profession. Here we can more clearly see the potential unfairness that results from tax evasion. Some professions are seemingly able to 'get away' with more than others.

3.8.1 Standard model of tax evasion

We begin our efforts to model tax evasion with the standard model, based around expected utility maximization. To explain the model let us look at Alan's choices. Suppose that Alan earns an income of $\$M$. He must decide how much of that income to report to the tax authorities. For every dollar he reports he is taxed at rate t. Every dollar he doesn't report is his to keep, provided he is not caught evading taxes. With probability p, the tax authorities will audit Alan and discover all his unreported income. For every dollar he did not report he is fined $f > t$.

Alan must decide how much income to report. Equivalently, he must decide amongst an array of prospects. He could choose to be fully compliant with the tax authorities, meaning he chooses prospect $(1, \$M(1 - t))$. In this case, it doesn't matter whether he is audited or not, because he did not break the rules. At the other extreme he could try to evade all taxes, meaning he chooses prospect $(1 - p, \$M; p, \$M(1 - f))$. In this case, if he is not audited then he avoids paying tax, but if he is audited he has to pay a fine that is more than his taxes would have been. In general, let $\$ R$ denote the amount of income Alan chooses to report and $\$ D = M - R$ the amount that he does not report. Then, he chooses prospect $(1 - p, \$R(1 - t) + D; p, \$R(1 - t) + D(1 - f))$. If he is not audited he pays taxes only on reported income. If he is audited then he pays taxes on reported income and a fine on unreported income. The expected value of the prospect is

$$V(R, D) = R(1 - t) + D(1 - pf).$$

To see more clearly how things work, let's put in some realistic numbers. Suppose that Alan has a wealth of $1 million and has an income, on top of that, of $100,000. The tax rate is 40 percent, $t = 0.4$, the probability of an audit is 2 percent, $p = 0.02$, and

Table 3.22 A summary of the outcome if Alan reports all, half or none of his income when $t = 0.4$, $p = 0.02$ and $f = 3$. Expected value is increasing in the amount not reported, but not reporting increases the risk of a big fine.

Income reported	Taxes paid	Fine if audited	Expected value	Total wealth	
				Not audited	Audited
100,000	40,000	0	60,000	1,060,000	1,060,000
50,000	20,000	150,000	77,000	1,080,000	930,000
0	0	300,000	94,000	1,100,000	800,000

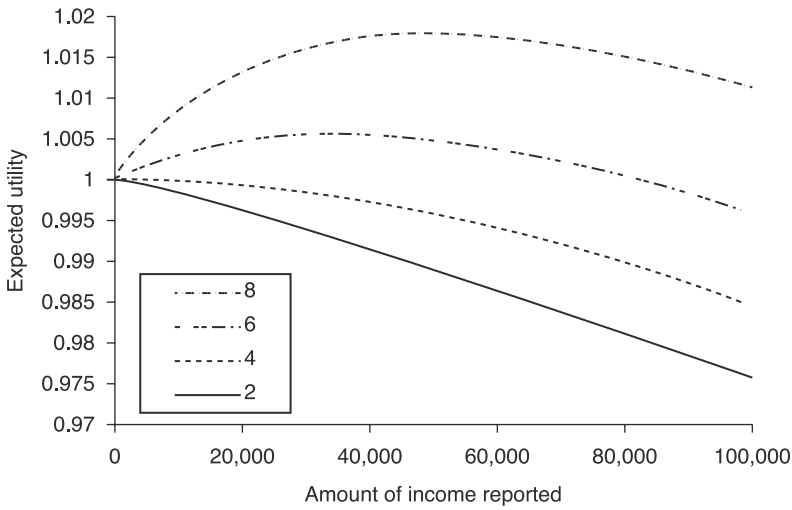

Figure 3.21 Expected utility (rescaled) for different amounts of income reported and for different values of the risk aversion parameter r_u^r.

the fine is 300 percent of unreported income, $f = 3$. Table 3.22 looks through three levels of reported income that Alan could choose. He is supposed to pay \$40,000 in tax. By trying to avoid this tax he increases expected value. This is primarily because it is unlikely he will be audited. Avoiding tax does, however, carry the risk of a big fine in the case that he is audited.

To see what Alan will choose we could calculate expected utility. I will assume that Alan has a constant relative risk aversion utility function with risk aversion parameter r_u^r. This means that the expected utility of the prospect when \$$R$ is reported is

$$U(R, D) = (1 - p)\frac{(R(1 - t) + D)^{1-r_u^r}}{1 - r_u^r} + p\frac{(R(1 - t) + D(1 - f))^{1-r_u^r}}{1 - r_u^r}.$$

This is easy enough to work out given specific values. Figure 3.21 plots expected utility as a function of reported income for different values of the risk aversion parameter. Most realistic is a risk aversion parameter of two. We can see in this case

that Alan will not report any income, because his expected utility is decreasing in the amount of income that he reports. If Alan is more risk-averse then his expected utility is maximized by reporting some of this income. Even for the implausibly large risk aversion parameter of eight, however, he would report only half his income.

This specific example illustrates a general problem with the standard model of tax evasion. For plausible parameter values it predicts that everybody should be doing a lot of tax evasion. It is difficult to reconcile this prediction with the largely honest reporting that we observe in many countries; many evade taxes but many do not. Put another way, the standard model says that the tax authorities are not tough enough. The probability of audit and the size of fine are predicted to be an insufficient deterrent against tax evasion. Again, this prediction doesn't seem to fit the facts. So, what are we missing?

3.8.2 Behavioral theories of tax evasion

In Chapter 2 we looked at evidence that people do not like to lie or deceive. Maybe this can explain why people honestly report income. It may be part of the story, but it is unlikely to be the whole story. This is because a person must have a pretty big aversion to lying in order to file an honest tax return. If Alan, for example, does not report any income he gains an expected $34,000. A lot of people would be willing to tell a lie for $34,000!

For a more likely explanation of honest reporting we can look at things such as disappointment, reference dependence and the overweighting of small probabilities. To illustrate, let's see what happens if we apply the reference-dependent utility approach. To do that, we need to think about expectations and the reference point. I will compare three different scenarios.

First, suppose that Alan never pays his taxes and resents having to give money to the government. In this case his reference point is a final wealth of $1,100,000. Paying any amount of tax is interpreted as a loss. Following the approach of section 3.3, we shall first rank outcomes. The worst outcome is clearly that he gets audited, and the best outcome is that he does not get audited. So, the decision weight on being audited will be

$$w_1 = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{\frac{1}{\delta}}}$$

and the decision weight on not being audited will be $w_2 = 1 - w_1$. If we put in values of $p = 0.02$ and $\delta = 0.69$ then we get $w_1 = 0.06$ and $w_2 = 0.94$. So, a 2 percent probability of being audited gets converted into a 6 percent decision weight, because of the overweighting of small probabilities.

This overweighting of the probability of being audited should push Alan more towards reporting his income, but is it enough? The next thing to do is to look at his utility function. I'll assume the utility function is given by

$$u(x) = \begin{cases} \sqrt{x} + (x-r)^{0.88} & \text{if } x \geq r \\ \sqrt{x} - 2.25(r-x)^{0.88} & \text{if } x < r \end{cases},$$

where x is final wealth and r is the reference point. Putting in $r = 1{,}100{,}000$ we can work out expected utility for any prospect. Figure 3.22 shows what we get. And you

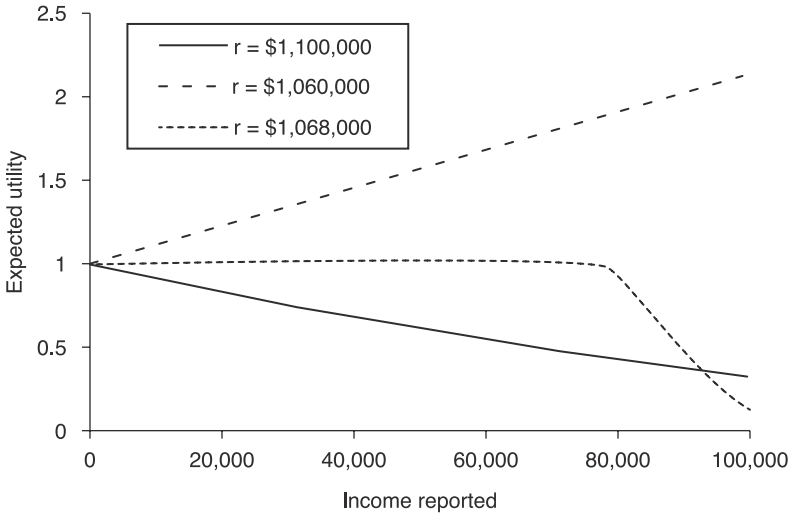

Figure 3.22 Expected utility (rescaled) for different amounts of income reported depending on the reference level.

can see that Alan still maximizes his utility by not reporting any income. So, the overweighting of probabilities is not enough in itself.

Now let us suppose that Alan always pays his taxes and believes in big government. In this case, it is natural to think of his reference point as a final wealth of $1,060,000. Alan could potentially gain by evading tax. But the only way he can lose is to evade tax and be caught. Given Alan's desire to avoid losses, this change in reference point makes a huge difference to his incentive to report taxes. The decision weights also change. The decision weight on being audited remains the same as before at 0.06. Given that being not audited now counts as a gain, however, the decision weight will be

$$w_2 = \frac{(1-p)^\gamma}{(p^\gamma + (1-p)^\gamma)^{\frac{1}{\gamma}}}.$$

Setting $p = 0.02$ and $\gamma = 0.01$ gives $w_2 = 0.87$. Notice how decision weights no longer need to add to one because being audited is a loss while being not audited is a gain. This effectively means that the decision weight on being audited becomes even higher. We can see in Figure 3.22 that these effects combined mean that Alan maximizes his utility by reporting all his income. We have a viable story of why Alan may honestly pay his taxes.

I want to consider a final possibility. This is arguably the most realistic in looking at tax evasion. Suppose that Alan has a salary of $80,000 and is used to paying tax on that. In addition, he earns $20,000 'on the side'. For example, he does a bit of consultancy, gets royalties from books he wrote or does odd jobs for friends and neighbors. Paying tax on this additional $20,000 seems to Alan like a loss. In this case, the natural reference point is the $1,068,000 that results from paying a $32,000 tax on his salary. The only way Alan can avoid a loss is to not report the $20,000 he earns on the

side. In this case, we can see in Figure 3.22 that we get a kink in the expected utility at $80,000. While it is not immediately apparent from Figure 3.22 given the flatness of the utility function, it turns out that Alan's utility is maximized when he reports around $50,000 of income. The more general point, however, is that he will not report more than the $80,000 that he is 'willing' to pay tax on.

3.8.3 Taxes and reference points

We have seen that reference-dependent expected utility can explain why Alan, or anyone else, would honestly report their income. But we have also seen that the reference point is crucial. If Alan is used to paying all his taxes then he maximizes his utility by paying taxes. If he is used to not paying taxes then he maximizes his utility by not paying taxes. To meaningfully apply the theory we, therefore, need to think carefully about what the reference point will be in particular contexts.

As we have already seen several times in the book so far, the reference point is likely to depend on framing. Suppose, for instance, that Alan works for a big employer who automatically takes tax out of his income. In this case the reference point is surely one of paying taxes. Alan essentially doesn't notice that the money going on taxes disappears from his pay. Or, relating back to the endowment effect, he never owns the money to lose it. Contrast this with the situation in which Alan is self-employed. Now all the money he earns passes through his hands, and it is up to him to pay his taxes. The framing is very different, and this may be reflected in a reference point of tax evasion.

Another side to the story is desire to pay taxes. If Alan thinks his tax money will be used wisely then that may increase the reference amount of income that he reports. Similarly, if Alan dislikes being threatened and harassed about taxes then, the more pressure the authorities put on him, the lower may be the reference amount of income that he reports. Alan may be more likely to honestly report his income if the government positively promotes the societal benefits of paying taxes than if it threatens him for not paying. Changing the frame can potentially do a lot more than changing the probability of audit.

A final framing effect I want to discuss is that of social influence. We shall look at social influence a lot more in Chapter 5, but it is worth briefly touching on the issue here. The basic idea is that a person's reference point may be influenced by what others are doing. So, if all Alan's friends and colleagues are evading taxes, this may lower the reference amount of tax that Alan reports. Similarly, if all his friends honestly report, then it may increase the reference amount of tax he reports. There are different reasons why we may observe such influence. One possibility is that Alan simply likes doing the same as others. Another possibility is that he learns from others the trade-offs of evading taxes. The outcome, however, is likely to be the same whatever the cause, in that we shall observe social reference points.

3.8.4 Tax evasion in the laboratory

As I have already said, it is not easy to measure tax evasion. It is not the kind of thing that people openly admit to doing! This makes it very difficult to know how a change in, say, the probability of audit affects the level of tax evasion. The experimental laboratory

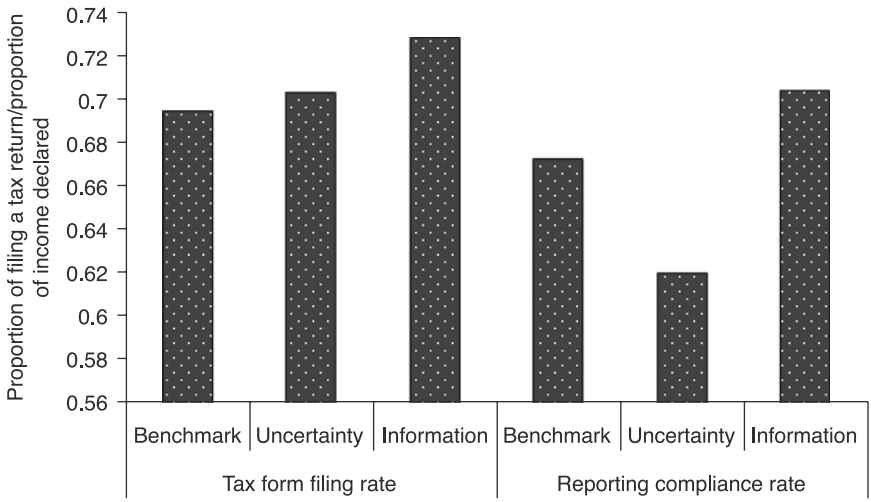

Figure 3.23 The proportion of subjects who filed a tax return (tax form filing rate) and the proportion of income that is declared (reporting compliance rate), by treatment. The filing rate and compliance rate are highest in the information treatment.
Source: Alm *et al.* (2010).

is, therefore, an invaluable place to learn more about tax evasion. In the laboratory we can directly measure evasion and much more easily see what effect things like the probability of audit have on reported income (see Research methods 3.9).

To illustrate what we can learn from tax evasion experiments, I will talk about a study by Alm and co-authors (2010). This is an interesting study for us to look at because it highlights another crucial aspect of tax evasion. We have taken it as given, so far, that Alan knows how much tax he is supposed to pay. The reality is almost always that a person does not know how much tax they are supposed to pay! For just about any country you choose to look at, the tax code will be impenetrable. Different income is taxed at different rates, there are deductibles, entitlements to benefits, tax credits, etc. This complexity can make filing a tax return into a nightmare.

In the study by Alm and co-authors, subjects had to pay tax on income earned during a simple task. To capture some of the complexity of tax reporting, a subject could claim a tax deduction (a reduction in taxable income) and a tax credit (extra money that declined with the level of income). If someone did not file a tax return then they paid no tax, but also missed out on claiming any tax credit. In a benchmark treatment, subjects were made fully aware of the rules for deduction and tax credit. In an uncertainty treatment, subjects were left guessing how much tax deduction and tax credit they could claim. Only if they were audited would they find out exactly how much tax they needed to pay. In an information treatment, subjects were left guessing how much income they should report but could press a button to find out exactly how much they should report.

One could argue that the uncertainty treatment most closely reflects the reality for many individuals filing their tax return. There is almost inevitably a little bit of

guesswork in knowing how much tax should be paid. This could result in a person paying more tax than they are supposed to, in order to avoid any chance of a fine, or paying less, because of annoyance with the tax system. The latter possibility seems more likely, and so there is an increasing recognition that tax authorities need to provide services to taxpayers. Their job is not only to audit and catch tax evaders but also to help compliant individuals more easily declare their income. This translates into good, user-friendly online systems, helpdesks that work, and so on.

Figure 3.23 summarizes the main results of the study. It allows us to compare across treatments the proportion of subjects who filed a tax return and the proportion of income that was declared. We can see that the filing rate and compliance rate were higher in the information treatment. They are not dramatically higher, but a small increase in reported income can add up to a large change in tax revenue when aggregated over millions of taxpayers.

This result may seem bizarre. The benchmark and information treatments are basically identical. All that differs is the frame. So, making the tax code more complex and then offering a solution to the problem increases compliance! The lesson to take away from this study is not, however, that we should be increasing the complexity of the tax code; tax codes are complex enough as it is. Instead, we learn that tax compliance can increase if the tax authorities provide a service to taxpayers. Moreover, this effect may be largely due to framing. Alan, for instance, is more likely to declare his income if he feels he gets something for his money. A nice, online interface that makes it effortless to file a tax return may convince Alan that he does get something for his money. A clunky system that doesn't work might leave Alan infuriated and unconvinced his taxes are worth paying.

This study on tax services is only one of many. Other studies have looked at how tax compliance varies with the probability of audit, size of fine, etc. The study nicely illustrates, however, the role that experiments can play in studying tax evasion. Improving tax authority services is not simple or cheap. Experiments such as this can inform on whether it is worth doing, and how best to do it. More generally, we have seen in this section that understanding tax compliance requires taking on board many of the things we have looked at in the last two chapters. Framing effects, endowment effects, reference points, weighting of probabilities are all key in the decision to evade taxes or not.

❖ RESEARCH METHODS 3.9 EXTERNAL VALIDITY, SOMETHING IS BETTER THAN NOTHING

Tax evasion experiments clearly raise issues of external validity. In a laboratory experiment we can have a student who has never filed a tax return in his life making decisions about $10 he earned in five minutes on a simple game. In the real world we can have a seasoned adult making decisions about $100,000 earned through a year of hard work, facing jail if he misreports. How can the one possibly tell us anything about the other?

A first response to this concern is to say that something is better than nothing. In order to inform policy-makers it is really important that we have some clear evidence to back up the theory. Improving tax authority services, for instance,

costs millions of dollars and so should not be done on a whim. If it is difficult, or impossible, to get evidence from the real world then we have to rely on laboratory studies. This is a strong argument. And it applies not only to tax evasion but to many areas of policy.

Difficulty in obtaining field data is not in itself, however, enough to justify laboratory experiments. We also need to have some confidence that behavior in the laboratory is something like behavior in the real world. Otherwise we are wasting our time, or, worse, misinforming policy-makers. This is why tax evasion experiments are one of the few occasions in experimental economics when an explicit frame is used. In the study by Alm and co-authors, subjects earned money, were asked to file a tax return and were given an interface that was not unlike filing a tax return.

That may not be enough, however, and checking external validity is impossible if we cannot accurately measure real tax evasion. There are no easy answers to such problems. This does not mean experiments are not useful, but it does mean we should not become overly reliant on them. We need good theory to complement experimental studies.

3.9 Legal settlements

Economics and law may seem like two very different subjects. There is a considerable amount of overlap between the two, however. Tax evasion illustrates one reason for this: many crimes, such as tax evasion, theft, smuggling, corruption and fraud, are motivated by financial gain. So, committing the crime is an economic gamble, and, as we have seen in the previous section, can be modeled as such. Other crimes, such as vandalism, minor assault, negligence, workplace discrimination, careless driving and drink driving, are not motivated by financial gain. The likely punishment is a financial one, however, in the form of a fine or of compensation to the victim. In this section we shall look at how behavioral economics can inform on the 'bargaining' process between the prosecution and defense.

To guide us through the issues I will consider the following example. Alan has been accused by one of his employees, Anna, of sex discrimination in the workplace. With probability p Alan will be found guilty, in which case he will have to pay compensation of \$5,000 to Anna. With probability $1 - p$ Alan will be found not guilty and have to pay no compensation. Suppose that both Alan and Anna start with wealth of \$100,000. They would like to reach an out-of-court settlement, and thereby avoid legal costs and emotional strain. But will they be able to agree?

Looked at from an individual point of view, Alan faces prospect (p,\$95,000; $1 - p$,\$100,000) and Anna faces prospect (p,\$105,000;$1 - p$,\$100,000). The fascinating thing, however, is the connection between the two prospects. Alan's loss is Anna's gain. In order to reach agreement they will, therefore, need to find a settlement that *both* prefer compared to their existing prospect. This provides an ideal opportunity to introduce a key feature of prospect theory that I have so far neglected: the fourfold pattern of risk attitudes.

3.9.1 Fourfold pattern of risk attitudes

Before we delve into some equations let me ask you what you would do. Here are some scenarios.

Scenario A: Imagine you are in the place of Alan. There is a 5 percent chance you will be found guilty and have to pay $5,000. Anna offers to settle out of court for a payment of $250. Would you agree to settle?

Scenario B: Now, imagine you are in the place of Anna. There is a 50 percent chance you will win the case and get $5,000. Alan offers you $2,500 to settle out of court; would you agree?

Scenario C: Still in the place of Anna. There is a 5 percent chance you will win the case and get $5,000. Alan offers you $250 to settle out of court; would you agree?

Scenario D. You are back in the place of Alan. There is a 50 percent chance you will lose the case and have to pay $5,000. Anna offers to settle out of court for a payment of $2,500; would you agree?

Any set of answers is possible. But, if you are like most people, your answers will have fitted the fourfold pattern of risk attitudes summarized in Table 3.23. You will have settled in scenarios A and B and not settled in scenarios C and D. The predictions in the case of B and D should need little explanation; in section 3.1.3 we have already discussed that people are typically risk-averse over gains and risk-loving over losses. The more interesting cases are A and C. These cases are characterized by the probability of a guilty verdict being relatively small. Why would this matter?

We have seen that people overweight small probabilities. So, when the probability of a guilty verdict is only 5 percent, Alan and Anna will likely overweight its probability. This overweighting of probability acts as a counterweight to the natural tendency to be risk-averse over gains and risk-loving over losses. Anna, for instance, will overweight the probability of winning the $5,000 and so be reluctant to settle out of court. Alan, by contrast, will overweight the probability of losing the $5,000 and be happy to settle out of court. Low probability events can, thus, cause a switching of risk attitudes, as summarized in Table 3.23.

To formalize the idea, let us work through a model of reference-dependent utility. The overweighting of probabilities is crucial, and so I will start with this. Following the logic of section 3.3, we can work out the decision weight Anna and Alan put on a guilty verdict. Anna stands to gain from going to court, and the best outcome is that she wins the case. So, the decision weight on winning the case will be

Table 3.23 The fourfold pattern of risk attitudes in litigation.

	Low probability	**Medium to high probability**
Gains	**Risk-loving** (scenario C) Will go to court unless offered a generous settlement	**Risk-averse** (scenario B) Happy to settle out of court
Losses	**Risk-averse** (scenario A) Happy to settle out of court	**Risk-loving** (scenario D) Will go to court unless offered a generous settlement

$$w_2 = \frac{p^{\gamma}}{(p^{\gamma} + (1-p)^{\gamma})^{\frac{1}{\gamma}}}.$$

The decision weight on not winning is $w_1 = 1 - w_2$. Table 3.24 illustrates how the decision weight varies with p when $\gamma = 0.61$. We see that a 5 percent chance of winning the case is converted into a 13 percent decision weight. Alan stands to lose from going to court, and the worst outcome is that he is found guilty. So, the decision weight on losing the case will be

$$w_1 = \frac{p^{\delta}}{(p^{\delta} + (1-p)^{\delta})^{\frac{1}{\delta}}}.$$

The decision weight on winning the case is $w_2 = 1 - w_1$. Setting $\delta = 0.69$, we get the decision weights given in Table 3.24. Notice how Alan's decision weights differ from those of Anna.

We can look next at expected utility. I'll assume that Alan and Anna have the utility function

$$u(x) = \begin{cases} 10\sqrt{x} + (x-r)^{0.88} & \text{if } x \geq r \\ 10\sqrt{x} - 2.25(r-x)^{0.88} & \text{if } x < r, \end{cases}$$

where x is final wealth and r is the reference point. The natural reference point here is a not guilty verdict, meaning r = \$100,000. Observe that there is nothing dodgy about the utility function. Alan and Anna have diminishing marginal utility of wealth, are risk-averse for gains and are risk-loving for losses. So, any switching of risk attitudes cannot be caused by the payoff function.

Table 3.24 details the expected utility of Anna and Alan if they go to court. It also details their utility from settling out of court with a payment equivalent to expected value. We can see that Anna's utility is higher from going to court if the probability of a guilty verdict is low. Alan's utility is higher from going to court if the probability of a guilty verdict is high. We get the predicted fourfold pattern of risk attitudes.

Table 3.24 Outcomes of going to court or settling out of court.

	Probability of Anna winning the case			
	0.05	**0.10**	**0.25**	**0.50**
Expected value	\$250	\$500	\$1,250	\$2,500
Anna's decision weight on guilty verdict	0.13	0.19	0.29	0.42
Alan's decision weight on guilty verdict	0.11	0.17	0.29	0.45
Anna's expected utility if go to court	3410	3512	3708	3951
Anna's expected utility if fair settlement	3295	3407	3717	4179
Alan's expected utility if go to court	2702	2460	1950	1288
Alan's expected utility if fair settlement	2868	2620	1947	922

3.9.2 Frivolous litigation

The switching of risk attitudes due to an overweighting of small probabilities means that we can expect very different outcomes in a case depending on the probability of a guilty verdict. If the probability of a guilty verdict is moderate to high then Anna's risk aversion would mean she is more willing to settle out of court than Alan. So, the defense has more bargaining power. Any settlement is likely to be for less than expected value. For example, if $p = 0.5$, the expected value of going to court is \$2,500, but Anna may be willing to settle out of court for, say, \$2,000. Of course, the amount for which Anna is willing to settle may be too much for Alan, and so they take their chances in court. The point, though, is that, if they reach a settlement, it will be favorable to Alan.

If the probability of a guilty verdict is low then Anna and Alan's overweighting of the chance of a guilty verdict means that Anna is less willing to settle out of court than Alan. The prosecution has more bargaining power. Any settlement is likely to be for more than expected value. For example, if $p = 0.05$, the expected value of going to court is \$250, but Alan may be willing to settle out of court for, say, \$400. Again, Anna and Alan may still take their chances in court. But this time, if they reach a settlement, it will be favorable to Anna.

A study by Fobian and Christensen-Szalanski (1993) illustrates this fourfold pattern of risk attitudes in action. They asked subjects to consider genuine medical liability cases. Subjects were also given advice 'from their lawyer' on the probability of winning the case and the amount of money they would win. The probability was either 0.1 or 0.9 and the amount was either \$100,000 or \$1 million. Figure 3.24 details the average amount of money that the prosecution and defense were willing to accept and pay to settle out of court. Focus for now on the case when $p = 0.1$ and there is no ambiguity. We see that the defense was willing to pay just over

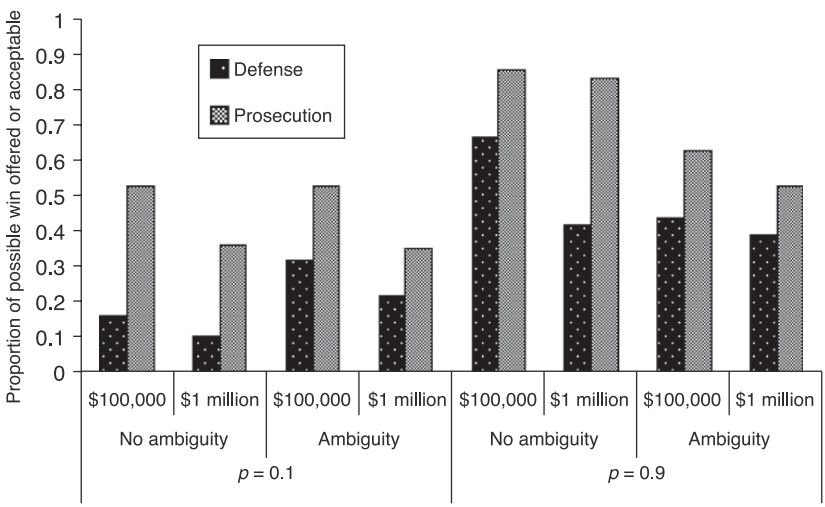

Figure 3.24 The amount of money the defense was willing to offer and the prosecution was willing to accept for settling out of court (as a proportion of the claim amount).

Source: Fobian and Christensen-Szalanski (1993).

Table 3.25 A summary of outcomes of litigation.

Probability of a guilty verdict	Best bargaining position	If settled out of court	Implication
Moderate to high	Defense	Less than expected value	In genuine cases the victim gets less than a fair amount
Low	Prosecution	More than expected value	Frivolous litigation is encouraged

expected value; they were risk-averse. The prosecution, by contrast, demanded way over expected value; they were risk-loving. Consider now the case when $p = 0.9$. This time the defense was willing to pay a lot less than expected value; they were risk-loving. The prosecution was willing to settle for slightly less than expected value; they were risk-averse.

These results perfectly illustrate the fourfold pattern of risk attitudes. They also illustrate the difficulty for parties to reach agreement. An out-of-court settlement is possible only if the defense offers more than the prosecution will accept. We can see in Figure 3.24 that the two parties are often a long way from agreement! So, these cases are going to court. In the ambiguity treatments the lawyer (while still predicting $p = 0.1$ or 0.9) expressed a lack of confidence in his belief about the probability of winning because of uncertainties with the case. This brought the claims a bit closer together, which is important, because ambiguity is almost inevitable in legal cases. Is it good news, however, that agreement is more likely?

Table 3.25 summarizes what we have seen so far. The fourfold pattern of risk attitudes lets us make sense of out-of-court settlements. A more normative question is whether we should use this knowledge to advocate reform of the legal system. You might argue that if Anna and Alan reach an out-of-court settlement then they do so voluntarily, and so we should not intervene. Things are not necessarily so simple, however. A genuine case is likely to fit the category of moderate to high probability of a guilty verdict, and we see that the prosecution will settle for less than expected value. So, people who genuinely are discriminated against will walk away short-changed. Frivolous, non-genuine cases are likely to fit the category of low probability and we see that the prosecution gets more than expected value. So, people who make up claims walk away with money! In both cases one can motivate intervention, and it is clear from Figure 3.24 that ambiguity may make matters worse.

The fourfold pattern of risk attitudes not only gives us a means to understand out-of-court settlements, therefore, it also gives us a mandate for policy change. Frivolous litigation, for instance, is particularly costly for the court system. If we know why it happens then we are in a better position to do something about it. In Chapter 11 we shall look in detail at how framing can be used as a policy tool to change behavior. In this instance, a different frame may avoid the overweighting of small probabilities, or turn perceived losses into gains. This almost certainly means fighting against a 'compensation culture' that encourages a perception of easy money to be made from frivolous claims. In Chapter 5 we shall see the difficulties of overcoming such biased beliefs.

3.10 Further reading

The survey paper by Starmer (2000) is a great place to start. Much of the material in this chapter owes its origins to the work of Robert Sugden and co-authors. Some of the relevant papers are Loomes and Sugden (1982, 1983, 1986), Loomes, Starmer and Sugden (1991) and Schmidt, Starmer and Sugden (2008). Most of the rest of the material owes its origins to Kahneman and Tversky, including their 1979 and 1992 papers I mentioned earlier (much of the suggested reading for Chapter 2 is relevant here as well). I also want to mention Rabin (2000), who was the first to appreciate the issues in section 3.1.2. On the issue of measuring risk aversion, the paper by Holt and Laury (2002) is a classic reference. For more on prospect theory, see the review article by Barberis (2013b). A lot of work has looked at estimating the parameters of prospect theory; see, for instance, Harrison and Rutström (2009). One issue emphasized in recent research is the importance of taking into account the weighting of probabilities, rather than just focusing on loss aversion; see Fehr-Duda and Epper (2012) for a review of the literature and Barseghyan and co-authors (2013) for another example looking at insurance. See Barberis (2013a) for a discussion on the closely related topic of tail events, which are events that happen with very small probability. Recall that the fourfold pattern of risk attitudes predicts risk-loving behavior for events that happen with small probability; Ebert and Strack (2015) show that this can have extreme implications for gambling and financial trading. For more on tax evasion, see the review articles by Alm (2012) and Hashimzade, Myles and Tran-Nam (2013).

3.11 Online materials

On the companion website you will find more details on the example of Alan buying insurance. You will also have access to flashcards, quizzes and answers to the review questions.

3.12 Review questions

3.1 Explain why the indifference curves in a probability triangle diagram are straight lines if preferences satisfy expected utility theory. [Hint: we can write the expected utility of a prospect as $U(x) = p^1u(x^1) + (1 - p^2 - p^3)\ u\ (x^2) + p^3u(x^3)$.]

3.2 Show why it is inconsistent with expected utility that most people choose prospect F over E, and prospect G over H, when the prospects are as in Table 3.5.

3.3 Using the model of disappointment with $\theta = 0.002$, consider the following three prospects: A = (0.5, \$2,400; 0.5, \$0), B = (0.7, \$2,400; 0.3, \$0) and C = (0.3, \$2,400; 0.7, \$0). Work out the utility of each prospect and comment on the result.

3.4 Using prospect theory, say whether a person would prefer prospect I or prospects J to N from Table 3.6.

3.5 In 1963 Paul Samuelson wrote about a colleague who said that he would turn down the prospect (0.5, –\$100; 0.5, \$200) but would accept 100 such prospects. Suppose that his utility function is $u(x) = x - w$ if $x \geq w$ and $u(x) = -2.5(w - x)$ if $w > x$, where w is his wealth. Show why he turned down the prospect.

Now, imagine two prospects will be done in turn and he will adjust his wealth level after each prospect. Show that he should turn down the two prospects. Finally, imagine that he adjusts his wealth level only after seeing the outcome of both prospects. Should he take on the two prospects? Should he take on 100 prospects?

3.6 What do you think is the relevant reference point of a prospect? How might the certainty effect be related to reference dependence?

3.7 One set of prospects considered by Loomes, Starmer and Sugden (1991) was the following: A = (0, \$10; 0.6, \$3), B = (0.7, \$7.50; 0.3, \$1) and C = (1, \$5). What would you choose between A and B, B and C, and A and C? How can regret theory help us explain choices in this case?

3.8 What would happen to the equity premium if investors were less loss-averse or the evaluation period was longer? How often should you evaluate your investments?

3.9 Is regret theory the same as a model of reference-dependent utility with a stochastic reference point?

3.10 Are preference reversals caused by intransitive preferences or procedural invariance?

3.11 I argued in section 3.7 that, applying the standard formulation of prospect theory, with $\lambda = 2.25$, $\beta = 0.88$ and $\gamma = 0.61$, a person would pay \$110 for an extra deductible of \$500 if the claim rate was 4 percent. How much would a person be willing to pay to reduce the deductible by \$165 if the claim rate is 25 percent? Comment on the numbers in Table 3.20.

3.12 Is expected utility theory of any practical relevance?

3.13 Carefully explain how we can estimate a person's risk aversion by their choice of insurance deductible.

3.14 Why may a person's decision to evade taxes be influenced by what his friends and colleagues are doing?

3.15 Carefully explain why the overweighting of small probabilities can lead to tax evasion and frivolous litigation.

3.16 In many court cases there is no clear-cut defense and prosecution (e.g. divorce disputes). What does the fourfold pattern of risk attitudes predict in such cases?

3.17 If loss aversion can lead to risk aversion, how can we tell the difference between loss aversion and risk aversion?

3.18 Is prospect theory an example of over-fitting?

Chapter 4

Choosing when to act

> The only way to get rid of a temptation is to yield to it.
> I can resist everything except temptation.
>
> Oscar Wilde

Time is important in most economic decisions because the choices we make will have future consequences. Should a person do the shopping today, or tomorrow? Should she buy a new TV, or save for retirement? Should she look for a new job? Should she go to university and get a qualification? Should she eat healthily and join a gym? The answer to all these questions involves weighing future benefits and costs with present benefits and costs. We need to think about how people do that, and that's the focus of this chapter.

4.1 Exponential discounting

To illustrate some of the issues, consider Maria planning when to do her homework. The homework was set on Friday and must be handed in on Monday morning. She can do it on Friday, Saturday, Sunday or Monday. Table 4.1 gives her day-by-day utility, depending on when she does the homework. For example, if she does it on Friday she pays a cost on Friday, gets to enjoy Saturday and Sunday, and on Monday gets to know how well she did. Doing the homework on a weekend is more costly but also means she does it better. When should Maria do the homework?

Clearly, she has to weigh up the benefits and costs of doing the homework on each possible day. Before we see how she might do this, it is worth introducing some general notation for thinking about choice over time. We can think of time as running from period one to period T. In the example, period one is Friday and period $T = 4$ is Monday. Maria needs to plan what she will do in each period between now and T, and, given a plan, she can work out what her utility will be in each period. We can think of the utility in a period as determined by the same utility function that we came across in Chapters 2 and 3, and I will use u_t to denote the utility in period t. (Just to clarify, in some parts of the book I have used u_i to denote the utility of person i, but not in this chapter.) In the example, a plan is when to do the homework, and we see, for example, that the utility in period three is $u_3 = 10$ if Maria plans to do the homework on Friday.

Table 4.1 The day-by-day utility for Maria of doing her homework on different days.

Plan	Utility			
	On Friday	**On Saturday**	**On Sunday**	**On Monday**
Do it Friday	–5	5	10	4
Do it Saturday	0	–5	10	10
Do it Sunday	0	5	–5	10
Do it Monday	0	5	10	–5

Table 4.2 The inter-temporal utility of each plan for three different discount factors.

Plan	Inter-temporal utility		
	$\delta = 1$	$\delta = 0.9$	$\delta = 0.7$
Do it Friday	14	10.5	4.7
Do it Saturday	15	10.9	4.8
Do it Sunday	10	7.7	4.5
Do it Monday	10	9.0	6.7

To try and model Maria's choice, we can use an **inter-temporal utility function** that combines utility from each period to one measure of overall utility. The simplest way to do this would be to just add together the utility from each time period. Generally, however, she might want to **discount** – that is, give less weight to future utility. This suggests a **utility function with exponential discounting**. The inter-temporal utility of getting u_1 in period 1, u_2 in period 2, and so on is then

$$u^T(u_1, u_2, ..., u_T) = u_1 + \delta u_2 + \delta^2 u_3 + ... \, , + \delta^{T-1} u_T = \sum_{t=1}^{T} \delta^{t-1} u_t, \tag{4.1}$$

where δ is a number called the **discount factor**. Inter-temporal utility is, therefore, a simple weighted sum of the utility in each period. If $\delta < 1$ then less weight is given to the utility in a period the further away that period is. So, the smaller is δ, the more future utility is discounted.

If we know the utility in each period of each plan then we can work out the inter-temporal utility of each plan. A prediction would be that Maria should choose the plan with highest inter-temporal utility. Table 4.2 illustrates what happens when we do this for three different discount factors. We can see that, the higher is δ, the more impatient Maria becomes. If $\delta = 1$, she is willing to sacrifice an enjoyable Saturday in order to get a higher mark on Monday. If $\delta = 0.7$, she is more impatient to enjoy herself and does not do the homework until Monday.

Exponential discounting is a very simple way to model choice over time. All we need to know is the discount factor, and then we can easily predict choice. For this

reason, exponential discounting is by far the most common way used in economics to model choice over time. What we need to do is ask whether it is sophisticated enough to capture all the things we observe. Before we do that, it is important that we fully understand what the model and, in particular, the discount factor imply, and so there are a couple of things I want to mention.

The first thing I want you to think about is the units of measurement: a period and utility. Maria's choice is going to have a stream of future consequences, and we somehow need to measure those consequences. What I have done is split the future into days and say what her utility will be on each day, but there is inevitably something arbitrary about how we split things up. For instance, I could have split the future into hours, and said what her utility will be in each hour; or into seconds; etc. One sensible approach is to split things up as Maria perceives them, and so, if she thinks in terms of what will happen each day, that is how we should split things. How Maria perceives things will likely depend on the context, however. This is already enough to suggest important context effects.

The second thing I want to clarify is the distinction between discount factor and discount rate. If the **discount factor** is 0.8 then '$10 worth of utility' next period is equivalent to $8 today. More generally, $10 next period is equivalent to $\delta$$10 today. The smaller the discount factor, the more impatient Maria is. Given a discount factor δ, we can work out a **discount rate** ρ using the relation

$$\delta = \frac{1}{1+\rho} \text{ or } \rho = \frac{1-\delta}{\delta}.$$

For example, if the discount factor is 0.8 then the discount rate is 0.25. In interpretation, a discount rate of 0.25 means Maria would require an interest rate of 25 percent to delay until next period. So, instead of $8 today she would want 1.25 × $8 = $10 next period. The higher the discount rate, the more impatient she is. It does not matter whether we use discount rate or discount factor, and so to be consistent I will use discount factor throughout this book. If you follow up the further reading, however, expect many to use the term 'discount rate'. [Extra] If you are wondering why the name 'exponential discounting': in continuous time, equation (4.1) becomes

$$u^T = \int_0^T e^{-\rho t} u_t,$$

where ρ is the discount rate.

4.1.1 The discount factor

To better understand exponential discounting, we need to see what values for the discount factor seem most appropriate. In the experimental lab the discount factor can be estimated by giving subjects questions of the form 'Would you prefer $100 today or $150 in a year's time?'. If they answer '$100 today' then the discount factor is smaller than 0.66 (but see Review question 4.1!). I will look next at a study by Benzion, Rapoport and Yagil (1989), in which subjects were asked questions a bit

like this. More precisely, subjects were asked questions of the four basic types given below. As you look through these questions, please think about what your answer would be.

- **Postpone receipt:** You have just earned $200 but have the possibility to delay receiving it by one year. How much money would you need to get after a year in order to want to delay payment?
- **Postpone payment:** You need to pay back a debt of $200 but have the possibility to delay payment by one year. How much money would you be willing to pay back after a year if payment is delayed?
- **Expedite receipt:** You will get $200 in one year but have the chance to receive the money immediately. How much money would you accept now rather than have to wait a year?
- **Expedite payment:** You need to pay back a debt of $200 in one year but can pay it now. How much would you be willing to pay now rather than pay off the debt after one year?

In specific questions the time period was changed from six months to four years and the amount of money from $40 to $5,000. Figure 4.1 summarizes the implied annual discount factor that we get from subjects' responses to each question. There is a lot

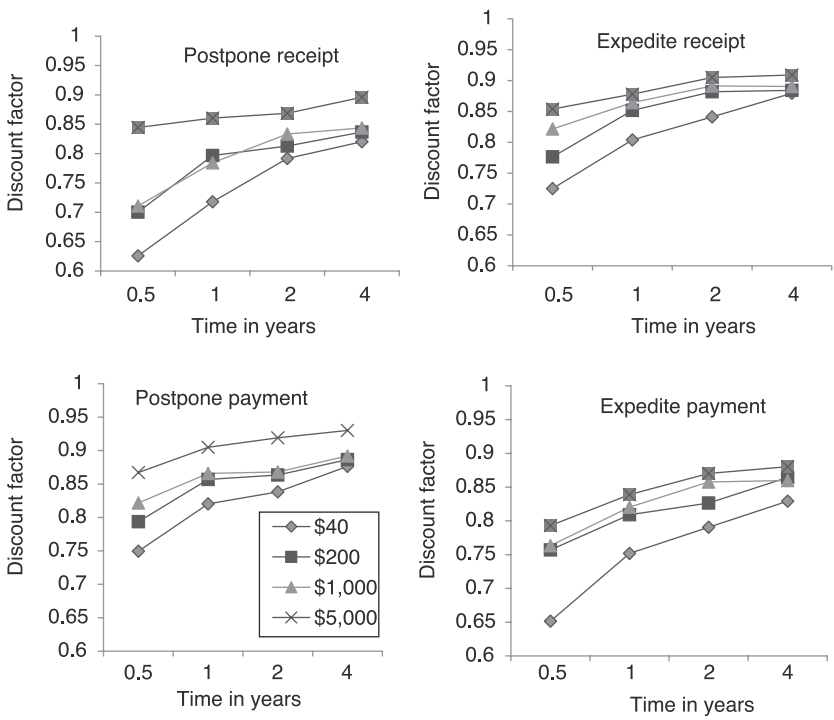

Figure 4.1 Estimated discount factors for four types of choices.
Source: Benzion, Rapoport and Yagil (1989).

to look at in Figure 4.1. One thing stands out immediately, however, and that's that the discount factor appears to depend a lot on context. We can see this in how the discount factor varies a lot depending on the length of time, amount of money, payment versus receipt, and expedite versus postpone.

In looking in a little more detail at Figure 4.1 I want to point out five things. First, notice that the discount factor is relatively low, around 0.8 to 0.9, suggesting that the subjects were relatively impatient. Next, notice that the discount factor is higher, the longer it is necessary to wait, suggesting **short-term impatience**. For example, subjects wanted on average almost as much compensation to postpone receipt by six months as to postpone by four years. A third thing to note is that, the larger the sum of money, the larger the estimated discount factor, suggesting people are more patient for larger amounts. This is called the **absolute magnitude effect**.

Next I want to compare payment versus receipt and loss versus gain. If you look from top to bottom in Figure 4.1 then you can see a **gain–loss asymmetry**, whereby the estimated discount factor is smaller for gains than for losses. For example, the discount factor is higher in the case of postponing payment (which would require a loss of money) than for postponing receipt (which would involve a gain of money). Similarly, the discount factor is higher in the case of expediting receipt (which would require a loss of money) than for expediting payment (which would involve a gain of money). This is consistent with loss aversion, because it suggests that subjects were reluctant to lose money in order to postpone or expedite.

The final thing I want you to do is compare postponing versus expediting. If you look from left to right in Figure 4.1 then you can see a **delay–speed-up asymmetry**, whereby the estimated discount factor is higher to postpone than to expedite payment, and higher to expedite than postpone receipt. This is because subjects were willing to pay relatively less money to postpone a payment than they demanded to expedite payment. Similarly, they demanded relatively more to postpone receipt than they were willing to pay to expedite receipt.

To really work well, exponential discounting requires that the discount factor should not change depending on context. In other words, in Figure 4.1 we should see overlaying horizontal lines all at the same discount factor. Clearly, we do not, and this is not good news for anyone wanting to use exponential discounting. By now, however, you should be getting used to the idea that context matters. Indeed, the gain–loss asymmetry is consistent with loss aversion, and the delay–speed-up asymmetry is suggestive of mental accounting.

The existence of such large context effects does not mean that exponential discounting is necessarily a bad way to model choice. What it does mean is that we have to be very careful in thinking about what the appropriate discount factor should be for a particular situation. There is no such thing as 'Maria's discount factor'. Instead, Maria will likely have different discount factors for different things. For example, she might be impatient to repay $10 she borrowed from one friend, but more patient to get back a CD she lent to another friend. Before we explore the implications of this there is one more important context effect that we have yet to consider.

❖ RESEARCH METHODS 4.1 EMPIRICAL VERSUS LAB EVIDENCE

Inter-temporal choice is one area where the experimental lab does seem inadequate to answer many of the questions we are interested in, because we want to know how people trade off money over relatively large time periods. How can we create delay over relatively large time periods in the experimental lab? One option is to use hypothetical scenarios, as in the study by Benzion, Rapoport and Yagil. This leads to the objection, however, that subjects are not making real choices. If we do use real choices then we cannot realistically delay payment beyond a few months, and certainly not years. It is also questionable whether subjects would find delayed payment credible.

This means that empirical studies in which we observe people making choices with long-term trade-offs are very useful. We can illustrate with a study by Warner and Pleeter (2001). In 1992 the US Department of Defense needed to reduce the size of the US military. It offered selected servicemen two possible packages for leaving the military: (1) a lump sum payment of around $25,000 for privates and NCOs, and $50,000 for officers; and (2) an annuity that would pay some fraction of current basic pay for twice the length the person had worked for the military.

Serviceman had, therefore, the choice between a one-off payment and deferred payments. For a particular serviceman it is possible to work out the discount factor such that the serviceman should be indifferent between the two options. Formally, we find the δ^* such that

$$LS(1 - T_{LS}) = \sum_{t=1}^{2YOS} \delta^{*t-1} A(1 - T_A),$$

where LS is the lump sum payment, A the annuity payment, YOS the years of service and T the respective tax rates. A serviceman should choose the lump sum payment if and only if they are more impatient than is implied by δ^* – i.e. only if their discount factor is $\delta \leftarrow \delta^*$. With data on what a person chose and by calculating δ^*, we can therefore estimate a person's discount factor. This is what Warner and Pleeter did. Estimates of the mean discount factor are around 0.85.

The estimate of 0.85 is consistent with what we see in Figure 4.1, which provides reassurance that the numbers obtained from the study by Benzion, Rapoport and Yagil are not unrealistic. I think this nicely illustrates how empirical studies can complement experimental studies. Experimental studies allow us to investigate relatively easily the comparative statics of changing things such as the amount of money and time frame. This is simply not possible with empirical studies, because events such as the US military downsizing are rare and give only a point estimate of the discount factor. Empirical studies do, though, serve as a useful check that results obtained in the lab are meaningful.

4.1.2 The utility of sequences

Imagine now that, instead of asking someone whether they would prefer $100 today or $150 in a year's time, we ask them whether they would prefer $100 today and $150 in a year's time, or $150 today and $100 in a year's time. This gets us thinking how people interpret **sequences** of events. To illustrate how people tend to respond to such questions, I will look at how subjects answered two questions asked in a study by Loewenstein and Prelec (1993).

In the first question people were asked to imagine two planned trips to a city they once lived in but do not plan to visit again after these two trips. During one trip they need to visit an aunt they do not like, and in the other visit friends they do like. Those asked were given three scenarios for when the trips might be (e.g. one this weekend and the other in 26 weeks) and asked to say who they would rather visit first. The results are in Figure 4.2.

One thing to pick out from Figure 4.2 is a **preference for an improving sequence** in which subjects visit the aunt first and then friends later. This might seem natural, but is the opposite of what exponential discounting would predict. If someone uses exponential discounting, with a discount factor less than one, then they would maximize their inter-temporal utility by visiting the friends first because they are impatient for higher utility. This impatience does become apparent when the gap between the trips is made bigger, because then more subjects wanted to visit friends first.

In the second question I want to look at people who were asked to choose when over the next three weeks they would like to eat out at restaurants called Fancy French and Fancy Lobster. The results are in Table 4.3. When asked to choose

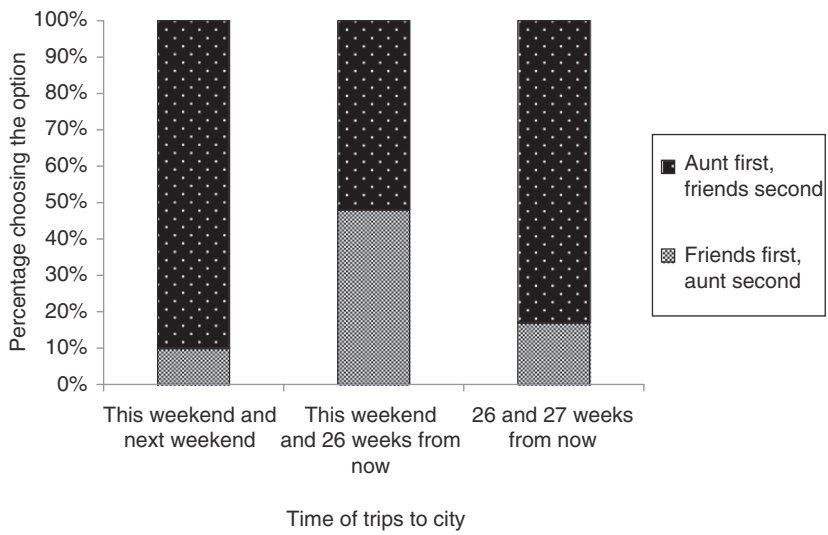

Figure 4.2 Choice of when to visit an aunt and friends. If the two trips were close together, most subjects chose to visit the aunt first. If the two trips were far apart, more subjects chose to visit friends first.
Source: Loewenstein and Prelec (1993).

Table 4.3 Choice of when to eat at a restaurant. When asked to choose between options A and B, most subjects prefer to delay eating at Fancy French restaurant. When asked to choose between options C and D, most subjects prefer not to delay eating at Fancy French restaurant.

Option	This weekend	Next weekend	Two weekends away	Choices (%)
A	Fancy French	Eat at home	Eat at home	16
B	Eat at home	Fancy French	Eat at home	84
C	Fancy French	Eat at home	Fancy Lobster	54
D	Eat at home	Fancy French	Fancy Lobster	46

Source: Loewenstein and Prelec (1993).

Table 4.4 Maria deciding when to visit an aunt and friends. Events in weeks 1 to 4 and weeks 24 to 28 are bundled together as a sequence. Events in weeks 24 to 28 are seen as disjoint from those in weeks 1 to 4 and discounted with a factor of 0.8.

Events		Inter-temporal utility
Weeks 1–4	Weeks 24–28	
Visit aunt then friends		10
Visit friends then aunt		5
Visit aunt		0
Visit friends		15
Visit aunt	Visit friends	12
Visit friends	Visit aunt	15
	Visit aunt then friends	8
	Visit friends then aunt	4

between options A and B, we see the same preference for an improving sequence that we saw for the previous question. When asked to choose between options C and D, however, many subjects preferred to spread the good events over the sequence. Again, this might seem natural, but is not consistent with exponential discounting. This switch in choices should not happen, because the utility in period three should not affect the optimal choice in periods one and two. Clearly, it may do.

Overall, therefore, we see a preference for improving sequences coupled with a preference for spreading events evenly throughout the sequence. None of this seems consistent with exponential discounting. The only way to make it consistent is to distinguish sequences from disjoint events. This can work. For instance, we could say that the relevant time period for deciding when to visit the aunt and friends is one month. If, therefore, the two trips are consecutive weekends then they should be bundled together and seen as a sequence, while if they are 26 weeks away they should be seen as disjoint events and discounted. Table 4.4 illustrates how this might work to explain what we see in Figure 4.2.

This approach works perfectly if we can distinguish a length of time and say that events happening over a shorter period of time are a sequence and everything happening over a larger period of time are separate events. In general, however, this

is not going to be easy, because we need to know how the person perceives things. Does Maria perceive the events before she hands in her homework as separate, or a sequence? Will she be thinking 'Today I want to play sport so other things can wait till tomorrow' or 'If I do my homework today then tomorrow I can go out and play sport'? It will likely affect her choice, but it is not obvious how she might think. There is certainly no simple rule to say when something is a sequence or not. It is more likely to depend on the context.

Clearly, context will matter in thinking about and modeling inter-temporal choice. What I want to do now is question what the implications of this may be, and whether we need a different model from that of exponential discounting in order to capture it.

4.2 Hyperbolic discounting

We clearly see in Figure 4.1 that the discount factor is larger for longer time intervals. This means that people are impatient over the short term but more patient over the long term. To give another example, in a study by Thaler (1981), subjects were asked the amount of money they would require in one month, one year or ten years to make them indifferent to receiving $15 now. The average responses of $20, $50 and $100 seem entirely sensible, but imply an annual discount factor of 0.22, 0.45 and 0.84, respectively. Relatively speaking, therefore, subjects were asking a lot to wait one month but not very much to wait ten years. Such decreasing impatience is called **hyperbolic discounting**.

One way to capture hyperbolic discounting is to modify the model of exponential discounting and allow for different discount factors in different periods. In equation (4.1) we assume exponential discounting where payment in t time periods from now is discounted by an amount

$$D(t) = \delta^{t-1}.$$

One of many possible alternatives is to assume

$$D(t) = \frac{1}{1 - \alpha(t-1)}, \tag{4.2}$$

where α is a parameter that can capture changes in the discount factor over time. I will call this a **model of hyperbolic discounting**. In either case, equation (4.1) is generalized to

$$U^T(x_1, x_2, \ ..., x_T) = \sum_{t=1}^{T} D(t)x_t. \tag{4.3}$$

To illustrate, Figure 4.3 plots $D(t)$ and the annualized discount factor assuming exponential discounting with factor $\delta = 0.85$ and hyperbolic discounting with $\alpha = 0.25$. We see that hyperbolic discounting does give a higher discount factor for longer periods and looks a bit more like what we saw in Figure 4.1.

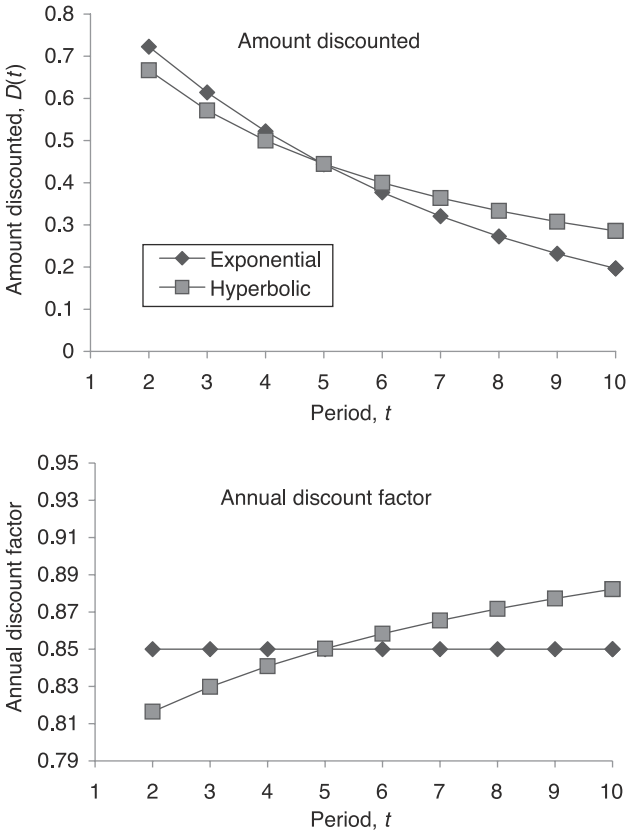

Figure 4.3 Hyperbolic discounting compared to exponential discounting. With hyperbolic discounting, the discount factor is lower for shorter time periods.

We are, therefore, able to accommodate a changing discount factor without too much change to a standard exponential discounting model. But are we really capturing decreasing impatience? To explain why we might not be, consider the following set of choices.

Do you want $100 today or $110 tomorrow?
Do you want $100 in 30 days or $110 in 31 days?

When asked such questions, many people choose the $100 today and $110 in 31 days. This can be consistent with hyperbolic discounting. The potential problem comes if we ask the same question after 30 days and get the same answer. That is, suppose Maria says she would prefer $110 in 31 days to $100 in 30 days, but then after 30 days says that she prefers $100 today to $110 tomorrow. This is not consistent with the model of hyperbolic discounting introduced above, because the model requires choices to be consistent over time. Maria should not change her mind. If she says that she prefers $110 in 31 days' time to $100 today, then after 30 days she should still prefer $110

tomorrow to $100 today. If people do change their mind then we are going to need a different model of hyperbolic discounting.

[Extra] To illustrate this with some numbers, let me first explain why it can be consistent with a model of hyperbolic discounting to choose the $100 today and $110 in 31 days' time. Suppose the time period is a day and $D(2) = 0.9$, $D(30) = 0.85$ and $D(31) = 0.84$. Then $110 tomorrow is worth $99 today, $100 in 30 days is worth $85 today and $110 in 31 days is worth $92.4 today, so it makes sense to choose $100 today and $110 in 31 days' time. What happens after 30 days? Given that $D(30) = 0.85$ and $D(31) = 0.84$, we should, after 30 days, calculate a revised discount factor of $D(2) = 0.99$. With this value of $D(2)$ it is optimal to choose the $110 tomorrow rather than $100 today. Note, however, that the discount factor has to change a lot over the 30 days, from 0.9 to 0.84 to 0.99. It would seem a remarkable coincidence that we asked Maria her choices just at the point in her life when her discount factor is so turbulent! There is almost certainly a more plausible, alternative explanation.

4.2.1 Quasi-hyperbolic discounting

With exponential and hyperbolic discounting, a time period should be interpreted as a specific date. For instance, in the homework example we thought of period one as Friday and period four as Monday. If today is Friday then the discount factors $D(t)$ are specific to that day, and on Saturday we need to update them. For example, if $D(2) = 0.9$ and $D(3) = 0.9$ on Friday then Maria does not discount between Saturday and Sunday, and on Saturday we should get $D(2) = 1$. Consequently, in a model of hyperbolic discounting, decreasing impatience means that a person gets less impatient as they get older, even if only by a few days.

An alternative interpretation of decreasing impatience is that a person is more impatient for short-term gains relative to long-term gains, and this has nothing to do with age. We can capture this by interpreting $D(2) = 0.9$ and $D(3) = 0.9$ as constant over time and showing how much future amounts are discounted relative to today. So, if today is Friday then $D(2) = 0.9$, and if today is Saturday then $D(2) = 0.9$. In this interpretation, Maria chooses the $100 today because she is always impatient for a possible immediate gain. The distinction is illustrated in Figure 4.4.

If it makes more sense to think of the discount factor as relative to today rather than calendar time, we say that there are **present-biased preferences**. In this case, Maria does not postpone her homework today, or choose the $100 today, because today is Friday, 1 May, and on Friday, 1 May, she is impatient; she is always impatient for immediate gains.

To model present-biased preferences we can reuse equation (4.2), but now interpret t as how many time periods from today rather than a specific point in time. But we still need to capture decreasing impatience. A simple way to do this is to use a model of **quasi-hyperbolic discounting** or (β, δ) **preferences**, where

$$U^T(x_1, x_2, ..., x_T) = x_1 + \beta \sum_{t=2}^{T} \delta^{t-1} x_t \tag{4.4}$$

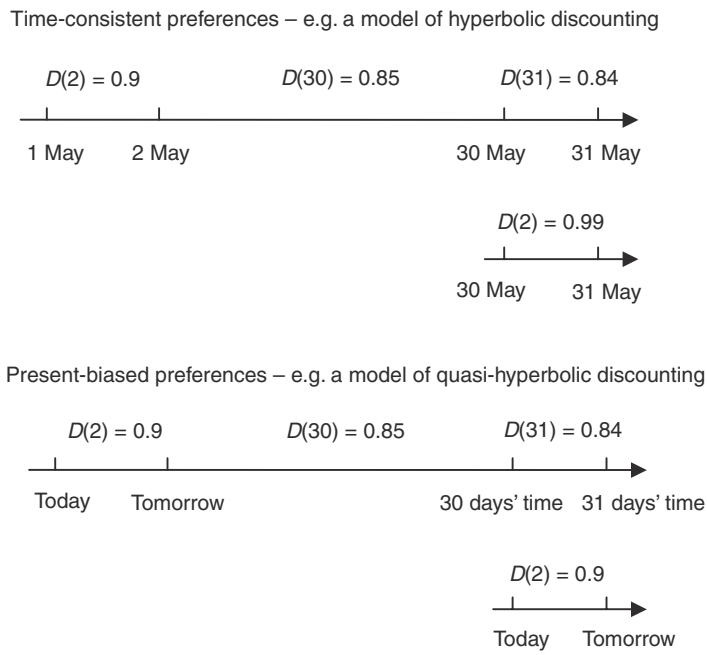

Figure 4.4 The difference between time-consistent and present-biased preferences. In the model of hyperbolic discounting, period refers to a specific moment of calendar time. In the model of quasi-hyperbolic discounting, period refers to a delay from now.

and where β is some number between zero and one. The β is crucial here because it measures **present bias**. If $\beta = 1$ then there is no bias, and equation (4.4) is equivalent to (4.1). If $\beta < 1$ then more weight is given to today than to the future. This means that there is a present bias and decreasing impatience. For instance, if $\beta = 0.7$ and $\delta = 0.99$ then \$110 in 31 days is always preferred to \$100 in 30 days, but \$100 today is always preferred to \$110 tomorrow.

To see how (β, δ) preferences work, we can return to the example of Maria deciding when to do her homework. Table 4.5 summarizes the payoffs that she gets if she does or does not have a present bias. First, look at what happens if $\beta = 1$, and so she does not have a present bias. On Friday she should plan to do her homework on Saturday, and on Saturday she also thinks this way, and so would do the homework. The same is true when $\beta = 0.9$. Next, look at what happens if $\beta = 0.8$, and so the present bias is relatively large. On Friday she would plan to do the homework on Saturday. The interesting thing is that, when Saturday comes, she would rather do it on Monday.

We say that there is a **time inconsistency** if someone plans to do something in the future but subsequently changes her mind. Planning to do the homework on Saturday but then on Saturday deciding to do it on Monday, and planning to wait an extra day for \$110 but when the time comes taking the \$100, are examples of time inconsistency. There cannot be time inconsistency in a model of exponential or hyperbolic discounting, but there can be in a model of quasi-hyperbolic discounting.

Table 4.5 Quasi-hyperbolic discounting in the homework example. If the present bias is high, Maria plans to do the homework on Saturday, but on Saturday would rather do it on Monday.

Plan	$\beta = 1, \delta = 0.9$		$\beta = 0.9, \delta = 0.9$		$\beta = 0.8, \delta = 0.9$	
	On Friday	**On Saturday**	**On Friday**	**On Saturday**	**On Friday**	**On Saturday**
Do it Friday	10.5	—	9.0	—	7.4	—
Saturday	10.9	12.1	9.8	10.4	8.7	8.7
Do it Sunday	7.7	8.6	7.0	8.2	6.2	7.9
Monday	9.0	10.0	8.1	9.5	7.2	9.0

Time inconsistency is potentially very important in terms of welfare and policy because it suggests that people plan to do something but then end up doing something else. I am, therefore, going to spend some time on the issue in Chapters 10 and 11 when we look at welfare and policy. At this point I want to show you how time inconsistency can have important consequences for behavior.

4.2.2 The consequences of time inconsistency

The consequences of time inconsistency will depend on whether people know they are time-inconsistent or not. For example, does Maria realize on Friday that, if she plans to do the homework on Saturday, when Saturday comes she might think differently? This could clearly be important in deciding what she does do on the Friday.

We say that someone is **naïve** if they are unaware that they have present-biased preferences. In this case, Maria would plan on Friday to do the homework on Saturday and not expect to think differently on Saturday. By contrast, we say that someone is **sophisticated** if they know that they have present-biased preferences. In this case, Maria will realize on Friday that if she leaves the homework until Saturday she will probably end up doing it on Monday.

I'll look first at what happens if Maria is naïve. It's easiest to start with cases in which the costs of an activity precede the benefits, as in the homework example. We saw in the example that on Friday Maria will expect to do the homework on Saturday, but on Saturday might decide to do it on Monday. This means that she can end up doing the homework later than she expected and later than she would have done without a present bias. To delay doing something in this way is called **procrastination**. Maria procrastinates because she puts off doing the costly thing. What happens if the benefits come before the costs?

To help illustrate, I will use a second example. Imagine it costs \$10 to go to the movies, and Maria has \$11 spending money. There are movies on Friday, Saturday and Sunday. Table 4.6 shows that she will prefer the movie on Sunday to that on Saturday and that on Saturday to that on Friday. Table 4.7 summarizes her inter-temporal utility (these tables can be compared to 4.1 and 4.5). If Maria has time-consistent preferences then she would plan to go to the movie on Sunday. If she has present-biased preferences then on Friday she would to plan to go on Sunday, but on Saturday would change her mind and go that day. Again, we see a time inconsistency. The difference is that this time Maria does something earlier than she expected and earlier than she would have done with no present bias. To bring forward something in this way we can call to **preproperate**.

Table 4.6 Payoffs from Maria going to watch a movie.

Plan	Payoff		
	On Friday	**On Saturday**	**On Sunday**
Go on Friday	5	0	0
Go on Saturday	0	6	0
Go on Sunday	0	0	8

Table 4.7 Inter-temporal utility in the movie example with time-consistent and present-biased preferences.

Plan	$\beta = 1, \delta = 0.9$		$\beta = 0.8, \delta = 0.9$	
	On Friday	**On Saturday**	**On Friday**	**On Saturday**
Go on Friday	5.0	—	5.0	—
Go on Saturday	5.4	6.0	4.3	6.0
Go on Sunday	6.5	7.2	5.2	5.8

Table 4.8 A summary of the two examples: choice and utility, calculated as inter-temporal payoff on Friday with no present bias.

	Delayed benefits (homework)		**Delayed costs (movie)**	
	Choice	**Payoff**	**Choice**	**Payoff**
Time-consistent	Saturday	10.9	Sunday	6.5
Sophisticated	Saturday	10.9	Friday	5.0
Naïve	Monday	9.0	Saturday	5.4

In both the homework and the movie examples, Maria is impatient for benefits. In the homework example this causes her to put off something that is costly. In the movie example it causes her to do early something that is pleasurable. Will someone who is sophisticated and knows they are time-inconsistent avoid such problems? The answer is a bit surprising.

The surprise is not in the homework example. If Maria knows that on Saturday she will delay until Monday, then on Friday she knows the real choice is between doing her homework now, giving utility 7.4, or on Monday, giving utility 7.2, so she will do the homework on Friday and behave as if time-consistent. Being sophisticated thus allows Maria to avoid any time inconsistency. So far, so good. But what about the movie example? In this case, Maria knows that if she does not go on Friday she will go on Saturday. On Friday she therefore knows that the real choice is between going today or going tomorrow. She will go today. Being sophisticated, therefore, means that Maria preproperates more than if she was naïve. She knows that she will not be able to resist going to the movies, and so goes even earlier!

Table 4.8 summarizes the choices made. We see that, in the case of delayed benefits, sophistication helps overcome the problems associated with present bias, but

in the case of delayed costs it makes things worse. It is not obvious what is better or worse in terms of utility because it is not clear whether to take into account the present bias or not. What we do see, however, is that a present bias has important implications for choice, even if someone is sophisticated.

The movie example shows that being sophisticated can make things worse, because Maria anticipates her future present bias. This misses part of the story, however, because someone who is sophisticated may use **commitment** to constrain her future choice. For instance, Maria could pre-order her movie ticket on Friday for Sunday. She would be willing to pay to do such a thing because it means she is committed to making the best choice. If she was naïve, she would see no need to do so. Someone who is sophisticated will thus be looking for ways that they can commit and avoid future temptation.

4.2.3 Temptation and self-control

The model of quasi-hyperbolic discounting is a nice, simple model that has allowed us to capture important aspects of inter-temporal choice. One thing that it does not capture so well is the potential benefits of commitment. An approach suggested by Gul and Pesendorfer (2001) provides us with a more elegant way to capture this. To illustrate how it works I will continue the homework example.

On Saturday Maria can either do the homework or play sport. What I want you to imagine now is the possibility that, on Friday, Maria can commit to what she will do on Saturday. So, she can commit to doing the homework or to playing sport by, for example, arranging things with friends. This means that on Friday Maria has three choices: commit to doing the homework on Saturday, commit to playing sport on Saturday or not commit and decide what to do on Saturday.

If Maria has time-consistent preferences then the possibility to commit is irrelevant. If she thinks it is best to do the homework on Saturday then she will be indifferent between committing to do the homework on Saturday and not committing, because she will do the homework on Saturday either way. If Maria has present-biased preferences then things are different.

If Maria thinks it is best that she does the homework on Saturday but knows that on Saturday she may not do it then it may be good for her to commit on the Friday to doing the homework on Saturday. What's interesting is that we can now distinguish two distinct reasons for her to commit. The more obvious one is that Maria knows on Saturday she will choose to play sport and wants to stop this happening. Looking back at Table 4.5, we can see this makes sense if $\beta = 0.8$. The second possibility is that Maria knows she will do the homework on Saturday but only after overcoming the temptation to play sport. That is, she will do the homework but only after exercising **self-control**, which is psychologically costly. She may be able to avoid this cost by committing on Friday to do the homework on Saturday. This fits the case when $\beta = 0.9$.

We can formalize this in a simple model. We already have u_2 to measure the utility of Maria on Saturday. If she does her homework she gets utility $u_2(H) = -5$, and if she plays sport she gets utility $u_2(S) = 5$. We are now going to add to this a cost of temptation. The cost of temptation should reflect Maria's psychological cost of not doing something she might have wanted to do. So, let $m(H)$ denote the **temptation**

to do homework and $m(S)$ the temptation to do sport. The **cost of temptation** from doing action a is then calculated as follows:

$$C(a) = m(a) - max_{b=H,S} m(b).$$

For instance, if the temptation to do sport is greater than that of doing homework, then $m(S) \geq m(H)$, and so the cost of temptation from doing the homework is $C(H) = m(H) - m(S)$. This is the psychological cost of doing the homework when Maria would rather be playing sport. The cost of temptation from doing sport is $C(S) = m(S) - m(S) = 0$. So, there is no cost of temptation from doing the thing she wants to do.

What if we add together Maria's basic utility plus the cost of temptation? Suppose that she has not committed on Friday to do something on Saturday. Then her overall utility on Saturday can be of two forms. First, Maria could exercise self-control and do the homework. In this case, her payoff should take account of the psychological cost of overcoming temptation. A general formulation for doing action a would be

$$u_2^m(a) = max_{a=H,S} \{u(a) + m(a)\} - max_{b=H,S} m(b). \quad (4.5)$$

In our specific example, this reduces to $u_2^m(H) = u(H) + m(H) - m(S) = u(H) + C(H)$. The second possibility is that Maria has not exercised self-control and so chooses the action that gives highest utility. In this case we can write

$$u_2^m(a) = max_{a=H,S} u(a) \text{ subject to } m(a) \geq max_{b=H,S} m(b). \quad (4.6)$$

In our example, this reduces to $u_2^m(a) = u(S)$. We can, therefore, distinguish the two reasons that Maria may want to commit.

Equations (4.5) and (4.6) might look a little over the top for what we need. The neat thing, though, is that, if Maria would want to pre-commit, we can say for certain that her overall utility must be like that in (4.5) or (4.6). More generally, if someone wants to be able to pre-commit then we must be able to represent their preferences with a utility function u and temptation function m in this way (see Research methods 4.2). This means that a desire to pre-commit must come hand in hand with temptation.

❖ RESEARCH METHODS 4.2 AXIOMS AND BEHAVIOR

Gul and Pesendorfer analyzed a much more general setting than that which we have discussed here and obtained a stronger result than I have suggested so far. They showed that preferences satisfy four 'intuitive' axioms if and only if preferences can be represented by functions such as u,m and u^m. As well as being a beautiful mathematical result, this is also an interesting illustration of the axiomatic approach to modeling choice.

In an axiomatic approach a researcher basically questions what properties of preferences are consistent with a particular way of modeling choice. Gul and Pesendorfer look at four axioms. Informally, these are that preferences (1) be complete and transitive, (2) be continuous, (3) satisfy independence, in the sense that a third option would not change the choice between two others, and (4) satisfy a preference for commitment, meaning that if someone would rather they did one action than another then they prefer to commit to doing that action. Their results show that, if you think a person's preferences satisfy these axioms, you have to think functions u,m and u^m are a good way to model choice. Similarly, if you think u,m and u^m are a good way to model choice, you have to think these axioms are reasonable.

The axiomatic approach has a very long history in economics. In principle, it seems entirely consistent with behavioral economics in trying to match plausible assumptions on preferences with observed choices. In reality, something of a divide has emerged between behavioral economics and the axiomatic approach. This is probably because axioms that seemed sensible have not held up to scrutiny in the experimental laboratory, transitivity being one example. The axiomatic approach does, however, give definitive answers about when ways of modeling choice are appropriate, and so it clearly does have its place, as Gul and Pesendorfer show.

To illustrate, Table 4.9 follows through what Maria's preferences may look like. If there is no present bias then there is no temptation, and so it is natural to think that $m(S) = m(H) = 0$. It does not matter whether she commits or not. If there is a relatively small present bias, $\beta = 0.9$, then we know that Maria will do the homework on Saturday. She still, though, gains from pre-committing to sport. In the example, I have set $m(S) = 0.5$ and $m(H) = 0$. So, by pre-committing to do her homework she can avoid the 0.5 cost of temptation and increase the utility she gets on Saturday. If there is a larger present bias, $\beta = 0.8$, then Maria will delay doing the homework unless she commits to doing it on Saturday. This is motive enough to commit.

This model allows us to nicely capture the benefits of commitment. In particular, it illustrates how someone may make consistent choices yet still desire commitment,

Table 4.9 The utility of Maria on Saturday. The basic utility u^T is the same as in Table 4.5, the temptation is given by m and the overall utility, taking into account the cost of temptation, is given by u^{Tm}. By committing to the homework, she avoids the cost of temptation when $\beta = 0.9$ and avoids procrastinating when $\beta = 0.8$.

	$\beta = 1, \delta = 0.9$			$\beta = 0.9, \delta = 0.9$			$\beta = 0.8, \delta = 0.9$		
	u^T	M	u^{Tm}	u^T	m	u^{Tm}	u^T	m	u^{Tm}
Committed to homework	12.1	0	12.1	10.4	0	10.4	8.7	0	8.7
Committed to sport	10.0	0	10.0	9.5	0	9.5	9.0	0	9.0
Decide on Saturday	12.1	0	12.1	10.4	0.5	9.9	9.0	0	9.0

in order to avoid the psychological costs of temptation. Altogether, therefore, we now have some useful ways to model decreasing impatience and present bias. We have addressed only one of the issues raised in section 4.1, however. Now it is time to consider some of the others.

4.3 Loss aversion and sequences

What I want to do now is look back on things such as the delay–speed-up asymmetry and see if we can better understand them. In doing so, I will draw heavily on work by Loewenstein and Prelec (1992, 1993).

4.3.1 Reference dependence

We can begin by applying the idea of reference dependence to inter-temporal choice. This means writing the inter-temporal utility function as

$$U^T(x_1, x_2, ..., x_T) = \sum_{t=1}^{T} D(t)v(x_t - r_t),$$

where $D(t)$ is the discount factor, $v(x)$ is some value function and r_t is the reference point in period t. We shall not say much about the discount factor, because everything we did in section 4.2 is relevant here. Our focus, therefore, will be on the value function.

As we saw in Chapters 2 and 3, the most important thing about the value function is that it measures consumption relative to some reference point. So, instead of thinking of x_t as the payoff in period t, we need to think of $x_t - r_t$ as the loss or gain relative to the reference point. This raises familiar questions about what the reference point should be, but we can assume for now that the reference point is current income. What does the value function need to look like for us to capture the behavior I talked about in section 4.1.1?

To answer that question, suppose that Maria is indifferent between getting $\$q$ now rather than $\$y > q$ in some future period t. This means $v(q) = D(t)v(y)$. Recall that the **gain–loss asymmetry** implies that she should be relatively reluctant to postpone payment compared to postponing receipt. So, she should prefer to pay $\$q$ now instead of having to pay $\$y$ in period t. For instance, if she is indifferent between getting \$10 today or \$11 tomorrow, she will prefer to pay \$10 today than \$11 tomorrow. This means $v(-q) > D(t)v(-y)$. The gain–loss asymmetry requires, therefore, that

$$\frac{v(q)}{v(y)} < \frac{v(-q)}{v(-y)}. \tag{4.7}$$

Recall that the **absolute magnitude effect** means that, the larger the amount, the more patient she is. So, Maria should prefer getting $\$\alpha y$ in period t to $\$\alpha q$ now, if

$\alpha > 1$ captures some proportional increase in the amounts to be paid. This means that $v(\alpha q) < D(t)v(\alpha y)$. The absolute magnitude effect requires, therefore, that

$$\frac{v(q)}{v(y)} > \frac{v(\alpha q)}{v(\alpha y)}. \tag{4.8}$$

Finally, we come to the **delay–speed-up asymmetry**. Speed-up is naturally thought of as expediting a receipt, in order to get $\$q > 0$ now rather than $\$y > 0$ in some future period t. One can think of the reference level as zero now and $\$y$ in the future. By expediting the receipt, Maria loses, therefore, $\$y$ in the future but gains $\$q$ now. If she is indifferent between doing this, then $-D(t)v(-y) = v(q)$. We can estimate a discount factor for speeding up t periods using $q = \delta^{speed\ up}\, y$ to give

$$\delta^{speed\ up} = \frac{q}{y} = \frac{v^{-1}(-D(t)v(-y))}{y}.$$

Delay is naturally thought of as delaying a receipt so as to get $\$y > 0$ at some future period t rather than $\$q > 0$ now. This time we can think of the reference level as $\$q$ now and zero in the future. By delaying receipt, Maria loses $\$q$ now but gains $\$y$ in the future. If she is indifferent to do this, then $-D(t)v(y) = v(-q)$. We can estimate a discount factor for delaying t periods using $q = \delta^{delay}\, y$ to give

$$\delta^{delay} = \frac{q}{y} - \frac{v^{-1}(-D(t)v(y))}{y}.$$

Suppose we compare these two discount factors and, for simplicity, assume that $D(t) = 1$. Then $\delta^{speed\ up} > \delta^{delay}$ if

$$v^{-1}(-v(-y)) > -v^{-1}(-v(y)).$$

If you stare at this long enough you might see that we require $-v(-y) > v(y)$, but this is what we get with loss aversion. So, all we need to explain the delay–speed-up asymmetry is loss aversion.

A model with reference-dependent preferences can, therefore, capture a lot of the things we observe in inter-temporal choice. Unfortunately, things do not work out as perfectly as we might hope, because the formula for the value function we looked at in Chapter 3 (see equation (3.7)) does not satisfy either relation (4.7) or (4.8)! We could easily tweak things and sort that out, and conclude, as in Chapters 2 and 3, that reference dependence helps a lot in understanding behavior. (In doing that I should give a reminder, from Chapter 1, of the need for care in adapting, or tweaking, a model to fit the facts.)

As you might expect by now, we do need to think a bit more about the reference point. To illustrate why, let's return to the homework example and suppose Maria must do the homework on Saturday or Monday. Furthermore, suppose that it is rare

Table 4.10 The homework example with reference-dependent preferences. If the reference point is to play sport, doing the homework on Saturday feels like a loss, but, if the reference point is to visit an aunt, doing the homework on Saturday feels like a gain.

	Utility		
	On Saturday	**On Sunday**	**On Monday**
Payoff x_t if do homework Saturday	–5	10	10
Payoff x_t if do homework Monday	5	10	–5
Reference point is to play sport			
Reference payoff r_t	5	10	0
Reference-dependent utility if do homework Saturday	–10	0	10
If do homework Monday	0	0	–5
Reference point is to visit aunt			
Reference payoff r_t	–10	10	0
Reference-dependent utility if do homework Saturday	5	0	10
If do homework Monday	15	0	–5

for homework to be set over the weekend, and so this is not her reference point. Instead, her reference point is to doing something on Saturday. One possibility is that on a Saturday she usually plays sport. Another is that she usually has to visit an aunt she does not like. This gives the reference-dependent utility in Table 4.10.

Irrespective of the reference point, Maria faces the same trade-off between a relatively large drop in utility on Saturday, a difference of ten, for a gain on Monday, a difference of 15. The reference point will matter in how she perceives this trade-off. Let's go through each reference point in turn and so how things will look for Maria on Friday. If she normally plays sport on Saturday then doing the homework on Saturday will feel like a loss. This might make her reluctant to do the homework. If

$$D(2)v(-10) + D(4)v(10) < D(4)v(-5)$$

then she will leave the homework until Monday. If she normally visits her aunt on Saturday then doing the homework on Saturday will feel like a gain. This might make her more likely to the homework. If

$$D(2)v(5) + D(4)v(10) > D(2)v(15) + D(4)v(-5)$$

then she would want to do the homework on Saturday.

To focus on the issue at hand, suppose that $D(2) = D(4) = 1$. Then we need to know whether $v(-10) + v(10) < v(-5)$ and whether $v(5) + v(10) > v(15) + v(-5)$. These equations may or may not be satisfied but suppose, by way of illustration, that $v(5) = 3, v(10) = 5, v(15) = 6, v(-5) = -7$ and $v(-10) = -13$. Then she would not do her homework on Saturday if the reference point is playing sport, but she would do her homework on Saturday if the reference point is to visit her aunt. The reference point and context can, therefore, influence behavior.

4.3.2 Preferences for sequences

What I want to look at next is how to model preferences over sequences. For now we shall ignore the question of when a set of events is viewed as a sequence rather than separate. So, what we need to do is measure the overall utility of a sequence of payoffs such as $\{x_1, x_2, \ldots, x_T\}$.

One way we do so is to compare this sequence with one that gives the same payoff in each period. That is, we compare the actual sequence to the smoothed sequence $\{\bar{x}, \bar{x}, \ldots, \bar{x}\}$, where

$$\bar{x} = \frac{1}{T}\sum_{t=1}^{T} x_t.$$

For every period t, let

$$d_t = t\bar{x} - \sum_{t=1}^{t} x_i$$

be the difference in utility, up to and including period t, between the smoothed sequence and actual sequence. We then get a nice way to think about the utility of a sequence by saying that utility is given by

$$U^q(\{x_1, x_2, \ldots, x_T\}) = \sum_{t=1}^{T} x_t + \beta\sum_{t=1}^{T} d_t + \sigma\sum_{t=1}^{T} |\, d_t \,|. \tag{4.9}$$

To see why this is a nice way to think about things I need to introduce two further concepts.

We say that the **anticipated utility** (AU) of receiving sequence $\{x_1, x_2, \ldots, x_T\}$ is given by

$$AU = \sum_{t=1}^{T} (t-1)x_t.$$

The story is that if Maria is going to receive payoff x_t in period t then she spends $t-1$ periods anticipating this. So, she gets $(t-1)x_t$ anticipated utility. Summing this over all periods gives total anticipated utility.

Following a similar logic, we say that the **recollected utility** (RU) of the sequence is

$$RU = \sum_{t=1}^{T} (T-t)x_t.$$

If Maria received payoff x_t in period t then she has $T-t$ periods to recollect it, and so gets $(T-t)x_t$ recollected utility.

A bit of rearranging shows that

$$\sum_{t=1}^{T} d_t = 0.5(AU - RU).$$

So, positive values of Σd_t are associated with more anticipated than recollected utility. This means that a high Σd_t goes hand in hand with an improving sequence in which the x_ts increase with t.

Now we can go back to equation (4.9) and make some sense of it. We can see that a $\beta > 0$ means a preference for an improving sequence and a $\beta < 0$ a preference for a worsening sequence. We also see that $\sigma < 0$ indicates a preference for a smoothed sequence while $\sigma > 0$ indicates a preference for a one-sided sequence. To illustrate what this means, Table 4.11 has some example sequences. If σ is small and β positive then Maria would prefer sequence C with lots of anticipation utility. If σ is negative and β positive then she may prefer sequence E, which still has anticipation utility but is more smoothed over time.

The evidence we looked at in section 4.1.2 would suggest that $\beta > 0$ and $\sigma < 0$. With the help of Table 4.12 we can see what this means in the homework example. Doing the homework on Saturday gives the most improving sequence, but not a very smooth one. To do the homework on Friday or Sunday may, therefore, be optimal. Indeed, if σ is very negative, Monday may be optimal, as it gives the smoothest sequence.

Table 4.11 Preference over sequences of outcomes with the anticipated and recollected utility.

	Period					AU	RU	Σd_t
	1	2	3	4	5			
Sequence A	1	1	1	1	1	10	10	0
Sequence B	5	0	0	0	0	0	20	−10
Sequence C	0	0	0	0	5	20	0	10
Sequence D	2	1	1	1	0	6	14	−4
Sequence E	0	1	1	1	2	14	6	4
Sequence F	2	0	1	0	2	10	10	0

Table 4.12 Preferences over sequences in the homework example.

Plan	Payoff				AU	RU	Σd_t
	On Friday	On Saturday	On Sunday	On Monday			
Do it Friday	−5	5	10	4	37	5	16
Do it Saturday	0	−5	10	10	45	0	22.5
Do it Sunday	0	5	−5	10	25	5	10
Do it Monday	0	5	10	−5	10	20	−5

The homework example brings us back to the troubling question of when period payoffs are seen as part of a sequence and when they are not. Will Maria think of the weekend as one sequence of events or think of each day as separate from another? We see now that this might affect her choice. Unfortunately, there is no simple rule as to when a person will think of periods as separate or part of a sequence. The further apart are two events, the more likely it intuitively seems that they will be seen as two separate events rather than as a sequence of events. This logic can only take us so far, however, as we shall see shortly in looking at habit formation.

4.4 Time and risk

The focus throughout the chapter so far has been on time. In this final section, before we look at some applications, I want to briefly illustrate the ways in which risk inevitably starts to enter the picture whenever we think about time. To get us started, consider Maria doing her homework. Before she does the homework she cannot tell for sure how easy it is going to be or how long it will take her. Doing the homework is, therefore, best thought of as a prospect. Recall (see Chapter 3) that a prospect is a list of probabilities that things will happen together with a payoff if they do happen. In this case, we have the probability the homework is easy or hard. Similarly, her payoff from not doing the homework may depend on a host of unpredictable events, such as the weather is sunny and whether her favorite football team wins. This is also best thought of as a prospect.

The simplest way to model risk and time is to combine expected utility (see equation (3.1)) and exponential discounting (equation (4.1)) to get a model of **discounted expected utility**. The basic idea here is that in each time period t Maria faces a prospect $A^t = (p_1{}^t, x_1{}^t; p_2{}^t, x_2{}^t; \ldots; p_n{}^t, x_n{}^t)$, where there are n possible outcomes. Her expected utility in period t can be calculated using the formula

$$U(A^t) = \sum_{i=1}^{n} p_i^t u(x_i^t).$$

Her inter-temporal utility over the T periods can then be calculated using exponential discounting to give

$$U^T(A^1, A^2, \ldots, A^T) = \sum_{t=1}^{T} \delta^{t-1} \left(\sum_{i=1}^{n} p_i^t u(x_i^t) \right).$$

With this approach, the numbers in Table 4.1 should be thought of as Maria's expected utility on, say, Friday if she does the homework on Friday.

At first sight it may look as though discounted expected utility allows us to treat risk and time separately, because we focus on risk within each period and then look at time over periods. Things are not so simple, however. To appreciate why, let us consider the simple case of Maria deciding whether she wants \$100 today or \$110 tomorrow. Notice that in this case there is no risk at all. If she chooses the \$100 today then her discounted expected utility over the two days would be $u(c + 100) +$

$\delta u(c)$, where c is the amount she would normally consume in a day and δ is her discount factor. This formulation recognizes that she gets to spend an extra \$100 today. If she waits until tomorrow then her discounted expected utility would be $u(c) + \delta u(c + 110)$. In this case she gets to spend an extra \$110 tomorrow.

Maria should choose the money today if

$$u(c + 100) + \delta u(c) > u(c) + \delta u(c + 110).$$

So what should she choose? If she is risk-neutral then her utility function is linear and so we get the simple condition that she should choose the money today if $c + 100 + \delta c > c + \delta (c + 110)$, which simplifies to the condition $\delta < 10/11$. Now suppose that Maria is risk-averse. To be more specific, suppose that $u(x) = \sqrt{x}$. Table 4.13 details her discounted expected utility if she takes the money today or tomorrow for four different values of the discount factor. In both cases we set c = \$100. For comparison, we also have the discounted expected utility for the case she is risk-neutral. The crucial thing to take from Table 4.13 is that risk aversion means Maria is more inclined to take the money today. For instance, if $\delta = 0.92$ then a risk-neutral Maria would wait until tomorrow but a risk-averse Maria would take the money today.

Risk aversion makes Maria less willing to wait because it means she values less the extra \$10 she gets from waiting. To explain this point, recall that risk aversion means the utility function is concave. That means Maria values less going from \$100 to \$110 (or \$200 to \$210) than she would going from \$0 to \$10. If she values less the extra \$10 she gets from waiting then it is no surprise she is less willing to wait. Risk aversion, therefore, can make people appear more impatient than they really are.

If we ignore risk, as is typical when looking at time, then we essentially assume people are risk-neutral. This means that the 'implausibly large' discount factors we observe in the lab (see section 4.1.1) may be partly down to risk aversion rather than impatience. A study by Andersen and co-authors (2008) allows us to quantify this effect. They measured both the risk preferences and time preferences for a representative sample of the Danish population. They find that, if subjects were assumed to be risk-neutral, the estimated discount factor would be 0.80. Subjects were not risk-neutral, though. Once risk aversion is taken into account the discount factor changes to a slightly more plausible 0.91. This discount factor still seems too low. Things are not as bad as they look, however, when we do not take into account attitudes to risk.

We have just seen that risk aversion can influence choice over time even if that choice involves no risk. Now let us start to add risk into choice. A study by Andreoni

Table 4.13 Discounted expected utility if taking the money today or tomorrow for different discount factors and optimal choice.

δ	Risk-neutral			Risk-averse		
	Today	**Tomorrow**	**Choice**	**Today**	**Tomorrow**	**Choice**
0.9	290	289	Today	23.14	23.04	Today
0.91	291	291.1	Tomorrow	23.24	23.19	Today
0.92	292	293.2	Tomorrow	23.34	23.33	Today
0.93	293	295.3	Tomorrow	23.44	23.48	Tomorrow

and Sprenger (2012) nicely demonstrates the issues that can arise. In the study, subjects were given a series of choices in which they had to allocate 100 tokens between two dates. The first date was seven days away and the second date was either 35 or 63 days away. In both cases, therefore, payment would be delayed, but for the second date it was more delayed (see Research methods 4.3). This allows a measure of preferences over time.

❖ RESEARCH METHODS 4.3 DELAYED PAYMENTS IN THE LAB

The typical lab experiment involves subjects being paid at the end of the session. If we want to study preferences over time then it is necessary to delay payment. For instance, in the study by Andersen and co-authors payment was delayed by up to two years. This raises some complex logistical problems, because the experimenter needs to keep track of how much a subject won and then arrange for the appropriate payment to arrive at the set date. And this needs to be done in a way that preserves as much anonymity as possible. A bigger concern, however, is that subjects might doubt the money will be paid as promised. Such doubts would add an extra, unwanted influence on choice.

To deal with these problems requires a careful experimental design, and the two studies by Andersen and co-authors and Andreoni and Sprenger are great examples of how this can be done. In both studies there was never a choice to take the money today. Instead, subjects had to choose between getting the money in one month or one week versus getting it in, say, six months or nine weeks. This approach completely rules out the temptation for subjects to take the money now. This is called **front end delay**.

We still need to instill in subjects confidence that the money will be paid as promised, though. In the study by Andersen and co-authors, payment was guaranteed by the Danish Ministry of Economic and Business Affairs and payment was made by automatic bank transfer. This should instill confidence. In the Andreoni and Sprenger study, all the subjects who took part lived on campus, where they had access to a personal mailbox. Subjects were given two envelopes, on which they wrote their address, the payment they were due and the date it was due. Payments were then sent by check through the campus mail service. Again, this should give subjects confidence the money will arrive as promised.

The crucial twist to the study is that payment was subject to risk. For instance, in one set of choices there was only a 50 percent chance that each payment would be sent. A role of the dice determined if the earlier payment, due in seven days, would be sent and another role of the dice determined if the later payment, due after 35 or 63 days, would be sent. For other choices these probabilities varied. To give a second example, for some choices there was an 80 percent chance the earlier payment would be sent and a 100 percent chance the later payment would be sent. This

particular scenario is interesting because it reverses the intuitive logic subjects may have that earlier payments are more secure than delayed payments.

Figure 4.5 compares choice in a setting in which there is no risk, (1, 1), to the case in which each payment has a 50 percent chance of being sent, (0.5, 0.5). The figure plots the amount of money, up to a maximum of $20, allocated to the earlier payment as a function of the return from delay. As we would expect, the higher the return from delay, the less subjects allocate to the earlier payment. More interesting for our purposes is the comparison between the (0.5, 0.5) and (1, 1) setting.

There are two key things to pick out from Figure 4.5. First, if a subject was maximizing dynamic expected utility then we would expect their choices to be exactly the same in the (0.5, 0.5) setting as the (1, 1) setting. Clearly, they are not. For the most part subjects were less willing to delay when there was only a 50 percent chance of getting the payment. Put another way, when payment was certain subjects displayed more patience to get the higher, delayed payment. This would appear similar to the certainty effect we looked in section 3.1.1 when discussing the Allais paradox.

The second thing to take from Figure 4.5 is the crossover effect for small values of the interest rate. If there was no extra payment for waiting – i.e. the interest rate was zero – then subjects were more willing to delay when there was only a 50 percent chance of getting the payment. At an intuitive level, this can be explained as a form of diversification. Given that both earlier and later payments independently occur with a 50 percent chance, it seems sensible to split money across the two payments so as to minimize the chance of getting nothing. It turns out, however, for

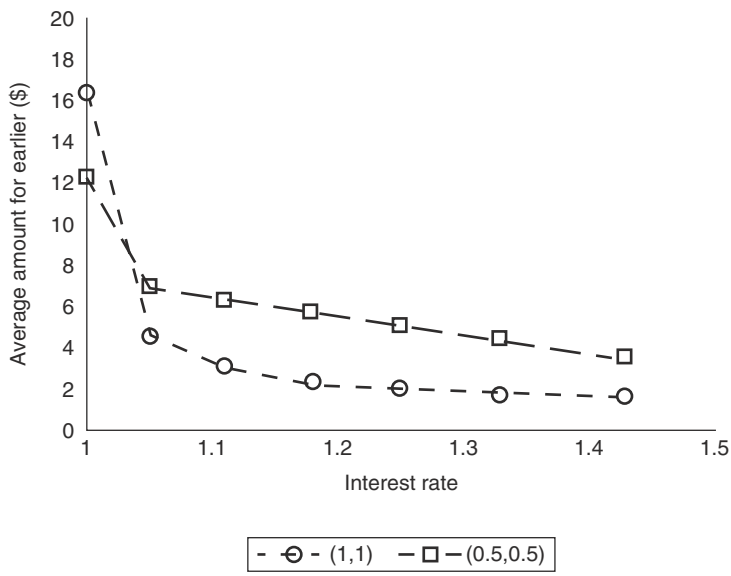

Figure 4.5 The amount of money allocated to the earlier payment, for different levels of interest rate, comparing certain payment (1, 1) with a 50 percent chance of payment (0.5, 0.5).
Source: Andreoni and Sprenger (2012).

reasons I will not delve into here, that the crossover effect we observe in Figure 4.5 is hard to reconcile with standard models of choice with risk.

Before I elaborate on this point in more detail, let us look at some more data. Figure 4.6 plots the amount that subjects allocated to the earlier choice for four different combinations of risk. Consider, first, the left-hand side of the figure, in which the earlier payment is more likely to be sent than the later payment. Specifically, we compare a setting in which the earlier payment is certain and the later payment is sent with 80 percent chance, (1, 0.8), with a setting in which the earlier payment is sent with 50 percent chance and the later payment with 40 percent chance, (0.5, 0.4). Here we see that subjects were less patient when there was a higher chance of receiving payment. This is consistent with the extreme left point on Figure 4.5.

Now look at the right-hand side of Figure 4.6. Here the earlier payment is less likely to be sent than the later payment. Specifically, we are comparing a setting in which the later payment is certain (0.8, 1) to one in which there is a 50 percent chance of the payment being sent (0.4, 0.5). Here we see subjects were more patient when there was a higher chance of receiving payment. This is consistent with the right-hand side of Figure 4.5. Crucially, this reinforces the idea that subjects appear to have a preference for certainty.

As I have already alluded to above, the results in Figures 4.5 and 4.6 are difficult to reconcile with discounted expected utility. Given the difficulties of modeling choice with risk and choice with time that we have looked at in the last two chapters, however, it should not come as much of a surprise that it is complicated to model choice with risk and time. Recent work by Epper and Fehr-Duda (2015) provides something of a breakthrough in showing that a form of rank-dependent probability weighting

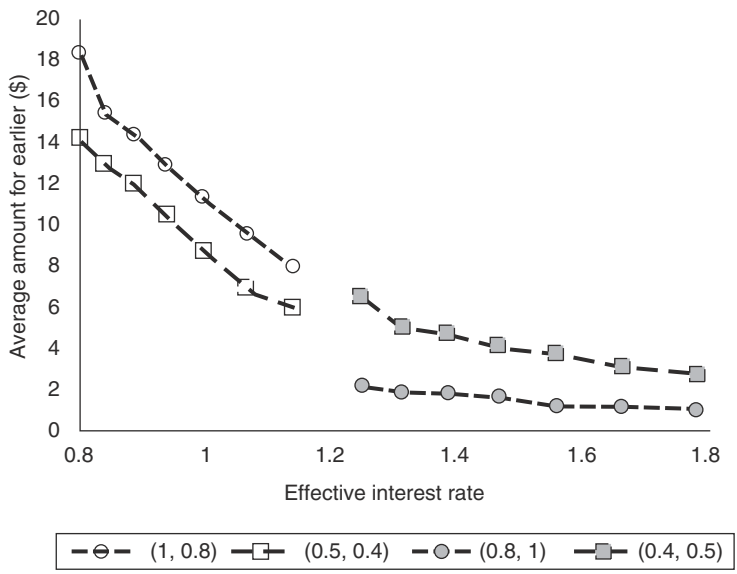

Figure 4.6 The amount of money allocated to the earlier payment, for different levels of interest rate, comparing four different levels of risk.
Source: Andreoni and Sprenger (2012).

(see section 3.3) can explain the choices in Figures 4.5 and 4.6. Key to this approach is the idea that subjects look at their portfolio of risk rather than just focusing on individual gambles (see section 3.3.3).

[Extra] It is useful to work through why rank-dependent weighting makes a difference, not least because it provides an interesting illustration of how rank dependence works. Recall that with rank-dependent probability weighting outcomes are ordered from the worst, x_1, to the best, x_n. Decision weights, in the case when there are only gains, are then given by

$$w_i = \pi (p_i + \ldots + p_n) - \pi (p_{i+1} + \ldots + p_n)$$

for all i where function π weights probabilities. With this in mind, consider the (0.5, 0.5) setting in which each payment is made with 50 percent chance. There are four possible outcomes in this case, all of which happen with probability 0.25: both the earlier and later payments are sent, only the earlier payment is sent, only the later payment is sent or no payments are sent. The best outcome, clearly, is that both payments are sent. This has decision weight $w_1 = \pi(0.5 \times 0.5) = \pi(0.25)$. If you look back at Figure 3.8, it is reasonable to set $\pi(0.25) = 0.25$, meaning that decision weight equals probability. Suppose that the second best outcome is to get the earlier payment but not the later payment. This has decision weight $w_2 = \pi(0.5) - \pi(0.25)$. It is reasonable, again with reference to Figure 3.8, to set $\pi(0.5) = 0.3$. But that would mean $w_2 = 0.05$. Crucially, this means that the second best outcome is being underweighted, because its decision weight of 0.05 is a lot less than the actual probability of 0.25. The third best outcome is to get the later payment and not the earlier payment. This has decision weight $w_3 = \pi(1 - 0.25) - \pi(0.5)$. Setting $\pi(0.75) = 0.4$ would give $w_3 = 0.1$. The third best outcome is also, therefore being underweighted. Finally, the worst outcome is that neither payment is sent. This has decision weight $w_4 = 1 - \pi(0.75)$, giving $w_4 = 0.6$. Note how this worst outcome is being considerably overweighted. With such decision weights, it is easy to explain why people would want to diversify their allocations across the two periods.

To finish the discussion, let us bring the focus back on time. One key implication of the results from the Andreoni and Sprenger study is that experimental subjects react strongly to certainty. This could have implications when we come to measure the discount factor. In particular, when we elicit choice over time it would seem natural for experimental subjects to think of money today as certain and money in the future as less certain. This would bias them towards taking the money today. Common ways of eliciting choice over time may, therefore, overestimate short-term impatience and present bias. Recall that we said risk aversion can also lead to an overestimate of impatience. For a full understanding of people's attitude to time it is, therefore, crucial to take account of attitudes to risk as well.

4.5 Summary

I began by looking at a utility function with exponential discounting. The basic idea here is that inter-temporal utility, and choice, can be captured by assuming that people discount the future at a constant rate. If true, this would make our life easier. Unfortunately, we saw that preferences seem to depend a lot on context, and

exponential discounting is not well placed to pick this up. For example, we observe things such as a gain–loss asymmetry, delay–speed-up asymmetry, absolute magnitude effect and preference for improving sequences.

In looking at alternatives to exponential discounting we focused first on the fact that people seem impatient for short-term gain. We saw that there are two distinct ways to think about this. We can think in terms of a model of hyperbolic discounting with time-consistent preferences, or we can think in terms of quasi-hyperbolic discounting with time-inconsistent or present-biased preferences. The latter looked the better way to go.

Present bias, and time inconsistency, raise some interesting possibilities. A person may procrastinate by putting off doing something costly, preproperate by doing early something pleasant or fight to overcome temptation. What happens can depend on whether the person is naïve or sophisticated in knowing about their present bias.

We next added reference dependence and loss aversion to the mix and showed this can explain things such as the delay–speed-up asymmetry and gain–loss asymmetry. We also saw that perceptions of the reference point can influence choice.

Next we looked at how we can capture the utility from a sequence of events using anticipated and recollected utility.

Finally, we added risk into the mix. We showed that risk aversion can result in people appearing less patient than they really are. We also saw that when there is choice involving risk and time there is a crossover effect in which people tend towards certain outcomes even if those certain outcomes involve delay. This means that common ways of eliciting the discount factor may overestimate present bias.

A recurring theme in Chapters 2 and 3 was that context mattered. It would influence perceptions, which would then influence choice because of things such as the reference point. In thinking about inter-temporal choice we can come up with many other good reasons why context will matter. For instance, whether someone perceives a series of events as a sequence or separated can have a huge impact on behavior: in the one they are impatient for gain and in the other they want to delay for the future.

Context effects will be less of a theme in the rest of the book (though they will crop up again in Chapter 7). So, before we move on, I do want to emphasize their importance. Most economic choices involve choice arbitrariness, risk and time, all mixed up together. The potential context effects are, therefore, huge, and behavior will depend a lot on the context, and perceptions of the context. It's vital to be aware of and try to account for this. It's a key part of being a behavioral economist. What this means in practice is questioning things such as what is the reference point, and what is a loss or a gain, in the different contexts that you are interested in understanding. This is a matter of judgment.

4.6 Borrowing and saving

In Chapter 2 we spent some time looking at the life cycle hypothesis and saving. Clearly, these are inter-temporal issues, and so it is natural to think back to them now that we have a means to model inter-temporal choice. In doing so there are two new issues I want to pick out: habit formation and consumer debt.

4.6.1 Saving equals growth or growth equals saving?

A high savings rate usually goes together with a high growth rate. What is not so clear is why. A standard life cycle model of consumption predicts that higher expected income should lead to less saving, because the higher expected income means less income needs to be saved now in order to smooth future consumption. The prediction, therefore, would be that high saving causes high growth, and not the other way around. A lot of evidence, however, tends to suggest that things do go the other way around. It seems that high growth tends to precede and cause high saving.

Habit formation is one possible way to account for this, as shown by Carroll, Overland and Weil (2000). I am going to look at habit in more detail in Chapter 10, but the basics are all we need here. In a **model of habit formation**, the utility of today's consumption depends on past consumption. One possible reason for this is that a person wants to maintain consumption at past levels. They have got used to a certain standard of living and would rather keep to, or improve upon, that standard. This fits with the idea that people want an improving sequence. People want to see their income and standard of living improving and not getting worse.

To capture this, assume that Maria has a habit level of consumption h. The habit evolves over time according to

$$\frac{dh}{dt} = \rho(c - h),$$

where c is consumption and ρ determines the extent to which habit depends on the recent past. The larger ρ is, the more quickly Maria's habits adapt to recent consumption.

A simple utility function we could use would be

$$u(c, h) = \frac{c}{h^{\gamma}},$$

where γ measures the importance of habit. If $\gamma = 0$ then habits are irrelevant and we have a standard model in which utility depends on consumption. If $\gamma = 1$ then utility depends only on whether consumption is more or less than the habit level. For intermediate γ, utility depends on both the absolute level of consumption and whether it is more or less than the habit level. For example, if $\gamma = 0.5$ then consumption of two with a habit level of one gives the same utility as if both consumption and the habit level are four.

It is possible to show, with some reasonable assumptions, that savings will increase with income if and only if the importance of habit γ is sufficiently large. Thus, habit can cause high savings if there is a high growth rate of income. The intuition for this is that a habit level of consumption means less reason to devote any increase in income to consumption. As long as consumption is above the habit level and the standard of living is rising, Maria will get relatively high utility. So, she does best to increase her consumption slowly over time and always have a rising standard of living. What she would not want to do is devote all the increased income to consumption. This would risk a subsequent fall in consumption below her new habit level, with a consequent fall in living standards and utility.

Table 4.14 Two sequences of consumption compared in a model of habit formation. Consumption schedule B gives the higher utility.

	Period				Total utility
	1	2	3	4	
Income	15	15	30	100	
Consumption A	40	40	40	40	
Habit A	15.0	27.5	33.75	36.9	
Utility A	10.3	7.6	6.9	6.6	31.4
Consumption B	25	35	45	55	
Habit B	15.0	20.0	27.5	36.0	
Utility B	6.5	7.8	8.6	9.1	32.0

Table 4.14 illustrates how this may work with $\gamma = 0.5$ and $\rho = 0.5$. We can see that income is going to increase over the four periods. We then can compare a smoothed consumption schedule A to an increasing schedule B. The increasing schedule gives higher utility because Maria does relatively well compared to what she has become used to. It also guards against any potential fall in the income she might get in period four.

We see, therefore, that habit and a preference for an improving sequence can explain why people save rather than spend, even if they expect future income growth. That people may think about their lifetime earnings as a sequence also highlights the problems of trying to distinguish when events are seen as separated rather than as a sequence. Here we are thinking of someone potentially seeing lifetime earnings as a sequence!

4.6.2 Why save when you have debts?

One of the more puzzling aspects of saving behavior is that many people borrow on credit cards at a high interest rate to buy consumer goods while simultaneously saving for the future at lower interest rates. To give some figures, in March 2017 US consumer debt, made up primarily of credit card debt, was over $12 trillion, and around 10 percent of disposable income went to paying off this debt. At the same time the savings rate was around 5 percent of disposable income. One might argue that some people are saving while others are borrowing, but this does not seem to be the case; many people are both saving and borrowing. Why?

In trying to explain this puzzle the key thing to recognize is that people typically borrow for the short term (i.e. to buy a new dress) while saving for the long term (i.e. to buy a new house or retire). This distinction between the short term and the long term leads naturally to thoughts of quasi-hyperbolic discounting and time inconsistency. Several studies have indeed shown that this can help explain what is going on. For example, Laibson, Repetto and Tobacman (2007) try fitting the (β, δ) preferences model to data from US households. Having taken account of risk attitudes, they find that the model with parameters $\beta = 0.90$ and $\delta = 0.96$ fits the data well. This equates to a long-run discount rate of 4.1 percent and a short-term rate of 14.6 percent.

Why does a model with (β, δ) preferences work? First, a consumer with a long-run discount rate of 4.1 percent would be happy to save over the long term at a

savings rate of, say, 5 percent per year. Second, a consumer with a short-term discount rate of 14.6 percent would be happy to borrow money to buy something now if she can pay that back at a rate of, say, 14 percent per year. Basically, the consumer wants to save but because of time inconsistency may also be unwilling to wait to buy that new dress.

Clearly, banks and credit card companies can profit from this, which brings us nicely on to the next topic.

4.7 Exploiting time inconsistency

We have seen that present-biased preferences and time inconsistency can result in someone procrastinating. They keep on delaying something they need to do. Could a company exploit this? To see why this question is pertinent, I am going to look at a study by Della Vigna and Malmendier (2006). The study was published with the title 'Paying not to go to the gym', which somewhat gives away the punchline.

4.7.1 Time inconsistency and consumer behavior

DellaVigna and Malmendier looked at attendance data for over 7,000 health club members at three health clubs in New England between April 1997 and July 2000. People going to the gym had four basic options: (1) pay $12 per visit, (2) pay $100 every ten visits, (3) sign a monthly contract with a fee of around $85 per month and (4) sign an annual contract with a fee of around $850. One difference between the monthly and annual contract is that the monthly contract was automatically renewed, while the annual contract was not. Those on a monthly contract could cancel at any time, but needed to do so in person or in writing. Those on an annual contract needed to sign up again at the end of the term, or their membership would stop.

This difference between the monthly and annual contract is a difference of what happens if the member does nothing, or is in the **default position**. The default position with a monthly contract is that a member needs to **opt out** of being a member. With the annual contract a member needs to **opt in** to being a member. Why should we worry about this distinction? Possibly we should not, because it costs very little to make a phone call, or write a letter, to opt in or opt out. There is potential for time inconsistency, however. A person may want to cancel, or rejoin, but put it off until tomorrow, and then put it off until the next day, and so on. Similarly, a person may think they will go to the gym often, but typically find some excuse to go tomorrow rather than today. Let's look at the data.

Any of the four payment options could make sense to a particular consumer. They may go rarely, so the pay per visit makes sense; they may know they will go often, so the annual contract makes sense; or they may be unsure how often they will go, so prefer the flexibility of the monthly option. The interesting question is whether consumers do choose the option that is best for them. The study finds that, basically, they do not. To give a first illustration of this, Figure 4.7 plots the average price per attendance of new members. We see that in the first six months of membership the average cost per visit was relatively high. Indeed, 80 percent of monthly members would have been better off paying for every ten visits.

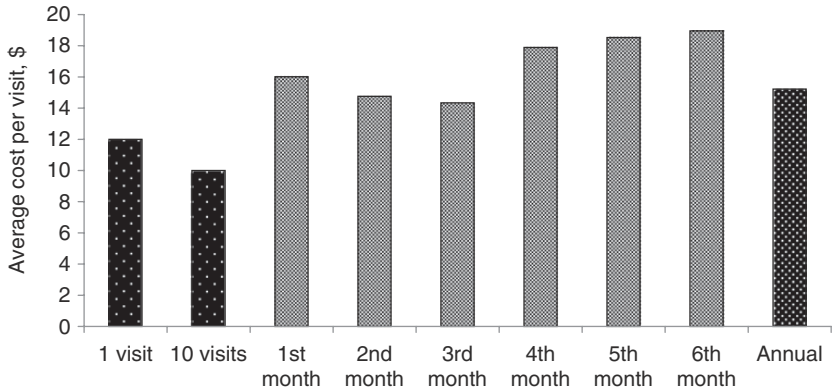

Figure 4.7 Average cost per visit to the gym for the first six months of monthly membership or year of annual membership. Monthly members pay on average more per visit than if they had paid per visit.
Source: DellaVigna and Malmendier (2006).

Figure 4.7 is suggestive of customers not choosing the best option, but does not capture the whole picture. That is because we need to see whether people learn and readjust. It may have been optimal for someone to sign a monthly contract but then switch to paying per visit once they realized how much (or how little) they used the gym. There is little evidence of readjustment, however. This is particularly the case for those on the monthly contract. Many monthly members appeared to consistently overestimate their future usage and put off canceling their membership. Particularly telling is the average lag of over two months between the last time a person attends the gym and the moment they cancel membership! Such a delay is costly. Looking over the whole period of membership, the average loss for someone on the monthly contract was $614. For those on the annual contract it was just $1.

In trying to make sense of these results, DellaVigna and Malmendier suggest the most likely explanations are overestimation of future usage, time inconsistency and naïvety. A model of naïve consumers with (β, δ) preferences where $\beta = 0.7$ and $\delta = 0.9995$ can explain the data, including the average delay of over two months in cancelation.

4.7.2 Firm pricing

In the gym example, the loss made by a consumer on the monthly contract was the firm's gain. This raises the question of whether firms use a pricing strategy that maximizes the profit they can make from consumers. To answer this question we first need to know how firms should price if consumers have a present bias. To do so I will work through a model of gym membership based on that by DellaVigna and Malmendier (2004).

There are three periods. In period one the gym proposes a **two-part tariff** (L, p) that consists of a **lump sum membership fee** L and a **user fee** p that is paid every visit. If Maria joins the gym then in period two she pays the membership fee and decides

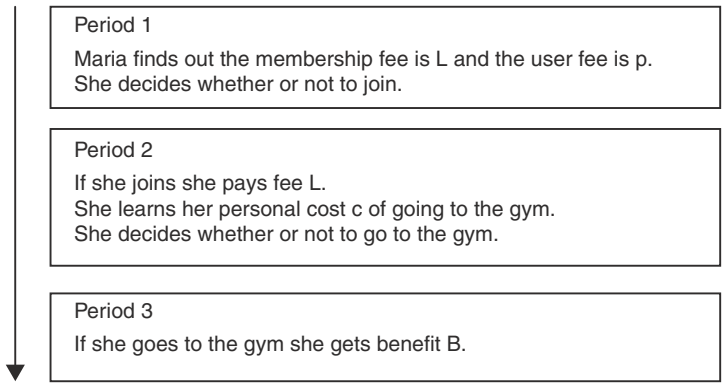

Figure 4.8 The timeline of the model of gym pricing.

whether or not to use the gym. If she uses the gym then she pays the user fee, *p*, plus some personal cost to attend, *c*. Maria does not know the personal cost until after joining. If she does go to the gym then in period three she receives some benefit, *B*. Figure 4.8 summarizes what happens. The key to the model is the delays. There is a delay between joining and using the gym, so Maria needs to predict future usage when joining. There is also a delay between using the gym and benefiting from having used the gym.

This delay gives scope for present bias to matter. So, let's assume that Maria has (β, δ) preferences, but thinks she has $(\hat{\beta}, \delta)$ preferences. If $\hat{\beta} > \beta$ then she has a present bias but underestimates the strength of the bias.

To see what Maria will do we first need to imagine that we are in period one and Maria is predicting what she will do in period two. She expects that when period two comes she will think the payoff from using the gym is $\delta\hat{\beta}B - p - c$, the discounted benefit minus the cost. She therefore expects that she will use the gym if $c < \delta\hat{\beta}B - p$. From the perspective of period one the benefit from using the gym is $\delta\hat{\beta}(\delta B - p - c)$. So, if Maria signs the contract, she should expect net benefit

$$NB = \delta\hat{\beta}\left(-L + \int_0^{\delta\hat{\beta}\beta - p} (\delta B - p - c)dc\right). \tag{4.10}$$

The integral term captures her uncertainty about what the personal cost of using the gym could be.

The next thing we need to do is question what Maria will actually do in period two. She will use the gym only if $c < \delta\beta B - p$. If, therefore, $\hat{\beta} > \beta$ she is biased in overestimating her likelihood of going to the gym, and because of this overestimates the benefits of joining the gym. This is consistent with what we saw in Figure 4.7.

Now we can consider the firm. Suppose the firm pays production cost *K* whenever anyone joins and per-unit cost *a* if anyone uses the gym. The firm knows that a consumer will use the gym in period two if $c < \delta\beta B - p$. The objective of the firm is to choose L and p to maximize profits:

$$max_{L,P}\left(L - K + \int_0^{\delta\beta B - p} (p - a)dc\right).$$

Table 4.15 The consequences of optimal pricing in the gym example with present-biased preferences.

Preferences	Consumer	Firm
No present bias, $\beta = 1$	Correctly predicts when will use the gym	Makes a profit on membership fees Makes no profit when members use the gym
Present bias and sophisticated, $\beta = \beta$	Knows will possibly not use the gym because of present bias	Needs to charge a lower user fee to attract customers Makes a loss when people use the gym but this is compensated for by a higher lump sum fee Does not profit from consumers' present bias
Present bias and naïve, $\beta < 1$ and $\hat{\beta} = 1$	Does not realize will possibly not go to the gym because of a present bias	Charges a lower user fee to exploit customers Makes a loss when people use the gym but this is compensated for by a higher lump sum fee and some customers not using the gym Profits from consumers' naïvety

The one caveat is that the net benefit NB must be greater than zero, or no one would join. If you solve this maximization problem (see below) then the optimal price is

$$p = a - \delta B(1 - \hat{\beta}) - \delta B(\hat{\beta} - \beta). \tag{4.11}$$

I'll look at this in three stages, summarized in Table 4.15.

If customers have no present bias, $\beta = 1$, the firm should set the per-usage fee equal to the cost, $p = a$. This result should be familiar to those who have studied price discrimination in a microeconomics course. The optimal thing for the gym to do is charge as high a membership fee as customers will pay and then make zero profit from customer use.

If customers do have a present bias, $\beta < 1$, then the firm should set the per-usage fee below the cost, $p < a$. This means the gym makes a loss every time a customer uses the gym! Why would that make sense? It depends on whether the customer is sophisticated or naïve. If Maria is sophisticated, she knows that, because of her present bias, she may put off going to the gym. Thus, she wants something that makes it more likely she will go. A lower user fee serves that purpose. Maria would, therefore, be willing to pay a higher membership fee in order to have a lower usage fee. The gym does not profit from Maria having a present bias, but it does need to charge different prices.

If Maria is naïve then she may overestimate how much she will use the gym. The lower the usage fee, the more she thinks she will use the gym, and consequently the higher the membership fee she is willing to pay. In this case, the gym can profit from Maria's present bias. By offering a low usage fee it can charge a high membership fee and gain on people who join but subsequently do not use the gym. The greater the naïvety of consumers, the more profit the gym can make.

[Extra] To see how to get equation (4.11): if we set $NB = 0$, we can substitute in for L from equation (4.10) to rewrite the firm's problem as

$$max_p \left(\int_0^{\delta\hat{\beta}B - p} (\delta B - p - c)dc - K + \int_0^{\delta\beta B} (p - a)dc \right).$$

This can be rewritten

$$max_p\left(\int_0^{\delta\beta B-p}(\delta B-a-c)dc-K+\int_{\delta\beta B-p}^{\delta\hat{\beta}B-p}(\delta B-p-c)dc\right)$$

and

$$max_p\left((\delta\beta B-p)(\delta B-a)+\delta B(\hat{\beta}-\beta)(\delta B-p)-K-\int_0^{\delta\hat{\beta}B-p}cdc\right).$$

Evaluating the last integral and differentiating with respect to p gives the result!

There are two important lessons from this simple model. First, firms should take account of present bias when setting an optimal price. This is particularly pertinent if consumers are sophisticated. In this case, the firm cannot profit from the present bias but does need to price differently to make the same profits as without a present bias. Second, firms can profit by exploiting naïve consumers. They can charge a high membership fee and offer a low usage fee, knowing that some customers will not take advantage of the lower usage fee.

The data we have for gym membership would suggest that firms are aware of this. The user fee for members was zero, which must be below the gym's per-unit cost. We also did see underuse of the gym. Let's look at another example.

4.7.3 Choosing the correct calling plan

If the title of DellaVigna and Malmendier's paper, 'Paying not to go to the gym', gives away the punchline, the title of a paper by Miravete (2003), 'Choosing the wrong calling plan?', leaves things more open.

Miravete looked at consumers' choice of telephone calling plans following a tariff change by company South Central Bell in 1986. Customers had the choice between a flat-rate tariff of $18.70 per month and a measured tariff of $14.02 per month plus call charges. Such choices (and utility payments in general) are often suggested as fertile ground for present bias. That's because it may be best for a customer to switch tariff, and she may know that, but delay doing so because of present bias. This study aimed to put that anecdotal hunch to the test.

Key to the study was a survey of 5,000 customers that included questions on expected telephone usage. If we know how much customers expected to use the telephone then we can hope to identify customers who should have expected the other tariff to be less costly, but did not switch.

Around 30 percent of the consumers surveyed chose the measured option and 70 percent stayed with the flat rate. So, did consumers make the correct choice and, more importantly, did they switch if it looked as if they had initially made a wrong choice? Table 4.16 provides some relevant data to see what happened. There was a high proportion of people over- or underestimating the number of calls they would make. This should feed through to some making the wrong choice of calling plan. Of interest to us is how many of those who were in the wrong plan in October had switched by December to the better plan. Overall,

Table 4.16 Choices of calling plan, estimates of usage and the proportion of customers making the wrong choice.

Choice in October	**Flat**	**Flat**	**Measured**	**Measured**
Choice in December	**Flat**	**Measured**	**Flat**	**Measured**
Number of customers	953	43	41	375
Underestimated calls by 20% or more	26%	28%	32%	33%
Overestimated calls by 20% or more	59%	49%	61%	49%
Made wrong choice in October	11%	44%	100%	57%
Made wrong choice in December	6%	7%	0%	67%

Source: Miravete (2003).

around 40 percent of customers were in the wrong plan in October, falling to 33 percent in December. There is also evidence that those with the most to gain from switching were the ones to switch.

This study gives an interesting contrast to the one on gym membership. Recall that, with gym membership, it was those opting for the 'flat' monthly tariff that were most often making the wrong choice. In this telephone example those choosing the flat rate are on average making the right choice. It is those choosing the 'variable' per-use tariff that are most often making the wrong choice. It would be hard, therefore, to argue that people are on average biased towards a flat versus variable tariff, or vice versa. It just seems that people are bad at predicting future usage and therefore end up making the wrong choice.

Present bias can give a plausible explanation for all of this. People might overestimate how much they will use the gym and underestimate how much they will use the telephone because of time inconsistency. Furthermore, they may put off changing plan once they have realized it's wrong because of procrastination. Present bias gives us a good story, but is it the whole story?

Probably not. In the gym case, those who are paying too much do eventually change their membership. Similarly, in the telephone case the number of customers choosing the wrong tariff does fall over time. What appears to be time inconsistency and procrastination may, therefore, just be people taking time to learn their preferences and make the best choice. This is not to say that present bias is not important; it is just to highlight that other things are clearly going on as well. These other things will be a theme of the next two chapters, as we look at how people interpret new information and how they learn over time.

4.8 Environmental economics

Up until now we have focused on an individual making choices about her future. For example, when should Maria do the homework, or go to the cinema or visit the aunt she does not like? Many important economic decisions, however, involve decisions that will affect future generations. For example, if Maria builds up too much debt then her children, and grandchildren, may be expected to pay it off. In this section I want to briefly look at how we can understand such long-term trade-offs.

In doing so I shall focus on environmental protection. This is natural, given that the most pressing examples of long-term trade-offs concern the environment. Many scientists paint a depressing picture of the future unless we reduce or control carbon emissions. Similar fears surround many related issues, such as deforestation, overfishing, radioactive waste, pollution of rivers, etc. I will use, therefore, the example of protecting the environment to illustrate the more general ways that behavioral economics can inform debate on long-term trade-offs.

4.8.1 Inter-generational discount factor

Debate on climate change was blown open in 2006 by Nicholas Stern's *The Economics of Climate Change. The Stern Review*, commissioned by the UK government, advocated immediate action to reduce CO_2 and other greenhouse emissions. It warned of dire long-run consequences unless we acted *now*. This advice was in stark contrast to conventional economic thinking. Most economists advocated a **climate policy ramp** in which policies to slow global warming would increasingly ramp up over time. So, why did *The Stern Review* reach such a radical conclusion?

To help explain the issues, consider the entirely made-up numbers in Table 4.17. These show the utility of an average person now and after 50 to 200 years, depending on the climate change policy enacted. If we do nothing we risk catastrophe. The climate policy ramp advocates gradual action in the future. If we act immediately then we lower current utility for the sake of the future. Which policy should we choose?

Table 4.17 is similar to Tables 4.1 and 4.6, which detailed Maria's utility from doing her homework or going to the movies. The same techniques we used to analyze those choices can be used to decide on a climate change policy. But there is one fundamental difference. As Maria decides when to do her homework she is trading off her utility now for her utility in the future. We need to know her personal discount factor, present-day bias, preference for an increasing sequence, and so on. When society decides on a climate change policy we are trading off the utility of the current generation for the utility of future generations. This means we need some way of making inter-generational choices. We need a discount factor and preference over sequences that measure society's long-term desires (see Research methods 4.4).

Table 4.17 The long-run consequences of four climate change policies.

Policy	Utility				
	Now	**50 years**	**100 years**	**150 years**	**200 years**
Do nothing	100	100	0	0	0
Climate policy ramp	100	95	80	90	100
Immediate action	90	90	100	100	100

❖ RESEARCH METHODS 4.4 INTER-GENERATIONAL LAB EXPERIMENTS

I have already remarked that long-run trade-offs are hard to measure in a laboratory experiment (see Research methods 4.3). Looking at inter-generational trade-offs would seem even more impossible. Economists are an ingenious bunch, however, and inter-generational experiments have been run. Given that such experiments involve games that we are not going to look at until later in the book, I will not go through them in detail here, but the basic design can be illustrated with a study by Chaudhuri, Schotter and Sopher (2009).

Groups of eight subjects were brought into the lab and played the weakest link game for ten rounds (see Chapter 6 for more on this game). Once the ten rounds were up, each subject was replaced by a new subject, her laboratory descendant. The new group of eight subjects played the weakest link game for a further ten rounds. This continued for up to nine generations. When a subject handed over to her descendant she was allowed to leave a message. Depending on the treatment, the descendant was allowed to see this message and the recent history of her parent. The final thing to point out is that a subject was paid based on her own performance and the performance of her descendant. A subject thus had an incentive to send a good message.

Experiments such as this cannot tell us much about inter-generational discount factors. They can, however, tell us some interesting things about how advice is passed down and acted on from generation to generation. For example, subjects in the study tended to send pessimistic advice, which the next generation initially ignored but soon followed. The outcome was inefficiency. In terms of protecting the environment, this suggests that we need to be wary of a 'Don't bother, it's hopeless' message being passed down the generations.

There are three basic approaches to tackling this problem. The approach taken by *The Stern Review* sees deciding on an appropriate discount factor and preference over sequences as a philosophical and ethical issue. What we know about personal discount factors is deemed irrelevant, because they are person-specific. We need to think of future generations, and *The Stern Review* argues that one generation should not unduly advantage itself at the cost of future generations. It thus advocates a discount factor near one and a smoothing of utility over time. Indeed, it argues that the only reason the discount factor should differ from one is because of the threat of extinction; we need to take account of the fact that a meteor might wipe out humanity before 200 years are up.

Table 4.18 shows that, when we put in a discount factor of 0.999, immediate action is the optimal policy. The high discount factor means we should make immediate sacrifice for future generations. It did not take long, however, for economists to criticize such a high discount factor.

An alternative approach is based on revealed preference. This approach basically says that we should look to personal discount factors when deciding policy because

Table 4.18 The discount factor for which each policy is optimal, and the inter-temporal utility of each policy for different discount factors.

Policy	δ when optimal	Inter-temporal utility			
		δ = 0.999	0.95	0.9745	Uncertain
Do nothing	0.945 or less	195.1	107.7	127.5	151.4
Climate policy ramp	0.945 to 0.996	422.1	107.8	134.6	264.9
Immediate action	0.996 to 1	434.0	97.6	124.9	265.8

these reveal how much we actually care about future generations. If Maria does not save for the future then she reveals that she does not care too much about her children and grandchildren! Moreover, this approach argues that decisions to invest in climate change abatement technology should be based on current rates of return.

Table 4.18 shows that, if we put in a discount factor of 0.95, then the climate policy ramp is optimal. If the discount factor is any lower then doing nothing is optimal. The radical conclusions of *The Stern Review* follow, therefore, from the relatively high weight put on the future. Crisis over!

Well, the material in this chapter should leave you somewhat skeptical of the revealed preference approach. We have seen that personal discount factors are highly context-dependent. We have also seen that people are time-inconsistent and happy to pre-commit. They borrow at high interest rates while saving at low interest rates and buy gym memberships they do not need. Against this backdrop, it seems hard to argue that current rates of return can be taken all that seriously when deciding on the optimal inter-generational discount factor. Indeed, the climate policy ramp looks remarkably like present bias.

Fortunately, there is a third approach that can help reconcile views. This approach takes as its starting point the huge uncertainty about climate change. If we do not know what discount factor to choose then let us build that into our thinking. This might sound like fudging the issue, but it turns out not to. To see why, let us say the appropriate discount factor is either 0.999 or 0.95. We put 50 percent probability on each possibility. At first glance this may seem equivalent to a discount factor halfway between 0.999 and 0.95, but Table 4.17 makes clear that it is not. If we are uncertain about the discount factor, then immediate action is the optimal policy. Basically, some probability that we should use a high discount factor is enough to advocate immediate action.

All three approaches summarized above have their critics, and it is up to you to decide which approach makes most sense to you. The point I would emphasize is that behavioral economics can inform your opinion, particularly about the reliability of using revealed preference. I want to show next that behavioral economics can also help us tackle climate change and reduce CO_2 emissions.

4.8.2 Reducing CO_2 emissions

Just about all economists agree that climate change requires action at some point in the future. That will require people to change their behavior. We need to pollute less, consume less, and so on. The standard economic model suggests that a change

in behavior will come about only when there are changes in prices. Attention has, accordingly, focused on things such as a carbon tax or subsidies for renewable energy. Behavioral economics points to many ways in which behavior can be changed without disturbing economic fundamentals, however. I will look at two examples to illustrate the point. In Chapter 11 I shall look at the general approach to behavior change in more detail.

Suppose that Maria is persuaded by the need to protect the environment and is committed to reducing her CO_2 emissions. She needs a new car and decides to buy an environmentally friendly one. Suppose that she drives 10,000 miles per year for work, and this cannot be changed. The standard model of car she would like to buy does 15 miles per gallon and costs $20,000. It is possible to buy alternative models that differ only in fuel efficiency, ranging from 20 miles per gallon up to 60 miles per gallon. How much should Maria pay for greater fuel efficiency?

Note that this question is partly about saving money. By buying a more fuel-efficient car Maria not only helps the environment but also can save money on fuel. It is potentially a win-win situation in which no sacrifice is needed in order to reduce CO_2 emissions. But this is going to happen only if Maria makes the right choice. A study by Larrick and Soll (2008) put subjects in the position of Maria and asked them how much they would pay for the fuel savings. Figure 4.9 summarizes the results. We can see that the subjects' willingness to pay followed a linear pattern. They were willing to pay as much to reduce fuel consumption from 45 to 55 miles per gallon as they were to reduce consumption from 25 to 35 miles per gallon. This is a basic misunderstanding of how the miles-per-gallon ratio works, however.

To see things more clearly, it is convenient to convert miles per gallon into gallons per 100 miles. A change from 25 to 35 miles per gallon means a reduction from 4 to 2.86 gallons per 100 miles. A change from 45 to 55 miles per gallon means a reduction from only 2.22 to 1.81 gallons per 100 miles. Remember that Maria was going to do 10,000 miles per year. So, buying the model that does 35 rather than 25 miles per gallon saves her 114 gallons of fuel. Buying the model that does 55 instead of 45 miles per gallon only saves 40 gallons of fuel. The linear willingness to pay schedule that we see in Figure 4.9 overestimates the benefit of going from 45 to 55 miles per gallon, while it underestimates the benefit of going from 25 to 35.

The bias we see here results from the use of simple heuristics. The framing in terms of miles per gallon makes most people think that any saving in ten miles per gallon is equivalent. If things are reframed in terms of gallons per 100 miles then people will be far less biased. The message, therefore, is simple. If we want people to make informed decisions we should use the amount of fuel per unit of output (as Europe typically does) rather than the amount of output per unit of fuel (as the United States and United Kingdom typically do).

Having bought a fuel-efficient car, Maria now looks to cut down on her domestic energy usage. Her local utility company starts sending her Home Energy Report letters from a company called OPOWER. The report compares her energy use to that of 100 similar neighbors. She is rated as 'Great', 'Good' or 'Below average' and offered tips on energy reduction. Will this motivate her, or inform her, to reduce energy consumption?

Allcott (2011) evaluated the consequences of sending such letters to over half a million energy customers across the United States. In households that received the letters energy consumption was reduced by around 2 percent. That may not sound

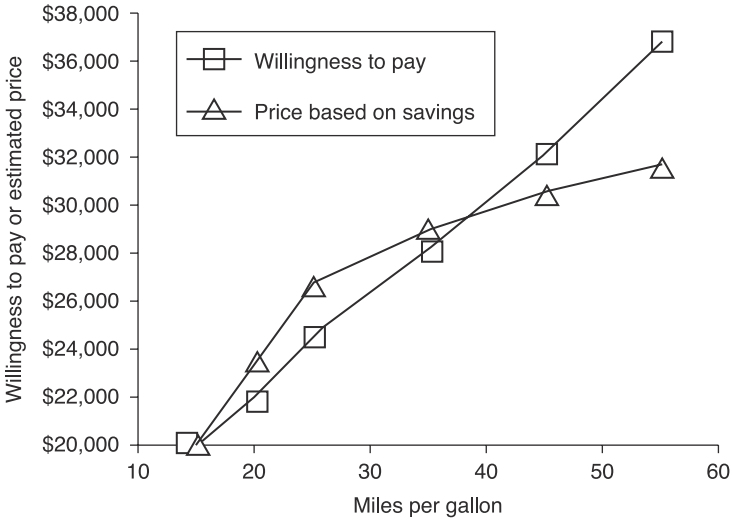

Figure 4.9 Willingness to pay for a more fuel-efficient car and the estimated price taking into account predicted fuel savings.
Source: Larrick and Soll (2008).

much, but it is a large reduction. For example, we can estimate that a price rise of around 11 to 20 percent would be needed to bring about such a fall in usage. Most of the reduction was caused by day-to-day changes such as turning off lights and adjusting thermostats. Again, we see that Maria can help the environment and save money at the same time.

These two examples illustrate once more the importance of framing and context effects. The new thing I want to emphasize is how such effects are amplified by procrastination. When it comes to reducing gas emissions, procrastination is potentially a very dangerous thing. If Maria, and everyone else, continually puts off protecting the environment until tomorrow, then it may not be long before there is no environment left to protect. This point makes clear that, for all the talk of *The Stern Review*, climate policy ramps and inter-generational discount factors, gas emissions will be reduced only if individuals change behavior. Procrastination is a major reason that people will not change behavior, even if they 'want' to.

Climate change policy, therefore, needs to look for ways of overcoming procrastination. A change in the way decisions are framed may help do that. For example, Maria may partly put off getting a fuel-efficient car because she underestimates the gains she would make from doing so. Similarly, receiving a letter with tips on energy reduction may help overcome her procrastination on reducing home energy usage. Small changes in framing can have big, long-run consequences!

4.9 Further reading

The survey papers by Frederick, Loewenstein and O'Donoghue (2002) and Cohen and co-authors (2016) are a good place to start. The pieces by O'Donoghue and

Rabin (1999, 2000, 2006) are recommended reading on time inconsistency, as is the pioneering article by Strotz (1956). Note that we will return to time inconsistency in Chapter 11, and so much of the further reading there is also relevant here. For an interesting look at the desire for wages to increase over time, see Frank and Hutchens (1993). For more on applying behavioral economics in industrial organization, see Ellison (2006) and Spiegler (2014). For more on behavioral environmental economics, see Shogren and Taylor (2008), Allcott and Mullainathan (2010) and Croson and Treich (2014).

4.10 Online materials

On the companion website you will find more details on Maria's homework example and Maria's movie example. You will also have access to flashcards, quizzes and answers to the review questions.

4.11 Review questions

4.1 Inter-temporal utility (as defined in equation (4.1)) is about measuring streams of utility over time. To make life easier for themselves, however, economists normally think about streams of money over time. This should be OK if we think in terms of, say, '$10's worth of utility', but could be problematic if we really mean 'a $10 note'. To illustrate, if someone says they are indifferent between $150 in one year's time and $100 today, what is the discount factor, when the utility function is $u(x) = x$ and $u(x) = \sqrt{x}$ and x is money? Think about the implications of this.

4.2 Explain the difference between time-consistent and present-biased preferences. Why is this distinction not important in a model of exponential discounting?

4.3 In section 4.3 I showed that $\delta^{speed\text{-}up} > \delta^{delay}$ in the case of a receipt. Show that $\delta^{speed\text{-}up} < \delta^{delay}$ in the case of a payment.

4.4 Looking back at Table 4.10, what would happen if Maria's reference point was to do her homework on Saturday, and what if it was to do it on Monday? How is this related to the concept of personal equilibrium that we looked at in the last chapter?

4.5 Looking back over this chapter and the previous two, come up with examples of context effects.

4.6 What is the expression for $D(t)$ with quasi-hyperbolic discounting?

4.7 What context effects do you think make it more likely that someone will think of a series of events as a sequence rather than separated?

4.8 What are the implications of habit formation for the life cycle hypothesis?

4.9 Carefully explain why risk aversion can make a person seem impatient.

4.10 If a person is risk-loving, will they appear more patient?

4.11 What evidence do we have that people do not maximize discounted expected utility?

4.12 What is the relationship between a model of habit formation and reference-dependent utility?

4.13 Discuss reasons why a person may simultaneously be borrowing money and putting money into savings.
4.14 Why do many firms, such as gyms, charge a two-part tariff with a zero user fee?
4.15 Is there a difference between firms exploiting a customer's bias and changing their strategy to take account of a customer's bias?
4.16 What strategy can firms use to take account of other cognitive biases and heuristics that we have looked at, such as loss aversion and risk aversion?
4.17 How can people overestimating how much they will use the gym and underestimating how much they will use the telephone both be caused by time inconsistency?
4.18 Compare the three approaches, looked at above, for deriving an intergenerational discount factor. Why is it so difficult to agree on the best approach?
4.19 To what extent do you think that present bias is a factor in combating climate change?

Chapter 5

Learning from new information

> He is accelerating all the time. That last lap was run in 64 seconds compared to the one before in 62 seconds.
>
> David Coleman, sports commentator

A theme in Chapters 2 and 3 was uncertainty and risk. For instance, in Chapter 2 I talked about Anna not knowing whether she would prefer the Nutty breakfast cereal or the Honey cereal. In Chapter 3 I talked about Alan deciding whether or not to fully insure his car without knowing whether his car would be stolen. New information should remove some of this uncertainty. For example, if Anna tastes the breakfast cereals, or Alan finds out the number of cars stolen in his area in recent years, then they should be able to make more informed decisions. This looks like a good thing.

It should, therefore, be a good thing that we are continually bombarded with new information from the TV, radio, internet, newspapers, friends, family, colleagues and our own experiences. New information is useful, however, only if it is used wisely. There can be no doubt that new information influences behavior in significant ways. Why else would firms spend so much money on advertisements? Why else would someone ask friends, or look on the internet, for advice before making a decision? What we need to think about is how wisely people use new information, and the consequences this has for how they behave. That will be the focus of this chapter.

5.1 Bayesian updating and choice with uncertainty

Imagine someone called John, who is thinking of buying a new car and has narrowed it down to two choices, a 'Sporty' car or a 'Comfort' car. Either car is approximately the same price, but John does not know which will give him higher utility. In the process of deciding, he tries to become better informed by asking his friends, doing a test drive, searching for advice on the internet, and so on. What should he do with this new information?

The benchmark for what he should be doing is **Bayesian updating**. Let p denote the **prior probability** that John puts on the Sporty giving the higher utility. Any new information he gets we call a **signal**. For example, a signal might be his friend saying 'Buy the Comfort' or him doing a test drive and 'liking the Sporty'. Let θ denote the **signal precision** or probability that he would have got the signal if Sporty does give the higher utility. The value of θ should capture things such as the trust that John puts

in his friend's opinion or expertise, or how relevant he thinks a test drive is. If John uses **Bayes' rule** then the **posterior probability** that he should put on the Sporty giving the higher utility is

$$\Pr(\text{Sporty is better} \mid \text{signal}) = \frac{\theta p}{\theta p + (1-\theta)(1-p)}. \tag{5.1}$$

For example, if John started relatively indifferent between the two, say, $p = 0.4$, but his friend said that Sporty is better and he trusts his friend's opinion, say, $\theta = 0.9$, then he should put posterior probability 0.86 on the Sporty being better for him. This new information changes his beliefs, and in doing so may change his behavior from buying the Comfort to buying the Sporty.

Bayesian updating is what John should do if he wants to make the most informed decision. For every new piece of information he gets, it makes sense to use Bayes' rule. We have, therefore, a prediction for how John should use any new information. Bayes updating can, however, be a bit more complicated than equation (5.1) might suggest. To explain why, it is useful to relate John's choice to what we did in Chapters 2 and 3.

We can think of John as not knowing what car will give him higher utility, and so he searches for more information to become better informed. He is never going to be completely sure what is better, but new information allows him to update his beliefs and eventually make a more informed decision. Throughout this process John is making risky choices, particularly when he finally buys a car. For example, we can think of there being a 'Sporty prospect' and a 'Comfort prospect', with each prospect detailing the possible outcomes from buying that car. New information allows John to change the numbers in one or more of the prospects. For example, if he test-drives the Comfort he can update the Comfort prospect. Once he has updated the Comfort prospect he can then update his beliefs that Comfort is better. Table 5.1 illustrates how this might work in practice if John test-drives the Comfort.

Framing things in this way, we can see two potential issues. First, it becomes complicated to update beliefs. When I introduced Bayes updating, the focus was on the probability that the Sporty is better for John. Now we are starting to get down to finer details about each car. This seems natural but means that John has some work to do in order to turn new information into updated beliefs. He has to update his beliefs about the merits of each car, and then update his belief about which car is better. This can easily get quite complicated. Can John realistically be expected to do it? We will get back to that question in section 5.2.

A second issue is that of ambiguity and uncertainty. At the start of Chapter 3 I mentioned the distinction between risk and uncertainty, and you might recall that if a person has a choice with uncertainty then they may not know the probability of possible outcomes. John can presumably predict that he could like the Comfort, not like the Comfort, and so on. What would be much harder for him to do is to put specific numbers on, say, the probability that he will like the Comfort, and the value he would put on it if he did like it. This means that John's choice of car is a choice with uncertainty. It also means we need to think about where the numbers in Table 5.1 could come from, because without them we cannot possibly expect John to use Bayesian updating. We need, therefore, to think about choice with uncertainty.

Table 5.1 John updates his beliefs after a test drive. Comfort is better if it is worth $12,000. Initially he puts probability 0.7 on this being the case, so Sporty is better with probability 0.3. He test-drives Comfort and does not like it. This lowers the probability that he puts on Comfort being worth $12,000 down to 0.6.

Event	John's opinion	Summary
Initial prior	Sporty is worth $10,000 with probability 0.9 and $9,000 with probability 0.1 Comfort is worth $12,000 with probability 0.7 and $4,000 with probability 0.3	Sporty is better with probability 0.3
Test-drives Comfort and gets signal 'Did not like Comfort'	The signal precision (i.e. probability that this would happen if Comfort is worth $12,000) is 0.39 He updates his beliefs about Comfort to being worth $12,000 with probability 0.6 and $4,000 with probability 0.4	Sporty is better with probability 0.4

5.1.1 Models of choice with uncertainty

I shall start by talking about a relatively simple example to illustrate the most important issues, and then relate things back to John buying a car. Imagine a box containing ten balls. One of the balls is going to be drawn randomly from the box, and you have to bet on what color the ball will be. If you get it correct, you win $2. Now contrast these two different possibilities.

Box 1: Contains five red balls and five black balls.
Box 2: Contains an unknown number of red balls and black balls. It could be zero red and ten black, ten red and zero black, or anything in between. Which box would you rather bet on?

Box 1 provides a simple choice with risk like that we looked at in Chapter 3. If John bets on black, he faces the prospect (0.5, $2; 0.5, $0), so with a 50 percent chance he will be correct and win the $2. His expected utility is then

$$U(\text{black}) = 0.5u(w + 2) + 0.5u(w),$$

where w is his wealth. What about box 2? This time I cannot write down the prospect, because I cannot say what probability John has of winning if he bets on black. The probabilities are **ambiguous**. If the probabilities are ambiguous then we cannot work out John's expected utility, use prospect theory or do anything else we did in Chapter 3. We need something new.

The simplest thing we can do is use **subjective expected utility**, in which we go with what John thinks will happen. If John thinks that box 2 will contain four red balls and six black balls then his subjective probability of drawing a black ball is 0.6 and his subjective expected utility is

$$U(\text{black}) = 0.6u(w + 2) + 0.4u(w).$$

That looks like a simple solution! The problem is that John could have any beliefs he wanted and not be wrong. He could believe that none of the balls are black, all of

them are black, and we have no way to tie down what his beliefs should be or are likely to be.

Consider now another possible box.

Box 3: A number between zero and ten will be randomly drawn and determine the number of red balls in the box. The rest of the ten balls will be black.

Does this box provide a choice with risk or with uncertainty? The correct answer is that it provides a choice with risk. It does look different from box 1, however. That's because it involves a compound lottery. A **compound lottery** is when there are two or more consecutive lotteries to determine the outcome. Box 3 is a compound lottery because there is a lottery to determine how many balls are in the box, and then a lottery when John comes to pick a ball. Just to confuse, the final stage, when John picks a ball, is typically called the **first-order lottery** or **second-stage lottery**. The preliminary stage, when the number of balls in the box is determined, is called the **second-order lottery** or **first-stage lottery**. Two stages are enough in this case, but there can be three or more stages in a compound lottery.

To try and capture the lottery of box 3 we could write down one long prospect. For instance, if John chooses black, he faces prospect (0.09, $2; 0.09, (0.9, $2; 0.1; $0); 0.09, (0.8, $2; 0.2, $2); …), where we recognize that there is a one-in-11 chance that all the balls are black and John wins $2, a one-in-11 chance that nine of the balls are black and John wins $2 with probability 0.9, and so on. This looks a bit messy. We could, therefore, try **reducing the compound lottery** to a simple lottery. To do this, we work out the overall probability that a ball will be black or red. If you do this, you should find that there is a 50 percent chance a ball will be black and a 50 percent chance it will be red. This suggests betting on black with box 3 gives the same prospect (0.5, $2; 0.5, $0) as it did with box 1!

It is not clear, however, that we should be reducing compound lotteries to a simple lottery. For instance, we could do something called **recursive expected utility** and work out the expected utility of each lottery and combine them. To do this, we start with the first-order lottery, when John picks a color.

If there is one black ball and nine red balls, his expected utility of betting on black is

$$U(\text{black} \mid \text{one black}) = 0.1u(w + 2) + 0.9u(w).$$

We need to do this for all possible combinations of black and red balls. We then go to the second-order lottery, when the number of balls is determined. For reasons we do not need to worry about, Kilbanoff, Marinacci and Mukerji (2005) show that his expected utility can be written

$$U(\text{black}) = 0.09\varphi\,(U(\text{black} \mid \text{no black})) + 0.09\varphi\,(U(\text{black} \mid \text{one black})) + \ldots,$$

where φ is some function that captures preferences over second-order lotteries.

The key to this way of doing things is that we may want to distinguish risk attitudes towards first- and second-order lotteries. In particular, we shall see that there are good arguments for why John may be more risk-averse in a second-order lottery than a first-order lottery. If so, we should not be reducing compound lotteries but,

instead, using recursive expected utility. It also means that John will be more reluctant to bet on box 3 than box 1, because, essentially, he does not like compound lotteries.

It is now interesting to go back and look again at box 2. When the probabilities are ambiguous, as they are in box 2, it seems more natural to think in terms of a compound lottery. That is, it seems natural to think of John having subjective beliefs about an additional stage in which balls are added to the box. For example, he might put subjective probability 0.5 on there being five red and five black balls, and probability 0.25 that all the balls are black, and so on. It may, therefore, be most apt to use **recursive subjective expected utility** to think about choice with uncertainty. To see whether that is the case we need to look at some data.

5.1.2 The Ellsberg paradox

I am going to look at the results of a study by Halevy (2007). In the study, subjects were asked to bet on boxes 1 to 3, and a fourth box.

Box 4: A fair coin will be tossed to determine whether all the balls in the box are black, or all of them are red.

The interesting thing about box 4 is that all risk is resolved in the first stage. So, we have box 1, for which risk is resolved in the second stage; box 4, for which it is resolved in the first stage; and box 3, which is a true compound lottery. If people are more risk-averse about second-order risk, we should be able to pick this up.

In the study, rather than just have subjects bet on a color, subjects were given the chance to sell bets by saying the amount they were willing to accept to forgo the bet. Figure 5.1 plots the average, minimum amount subjects were willing to accept. Remember they would win $2 if they won the bet, so a risk-neutral person would want $1. The most important thing to pick out is the greater willingness to sell a bet

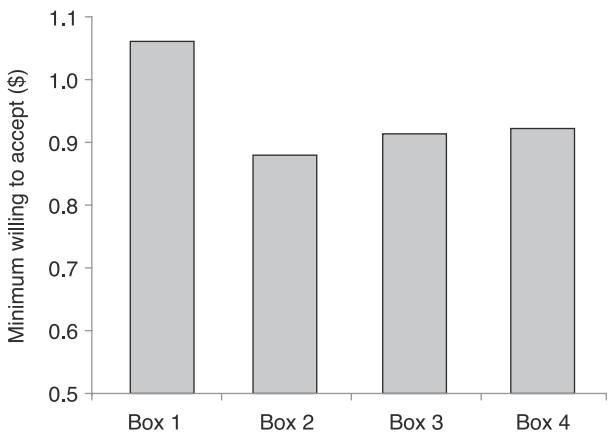

Figure 5.1 The willingness to sell a bet on each of the boxes. The Ellsberg paradox is why subjects want less to sell a bet on box 2 than on box 1.
Source: Halevy (2007).

Table 5.2 The proportion of subjects fitting each model of preferences and the average willingness to sell.

Model of preferences	Proportion of subjects (%)	Willingness to accept forgoing bet			
		Box 1	Box 2	Box 3	Box 4
Subjective expected utility	19.2	1.03	1.03	1.03	1.03
Recursive expected utility	40.3	1.04	0.78	0.96	0.79
Recursive non-expected utility	39.4	1.11	0.91	0.85	1.07
Maxmin expected utility	0.1	0.90	0.80	0.90	0.90

Source: Halevy (2007).

on box 2 than a bet on box 1. This difference was first pointed out by Daniel Ellsberg (1961), and so goes by the name of the **Ellsberg paradox**. (For another Ellsberg paradox, see Review question 5.2.) The paradox is why people do not like betting on box 2. It seems that people are **ambiguity-averse** in preferring a choice with risk to one with uncertainty. The results for boxes 3 and 4 suggest, however, that we need to qualify this somewhat, because subjects also seem averse to a second-order lottery.

To delve a little deeper, Halevy looks at the choices of individual subjects and tries to classify those choices. Table 5.2 summarizes his findings. Let's go through the table row by row.

First, there were subjects willing to sell a bet on any box at the same price. They did not seem to be bothered by ambiguity or a second-order lottery, and so their choices seem consistent with subjective expected utility maximization.

Second are those who were more willing to sell a bet on box 4 than on box 3 than on box 1. These subjects seemed to be risk-averse about a second-order lottery, and so their choices are most consistent with recursive expected utility.

Third are those who were willing to sell a bet on boxes 1 and 4 at the same price but more willing to sell a bet on box 3. It does not look as though these subjects were risk-averse about a second-order lottery, otherwise they would be more willing to sell a bet on box 4. Instead, it seems they disliked the compounding of lotteries with box 3. I will not go into the details, but this is consistent with recursive non-expected utility, in which a person uses rank-dependent utility (as introduced in Chapter 3).

Finally, we come to those who simply do not like ambiguity. In his original paper, Ellsberg suggested that, when there is ambiguity, people are pessimistic and go with the worst thing that can happen. Hence we get a story of maximizing the minimum. Only one subject in this study behaved like this, in being more willing to sell box 2, but not boxes 3 or 4.

Let's go back now and think about John deciding what car to buy. First of all, we now know that we should view the numbers in Table 5.1 as John's subjective beliefs about what he thinks might happen. We can also now see that it is natural to think of his choice as a compound lottery in which, say, the quality of each car is determined in a second-order lottery and the car he would prefer is determined in a first-order lottery.

The subjective nature of John's beliefs means it will be very hard for us to know why John would choose one car over another and whether he made the better choice. What we can see, however, in Table 5.2 is that the majority of people do

not seem to like compound lotteries, or are ambiguity-averse. This does have clear implications: John may be reluctant to gamble on which car is better, and so he may put off buying a car; he should also be eager to become better informed about each car in order to reduce the ambiguity or second-order risk that he faces. It is this last possibility I want to focus on first.

5.2 Two cognitive biases

If John does not like ambiguity or second-order risk then he needs to become better informed, and so he should be on the lookout for new information. What we need to ask is whether he will use that information wisely. Will he, for example, use Bayes' rule? To do so would be quite complicated, and so it would not be a huge surprise to see some mistakes and biases in how he does interpret new information. In this section I will look at two of the more common biases that we seem to observe: the confirmatory bias and the law of small numbers.

5.2.1 Confirmatory bias

Suppose John asks his friend's advice and is told: 'The Comfort is a great car, but the Sporty is more eco-friendly.' This may be useful information, but it is not obvious how John should interpret it. The basic problem is that the information is ambiguous, in the sense that any interpretation of it will be subjective, and different people might interpret it in different ways. Some might focus on '[t]he Comfort is a great car' and others on 'the Sporty is more eco-friendly'.

This ambiguity is common in a lot of information we get and can easily lead to a biased interpretation of information. We say that there is **confirmatory bias** if a person has a tendency to interpret ambiguous, new information as consistent with his initial beliefs and opinions. Two psychology studies will help illustrate.

In a study by Lord, Ross and Lepper (1979), subjects were asked to read two reports on the deterrent effect of the death penalty and then asked to use their own 'evaluative powers' in thinking whether the reports provided support for or against the death penalty. They were initially given a brief description of the first report and asked questions about their attitudes and change in attitudes having seen the report. They were then given a more detailed description and asked questions about how well they thought the research in the report had been conducted and how convincing it was. The whole process was then repeated for the second report.

Some weeks before doing this experiment subjects had filled in a questionnaire in which half the subjects invited to the experiment had appeared proponents of capital punishment and the other half opponents. The reports given to subjects could also be classified as having a pro- or anti-deterrence conclusion.

The main results of the experiment are summarized in Table 5.3. We see that brief descriptions of a study changed attitudes in the way we would expect. After a detailed description, however, attitudes diverged. Proponents of the death penalty rated pro-deterrence studies as convincing and well conducted, while anti-deterrence studies were not thought convincing or well conducted. Vice versa for opponents.

Table 5.3 Evidence of confirmatory bias in attitudes to capital punishment. Mean evaluations on a –8 to +8 scale, where +8 means agree with capital punishment or convincing and well done and –8 means disagree with capital punishment or not convincing and not well done.

Question	Report	Evaluation	
		Proponents	Opponents
Brief description changes my attitude to capital punishment	Pro-deterrence	1.3	0.4
	Anti-deterrence	–0.7	–0.9
How convincing the research was	Pro-deterrence	1.4	–2.1
	Anti-deterrence	–1.8	0.1
How well the research was conducted	Pro-deterrence	1.5	–2.1
	Anti-deterrence	–1.6	–0.3
Detailed description changes my attitude to capital punishment	Pro-deterrence	0.8	–0.9
	Anti-deterrence	0.7	–0.8

Source: Lord, Ross and Lepper (1979).

Table 5.4 Evidence of confirmatory bias in assessing the ability of a child. Predicted grade placement goes down after watching a video if the child was perceived to be from a low-income background, and goes up if she was perceived to be from a middle-class background.

Watched the video or not?	Information provided	Grade level placement		
		In mathematics	In reading	In liberal arts
Did not	Low-income	3.98	3.90	3.85
	Middle-class	4.30	4.29	4.03
Did	Low-income	3.79	3.71	3.04
	Middle-class	4.83	4.67	4.10

Source: Darley and Gross (1983).

Thus, subjects had a tendency to rate positively evidence that was consistent with their opinion, and rate negatively evidence inconsistent with their opinion! This is **confirmatory bias**. The consequence is a **polarization of attitudes**, with proponents of the death penalty becoming more in favor and opponents less in favor, despite having read the same evidence.

In the second study, by Darley and Gross (1983), subjects were asked to evaluate the academic capabilities of a fourth-grade child. All subjects were given some basic information about the child. Half the subjects were given the impression that she was from a low-income background and the other half that she was from a middle-class background. Half the subjects then watched a video of the child answering questions. The video was designed to be relatively uninformative about the child's ability. Subjects were finally asked to evaluate the child's capabilities.

The results of the experiment are summarized in Table 5.4. We see that those who did not watch the video and believed she was from a middle-class background had a higher expectation of ability than those who believed she was from a low-income background. Whether or not this might be justifiable is not our main concern.

Of interest to us is that watching the video lowered perceptions of ability if the child was believed to be from a low-income background but increased perceptions of ability if she was believed to be from a middle-class background. How subjects interpreted the same video depended, therefore, on initial information about the child. Those primed to think she might be low ability had their belief confirmed by the video, and those primed to think she might be high ability had this different belief confirmed by the same video! This is confirmatory bias.

In neither of these studies were subjects asked to make economic choices. Clearly, however, confirmatory bias can potentially influence economic behavior. If John was favoring the Comfort then confirmatory bias suggests the main thing he will take from his friend's advice is '[t]he Comfort is a great car'. This could mean that he is overconfident when making his choice. To explore this further it is useful to have a model, and for that I will turn to the work of Rabin and Schrag (1999).

5.2.2 A model of confirmatory bias

Imagine that every week John reads a magazine or looks up some review on the internet and gets a signal of which car is better. The signals are correct with probability $\theta \in (0.5,1)$. So, if the Comfort is better for John, then with probability θ a review will suggest the Comfort is better, and with probability $1-\theta$ it will suggest the Sporty is better.

To capture confirmatory bias we assume that John does not always correctly interpret signals. To see how this works, suppose that he currently believes the Comfort is better. If he reads a review that suggests the Comfort is better, he always correctly interprets the review as supportive of the Comfort. But, if he reads a review that suggests the Sporty is better, he incorrectly interprets the review as supportive of the Comfort with probability q. This captures confirmatory bias, because John has a bias to interpret information as consistent with his prior belief. If he thinks that the Comfort is better, he is likely to think a review is supportive of the Comfort even though it is not. Similarly, if his prior belief is that the Sporty is better, he is likely to think a review is supportive of the Sporty even though it is not. The larger is q, the larger the bias. If $q = 0$ there is no bias, because all signals are correctly interpreted. If $q = 1$ then there is complete bias, because whatever John reads he will think is supportive of what he believes.

Suppose John uses Bayes' rule based on the signals he thinks he has received. If he started indifferent between the two cars, and after t weeks thinks he has read S_t reviews supportive of Sporty and C_t reviews supportive of Comfort, then his posterior belief should be

$$prob(\text{Sporty is best}\,|\,S_t, C_t) = \frac{\theta^{S_t}(1-\theta)^{C_t}}{\theta^{S_t}(1-\theta)^{C_t} + (1-\theta)^{S_t}\theta^{C_t}}.$$

What are the consequences of confirmatory bias?

We can first look at whether it leads to John having wrong beliefs more often. By 'wrong beliefs' I mean John thinking one car is better when actually the other car is better. Table 5.5 details the proportion of times he has the wrong belief after reading ten or 100 reviews. Remember that if $q = 0$ there is no confirmatory bias, and so we

Table 5.5 The probability of an incorrect belief (%) after reading ten or 100 reviews. The higher the confirmatory bias q, the higher the probability that John has wrong beliefs, and the less difference it makes if he reads more reviews.

q	θ			
	0.6	**0.7**	**0.8**	**0.9**
0	25,2	7,0	2,0	0,0
0.1	25,5	10,0	2,0	1,0
0.3	32,27	13,5	6,0	2,0
0.5	35,34	20,18	9,7	5,3
0.7	36,36	24,24	15,15	7,7
0.9	40,40	27,27	19,19	10,10

see the unbiased probability of John being wrong. Confirmatory bias has two consequences: (1) John has wrong beliefs more often, as seen by the larger proportion of errors as q increases; (2) he does not correct these beliefs over time, as seen by the error rates being similar after reading 100 reviews and ten reviews, the higher is q.

We can see that these two effects combine to make wrong beliefs much more likely. To get some intuition for why this is the case, suppose that the Sporty is the better car for John and he receives signals C,C,S,C,S,S,S, so the first two reviews are supportive of the Comfort, the next of the Sporty, and so on. If John has no confirmatory bias he will correctly think that three reports are supportive of the Comfort and four of the Sporty, and four beats three, so – overall – the Sporty looks better. If he has confirmatory bias, however, he may become biased by the first few reports that are supportive of the Comfort. Specifically, he will correctly think the first, second and fourth reviews were supportive of the Comfort. So, he needs only to wrongly think one of the third, fifth, sixth or seventh reviews was supportive of the Comfort and he will have wrong beliefs. The higher is q, the more this becomes likely.

Confirmatory bias can, therefore, lead John astray. This means we might want to be careful of anything he tells us. Suppose, for instance, John tells you that 'I think the Comfort is best because 11 of the 15 reviews I have read are supportive of the Comfort'. What should you believe if you suspect confirmatory bias? It is possible to show that you should believe that John is correct but be less confident than he is. For instance, if he says that 11 of the 15 reviews he has read are supportive of the Comfort then it would be sensible for you to think that around nine or ten were actually supportive. This is because the confirmatory bias will, on average, make him overestimate the number of times he read a supportive review.

Next, suppose John tells you: 'The Comfort is better, but I have recently read some reviews supportive of the Sporty,' or he says: 'The Comfort is better, but I used to think the Sporty was.' In the first case you should probably put less faith in his belief, because there are signs that confirmatory bias is causing him to overestimate Comfort. Indeed, if there is a strong confirmatory bias, the fact that he has recently read some reviews supportive of the Sporty might imply you should think the Sporty is better. In the second case, we have cause to be confident in John's belief. That he recently changed his mind, despite confirmatory bias, suggests you should be more confident than John that the Sporty is better.

Things are now starting to get a bit messy and complicated. Indeed, we now see that John should be thinking about the confirmatory bias of the people writing the reviews that he is reading! The main point, however, is simple: confirmatory bias means people may not correctly interpret new information, and so we may want to be skeptical of what they say they have learnt.

5.2.3 Law of small numbers

Suppose (continuing the previous thread) John tells you that 75 percent of the reviews he has read are supportive of the Comfort. If you do not suspect confirmatory bias, how should you interpret this information?

A Bayesian updater wouldn't do anything without asking first how big the sample size was. That three out of four reviews were supportive of the Comfort should be much less persuasive than hearing that 15 out of 20 reviews were supportive. For instance, if $\theta = 0.6$, a Bayesian updater would put probability 0.69 on the Comfort being better for John in the first case, and 0.98 in the latter.

There is evidence, however, that many people fail to recognize the importance of sample size. They expect small samples to be more representative of the population than they likely are. This is the **law of small numbers**. There are two sides to the law of small numbers that we need to look at. I'll start with something called the gambler's fallacy.

If John expects a small sample to be representative of the population then he may expect random events to be self-correcting. For example, if ten tosses of a coin have been heads, there may be a tendency to think the next one 'must be tails'. This is known as the **gambler's fallacy**. We see interesting evidence for the gambler's fallacy in betting on lottery numbers. To illustrate, I will look at studies by Clotfelter and Cook (1991) and Terrell (1994), which analyzed data from the state lotteries of Maryland and New Jersey. In both lotteries a better must correctly guess a three-digit winning number. In Maryland any winner gets $500. In New Jersey a prize fund is split between all those who guess the right number, meaning that a person wins more the fewer guess the winning number.

In both Maryland and New Jersey there was clear evidence of people betting less on a number that had recently won. This is consistent with the gambler's fallacy: 'If number 525 has just won, it cannot win again.' In Maryland the gambler's fallacy does not affect expected winnings because all numbers are equally likely and every winner gets $500. In New Jersey, however, the gambler's fallacy does matter, because winnings are higher the fewer pick the number. If others are unlikely to pick a number that has just won then that is actually a good number to pick.

Table 5.6 demonstrates that winners in New Jersey did indeed get more when betting on numbers that had recently won. The difference is large, with a person choosing a number that won in the previous week winning on average $89, or 34 percent more than the average winner! The gambler's fallacy thus lowered a person's expected winnings. This suggests that we should see less evidence of the gambler's fallacy in New Jersey, where it did reduce earnings, than in Maryland, where it could not. Table 5.6 partly supports this by estimating the payout that there would have been in Maryland if winners were paid as in New Jersey. There we see slightly larger differences in potential winnings, but only slightly larger. That the gambler's fallacy

Table 5.6 Payouts to winners in state lotteries. Fewer people bet on numbers that have just won, so the payout is higher in New Jersey, where the winnings are shared, and would be higher in Maryland, if the winnings were shared.

Winning number repeating	Average payout in New Jersey	Estimated average payout in Maryland
Within 1 week	$349	$396
Between 1 and 2 weeks	$349	—
Between 2 and 3 weeks	$308	—
Between 3 and 8 weeks	$301	$382
Not within 8 weeks	$260	$289

Source: Terrell (1994).

Table 5.7 Evidence (or lack of it) for the hot hand fallacy. Supporters and players expected performance in the recent past to be a predictor of future performance – but it was not.

	Probability of making a field goal shot	
	If just missed one (%)	If just made one (%)
What supporters expected	42	61
What players expected	50	63
What the data says	54	51

Source: Gilovich, Vallone and Tversky (1985).

would reduce winnings did not, therefore, seem to make much difference to the strength of the bias.

The flip side to the gambler's fallacy is the **hot hand fallacy**. With this there is an expectation that streaks will continue, and it has particularly come to be associated with professional sport. For instance, a basketball player who is 'on form' or 'on a roll' is expected to keep on shooting well.

In a famous study, Gilovich, Vallone and Tversky (1985) put the hot hand fallacy to the test by looking at the performance of professional and college basketball players. This included following the fortunes of the Philadelphia 76ers during the 1980 to 1981 season. As Table 5.7 shows, supporters and players expressed a belief in a hot hand. They thought that a player would be more likely to make a shot if he had just made one. The data on actual performance painted a different picture. There was no evidence for a hot hand. This study has now been replicated in other sports, with similar results. Sports fans, it seems, are not good at judging whether a sequence is random or not!

At first sight, the gambler's fallacy seems very different from the hot hand fallacy. In the one, 'Five heads means the next must be tails', while, in the other, 'Five successes must mean he will be successful again'. It is important, therefore, to understand why the gambler's fallacy and hot hand fallacy can both result from the same law of small numbers. The key difference is whether the probability of each outcome is known.

The gambler's fallacy can occur when the probability of the outcome is known. It occurs if someone trying to predict the next draw from a sequence expects a small

sequence to match the known probability. For example, we know that a fair coin has a 50 percent chance of landing heads, and so John might expect 50 percent of coin tosses to be heads. This means he would think five heads in a row is unlikely, and likely to be balanced out by a subsequent tail. Similarly, a lottery number should not occur two weeks in a row.

The hot hand fallacy can occur when there is uncertainty over the probability of the outcomes. Someone trying to predict the probability of the next draw expects a small sequence to be a good predictor of what will happen next. So, a basketball shooter who has got five in a row must be on form and likely to get the next. Similarly, if John has read four reviews in a row supportive of Comfort, he might expect the next one will also be supportive of Comfort.

One interesting way to illustrate how these two fallacies can fit together is to look at sequences of random events, such as coin flips. Here are three possible sequences of coin flips. Which do you think are really generated by a coin?

hhTThTThhhTTThhTTThTThTTTThhhhhThTThThTTThThhTTTh
hTThhhTThTThThhThTTThhThhhThTThhThTTThhThhTTThThhT
hhhhhTThThhTThhhhhhhhThTTThTThhTThThhTTThThhhhhTT

The top and bottom really were random and the middle one a figment of my imagination. To try to understand why you might think the middle one is random it is useful to use the **proportion of alternations**. For a sequence of n signals you need to count the number of **streaks**, or unbroken sub-sequences, r and then calculate

$$\text{Proportion of alternations} = \frac{r-1}{n-1}.$$

If the sequence is random, the proportion should be 0.5, and in the top and bottom sequences it is 0.49 and 0.43. The evidence, however, is that many people think sequences with a proportion of around 0.6 are actually random, and in the middle sequence the proportion is 0.59.

That we expect the proportion of alternations to be high is consistent with the gambler's fallacy; we do not expect five heads in a row. That we think sequences with a low proportion of alternations cannot be random is consistent with the hot hand fallacy; if we do see five heads in a row then the coin cannot be fair. (Basketball fans may be interested to know that Gilovich, Vallone and Tversky (1985) find the proportion of alternations of professional basketball players was very close to that expected if the sequence of hits and misses was random.)

In looking at the consequences of the law of small numbers it is useful to have a model. This time we turn to Rabin (2002).

5.2.4 A model of the law of small numbers

Imagine again John reading car reviews. As before, let θ be the probability that a review is supportive of the Comfort and $1-\theta$ the probability that a review is supportive of the Sporty. John wants to know θ because then he can better know which car is better.

First, let's think about the gambler's fallacy. In this case, it is most apt to think that John is confident that he knows θ. The bias comes when he tries to predict what a review will say. We can capture this by assuming that John thinks signals are randomly drawn **without replacement** from a box containing N balls. On θN of the balls are written 'Comfort is better' and on $(1-\theta)N$ are written 'Sporty is better'. I have emphasized the 'without replacement' because that is what causes the bias. If you substitute 'with replacement' then there would be no bias. To see why, suppose that $N = 2$ and John thinks $\theta = 0.5$. Then John thinks the box contains one review supportive of Sporty and one supportive of Comfort. So, if he reads a review supportive of Sporty he would think the next report must be supportive of Comfort. If $N = 4$ then there are two reviews supportive of Sporty and two supportive of Comfort. If he reads a review supportive of Sporty he would think there is a two-in-three chance that the next report will be supportive of the Comfort.

This nicely captures the gambler's fallacy, because John expects that a sample of size N will be representative. The smaller is N, the more biased he is. There is a problem with the model, however: what happens when he has read more reviews than there are balls in the box? To overcome this, we assume that the box is replaced after he reads every other review. This makes the model work in a way that we can still capture the gambler's fallacy. Can it also capture the hot hand fallacy?

To capture the hot hand fallacy we need to add some uncertainty about θ. Suppose John thinks that $\theta = 0.25$, 0.5 or 0.75 are equally likely. He then reads two reviews in a row that are supportive of the Comfort. How will he update his beliefs about θ? Table 5.8 compares what he will think if he is unbiased and biased. We see that when he is biased he ends up overestimating the probability that reviews are supportive of Comfort. Indeed, he infers that it is impossible that reviews are more likely supportive of Sporty. This example is representative of what will happen more generally. John will overestimate how much he can learn from reading a small number of reviews. This is the hot hand fallacy.

One implication of the law of small numbers is that John's beliefs are too sensitive to the signals he gets. Whether he reads good reviews about the Comfort or the Sporty will lead to a bigger swing in his beliefs than is justified. This can have a knock-on effect in how much variation he perceives in the quality of cars. For example, suppose he compares other pairs of cars besides Sporty and Comfort. If $\theta = 0.5$ for all the pairs, then in around half the pairs he is likely to read two reviews supportive of one of the cars. The hot hand fallacy means that John will read too much into this and come away believing there is more variation in cars than there actually is. In other

Table 5.8 Beliefs when there is the law of small numbers. If John is biased, he overestimates the significance of two C signals.

		θ		
		0.25	**0.5**	**0.75**
The likelihood of two C signals	Unbiased, N large	0.0625	0.25	0.5625
	Biased, N = 4	0	0.166	0.495
Updated beliefs after two C signals	Unbiased, N large	0.071	0.286	0.643
	Biased, N = 4	0	0.251	0.749

words, some cars are naturally going to have a run of bad or good reviews, but John will overinterpret this as due to differences in quality rather than chance.

5.2.5 Generating random sequences

One interesting implication of the law of small numbers is that people may be bad at generating random sequences. For example, Rapoport and Budescu (1997) asked subjects to imagine a sequence of 150 draws with replacement from a deck of cards with five black and five red cards. They were asked to call aloud the sequence. Table 5.9 summarizes how subjects were biased towards switching too often and balancing out the sequence. If the sequence really is random, the probability of calling out black should always be 50 percent. We see, however, that subjects called out black more often if they had just called out red, and called out black less often if they had recently called out black more often.

It might seem a bit artificial to ask a person to generate a random sequence, but there are strategic situations in which this is a useful skill. One such setting is competitive sports, when a player wants to keep an opponent guessing. For example, in tennis a server does not want his opponent to be able to predict where in the court he will serve. Similarly, in soccer or hockey a penalty taker does not want the goalkeeper to know which side he will shoot. We might think that professional sportsman would learn to generate random sequences when competing. Do they?

One study, by Walker and Wooders (2001), looked at serving in ten important tennis matches and analyzed how often players served left versus right. There are two things that it is interesting to check with the data: whether a player was serving left and right with a probability that maximized his chance of winning, and whether sequences of serves were random. They find that players did serve left and right with appropriate probabilities but, on average, there was too much switching for the sequences to be random. To illustrate, Table 5.10 summarizes the famous 1982 Wimbledon men's singles final between Jimmy Connors and John McEnroe. Connors won the match, but that would be hard to guess from the data! McEnroe won roughly as many points to the left and right, showing that he was serving left and right with the correct probability. He also had an appropriate number of streaks (or proportion of alternations), indicating that the sequence of serves looked random. Connors, by contrast, should have served more to the right from the ad court and

Table 5.9 Probability of calling out a sequence of blacks and reds in a deck of cards with five black and five red cards.

Last few calls	Probability of calling out black (%)
Red	58.5
Black, red	46.0
Black, black, red	38.0
Black, black, black	29.8

Source: Rabin (2002).

Table 5.10 Data from the 1982 Wimbledon men's singles final. If maximizing their payoff, players should win approximately as many points to the left and right. A high number of streaks tells us if players were switching from left to right more often than a random sequence.

Player	Court	Serves			Points won		Streaks	
		Total	Left (%)	Right (%)	Left (%)	Right (%)	Number	Probability if random
Connors	Ad	78	41	59	50	70	49	0.001
	Deuce	91	84	16	67	53	31	0.042
McEnroe	Ad	71	45	55	72	62	36	0.563
	Deuce	79	44	56	69	68	36	0.848

Source: Walker and Wooders (2001).

more to the left from the deuce court than he did. He also had too many streaks for the sequence to look random, suggesting he was switching from left to right too often.

We see, therefore, that learning and experience does not necessarily mean that people can generate random sequences. It does seem clear, however, that people become better at generating random sequences with experience. For example, Walker and Wooders compared the performance of the tennis players with subjects in a previous lab experiment, and the data from the tennis players was starkly different from that of the lab subjects. Less clear is whether people experienced at generating random sequences in one setting, such as a tennis game, are able to carry that skill over to a different setting, such as a game of poker.

5.2.6 Do biases matter?

To briefly summarize what we have seen in this section, it is helpful to first recall how Bayesian updating requires both correctly interpreting signals and updating probabilities according to Bayes' rule. Confirmatory bias means that people are biased in interpreting signals. The law of small numbers means that people do not update probabilities according to Bayes' rule. Overall, therefore, we get a somewhat pessimistic picture of people's ability to use new information appropriately. And the bad news does not stop there, because there are many other biases that I have not mentioned here (some of which we shall see later).

Do such biases matter? Some important consequences of the biases we have seen include: (1) overconfidence and wrong beliefs; (2) a belief that there is more variation in a population than there actually is; and (3) a polarization of views, despite people seeing similar information. These look like pretty important things. For example, overconfidence can cause John to make a relatively uninformed choice of car and not choose the car that is best for him. Polarization can lead to prejudice and discrimination. It is clear, therefore, that we do need to take seriously possible biases in interpreting new information. This will be further illustrated when we look at health care and the stock market later in this chapter.

5.3 Learning from others

One issue that we need to consider in a bit more detail is how to interpret information that comes from other people. Often this is how we get new information, such as John asking his friend or reading a magazine review. The complication I want us to think about is that John sees only the choice or decision of the other person and not the reason for that choice. For example, if John sees a friend driving the Sporty car, he can never really know why his friend chose it. Instead, he has to try and infer his friend's motives and then question whether this tells him anything useful.

When trying to infer motives many complicating questions arise, such as the following. (1) How biased is my friend? (2) Does he have similar preferences or tastes to me? And (3) is he deliberately trying to manipulate me? We need to have some answers to all these questions.

5.3.1 To conform or not

To get us started, imagine that John test-drives both the Sporty and Comfort and decides he prefers the Sporty. Before he makes a final decision, however, he goes out and looks at what others are driving. If he sees lots of others driving Comfort, should he change his mind?

To answer that question it is useful to work through a simple model. Suppose that everyone has the same preferences. So, either the Comfort is the better car for everyone or the Sporty is the better car for everyone. Let β be the probability that Comfort is the better. Before anyone buys a car they test-drive both cars and form an independent opinion of which car is better. If Comfort really is the better car there is a θ_C chance a person will prefer Comfort on the test drive, and if Sporty is the better car there is a θ_S chance they will prefer Sporty on the test drive. To see what happens I will set $\beta = 0.5$ and $\theta_C = \theta_S = 0.67$ and work through the possible choices. Table 5.11 summarizes what is going on.

Anna is first to buy a car. She has no choice but to buy the car she preferred on the test drive. Imagine she preferred the Comfort and bought it.

Alan is the next to buy a car. He does the test drive and sees that Anna has bought the Comfort. If he preferred the Comfort on the test drive, then he has an easy choice. If he preferred the Sporty, he has a more difficult choice; it makes sense,

Table 5.11 An information cascade. Anna and Alan both prefer Comfort when they test-drive it, so buy it. Emma prefers Sporty but can see that the other two have bought Comfort, so also buys Comfort. Same for John. If three friends each prefer Sporty, their three signals outweigh the two choices of Anna and Alan.

	Order to buy a car					
	1st	2nd	3rd	4th	5th–9th	10th
Person	Anna	Alan	Emma	John	—	3 friends
Signal	Comfort	Comfort	Sporty	Sporty	—	3 Sporty
Choice	Comfort	Comfort	Comfort	Comfort	Comfort	Sporty

however, to choose the Sporty. Imagine that Alan actually preferred the Comfort on the test drive and bought it.

The choice of both Anna and Alan tells a lot about what car they preferred on the test drive. In particular, it should be possible to infer their signal from their choice. If Anna buys the Comfort she must have preferred the Comfort on the test drive. If Alan buys the Sporty he must have preferred the Sporty on the test drive. This means that observing what they bought is useful information.

Emma is the third person to buy a car. She can see that two people have bought the Comfort and from this infers that two people must have preferred the Comfort on the test drive. What does that mean? Even if Emma does a test drive and prefers the Sporty, it is still two against one in favor of the Comfort. Using Bayes' rule, these two observations should outweigh her one signal. She should choose Comfort.

We are arguing that, irrespective of what she thinks on the test drive, Emma should choose Comfort. This means that she should ignore her own signal. Equivalently, her choice tells us nothing about her signal. We call this an **information cascade**. The same is true of John and all those after him. They should buy Comfort.

Finally, suppose that three friends decide to all buy a car simultaneously. They each test-drive the cars and say what they prefer. Imagine that all three preferred the Sporty. Do these three signals that Sporty is better outweigh seeing nine people driving Comfort? Yes, they do, because the friends should know that they cannot infer anything from Emma, John and those who follow buying Comfort. They bought Comfort only because Anna and Alan did. So, it is three signals versus two observations, and they should buy the Sporty.

What we see here is that an information cascade can be **fragile**. That three, four, five or more bought Comfort conveys no new information. Thus, someone with a more precise signal may want to go with their signal rather than what others have been observed doing.

What can we learn from this model? First of all, we see that a person's choice may convey a lot of information about what they know. This is the case for Anna, Alan and the three friends. We also see, however, that a person's choice may convey very little information about what they know, because they are conforming to others. This is the case for Emma and John. It is very important, therefore, to try and work out how much information someone's choice conveys before drawing too many inferences from it. In practice, this can be complicated to do, particularly if $\theta_C \neq \theta_S$ or different people have different θ_S. But it is more important for us to start questioning how people behave when faced with such situations.

5.3.2 Cascade experiments

There have now been a variety of experimental studies to see whether information cascades happen as predicted. Anderson and Holt (1997) were two of the first to do so, and Figure 5.2 summarizes some of their results. There are two things that these results try to pick up. The first is how close subjects got to the payoff they would have achieved using Bayes' rule. This is measured by actual efficiency. When $\theta_S = \theta_C = 0.67$ we see that subjects were pretty close to using Bayes' rule, but in the more complicated game, in which $\theta_S = 0.86$, $\theta_C = 0.71$, there is a bigger gap. The next

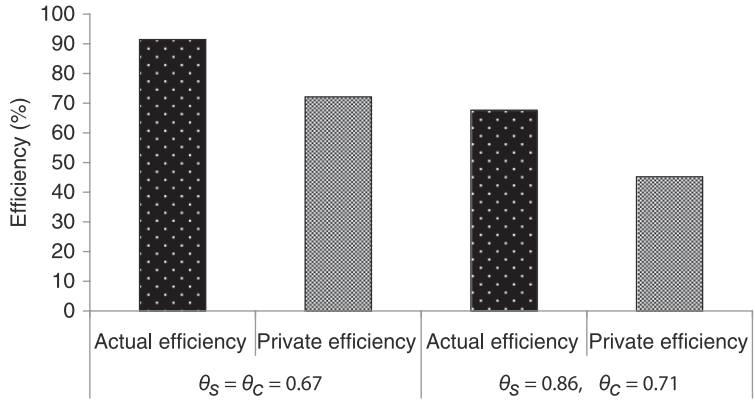

Figure 5.2 A cascade experiment. The payoff that subjects got relative to what they would have got if they had followed Bayes' rule, actual efficiency. The payoff they would have got if they had followed their own signal relative to what they would have got with Bayes' rule, private efficiency.
Source: Anderson and Holt (1997).

question is how well subjects would have done by just following their own signal. This is measured by private efficiency. We see that in both cases subjects did a lot better than they would have done from following their own signal. It is clear from this, and other studies, that people learn from observing others – actual efficiency is above private efficiency – but not everyone uses Bayes' rule – actual efficiency is less than 100 percent.

If John knows that others are likely to be biased in some way then he should rethink what choices say about signals. For example, if he suspects that Emma will follow her own signal and ignore that Anna and Alan chose Comfort, her choice is more informative than it would be if she was following Bayes' rule. We can test whether subjects do take this into account by looking at choices relative to the **empirical optimum** (see Research methods 5.1). For example, we can look at all situations in which a subject has observed three others choosing C. On average, does a subject in this situation get a higher payoff choosing S or C? We call the answer the empirical optimum, and then ask how many subjects follow the optimum.

❖ RESEARCH METHODS 5.1 EMPIRICAL OPTIMA VERSUS NASH EQUILIBRIUM

A Nash equilibrium details a strategy for each person such that it is optimal for a person to follow his strategy if all other people follow theirs. That is, it consists of a strategy s_i for each person such that he can do no better,

$$u_i(s_i, s_{-i}) \geq u_i(x, s_{-i})$$

by playing another strategy x if others play strategies s_{-i}. In a cascade experiment there is one Nash equilibrium, and so it is simple to say whether a subject behaved consistently with the Nash equilibrium.

The problem is that, if one person does not behave consistently with Nash equilibrium, because of some bias, then it may no longer be optimal for others to follow their Nash equilibrium strategy. For example, Anna is supposed to buy the car she preferred on the test drive. If she randomly decides what to buy then it is optimal for others to ignore what she does. So, if Anna does not follow her Nash equilibrium strategy, neither should anyone else!

This makes it difficult to judge whether people are behaving optimally. Does a person deviate from his Nash equilibrium strategy because he is biased or because he expects others to be biased?

This will be a recurring question over the next few chapters, and calculating the **empirical optimum strategy** is one way to address it. First we calculate the strategies that others typically follow. Denote by e_i the strategy the person usually plays. We then ask whether person i is maximizing his payoff,

$$u_i(e_i, e_{-i}) \geq u_i(x, e_{-i})$$

given that others are playing e_{-i}. If the answer is 'Yes' then we say that i is playing the empirical optimum strategy. If all people play the empirical optimum strategy we get back to a Nash equilibrium.

Weizsäcker (2010) did this for a collection of previous experimental studies (see Research methods 5.2), and the results are summarized in Table 5.12. Let's focus first on the choice rate. This tells us that on those occasions when subjects would have maximized their expected payoff by not following their own signal they did so only

Table 5.12 Did subjects play the empirical optimal strategy in cascade experiments? The choice rate tells us the percentage of times subjects played the optimal strategy. The success rate tells us the percentage of times they chose the best option.

Some possible scenarios	Choice rate	Success rate	
		Actual	Empirical optimal
Should not follow own signal	44	53	64
Should follow own signal	91	73	75
Signal different from majority choice	—	54	—
Signal same as majority	—	89	—
Signal different from what many and all have previously chosen	—	73	

Source: Weizsäcker (2010).

44 percent of the time. The conclusion is clear, therefore: subjects followed their own signal more often than they should have done. By looking at the success rate we can see how much this lowers their payoff. When it is optimal for them to not follow their own signal they choose the best option 53 percent of the time, but if they were to follow the empirical optimal strategy they would choose the best option 64 percent of the time.

This drop in success rate shows that subjects did not learn enough from what others were doing. The second part of Table 5.12 shows us that subjects were successful in learning from others only if a lot had all chosen the same thing. Put another way, for the average subject to not follow his signal there needed to be a more than 67 percent chance that not following his signal was optimal. This ratio should have been 50 percent. We need to ask why subjects were reluctant to learn more from the choices of others.

❖ RESEARCH METHODS 5.2 META-DATA ANALYSIS

A typical experimental study will involve around 50 to 150 participants. If each participant makes between ten and 20 decisions, this gives 500 to 3,000 data points. This should give enough statistical power to test some important hypotheses, but not all. Ideally, we would want more data points and more variety of locations and subjects. One way around this is to be a bit more inventive in recruiting subjects, such as using the internet or newspapers. Another way is to compile meta-data sets in which results from similar studies are combined. This is what Weizsäcker (2010) did for his study. The results from 13 cascade experiments, including that by Anderson and Holt, were pooled to give almost 30,000 data points. With such a large data set it is possible to ask questions that cannot be asked with a single study.

5.3.3 What happened to conformity?

The basic conclusion from the cascade experiments is that subjects follow their own signal more often than they should. This is interesting, because it seemingly goes against psychological evidence that people tend to conform to the choices of others more than they should. In famous experiments by Soloman Asch, for instance (see Asch 1956), subjects given a very simple visual task seemingly conformed to the choices of others (see Research methods 5.3). So, what role does conformity play? A study by Goeree and Yariv (2015) shows that conformity can matter. In their study, subjects had to choose between S and C, as in a cascade experiment, but without seeing an informative signal. Specifically, the first three subjects made their choice seeing only the choices of those before them. This would equate to Anna, Alan and Emma choosing without the benefit of a test drive. They seemingly have no option but to randomly choose a car. Subsequent subjects were then given the choice to see a private informative signal or to see the choices made by those who had not received a signal. So, John could either do a test drive or find out what Anna, Alan and Emma had chosen.

❖ RESEARCH METHODS 5.3 DECEPTION IN EXPERIMENTS

In Chapter 1 I briefly talked about the issue of deception. In economic experiments, everything that subjects are told should be true (though there might be things they are not told). In a psychology experiment it is not considered so bad to mislead or deceive subjects about the true experiment. This is apparent in the studies by Asch.

In these experiments, subjects were asked to say which of three lines was the same length as a fourth. In a control experiment only one subject ever gave a wrong answer, so this was a very simple task. In the real experiment there was one subject and a number of **confederates**, people told what to do by the experimenter. In turn, the confederates would say which line they thought was the same length, and then the subject would say. Note that the subject did not know that the others were confederates, hence the deception. On some occasions the confederates were told to give the wrong answer, with the objective of seeing whether the subject would conform to the obviously wrong choice of others. Three-quarters of subjects did conform at least once during the experiment.

As I have already said, the deception in this experiment was to give the impression that confederates were genuine subjects. Economists do not like deception, for a number of reasons. A most basic concern is whether choices will be reliable if the subject fears deception, or the situation appears implausible. In the Asch experiments it is simply not possible that a string of people would have given the wrong answer. The subject is, thus, faced with a non-credible situation that can create a sense of discomfort and unease because 'something is not right'. How reliable are choices made in such a context? The study by Goeree and Yariv does not raise such concern because all choices were made by genuine subjects.

If a subject's objective is to choose the best option then this is a simple choice: they should get the informative signal. As Figure 5.3 shows, however, this is not what a significant proportion of subjects did. Two treatments were considered. In the first, a subject is paid solely for choosing the best option. In this case, 34 or 50 percent of subjects, depending on the stakes, chose to see what others did and 84 or 88 percent of those copied the majority choice. In the second treatment, subjects were paid if the majority of subjects chose the best option. Here there is at least some rationale for looking at others' choices in order to go against the majority and balance out the uninformed choices of others. We see that most subjects still conformed to the majority choice, however. The end result of this was lower success rates. In the high stakes treatment, subjects missed out on an average of $16.80 each, compared to the optimal strategy!

Overall, therefore, we get a mixed picture: the cascade experiments suggest people often overweight private information, while this study and others suggest people

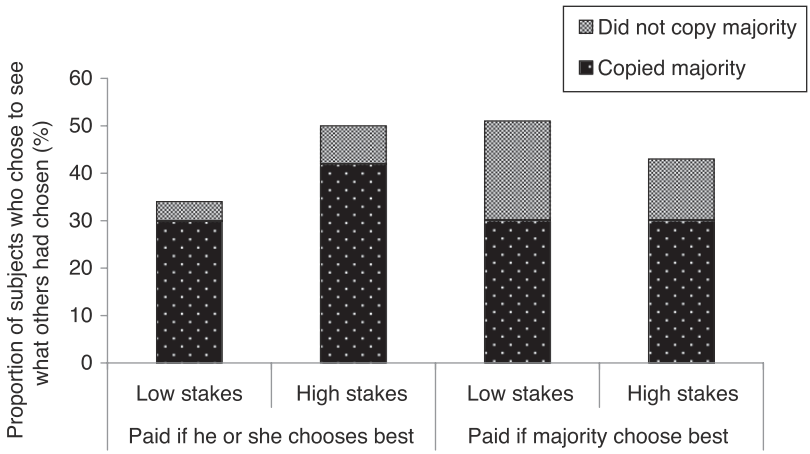

Figure 5.3 The proportion of subjects who chose to see what others had done and then copied or not the majority.
Source: Goeree and Yariv (2015).

conform too much to the choices of others. These two things need not be incompatible, but just how much choices are determined by own information, or those of others, remains an open question.

5.3.4 Signaling games

Before we finish this section there is one important complication I want to mention. In the cascade experiments Anna, Alan and everyone else have no incentive other than to try and buy what they think is the best car. Sometimes, however, a person may have an incentive to try and influence others through her actions. For instance, if Maria's job is to sell Comfort cars then it is not a surprise that she drives one. It would also not be a surprise if she says to John that '[i]t is a fantastic car, far better than Sporty'. She might be telling the truth, but she clearly has an incentive to say such a thing, even if it is not the truth.

To model such situations we can use a signaling game. The basic idea of a **signaling game** is that one person, called a **sender**, does some action, such as saying how wonderful the Comfort car is, and another person, called the **receiver**, tries to infer information from that, such as how good Comfort really is. Signaling games are very useful for modeling strategic interaction, and so we shall come across them a few times in the next couple of chapters. Here I shall look at a study by Brandts and Holt (1992) to illustrate the issues that arise.

Rather than talk of car saleswomen, I want you to imagine signaling in the labor market. Figure 5.4 schematically depicts the game in a game tree. It looks a bit messy but is not that bad if we think one stage at a time. Start at the middle. Chance decides whether the worker, John, is low ability or high ability. John then chooses how hard to work in his education. He can work hard and get A grades, taking him to the left of Figure 5.4, or work less hard and get B grades, taking him to the right of the figure. Knowing the grades of John, but not knowing his innate ability, the employer, Anna,

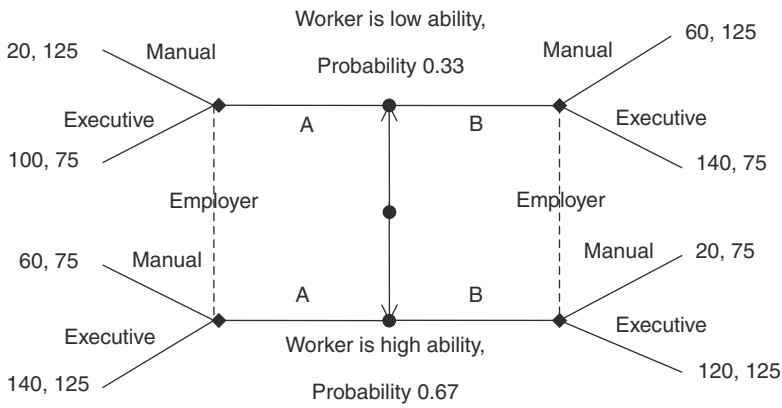

Figure 5.4 The game tree of a signaling game. Chance decides the type of the worker. The worker then chooses A or B and the employer manual or executive. The worker knows his type but the employer does not.

must decide whether to give John an executive job or a manual job. The numbers give the payoffs to John and Anna for any combination of decisions.

The payoffs mean that, *ceteris paribus*, John prefers the executive job to the manual job, but prefers to get A grades when high ability and B grades when low ability. Anna would like to put a low-ability worker in a manual job and a high-ability worker in an executive job.

What makes the game interesting is that John knows whether he is high or low ability but Anna does not. She would like to know, and so has to try to infer something from John's actions. John, however, has an incentive to appear as though he is high-skilled. He potentially, therefore, has an incentive to mislead. It is this incentive to mislead that has been absent from everything we have done so far in this chapter.

There are two Nash equilibria in this game, one given the technical name of intuitive and the other unintuitive. The **intuitive equilibrium** is that John works hard and gets an A grade whether he is high or low ability and is then given an executive job; if John did get a B grade, he would be given a manual job. The **unintuitive equilibrium** is that John works less hard and gets a B grade whether he is high or low ability and is given an executive job; if John did get an A grade, he would be given a manual job.

In both cases there is **pooling**, in the sense that John will do the same irrespective of his type. This means that the employer will not be able to distinguish whether John is high or low ability. The reason that one of the equilibria is called 'intuitive' and the other 'unintuitive' should, hopefully, need little explanation. It seems far more natural that John should have to sacrifice a little when he is of the low-ability, undesirable type than when he is of the high-ability, desirable type. A priori, however, we cannot rule out the unintuitive equilibrium.

Table 5.13 summarizes results from the study by Brandts and Holt in which subjects played this and similar games. In the top row we see that subjects did not initially play either the intuitive or unintuitive equilibrium. Instead, high-ability workers chose A and low-ability workers chose B. Employers did behave consistently with the intuitive equilibrium, however, by assigning workers with an A grade to executive jobs and

Table 5.13 Play in a signaling game (%). In the standard game, behavior converges towards an intuitive equilibrium. In the revised game, behavior is closer to that of the unintuitive equilibrium.

Game	Round	Worker choice given type		Employer choice given worker choice	
		A if high	A if low	Executive if A	Manual if B
Standard game	1–4	100	21	98	74
Standard game	9–12	100	75	98	60
Revised game	1–6	39	52	45	15
Revised game	7–12	15	38	56	13

Source: Brandts and Holt (1992).

those with a B grade to manual jobs. Workers learnt from this, and by rounds 9 to 12 both types of worker were choosing A.

In some subsequent experiments the payoffs of the game were revised to make workers more indifferent whether they chose A or B but to give employers more incentive to predict worker type. One objective was to see whether the unintuitive equilibrium would ever emerge. To see whether it did, we look to the bottom two rows of Table 5.13. Intriguingly, we do see that most workers choose B and high-ability workers are less likely to choose A than low-ability workers. This is consistent with the unintuitive equilibrium. Employer behavior is not consistent, however, because those getting an A grade are still more likely to be allocated an executive job. We do not, therefore, see the unintuitive equilibrium, but are far from seeing the intuitive equilibrium either.

This study by Brandts and Holt illustrates some of the more general lessons that have emerged from looking at signaling games. The most important lesson is that people do seem to behave strategically and expect others to behave strategically. For instance, the low-ability worker may work hard to get A grades in order to appear high ability, but the employer will probably realize this and be skeptical whether those with A grades are high ability. Similarly, the car saleswoman may say 'Comfort is best', even if she knows it is not, but the buyer is likely to be skeptical of such claims.

It seems, however, that people do need a bit of experience before they start behaving strategically. At first, the employer may think that all those with A grades must be hard-working, and John may believe the car saleswoman's claims. With experience they learn to be more skeptical. One consequence of such learning is that it becomes hard to predict what will happen. We can predict that people will learn, but we cannot predict quite what they will learn; it could depend on all sorts of random things. This makes it hard to know for sure whether the intuitive or unintuitive equilibrium will be the one to emerge in a particular situation.

In the next chapter we shall explore these issues of strategy and learning in a lot more detail. The key thing I want you to see at this point is how complicated it can be to try to infer information from someone's actions, particularly if there is an incentive for them to mislead. Things get even more complicated when we realize that confirmatory bias, the law of small numbers and conformity are all clearly relevant in signaling games. For example, an employer who has some workers who got A grades and some B grades might suffer from confirmatory bias, and think that those with A grades are more productive than they really are. In addition, because of the law of small numbers, and a few bad experiences with workers who got B grades, he may

hire only those with A grades. Learning from new information is complicated, and learning from what others do is even more complicated!

5.4 Summary

We started by looking at Bayesian updating as the most appropriate way to learn from new information.

I next suggested that choices often involve compound lotteries or ambiguity. This led us to look at recursive expected utility and ambiguity aversion. We saw that most people seem to be averse to ambiguity and/or second-order risk. This means that people should be eager to get new information.

We also saw, however, that people can be biased in how they interpret new information. Confirmatory bias means someone can misinterpret new information and the law of small numbers means they can infer too much from new information.

Next, we saw that people can also be biased when learning from what others do. In some situations people seem to learn less from others than they could have done, and in other situations they conform more to others than seems sensible.

Finally, we saw that if one person has the incentive to mislead another then we can use a signaling game to capture this, but things can quickly become quite complex. It therefore becomes hard to know how a person should interpret new information.

Overall, this presents a somewhat depressing picture of how wisely people use new information. Are things that bad? On the one hand, I think they are; the assumption that people use Bayes' rule I would comfortably put as the dodgiest assumption in the standard economic model. On the other hand, things are probably not so bad; in particular, people can still make good choices even if they do not use information wisely. One reason for this is the aversion to ambiguity and second-order risk. Such an aversion may mean that people simply avoid making choices when they are poorly informed, and we saw some evidence for this in Chapters 2 and 4.

5.5 Health care

Health care is one area of life in which information is fundamentally important. To see why, let's think about the decisions that are made. A patient will decide things such as: whether to take out private medical insurance, whether to seek treatment if ill, where to go for treatment, what to tell the doctor or nurse, whether to follow their recommendations and whether to seek a second opinion. Doctors and other medical practitioners decide things such as: which patients to prioritize, what tests to perform, what treatments to recommend, and what to tell a patient. Each of these decisions requires good use of available information, and the stakes could not be higher, because lives are involved. What evidence do we have that patients and practitioners make these decisions wisely? I will look at patients first and practitioners second.

5.5.1 Patients

Recent years have seen dramatic increases in the amount of information available to patients in many Western countries, the internet being one reason for this. For example, patients in the United Kingdom can go to a website (www.nhs.uk), put in

their address and instantly find a wealth of information about doctors and hospitals in their area. Everything is rated, from the quality of hospital food to mortality rates. Similar information is available in the United States (www.HospitalCompare.hhs.gov). Do people use such information?

The evidence suggests not. The Kaiser Family Foundation routinely surveys people in the United States about attitudes to health care. In the 2008 survey only 30 percent of people had seen information comparing health care providers in the last year and only 14 percent used this information. Instead, people were more likely to rely on family and friends; an earlier survey found that 59 percent thought that the opinions of family and friends are a good source of information, compared to just 36 percent who thought that family and friends do not have enough knowledge and experience to advise. Figure 5.5 gives some more evidence to think about. Here we see a strong, but declining, desire for familiarity rather than higher expert ratings.

One might argue that family and friends are a potentially good source of information. The evidence suggests otherwise, however. For instance, patients are good at judging aspects of quality such as whether the doctor was respectful, attentive and clear in explaining issues and had a clean and efficient office. Unfortunately, they are not so good at judging aspects such as whether the doctor supplied appropriate, evidence-based treatment! It also seems that patients have high levels of trust in practitioners and are willing to put errors down to things beyond the practitioner's control, even if it is the case that the practitioner is stressed and overworked.

The evidence does, therefore, point towards the biases we looked at in section 5.2. For instance, we see evidence for the law of small numbers, given that patients are content to rely on past experience and the experience of friends; such experience is likely to be a fairly small sample on which to draw inferences. We see evidence for confirmatory bias in the trust that patients put in practitioners and treatments; if they think their doctor is good or medication is working then ambiguous information may be interpreted in a way to support this belief. The considerable ambiguity of

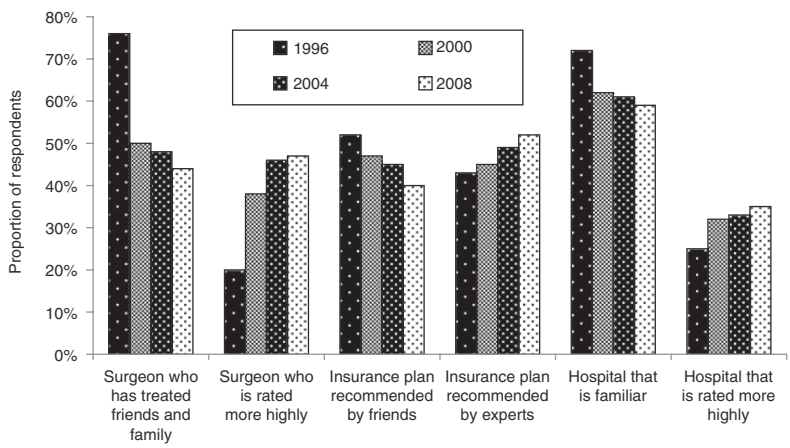

Figure 5.5 Preferences of respondents in surveys of patients in 1996, 2000, 2004 and 2008. Most people prefer things that are familiar or recommended by friends and family, rather than things recommended by experts.
Source: Kaiser Family Foundation (2008).

information about health care exacerbates these two biases. For instance, if a patient gets better, was that because of the medication or nature taking its course? The answer is ambiguous enough that a patient could believe whatever they want.

That we observe the law of small numbers and confirmatory bias is not so surprising. What is more surprising is that patients have alternative sources of information that they seem to ignore. Why would patients choose not to use the wealth of information that is available about health care based on large samples and scientifically based experiments?

One possibility is the availability heuristic. The **availability heuristic** could be summarized: if you remember something, it must be important. One consequence of this heuristic is that more vivid and memorable experiences can seem more important. For instance, John is more likely to remember his grandfather living to 100 despite smoking all his life than he is to remember the statistics on lung cancer. Similarly, he is more likely to remember the great treatment he got at a hospital than a report showing that the average performance of the hospital was poor.

A nice illustration of the availability heuristic is provided by Lichtenstein and co-authors (1978). In the study they looked at whether subjects knew the likely causes of death in the United States. In one part of the study subjects were told that, on average, 50,000 people per year die in the United States because of motor vehicle accidents and were then asked to say how many they thought died from 40 other possible causes. Table 5.14 gives some of the results. There is a primary bias to

Table 5.14 Approximate number of deaths in the United States and estimated number of deaths from different causes.

	Actual number	Estimated number	Newspaper inches
Fireworks	6	331	0
Whooping cough	15	171	0
Venomous bite or sting	48	535	0
Tornado	90	688	153.5
Lightning	107	128	0.8
Non-venomous animal	129	298	33.8
Flood	205	863	41.8
Pregnancy, childbirth and abortion	451	1,932	0
Appendicitis	902	880	0
Electrocution	1,025	586	42.2
Train collision	1,517	793	0
Asthma	1,886	769	1.9
Firearm accident	2,255	1,623	28.2
TB	3,690	966	0
Fire and flames	7,380	3,814	320.7
Drowning	7,380	1,989	247
Leukemia	14,555	2,807	14.8
Homicide	18,860	8,441	5,042.9
Accident falls	17,425	2,585	124.8
Breast cancer	31,160	3,607	0
Diabetes	38,950	2,138	0
Motor vehicle accidents	55,350	50,000	1,440.5
Lung cancer	75,580	9,723	35.9
Heart disease	738,000	25,900	303.4

Source: Lichtenstein *et al.* (1978).

overestimate the number dying from less likely causes. More interesting is a secondary bias in which two things that actually cause a similar number of deaths are thought to cause very different numbers of deaths. They find that this bias correlates with things that might influence availability, including newspaper coverage and whether the subject had direct experience of someone dying from the cause.

When put together, the law of small numbers, confirmatory bias and the availability heuristic present a depressing picture of how wisely patients may use information and make choices. One particular situation in which the negative consequences of such biases are abundantly clear is health scares. For example, a study published in 1998 in the United Kingdom suggested that a vaccine against measles, mumps and rubella (the MMR vaccine) caused autism in children. For one reason or another, this story hit the headlines, and, despite all the experts saying that the MMR vaccine was safe, the number of parents who gave their children the vaccine fell dramatically. By 2009 the United Kingdom had one of the highest rates of measles in Europe. Given that one in 15 children suffers complications from having measles, this is a serious problem that could have been avoided. An inability to correctly interpret new information caused this health scare.

5.5.2 Practitioners

That patients may be biased is not great, but surely we can expect better of medical practitioners? Unfortunately, the evidence here is not good either.

The possibilities for bias are not hard to imagine. For instance, we might imagine a doctor is reluctant to use a treatment that did not save one of his previous patients. More generally, we might imagine a doctor relies more on his own experience with treatments than the results of scientific tests. The law of small numbers, confirmatory bias and the availability heuristic all come into play again. And now we have a new factor to add into the mix: conformity. If a doctor is in a hospital where other doctors use a particular treatment, will he be the one to do something different? But this is just conjecture. What does the data say?

One thing that it is very easy to find data on is whether patients receive the same treatment for the same illness in different parts of the country. To a rough approximation, they should. Often they do not. Table 5.15 gives some examples of the

Table 5.15 The highest and lowest rates of performing a treatment (per 10,000 people) in different areas of the United States.

Procedure	Highest rate	Lowest rate
Injection of hemorrhoids	17	0.7
Knee replacement	20	3
Carotid endarterectomy	23	6
Bypass surgery	23	7
Heart catheterization	51	22
Hip replacement	24	8
Appendix removal	5	2
Hernia repair	53	38

Source: McCall (1996).

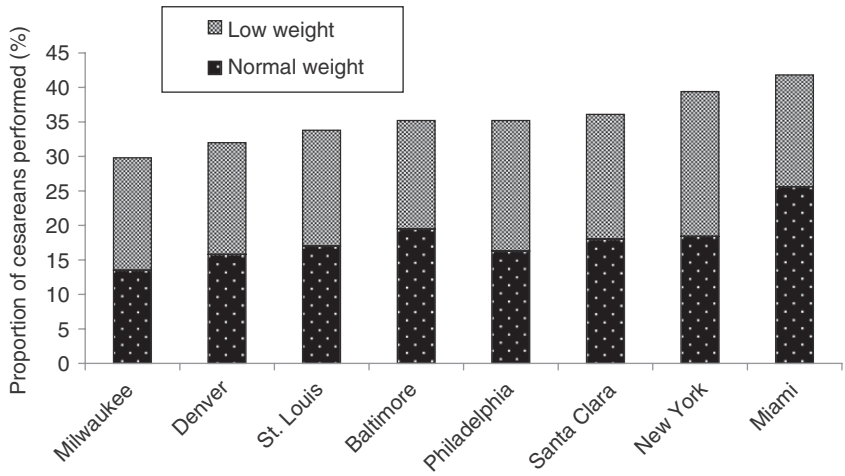

Figure 5.6 Cesarean rates in different counties of the United States for normal-weight and low-weight babies.
Source: Baicker, Buckles and Chandra (2006).

differences in the number of people treated in different areas. Such variation is known as **small area variation**. To give a more specific example, Figure 5.6 summarizes some data from a study by Baicker, Buckles and Chandra (2006) into the prevalence of cesarean sections in different US cities. The variation may not look that big, but, given the number of births in these cities, there is no way that this could be random variation.

If it is not random variation, then what is it? One suggestion is that practitioners in certain areas conform to the norms of their peers or learn from their peers. This can be exaggerated by the education of new doctors if different teaching hospitals emphasize different methods. For instance, Phelps and Mooney (1993) quote a physician from New Haven: 'The academic flavor in Boston, the teaching atmosphere, has a much stronger tradition of bringing people into the hospital… [When Boston-trained physicians relocate in New Haven] they bring their bad habits with them, but peer pressure changes that.'

Small area variation cannot be efficient. Treatment is expensive, and so either patients are getting undertreated in some regions or overtreated in others. In a thorough review published in 2001, the Institute of Medicine gives a fairly damning verdict: 'Americans should be able to count on receiving care that meets their needs and is based on the best scientific knowledge. Yet there is strong evidence that this frequently is not the case.' They do indeed find evidence of under- and overtreatment. They also find that physicians regularly departed from best practice in treating conditions such as type 1 diabetes, depression and asthma, and were reluctant to use drugs widely shown to have improved health outcomes.

My objective is not to scare you about the state of the US health system, or any other health system. But we do need to be realistic. Health care practitioners are no better or worse than anyone else, and so they are likely to be biased in interpreting information. Recognizing this provides a means to progress. If we have the tools to

understand biases, and we have seen some in this chapter, then we can potentially mitigate the problems they cause.

5.6 Bubble and bust

Bubbles and crashes are familiar in many markets, such as the housing and financial markets. It is headline news when markets crash, and bad news for many. This is why crashes are often etched into the history books as a black day: 'black Friday' on September 24, 1869, 'black Thursday' and 'black Tuesday' on October 24 and 29, 1929, 'black Monday' on October 19, 1987, 'black Wednesday' on September 16, 1992, and 'black week' beginning October 6, 2008! What I want to look at in this section is why we get such extreme swings in the prices of goods.

I will primarily focus on stock markets. The **efficient markets hypothesis** says that the price of stock should be equal to its fundamental value. The **fundamental value** is, basically, the expected future returns from the stock, in the form of dividends. We cannot predict for certain what these future returns will be, and so the fundamental value is uncertain. The share price should, however, take into account all the information available at the time. It follows that the share price should change only if there is some new information about fundamental value.

So, information is key. Given, therefore, that new information is likely arriving all the time, is it any surprise that prices fluctuate? This turns out to be a tricky question. The consensus, however, is that the fluctuations we observe in share prices seem too large to be caused solely by new information.

To explain: mathematical logic says that the fluctuations in the fundamental value must exceed fluctuations in the expectation of the fundamental value. This means that share prices should be less volatile than fundamental value. The intuition for this is that, if people do not know all the relevant information, the share price should vary less, on average, than it would have done if they had known everything. What we observe is the opposite. To illustrate, Figure 5.7 plots the price of the Standard & Poor's (S&P) 500 index, compared to the estimated fundamental value. Prices clearly fluctuate more than the fundamental value.

There are lots of reasons to question the data in Figure 5.7. For instance, the fundamental value at a particular time is found by looking at the discounted value of dividends that accrues after that time, and we could argue what discount rate should be used. The data also looks at a composite index, rather than individual assets. Even if we take this into account, however, the evidence points to more volatility in share prices than we might expect. We observe **bubbles**, when price exceeds fundamental value, **busts**, when price is below fundamental value, as well as sudden rises and falls in prices.

Before we get on to why this happens, it's useful to step into the experimental lab to see whether we observe bubbles and busts there.

5.6.1 Bubbles in the lab

In Chapter 2 we saw that double-auction markets appeared to work well, in that the price and quantity traded was close to the efficient amount with supply equaling demand. In those experiments, subjects were either a buyer or a seller, knew their

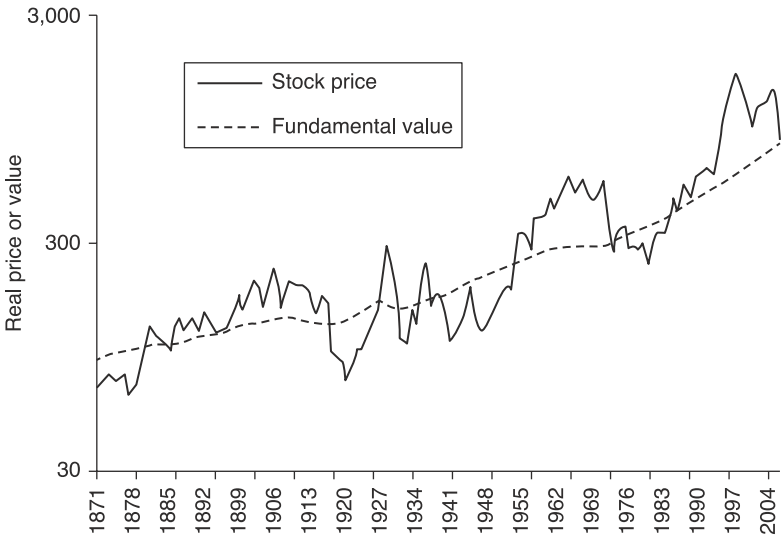

Figure 5.7 Price of the S&P 500 index compared to the estimated fundamental value. The fundamental value at a particular time is calculated by looking at what the dividends were in the following years.
Source: Shiller (2003).

value for the good and made their trades one at a time. This is a good representation of many markets, such as the labor market or new car market, but is not so good a representation of many other markets, such as the stock market. In a stock market we observe traders who sometimes buy and sometimes sell, do not know the value of the good they are trading and trade over consecutive periods. What happens when we add such things to experimental markets?

To illustrate, I shall look at some results of a study by Smith, Suchanek and Williams (1988), in which subjects took part in an experimental asset market. To explain how these experiments worked I will talk though an example, with Figure 5.8 to help. There are three subjects or traders. At the start of the experiment each trader is given an initial amount of money and some units of the asset. There are then 15 market periods, lasting at most 240 seconds each. In each period traders can sell any units of the asset they hold, or buy units if they have sufficient money. This buying and selling takes place via a double auction. At the end of each period the asset yields a randomly determined amount of money called a dividend. Traders receive this dividend and the next period begins. Thus, traders can accumulate money and assets over the 15 periods. At the end of the fifteenth period traders take home any money they have accumulated, and the assets become worthless.

Given that assets will be worthless by the end of the experiment, the primary reason to hold an asset is the dividend it may yield. Suppose that the dividend could be \$0.00, \$0.08, \$0.16 or \$0.40 with equal probability. The expected dividend is \$0.16 per period, and so the expected value of holding the asset in period one is 15 × \$0.16 = \$2.40, the expected value in period two is 14 × \$0.16 = \$2.24, and so on. By the final period the expected value is just \$0.16. The maximum dividend is \$0.40

Start of the experiment	Anna is given $2.80 and 4 units of the asset. Alan is given $7.60 and 2 units. Emma is given $10.00 and 1 unit.
Market period one	Anna sells a unit to Alan for $4.00 and Emma for $3.80.
End of period one	Anna has $10.60 and 2 units of the asset. Alan has $3.60 and 3 units. Emma has $6.20 and 2 units. The asset yields a dividend of $0.40.
Start of period two	Anna has $11.40 and 2 units of the asset. Alan has $4.80 and 3 units. Emma has $7.00 and 2 units.

Figure 5.8 A description of an experimental asset market that can give rise to bubbles.

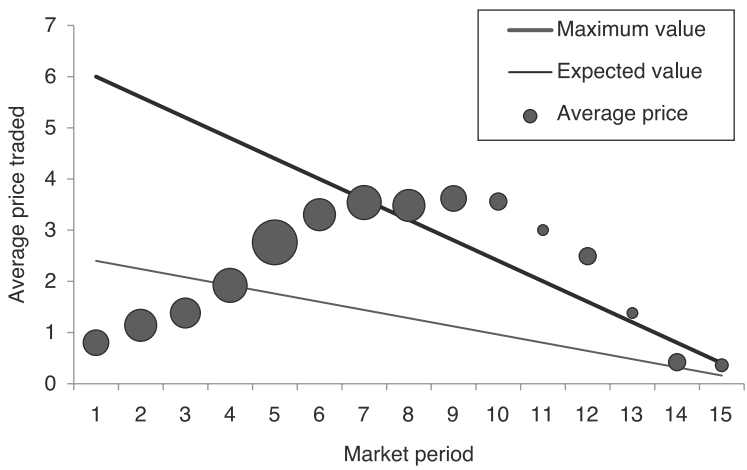

Figure 5.9 A bubble in an experimental asset market. The size of the average price circle indicates the number of trades done. The efficient markets hypothesis would suggest average price should track expected value or, at least, not go above maximum value.
Source: Smith, Suchanek and Williams (1988).

per period, and so the maximum value of holding the asset in period one is 15 × $0.40 = $6.00, the maximum value in period two is 14 × $0.40 = $5.60, and so on.

Figure 5.9 illustrates what happened in an experimental market doing the same as we have described but with 12 rather than three subjects. What we see is a spectacular bubble and crash! The price rises in the first eight periods with relatively large

amounts of trading. It reaches a price far higher than can be justified by the value of the asset, even on the most optimistic hopes of future dividends. Then, after period ten, the volume of trade slows down and the price crashes towards zero.

What we observe in Figure 5.9 is startling, and to my mind one of the most fascinating insights from behavioral economics to date. This was a very simple market to understand, with a minimal amount of uncertainty, yet we observe trading at prices well below and well above the fundamental value of the asset. If we get such bubble and bust in so simple an environment then no wonder we see bubble and bust in more complicated financial and goods markets.

You might think that bubbles would stop happening if we change the market institution. For example, what if we allow short selling, when a trader can sell assets they do not own? Or what about futures trading, when traders can trade, say, period eight assets ahead of schedule? The basic conclusion seems to be that institutions sometimes dampen bubbles, but they do not stop them happening. Bubbles, it seems, are hard to stop, even in these simple experimental markets (though see Research methods 5.4).

❖ RESEARCH METHODS 5.4 WHAT IF SUBJECTS DO NOT UNDERSTAND INSTRUCTIONS?

A fairly basic requirement if an experiment is to go well is that subjects understand the instructions. If subjects are confused then, clearly, their behavior may not reveal how they would behave if they actually knew what was going on. It is vital, therefore, that the experimenter make instructions as clear as possible. Things may not be so simple in practice, however, as a pair of studies by Huber and Kirchler (2012) and Kirchler, Huber and Stöckl (2012) illustrate.

One counter-intuitive aspect of the standard experimental asset market design is that the fundamental value of the asset declines each period. Indeed, the asset becomes worthless by period 15. If we buy an asset in real life we would not expect it to quickly become worthless. Maybe, therefore, subjects were confused about fundamental value, and this explains why we observe bubbles?

Table 5.16 summarizes the outcomes in eight treatments with different instructions or a different market design. Detailed are the relative absolute deviation (RAD) and relative deviation (RD) of price from fundamental value (FV). The baseline treatment is similar to that we have looked at so far (in Figures 5.8 and 5.9). Here there is a large RAD, indicating that prices differed a lot from fundamental value, and there is a positive RD, indicating that prices were higher than fundamental value. So, we observe bubbles.

In the constant-FV treatment the experiment design is tweaked so that the fundamental value of the asset does not fall. Here we observe a smaller RAD and negative RD. So, no bubbles. In the graph treatment subjects were provided

Table 5.16 The relative absolute deviation and relative deviation of prices from fundamental value in eight asset market treatments.

Treatment	RAD	RD
Baseline	0.41	0.30
Constant FV	0.08	–0.06
Graph	0.23	–0.07
Gold mine	0.16	–0.07
Constant cash ratio	0.31	–0.04
Constant FV and cash ratio	0.03	–0.03
Rekindle with gold mine	0.16	–0.06

with a graph of fundamental value so that they could see it was declining. No bubbles. In the questionnaire treatment subjects were asked to give an estimate of fundamental value at the start of each market period. No bubbles. In the gold mine treatment subjects were given the frame of a depletable gold mine to make the idea of declining fundamental value more intuitive. No bubbles. This holds even if the market is rekindled (see the next section to see what this means).

So, can bubbles in the lab simply be dismissed as due to confusion? A theme throughout the book is that framing matters. It is not, therefore, too much of a surprise that changes in the way a situation is framed might change how people behave, or expect others to behave. The difficulty we face is to distinguish genuine framing effects from confusion. Our task is made even more difficult when looking at bubbles, because we actually want to study confusion! Consider, for example, a trader who bought at the peak of the dot-com bubble. This trader must surely have been confused as to the fundamental value of tech companies. Again, it is not too much of a surprise that if we inform subjects enough we can eliminate bubbles, particularly as double auctions are very efficient (see section 2.5).

In order to study bubbles in the lab we may, therefore, want an element of confusion. This is illustrated by the two treatments in Table 5.16 that I have not yet mentioned. One aspect of the standard asset market experiment is that the amount of cash increases through the periods as dividends are paid back into the market. What happens if we keep the cash ratio constant? This, of itself, is enough to stop bubbles happening. That means that bubbles are not solely caused by confusion about a declining fundamental value.

5.6.2 Experience and bubbles

One thing that might stop bubbles happening is experience. When the price crashed, in Figure 5.9, some of the subjects lost money. Will they be reluctant to fall for the same trap a second time? Figure 5.10 plots what happened when nine subjects from

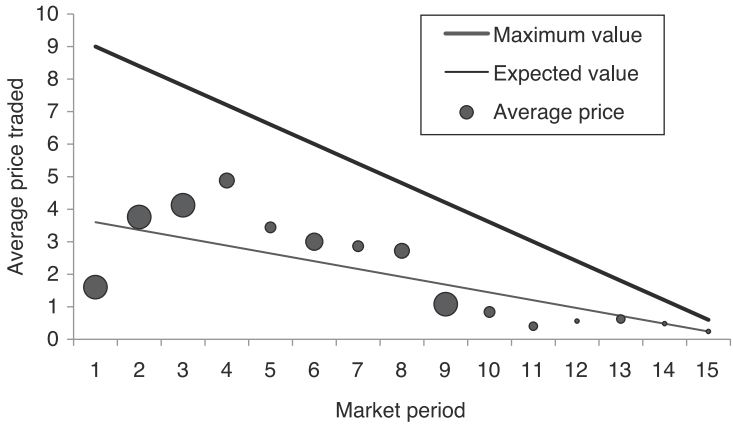

Figure 5.10 The second time that the group takes part in an experimental asset market the bubble is much reduced.
Source: Smith, Suchanek and Williams (1988).

the session that generated the data in Figure 5.9 came back to take part in a second asset market (see Research methods 5.5). As we can see, the bubble and bust are much reduced. In general, it seems that experience is enough to stop bubbles in some situations. For instance, bubbles are rare if subjects have taken part in two previous asset markets in the same group. It is clear, however, that experience in itself is not enough to always stop bubbles happening.

❖ RESEARCH METHODS 5.5 EXPERIENCED TRADER VERSUS EXPERIENCED GROUP

In most experiments we want subjects who have no prior experience. They have not taken part in the experiment before, they do not know the other subjects, they have not studied the experiment in an economics class, and so on. Sometimes, though, it is useful to have something different. In the study by Smith, Suchanek and Williams all subjects had taken part in a double auction before taking part in this study. There may be practical advantages to this. For instance, it is not necessary to spend so long on the instructions, and still be unsure how much subjects really understood. There are also interesting economics issues that can be addressed with experienced subjects, such as whether experienced traders are less susceptible to bubbles.

There are various ways that experience can be brought to the lab, as the study by Smith, Suchanek and Williams demonstrates. In one experimental session, businesspeople from the Tucson community were subjects. In this case subjects were experienced in 'real-world' trading. In some sessions, nine of the 12 who had taken part in an earlier session took part again. In this case the subjects had experience of trading, and of trading with other

people in that group. In two further sessions, subjects were given experience in an environment in which bubbles could not happen before being put in an environment in which bubbles could happen. In this case there is experience of the mechanisms of the market but not of the possibility of bubble and bust. Finally, other sessions involved a combination of inexperienced subjects and subjects who had taken part in at least two previous asset markets and performed well. Here it is possible to combine subjects who had previous experience of bubbles with ones who had no previous experience of bubbles. There are clearly lots of possible combinations and permutations – all of them interesting.

To illustrate, I will look at a study by Hussam, Porter and Smith (2008). Subjects were invited to take part in three experimental sessions. In the first session they were all inexperienced. In the second session everything was kept the same as the first session, both the market and the other subjects. In the third session the market was changed (e.g. the expected dividend was changed) and subjects were put into a new group. So, subjects were given time to learn and then put in a very different environment. Would bubbles reappear in the new environment? As a control it is necessary to have some subjects who interacted together in all three sessions. This gives three treatments, as shown in Table 5.17.

A snapshot of the results is provided in Figure 5.11. The figure shows what happened in three sessions, one from each treatment, after subjects had become experienced. In the baseline treatment we see no bubble, and so it looks as if experience does matter. In the rekindle treatment, however, we see a bubble, and so it seems that bubbles can reappear if the environment changes. This could be due to a change in the market or a change in group, and the extreme treatment suggests it is a bit of both. Here we also see a bubble that is as big but lasts less time.

It is a little bit cheeky to give the data from just three sessions, but I chose these sessions – trust me – because they illustrate the more general pattern we observe. Experience in the same environment reduces the duration of bubbles but not their amplitude. Sometimes this is enough to stop bubbles happening, but not always. Experience, therefore, does not stop bubbles happening. Indeed, nothing seems guaranteed to stop bubbles happening (see Research methods 5.6).

Table 5.17 The three different treatments in an experiment with experienced subjects.

Treatment	1st session	2nd session	3rd session
Baseline	Basic market	Basic market	Basic market, same group
Extreme	Extreme market	Extreme market	Extreme market, same group
Rekindle	Basic market	Basic market	Extreme market, different group

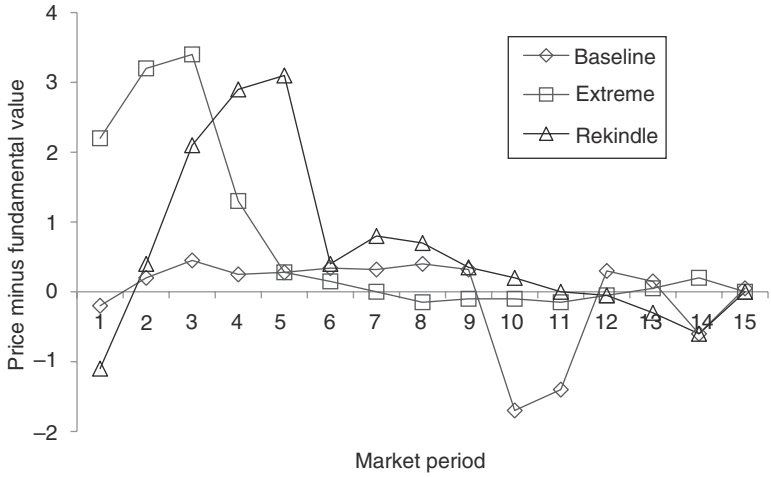

Figure 5.11 The deviation of price from fundamental value when subjects have already experienced two markets. In the baseline market there is no bubble. In the extreme market there is a bubble. In the rekindle market there is a bubble of longer duration.
Source: Hussam, Porter and Smith (2008).

❖ RESEARCH METHODS 5.6 WHY DO EXPERIMENTS?

The experimental work on asset market bubbles provides a good example of why experiments are worth doing. Bubbles and crashes are ubiquitous and matter, so we need to know why they happen. But trying to discern this from real markets seems like a hopeless task, because they are so complicated. In the experimental laboratory we can simplify down to something that is manageable.

I think that, to most people, it is surprising how easily bubbles can be created in very simple markets, even with experienced traders. It is this surprise that shows why experiments are worth doing. Economics is full of very complicated theories for why we might get bubbles. These might be correct, but we observe bubbles in experiments with markets far too simple for these theories to make sense. So, something more basic is going on, even if that is simple confusion (see Research methods 5.4). Realizing that seems fundamental to making any progress in truly understanding why bubbles happen.

This logic will work only if the bubbles we get in the laboratory resemble those in real markets, however. Fortunately, they appear to. For instance, Henker and Owen (2008) looked at data from the Australian stock exchange, and find similarities with the laboratory data.

5.6.3 Explaining bubbles

Having seen bubbles, we now need to try and work out why they happen. Let's start by saying that the bubbles we have seen are not necessarily damning evidence against the efficient market hypothesis. That is because the trading price does eventually equal the fundamental value. For instance, it gets there by period 14 in Figure 5.9, period 13 in Figure 5.10 and period 12 in Figure 5.11. So, we could argue that subjects just take time to learn. This is particularly plausible if subjects start with different initial beliefs and it takes time for these beliefs to converge.

There are good reasons to think that bubbles are caused in part by biases in interpreting new information, however. The law of small numbers would suggest that, if the price has gone up three or four periods in a row, it should keep on rising. Confirmatory bias would suggest that, if people expect prices to rise, they will interpret a slight fluctuation in price as a signal that this is going to happen. Putting these and other biases together gives something called feedback theory. The basic idea of **feedback theory** is that prices rise because people expect them to rise, and they expect them to rise because they have been rising! A rising price translates, therefore, into a rising price.

Feedback theory can explain bubbles, and comes about because people are biased by what they see happening to the price. What it cannot do is explain the crash. Why do prices not keep on rising? Well, there are only so many investors, and so prices cannot keep on going higher and higher, because eventually no one is left to buy. In the laboratory, a crash is usually preceded by a period with fewer trades (see Figures 5.9 and 5.10). Basically, therefore, bubbles run out of steam. In real markets there is a less pronounced drop-off in volume than in the laboratory, presumably because real markets are not going to end after a known 15 periods, but the logic still holds.

To give a bit more substance to the feedback theory, I want to finish with an interesting study by Barber and Odean (2008). They looked at data from individual and institutional investors to see whether the news matters. The news might matter because of various reasons, but one is the availability heuristic. When choosing where to invest, people might remember seeing a stock in the news and choose that one. If stocks get into the news because their price has increased this can cause a feedback of prices rising because prices are rising.

One key to Barber and Odean's empirical strategy is the difference between buying and selling. When investing, the investor has thousands of stocks to choose from, but when he comes to sell, he has only those stocks he has invested in to look at, and these probably number at most 100. So, we should expect investors to be more influenced by the news when they buy than when they sell. For institutional investors this is less likely to matter, because they can short-sell stocks they do not own. More generally, we might expect professional, institutional investors to be less influenced by the news.

If we expect buying to be influenced by the news and selling to not be then the buy–sell imbalance is a useful measure. For a particular stock on a particular day this is calculated as follows:

$$\text{Buy–sell imbalance} = \frac{\text{Number of buys} - \text{Number of sells}}{\text{Number of buys and sells}}.$$

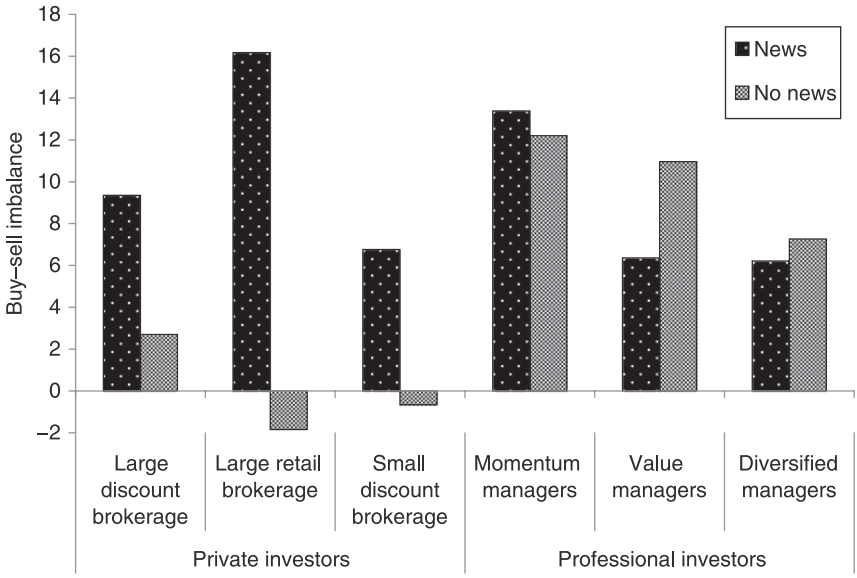

Figure 5.12 The buy–sell imbalance for different types of investors depending on whether the stock appeared in the news or not.
Source: Barber and Odean (2008).

Barber and Odean consider various measures of a stock being in the news, but we will just consider one: whether the stock appeared on the daily news feed from the Dow Jones News Service on a particular day. Figure 5.12 summarizes the result. We see that, for private investors, there is a large difference in the buy–sell imbalance depending on whether the stock was in the news or not that day. For institutional investors, there is no such difference. It would appear, therefore, that people are influenced by what stocks have appeared on the news when deciding where to invest. Interestingly, this effect becomes even more pronounced if the stock went down in value that day (i.e. the news was likely to be bad) than if it went up!

That people are more likely to buy stocks that have appeared on the news brings us nicely on to something called the beauty contest, which we shall look at in detail in the next chapter. A slightly different form of beauty contest is voting in elections.

5.7 Voting in elections

Hillary Clinton's campaign to become US president in 2016 cost well over $1 billion. This is a staggering amount of money. Was it worth it? An argument to defend such spending might go something like this. It is crucial we elect the best candidate for office, and so all the money spent on political rallies, TV adverts, campaign leaflets, and the like is essential to inform the electorate. This argument might have some merit if the electorate are ready to be informed, but we have seen plenty of examples in this chapter that people have problems coping with new information. So, do people form an informed opinion before going out to vote?

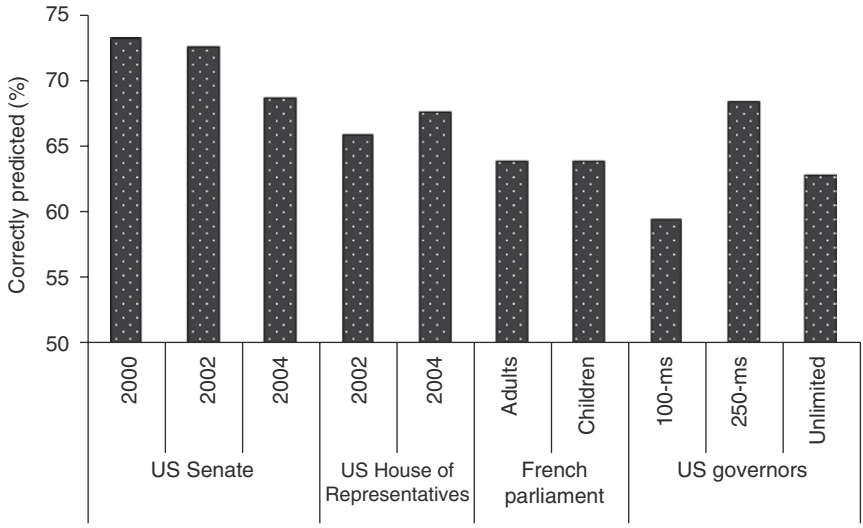

Figure 5.13 The proportion of outcomes that were correctly predicted by looking at photos of the two candidates. In the US gubernatorial elections, subjects were given either 100ms (a tenth of a second), 250ms or unlimited time.
Sources: Todorov *et al.* (2005), Ballew and Todorov (2007) and Antonakis and Dalgas (2009).

There are now a collection of studies that provide a reliable way of predicting election outcomes. The method is to show people photos of the two candidates for a second or less and ask them who they think will win. You do not need to give any information about political persuasion, or political views, experience, etc. One glance is all it takes!

To illustrate how this works, consider first a study by Todorov and co-authors (2005). They showed subjects the faces of rival candidates for one second and asked them to say who was more competent. Inferences of competence were far better than chance at predicting the election outcomes for both the US Senate and House of Representatives; the results are given on the left-hand side of Figure 5.13. They also find that the inferred difference in competence was able to predict the margin of election victory. Clearly, such findings are controversial, but they seem robust. For instance, judgments on other traits, such as intelligence, leadership, honesty, trustworthiness, charisma and likability, do not predict the election outcome. Only competence seems to matter.

In order to demonstrate the robustness of these findings, a study by Antonakis and Dalgas (2009) is particularly worth a mention. They showed that children aged 13 or below were able to predict the outcomes of French parliamentary elections (see Figure 5.13 for the results). This shows that the ability to predict elections is innate and not a quirk of the US electoral system (the children in the study were from Switzerland). In addition, it is gut instinct that seems to matter, as shown in a study by Ballew and Todorov (2007). They gave some subjects 100ms, or one-tenth of a second, to look at the photos. Others got 250ms, or a quarter of a second, and yet others got as long as they wanted. As shown on the right-hand side of Figure 5.13, a quarter of a second was enough. Having time to think decreased predictive power!

What to make of such findings? They do not imply that people go to vote without taking account of what the candidates say and stand for. Many people take voting as a serious responsibility. Recall, though, the distinction between system 1 and system 2 brain processes introduced in Chapter 2. Making a judgment on the basis of appearance is clearly a system 1 process: it is automatic and based on spontaneous judgment. When we go out to vote, or listen to a candidate speak, system 2 processes begin to work as we reason through the pros and cons of each candidate. Crucially, however, this reasoning is inevitably influenced by the perceptions and intuition gained from first impressions. System 2 starts from where system 1 leaves it. So, a candidate who does not look competent has some catching up to do.

Will he be able to catch up? We have primarily focused on system 2 reasoning processes throughout this chapter. What we have seen suggests that it may be very difficult for a candidate to play catch-up. Take confirmatory bias, for instance: if a candidate is inferred as less competent then any new information can easily be interpreted as confirming initial impressions. Those inferred as competent will appear more competent than they really are and those inferred as not competent will appear less competent than they really are. Or, consider conformity: if a voter talks to friends about the election, he may be swayed by their views, even though those views are based on first impressions. Do not expect, therefore, thoughtful reasoning to overcome spontaneous inference. Moreover, do not expect the better candidate to win the election, and do not expect campaign spending to make much of a difference to the outcome.

The more general point to take away from this discussion is that bias can happen in both system 1 and system 2 brain processes. This is an important point to keep in mind, because it is common to think that system 1 processes are the cause of all bias. To illustrate the distinction, let us look at two examples. In Chapter 2 I discussed the endowment effect, whereby physically handling a mug made people value it more. This bias can be put down to system 1 processes; the more a person reasons things through, the smaller the bias is likely to become. The picture is, thus, one of intuition creating a bias that reasoning corrects. Compare this to confirmatory bias in judging electoral candidates. Intuition might suggest the candidate looks good, and confirmatory bias makes him look even better. In this case, intuition creates a bias that reasoning amplifies. Reasoning can, therefore, make matters worse!

5.8 Further reading

Tversky and Kahneman (1974) provide an interesting introduction to the psychology of choice under uncertainty. Camerer and Weber (1992) provide a comprehensive survey of economic research on uncertainty and ambiguity. Rabin (1998) gives a good overview of many of the biases we have looked at. Oskarsson and co-authors (2009) provide a comprehensive survey of the literature on the gambler's fallacy and hot hand beliefs. Also worth a read is Kahneman and Tversky (1973). One consequence of the law of small numbers and confirmatory bias can be prejudice and stereotyping. The book by Brown (1995) provides a fascinating and thought-provoking introduction to these issues. Bikhchandani, Hirshleifer and Welch (1998) provide an introduction to information cascades. Frank (2007) is a great starting point for more on health care. The book by McCall (1996) is also worth a read. For more on bubbles and stock

market movements, I would recommend the excellent surveys by Caginalp, Porter and Smith (2000), Porter and Smith (2003) and Shiller (2003). See also Dufwenberg, Lindqvist and Moore (2005) for the first study to look at experience and bubbles. For another take on the role of confusion in experiments, see Andreoni (1995).

5.9 Online materials

On the companion website you will have access to flashcards, quizzes and answers to the review questions.

5.10 Review questions

5.1 Why might people dislike compound lotteries or uncertainty?

5.2 Ellsberg (1961) suggested another box in which there are 90 balls. You know that 30 of the balls are red and the other 60 are some mix of black and yellow. One ball will be randomly drawn from the box. Would you prefer to 'bet the ball is red' or 'bet the ball is black'? Would you prefer to 'bet the ball is red or yellow' or 'bet the ball is black or yellow'? Were your choices consistent with subjective expected utility maximization?

5.3 Why might someone who is ambiguity-averse delay making a decision? How does this relate to procrastination because of a present bias?

5.3 Think back through some recent decisions you have made. Were they choices with risk or uncertainty?

5.4 There are no experiments to test for confirmatory bias in economic behavior. Design one. And then do one to test for the law of small numbers.

5.5 Why might people be overinfluenced by what others do in some situations, and underinfluenced in others?

5.6 Carefully explain the difference between a signaling equilibrium and pooling equilibrium in a signaling game.

5.7 Can we trust people's opinions? Why might people trust friends over expert advice?

5.8 Consider reasons why confirmatory bias and the law of small numbers may matter in health care.

5.9 How can we address the problems caused by biases in the interpretation of new information in health care, and more generally?

5.10 Carefully explain why small area variation is evidence of inefficiency in health care?

5.11 When should someone buy an asset that is rising in price? When should they sell it?

5.12 Why do we see bubbles and crashes in the housing market?

5.13 A standard asset market experiment involves choice with risk because the probability distribution of future dividends is known. What do you think would happen if the probability distribution over future dividends was uncertain?

5.14 What do you think happens if some traders have experience of bubble and crash but some do not?

5.15 Given that double auctions are known to be efficient in the lab, is it possible to create bubbles in the lab without subjects being confused?
5.16 How does the buy–sell imbalance relate to the ostrich effect?
5.17 Come up with a list of reasons why a person's opinion of electoral candidates may be biased. On what basis should the person decide whom to vote for?
5.18 Provide examples from health care, finance and voting in which the availability heuristic may influence decisions.
5.19 Will the rise of social media lead to more or less bias in voting decisions?
5.20 In 2017 the phrase 'alternative facts' was used following disagreement over attendance at President Trump's inauguration. Why may it be difficult to resolve conflicting opinions and reach a consensus on 'the truth'?

Chapter 6

Interacting with others

> For it is, so to speak, a game of Snap, of Old Maid, – a pastime in which he is victor who says *Snap* neither too soon nor too late, who passes the Old Maid to his neighbour before the game is over, who secures a chair for himself when the music stops.
>
> John Maynard Keynes, *General Theory of Employment, Interest and Money*

Economics is fundamentally about exchange and interaction between people. We can capture some of that by saying that people interact through markets, but a lot of the time that is not enough. We need to think how people interact with each other more directly, and this will be the theme of the next two chapters. In this chapter I want to focus on strategic behavior, which primarily means focusing on how people try to predict what others might do. For example, when should a trader sell her investments to ride the bubble but avoid the crash? What route should she take to work to avoid traffic jams? How much effort should she put into keeping her office tidy? And what should she choose for lunch at the canteen? In all these examples, the trader would do best to predict what others will do, whether they are fellow traders or the canteen chef, before making her choice. In this chapter I want to look at how she can do this.

6.1 The beauty contest

In this first section I want to set the scene a little for the rest of the chapter by looking at a game called the beauty contest. In the **simple beauty contest**, people are independently asked to say which out of a set of faces is the most beautiful. The face with the most votes is declared the most beautiful, and anyone who voted for the most beautiful face is eligible for a prize. John Maynard Keynes said that this contest was a useful way to think about how stock markets behave. His basic argument was that to pick the most beautiful face is a relatively naïve strategy. More sophisticated is to pick the face you believe most others will pick. This maximizes your chances of winning the prize. The stock market is similar, except that we substitute beautiful companies for beautiful faces. The fundamental value of a company may, therefore, be less relevant in determining the movement in share price than popularity.

We are going to look in more detail at a slightly different version of the beauty contest. To motivate this version, imagine that someone called Emma holds shares in

a company called FlyMe. At the moment the share price is increasing, which is great, but it is also well above fundamental value. Emma knows the price is probably going to crash, but she wants to hold on to the shares while the price is still rising. The decision she has to make is how long she is going to hold on to the shares before she sells them.

What she will probably do is track the share price and see what happens. That does not make for a very simple story, however. So, we are going to assume instead that she has to decide today when she will sell the shares, and after that she cannot change her mind. She could decide to sell them anytime between today, zero days and 100 days from now. Suppose that everyone else who holds shares in FlyMe also has to make the same decision. When is the optimal time to sell?

The 'it' in the quote with which I started this chapter is 'investing in the stock market'. Keynes explains that a trader needs to be 'neither too soon nor too late', and we saw this when looking at bubbles and crashes in the last chapter. The crucial thing is that what is too soon or too late will depend on others, because the optimal time to sell will depend on how long others delay. A simple way to capture this is to assume that the optimal time to sell depends on the average number of days that other investors choose to delay.

This is what happens in a **p-Beauty Contest**, in which the objective is to be as close as possible to p times the average chosen by others. A p less than one makes most sense, because then Emma does best to wait less than the average, and so is selling before most others. For example, if the average trader delays 30 days and $p = 0.67$, the optimal delay is 20 days. This way Emma does not sell too early, to avoid the entire rise in price, but sells early enough to avoid the crash. To understand what Emma should do in this game (given that she does not know what others are doing), we need to introduce some basic concepts from game theory.

6.1.1 Strategy and Nash equilibrium

One thing we shall talk a lot about in this chapter is strategy. Recall (see Chapter 1) that a **strategy** is best thought of as a person's plan of what they will do. In the simple beauty contest, Emma's strategy is which face to choose. In a p-beauty contest, Emma's strategy is how long to delay. Emma's payoff will depend on her strategy, and the strategy of others.

When it comes to determining payoff in the p-beauty contest there are different possibilities. In the most commonly used version the investor closest to the optimal delay earns a positive payoff and everyone else gets nothing. An alternative version will be more useful for our purposes, in order to illustrate the concepts that follow. In this version the investor gets a higher payoff the closer she is to the optimal delay. For example, in a 0.5-beauty contest with n investors, we could write the payoff of Emma as

$$u_E(s_E, s_{-E}) = 100 - \left| s_E - 0.5 \times \left(\frac{s_1 + s_2 + \ldots + s_n}{n} \right) \right|,$$

where s_E is her strategy and s_{-E} is a shorthand way of writing down the strategies of everyone but her. She wants, therefore, to be as close as possible to 0.5 times the average, and if she gets it spot on her payoff is 100.

One of the most basic concepts in game theory is that of a dominated strategy. To illustrate the idea, let's think about the 0.5-beauty contest. Given that investors can delay at most 100 days, the highest the average delay can possibly be is 100 days. This means that the highest the optimal delay can possibly be is 0.5 × 100 = 50 days. Does it, therefore, make sense for Emma to delay more than 50 days? No; by delaying 50 days she is guaranteed doing at least as well as if she delays 51, 52 or more days. This does not mean she should necessarily want to delay 50 days, but it does mean she would be unwise to delay more than 50 days; the strategy of waiting 50 days dominates any strategy of waiting more than 50 days.

More generally, we say that a strategy *l* **dominates** a strategy *j* for Emma if she always gets a higher payoff from *l* than from *j* – that is,

$$u_E(l,s_{-E}) > u_E(j,s_{-E})$$

for any s_{-E}. The prediction is that people should not play a dominated strategy. Intuitively, this makes sense, because strategy *l* is guaranteed to give a higher payoff than *j*, so why play *j*?

Now, suppose that Emma believes no one would ever play a dominated strategy. For instance, she does not expect anyone to delay more than 50 days in a 0.5-beauty contest. If no one is going to delay more than 50 days then the optimal delay cannot be more than 0.5 × 50 = 25 days. This means that delaying 25 days seems better than delaying more than 25 days. We say that the strategy to delay more than 25 days is **one-step dominated** by a strategy to wait 25 days. The one step is there to acknowledge that Emma needed to go through one step of reasoning to get to this; first she eliminates anyone waiting more than 50 days, and then she eliminates waiting more than 25 days. But why stop there? If Emma believes that no one would ever play a one-step dominated strategy and delay more than 25 days, then there is no sense in delaying more than 12.5 days. The strategy to delay more than 12.5 days is **two-step dominated** by a strategy to wait 12.5 days.

More generally, we say that strategy *j* is **one-step dominated** by strategy *l* for Emma if she always gets a higher payoff from *l* than from *j* if no one plays a dominated strategy – that is,

$$u_E(l,s_{-E}) > u_E(j,s_{-E})$$

for all s_{-E} where no one plays a dominated strategy. Two-step dominates and three-step dominates are defined in an analogous way.

To continue in this way is called **iterated deletion of dominated strategies**. There is a logic that says people should play only strategies that **survive** iterated deletion of dominated strategies. What strategies survive iterated deletion of dominated strategies in the 0.5-beauty contest? There is only one, and that is to delay zero days. We get, therefore, a very specific prediction of what should happen in this game.

This prediction is an example of a Nash equilibrium. Recall that a **Nash equilibrium** is a list of strategies s_1 to s_n for each person such that no one could change their strategy and do better. That is,

$$u_i(s_i, s_{-i}) \geq u_i(j,s_{-i})$$

for all people *i* and any strategy *j*. A dominated strategy cannot be part of a Nash equilibrium, and so, if there is only one strategy that survives iterated deletion of dominated strategies, that strategy must be a Nash equilibrium strategy. In the 0.5-beauty contest, for example, the unique Nash equilibrium is that everyone chooses zero.

If, as in the 0.5-beauty contest, there is only one strategy that survives iterated deletion of dominated strategies then we get a nice and simple prediction. Unfortunately, things do not always work out so nicely. Sometimes eliminating dominated strategies does not get us very far, and we end up with multiple Nash equilibria. We have already come across this in the previous chapter, where there was an intuitive and an unintuitive equilibrium in a signaling game.

To give another example, think about a 1.5-beauty contest, in which it is optimal to choose 1.5 times the average. What should Emma do if everyone else chooses zero? The best she can do is to also choose zero, because $0 = 1.5 \times 0$ and she will be spot on the optimal delay. What should Emma do if everyone else chooses 100 days? The best she can do is to also choose 100, because that is as close as she can get to the optimal delay of 150 days. 'Everyone delays zero days' and 'Everyone delays 100 days' are both Nash equilibria, and we can tell from this that there is no dominated strategy in this game.

If there are multiple Nash equilibria, we get a less precise prediction of what will happen. But we still get a prediction. The natural question is whether dominance and Nash equilibrium give us good predictions.

6.1.2 Choice in a beauty contest

To recap: in a p-beauty contest, the Nash equilibrium is for everyone to choose zero if $p < 1$, and either for everyone to choose zero or everyone to choose 100 if $p > 1$. Is this what people actually do? Figure 6.1 gives some data from a study by Nagel (1995). It shows what numbers subjects chose the first time they played a p-beauty contest. A couple of things stand out from this data. On the one hand, it looks a mess, because there is a lot of heterogeneity in what subjects did, and very few chose the Nash equilibrium strategy. On the other hand, there are clear peaks in the data that change as *p* varies. So, people were not playing the Nash equilibrium, but they were seemingly responding in a systematic way to the incentives of the game.

Why does the Nash equilibrium turn out to be such a bad predictor in this game? One thing we can do is to see whether people play dominated strategies or not. Let's say that a person is **type D-0**, if they never play a dominated strategy, but may play a one-step dominated strategy. Similarly, someone is a **type D-k** if they never play a k-step dominated strategy, but may play a $k + 1$-step dominated strategy. By looking at Figure 6.1 we can start to work out how many subjects seemed to be of each type. A study by Ho, Camerer and Weigelt (1998) does this in a more structured way by looking at four different beauty contests. Their results are summarized in Table 6.1, where lower bounds are put on the proportion of subjects of each type. They are lower bounds because someone who does not realize that a strategy is dominated may still not play it, by accident.

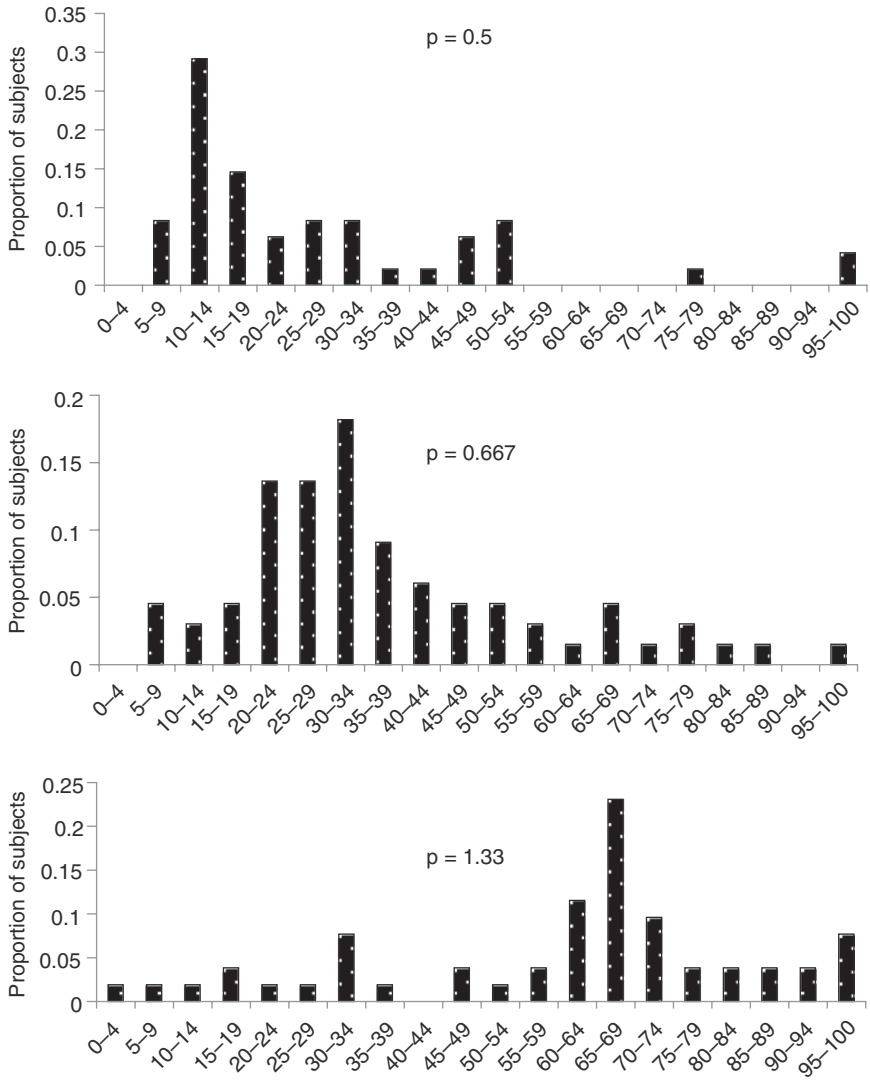

Figure 6.1 Choices the first time that subjects play a p-beauty contest.
Source: Nagel (1995).

The results in Table 6.1 are broadly consistent with what we see in many other games. Relatively few people play dominated strategies, but most people go through only one, two or three steps of iterated deletion of dominated strategies. To say that people do not play dominated or iterated dominant strategies does, therefore, give us some power to predict what people will do, but we are going to need more than this to fully understand what is going on (see also Research methods 6.1). For example, looking at Figure 6.1, we have an explanation for why so few people choose more

Table 6.1 Proportion of subjects of each type in four p-beauty contests.

Game	Type (%)					
	Played dominated	**D-0**	**D-1**	**D-2**	**D-3**	**D-4 +**
$p = 1.3$	15.7	36.4	36.1	11.8		
$p = 1.1$	4.3	3.2	5.1	9.8	35.0	42.6
$p = 0.9$	4.4	2.6	8.4	6.2	12.0	63.4
$p = 0.7$	15.0	23.2	18.9	12.5	13.9	16.5

Source: Ho, Camerer and Weigelt (1998).

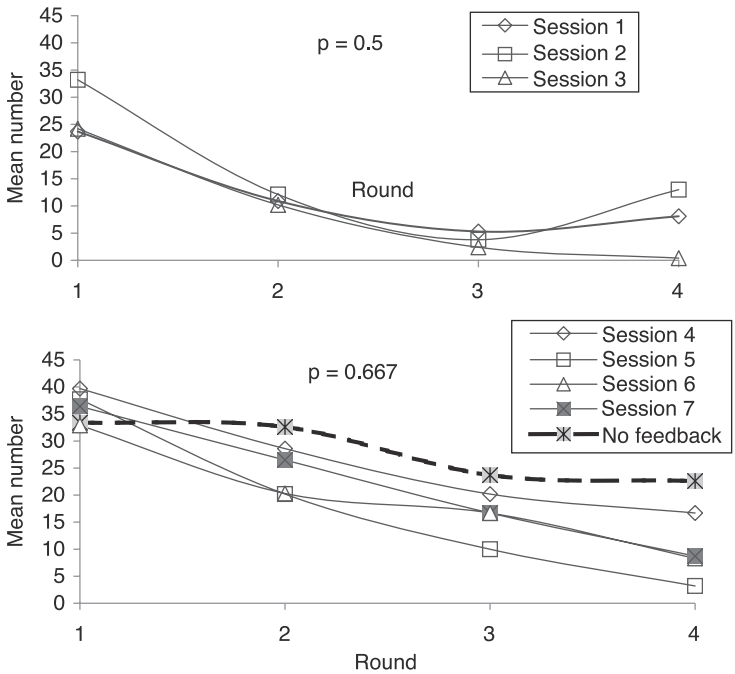

Figure 6.2 The mean number chosen when playing a p-beauty contest four consecutive times, for each of the experimental sessions.
Sources: Nagel (1995) and Weber (2003).

than 50 days when $p = 0.5$ or 0.667, or fewer than 50 when $p = 1.33$. We are still lacking an explanation for the choices people did make. Arguably, though, this is natural enough, because the focus in looking at dominated strategies is on what strategies people will not use, and not on what they will.

6.1.3 Learning in a beauty contest

The next issue I want to look at is what happens when the same game is played a second, third and fourth time. Figure 6.2 illustrates what happened in the study by Nagel. There is a clear downward trend in the average number chosen. This suggests

that subjects learnt from experience, and behavior is converging towards the Nash equilibrium. The Nash equilibrium becomes a better predictor, therefore, if people have some experience with the game. Just to illustrate how intriguing this issue can be, in Figure 6.2 I also included data from a no-feedback treatment from a study by Weber (2003). In the experiments by Nagel, subjects were told after each round what the winning number was before they played again. Most subjects guessed too high and so lowered their choice next round, making things converge towards the Nash equilibrium. In the no-feedback treatment, subjects were not told the winning number before playing again. Subjects could not, therefore, have known that they guessed too high. Still, however, the choice decreased over the rounds!

This discussion of the beauty contest helps motivate the three key themes I want to focus on in the rest of the chapter. First, I want us to think more about what people do the first time they play a game, such as the beauty contest. How can we explain the data in Figure 6.1? Second, I want us to think more about how people learn when playing a game several times. How can we explain the data in Figure 6.2? Third, I want us to think more about the relevance of the Nash equilibrium, particularly in situations when there are multiple Nash equilibria. Once we have done all that we should have a good picture of how people behave in strategic situations.

6.2 Playing for the first time

I will begin by looking at how we can model what people do when they play a game for the first time. A useful starting point for doing that is something called level-k thinking.

6.2.1 Level-k thinking

The basic idea behind level-k thinking is that someone, such as Emma, forms an expectation of what others will do and tries to be 'one step ahead of them'. In a p-beauty contest this is nice and easy to apply: Emma is assumed to form an expectation of the average delay of others and then guesses p times that average. For example, if she thinks the average will be 50 she will guess $50p$, and if she thinks the average will be $50p$ she will guess $50p^2$.

In a general model of **level-k thinking** we start with people of **type L-0**, who are assumed to choose without thinking too much about what others will do (more on this below). We then say that someone such as Emma is of **type L-1** if she thinks that all others will be of type L-0, and maximizes her payoff given this belief. We say that she is of **type L-k** if she thinks that all others will be of type L-k-1, and maximizes her payoff given this belief.

The type L-0s are crucial here because they determine what the L-1 types will do, which determines what the L-2 types will do, and so on. Unfortunately, there is no simple rule for saying what the L-0 types will do. Two possibilities are: (1) a **type L-0 randomly chooses** a strategy because they are maybe confused, misunderstood the game or are simply uninterested; and (2) a **type L-0 chooses some salient or focal strategy**, such as the Nash equilibrium. In neither case does a type L-0 take account of what others will do. We shall see, however, that a different assumption

Table 6.2 Level-k thinking in the p-beauty contest if type L-0s choose randomly, compared to choosing the Nash equilibrium.

Type	Type L-0s choose randomly			Type L-0s choose Nash equilibrium		
	p = 0.5	p = 0.67	p =1.33	p = 0.5	p = 0.67	p = 1.33
L-0	50	50	50	0	0	100
L-1	25	33.5	66.5	0	0	100
L-2	12.5	22.45	88.44	0	0	100
L-3	6.25	15.03	100	0	0	100

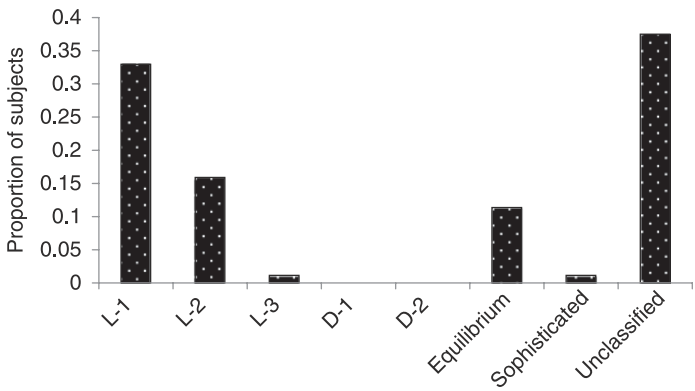

Figure 6.3 Classifying subjects by type after they had played 16 different games. Many subjects were best classified as of type L-1 or L-2.
Source: Costa-Gomes and Crawford (2006).

about type L-0s can lead to very different conclusions about what everyone else will do.

To illustrate, we can work through the p-beauty contest with the help of Table 6.2. If type L-0s choose randomly then, on average, they will choose 50. This means that Emma will choose $50p$ if of type L-1 and $50p^2$ if of type L-2, which equals 25 and 12.5 in a 0.5-beauty contest. Now, suppose, by contrast, the L-0 types choose the Nash equilibrium of 0 or 100. Maybe, for instance, they have just gone to a game theory class and been told this is the equilibrium. Then Emma should also choose the Nash equilibrium, and we get a very different prediction. Clearly, some thought needs to go into what the type L-0s will do, and different situations will suggest different ways they might behave. Having done this, however, we have a simple way to make predictions of what people will do. Are these predictions any good?

You may already have noted that the spikes in Figure 6.1 at 12.5, 33.33 and 66.67 suggest that we are on the right track. A study by Costa-Gomes and Crawford (2006) provides a more detailed test. In the study, 88 subjects played 16 two-player guessing games similar to a beauty contest. The 16 games were chosen to try and distinguish how subjects were reasoning and classify them as of type L-k, D-k, someone who plays the Nash equilibrium or someone who is sophisticated (which I shall explain below). Figure 6.3 summarizes the main results.

We see that many subjects were best classified as of type L-1 or L-2. It seems, therefore, that level-k thinking does a relatively good job of capturing what people do. By contrast, no one was best classified as of type D-1 or D-2. This is not inconsistent with what I said in section 6.1.2 about dominated strategies, because a level-k thinker will never play a dominated strategy. So, we are still correct in saying that most people do not play a dominated strategy; we just have a sharper prediction now of what they will play. Before we question more how good a prediction that is, I should explain what is meant by someone being sophisticated.

❖ RESEARCH METHODS 6.1 SEEING HOW SUBJECTS CHOOSE

In behavioral economics we are often interested in how people decide what to do. Experiments typically show us only what people do, however, and not why they did it. It is routine to add questionnaires to the end of an experiment asking subjects why they chose as they did. But it is also routine to ignore such questionnaires as unreliable, because subjects may not know why they did what they did. Technology is starting to offer more possibilities.

The study by Costa-Gomes and Crawford used a program called Mouselab that can track where subjects click or put the cursor on the computer screen. This makes it possible to know what subjects would have known when making their choice and what they wanted to know before making their choice. This is useful in trying to understand how people decide, and in Figure 6.3 I report results that make use of this data. More technologically demanding ways of recording what subjects do include eye tracking (see Research methods 2.5) and brain imaging (see Chapter 9).

6.2.2 Sophisticated beliefs

One slightly concerning aspect of k-step elimination of dominated strategies and level-k thinking is that, the higher *k* is, the worse a person will typically be at predicting what others will do. This is because a type D-k believes that all others will be of type D-k-1 and a type L-k believes that all others are of type L-k-1. The higher k is, the less sensible this might be. For example, in a 0.5-beauty contest, if Emma is of type D-3, she would think that no one will choose more than 12.5 and, if of type L-3, would expect the average choice to be 12.5. These are relatively low numbers. Is it reasonable that Emma expects everyone else to be one level lower than her?

It might be, but, equally, it might not. Indeed, it seems more intuitive for me that she would expect some to be, say, level-0 thinkers, some to be level-1 thinkers, some to play the Nash equilibrium, and so on. This leads us to the idea of sophisticated and semi-sophisticated types.

To illustrate, we can look at the **cognitive hierarchy model** of Camerer, Ho and Chong (2004). The model starts with type L-0s and L-1s exactly as I defined them above. But now we need to keep track of how many people are of each type, so let

$f(0)$ be the proportion of type L-0s in the population, $f(1)$ be the proportion of type L-1s and $F(1) = f(1) + f(0)$ be the proportion of type L-0s and L-1s.

We then say that someone is of **type C-2** if she expects that proportion $f(0)/F(1)$ of others will be of type L-0 and proportion $f(1)/F(1)$ will be of type L-1, and maximizes her payoff given this belief. Thus, if Emma is of type C-2, she expects that some will be of type L-0 and some will be of type L-1. This is different from a type L-2 who expects all others to be of type L-1. Furthermore, if she is of type C-2, she correctly predicts the proportion of type L-0s relative to type L-1s. What she does not do is account for the possible existence of type L-2s, C-2s or beyond.

But there is no need to stop with a type C-2. Let $f(2)$ be the proportion of type C-2s and $F(2) = f(1) + f(2) + f(3)$ be the proportion of types L-0, L-1 and C-2 in the population. Emma is of **type C-3** if she expects proportion $f(0)/F(2)$ to be type L-0, proportion $f(1)/F(2)$ to be type L-1 and proportion $f(2)/F(2)$ to be type C-2, and chooses the optimal strategy given this belief. Continuing in this way, we can define type C-4, and so on. Table 6.3 illustrates how this works in the 0.5-beauty contest.

We can say for sure that, if Emma is of type C-3, she will be better at predicting what others will do than someone of type C-2. Similarly, if she is of type C-4, she will be better than someone of type C-3. The higher, therefore, k is, the better Emma is at predicting what others will do, creating a natural cognitive hierarchy.

Someone of type C-k is still only **semi-sophisticated**, however, because they fail to predict people at the same cognitive level as themselves (or higher). The final step is to consider someone who is **sophisticated** and correctly predicts what all others, including other sophisticated people, will do. That is, if Emma is sophisticated, she will correctly predict the proportion of people of type D-k, L-k and C-k, the proportion of people who play the equilibrium, are sophisticated, and so on. This means that she will play the empirical Nash equilibrium and best respond to what others are doing (see Research methods 5.1).

The issue of sophisticated and semi-sophisticated people brings into sharp focus the relative proportions of each type in the population. That's because we need to clarify what it means to correctly predict the proportions of people at lower, or the same, cognitive levels. For Emma to correctly predict the actual proportions would be asking too much, unless she is telepathic! More sensible is that she predicts the average proportions that one might expect to see. We need, therefore, to think what these proportions might be.

Camerer, Ho and Chong suggest that, the higher k is, the less likely there are to be people of type C-k. One way to capture this is to assume that

$$\frac{f(k)}{f(k-1)}\frac{1}{n} \propto \frac{1}{k}.$$

This implies that $f(k)$ is characterized by a Poisson distribution with parameter τ, and we can write

$$f(k) = \frac{e^{-\tau}\tau^{k}}{k!}.$$

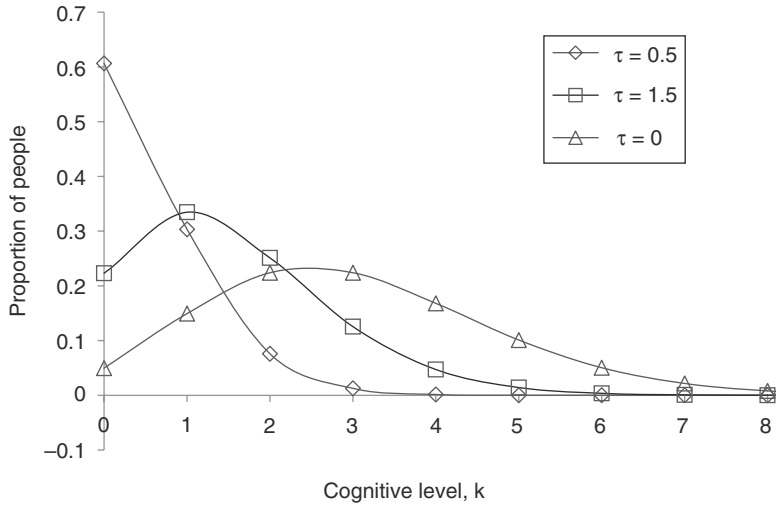

Figure 6.4 The proportion of people at each cognitive level in a cognitive hierarchy.

Figure 6.4 illustrates what proportion of people will be of each type for three different values of τ. We see that, the higher it is, the more levels of reasoning the average person is expected to do. To put this in some perspective, if $\tau = 1.5$ then someone of type C-2 would expect 40 percent of others to be type L-0 and 60 percent to be type L-1.

The neat thing about this model is that everything now depends on just one parameter. If we know τ then we know how many people to expect of each type, and if we know how many people to expect of each type we can predict what everyone will do. With just one parameter we can estimate the full distribution of choices. Fantastic, if it works!

Camerer, Ho and Chong estimated τ for a variety of games, including the p-beauty contest, and suggested $\tau = 1.5$ as a reasonable value to assume. Without going into the details, the cognitive hierarchy model with $\tau = 1.5$ seems to work quite well for some games, and so can give us a very simple and parsimonious way to model choices. For other games the model seems a bit too restrictive, and we need to better capture the mix of L-k, D-k, C-k and sophisticated types. This can easily be done, though, by changing assumptions about the distribution of types. We can, therefore, create models that are as simple or complicated as needs dictate.

Let's not get too carried away just yet, however. If you look back at Figure 6.3 you will see that the biggest spike is over 'unclassified'. With the notions of type D-k, L-k and C-k we are able to better understand and model what a lot of people seem to be doing, but we are still some way from a complete picture. We have also yet to tackle the problem of multiple equilibria. It is to this problem that I now turn.

6.2.3 Focal points

Consider the simple beauty contest that I introduced towards the start of the chapter. Recall that people were asked to pick which of a set of faces they thought most

beautiful, and anyone who voted for the most popular face was eligible for a prize. If we know that Emma is of type L-k, or D-k or sophisticated, can we predict what she will do? Not really; at least, not without doing a bit of work first. That's primarily because there are multiple Nash equilibria.

To explain: in the simple beauty contest, if everyone chooses, say, face two, then the optimal thing that Emma can do is to also choose face two. The same if everyone chooses face one, three, etc. There are, therefore, as many Nash equilibria as faces (plus some ones we don't need to worry about). Basically, it does not matter to anyone what they choose; they just want to choose the same as everyone else. This is called a **pure coordination game**. If people are not allowed to communicate with each other, then it can seem a pretty hopeless task to coordinate in a pure coordination game. Is Emma supposed to just choose one of the faces at random and hope for the best?

Maybe not, if there exists a focal point. Thomas Schelling highlighted the remarkable power of focal points in his book *The Strategy of Conflict* (Schelling 1990). For example, Schelling asked people to imagine that they had to meet someone in New York, but were unable to communicate with that person. Where should they go in the hope of meeting? New York is a big city, so this seems like a difficult problem, but most people named the same place: Grand Central Station. As a further example, Schelling asked people to independently choose between the words 'Heads' or 'Tails' trying to match someone else's choice. Most people (36 out of 42) said 'Heads'.

The point Schelling was trying to make is that people are able to coordinate much better than can be due to pure chance. So, Emma does not need to pick a face randomly. But she does need to know the secret of how to coordinate. There are, at least, three plausible explanations for what this secret is.

The **primary salience hypothesis** is that people choose the option that is most salient to them, and, because the same option is salient to many people, people coordinate by accident. For example, people might have said Grand Central Station simply because it is their favorite place in New York. Similarly, Emma could just say which face she thinks is most beautiful and expect to coordinate with others.

The **secondary salience hypothesis** is that people expect others to use primary salience and so choose the option they think will have primary salience for others. For example, if Emma's favorite place in New York is the Yankee Stadium, but she expects the favorite place of most others to be Grand Central Station, then using secondary salience she will choose Grand Central Station. Similarly, if she really likes face three but thinks that others will like face two, she should choose face two.

Secondary salience may sound like level-k thinking. That's no accident, given that it basically assumes everyone will be of type L-1. It does a little bit more than this, however, because it ties down what someone of type L-0 will do. Someone of type L-0 is assumed to (3) choose their favorite option, or the option that is primary salient. Recall that, when discussing the p-beauty contest, we assumed that type L-0s would choose randomly or choose the equilibrium. We now have, therefore, a third possible assumption for what type L-0s might do. The assumption that it seems most apt to use will likely depend on the context. For instance, it is hard to think of a 'favorite option' in the p-beauty contest.

The third possible secret to coordination is the Schelling salience hypothesis. The **Schelling salience hypothesis** is that people ignore what is primary or secondary salient and look for some key or clue to how to coordinate. For example, Grand

Table 6.3 Level-k thinking in the 0.5-beauty contest compared to a model of cognitive hierarchy. For example, a type C-2 believes 37.5 percent of others are type L-0 and 62.5 percent are type L-1, so chooses 17.19.

Level-k thinking		Cognitive hierarchy				
Type	**Choice**	**Type**	**Proportion (%)**	**Type C-2 beliefs (%)**	**Type C-3 beliefs (%)**	**Choice**
L-0	50	L-0	30	37.5	33.3	50
L-1	25	L-1	50	62.5	55.6	25
L-2	12.5	C-2	10	0	11.1	17.19
L-3	6.25	C-3	10	0	0	16.23

Table 6.4 The answers subjects gave, and the proportion giving each answer, to two questions. The answers depend on whether the subjects did or did not have an incentive to match with others.

Question	No incentive to match		Incentive to match	
	Response	**Proportion**	**Response**	**Proportion**
'Write down any year'	1971	8.0	1990	61.1
	1990	6.8	2000	11.1
	2000	6.8	1969	5.6
	1968	5.7		
'Write down any positive number'	7	11.4	1	40.0
	2	10.2	7	14.4
	10	5.7	10	13.3
	1	4.5	2	11.1

Source: Mehta, Starmer and Sugden (1994).

Central Station may not be the favorite place of anyone, but somehow stands out as the best choice. Similarly, no one may like face four but it stands out, because, say, it is the only male face. Schelling salience may not look much like level-k thinking, but it is not really that different. That's because we can again think of it as assuming that people will be of type L-1. The difference now is that someone of type L-0 is assumed to (4) choose the **focal point** or the option that stands out the most. We do, of course, need to think what the focal point might be, but let us look at some examples first.

In a study by Mehta, Starmer and Sugden (1994), subjects were asked questions such as 'Write down any year, past, present, or future' and 'Write down any positive number'. Some subjects were given no incentive to match the choices of others, while others were given the objective of trying to match the choices of others. Table 6.4 summarizes the results for two of the questions. The answers of those with no incentive to match should tell us what was primary salient. For example, when asked to 'Write down any year…', many chose their year of birth. There was a lot of variety in the answers, and we can infer that coordination due to primary salience would be low. There is much more consistency in the responses of subjects who were given an incentive to match.

These examples illustrate the ability of people to coordinate when needed. That the incentive to coordinate changed responses so much also means that we can rule out primary salience as an explanation for coordination. The data points more towards Schelling salience or secondary salience. Distinguishing between these two is much harder, however. In the 'Write down any positive number' example it looks as if Schelling salience is at work, because many chose number one when they had an incentive to coordinate and very few without this incentive. In the 'Write down any year…' example, it is harder to tell. The majority of students chose the then current year, 1990, when they had an incentive to coordinate, and so we might think of this as a focal point. But 1990 was also a relatively popular response for those with no incentive to coordinate, so secondary salience also looks plausible.

More generally, it seems that an ability to coordinate can come from either secondary salience or Schelling salience. Which fits better seems to change, depending on the context. One thing is clear, however, which is that the notion of focal point and Schelling salience is compelling. In some situations we are able to coordinate with others because there exists some focal option. If we are going to predict what people do in such situations then we need to be able to predict the focal point. This, though, is no easy task (see Research methods 6.2).

❖ RESEARCH METHODS 6.2 ALGEBRA HAS LIMITS?

In particular situations it can be 'obvious' to most people what the focal point is. The problem is that, to make good economic predictions, we need to abstract away from particular situations to get a more general model. This is hard to do. It is hard to say in general why, and when, 'things stand out' or 'seem obvious'. This presents a particular problem for economists, who have an insatiable desire to capture things in models with equations. The idea of a focal point may be something that equations simply cannot capture. Framing effects may be another. Does this matter?

It should not. Just because we cannot write down a nice equation does not mean we should not take focal points or framing seriously or be able to model them. We just need to be content with a more descriptive theory. Kahneman comments:

> [P]sychological theories of intuitive thinking cannot match the elegance and precision of formal normative models of belief and choice, but this is just another way of saying the rational models are psychologically unrealistic.
>
> (Kahneman 2003: 1449)

Sugden (1995) has proposed one means to try to capture focal points by introducing the idea of a **collectively rational recommendation** (CRR). The basic idea

behind a CRR is to imagine two people being recommended rules by which to make a choice; if one or more of them do not use a recommended rule, then they both get strictly less than if they used the rules. To illustrate the concept, consider these three different scenarios.

Scenario A: Emma and Alan are independently handed five discs that are identical except for the numbers one to five written on the underside of the discs. Both are asked to choose one of the discs with the objective of matching the choice of the other. They cannot see the numbers but an independent witness can, and can therefore verify whether they choose the same disc. The discs are presented to each of them in a random order.

In this scenario there is simply no way that Emma and Alan can improve their chances of coordinating. For example, if both use the recommendation 'Pick the one furthest to the left', the random way that they are presented means the chances of matching are still only 20 percent. In this game any recommendation is a CRR.

Scenario B: Now suppose that the discs have the numbers written visibly on the topside.

In this scenario it is much easier for them to coordinate. For instance, if they both use the recommendation 'Choose disc one' then they coordinate for sure. The problem is that they need to coordinate on a recommendation. This means that there is no CRR. For example, recommending that they 'Choose disc one' cannot be a CRR because they could both ignore the recommendation, use their own rule, 'Choose disc two', and do just as well.

Scenario C: Finally, suppose that the numbers are written on the non-visible underside but two discs are blue and three are red.

This scenario combines elements of both scenarios A and B. Specifically, there is the chance to coordinate (as in scenario B) using color, but this will not guarantee coordination (as in scenario A) because there is more than one disc of each color. The CRR in this scenario is for both to 'Choose a blue disc'. This does not guarantee coordination but does mean a better chance of success than both choosing red. In this scenario, blue is a focal choice.

Can we use the concept of a CRR to capture why so many people choose Grand Central Station and 'Heads'? Yes. That is because in many coordination games it turns out that the only CRR is to recommend 'Choose the option most frequently mentioned in everyday life'. This is basically because a lot more people are likely to agree on what is the most frequently mentioned option than to agree on what is, say, the tenth or 11th or 235th most frequently mentioned option. The chances of coordination are thus maximized if everyone chooses the most frequently mentioned option. Closely related is the **recognition heuristic** of 'If in doubt, choose the one you recognize, or are most familiar with'. For example, subjects in the study by Mehta, Starmer and Sugden were also asked to name a mountain, and 89 percent of subjects named Everest.

The notion of a focal point is a hard one to tie down, and the concept of a CRR is going to make sense only in some situations. It does, though, give us some insight. And predicting what people will do in pure coordination games is inevitably going to be very tough, given that people are completely indifferent to what they do as long as they coordinate. In games with a bit more structure, when people are not so indifferent to what they do, our task may be a little easier.

6.2.4 Equilibrium refinement

So far we have looked in some detail at the p-beauty contest, in which there is a unique Nash equilibrium, and the simple beauty contest, in which there are many Nash equilibria and little to choose among those equilibria. Most games that are of interest to us lie between these two extremes. They have multiple equilibria but those equilibria differ in some systematic way. What I have done here is to select three games in order to illustrate some concepts that are useful in analyzing how people behave and reason in such instances.

Let's start by looking at the game depicted in Table 6.5. If you have never seen a matrix game before, the idea is very simple. Alan and Emma both need to decide at the same time whether to choose high or low effort. The numbers indicate the possible payoffs. So, if they both choose high effort they get $13 each. In section 6.6 I am going to look in more detail at why these payoffs fit a story of high or low effort. For now, I want you to think what Alan should do.

There are two Nash equilibria in this game (plus another one I will ignore). There is one in which they both choose low effort and another in which they both choose high effort. You could argue that the equilibrium in which they put in high effort stands out as more salient because it gives a higher payoff. To capture this, we say that a **Nash equilibrium is Pareto-dominant** if it gives a higher payoff to everyone than any other Nash equilibrium. The high-effort Nash equilibrium is Pareto-dominant in this game. Will Emma and Alan coordinate on the Pareto-dominant Nash equilibrium?

If Alan chooses low effort he is guaranteed $7, while if he chooses high effort he could get $13, but might also get $0. If he is worried about getting $0 then he may prefer to play it safe and guarantee himself $7. To capture this, we say that a **Nash equilibrium is risk-dominant** if each person is choosing the strategy that would maximize their payoff if they cannot predict what others will do. To see how this works, suppose that Alan cannot predict what Emma will do and so says there

Table 6.5 A matrix game. For example, if Alan chooses low effort and Emma chooses high effort, Alan gets $7 and Emma gets $0. The Nash equilibrium of low effort is risk-dominant and high effort is Pareto-dominant.

		Emma	
		Low effort	**High effort**
Alan	Low effort	$7, $7	$7, $0
	High effort	$0, $7	$13, $13

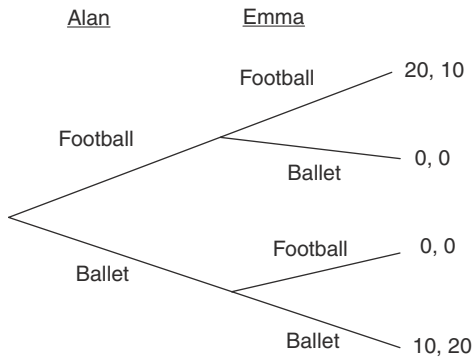

Figure 6.5 A sequential battle of the sexes game. Alan chooses to buy football or ballet tickets and then Emma decides whether to go to the football or ballet. For instance, if they both choose football, Alan gets payoff 20 and Emma payoff 10.

is a 50 percent chance she will choose high effort and a 50 percent chance she will choose low effort. Given this, his payoff from choosing low effort is \$7 and that from high effort is 0.5 × \$0 + 0.5 × \$13 = \$6.5; \$7 beats \$6.5, and so Alan should choose low effort. The low-effort Nash equilibrium is risk-dominant.

We see that there are good arguments why Alan should choose high effort, and good arguments why he should choose low effort. That's the problem of multiple equilibria! In this case there is a tension between the Pareto-dominant equilibrium, which gives a high payoff, and the risk-dominant equilibrium, which is less risky. It is only by observing what people do that we can hope to pull apart when people favor one or the other. We will come back to this in section 6.7 and look at some data.

I next want us to look at sequential games. In a **sequential game** one person chooses before another. To illustrate, we can look at versions of the battle of the sexes game. Figure 6.5 summarizes a standard version. In this case, Alan decides whether to buy tickets for the football or ballet. His friend, Emma, sees what tickets he bought and then decides whether to go to the football or ballet. The numbers indicate the respective payoffs, with Alan first. You can see that they both want to go together but Alan prefers football and Emma prefers ballet. What should Alan choose?

In this game there are three Nash equilibria. In describing the Nash equilibria we need to be a little careful, because – recall – a strategy should be a complete plan of what to do. So, Emma needs to say what she will do if Alan buys football tickets and what she will do if he buys ballet tickets. The simplest strategy is when Emma says: 'If he buys football tickets I go to the football and if he buys ballet tickets I go to the ballet.' If she does this and Alan chooses football, we have our first Nash equilibrium.

In this case, it does not really matter what Emma would have done if Alan had bought a ballet ticket. There is, therefore, a second Nash equilibrium, in which Alan buys a football ticket and Emma says: 'I will go to the football if he buys football or ballet tickets.' The final equilibrium is similar. This time Emma says: 'I am going to the ballet whether he buys football or ballet tickets.' You could think of this as a threat that she will go to the ballet. In this case, the best thing Alan can do is buy a ballet ticket,

Table 6.6 The Nash equilibria in the sequential battle of the sexes.

Equilibrium	Alan's strategy	Emma's strategy
1	Football	Football if Alan chooses football and ballet if he chooses ballet
2	Football	Football no matter what Alan chooses
3	Ballet	Ballet no matter what Alan chooses

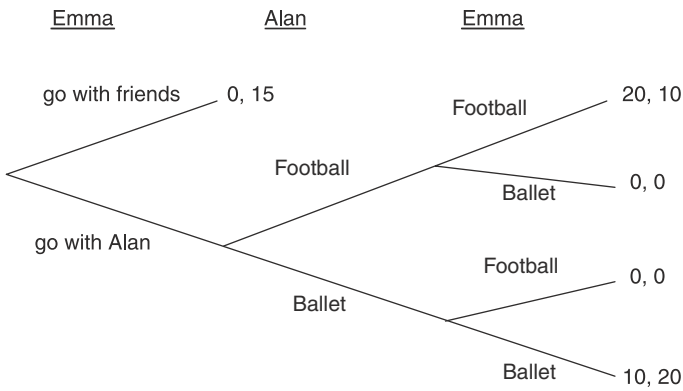

Figure 6.6 A sequential battle of the sexes game with an outside option. Emma decides first whether or not to go with friends or with Alan.

and this gives us our third Nash equilibrium. The three equilibria are summarized in Table 6.6. How can we choose amongst these three equilibria? One of them probably stands out as more sensible. That's the first one, in which Emma does what Alan does. The second equilibrium seems strange, given that Emma would go to the football if Alan buys ballet tickets! The third equilibrium is more intuitive, but there is a problem. The threat of Emma to go to the ballet if Alan buys football tickets is not a **credible threat**, because she would get a higher payoff going to the football.

One technique we can use to justify the first equilibrium as more sensible is **backward induction**. We start with Emma's choice and ask what it is optimal for her to do. If Alan buys football tickets it is optimal she chooses football, and if he buys ballet tickets it is optimal she chooses ballet. We then assume that Alan knows this. If he knows that Emma will choose football if he buys football tickets then it is optimal for him to buy football tickets. Another way to describe this is to say that the first equilibrium is **sub-game perfect**, because, if we look at any part of the game, people are behaving optimally. The first equilibrium is the unique sub-game perfect equilibrium of this game. If we look at the third equilibrium, for instance, Emma is not behaving optimally in the part of the game where we imagine Alan has chosen football.

A complement to backward induction is forward induction. To motivate this, consider a modified battle of the sexes game, such as that in Figure 6.6. What I have done is give Emma the option of choosing whether or not to go out with friends before she contacts Alan and asks him if they should go somewhere together. How does this change things? If you use backward induction you will find that there is a unique sub-game perfect Nash equilibrium, and that involves Emma going with friends. If she goes

with Alan they will end up going to the football and she gets a payoff of 10, which is less than if she goes with friends. This looks like a bad outcome.

What would happen if Emma were to choose to go with Alan? Alan might reason that she could only possibly have chosen to go with him if she expects they will go to the ballet. That is the only thing that could justify her choice. If they go to the ballet he gets payoff 10, which is better than he would have got if Emma went with friends. This might, therefore, look like a good deal. This is an example of **forward induction**, when we start with the first choice and ask what will happen if a certain choice is made. Forward induction gives a Nash equilibrium in which Emma chooses to go with Alan and threatens to choose ballet whether Alan buys football or ballet tickets. The fact that she chose to go with Alan arguably makes the threat more credible.

Hopefully, you are getting the idea that there are lots of ways we can model how people might reason when playing games. Risk dominance, Pareto dominance, backward induction, sub-game perfection and forward induction are just some of the concepts that game theorists have dreamt up to deal with the complication of multiple equilibria. In the latter half of this chapter and the next we will look at lots of examples to see whether these concepts are any good at predicting what people actually do. All the concepts we have looked at so far, though, are called equilibrium refinements (see Research methods 6.3); to finish I want to give one example of equilibrium selection.

❖ RESEARCH METHODS 6.3 EQUILIBRIUM REFINEMENTS AND SELECTION

If there are multiple Nash equilibria it seems natural to ask if one equilibrium might be more likely to be observed than others. This has led to various suggestions of how to refine and select equilibria.

A Nash equilibrium is **refined** by adding an extra condition that must be satisfied. For example, a sub-game perfect Nash equilibrium is required to be a Nash equilibrium of every possible sub-game of the original game. A Pareto-dominant Nash equilibrium is required to give higher payoffs to everyone than any other Nash equilibrium. Refining the set of Nash equilibria can be done using ideas from behavioral economics. The problem is that there may be multiple or no equilibria that satisfy the refinement. So, the original problem of saying which equilibrium will be observed need not be resolved.

Equilibrium **selection** tackles the problem of multiple equilibria head-on by providing an algorithm to select a unique equilibrium for a general class of games. We shall see that quantal response equilibrium can be applied to do this. Clearly, this approach provides a specific prediction of what people will do. The problem is that the behavioral justification for the algorithm may have to give way to the analytical need to pick one equilibrium. Put another way, it may be natural to think of multiple or no equilibria as possible outcomes of the game, and equilibrium selection does not allow that.

6.2.5 Nash equilibrium with mistakes

In section 6.1 we saw that Nash equilibrium was a poor predictor of what people would do when they play the p-beauty contest for the first time. In section 6.2.3 we saw that the Nash equilibrium was a slightly better predictor in pure coordination games, but still far from perfect. Trust me: we are going to see before too long that the Nash equilibrium is not always as bad. Even so, it seems as though we should be looking at deviations from the Nash equilibrium. One way to do this is to assume that people sometimes make mistakes. There are different ways this can be done, but I will focus on one possibility called quantal response equilibrium, for which we need to thank McKelvey and Palfrey (1995).

The basic idea behind **quantal response equilibrium** (QRE) is that people choose their optimal strategy with some error. The main thing we need to know is that, at a QRE, the probability person i will choose strategy l is given by

$$\sigma_i^*(l) = \frac{e^{\lambda u_i(l,\sigma_{-i}^*)}}{\sum_j e^{\lambda u_i(j,\sigma_{-i}^*)}} \tag{6.1}$$

for all i and l, where λ is a precision parameter that measures the size of errors. This equation looks a bit messy but it does turn out to be a nice way to capture mistakes. That's because: if $\lambda = 0$ then everyone chooses at random, and, the larger is λ, the fewer errors people make. As λ tends to infinity, QRE converges to a Nash equilibrium.

To illustrate how to find a QRE, consider the matrix game in Table 6.7. The only Nash equilibrium here is what's called a Nash equilibrium in mixed strategies, when Alan and Emma are supposed to randomize what they do. To find the QRE, let p_L denote the probability that Emma chooses left and q_U denote the probability that Alan chooses up. Plugging this into equation (6.1), we get

$$q_U = \frac{e^{9\lambda p_L}}{e^{9\lambda p_L} + e^{\lambda(1-p_L)}} \quad \text{and} \quad p_L = \frac{e^{\lambda(1-q_U)}}{e^{\lambda q_u} + e^{\lambda(1-q_U)}}.$$

We now need to find the values of q_U and p_L such that both equations hold. This will depend on λ. For example, if $\lambda = 1$ then we get $q_U = 0.89$ and $p_L = 0.31$ (check it and see). Figure 6.7 plots the values of q_U and p_L for different values of λ.

One way to interpret the QRE here is that, because of mistakes, Emma is likely to choose left more often that she would do with the Nash equilibrium. Consequently,

Table 6.7 A matrix game

		Emma	
		Left	**Right**
Alan	Up	\$9,\$0	\$0, \$1
	Down	\$0, \$1	\$1,\$0

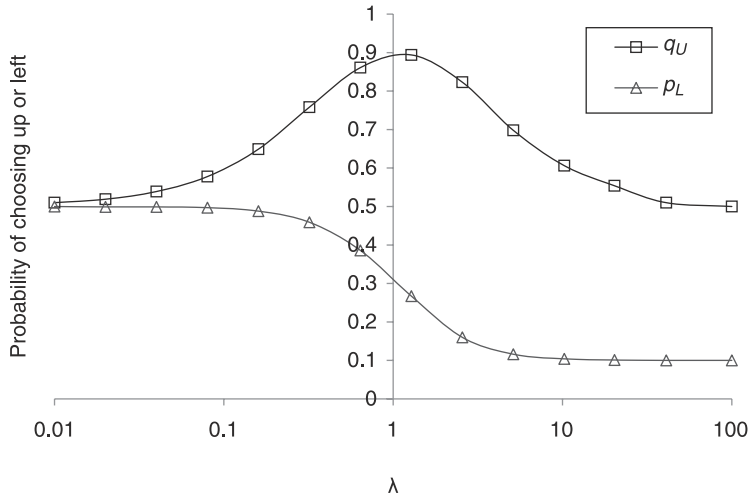

Figure 6.7 The QRE for the game in Table 6.7. The equilibrium depends on λ, with a high λ meaning fewer mistakes are made.

Alan should play up more often than with the Nash equilibrium. QRE thus gives us a prediction of where the Nash equilibrium might go wrong. The really clever thing, however, is that it also gives us a means to select equilibria. This is how it works: if = 0, there is a unique QRE. The trick is then to increase λ and keep track of how this QRE changes. Doing this gives a unique QRE for each λ and, ultimately (when $\lambda = \infty$), selects a unique Nash equilibrium. Analytically, this is very neat. Does it work?

McKelvey, Palfrey and Weber (2000) had subjects play the game in Table 6.7 and two others. In these other games the payoffs of Alan and Emma were four times as big. This increase in payoffs does not change the Nash equilibrium but it does change the QRE. Indeed, one of the main predictions of QRE is that people should make smaller errors, and be expected to make smaller errors, when payoffs are bigger.

McKelvey, Palfrey and Weber found that QRE did a lot better at predicting what subjects do than Nash equilibrium or an alternative of Nash equilibrium with mistakes. The news was not all good, however. To see why, Figure 6.8 plots the estimated distribution of λ for each of the games. The way to interpret this distribution is that some people make more mistakes than others, and the distribution shows the probability that a random person is best characterized by the different values of λ. The main concern is that the distribution of λ is different for each game. This makes it hard to apply QRE, because we cannot know what the distribution of λ will be for the particular game we are interested in.

Let's finish this section by focusing on the good news. We have now looked at lots of different ways to think about how people might play games. Most of them are complementary. For example, level-k thinking requires us to think what the type L-0s will do, and the equilibrium refinements and selections we have looked at give us some clue as to what they might do. What we need to do now is get a feel for which

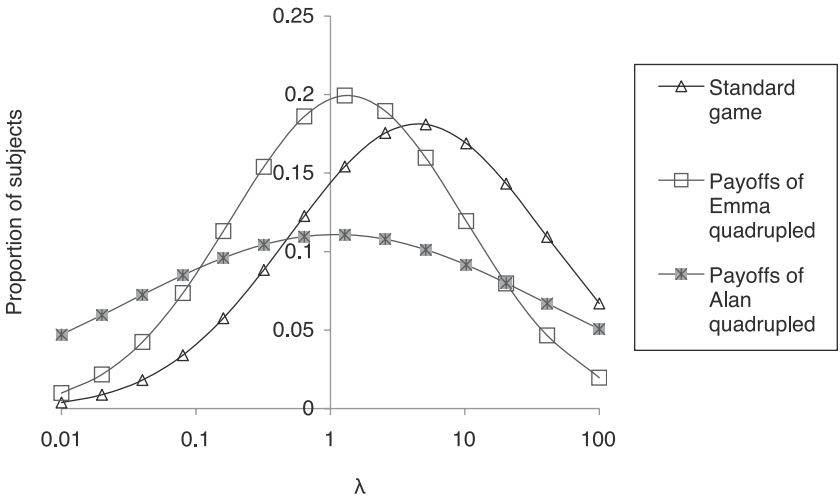

Figure 6.8 The estimated distribution of λ across subjects for each of the three games.
Source: McKelvey, Palfrey and Weber (2000).

way of thinking is best for particular games. Before getting started on that, however, I want to look at how people can learn from experience.

[Extra] If you are desperate to know where equation (6.1) comes from then here is one motivation. Suppose Emma's perception of the payoff she will get is slightly wrong. We denote by $u_E^P(l, s_{-E})$ the payoff she thinks she will get from choosing l if others choose s_{-E}. This will be her actual payoff plus an error – i.e.

$$u_E^P(l, s_{-E}) = u_E(l, s_{-E}) + \varepsilon_l,$$

where εl is the size of the error. Suppose that the size of the error ε_l is determined randomly according to probability distribution f. If Emma believes others will use strategy σ_{-E} then she will maximize her expected payoff by choosing l if $u_E^P(l, \sigma_{-E}) > u_E^P(j, \sigma_{-E})$ or $u_E(l, \sigma_{-E}) + \varepsilon_l \geq u_E(j, \sigma_{-E}) + \varepsilon_j$ for all j. The probability that she will choose l is, therefore,

$$\sigma_E^*(l) = \int_{R_l(\sigma_{-E})} f(\varepsilon) d\varepsilon, \tag{6.2}$$

where $R_l(\sigma_{-E})$ is the set of errors ε_l for which she would choose strategy l. A QRE is obtained when everyone has correct beliefs of what others will do. That is, strategy vector $\sigma^* = (\sigma_1^*, \ldots, \sigma^{*N})$ is a QRE if each σ^* satisfies (6.2) with $\sigma - i = \sigma^*_{-i}$. If f is something called the extreme value distribution then it is possible to derive equation (6.1). The main thing to take from this exercise is that we can justify QRE as a likely outcome if people are uncertain of their payoff function. In Chapter 2 I suggested people may be uncertain of their payoff function, so QRE has an intuitive appeal.

6.3 Learning from experience

We shall look next at what happens if Emma has the chance to play a game a few times and so learn from experience. In this context it is natural to think of the *probability* that she will play a strategy; she may, for instance, want to try a new strategy to see what happens, or stick with a strategy that has done well in the past. What we need to determine, therefore, is the probability $P^I(t)$ that Emma will play strategy I the t'th time she plays a game. There are two key elements to doing this.

First are initial beliefs. In the previous section we were effectively questioning what $P^I(1)$ is – that is, the probability Emma will play each strategy the first time she plays the game. What she plays will depend on her **initial beliefs**. For example, Emma may have no idea what her payoff function is, or what others will do, and so chooses a strategy randomly. In other situations, she may know her payoff function and, having read section 6.2, be sure what strategy she wants to choose. Initial beliefs are important, but primarily the subject of the last section, and so I will mention them only briefly. Instead, I want to focus on the second element of learning.

Once she has played the game Emma will almost certainly want to **update her beliefs** and change the probability that she chooses each strategy. In modeling how this works it is useful to measure the attractiveness of a strategy. I shall denote by $Q^I(t)$ the **attractiveness of strategy** I for Emma the t'th time she plays the game. What a learning model needs to do is specify the **learning rule** or **updating rule** whereby we can determine $Q^I(t)$ from $Q^I(t-1)$. That is, we need to know the rule determining how attractive each strategy will look for Emma.

Once we know the learning rule and the attractiveness of each strategy we can then work out the probability that Emma will choose each strategy. As you might expect, we will assume that more attractive strategies will be more likely to be chosen. The simplest way to do this is to set

$$p^I(t) = \frac{Q^I(t)}{\sum_j Q^j(t)}.$$

In other words, we assume that Emma chooses a strategy with a probability **proportional to its attractiveness**. This way is simplest but not necessarily best. Other ways include use of an exponential, power or probit formula. For example, using an **exponential** formula we get

$$p^I(t) = \frac{e^{\omega Q^I(t)}}{\sum_j e^{\omega Q^j(t)}}. \tag{6.3}$$

The parameter ω allows us to vary how sensitive Emma's choice is to the attractiveness of a strategy. The higher is ω, the more likely is Emma to choose the strategy with the highest attractiveness.

The focus of this section will mainly be on the learning rule, and there are three main models of learning that we need to look at.

6.3.1 Reinforcement learning

Abundant psychological evidence has shown that choices that led to good outcomes in the past are likely to be repeated. This is called the **law of actual effect**. For example, if Emma is playing the 0.5-beauty contest, delays 20 days and does well then she is likely to delay around 20 days the next time. Models of **reinforcement learning** or **adaptive learning** are designed to capture this. To illustrate how such models work, I shall work through a model based on that of Roth and Erev (1995).

Imagine that the t'th time Emma played a game she chose strategy I and got payoff $u_E(I,s_{-E}(t))$. Also imagine that she has an **aspiration level** of a (which you can think of in the same way as the aspiration level I introduced in Chapter 2). We then assume that she updates the attractiveness of each strategy using learning rule

$$Q^j(t+1) \propto \begin{cases} Q^I(t) + [u_E(I, s_{-E}(t)) - a] & \text{if } j = I \\ Q^j(t) & \text{if } j \neq I \end{cases}. \tag{6.4}$$

There are two key things about this rule. First, a strategy can become more attractive only if Emma plays it and it gives a higher payoff than her aspiration level. Second, the higher the payoff above the aspiration level, the more attractive it becomes next time.

This means that Emma will tend over time to choose those strategies that she has played in the past and gave her a relatively high payoff. She is less likely to choose those strategies that gave her a relatively low payoff. That looks like a reasonable plan. To see how it works, Table 6.8 illustrates what might happen as Emma plays the 0.5-beauty contest with an aspiration level of 90. Recall that her payoff could be between zero and 100, and an aspiration of 90 means she wants to be within 10 of the optimum. To keep things simple I assume that Emma considers waiting only 20 days or 30 days and the probability of choosing either is proportional to the attractiveness. After playing the game four times she has become more likely to wait 30 days (see also Review question 6.5).

This basic model can be adapted in various ways. One interesting extension is **local experimentation**. In this case, for every strategy I we associate a set of

Table 6.8 Emma using reinforcement learning in the 0.5-beauty contest. Her choice and the optimum are hypothetical, but, given this hypothetical experience, we can see what happens to the attractiveness and probability of her choosing each strategy.

Round	Attractiveness Q^I		Probability of choosing P^I		Outcomes	
	20 days	30 days	20 days	30 days	Chooses	Optimum
1	30	30	0.5	0.5	20	40
2	20	30	0.4	0.6	20	25
3	25	30	0.45	0.55	30	30
4	25	40	0.38	0.62	30	35

strategies $L(I)$ that we consider near to or similar to I. Assuming an aspiration level of zero, the updating rule is then modified to

$$Q^j(t+1) \propto \begin{cases} Q^j(t) + u_E(I, s_{-E}(t)) & \text{if } j = I \\ Q^j(t) + \varepsilon u_E(I, s_{-E}(t)) & \text{if } j \in L(I) \\ Q^j(t) & \text{otherwise,} \end{cases}$$

where ε measures the extent of local experimentation. The additional feature that this adds is that a strategy becomes more attractive if a nearby strategy earned a high payoff. Thus, Emma would be willing to, say, experiment with a delay of 29 days if 30 days has worked in the past.

The main advantage of reinforcement learning is that Emma need only know her own payoff. She just tries things and sees what happens. She does not need to know what others did, what their payoffs were, what payoffs she would have got for doing something else, etc. Reinforcement learning can, therefore, be very generally applied and has a good chance to capture learning when someone is relatively uninformed about a game. If Emma does know something about the game, however, reinforcement learning gives no scope for her to exploit what she knows. Belief-based learning, which I shall look at next, does.

6.3.2 Belief-based learning

As well as reflecting on past payoffs it may be appropriate to reflect on the past behavior of others and try to predict what they will do in the future. For example, if it would have been optimal for Emma to delay 25 days last time, then maybe she should delay 25 days next time. Models of **belief-based learning** capture this. To illustrate, I will work through a model based on that of Cheung and Friedman (1997).

The main thing we need to do is determine Emma's beliefs about what others will do – for example, how long she thinks others will delay. Let $B^c(t)$ denote the probability with which she thinks others will play strategy combination s^c_{-E} the t'th time the game is played. Beliefs will likely depend on what happened in the past, and so we assume that Emma keeps track of what others have done. Let $\#^c(t)$ denote the number of times that strategy combination s^c_{-E} has been chosen by others in the first t plays of the game. In a p-beauty contest this would require Emma to keep count of the number of times that the optimal delay was one day, two days, etc.

If $s_{-E}(t)$ denotes what others did in the t'th game, we assume that beliefs are updated using rule

$$B^c(t+1) \propto \begin{cases} \rho\#^c(t) + 1 & \text{if } s_{-E}(t) = s^c_{-E} \\ \rho\#^c(t) & \text{if } s_{-E}(t) \neq s^c_{-E}, \end{cases}$$

where ρ is a parameter. To understand this rule it is useful to think about two special cases. If $\rho = 0$, Emma would expect others to do in the future what they did last time. This is called **best reply** or **Cournot**, and means beliefs are based solely on the last time the game was played. If $p = 1$, Emma believes the probability that a strategy combination will be chosen by others in the future is equal to the relative frequency with

Table 6.9 Emma using belief-based learning in the 0.5-beauty contest with $\rho = 0.5$. Her beliefs keep track of what has happened in the past. The attractiveness is worked out using equation 6.5. For simplicity I give the attractiveness of only three possible strategies.

Round	Beliefs B^c			Attractiveness Q^l			Others' choice
	10	20	30	5	10	15	
1	0	0	0	–	–	–	30
2	0	0	1	90	95	100	20
3	0	1	0.5	140	147.5	145	10
4	1	0.5	0.25	170	168.75	162.5	20

which it has been chosen in the past. This is called **fictitious play**, and means beliefs are based on the average of what has happened in all plays of the game. The parameter ρ measures, therefore, how far into the past Emma wants to look. The lower ρ is, the more quickly the past is ignored, or the more quickly she forgets. Table 6.9 illustrates what Emma may do in the 0.5-beauty contest when $\rho = 0.5$.

From the beliefs $B^c(t)$ we need to get to attractiveness. We can do this by equating attractiveness with expected payoff. Given beliefs $B^c(t)$, Emma's expected payoff from choosing strategy l is proportional to

$$Q^l(t) = \sum_c u_E(l, s^c_{-E})B^c(t). \tag{6.5}$$

In this case, Emma is more likely to choose a strategy that would have been a relatively good strategy to have chosen in the past. If $\rho = 1$, she will choose the strategy that would have been best the last time the game was played. If $\rho = 0$, she will choose the strategy that would have been best averaging over all the times she has played the game. Again, these seem like sensible things to do. [Extra] Equation (6.5) is good enough for our purposes but if we substitute in for $B^c_i(t)$ we can get the updating rule

$$Q^l(t+1) \propto \rho Q^l(t) \sum_c \#^c(t) + u_E(l, s^c_{-E}(t)). \tag{6.6}$$

Crucial in the model is the value of ρ. In the study by Cheung and Friedman, they report experiments that allow us to estimate ρ for different subjects and different games (see Research methods 6.4 for more on this). The estimates are given in Figure 6.9 (and for completeness I also give estimates from Camerer and Ho (1999) that we shall discuss below). Overall, we see relatively low values of ρ, suggesting that the average player looks mainly to the recent past. There was a lot of heterogeneity, however. Of those subjects they could classify, Cheung and Friedman estimated 56 percent used Cournot (with $\rho = 0$), 33 percent used fictitious play (with $\rho = 1$) and 11 percent were adaptive (with $\rho \in (0,1)$). Most subjects, therefore, took a short-term view, but many took a long-term view. Crucially, though, the distribution of ρs was

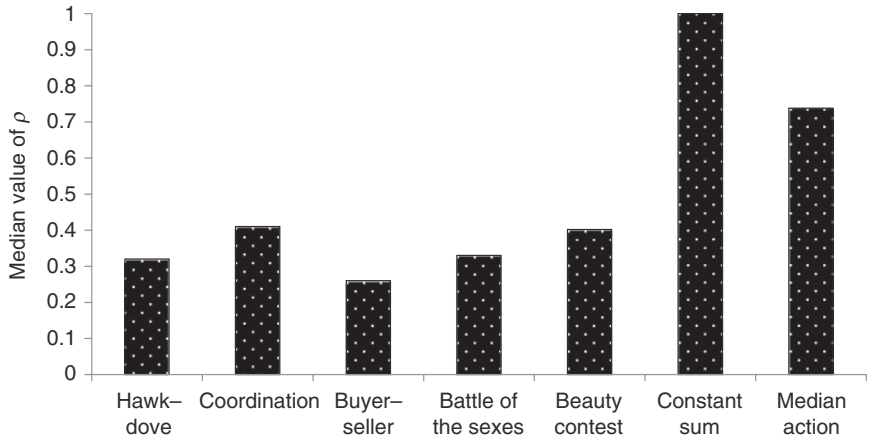

Figure 6.9 A model of belief-based learning. Estimates of the median ρ for different matrix games.
Sources: Cheung and Friedman (1997) and Camerer and Ho (1999).

similar across the games considered, which is useful for applying the model to predict what will happen in other games.

❖ RESEARCH METHODS 6.4 ESTIMATING LEARNING MODELS

A learning model is typically fitted to experimental data using **maximum likelihood estimation**. Basically, for any given set of parameters it is possible to calculate the likelihood of the experimental data given the particular set of parameters and model. For example, in the EWA model, we can fix values for δ, β, ϕ, $P^1(0)$, ..., $P^K(0)$ and #(0) and then say how likely a particular set of experimental data was, given the EWA model. We then search for the parameters that maximize the (log) likelihood to see how well a model can fit the data.

A real concern when doing this is that some models have more parameters than others and so a priori are likely to have a better chance of fitting the data. Put another way, a model with too many parameters may over-fit the data. There are various sophisticated ways to check for such things, but a more basic solution is to do an **out of sample forecast**. To do this, maximum likelihood estimates are derived using, say, 70 percent of the data and then the estimated model is used to predict what should happen in the other 30 percent of the data. If the model does a great job of fitting the original 70 percent but a poor job of predicting the other 30 percent, there is evidence of over-fitting.

A related issue is how to capture heterogeneity. To assume that all subjects will have the same parameters seems a little extreme, because we might imagine that some people learn more by reinforcement learning and others by belief-based learning, etc. If we allow different subjects to have different parameter values, however, we are likely to over-fit the data. A **latent class approach** allows there to be one, two or more types of people, with each person of a particular type assumed to have the same parameter values. The amount of data should be used to determine how many types seem reasonable. Camerer and Ho consider two latent types; Friedman and Cheung let each individual have different parameters.

Friedman and Cheung also provide interesting insight on the choices of subjects and whether subjects chose the strategies with highest attractiveness. They use equation (6.3) to estimate choices from attractiveness. Recall that a high *co* means that someone is more likely to choose attractive strategies. Figure 6.10 gives estimates of *co* for two of the games but four different treatments. In a mean matching treatment, each subject is matched to play every other subject in each round, while, in a random treatment, each subject plays one randomly selected subject per round. In a history treatment, the distribution of past choices in previous rounds is displayed, and, in a no-history treatment, it is not. In the treatments with mean matching and history, a subject has relatively more information on what others have done in the past. This should help form beliefs. Consistent with this we see, from the higher ω, that choice was more likely to be informed by beliefs with mean matching and history.

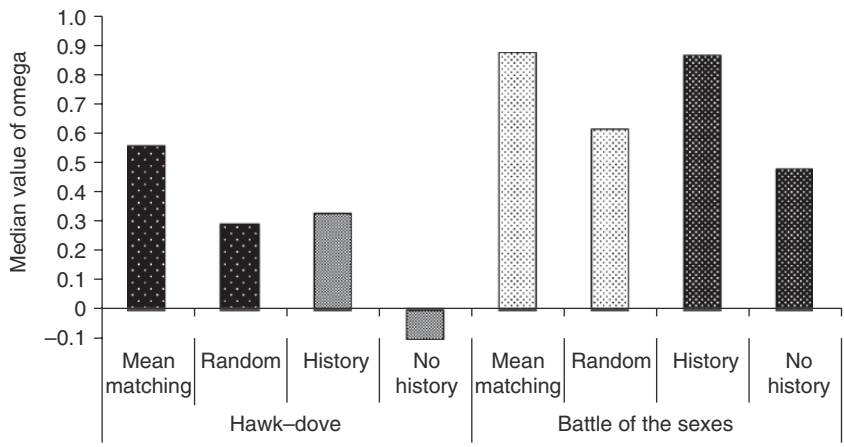

Figure 6.10 Median estimates of ω in a belief-based model of learning. The higher is *w* in situations with mean matching and history suggests subjects were more influenced by what others have done in the past in these situations.
Source: Cheung and Friedman (1997).

6.3.3 Experience-weighted learning

Reinforcement and belief-based learning nicely complement each other. Reinforcement learning means Emma looks at whether a strategy did give a high payoff; with belief-based learning she looks at whether it would have given a high payoff. Reinforcement learning means Emma updates beliefs about the payoff of each strategy; with belief-based learning she updates beliefs about what others will do. Is there any way to combine the two and get the best of both worlds?

Camerer and Ho (1999) show that this is possible with **experience-weighted attraction learning** (EWA). Similar to a belief-based model, we need a variable $\#(t)$ that keeps count of past experience and is updated using rule

$$\#(t) = \beta\#(t-1) + 1,$$

where β is some number called the **depreciation rate**. To see how $Q^l(t)$ is updated, suppose that in period t Emma chooses strategy l and others choose $s_{-E}(t)$. She gets payoff $u_E(l, s_{-E}(t))$. If she had chosen some other strategy j she could have got payoff $u_E(j, s_{-E}(t))$. Emma is assumed to give weight one to received payoff and δ to missed payoffs. The updating rule is then

$$Q^j(t+1) \propto \begin{cases} \varphi\#(t)Q^l(t) + u_E(l, s_{-E}(t)) & \text{if } j = l \\ \varphi\#(t)Q^j(t) + \delta u_E(j, s_{-E}(t)) & \text{if } j \neq l' \end{cases}$$

where φ is some number called the **decay rate**.

This looks a bit complicated, but you do not need to worry too much about that because the beauty of the EWA model is its ability to capture reinforcement learning and belief-based learning as special cases. It does so primarily through the parameter δ, which measures the relative weight given to missed payoffs rather than received payoffs. If $\delta = 0$, $\#(0) = 1$, $\phi = 1$ and $\beta = 0$ then we obtain a model of reinforcement learning like that in equation (6.4). That's because, when $\delta = 0$, Emma focuses exclusively on the payoff she did receive and ignores the payoffs she missed out on. If $\delta = 1$ and we set $\beta = \phi = \rho$, we obtain a model of belief-based learning similar to that of equation (6.6). That's because, when $\delta = 1$ Emma, focuses equally on the payoff she missed out on compared to the payoff she did receive. If δ is between 0 and 1, we get a mix of reinforcement learning and belief-based learning.

Camerer and Ho fitted the EWA model to a variety of data and compared fit to that with reinforcement and belief-based learning. Table 6.10 summarizes the estimates they find for the various parameters and Figure 6.11 plots an estimate of goodness of fit. As we would expect, EWA does relatively well. More interesting is how reinforcement learning does relatively well in the beauty contest and median action game, while belief-based learning does better in constant sum games. We can give an ex post rationalization for this but the main point is that people might learn in different ways in different games, and EWA is well placed to pick this up. Indeed, with the three models of learning we now have, we are well placed to better understand how people learn from experience. But is this enough?

Table 6.10 EWA parameter estimates.

Game	Learning	δ	β	ϕ	ω
Beauty contest	Reinforcement	0	0	1.38	0.22
	Belief-based	1	0.40	0.40	0.94
	EWA	0.23	0.94	1.33	2.58
Constant sum (4 different games)	Reinforcement	0	0	0.96–1.01	0.03–0.10
	Belief-based	1	0.99–1.00	0.99–1.00	0.46–1.81
	EWA	0.00–0.73	0.93–0.96	0.99–1.04	0.18–0.65
Median action	Reinforcement	0	0	0.93	1.19
	Belief-based	1	0.74	0.74	16.73
	EWA	0.85	0.00	0.80	6.82

Source: Camerer and Ho (1999).

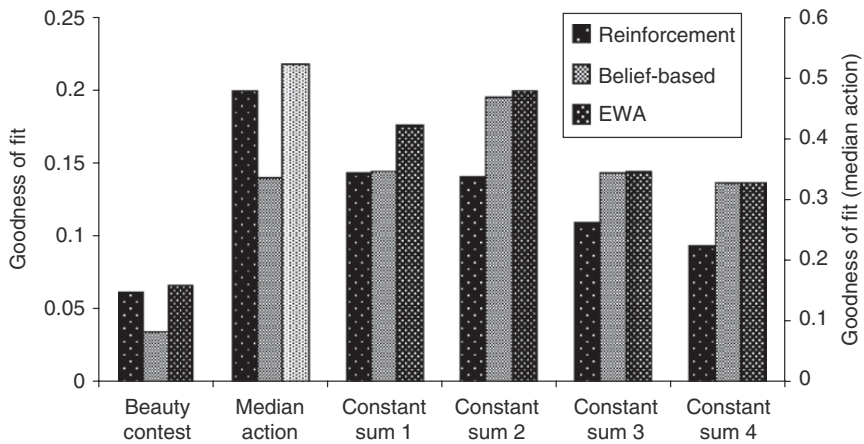

Figure 6.11 The adjusted R^2 for the three different learning models in six games.
Source: Camerer and Ho (1999).

6.3.4 Learning and prediction

A model of learning says what Emma will do, conditional on what has happened previously. This is great for looking back and trying to make sense of what she did. But is it so useful in trying to look forward and predict what she will do? One thing we could do is combine the ideas of sections 6.2 and 6.3, get a prediction of what will happen the first time a game is played and then use a learning model to see what will happen after, say, ten plays of the game. But how confident can we be in the answer? Looking back at Tables 6.8 and 6.9, which trace through what happens as Emma learns in the 0.5-beauty contest, we can see that the prediction we will get depends a lot on chance events, such as what Emma randomly chooses to do.

Another approach is to ask whether play will **converge**, and, if so, to what. To do this, we ask what happens after people have played the game, say, 1,000 times. If we always get the same answer then we say that play converges, and then use this as a prediction of what will happen. This puts the focus firmly on what will happen in the

long run. But maybe it also gives a good prediction of what will happen in the short run after, say, ten plays of the game. Unfortunately, this approach does not work so well for the three learning models we have seen so far. That's because play need not converge, and, if it does converge, it converges to a Nash equilibrium. So we are basically just using a kind of equilibrium refinement. There are alternatives, though, and in this section I will consider three.

I will start with something called **payoff sampling**. Suppose that, for each possible strategy, Emma samples the payoff she got from choosing that strategy on, say, six previous occasions. Remember: we are assuming Emma has played the game lots of times and so she can do this. She then chooses the strategy that gave her the highest payoff. This is a type of reinforcement learning and lets us predict what Emma will do. Indeed, we can find a **payoff sampling equilibrium** that tells us the probability that she will choose each strategy in the long run if everyone else is also using payoff sampling.

To illustrate how it works, consider the matrix game in Table 6.11. Suppose that, when Alan samples the payoff from choosing up, there were U_L times that he got 10 and $6 - U_L$ times he got 0. When he samples the payoff from choosing down he must have got 9 all six times. So, he will choose up if $10U_L > 9 \times 6 = 54$. This could happen only if $U_L = 6$. If ρ_L is the probability that Emma chooses left, this happens with probability $p_L{}^6$. So, letting q_U denote the probability that Alan chooses up we get

$$q_U = prob(U_L = 6) = p_L{}^6.$$

We can now do the same to derive p_L as a function of q_U. When Emma samples her payoff from choosing left, she gets 9. Suppose that when she samples her payoff from choosing right there were R_U times that she got 18 and $6 - R_U$ times she got 8. She will choose left if $9 \times 6 > 18\ R_U + 8(6 - R_U)$ or $6 > 10\ R_U$. This can be satisfied only if $R_U = 0$. Thus

$$p_L = prob(R_U = 0) = (1 - q_U)^6.$$

We now know the probability that Alan chooses up as a function of the probability that Emma chooses left, and vice versa. To obtain a payoff sampling equilibrium we look for the point where these two probabilities are consistent – that is, where $q_U = ((1 - q_U)^6)^6$ and $p_L = (1 - p_L{}^6)^6$. Values $q_U = 0.071$ and $p_L = 0.643$ work. This is a prediction of what Alan and Emma might do.

Before we see how good this prediction is, I want to go through the other two possibilities. Next up is **action sampling**. Imagine now that Emma samples what Alan

Table 6.11 A matrix game to illustrate payoff sampling.

		Emma	
		Left	**Right**
Alan	Up	\$10, \$9	\$0, \$18
	Down	\$9, \$9	\$9, \$8

chose on, say, seven previous occasions. She then chooses the strategy that would have maximized her payoff given what Alan did. This gives a form of belief-based learning. By finding an **action sampling equilibrium**, using similar techniques to those that we used to find a payoff sampling equilibrium, we get a prediction of what Emma will do. If we work through the game in Table 6.11 then we get the values $q_U = 0.057$ and $p_L = 0.664$.

[Extra] To see where these numbers come from. Alan should choose up if $10p_L > 9$ or $p_L > 0.9$. So, he should choose up if and only if he observed Emma choosing left all seven times he samples. Thus

q_U = *prob* (observes left seven times) = p_L^7.

Emma should choose left if $9 > 18\,q_U + 8(1 - q_U)$ or $q_U < 0.1$. So, she should choose left if Alan never chose up on the seven times she samples. Thus

p_L = *prob* (does not observe up) = $(1 - q_U)^7$.

To obtain an action sampling equilibrium we look for the point where these two probabilities are consistent. That is, $q_U = ((1 - q_U)^7)^7$ and $p_L = (1 - p_L^7)^7$. The values $q_U = 0.057$ and $p_L = 0.664$ work this time.

The final possibility I want to look at is called **impulse balance**. What we do here is calculate a **security level** of payoff that Emma can guarantee herself. For instance, in the game in Table 6.11, Emma can guarantee a payoff of \$9 by choosing left. Taking into account loss aversion, we then say that gains above the security level are treated as less than losses below the security level. For example, we might say that a loss counts twice as much as a gain. Table 6.12 shows what happens when we do this.

Table 6.12 Transforming the payoffs to take into account loss aversion. Losses relative to the security payoff of \$9 count twice as much.

		Emma	
		Left	**Right**
Alan	Up	10, 9	–9, 18
	Down	9, 9	9, 7

Table 6.13 Impulses after taking into account loss aversion. For instance, if Alan chooses up and Emma chooses left, he has no impulse to change strategy, but if he had chosen down he would have had an impulse of 1 to change strategy.

		Emma	
		Left	**Right**
Alan	Up	0, 9	18, 0
	Down	1, 0	0, 2

For instance, if Alan chooses up and Emma chooses right, the payoff of \$0 is \$9 short of the security level and so would feel to Alan like $9 - 2 \times 9 = -9$.

The final thing we do is to say that Emma would have an **impulse** in the direction of a particular strategy if choosing that strategy would have earned a higher payoff. The size of the impulse is given by the payoff difference. Table 6.13 shows the impulses in the example. An impulse of 0 means she has no incentive to change strategy, and, the higher the number, the stronger the impulse to change. There is an **impulse balance equilibrium** if these impulses to change strategy cancel out.

To see how this works, we can work through the example. Given that Emma chooses left with probability p_L and he chooses up with probability q_U, Alan's expected impulse from up to down is $18q_U(1 - p_L)$ and his expected impulse from down to up is $(1 - q_U)p_L$. At an impulse balance equilibrium we want these to be the same, so $18q_U(1 - p_L) = (1 - q_U)p_L$, or

$$q_U = \frac{p_L}{18 - 17p_L}.$$

Similarly, Emma's expected impulse from left to right is $9q_Up_L$ and from right to left is $2(1 - q_U)(1 - p_L)$. Setting these equal gives

$$p_L = \frac{1 - q_U}{3.5q_U + 1}.$$

Solving these two equations gives the equilibrium – in this case, $q_U = 0.1$ and $p_L = 0.67$. Impulse balance combines elements of both reinforcement learning and belief-based learning, because it favors strategies that give a relatively high payoff. The most appealing thing about it, however, is that it can capture the loss aversion we saw in Chapters 2 and 3. It says that people will be drawn away from strategies that are likely to cause a loss.

The nice thing about payoff sampling equilibrium, action sampling equilibrium and impulse balance equilibrium is that there is only one of them. They give us, therefore, a nice prediction of what should happen in the long run as people learn. Furthermore, they are relatively easy to calculate, because we do not need to model the learning process itself. But are the predictions any use?

In a study by Selten and Chmura (2008), subjects played 12 different games similar to that in Table 6.11. Figure 6.12 gives a measure of how good different equilibria were at fitting the data. We see that QRE, action sampling, payoff sampling and impulse balance did a much better job than the Nash equilibrium. This is a clear signal that we need to take seriously deviations from the Nash equilibrium. Whether the three equilibria we have looked at in this section are an improvement over QRE is a more subtle and controversial point. But it is safe to say that they do no worse.

An important question to ask in interpreting these results is how many times subjects played the games. There were 12 games and a subject played for 200 rounds, so they played each game around 16 or 17 times. The results are the same, though, if we focus on just the first 100 rounds. This is good news, because it means these equilibrium predictions look as though they work even if people have had little or no

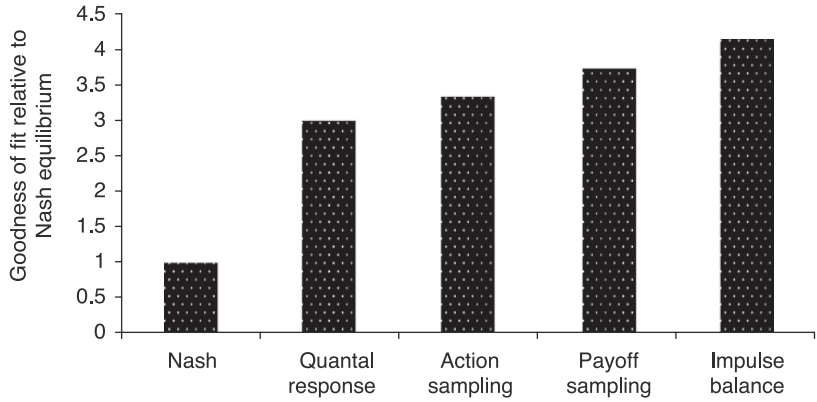

Figure 6.12 Assessing the goodness of fit of various equilibria relative to the Nash equilibrium.
Source: Brunner, Camerer and Jacob (2011).

chance to learn from experience. That does present something of a conundrum as to why they work so well, but maybe they just capture nicely how people reason when playing games. That seems a positive note upon which to finish the section.

6.4 Teams make decisions

The focus throughout the chapter so far has been on how individuals choose a strategy, or learn from experience. This reflects the norm in economics of focusing on individuals. The reality, however, is that many economic decisions are made by teams. A board of directors decides on the strategy of the firm. A committee of politicians and experts designs policy. A husband and wife decide whether to sell their house and move on. In all these examples the collective decision is made after team members have discussed things through with each other and reached a joint decision. The decisions that teams make could differ substantially from the decisions that individuals make. If so, an almost exclusive focus on individual behavior is misplaced!

In this section I want to begin to look in more detail at whether the decisions made by teams differ from those of individuals. I also want to illustrate how the theory we have learnt so far in the chapter can be applied to team decision-making. More examples will then follow over the next few chapters. Before we get started, let me clarify definitions, and in particular, clarify the distinction between a team and group.

When I introduced the notion of a game in Chapter 1 I used the example of a group of five people trying to finance a threshold public good. Why is this not an example of team decision-making – particularly if group members can discuss amongst each other, bargain, etc.? The answer is that each group member ultimately makes a decision independently of the others. If there are five members of the group, we observe five actions in the game, or five contributions to the public good. This is different from saying that a team of people reach a collective decision. If there are five members of the team, we observe only one action. For example, a household of husband, wife and three children may make one contribution to the public good.

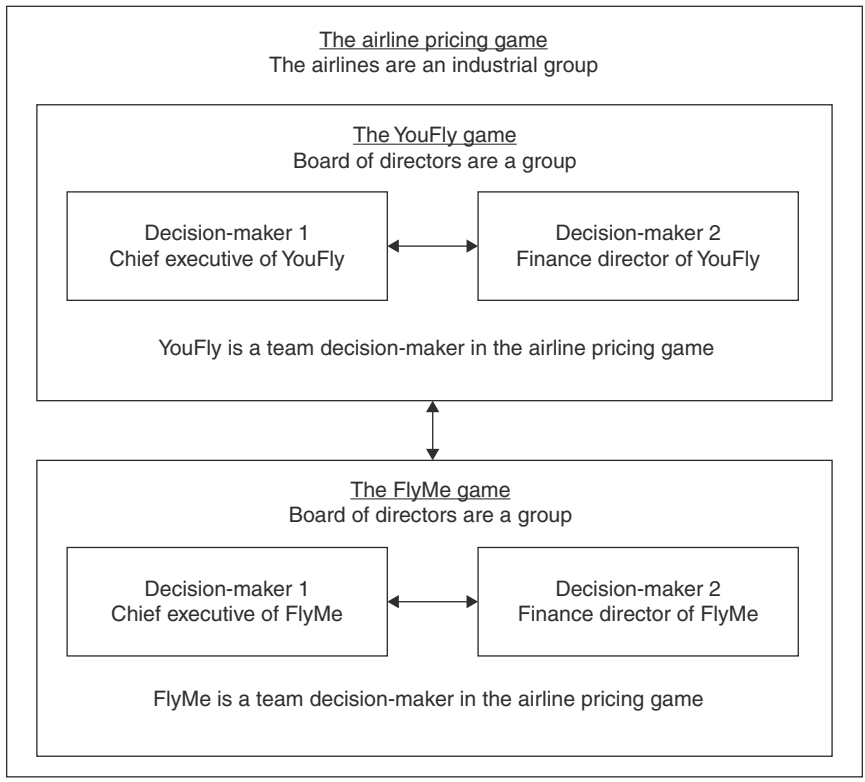

Figure 6.13 Within FlyMe and YouFly the board of directors play a game. The outcome of each game determines the strategy of each firm, or team, in the airline pricing game.

To avoid confusion I shall talk of decision-makers. For each **decision-maker** we will see an action. A decision-maker may be an **individual** or a **team** of people. Decision-makers may belong to a **group**. For example, we might have five decision-makers in a group trying to finance a public good with each decision-maker a household of two or more people. To give another example, we may have five decision-makers in an industrial group trying to make profit, with each decision-maker a company board of directors.

This distinction between groups and teams is convenient. But is it convincing? The reason I ask this question is that, when a team of people try to make a collective decision, they are playing a game. The husband and wife arguing over how much to contribute to the public good are playing a game. Similarly, company directors trying to pull the company one way or another are playing a game. The outcome of this game determines the action of the team. The picture is, therefore, somewhat like that in Figure 6.13. The figure shows a schematic of competition between two airlines, FlyMe and YouFly. Within FlyMe there is a game between directors over the direction the firm should take. Something similar is happening within YouFly. The outcomes of these two games determine pricing decisions that feed into the airline pricing game.

The key point to take from this is that team decision-making is a game in itself, with all the strategic interaction that this entails. It may be the whole game, such as a household deciding where to go on holiday; or it may be a small part of a bigger game, such as the directors of YouFly deciding on pricing strategy. Recognizing the strategic interaction between team members does not mean, however, that the distinction between teams and groups is not important. It is often very useful to treat the strategic interaction between team members as a black box that we do not need to worry ourselves about. For example, when analyzing the competition between FlyMe and YouFly we do not want to get bogged down looking at the dynamics within the boardroom. But for this to work we need general insight into how a team works and how we can appropriately model it.

6.4.1 Teams and the beauty contest

The p-beauty contest is an ideal place to begin comparing teams against individuals in a simple strategic environment. It also provides an interesting window onto financial markets where individual investors are routinely pitted against teams of investors. At the same time, for instance, as Emma is deciding when to sell her shares in FlyMe, you can be sure a fund managed by a team of investors is also deciding when to sell their shares. Does a team of investors outperform an individual?

To give some insight on how teams perform in the 0.67-beauty contest I will look at studies by Kocher and Sutter (2005) and Sutter (2005). In these studies a team consisted of two to four individuals, who have five minutes to discuss with each other face to face before making a joint decision of what to do.

The first thing we can look at is how teams versus teams compare to individuals versus individuals. Figure 6.14 summarizes the outcomes from four sessions. The

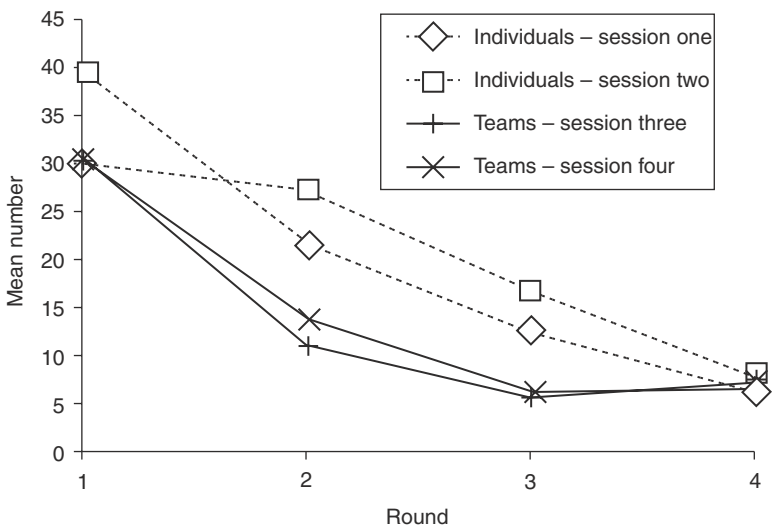

Figure 6.14 The mean number chosen when playing a 0.67-beauty contest four consecutive times, comparing individuals versus individuals and teams versus teams.
Source: Kocher and Sutter (2005).

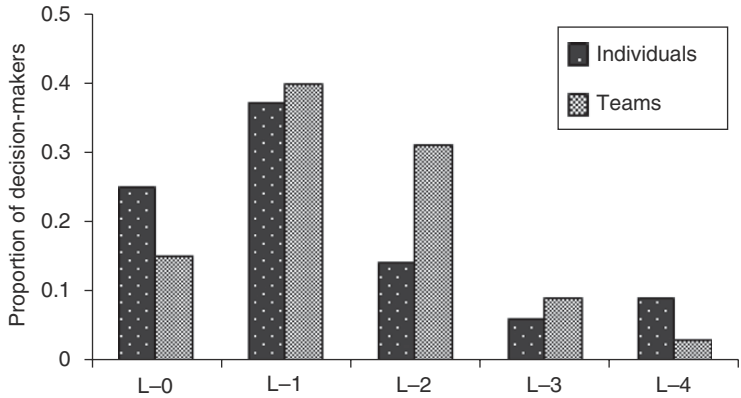

Figure 6.15 The proportion of individuals and teams of each type of level-k thinking.
Source: Kocher and Sutter (2005).

outcomes when individuals played with individuals (sessions one and two) are as we would expect (see Figure 6.2). Compare this with how teams played with teams (sessions three and four). There is no difference between individuals and teams in round one or round four. There is a difference between teams and individuals in rounds two and three, however. This suggests that teams learnt more quickly than individuals.

To explore things in a little more detail we can apply some of the models introduced earlier in the chapter. Let us start by looking at choices in round one. Figure 6.15 details the proportion of decision-makers, individuals or teams, that could be classified as level-k thinkers. The proportion of individuals who were type L-1 is similar to the proportion of teams that were type L-1. This can help explain why individuals and teams made similar choices in round one. It is noticeable, however, that there was a lower proportion of L-0 types and a higher proportion of L-2 types in teams.

In this setting, a type L-0 is someone who randomly chooses what to do. This presumably means they did not understand the game to some extent. A team is much less likely to fit this description, because one member of the team can explain the game to others. To be type L-0 we therefore need all three members of the team to misunderstand the game! Three heads are better than one. And, speaking much more generally, we typically find that teams perform better than individuals on **intellective tasks** in which there is a clear problem to solve. This makes intuitive sense, because team members can feed off each other and collectively work towards a solution. It needs just one member to solve the problem for the team to have solved the problem.

The p-beauty contest is more than an intellective task, however. It is also a **judgmental task**, because there is a need to predict what others will do. The general evidence is that teams are no better than individuals at solving such tasks. To give some intuition for why this makes sense, you need to recall issues covered in the previous chapter such as information cascades and conformity. One member of the team can easily drag others off in a strange direction. And, because there is no right or wrong answer, the team cannot tell it is going wrong. This can help explain why teams did not think too much ahead of individuals in the first round of the beauty contest.

Table 6.14 EWA parameter estimates comparing individuals and teams.

Decision-makers	Learning	δ	β	ϕ	ω
Individuals	Belief-based	1	0	0.70	1.06
	EWA	0.88	0	0.71	0.59
Teams	Belief-based	1	0	0.64	0.52
	EWA	1	0.09	0.64	0.53

Source: Kocher and Sutter (2005).

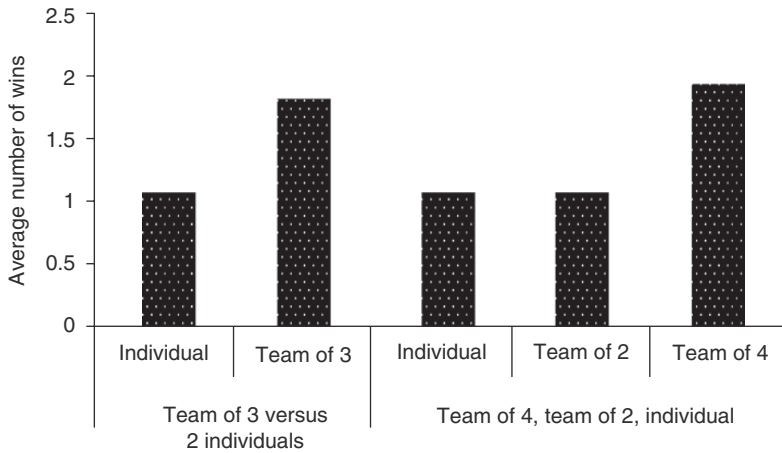

Figure 6.16 The average number of times (out of four rounds) that different types of decision-maker were closest to the optimal number.
Sources: Kocher and Sutter (2005) and Sutter (2005).

In the second round of the contest it is clear that teams reacted more to experience than individuals. To capture this, Kocher and Sutter fitted the EWA model to the data. The results are given in Table 6.14 (and can be compared with Table 6.10). Teams had a noticeably lower decay rate ϕ than individuals, meaning teams adapted more quickly. Teams also had a higher δ, meaning missed payoffs were weighted the same as received payoffs, as is the case in belief-based learning. This would suggest that teams more quickly learnt to think one step ahead. It takes individuals time to catch up.

If teams are one step ahead of individuals then they should earn a higher payoff. This makes it particularly interesting to see how teams perform against individuals. Recall that, until now, teams were pitted against other teams. The study by Kocher and Sutter (2005) also looked at what happens when a team of three play against two individuals. Sutter (2005) looked at what happens when a team of four, a team of two and an individual play against each other. Figure 6.16 details the average number of times that different types of decision-makers 'won' by being closest to the optimal number. The conclusion is clear: teams did better than individuals, and teams of four did better than teams of two. So, you would be better to trust your investment decisions to a team!

6.4.2 The sophistication of teams

The clear lesson from the p-beauty contest seems to be that teams outperform individuals. The victory is far from overwhelming, however. In particular, it rests on teams learning more quickly than individuals. In round one there is no big difference between individuals and teams. To broaden the picture, it is interesting to look at a study by Sutter, Czermak and Feri (2013). In the study, subjects played 18 different games, with some subjects playing as an individual and some in a team of three. Each game was played only once, and so the focus was very much on initial play rather than learning.

In all the 18 games played there was a unique Nash equilibrium, but the games differed in complexity. In some games there was a dominant strategy, in others not. In some games the Nash equilibrium could be found by iterated deletion of dominated strategies, in others it could not. In some games it took three steps of elimination of dominated strategies, in others one or two. Figure 6.17 shows the proportion of decision-makers who chose consistent with the Nash equilibrium for progressively more complex games.

The simplest situation is when a decision-maker has a dominant strategy and so one step of elimination of dominated strategies reveals the Nash equilibrium strategy. This situation is simple enough that most decision-makers, individuals or teams, played the Nash equilibrium strategy. At the other extreme are games that are not solvable by elimination of dominated strategies, or when three steps of elimination are required to reveal the Nash equilibrium. This situation appeared complex enough that few decision-makers, individuals or teams, played the Nash equilibrium.

The difference between teams and individuals comes when two steps of elimination of dominated strategies are required to reveal the Nash equilibrium. In these situations of intermediate complexity, teams are significantly more likely to choose

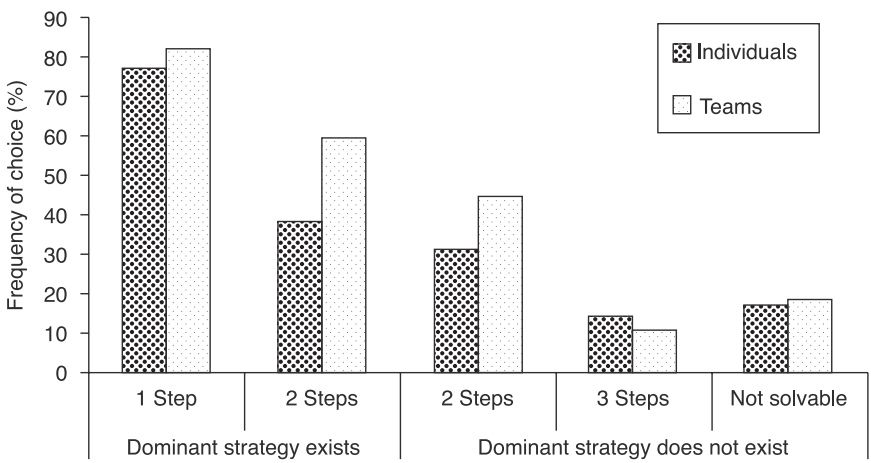

Figure 6.17 The proportion choosing the Nash equilibrium strategy depending on whether a dominant strategy existed or not and whether the game was solvable in one, two, three or no steps of iterated deletion of dominated strategies.
Source: Sutter, Czermak and Feri (2013).

the Nash equilibrium strategy. This is a fascinating result that allows us to pinpoint when teams can outperform individuals the first time they play a game. Teams are better than individuals provided the game is not too simple or too complex. Note that the p-beauty context would be counted as complex because more than three steps of elimination are needed to reveal the Nash equilibrium. So, it makes sense that we found no big difference between teams and individuals in the first round in that case.

The study by Sutter, Czermak and Feri also asked subjects what they thought the opponent would do. Specifically, the individuals or teams were asked what action they expected their opponent to choose – **first-order belief** – and what action they expected their opponent would expect them to choose – **second-order belief**. Eliciting such beliefs can tell us something about the sophistication of the decision-maker. For example, we can see if the action chosen is a best reply to the first-order belief. Consistent with what we see in Figure 6.17, teams appeared more sophisticated. Teams were more likely to choose the best reply to their first-order belief and more likely to expect the Nash equilibrium outcome. It is not by accident, therefore, that teams chose the Nash equilibrium strategy. Teams seemed to reason through the game more than individuals.

6.4.3 Are teams smarter?

As already mentioned, I shall periodically come back to the distinction between teams and individuals through the remainder of the book. We have already seen enough, however, to know that the behavior of teams can be different from that of individuals. Clearly, therefore, it is important to take seriously the fact that many economic decisions are taken by teams. To put things in context, two observations seem pertinent.

The first thing to observe is that teams appear to behave more like *Homo economicus* than do individuals. So, teams better fit the standard economic model. In the p-beauty contest, for instance, teams converged more quickly towards the Nash equilibrium. In the games we have just been looking at, teams were initially more likely to play the Nash equilibrium. Teams, therefore, fit somewhere between the 'rationality' of *Homo economicus* and the heuristically driven *Homo sapiens*.

But let us not get carried away. The second thing to observe is that teams are not much different from individuals. In the first and fourth round of the p-beauty contest, for instance, teams and individuals are indistinguishable. We have just seen that teams and individuals are also indistinguishable in relatively simple or complex games. On the scale of rationality between *Homo economicus* and *Homo sapiens*, therefore, teams are sitting somewhere next to *Homo sapiens*. This should not be a surprise given that teams are, of course, composed of individuals!

To draw this section to a close, I'm giving a reminder that it is often very useful to treat the strategic interaction between team members as a black box that we do not need to worry ourselves about. The fact that teams appear to behave similarly to individuals a lot of the time should reassure that we can get away with this. We, hopefully, do not miss much by approximating a team with an individual. We have also seen, however, that in some contexts teams can differ from individuals. It is important to keep this in mind, particularly in applications. See, for example, section 6.8 on monetary policy by committee.

6.5 Summary

In trying to understand strategic behavior we started with the notion of the Nash equilibrium and elimination of dominated strategies. A basic premise is that people should not play dominated strategies. A slightly more involved premise is that they should play a Nash equilibrium strategy. In some games this gives us a very nice prediction, in principle, of what people should do because there is a unique Nash equilibrium or a unique strategy that survives iterated elimination of dominated strategies.

Evidence on how people behaved in the p-beauty contest was enough to suggest, however, that the Nash equilibrium and elimination of dominated strategies are not enough to explain behavior. This led us to the ideas of level-k thinking and semi-sophisticated and sophisticated behavior. The focus here is to model the various ways that people might reason in strategic situations and not 'get as far as' the Nash equilibrium. On a related note, we looked at quantal response equilibrium, which tries to model what happens if people make mistakes.

In many contexts the notion of the Nash equilibrium and elimination of dominated strategies does not get us very far, because there are several Nash equilibria. That means that we need to think how people might reason which equilibrium 'makes more sense'. By looking at focal points, risk and Pareto dominance, and backward and forward induction, we have seen that there are various ways to do this.

We then looked at how people might learn with experience, contrasting reinforcement learning and belief-based learning. We also looked at how learning models can be used to predict what will happen after people have learnt.

Finally, we contrasted behavior in teams with that of individuals. When looking at teams we have two levels of strategic interaction, that within the team and that between teams. The ways we have seen to model individual behavior can be applied to model behavior within the team and between teams.

One thing you may have noticed by now is that modeling behavior in strategic contexts is not easy. There is a lot going on, because we have all the choice arbitrariness, risk and uncertainty that we had in Chapters 2 to 5 compounded by the interdependence between people. Furthermore, different games raise very different issues. For this reason, anyone expecting a unified theory of behavior in strategic contexts is probably going to be disappointed. Sometimes level-k thinking will work well, sometimes not; sometimes backward induction will work well, sometimes not. This, though, is not bad news. We now have lots of different ways to help us think about strategic behavior.

6.6 Auctions

A nice example of when strategy matters is bidding at an auction. We have already encountered auctions a few times, in Chapters 2 and 4, but now I want to look at them in a bit more detail. It seems apt to start by saying that auctions are one of the most important ways that goods are sold. Examples are as diverse as the selling of antique furniture or paintings at Sotheby's, livestock at a cattle market, cars, houses, goods on eBay, and governments selling tracts of land for oil drilling. Here I want to focus on two particularly interesting aspects of auctions: revenue equivalence and the winner's curse. In Chapter 11 we will look at some other aspects.

6.6.1 Revenue equivalence

There are various different types of **auction format**. I will look at four of the more common formats here (and a few more in Chapter 11).

In an **English auction** the auctioneer starts the bidding at a low price and incrementally increases the price until only one bidder is willing to pay that price. This is the most common auction format, and is what you would see, for instance, at Sotheby's. In a **Dutch auction** the auctioneer starts the bidding at a high price and incrementally lowers the price until someone is willing to buy at that price. This auction gets its name from the way tulips and other flowers are sold in Holland.

In a **first-price sealed bid auction** potential bidders are asked to submit simultaneously and independently the amount they are willing to pay for the item. The highest bidder wins and pays the price he bid. This is a standard format for government tendering.

Finally, in a **second-price sealed bid auction** everything is the same as the first-price sealed bid auction, except that the highest bidder has to pay only the amount bid by the second highest bidder. EBay is a combination of this and an English auction.

A natural question for anyone looking to sell using an auction is which type of auction will maximize the revenue from the sale. Auction theory offers a specific and intriguing answer to this question: the **revenue equivalence theorem** says that in many cases all four auction formats result in the same revenue if bidders behave optimally.

One case in which this theorem applies is an **independent, privately known value auction**. In such an auction different bidders may value the item differently, but each bidder knows her own value, and knows that her own value does not tell her anything about the likely values of others. This is the situation found in many common auctions, such as the auctioning of an antique, car, house or consumer item on eBay. For example, consider a painting up for auction that Emma is thinking of bidding for. She knows that different people will likely value the painting differently, some liking it and some not, but she knows only her own preferences, and not those of others.

If Emma decides she wants the painting, and values it at $1,000, how much should she bid? The optimal way to bid in this type of auction depends on the auction format.

In an English or second-price sealed bid auction, she should bid up to the amount she values the item – i.e. $1,000. There is no point bidding more because she might end up paying more than she thinks it's worth. There is no point bidding less because she will have to pay only the amount the second highest bidder is willing to pay. For example, if Alan is the second highest bidder at $950, she has to pay only $950 (even if she bid $1,000), and so there would have been no gain in bidding, say, $980 rather than $1,000.

Things are different in the Dutch and first-price sealed bid auctions. Now Emma needs to bid below the amount she values the item – say $980. This is called **shading the bid**. If she does not shade her bid then she can never make any profit, because she will have to pay what she bids. But she does need to bid high enough to have a good chance of winning. It is quite complicated knowing exactly how much to bid, given these competing incentives, but we do not need to worry too much about that here.

Table 6.15 Bidding in the Magic game. The Dutch auction appears to raise more revenue than the first-price auction. There is no difference between the English and second-price auctions.

	First-price		Dutch		Number of cards sold at higher price		
	Revenue	**Bidders**	**Revenue**	**Bidders**	**First-price**	**Neither**	**Dutch**
Collection 1	$431.25	32	$446.35	56	12	12	63
Collection 2	$327.05	42	$348.45	88	22	5	59
	Second-price		**English**		**Number of cards sold at higher price**		
	Revenue	**Bidders**	**Revenue**	**Bidders**	**Second-price**	**Neither**	**English**
Collection 3	$85.50	27	$79.50	40	3824	8	20
Collection 4	$517.05	43	$600.40	38		9	65

Source: Lucking-Reiley (1999).

Let's now go back to the revenue equivalence theorem. We can see now why the English and second-price auction should give the same revenue, and the Dutch and first-price should give the same revenue. What remains to be shown is that a first-price or Dutch auction gives the same revenue as a second-price or English auction. In a first-price auction Emma needs to shade her bid, and so bids less than she would in a second-price auction. If she wins a second-price auction, however, she has to pay only the amount bid by the second highest bidder. The neat thing is that these two effects cancel each other out: the highest bid in a first-price auction will, on average, equal the second highest bid in a second-price auction! Thus, all four auctions give the same expected revenue.

Well, that's the theory. Lucking-Reiley (1999) put the theory to the test in a field experiment that involved selling 'Magic cards' to online devotees of a game called 'Magic: The Gathering'. Four different collections of cards were independently auctioned using two of the four different auction formats. The results are summarized in Table 6.15. Collections one and two allow us to compare first-price and Dutch auctions. We see that the Dutch auctions raised more revenue than the first-price auctions, with more cards selling for a high price. Collections three and four allow us to compare second-price and English auctions. Here we see no significant differences in revenue. Overall, therefore, we see somewhat mixed support for the revenue equivalence theorem. A similar conclusion has been found in laboratory studies, but with different results. Indeed, in the laboratory, first-price auctions have been seen to raise more revenue than Dutch auctions, and second-price auctions more than English auctions (see Research methods 6.5).

Anyone selling an item might, therefore, want to think more closely about which auction format to use than the revenue equivalence theorem might suggest. The popularity of the English auction is possibly a signal as to what many sellers think is best. To progress beyond such conjectures we need to know more about how people typically bid in auctions. We will get some insight into this as we look at the winner's curse.

❖ RESEARCH METHODS 6.5 FIELD EXPERIMENTS

Given that auctions are so common and easy to take part in, they seem one natural setting for field experiments or source of field data. In his study, Lucking-Reiley needed to find a setting flexible enough that he could auction the same things using four different auction formats. The online community of 'Magic: The Gathering' offered such an opportunity, because there was already use of the English, first-price and Dutch auctions, so only the second-price auction would be new to bidders.

The usual potential advantages of field experiments over laboratory experiments are seemingly present in this study, namely that the subjects should be experienced in valuing the items and bidding for them. The results illustrate why this can prove important. That the Dutch auction raised more revenue than a first-price auction goes against what has been observed in laboratory experiments. The probable explanation is that more people wanted to take part in the Dutch auction. This, unpredictable, increase in the number of bidders could not have happened in a lab experiment, for which the number of bidders is fixed by how many subjects the experimenter invites to the lab.

6.6.2 Winner's curse

Recall that, in a privately known value action, Emma knew how much the item was worth to her but not to others. The opposite extreme is a **common value auction**. In this type of auction a bidder will not know how much the item is worth until she has got it, but does know that others will value it as much as her. An example of this would be drilling rights to a tract of land that may or may not have oil. Only having purchased the land and drilled for oil can it be known for sure whether there is oil, but, if there is, any oil company would have wanted to buy the land. A slightly more mundane example would be people entering a charity auction at the local fete not knowing what the mystery winning prize is.

In a privately known value auction the bidder who wins is likely to be the one who put the highest value on the item – for example, the one who most liked the painting. In a common value auction the bidder who wins is likely to be the one who was most optimistic about the value of the item. Optimism is sometimes a dangerous thing. If the bidder is more optimistic than others then she may be too optimistic and may pay more for the item than it is actually worth. This is called the **winner's curse**. Examples of the winner's curse abound. Oil companies have overpaid for land, telephone companies have overpaid for the rights to transmit signals, TV channels have overpaid for the rights to televise sporting events, and many have been disappointed by the prize they won at the local fete.

To look in more detail at why the winner's curse can happen, it is useful to first look at how a person should bid in a common value auction. So, imagine Emma is entering a charity English auction for a prize of unknown value. She likely gets some

signal of how much the prize might be worth. From this she gets an initial expectation of the value of the item. For example, she might have an initial expectation that the prize is worth $20. Surely she should bid up to $20? No. She should reason that, if she wins, she must have had the most optimistic initial expectation. For example, if she wins bidding $20 then everyone else must have thought the prize was worth less than $20. If everyone else thinks the prize is worth less than $20, maybe she was overoptimistic? She should lower her expectations, and so bid up to something less than her initial expected value.

With this we can now list some distinct reasons why Emma might suffer the winner's curse. (1) She could be a risk-loving person who bids high knowing that the probability that the prize is worth so much is small, but she is happy to take the risk; ultimately, the gamble does not pay off. (2) She bids optimally, and so lowers her bid below her initial expectation, but still proves unlucky, because the prize turns out to have a lower value than most expected. (3) She naïvely bids up to her initial expectation, and is unlucky. (4) She bids an irrationally high amount; this is most likely in a second-price auction, if she expects the second highest bid will be relatively low.

I now want to look at two studies. The first, by Crawford and Iriberri (2007), looks to explain the winner's curse using level-k thinking. The focus, therefore, is on the behavior of inexperienced bidders. The second, by Selten, Abbink and Cox (2005), uses impulse balance and so puts the focus more on learning from experience.

We know that in level-k thinking what we assume about type L-0s is crucial. Crawford and Iriberri consider two alternatives. First, that type L-0s bid randomly; or, second, that they bid naïvely their initial expectation. In different auction formats bidding behavior will differ, but Table 6.16 summarizes the main predictions of level-k thinking by comparing level-k bidding with the equilibrium bidding that would happen if everyone used the optimal strategy.

Table 6.16 Bidding strategies of a type L-k comparing what happens with two different assumptions about type L-0s.

Type	Random	Initial expectation
L-0	Bids randomly	Bids up to her initial expectation
L-1	Bids up to her value, and more than the equilibrium amount Because type L-0s are bidding randomly, that she wins tells her nothing about the expected values of others, so she sees no reason to lower her bid	Bids less than her value, and less than the equilibrium amount Because type L-0s are bidding more than in equilibrium, she thinks that if she wins she must have been very optimistic about the expected value, so she lowers her bid to compensate
L-2	Bids less than her value, and less than the equilibrium amount Because type L-1s are bidding more than in equilibrium, she thinks that if she wins she must have been very optimistic about the expected value, so she lowers her bid to compensate	Bids more than an L-1 type and possibly more than the equilibrium amount Because type L-1s are bidding less than in equilibrium, she thinks that if she did win she was not necessarily optimistic, so there is less need to lower her bid

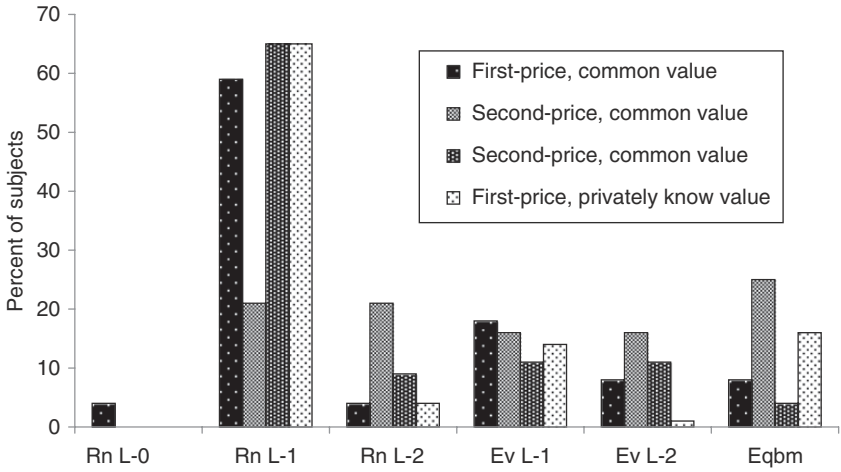

Figure 6.18 Percentage of subjects of each type in four auction experiments. Rn assumes type L-0s choose randomly and Ev means they choose initial expectation. *Source:* Crawford and Iriberri (2007).

We see that if people of type L-0 bid randomly then people of type L-0 and L-1 are particularly vulnerable to the winner's curse. If people of type L-0 bid initial expectations then people of type L-0 and L-2 are vulnerable. Figure 6.18 details the percentage of subjects that Crawford and Iriberri estimate were of each type in four different auctions. It's fairly clear that most subjects were classified as being type L-1 with an expectation that type L-0s bid randomly. These types are vulnerable to the winner's curse because they naïvely ignore the information that the bids of others convey. They fail to take into account that, if they win, they must have been relatively optimistic about the expected value.

If people are vulnerable to the winner's curse because of naïve bidding then we might expect that they should learn with experience. Things are not so simple, however, because Emma, say, could bid 'too much' but be lucky and win a great prize, or bid 'the optimal amount' and be unlucky because the prize is less than everyone expected. This makes it potentially difficult to learn from experience.

The study of Selten, Abbink and Cox uses impulse balance equilibrium and a closely related idea of learning direction theory to try and understand how people learn when playing auctions several times. The basic idea of **learning direction theory** is that someone goes with their impulse and moves in the direction of improved payoffs. In the context of an auction, the person who wins almost certainly bids more than they needed to and so will lower their bid subsequently. The person who does not win might regret not bidding more and so increase their bid subsequently. This gives an interesting contrast between overpaying and a missed opportunity.

Selten, Abbink and Cox find that learning direction theory and impulse balance equilibrium fit well how subjects behaved in the experiments. This means that bids did not converge over time to that predicted by equilibrium but, instead, tended to fluctuate as subjects overpaid some periods but then regretted not bidding more in other periods. This is despite playing 100 successive auctions. To see this a bit more clearly,

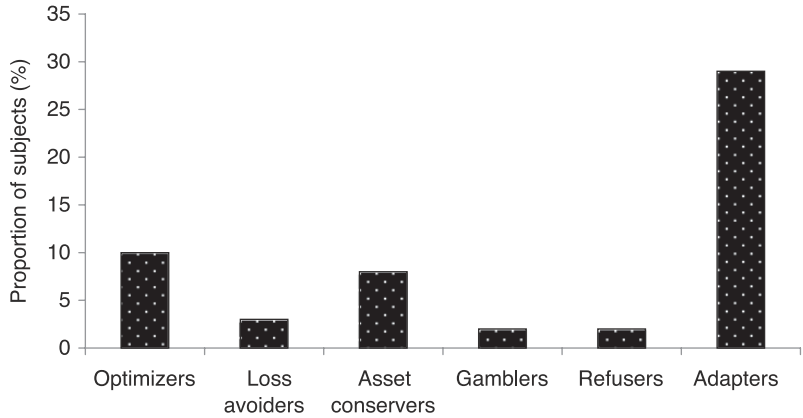

Figure 6.19 Categories of bidder. Adapters behaved consistently with learning direction theory, increasing their bid if they had previously missed out and lowering their bid if they had previously paid too much.
Source: Selten, Abbink and Cox (2005).

Figure 6.19 shows what happens if we try to put subjects into categories depending on how they bid. Clearly, a lot of subjects were adapters, and this is the category that most closely fits learning direction theory. Of the other categories, gamblers are worth a mention. Gamblers were seemingly interested in high gains and appeared to fall for the gambler's fallacy; if they overpaid one auction they tended to bid more next time, possibly thinking their luck must change.

Such gambling is the opposite of what is predicted by learning direction theory. So, learning direction theory cannot capture what everyone does, but it does do a good job of capturing what many do. It also gives us an explanation of why the winner's curse does not disappear with experience: naïve bidding sometimes pays off, and so people need not realize that it is not the best thing to do.

This is an important point, not only in looking at auctions, because it demonstrates that learning need not converge to Nash equilibrium. The reason it did not do so here is because payoffs were random. If payoffs are going up and down, through good luck and bad luck, more quickly than Emma learns the optimal way to bid, she might never learn the optimal way to bid. Auctions, therefore, are one area of life in which it pays to read the economic textbook and not rely on trial and error!

6.7 Learning to coordinate

The next thing I want to do is go back and look in a bit more detail at **coordination games** – games in which people want to coordinate their strategies to mutual benefit. Coordination is vital in many areas of life, whether it be the world deciding how to tackle climate change or two friends trying to meet in a crowded cafeteria. So, it is very interesting to look at coordination games and see what we can learn. In section 6.2.3 we saw some evidence that people can coordinate using focal points, and in section 6.2.4

we saw how people might use forward or backward induction to coordinate. In this section I want to ask more generally how well people can coordinate. In doing so we shall look at two very different types of coordination problem (see also section 6.9.2).

6.7.1 Weakest link games

In a **weakest link game** or **minimum effort game** a group of people independently must decide how much effort to put into a group task. Higher effort is costly, but the output of the group is determined by the smallest effort that someone puts in. A nice example to illustrate is going somewhere by plane. To get to her destination on time Emma would probably need lots of people to put in effort: the taxi driver to get her to the airport, the pilot to be there, the cleaner to clean the plane, the serviceman to put in the fuel, etc. If anyone puts in low effort, Emma may not leave on time. The plane is going nowhere, for instance, if the pilot turns up late.

Table 6.17 gives the payoffs that are usually used in weakest link game experiments. In this case people need to choose effort between 1 and 7. The best thing for the group and the individual, the Pareto-dominant Nash equilibrium, is for all members of the group to put in high effort. This gives the highest payoff of 13. The safest thing for the individual, however, the risk-dominant Nash equilibrium, is to put in low effort. This guarantees a payoff of 7. Choosing higher effort is risky, because it takes just one member of the group to choose low effort for the high effort to be wasted and costly.

The weakest link game has the tension between the Pareto-dominant and risk-dominant Nash equilibrium that we looked at in section 6.2.4. Indeed, the payoffs in Table 6.5 are a condensed version of Table 6.17. Ideally, we want people to coordinate on the Pareto-dominant Nash equilibrium.

Figure 6.20 gives a fairly typical picture of what happens in a weakest link game experiment. In this case there were four players who played together for ten rounds. We see that there was coordination failure. That subject 2 chose low effort in every round meant that the high effort of subjects 1, 3 and 4 was wasted. Over the rounds the effort of these three subjects falls as they learn effort is not going to pay. Figure 6.21 summarizes some aggregate results from a series of experiments by van Huyck, Battalio and Beil (1990). The dramatic rise in the proportion of subjects

Table 6.17 Payoffs in a weakest link game.

Your choice of effort	Smallest choice of effort in the group						
	7	**6**	**5**	**4**	**3**	**2**	**1**
7 (high)	13	11	9	7	5	3	1
6		12	10	8	6	4	2
5			11	9	7	5	3
4				10	8	6	4
3					9	7	5
2						8	6
1 (low)							7

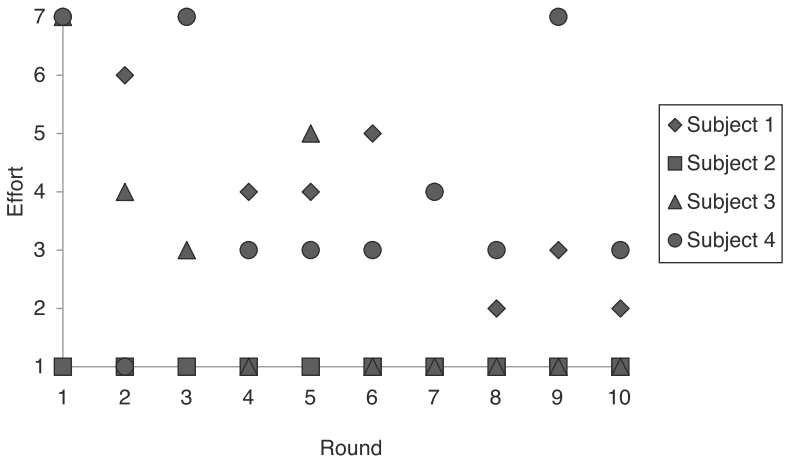

Figure 6.20 An example of coordination failure in a weakest link game. That subject 2 chose an effort level of 1 meant the high effort of the other three was wasted.
Source: Cartwright, Gillet and van Vugt (2013).

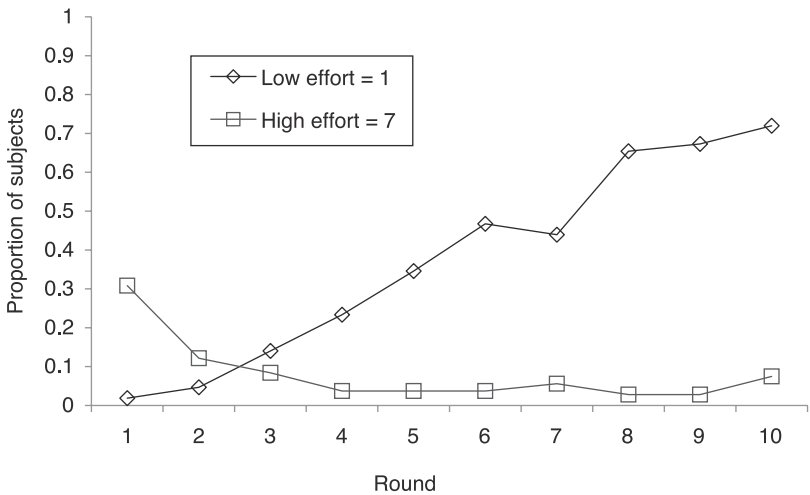

Figure 6.21 The proportion of subjects choosing 1 and 7 in weakest link games. There is a large rise in the proportion choosing low effort.
Source: van Huyck, Battalio and Beil (1990).

choosing low effort clearly demonstrates coordination failure. Nearly 80 percent of subjects are choosing low effort by round ten!

Such coordination failure is not good. With results like this I would not be betting on Emma's plane leaving on time. So, how can we make things better? Many suggestions have been tried out in the laboratory, with varying degrees of success. I shall look at some shortly (see also Research methods 6.6). First, however, it is useful to get a better idea of why there is coordination failure.

❖ RESEARCH METHODS 6.6 MONEY VERSUS REAL EFFORT

A common interpretation of choice in the weakest link game is a choice of effort. Experimental subjects are usually asked how much money to give, however. Do we get different results if subjects are asked to put in effort rather than give money? This is a question about the external validity of weak link game experiments. Bortolotti, Devetag and Ortmann (2009) go some way to answering this question by getting subjects to do a real effort task. Subjects were asked to sort and count, within a given time interval, a bag of one, two, five and ten cent euro coins. Group payment was based on the worst performance in the group. They found that there was significant coordination failure in the first few rounds, but increasing coordination over time. This suggests that real effort might be different from money, because of, say, social norms around working versus shirking. The effort task used was still very particular, however, and so concerns about external validity can never be completely removed. Moreover, significant coordination failure was still observed, just less than in standard experiments.

To do this we can use the models we looked at in sections 6.2 and 6.3. I will start by looking at initial choices. Costa-Gomes, Crawford and Iriberri (2009) looked at how well things such as level-k reasoning did at matching first-round choices in the weakest link game. Figure 6.22 gives the log likelihood of each model which is an estimate of goodness of fit. To give some interpretation to the numbers, the first two bars show what numbers we would get with a perfect fit and assuming subjects choose randomly. The clear winner is that subjects played the Pareto-dominant equilibrium of high effort. How can that be? How can people be choosing the Pareto optimum of high effort, yet we get coordination failure? The answer is that we need just one person to choose low effort to get coordination failure. Most subjects do choose high effort but there is typically at least one subject who chooses low effort, as in Figure 6.20. This causes coordination failure.

Next we can look at how subjects learn with experience. Crawford (1995) shows that a model of belief learning, similar to that discussed in section 6.3.2, can do a good job of explaining the change in behavior over time. In the model people updated their beliefs about what the minimum effort would be over time. A person who expected others to choose high effort but observed low effort would lower her future estimate of expected effort. If less effort was expected from others, she put in less effort herself. This means that the low effort of one person can drag down the effort of everyone in the group. That's also what we see in Figure 6.20.

The picture we get is one of coordination failure being caused by a small minority of people who choose low effort. Most do try high effort but the one or two people choosing low effort means that this is wasted. That suggests that we might be able to overcome coordination failure if we can stop this minority of people choosing low effort.

One way that seems to work quite well is to have people work in teams. Specifically, a study by Feri, Irlenbusch and Sutter (2010) looked at what happens

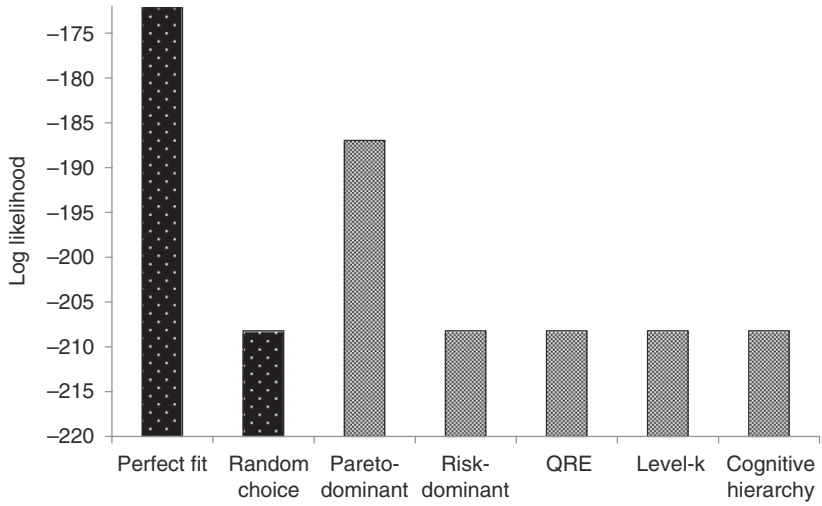

Figure 6.22 How well various models fit data from the weakest link game. Perfect fit and random choice give us an upper and lower bound on how good fit could be.
Source: Costa-Gomes, Crawford and Iriberri (2009).

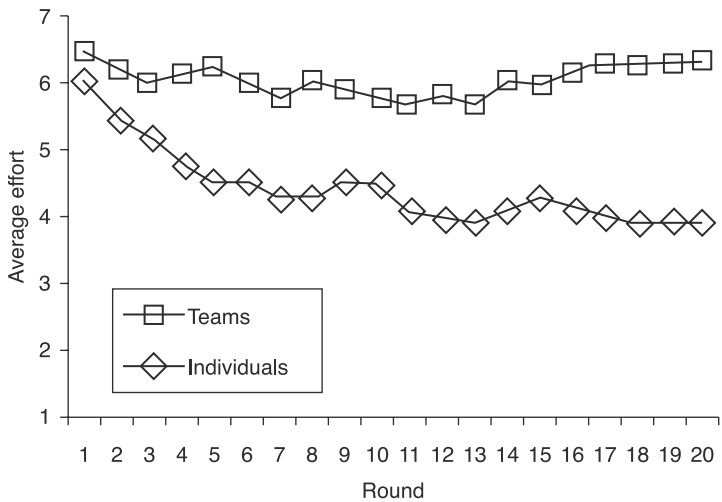

Figure 6.23 Average effort in a weakest link game comparing teams and individuals.
Source: Feri, Irlenbusch and Sutter (2010).

when decisions are made by teams of three. Figure 6.23 shows that average effort was a lot higher when decisions were made by teams than when they were made by individuals. More importantly, there is also no sign that average effort is declining over time. Teams, therefore, seem better able to coordinate on high effort. This is further evidence that teams can behave differently from individuals.

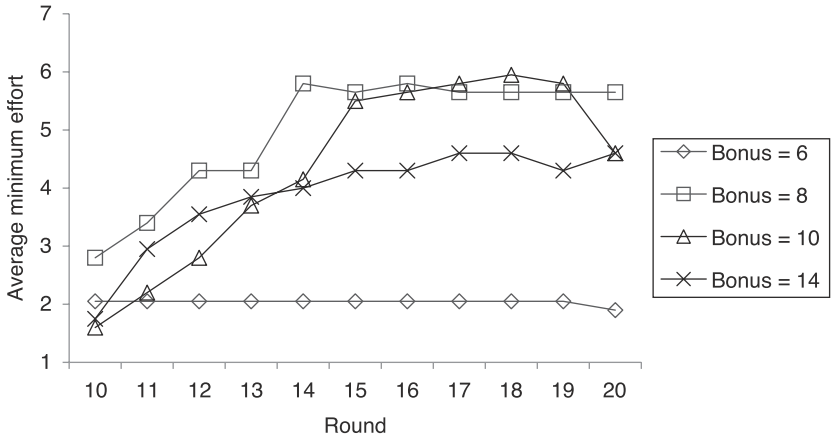

Figure 6.24 The minimum effort in the turnaround game. A change in the bonus to coordination allows groups to escape from low effort.
Source: Brandts and Cooper (2006).

The design of the study allows us to explore why this was. Before first interacting with fellow team members, each subject was asked what effort they intended to choose. These intentions did not differ from the choices made by individual decision-makers. So, being part of a team, in itself, made no difference. The difference comes from the interaction between team members. After this interaction many team members chose a higher effort than they had intended. So, team interaction does something to eliminate the minority of people choosing low effort. The way that the team reached a decision, whether by voting, leadership or consensus, seemed to make little difference (see section 6.8 for a similar result).

So, team decision-making seems to work. Is there any other way to keep effort high? I will look next at a series of studies by Brandts and Cooper. Not content with the challenge of maintaining high effort, they asked if coordination failure can be overcome. They did so by looking at a **turnaround game**, in which we can think of a manager trying to get a group of workers to increase effort. In the first ten periods the group plays the weakest link game and effort typically falls to a low level. In period 11 the manager steps in and tries things to increase effort. For example, the manager can talk to the group or increase the payoff incentives for high effort.

Let's look first look at what happens if the payoff incentives to high effort are changed. Figure 6.24 plots minimum effort from round ten onwards as the bonus from coordinating is increased from 6 to more than 6. The first thing to observe is that effort is low in period ten and does not improve if the bonus remains at 6. By contrast, if the bonus is increased, we observe a rise in effort. Interestingly, the change in the bonus to 8 increases effort more than a change to 10 or 14. So, it is not the size of the incentive that seems to matter, but that incentives were changed.

The most plausible explanation for the increased effort is that the change in incentives provides a focal event around which members of the group can escape from low effort. One critical part of this story is that subjects observe the effort of all members of the group, called **full feedback**. If they observe only the minimum effort

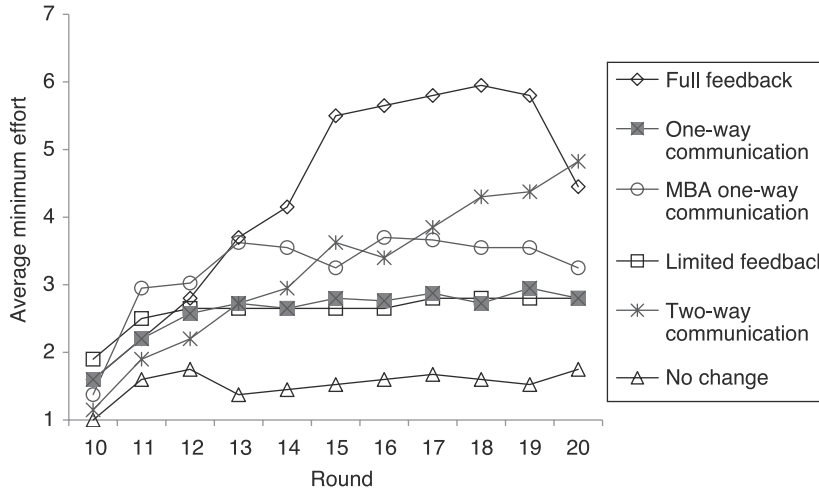

Figure 6.25 Average effort in the turnaround game. It is easier to overcome coordination failure with full feedback than with limited feedback, and with two-way communication than with one-way communication.
Sources: Brandts and Cooper (2007) and Cooper (2006).

in the group, called **limited feedback**, then things do not work so well. This is apparent in Figure 6.25, where we can compare full feedback and limited feedback. With limited feedback the increase in effort is much less. The story we get, therefore, is one in which the change in bonus is a focal event around which enough people increase effort to persuade others to increase effort. This time it is the subjects putting in high effort that drag up the effort of others.

Another possibility the manager could use is communication. With one-way communication, the manager can talk to the group and try to encourage high effort. With two-way communication, members of the group can talk to the manager, and perhaps suggest what might work well, and then the manager talks to the group. Figure 6.25 plots some results that allow us to see how well communication performs (see also Research methods 6.7). Overall, we see that one-way communication was not very successful, but two-way communication was. This masked the fact, however, that saying the right thing did usually work. Saying the right thing meant asking for a specific effort level and stressing the benefits of coordination. This usually led to an increase in average effort, when coupled with an increase in the bonus.

❖ RESEARCH METHODS 6.7 ARTEFACTUAL FIELD EXPERIMENTS

An **artefactual field experiment** is performed in the lab but with subjects likely to have **task-specific experience**. A study by Cooper (2006) provides an interesting example by comparing the performance in the turnaround game of undergraduate students to that of managers on an executive MBA program.

People on the program had at least ten years' work experience, including five years as managers, and had average annual earnings of over $120,000. The main findings of the study were that those on the MBA were very successful at turning things around and did so more quickly than students. They did so because they knew how to effectively communicate to people in the group better than students. These results offer reassurance that the turnaround game does capture things going on in the workplace.

It is possible, therefore, to get people to coordinate in a weakest link game, but it is not easy. It will be a similar story for the next type of game that I want to look at.

6.7.2 Threshold public good games

The basic idea of a **threshold public good game** is that a good will be provided for the benefit of everyone if and only if people contribute enough to exceed some **threshold**. For example, if a church roof needs to be replaced, and doing so costs $100,000, then members of the church somehow need to raise the $100,000 threshold. Or, if flatmates have to clean a flat to the standard their landlord requires, between them they need to put in time cleaning the flat to the threshold standard. In this game there is no trade-off between risk dominance and Pareto dominance. The best outcome is that they contribute enough to exceed the threshold. The problem now is a **conflict of interest** over how much each should contribute. For instance, one flatmate may do little to clean the flat, in the hope that another flatmate will do a lot.

The main question is whether people can coordinate by contributing enough, despite the conflict of interest. Figure 6.26 gives the results of a typical experiment. In this experiment subjects were put in groups of five, and were given 55 tokens. To produce the public good they needed to contribute between them 125 tokens (see Chapter 1 for more on this). We see that the group was not as good at coordinating as we might have hoped. In nine of the 25 rounds the total contribution was less than the threshold, and things got worse as time went on. This is primarily because subjects 2, 4 and 5 give less and try to free-ride on the contributions of the others.

Yet again, we see that people are not great at coordinating. Yet again, we need to ask what might help people coordinate better. One thing that might matter is the institution in place to collect contributions. For example, if contributions fall short of the threshold then we might be able to give people a **refund** on their contribution. Sometimes this is not possible; for instance, flatmates cannot get back the time they have spent cleaning. But sometimes it is possible; for instance, the church could give back donations if insufficient funds are raised. Another possibility is to give a **rebate** if contributions exceed the threshold.

Figure 6.27 gives for comparison the success rates of providing the public good observed in various experiments with various institutions. We want the success rate to be 100 percent, meaning that the public good is always provided. In the base case, when there is a refund, no rebate and contributions are made simultaneously, the success rate is a disappointing 40 to 60 percent. Can we do any better?

The most dramatic drop in the success rate comes from removing the possibility of a refund. In this case, the success rate drops to around 20 percent! So, refunds

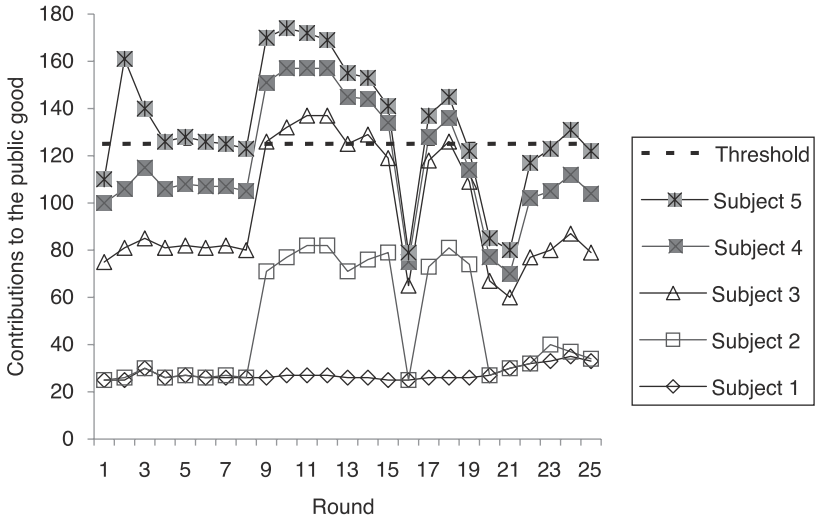

Figure 6.26 Cumulative contributions in a typical threshold public good experiment. The public good is provided if total contributions exceed 125.
Source: Alberti and Cartwright (2015).

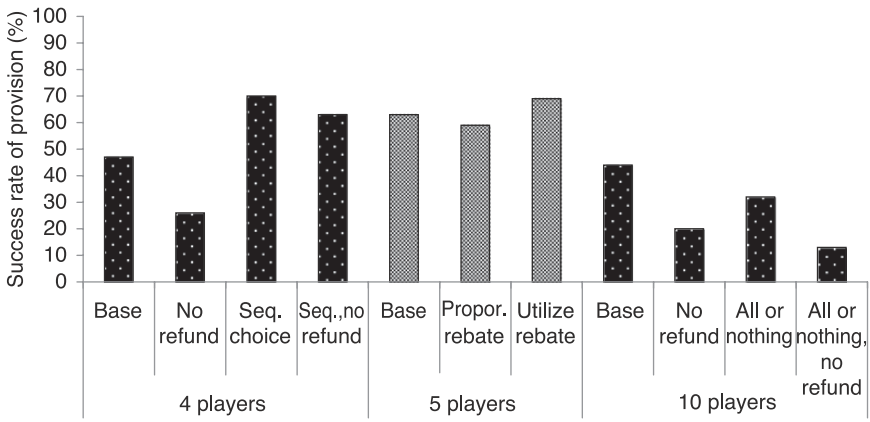

Figure 6.27 Success at providing a threshold public good depends on the institution.
Sources: Croson and Marks (2001), Marks and Croson (1998), Coats, Gronberg and Grosskopf (2009) and Cadsby and Maynes (1999).

work. The consequences of a rebate seem to depend on the type of rebate. I have given data for two types of rebate in Figure 6.27. With a proportional rebate, contributions above the threshold are given back in proportion to how much a person contributed. This does not seem to increase the success rate. What works better is a utilization rebate whereby excess contributions are used in some way to increase the value of the public good. An example of this type of rebate would be the church using any money it raises above the $100,000 to, say, paint the walls. Another thing

that seems to work well is sequential choice. In this case people can see contributions as they happen. For example, the church might put a chart on the wall to show people how close they are to the target.

I could go on talking about various other institutions that have been looked at, but I think that we have done enough to illustrate the main points. In both the weakest link game and the threshold public good game we have seen that it can be a problem to coordinate. We need to understand why, and ask how to make things better. The benefits of doing so could be huge: Emma gets her plane on time, the church gets its new roof, the flat gets cleaned, and much more besides. When looking at the weakest link game I focused on how people within or connected to the group could possibly make things better. When looking at the threshold public good game I have focused on how the broader institution can possibly make things better. Together, therefore, we have a variety of ways that might help people coordinate.

The main lesson, however, seems to be that coordinating is not easy; the best we have seen are success rates of around 70 percent. Fortunately, the next chapter is a bit more optimistic in showing that people can cooperate better than we might have expected. But before we move on I want to add a key extra ingredient into the mix.

6.7.3 Coordinating on networks

So far in this section, and chapter, we have focused on decision-making within a group. Often, however, decision-making takes place within a broader network, and this may have implications for how well people can coordinate. Consider, for instance, a large firm. Within that firm, colleagues will work together in small groups on specific projects. Any one of these projects could be a weakest link game or threshold public good game. Crucially, though, an employee will likely be involved, over time, in multiple projects with a large number of colleagues. I will look at two studies that illustrate the potential implications of this.

Let us look first at a study by Falk, Fischbacher and Gächter (2013). The network aspect of the study is depicted in Figure 6.28. Subjects were in a network with nine people and split into groups of three to play a weakest link game. Subject 5, for instance, was in group *B* with subjects 2 and 8 and a separate group *E* with

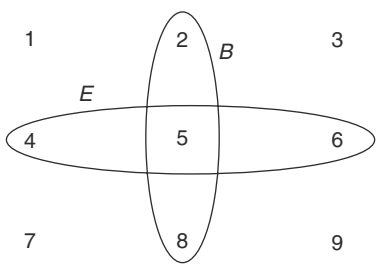

Figure 6.28 How nine subjects were allocated to groups within a network. Subject 5 was in two groups: group *E* with subjects 4 and 6, and group *B* with subjects 2 and 8.
Source: Falk, Fischbacher and Gächter (2013).

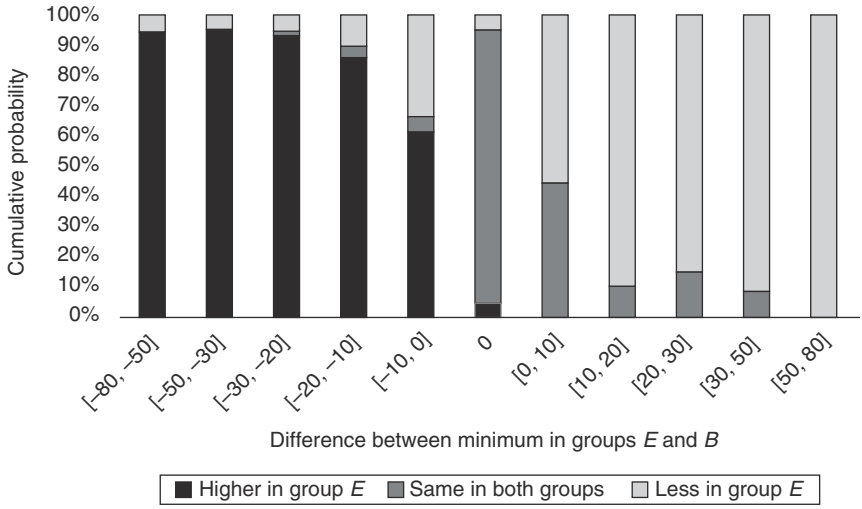

Figure 6.29 A subject is more likely to choose high effort in a group that has a higher minimum effort. Possible effort levels ranged from 20 (equivalent to 1 in Table 6.17) to 100 (equivalent to 7).
Source: Falk, Fischbacher and Gächter (2013).

subjects 4 and 6. Similarly, subject 4 was in the group with 5 and 6 and a separate group with subjects 1 and 7. Each subject, therefore, was involved in two overlapping coordination games.

One key question is whether behavior in one group influenced behavior in other groups. For instance, was subject 5 more likely to exert high effort when interacting with subjects 2 and 8 in group *B* if subjects 4 and 6 were choosing high effort in group *E*. The simple answer seemed to be 'No'. Instead, subjects coordinated their effort in each group with the others in that group. So, for instance, subject 5 might choose high effort in group *B* and low effort in group *E* to match the effort of others. This is illustrated in Figure 6.29, which plots the probability of a subject contributing more in one group or another depending on the relative minimum effort in each group. To interpret this figure, suppose that the minimum effort was higher in group *B* than group *E*; then we see that subject 5 would be much more likely to choose lower effort in group *E*. Only if the minimum was the same in both groups is it likely that the subject chooses the same effort in both groups.

The findings of this study illustrate the ability of people to adapt to their surroundings. They also highlight the importance of coordination within groups. We should not be surprised, for instance, to see a worker exerting high effort on a project with one set of colleagues and low effort on another project with a different set of colleagues. This tells us less about the attributes of the worker than it does the success the groups have had in coordinating on their respective projects. And, as we have seen above, there are various factors that can influence the likelihood of success.

For a different perspective on the role of networks, we can look at a study by Charness and co-authors (2014). The focus here shifts towards the role of network structure in facilitating coordination. In the study, a range of different networks were

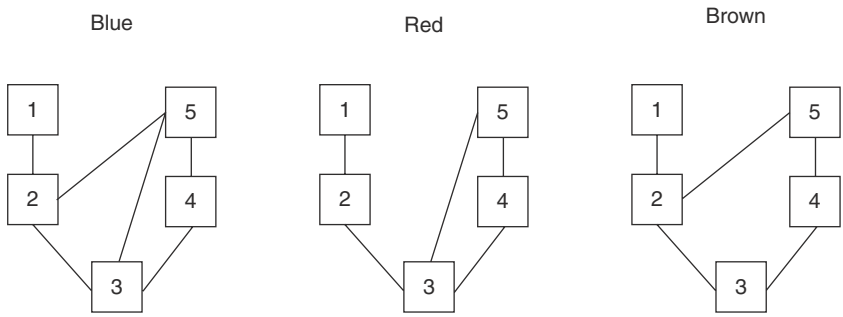

Figure 6.30 Three different network structures for five people.
Source: Charness *et al.* (2014).

considered, but it will be enough for us to focus on the three networks depicted in Figure 6.30. In each network there are five people who interact with all those they are linked to. For instance, in the Blue network subject 1 interacts with subject 2 while subject 2 interacts with subjects 1, 3 and 5. In interpretation, we can again think of colleagues in a firm working on projects with those they are linked to. Notice that the Red and Brown networks can be obtained from the Blue network by deleting the link between subject 5 and subject 2 or 3, respectively.

A coordination game then takes place on the network. Specifically, we can think of subjects as having a choice between high and low effort. If a subject chooses low effort then they get a guaranteed payoff of 50. If a subject chooses high effort then they get 33.33 for every subject they are linked to who also chooses high effort. To see how this works, let us look at the Blue network. Subject 1 is linked with only one other subject, and so the maximum payoff from high effort is 33.33. Thus, she has a dominant strategy to choose low effort. Subject 2, by contrast, is linked with two other subjects, ignoring subject 1. If both these subjects choose high effort then her payoff is 66.66, and she does best to choose high effort. If only one chooses high effort or neither chooses high effort then she does best to choose low effort.

Extending this logic, we can see that there are two Nash equilibria. There is an inefficient equilibrium where everyone chooses low effort and gets a payoff of 50. And there is an efficient equilibrium where subject 1 chooses low effort and subjects 2, 3, 4 and 5 choose high effort. In this latter case there are a cluster of subjects who benefit from coordinating on high effort. Interestingly, some of the subjects, namely 3 and 5, have more to gain from coordination, because they have more links in the cluster and so would get a higher payoff at the efficient equilibrium.

The Red and Brown networks similarly have an inefficient and efficient equilibrium. The nature of the efficient equilibrium differs across networks, however, as detailed in Table 6.18. Table 6.18 summarizes the efficient equilibrium and equilibrium payoffs for each of the networks. In the Red network we see that subject 2 should choose low effort, and so the cluster of subjects who have an incentive to choose high effort is smaller in the Red than Blue network. In the Brown network subject 2 does have an incentive to choose high effort but subjects 3 and 5 are no longer getting higher payoffs. This is because there are fewer links between subjects in the cluster.

Table 6.18 The efficient equilibrium (eqbm.) probability (%) of high effort and payoff compared to observed proportion (obser.) choosing high effort (%).

Subject	Blue			Red			Brown		
	Eqbm.	Payoff	Obser.	Eqbm.	Payoff	Obser.	Eqbm.	Payoff	Obser.
1	0	50	0	0	50	1	0	50	2
2	100	66.67	93	0	50	19	100	66.67	79
3	100	100	97	100	66.67	85	100	66.67	62
4	100	66.67	76	100	66.67	65	100	66.67	60
5	100	100	96	100	66.67	63	100	66.67	56

Table 6.19 The efficient equilibrium (eqbm.) probability (%) of high effort and observed proportion (obser.) choosing high effort (%) with incomplete information, where *p* is the probability of the Blue network.

Links	p = 0.2		p = 0.8	
	Eqbm.	Obser.	Eqbm.	Obser.
1	0	8	0	7
2	100	73	100	79
3	100	95	100	96

The greater clustering of subjects in the Blue network suggests that it should be easier to coordinate on high effort than in the Red or Brown networks. So, what happened? Table 6.18 details the observed proportion choosing high effort. In the Blue network we can see that subjects did indeed manage to coordinate successfully, with subjects 2, 3, 4 and 5 all choosing high effort most of the time. In the Red and Brown networks the proportion choosing high effort notably drops, even if it is still high enough that many groups coordinate on the efficient equilibrium.

We see here the potential importance of network structure to influence how well groups can successfully coordinate. In this instance, the more clustered the group, the easier it appears to be to coordinate. Network structure can be added, therefore, to the list of factors that can facilitate (or not) coordination. There is one other factor that I want to add to the list before we finish. Until now I have assumed that each subject knows the structure of the network. More realistically, a person might know who they are linked with but not know linkages in the network further afield. So, we can have incomplete information.

The study by Charness and co-authors provides insight into the potential role of incomplete information. In some treatments they made subjects aware of the Blue, Red and Brown networks but did not tell them which network they were in at any one time. Instead, subjects were told the probability of the Blue network and were told how many links they had in the network. Table 6.19 summarizes the main results. The main thing to observe is how the proportion of subjects choosing high effort is strongly influenced by the number of links they have. So, those with three links are more likely to choose high effort than those with two links. Again, network structure matters.

If we compare Tables 6.18 and 6.19, it is a moot point whether incomplete information has lowered efficiency. The proportion choosing high effort when the probability of the Blue network is 0.8 is slightly below that we observe when the Blue network is known for sure. The difference is not particularly pronounced, though. More generally, however, the evidence seems to suggest that incomplete information about network structure can be a hindrance to coordination.

6.8 Monetary policy by committee

The most important component of monetary policy is undoubtedly the interest rate. A change in the interest rate affects everyone, from savers, borrowers and investors to traders; it has long-term consequences for inflation, wages, jobs and growth. Indeed, it is hard to think of another routinely made policy decision that has such wide-ranging implications. So who would you choose to make the crucial decisions? I'll simplify your task by giving you some possibilities.

Traditionally, interest rates in the United Kingdom were decided by one individual, the Chancellor of the Exchequer (i.e. finance minister). In 1997 that all changed when responsibility for setting interest rates was devolved to the Bank of England's Monetary Policy Committee (MPC). The committee is composed of eight members, four from the Bank of England and four independent members, and meets once a month. Decisions are made by majority vote, with the governor's vote counting double. In the United States, interest rates have long been set by the Federal Open Market Committee (FOMC). The FOMC is similar to the MPC, but there are some notable differences. It has a few more members, meets a bit less often and requires a consensus or unanimous verdict.

So, the three basic possibilities are: (1) an individual, (2) a team that must reach a majority decision, (3) a team that must reach a unanimous decision. You could also mess around with the number of team members and frequency of decision. What is best? The fact that the United Kingdom and many other countries have changed to a committee system might lead you to believe that teams must be best. The change towards committee was not motivated by evidence that teams are best, however! Fortunately, a study by Blinder and Morgan (2005) gives us more evidence to go on. In so doing, it also provides a nice illustration of why the distinction between individuals and teams is important to take seriously.

The study compared individuals and teams in an explicit context of monetary policy. The economy was simulated using the standard model you can find in any macroeconomic textbook. Subjects were put in the role of a monetary policy-maker trying to keep unemployment at 5 percent and inflation at 2 percent. Each round of the experiment consisted of ten periods, which you can think of as ten months of policy-making. Every month the economy was subject to minor shocks. But, in one randomly selected month, there was a major shock. Roughly speaking, it was best to keep interest rates unchanged in the face of minor shocks and adjust in the case of a major shock. The skill is distinguishing a major shock from a minor shock (or an accumulation of minor shocks) and to move interest rates in the right direction by the right amount.

In each round the monetary policy was scored on how close unemployment and inflation were kept to target. Figure 6.31 shows how average score varied during the

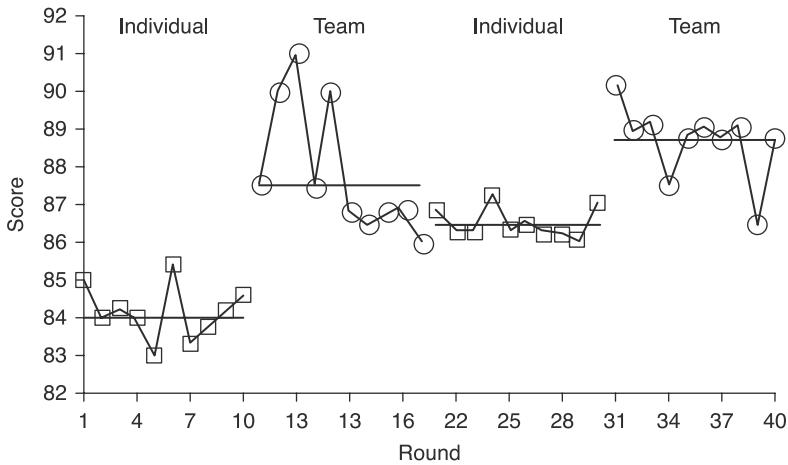

Figure 6.31 The score of monetary policy-makers trying to keep unemployment and inflation on target. Teams perform better than individuals.
Source: Blinder and Morgan (2005).

40 rounds of the experiment. In the first ten rounds, subjects acted as individuals. In the second set of ten rounds, subjects were put into teams of five. Some teams were told to decide by majority vote and others by consensus. We can see that the introduction of teams in round 11 led to an immediate jump in score. It is clear that this was not due to any learning effect, given the abruptness of the jump; teams simply performed better than individuals. In the third set of ten rounds, subjects went back to making decisions as individuals. The experience of having been in a team seemed to make subjects better monetary policy-makers! In the final set of ten rounds, subjects were again put back into teams of five. We see a further jump in performance.

We have already seen evidence in this chapter that teams are better than individuals. Each context is different, however, and so it is reassuring to see specific evidence that monetary policy by committee makes sense. Indeed, one interesting factor present in this context, which has been absent from everything we have looked at so far, is time. A particular concern of many macroeconomists was that decision-making by committee would lead to inertia, or a hesitation to act until sufficient evidence was accumulated. Such inertia could be costly, and seems especially likely if consensus is required. The study by Blinder and Morgan, however, found that teams changed the interest rate *more quickly* than individuals following a shock. Teams waited an average of 2.3 periods, or months, versus 2.45 periods for individuals. There was also no evidence that requiring consensus was any different from having a majority vote.

6.9 Industrial organization

It is natural in a chapter on strategic interaction to look at industrial organization, a field of economics that analyzes the firm and market structure. Just about any decision a firm makes should involve thinking through the reaction of others. For instance, when deciding on what price to charge, whether to build a new factory, how much

to advertise and whether to invest in research and development the firm needs to anticipate how rivals and customers will react. There are plenty of examples in this book of how behavioral economics can help us understand better the relationship between a firm and its customers (see, in particular, sections 4.6.2 and 7.6.4).

What I want to do in this section is focus on the relationship between a firm and its rival. Given that industrial organization is a huge field, I need to be somewhat selective in what I cover. The issues of limit pricing, market entry and quantity leadership are the topics I have chosen to briefly look at.

6.9.1 Limit pricing

To begin, let us imagine we have a market with a monopoly supplier. For example, suppose your local town has only one superstore. It is called BuyFood. Monopoly power gives BuyFood the chance to make maximum profit; but such profits may look attractive to rival firms. One such rival, called CheapFood, is deciding whether to enter the market by building its own superstore in the town.

The **incumbent monopolist**, BuyFood, would rather the **potential entrant**, CheapFood, did not enter the market. But is there anything BuyFood can do to deter entry? One possibility is called limit pricing. Essentially, it means that BuyFood precommits to low prices and thereby makes the market look unattractive to CheapFood. Low prices result in BuyFood making less profit, but it may still make more profit than it would have done if CheapFood had entered the market. Whether or not limit pricing can work will depend on how easily BuyFood can make the market look unattractive to CheapFood.

The timeline of a simple limit pricing game is depicted in Figure 6.32. The extra thing added into the mix is some uncertainty. BuyFood may be a low-cost seller, in which case it can easily lower prices without sacrificing profit. Or, it may be a

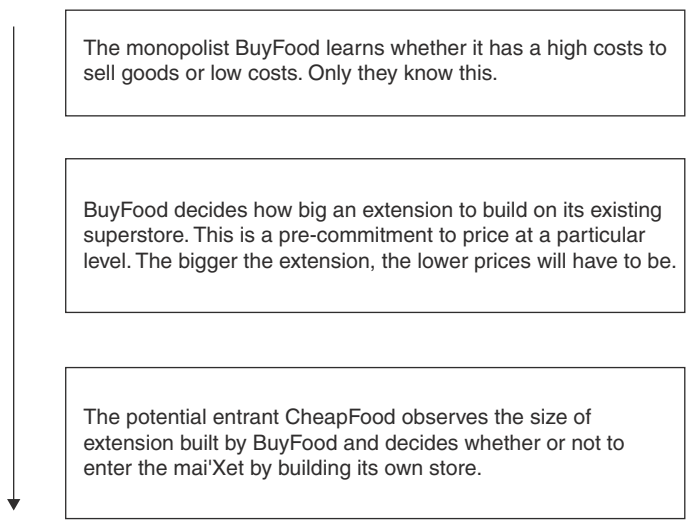

Figure 6.32 Timeline in a limit price entry game.

Table 6.20 Possible equilibria in the limit price entry game.

Scenario	Equilibrium	Monopolist	Entrant
Hard to enter	Pooling	Chooses a medium-sized extension irrespective of cost	Stays out
Easy to enter	Separating	Chooses a large extension if low cost and small extension if high cost	Stays out if a large extension and enters if a small extension

high-cost seller, in which case lower prices would require a big drop in profit. If CheapFood knew the costs of BuyFood then things would be simple: it would enter the market if BuyFood has high costs, but stay out of the market if BuyFood has low costs. But it is more realistic to assume that CheapFood does not know the costs of BuyFood.

The lack of information on cost turns the limit pricing game into a signaling game. The size of extension that BuyFood builds sends a signal to CheapFood about its likely costs. As we have already seen in section 5.3.4, things can get complicated pretty quickly when looking at signaling games. In addition, there are many different permutations of possible costs for BuyFood and ease with which CheapFood can enter. So, I will keep things simple by looking at the two possible equilibrium outcomes summarized in Table 6.20, based on two different scenarios.

In the first scenario, it is hard for CheapFood to enter the market – hard enough that the firm should enter only if it can be confident that BuyFood has high costs. What if BuyFood chooses the same size of extension irrespective of its costs? That way CheapFood cannot infer anything about the costs of BuyFood and stays out of the market. We get a pooling equilibrium. The key thing here is that BuyFood builds a bigger extension than it would ideally like when it has high costs. It does so to deter entry.

In the second scenario, it is easy for CheapFood to enter the market – easy enough that the firm will enter unless it knows for sure that BuyFood has low costs. In this scenario, pooling no longer makes sense, because BuyFood will need to signal that it has low costs. So, it will build a large extension if it has low costs and a small extension if it has high costs. CheapFood can, thus, infer the costs of BuyFood, and we get a separating equilibrium. This time BuyFood builds a bigger extension than it would ideally like when it has low costs.

The equilibria in Table 6.20 are a hint of how complex and subtle a limit pricing game can be. Industrial organization textbooks churn out the equations that help find the optimal strategy. But how realistic is it to expect managers to follow the strategy? We saw in the previous chapter how easily people can be biased when interpreting information. What kinds of biases can we expect when managers interpret the behavior of rival firms? It is questions such as these that illustrate the important role that behavioral economics can play in industrial organization.

Two reasons to put some faith in managers are that they have the chance to learn from experience, and typically work in teams. With this in mind, let us look at a study by Cooper and Kagel (2005), in which subjects played a limit price entry game. Some subjects played as individuals and some in teams of two. The main

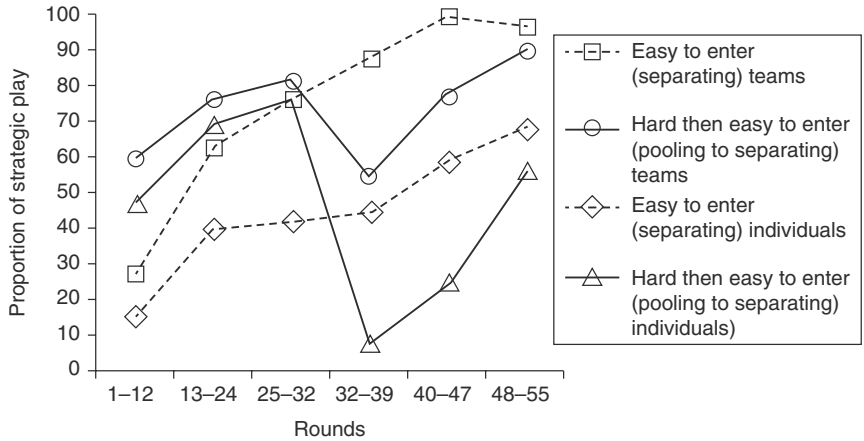

Figure 6.33 The proportion of strategic play in a limit pricing game. Each cycle lasted 12 or eight rounds.
Source: Cooper and Kagel (2005).

results are summarized in Figure 6.33. Focus first on the dashed lines. These tell us what happened when subjects repeatedly played a version of the game in which entering the market was easy. In this case, strategic play required the monopolist to build a large extension to signal low costs. In the first 12 rounds we see that few behaved strategically. With experience, the proportion of strategic play rose. Teams were also a lot more strategic than individuals. Experience and teamwork seem to do the trick!

Focus now on the solid lines. In the first 32 rounds subjects played a version of the game in which entering the market was hard. In this case, strategic behavior equated to a pooling equilibrium. We can see that there was a larger proportion of strategic play than when entering was easy. The really interesting thing, however, is what happens from round 33 onwards. At this point the game changed to one in which entering was easy. Strategic behavior thus equated to a separating equilibrium. The remarkable thing to note is that subjects adapted quickly to the new setting, and teams adapted particularly quickly. Indeed, teams in this 'crossover' treatment were much more strategic than teams in the comparator treatment were for similar levels of experience.

We have another example of how teams can outperform individuals. The interesting thing in this case is how experience in one setting translates into an improved performance in a different setting. Teams adapted to a change in the game more quickly than individuals did. Again, however, we should be careful to not get too carried away. It took a lot of experience before subjects were consistently behaving strategically. Moreover, there are other variants of the limit pricing game in which teams do not come close to strategic behavior even with experience.

The general lesson, therefore, is that firms will sometimes deviate from the optimal strategy. In this context, that means firms may not do all they could to deter potential entrants.

6.9.2 Market entry

In the limit pricing game an incumbent monopolist is pitted against a potential rival. Consider now a new market with no existing firms. To stick with the superstore example, you might think of a new town, or town that is rapidly expanding and is not currently served by any store. I am going to use the slightly different example of a new technology. Something called a Widget is invented that will revolutionize home entertainment. Imagine that you own a firm in the home entertainment business. Should you enter the Widget market?

Your first reaction might be to say: 'Of course I should enter, it is an exciting new opportunity and I should not miss out.' The problem with this line of reasoning is that other firms in the home entertainment business are probably thinking the same thing. And, if all firms enter the market, there will be oversupply, and someone is going to lose out. That someone might be you. Somehow, therefore, the 'right' number of firms need to enter the market. If too few enter, those in the market make bumper profits. If too many enter, some are going to make losses. We have a **market entry coordination game**. Firms need to coordinate on who enters and who stays out of the Widget market.

A market entry game is relatively complex, and it seems difficult to know how firms could coordinate on who enters the market. Subjects in the laboratory, however, seem good at solving this problem. To illustrate, I am going to look at a study by Rapoport, Seale and Winter (2002). This study is particularly interesting, because it considers firm asymmetry. For some firms it is less costly to start producing Widgets than for others. This asymmetry seemingly offers a simple solution to the coordination problem, namely those with the lowest cost to produce Widgets should be the ones to enter the market. This is not how things turned out!

First let us look at Figure 6.34, which illustrates how good subjects are at coordinating. In the games considered there were 20 firms in a group and the capacity of

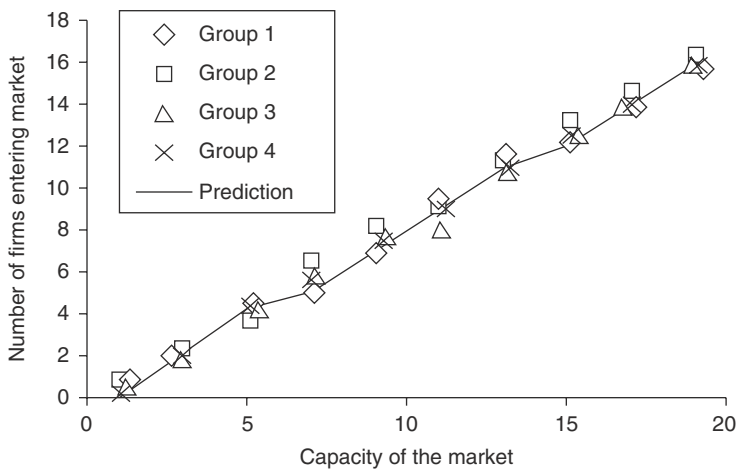

Figure 6.34 The predicted number of entrants in a market entry game compared to the average number of entrants in four groups.
Source: Rapoport, Seale and Winter (2002).

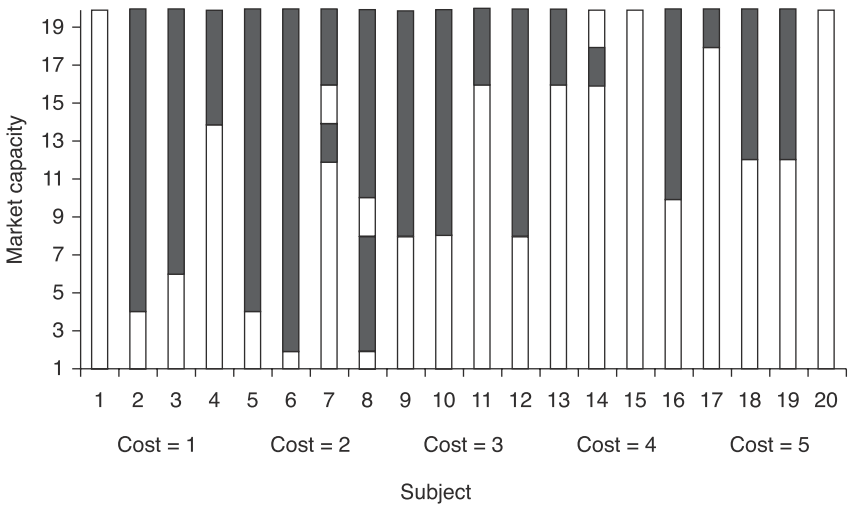

Figure 6.35 The behavior of subjects in a market entry game. The filled-in blocks show the subject entered, and the empty blocks that they stayed out.
Source: Rapoport, Seale and Winter (2002).

the market could be anything between one and 19. Over 100 rounds, subjects were exposed to all the possible capacities ten times. In Figure 6.34 we see the average number of firms entering the market compared to the equilibrium prediction. Every group is basically spot on the prediction! So, subjects learned to coordinate very well.

When we analyze how subjects learnt to coordinate, things get somewhat intriguing. Figure 6.35 shows what happened in one of the groups in the final rounds of the experiment. Subjects 1 to 4 had the lowest cost to enter and subjects 17 to 20 the highest cost. The filled-in blocks show for which market capacity a subject entered. Subject 2, for instance, stayed out of the market if capacity was one or three, but entered if capacity was five or more. Recall: we would predict that those with the lowest cost should be the ones to enter. That did not happen! Subject 1, for instance, always stayed out, while subject 18 entered if capacity was 13 or above.

Clearly, subjects found a way to coordinate other than through cost. Rapaport, Seale and Winter show that a model of reinforcement learning can explain how they do so. The key to the model is to equate strategy with the rule 'Enter the market if capacity is *x* or above'. With experience, a subject revises the value of *x*, meaning that, over time, subjects come to coordinate. This result reiterates the point that people can be good at coordinating, but not necessarily on the most efficient outcome. It also suggests that market entry may be based more on history and precedent than on costs. Firms with a history of entering new markets in the past may be the most likely to enter the market for Widgets, even if they are no good at producing Widgets.

On the topic of market entry, I will briefly mention a fascinating study by Goldfarb and Xiao (2011). They looked at entry in local US telephone markets after the 1996 Telecommunications Act. This Act opened up local telecommunications to competition and so created a swath of market entry games. Goldfarb and Xiao looked at how the strategic decision of firms related to the characteristics of the firm's manager.

They found that experienced and 'better educated' CEOs made better decisions. In particular, they entered markets with fewer competitors. Reassuringly, having an economics or business degree was a big component of being 'better educated'! Indeed, an economics or business degree was a substitute for experience.

Results such as these reiterate the basic point that firms are run by managers, and managers are inevitably going to be subject to all the cognitive biases and limited strategic reasoning that we have seen so far in this book. As mentioned in the previous section, the fact that decisions are made by teams of experienced managers is only going to go a little way to eliminating such bias and improving strategic reasoning. So, do not expect firms to be as strategic as they could be when entering the Widget or telecommunications market.

6.9.3 Quantity leadership

It would seem rude to discuss industrial organization and not mention Antoine Cournot, Heinrich von Stackelberg and Joseph Bertrand. These three names are etched in history for having given us the three classic models of firm competition. Conveniently ignoring the Bertrand model, I shall look at the Cournot and Stackelberg models. In both models, firms independently choose a level of output. To fix ideas, we can think of BuyFood and CheapFood deciding how big a superstore to have, or how many tins of tuna to order for resale. In the Cournot model, BuyFood and CheapFood choose their own quantity without knowing the quantity of the rival. In the Stackelberg model, a leader firm chooses its quantity, and, having observed this, the follower firm chooses its quantity.

Figure 6.36 illustrates how the models work. The two superstores choose to produce either a low, medium or high amount of output. The best for them, not for customers, is that they produce a low output and each get profit 72. This collusive outcome is unlikely, however, because each firm would have an incentive to increase output. The **Cournot equilibrium** is to produce medium output and each get profit 64. This is an equilibrium, because neither store can gain by producing more or less, given the output of its rival. Finally, we come to the two **Stackelberg equilibria**. If BuyFood produces a high output then CheapFood's optimal response is to produce low output. This is the best that BuyFood can do, and so is the Stackelberg equilibrium outcome. If CheapFood is the leader then we get the analogous equilibrium.

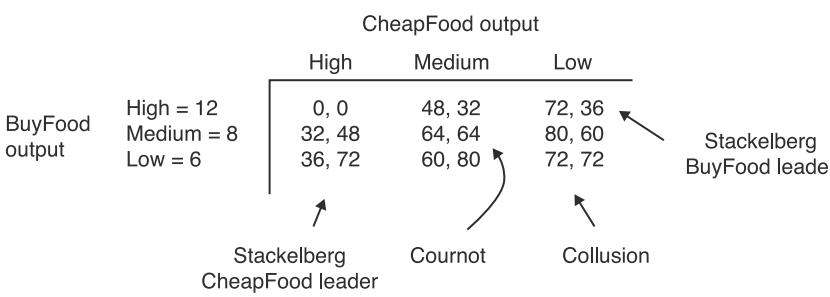

Figure 6.36 Profits in a quantity-setting game and classification of outcomes.

A study by Huck, Müller and Normann (2001) looked at whether the predictions of the Cournot and Stackelberg model are borne out in the laboratory. It also looked at the potential for collusion by comparing a setting in which subjects interacted with a random opponent each round with a setting in which subjects repeatedly interacted with the same opponent each round. Repeated interaction is much more conducive to collusion. In order to connect the results of the study with the numbers in Figure 6.36 we need to equate high output with a choice of 12, medium with eight and low with six. Subjects could choose any output between 3 and 15, but the equilibria were still as given in Figure 6.36.

The results of the study are summarized in Figure 6.37. Consider first the outcome when firms choose simultaneously, as in the Cournot model. The equilibrium prediction is that each firm produces output eight. This prediction is spot on. Output is slightly lower with repeated interaction, but not by much. This suggests little collusion. Look next at the outcome when firm 1 chooses before firm 2, as in the Stackelberg model. The equilibrium prediction is for firm 1 to produce 12 and firm 2 to produce six. This prediction is off the mark. Total output is near to that predicted, but firm 1 produces significantly less than predicted. I will come back to an explanation for this shortly. Again, output is slightly lower with repeated interaction, but not by much.

In the Cournot and Stackelberg models, 'fate' decides whether a firm chooses output with or without knowing the output of the other. In reality, firms have considerable control over when they decide on output. BuyFood, for example, may decide to delay extending its superstore until it sees what CheapFood does. This brings us on to models of **endogenous timing**. One of the many possible models of endogenous timing goes something like this. There are two periods, today and tomorrow. Today each firm can choose what output to produce or decide to wait and decide tomorrow. Tomorrow any firm that waited must choose output.

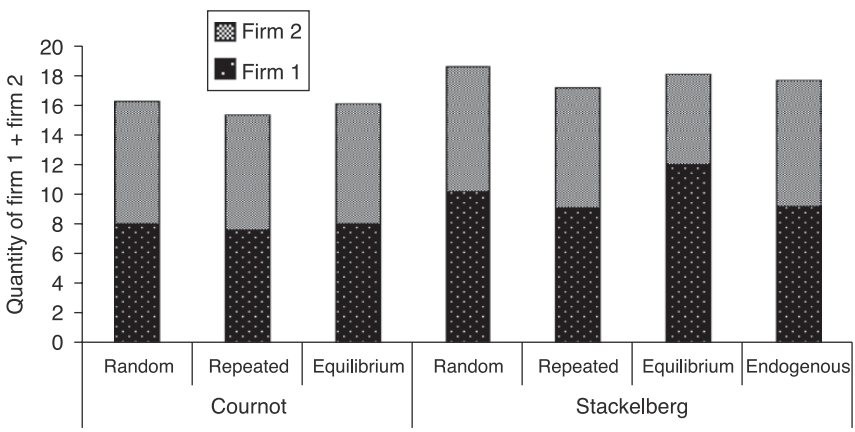

Figure 6.37 The output of firm 1 and firm 2 combined in Cournot and Stackelberg markets compared to the equilibrium.
Source: Huck, Müller and Normann (2001, 2002).

If both firms decide to wait until tomorrow then we should get the Cournot equilibrium outcome, because both firms are choosing simultaneously. But this means that both firms will not wait until tomorrow. If, for instance, BuyFood expects CheapFood to wait then it would do better to choose high output today and 'force' the Stackelberg equilibrium outcome. We can also rule out both firms choosing output today. If, for instance, BuyFood expects CheapFood to choose output today then it cannot possibly lose anything by waiting to see what output CheapFood will choose. The prediction, therefore, is that we should get the Stackelberg equilibrium outcome. One firm should lead and the other should follow.

Another study by Huck, Müller and Normann (2002) tests this prediction. And they do not find good support for it. One reason was that too few subjects waited. The prediction would be that 50 percent will wait, but only 39 percent did so. Another issue was that those choosing output today chose significantly less than the Stackelberg outcome. This finding fits the results of the earlier study. Indeed, the final bar of Figure 6.37 shows the average output of subjects that chose today (firm 1) and those that decided to wait (firm 2); endogenous timing seems to make little difference.

Both the studies we have looked at seem to suggest that the Stackelberg model is missing something. With both exogenous and endogenous timing we see systematic deviations from the predicted output. This is mainly because leaders chose less than they should. Why did they do that? Figure 6.38 shows the response of followers to the output of the leader. We can see that, if leaders chose a high output, followers responded by choosing medium to high output. They were supposed to choose low output. Leaders, therefore, were wise to not choose a high output! The question becomes why followers chose a larger output than predicted.

Look back to the numbers in Figure 6.36. If BuyFood is leader then the Stackelberg outcome gives CheapFood a lot less profit than BuyFood. By choosing medium instead of low output, CheapFood loses a little, but dramatically closes the

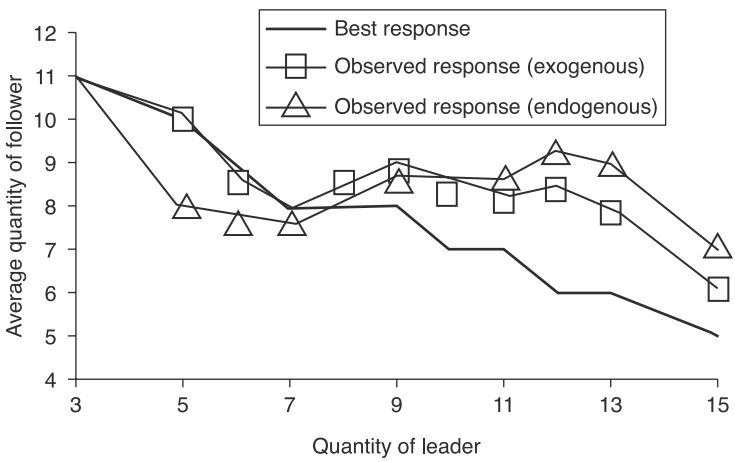

Figure 6.38 The best response and observed response of followers to the output of a Stackelberg leader.
Source: Huck, Müller and Normann (2002).

gap with BuyFood. If, therefore, CheapFood is willing to sacrifice some profit to make things more equal with BuyFood, it is no surprise that we see a systematic deviation from the Stackelberg equilibrium.

This brings us neatly onto the next chapter, on social preferences. In the chapter we shall discuss the overwhelming evidence that people do not like inequality. If this distaste for inequality extends to firm managers – and there is no reason why it should not – then we have good reason to question the predictions of the Stackelberg model. Firms that lead in the market may not be able to exploit their position quite as much as the model predicts.

6.10 Further reading

We have covered a lot of ground in this chapter, so there are plenty of sources for further reading. The book by Camerer (2003) is the most natural place to start. For a recent reinterpretation of behavior in the p-beauty contest, see Breitmoser (2012). Additional material on equilibrium selection and equilibrium refinements can be found in any good game theory textbook, of which there are many. For more on focal points, coordination games and strategic interaction more generally, the book by Schelling (1990) is definitely recommended. The further reading for Chapter 11 contains references to more on auctions, but I will mention the survey by Kagel and Levin (2016) here. Closely related to auctions are contests, and Dechenaux, Kovenock and Sheremeta (2015) provide a comprehensive review of the experimental literature. See Devetag and Ortmann (2007) for more on weak link games and Croson and Marks (2000) for more on threshold public good games. One issue I have ignored is that people rarely play the same game multiple times, and so more realistically have to learn from their own or others' experiences of playing 'similar' games. This raises questions about what is similar, and so on. Case-based decision theory (see Gilboa and Schmeidler 1995) and analogy-based expectation equilibrium (see Jehiel, 2005) are two potential ways to deal with this. For more on team decision-making, see Charness and Sutter (2012).

6.11 Online materials

On the companion website you will have access to flashcards, quizzes and answers to the review questions.

6.12 Review questions

6.1 When is the best time to sell an asset as a function of what other investors will do?

6.2 Suppose you are playing the p-beauty contest with a group of people who have only just been told about the game. Is what you have learnt in this chapter any help to you? What about in the simple beauty contest?

6.3 In Figure 6.2 we saw that the winning number in the 0.667-beauty contest fell over rounds even though no feedback was provided to subjects. Why might this be?

6.4 Does it make any difference if people can communicate with each other before playing a game?
6.5 Carefully explain the difference between the cognitive hierarchy model and k-level reasoning.
6.6 Does a sophisticated person behave consistent with the empirical optima (see Research methods 5.1)?
6.7 Does risk dominance or Pareto dominance seem more relevant to you in predicting what people will do? What about forward or backward induction?
6.8 Explain why sub-game perfection rules out non-credible threats. What about forward induction?
6.9 Look back at Table 6.8. I assumed an initial attractiveness of 30 for both strategies. What happens if the initial attractiveness of both is 10 or is 100? What has this got to do with initial beliefs?
6.10 Look back at Table 6.9. What would happen if $\rho = 0$ or $\rho = 1$?
6.11 I said that it does not always make sense to set the probability of choosing a strategy proportional to its attractiveness. Can you explain why?
6.12 How can we change the models of reinforcement learning and belief-based learning to take account of loss aversion?
6.13 What do you think are the consequences of the law of small numbers and confirmatory bias for learning in games?
6.14 Explain the difference between best reply learning and fictitious play.
6.15 Discuss why the strategy of teams may differ from that of individuals.
6.16 Given what we see in Figures 6.18 and 6.19, which auction format do you think will raise the most revenue?
6.17 Why do you think it matters if there is the possibility of a refund or rebate when contributing to a threshold public good?
6.18 How can leadership help people coordinate in the weakest link game and threshold public good game? You might want to think about forward induction.
6.19 Why might a person's position in a network influence the choices they make?
6.20 What factors might influence a firm's decision whether or not to enter a new market?
6.21 In a market with two firms in which firms sequentially choose quantity of output (i.e. the Stackelberg model), is there an advantage in being the first firm to choose quantity?

Chapter 7

Social preferences

> Charity collector: 'I want you to give me a pound, and then I go away and give it to the orphans.'
> Merchant banker: 'Yes.'
> Charity collector: 'Well, that's it.'
> Merchant banker: 'No, no, no! I don't follow this at all. I mean, I don't want to seem stupid, but it looks to me as though I'm a pound down on the whole deal.'
>
> Monty Python's merchant banker sketch

You may recall from Chapter 1 my saying that *Homo economicus* is both rational and selfish. Up until now we have primarily been questioning whether it makes sense to assume that people are as rational as *Homo economicus*. Now it's time to look at whether it makes sense to assume they are as selfish. That people behave as if selfish is one of the most basic assumptions in the standard economic model. The evidence suggests, however, that most people do care about others. Why else would millions of dollars be given to charity every year? Why else would someone feel envious of a work colleague earning a higher salary than them? If the utility of one person depends on the utility of another, then we say that there are **social** or **interdependent preferences**. In this chapter our task is to review the evidence for social preferences, look at how we can model them and look at what consequences they can have.

7.1 The experimental evidence for social preferences

There have been literally thousands of experiments suggesting that people have social preferences in one shape or another. In this section I am going to look at a selection of these experiments in order to illustrate the main things that we observe. At this point, I shall not try to explain why people behave as they do (that comes in sections 7.2 to 7.4), but you might want to be asking yourself this question when you look at the evidence. To try to help you through the maze of different games that I will look at, Table 7.1 gives a basic description of each game.

Table 7.1 A brief overview of the games I will look at.

Game	Stage 1	Stage 2
Dictator game	Proposer splits $10	—
Trust game	Investor can invest up to $10	Invested money is tripled and proposer can give back to the investor
Gift exchange game	Employer pays the worker a wage	Worker chooses how hard to work
Ultimatum game	Proposer splits $10	Receiver accepts or rejects
Dictator game with third-party punishment	Proposer splits $10	Third party can punish proposer
Linear public good game	Contribute to a public project	—
Linear public good game with punishment	Contribute to a public project	Can punish others after seeing contributions
Moonlighting game	Proposer invests or takes from responder	Responder gives or takes from proposer
Ultimatum game with responder competition	Proposer splits $10	Many receivers can try to accept
Ultimatum game with proposer competition	Many proposers say how they would split $10	Receiver can accept best offer

7.1.1 The nice side of social preferences

I shall begin with three games that show the nice side of social preferences. That is, they show people voluntarily giving money to others.

The dictator game is an obvious place to start. A **dictator game** involves two people: a proposer and a receiver. The **proposer** is given a sum of money, say $10, and asked how much of that $10 he would like to give to the **receiver**. There the game ends! The receiver gets what she is given and the proposer keeps the rest. You may well have played this game yourself in a restaurant you will never visit again. In that case you play the role of proposer, deciding what tip to leave the waiter before you leave the restaurant. If the proposer cares only about his own payoff then he should simply walk off with the $10 and give nothing to the receiver. What proposers actually do is quite different.

Figure 7.1 summarizes what happened in a series of comparable experiments by Forsythe and co-authors (1994), Hoffman and co-authors (1994) and Hoffman, McCabe and Smith (1996). Each treatment varied the way that the game was run (see Research methods 7.1 for more information), but we can ignore that here. The size of the bubble in the figure indicates what proportion of proposers offered a particular share to the responder, so, the larger the bubble, the more offered that amount.

By looking up from zero we can see that lots of proposers gave zero to the receiver; the proportion of subjects varies from around 20 percent in the standard experiments to 60 percent in the double-blind experiments. Clearly, however, not everyone gave zero, and lots of proposers gave money to the receiver. The amount they gave varies, but there were always at least some willing to give away half the money.

❖ RESEARCH METHODS 7.1 SINGLE- AND DOUBLE-BLIND: ANONYMITY VERSUS REALITY?

The dictator game is a **one-shot** game, in that it is played just once and anonymously. So, in principle, there is no point in a subject thinking about reputation or what others, including the researcher, might think about them. In reality, particularly when experiments are done on university campuses with students and faculty, can we expect subjects to behave as if the game is one-shot and anonymous? Or does it cross their mind that, '[i]f I take all the money the researcher will think that I am mean, and he will probably be marking my exam soon'?

In a series of experiments, Hoffman and co-authors (1994, 1996) varied the **social distance** between subjects and researcher to try to answer such questions. The results are summarized in Figure 7.1. I shall not explain all the nuances of the different treatments, but here are the basic differences.

In the standard and replication treatments, the instructions contain **sharing language**, such as 'has been provisionally allocated' $10 and 'the amount to

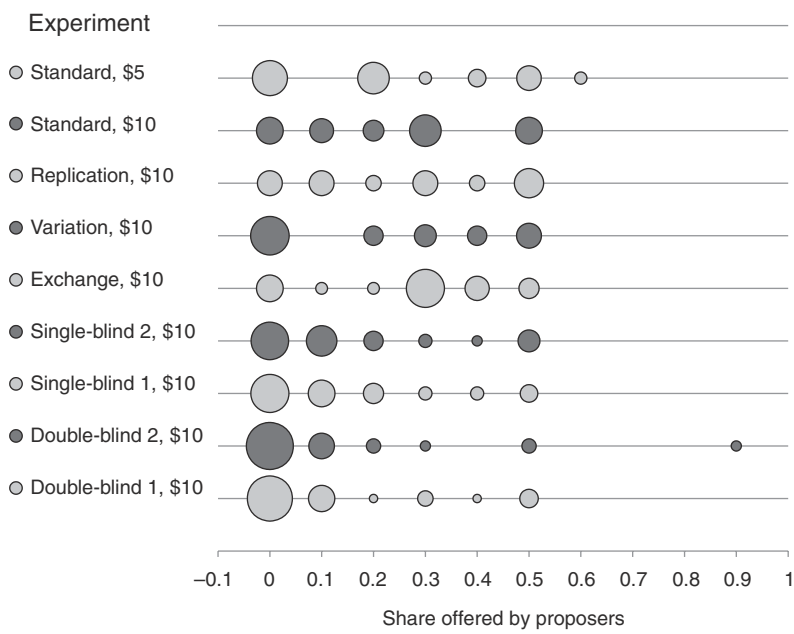

Figure 7.1 The amount that proposers gave in dictator experiments. The size of the bubble indicates the proportion of subjects offering each share. Many do give zero, but the majority give more than zero.
Sources: Forsythe *et al.* (1994), Hoffman *et al.* (1994), Hoffman, McCabe and Smith (1996) and Camerer (2003).

be divided'. In the variation treatment, the sharing language is dropped. In the exchange treatment, the game is framed as a seller (the proposer) setting a price that the buyer (receiver) must pay.

In the single-blind 2 treatment, once they have written down their choice, the proposers put any money they will give to the receiver into a sealed envelope, along with some blank pieces of paper (so that all envelopes have the same weight/size), and drop this envelope into a box. All proposers put their envelopes into the same box. Receivers randomly pick an envelope from the box. This guarantees anonymity between proposer and receiver and so is called **single-blind**. In the single-blind 1 treatment, a proposer no longer has to write down their choice but does let the researcher look inside the envelope before it is sealed.

In the double-blind 2 treatment, proposers put envelopes into the box and leave. Only then does the researcher look inside the envelopes and record how much was left. This guarantees anonymity between proposer and researcher, and so is called **double-blind**. In the double-blind 1 treatment the recording of offers is made by a randomly assigned student monitor (not the researcher) and two dummy envelopes containing no money are also put in the box. This means that a receiver who gets $0 cannot know whether there was a proposer who left $0 or they were just unlucky and got a dummy envelope. Anonymity between proposer and receiver and proposer and researcher is maximized.

The further we go down the line, the greater the social distance between the proposer and receiver, and the proposer and researcher. This should make the proposers more readily believe that the game is one-shot and anonymous. Figure 7.1 suggests that this is the case, given that we see less giving with more social distance. Note, however, that a significant proportion of proposers still give money, even in the double-blind 1 treatment. Care is also needed in creating too much social distance, because it could leave subjects doubting the experiment is real. For example, proposers might start doubting whether there really is a receiver waiting in another room to pick up the envelope!

The next game I want to look at is called the trust game. A **trust game** involves an investor and a proposer. Both investor and proposer are given $10. The **investor** is told that he can give as much of his $10 as he likes to the proposer. Any amount that he gives will be tripled in value before being given to the proposer. The proposer can then give as much money as she likes back to the investor.

As in the dictator game, a proposer who cares only about her own payoff should not give any money back to the investor. Investors, therefore, should not give money to proposers. This results in inefficiency, however, because it misses out on the tripling of money that is possible with investment. To illustrate, if the investor invests $0 then both investor and proposer get $10. If the investor invests all his $10 then the proposer would have $40 (the initial $10 plus the invested $30) and so could easily pay back the investor on his investment.

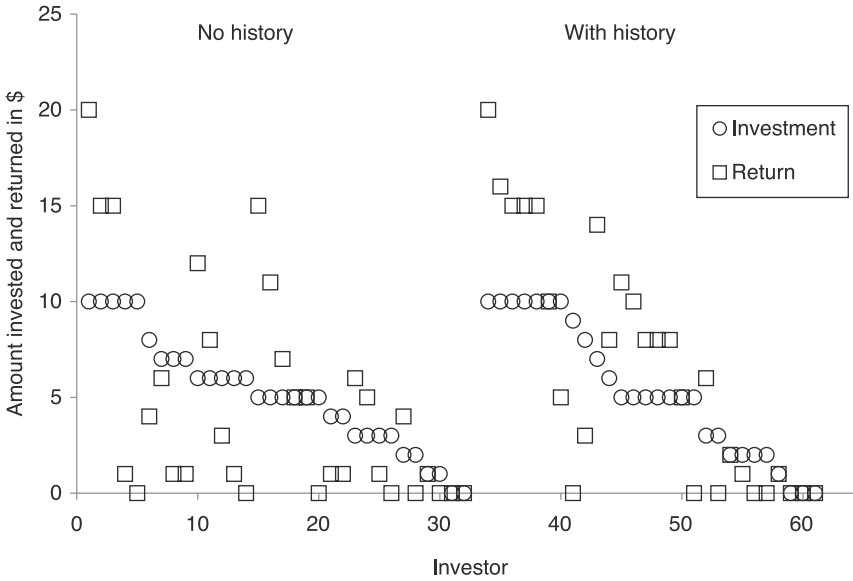

Figure 7.2 Investments and returns in a trust game with and without social history. Most investors invest something. Many proposers do not give anything back, but the majority do give something back.
Source: Berg, Dickhaut and McCabe (1994).

Figure 7.2 shows what subjects did in a study by Berg, Dickhaut and McCabe (1994). For now, let's ignore the distinction between 'No history' and 'With history'. The squares show the amount that different proposers returned to investors. As in the dictator game, we see that some proposers kept all the money for themselves, but many – the majority – gave at least some money back. The circles tell us how much investors gave to proposers. The majority of investors were willing to invest at least half their $10 with the proposer. This could be because the investors were happy to give money away to proposers. More realistically, however, it suggests that most investors were willing to trust that proposers would return the investment. An investor who gives to a proposer expects the proposer will give something back.

This brings into focus what is expected of others and whether others fulfill expectations. The distinction between 'No history' and 'With history' takes this further. In a 'No history' experiment, subjects simply got the instructions and played the game. In the 'With history' experiment, investors and proposers were told what happened in the earlier experiment (i.e. the 'No history' side). Things could have gone potentially two ways: investors could have been put off investing by seeing that some proposers kept all the money, or investors and proposers could have been influenced by the positive investing and returning. The latter effect seemed to come through in the experiment, with more money being returned. Proposers perhaps better understood that they were given money on the expectation that they would give a 'fair' share back.

The third game I want to look at is a gift exchange game. A **gift exchange game** is a slight variant on a trust game. There are two people, often thought of as worker and employer. The employer is given an amount of money, say $10, and chooses how

much of this to give to the worker, as a wage. Having observed her wage, the worker chooses how much effort to put into working. Higher effort costs the worker, but also means the employer receives higher revenues.

Cost and revenue are usually designed in such a way that higher worker effort is mutually beneficial – i.e. the extra effort would earn enough extra revenue that the employer could pay a high enough wage to offset the worker's effort cost. A worker motivated solely by her own monetary payoff would not put in effort, however, because it is too late to do anything about her wage, and so all it would do is lower her payoff. Given this, why should an employer pay a high wage? He should not.

To see what happens in a typical experiment, I will look at the results of a study by Gächter and Falk (2002). In this case, employers could choose a wage between 20 and 120 and workers choose effort between 0.1 and 1. Figure 7.3 plots wages and effort over ten rounds of play. In a one-shot experiment, workers and employers were randomly rematched each round, and so interaction was one-shot. In a repeated game experiment, the same workers and employers were matched in each round. As in the trust game, we see that, on average, employers were willing to give workers wages above the minimum, and workers reciprocated by choosing effort above the minimum.

To delve a little deeper, in the one-shot game Gächter and Falk classified a worker as a **reciprocator** if there was a positive correlation between the wage they were offered and the effort they chose to make, and an **egoist** if they chose the minimum effort most of the time. They found that 20 percent of workers were egoists and 53 percent were reciprocators. The majority of workers, therefore, reciprocated a high wage by higher effort.

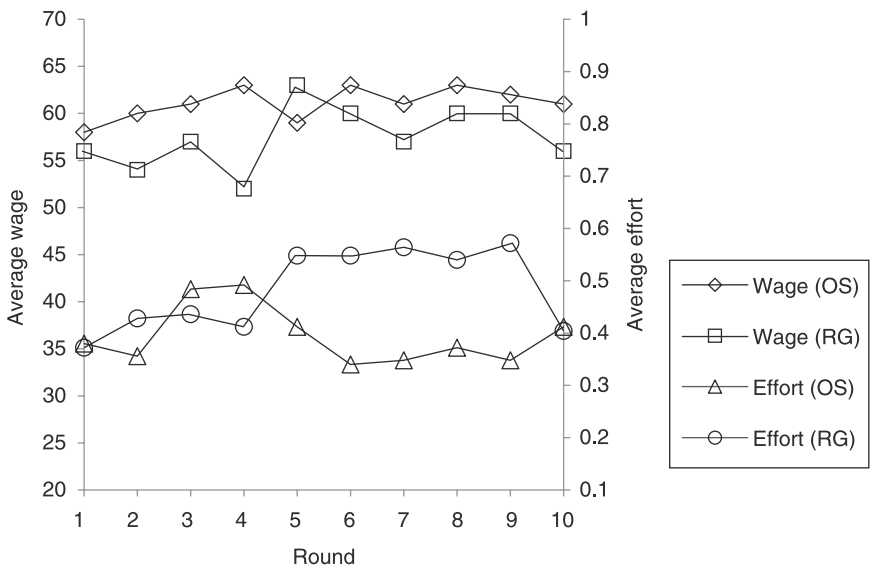

Figure 7.3 Wage and effort levels in a gift exchange experiment comparing one shot (OS) and repeated game (RG). The average wage and average effort are well above the minimum. Effort is higher in the repeated game, but falls in the last round consistent with imitation.
Source: Gächter and Falk (2002).

In the repeated game it becomes slightly harder to judge reciprocation, because a worker may want to build a reputation as someone who reciprocates a high wage. This way she can secure a high wage in the future. We could call such a worker an **imitator**, because they imitate a reciprocator. How can we tell a reciprocator from an imitator? We look at the last round. In this round an imitator no longer has any need to reciprocate a high wage, so she should choose minimum effort. A genuine reciprocator would still reciprocate a high wage by high effort. Gächter and Falk found that 20 percent of workers were imitators, 21 percent were egoists and 48 percent were reciprocators. What we learn new, therefore, is that some people realize that it may be in their interests to look like someone who gives or reciprocates.

To summarize what we have seen so far: many people are willing to give to others at a cost to themselves, expect others will reciprocate giving and realize it may be in their interests to look like someone who gives. But don't overlook the fact that many do not give at all.

7.1.2 The nasty side of social preferences

Now it is time to look at a more nasty side of social preferences: people taking money from others.

I will start by looking at the ultimatum game. The **ultimatum game** is a twist on the dictator game. As in the dictator game, a proposer is given a sum of money, say $10, and asked how much of that $10 he would like to give to the receiver. The twist is that the receiver must decide to either accept the offer, in which case the receiver gets what she is given and the proposer keeps the rest, or reject the offer, in which case both get zero.

If a receiver cares only about her own payoff, she should accept any positive offer, because if she rejects the offer she gets zero. To illustrate what typically does happen, Figure 7.4 gives the proportion of subjects who would have rejected offers in

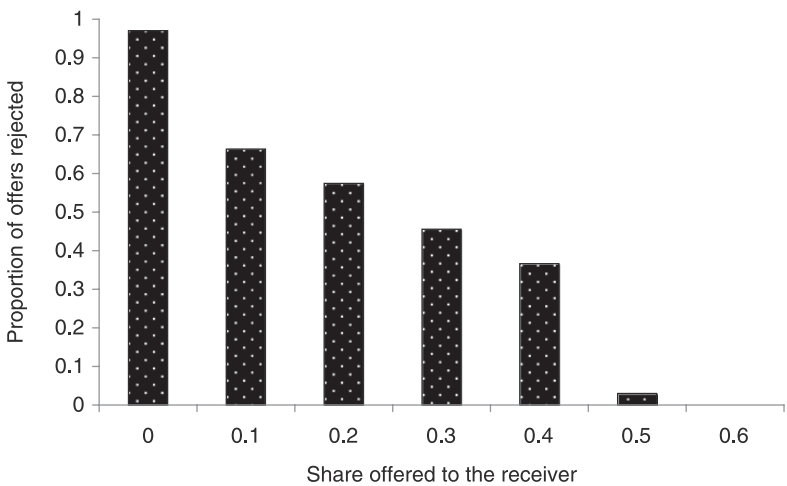

Figure 7.4 The proportion of offers rejected in an ultimatum game. Offers of a 0.5 share or better are rarely rejected, but offers of a less than 0.5 share are often rejected.
Source: Larrick and Blount (1997).

experiments by Larrick and Blount (1997). (See Research methods 7.2 for more on why I chose this study.) Clearly, a large proportion of people were willing to reject offers less than a 50:50 split. The evidence, therefore, is that some people are willing to sacrifice their own monetary payoff to decrease that of others.

❖ RESEARCH METHODS 7.2 FINDING OUT WHAT YOU ACTUALLY WANT TO KNOW: THE STRATEGY METHOD

In the ultimatum game, what we are particularly interested in is whether receivers reject offers. This presents a problem, however: if all the proposers propose a 50:50 split, we can never know whether or not the receivers would have rejected a low offer. One way around this is the **strategy method**, whereby responders are asked what they would do in any possible contingency. So, receivers have to say what they would do if they were offered $0, $1, $2, and so on. Only after this do they learn how much they have actually been offered. This way we can learn if a low offer would be rejected by the receiver.

Despite this big advantage, the strategy method is rarely used. I took the results for Figure 7.3 from one study that did use this method. More typical is the **game method**, in which receivers get to see the actual offer made by proposers before deciding what to do.

In principle, the strategy method and game method frame the problem in two different ways, so we might expect different behavior. It could also be argued that the more typical framing in real life is the game method. We get to see the waiter's service, for example, before we decide the tip. The evidence, however, suggests that there is relatively little difference in conclusion depending on what method is used. It is also not uncommon to see framing in real life similar to the strategy method. For example, a wage contract might specify what will happen for a variety of different possible effort levels.

The main lesson to be learnt from the ultimatum game is the rejection of positive offers. For completeness, however, we should have a look at what proposers offer. This will not take long, though, because the evidence is relatively easy to summarize: most proposers offer a 50:50 split. Figure 7.5 illustrates some representative experiments. In Chapter 9 we shall have more to say about ultimatum game giving across different cultures.

Another variant on the dictator game is the **dictator game with third-party punishment**. In this game, two people play the dictator game while being observed by a third person. This third person has no monetary interest in the game, but having observed the game has the option to reduce the payoff of the proposer. Specifically, for every $1 the third person gives up, she can lower the payoff of the proposer by $3. If she cares only about her own payoff, there is no reason for the third person to punish. But many do.

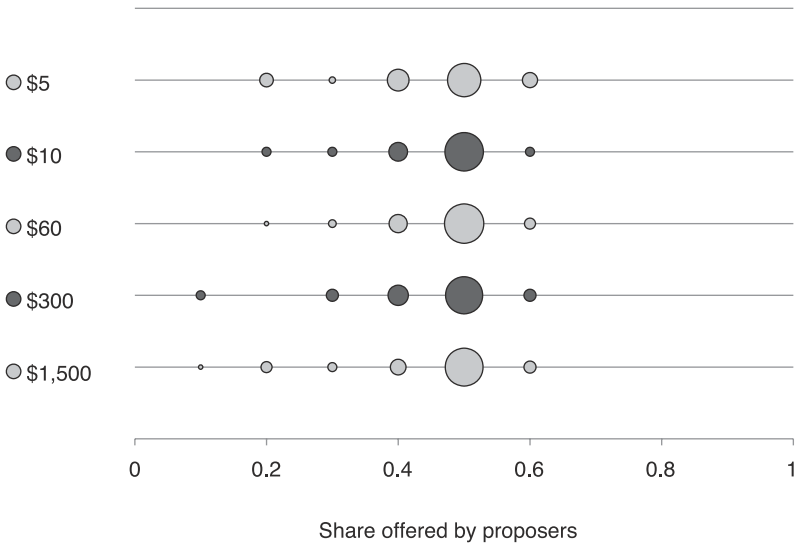

Figure 7.5 The share that proposers gave in ultimatum experiments when the amount of money given to the proposer ranged from $5 to the equivalent of $1,500.
Sources: Forsythe *et al.* (1994) and Slonim and Roth (1998).

❖ RESEARCH METHODS 7.3 DOES STAKE SIZE MATTER?

A guiding rule of economic experiments is that subjects should be paid according to their decisions. Otherwise, we cannot be sure that subjects take decisions seriously. But how much do subjects need to be paid to take things seriously? The evidence seems to be that subjects do need to be paid something (see Research methods 3.7), but paying them more does not necessarily make much difference. Or, at least, if it does make a difference it is not in a predictable direction. Small stake sizes should, therefore, be enough.

To illustrate, Figure 7.5 shows the results from experiments by Slonim and Roth (1998) in Slovakia, in which the money at stake was varied between $60 and $1,500. These can be compared with standard experiments by Forsythe and co-authors (1994) using a stake of $5 and $10. It is clear that stake size had no effect in these experiments for the amount that proposers gave. There was evidence that it had an effect on receivers, with them being less likely to reject the higher the stake. The effect was arguably still not large, however, with the rejection rate falling from 17.1 percent to 12.1 percent to 8.8 percent as the stake size went from $60 to $300 to $1,500.

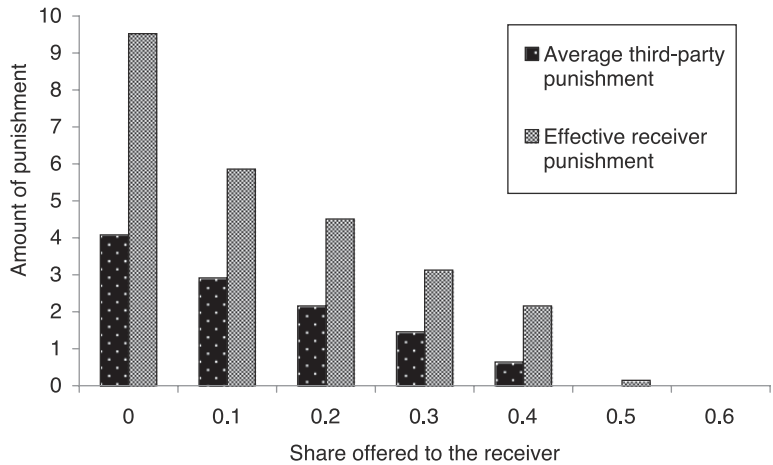

Figure 7.6 Punishment in a dictator game with third-party punishment compared to in the ultimatum game. Third-party observers punish low giving, but not as much as responders do.
Sources: Fehr and Fischbacher (2004) and Larrick and Blount (1997).

Figure 7.6 illustrates what happened in a study by Fehr and Fischbacher (2004). Focus first on the dark shaded columns. These show the amount that proposers were punished as a function of the share offered. We see that low offers are punished by the third-party observer. In fact, roughly 60 percent of observers were willing to punish low offers, and, the lower the offer, the more they punished. Many people, therefore, were willing to sacrifice their own payoff to lower that of someone they had observed giving too little to someone else. This is a step up from the ultimatum game, in which people punished people giving little to them. For comparison, the lighter shaded columns in Figure 7.6 show the effective punishment in the ultimatum game as seen in Figure 7.4. Third-party observers punish less than responders, but still a relatively large amount.

7.1.3 Reciprocity

Having seen something of the nice and nasty sides of social preferences, I am now going to consider two games in which both sides seem to combine: people voluntarily giving money to some and taking from others.

In a **linear public good game** there is a group of n people, typically four in experiments, who are each given some money, say \$20, and asked how much of this \$20 they wish to contribute to a group project. Any money contributed to the group project is multiplied by some multiplier M and shared equally amongst the people in the group. The key to the game is that $1 < M < n$, and so it is good for the group that a person contributes, but not in the material interests of the person to contribute.

For example, if $M = 1.6$ and someone such as Anna contributes \$1, each person in the group gets \$1.6/4 = \$0.4 and the overall return is \$1.6. This is clearly better than the \$1 contributed. So, the best for the group, or **Pareto-efficient** outcome,

Table 7.2 Some possible payoffs in a four-person linear public good game with endowment \$20 and multiplier $M = 1.6$. The group payoff is maximized if all people contribute \$20. Person 1's payoff is highest if she free-rides by contributing \$0 while all others contribute \$20.

Scenario	Contributions	Payoff of person				Total payoffs
		1	2	3	4	
A	\$20, \$20, \$20, \$20	\$32	\$32	\$32	\$32	\$128
B	\$0, \$0, \$0, \$0	\$10	\$10	\$10	\$10	\$40
C	\$0, \$20, \$20, \$20	\$44	\$24	\$24	\$24	\$116
D	\$10, \$10, \$10, \$10	\$26	\$26	\$26	\$26	\$104
E	\$0, \$10, \$10, \$20	\$36	\$26	\$26	\$16	\$104

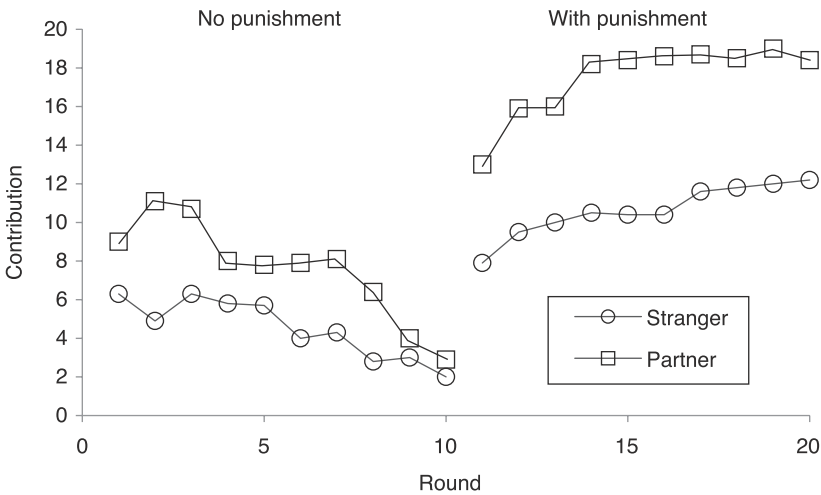

Figure 7.7 Contributions in linear public good games with and without punishment and with stranger and partner matching. Contributions fall over time without punishment and rise with punishment.
Source: Fehr and Gächter (2000a).

is that every person contributes everything they have. For each \$1 that Anna contributes, however, she gets back only \$0.4, and so it is not in her material interests to contribute. A person interested only in her own material payoff would, therefore, not contribute to the public project, and the **Nash equilibrium** is for no one to contribute anything. Table 7.2 illustrates this trade-off.

As a slight aside: in general, we say that there is a **social dilemma** whenever there is a difference between the Pareto-efficient outcome, which captures what is best for the group, and the Nash equilibrium, which captures individual incentives. Linear public good games, trust games and gift exchange games are social dilemmas.

Figure 7.7 shows what happened in a study by Fehr and Gächter (2000a). For now, focus on the left-hand side of the figure – i.e. the first ten rounds. In a partner

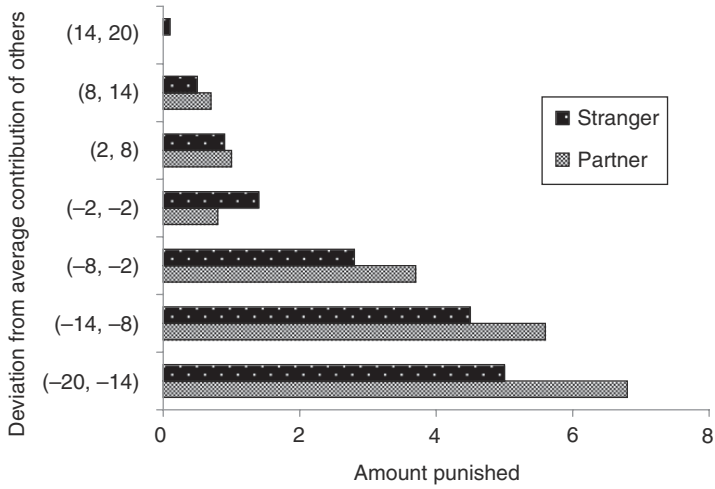

Figure 7.8 The amount of punishment depends on how much the subject contributed relative to the average contribution in the group. Those contributing less than the average are punished more.
Source: Fehr and Gächter (2000a).

treatment, the same four people played the game ten times. In a stranger treatment, each person was matched with three randomly chosen people each round. In both experiments, we see that subjects contributed in the first round but contributions fell over time. Given what we saw in the dictator game, it comes as no surprise that in the first period some subjects contributed zero and some contributed a positive amount. What is new here is the dynamic effect of falling contributions over time. This suggests that giving to others is not done unconditionally. In particular, a subject who contributes, but sees that others did not, has a tendency to lower their subsequent contributions. Maybe they do not want to give to the **free riders** who contribute zero.

Next consider a variant of the game called the **linear public good game with punishment**. The linear public good game is played, as explained above, but then a second stage occurs. In this second stage each person can pay an amount t to punish another person by lowering that person's monetary payoff by \$1. For example, if t = \$0.33 then, by paying \$1, Anna could lower the payoff of three other people by \$1 each or lower the payoff of one person by \$3. Given that punishment is costly and cannot possibly change the contributions that have already been made, no person, who maximizes their own material payoff, should ever punish. Punishment does not, therefore, change the Nash equilibrium or the fact that there is a social dilemma.

Punishment does change what people typically do, however, as we can see on the right-hand side of Figure 7.7. With punishment, there are high and increasing investments over time. Punishment makes a difference because subjects punish those who do not contribute. Figure 7.8 illustrates by plotting the amount of punishment received by a subject depending on how much less or more they contributed relative to the average contribution of others in the group. We clearly see that those who

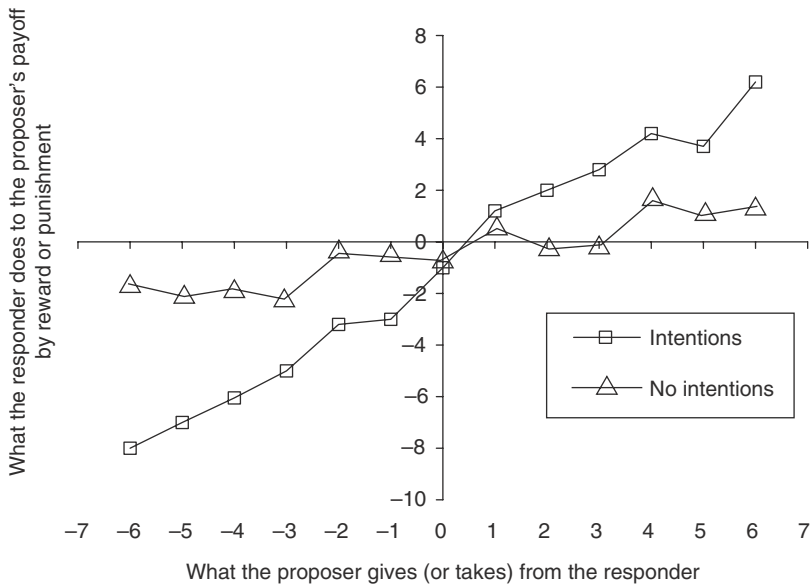

Figure 7.9 Responders' impact on proposers in the moonlighting game. If the proposer gives to the responder, the responder rewards the proposer. If the proposer takes from the responder, the responder punishes the proposer. If the computer randomly decides, the responder neither rewards nor punishes the proposer.
Source: Falk, Fehr and Fischbacher (2008).

contribute less are punished more. This is incentive enough for those who might want to free-ride to contribute.

In the linear public good game with punishment, we see people willing to give money to others by contributing, but also taking money from some by punishing. Giving and taking is thus conditional on others' behavior. A game that makes this more explicit is the moonlighting game.

The **moonlighting game** combines various elements of the games considered so far. There are two people: a proposer and responder. Both are given, say, $12. As in a trust game, the proposer can give up to $6 to the responder, and any money given will be tripled in value. The proposer can, however, also take up to $6 from the responder. Having seen what the proposer has done, the responder can reward the proposer by giving her some of her money. Or, she can punish the proposer by paying to reduce the money of the proposer. This game is the most flexible that we have considered so far, in that it allows giving, taking, punishing and rewarding.

The main lesson that we can learn from this experiment is summarized in Figure 7.9, with results from a study by Falk, Fehr and Fischbacher (2008). Of most interest to us, at this stage, is the intention treatment. Here, as in the trust game, we see that, if a proposer gives money, the responder reciprocates by also giving. As in the ultimatum and public good games, we see that, if a proposer offers a low share (takes money), the responder reciprocates by punishment. In a no-intentions experiment,

the choice of the proposer is made randomly by a computer, and we see much less evidence of reciprocation. We will come back to this in section 7.3.

7.1.4 Fairness and competition

I want to finish with some experiments showing that, in the face of competition, giving can change. To do so, I will revisit the ultimatum game. In an **ultimatum game with responder competition** there is one proposer and many receivers. Once a proposer has made an offer, each receiver says whether she would accept the offer. One of those who said she would accept the offer is randomly chosen to actually get the offer. In an **ultimatum game with proposer competition** there are many proposers and one receiver. All proposers make an offer, and the receiver can accept or reject the best offer made.

Competition would suggest that offers should be higher with proposer competition and lower with responder competition. Figure 7.10 shows that this was the case in experiments run by Roth and co-authors (1991) and Grosskopf (2003). This figure can be compared with Figure 7.5 to see that competition makes a real difference. Interestingly, however, competition does not eradicate giving. For example, with responder competition, offers are fewer than in the ultimatum game but still well

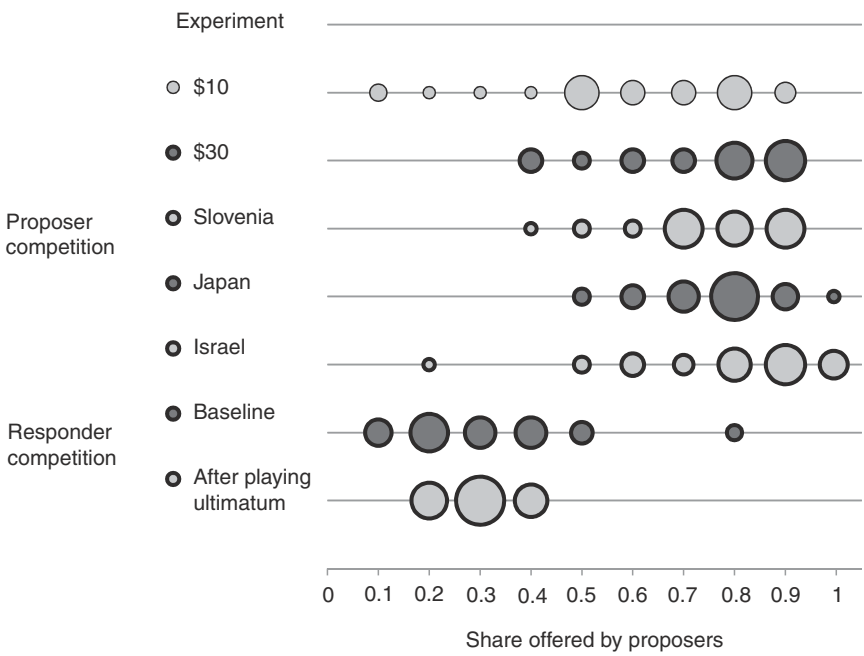

Figure 7.10 The share that proposers offer in ultimatum games with competition. Roth and co-authors find that the effect of proposer competition is the same across four countries, and Grosskopf finds that responder competition lowers offers even if subjects had played the standard ultimatum game beforehand.
Sources: Roth *et al.* (1991) and Grosskopf (2003).

above zero. It would seem, therefore, that proposers still wanted to give something, but gave less than in the ultimatum game.

7.1.5 The terminology of reciprocity

I have talked through a lot of different games and results, so it seems apt to summarize what we have seen. In doing so, I will introduce some useful terminology. We saw the following.

- Many people are willing to sacrifice their own monetary payoff to increase that of others (dictator game, trust game, public good game).
- Many people reciprocate the kind action of another to them by kindness of their own (trust game, gift exchange game, moonlighting game). This is **positive reciprocity**.
- Many people reciprocate the unkind action of another to them with punishment (linear public good game, moonlighting game). This is **negative reciprocity**.
- Many people show both positive and negative reciprocity (moonlighting game, linear public good game). This is **strong reciprocity**.
- Many people reciprocate the kind or unkind action of another to someone other than them (dictator game with third-party punishment). This is **indirect reciprocity**.
- There is considerable **heterogeneity** in desires for giving and reciprocity (all experiments), with many giving zero if they have the chance.

The clear picture we get is one of reciprocity, not of altruism or unconditional giving. An important question is whether such behavior is consistent with utility maximization. We know that it is inconsistent with utility maximization if a person cares only about her own payoff. But this does not mean it is inconsistent with utility maximization. Andreoni and Miller (2002) had subjects play a series of modified dictator games and asked whether choices were consistent with utility maximization (see Review question 7.3 for more details). They were, for all bar three of the 176 subjects.

This is important, because it shows that interdependent preferences, reciprocity, and the like need not be inconsistent with the standard economic story of utility maximization. Our task now, therefore, is to find a utility function that can capture the type of behavior described above. In doing this, it is useful to distinguish material payoff from utility. You should think of **material payoff** as the utility from money or goods that goes directly to the person; so far in the book that is what we have focused on. What we want to allow now is that utility depends on some social component as well as material payoff, giving us social preferences.

7.1.6 Social preferences and teams

Before we look at modeling social preferences it is interesting to make a brief diversion to look at giving by teams. Recall (see section 6.4) that teams are typically more like *Homo economicus* than individuals. Our focus in Chapter 6, however, was on strategic interaction. What I want to look at now is whether teams are less likely to give, punish, trust or reciprocate than individuals.

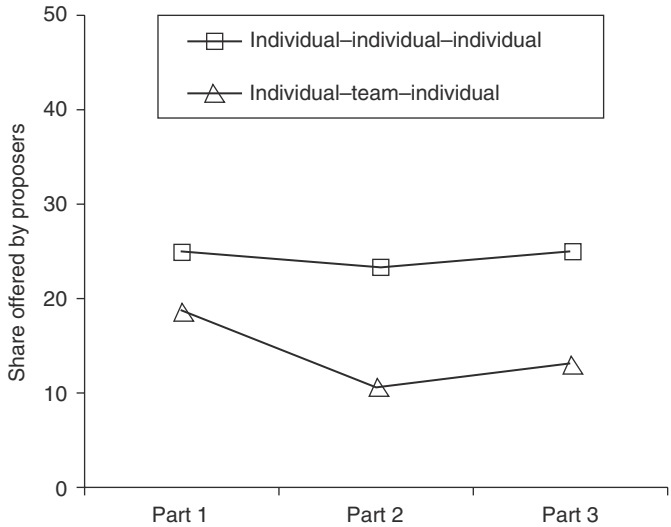

Figure 7.11 The share offered by proposers in the dictator game. In parts 1 and 3, all subjects acted as individuals. In part 2, some subjects acted as individuals and some as a team of three.
Source: Luhan *et al.* (2009).

Consider first a study by Luhan, Kocher and Sutter (2009) on the dictator game. Subjects played the game as individuals, then in teams of three, and again as individuals (see Research methods 7.4). This can be compared to a control treatment in which subjects played as individuals three times. Figure 7.11 summarizes the results. We see that the share offered by teams was less than that offered by individuals. Moreover, the experience of being in a team lowered the offer that many subjects made when they subsequently played as an individual. Even so, the differences are not huge.

❖ RESEARCH METHODS 7.4 LETTING SUBJECTS CHAT

One possibility is to let team members interact face to face when reaching a decision (see Research methods 2.6). This slightly diminishes anonymity, however, and when it comes to measuring social preferences this may not be a good thing. Another possibility is to let team members chat via computers. This retains anonymity, and was the method used by Luhan, Kocher and Sutter. The chat can subsequently be analyzed to give some insight into how teams reached a decision. For example, fairness was mentioned by 17 of the 30 teams. Perhaps surprisingly, 11 teams used fairness to justify a less generous transfer! Here is an excerpt from one piece of chat:

Team member 1: Obvious decision, don't you think?
Team member 3: Sure; all for ourselves.
Team member 1: I am no Good Samaritan.

Team member 3: Transfer 0.
Team member 3: Right, neither am I.
Team member 1: We are not responsible for [the receiver] are we?
Team member 1: Therefore 0.
Team member 1: OK. I'll enter it.
Team member 3: No. 2, do you share our opinion?
Team member 2: I think we should be fair.
Team member 3: I don't. Nobody has ever been fair to me.
Team member 1: 0 is fair.
Team member 2: Wait a second and don't enter anything.

After a long chat they shared 20 percent with the receiver.

Consider next studies by Bornstein and Yaniv (1998) on the ultimatum game and Kugler and co-authors (2007) on the trust game. In both studies direct comparison was made between individuals and teams of three. Figure 7.12 summarizes the results. We see that teams offered less in the ultimatum game, and invested and returned less in the trust game than individuals. Again, however, the differences are not huge.

These findings confirm and extend to social preferences the general picture that I painted in Chapter 6. Teams are more consistent with the standard economic model than are individuals. The behavior of teams is also, however, a lot closer to that of individuals than it is to that of *Homo economicus*. Teams still give, take and reciprocate. This reiterates the need for an appropriate way to capture social preferences.

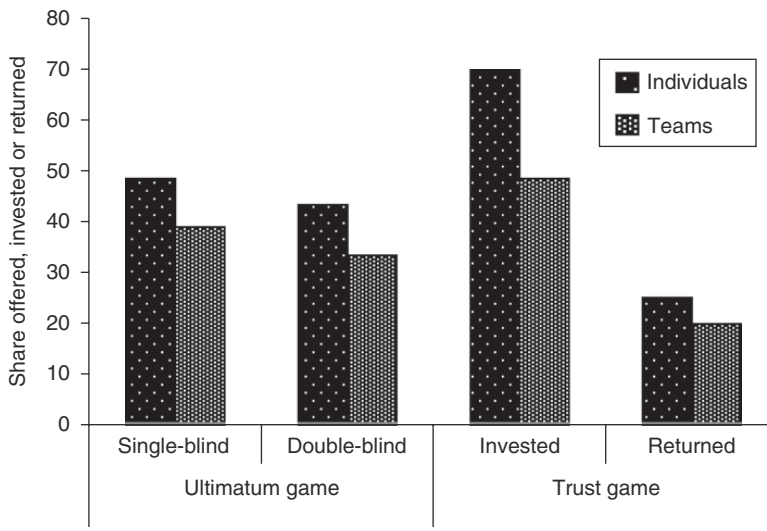

Figure 7.12 The share proposed in the ultimatum game, invested in the trust game and returned in the trust game comparing individuals and teams of three.
Sources: Bornstein and Yaniv (1998) and Kugler *et al.* (2007).

7.2 Inequality aversion

Reference dependence was a recurring theme in Chapters 2, 3 and 4. We saw that the perceptions of an outcome are often determined by whether it was above or below some reference level. We also saw that the reference level can be determined by a multitude of factors, such as past experience, an expectation of what would happen, and so on. A very natural reference point, which we only briefly touched upon in section 2.6, is what others are getting. For instance, if Alan observes his co-workers getting paid $30 an hour, he might expect to get paid $30 an hour. Similarly, if Anna knows her friend bought a box of Nutty cereal for $4, she might also expect to pay $4.

If the reference point is what we observe happening to others, doing better than others is a gain, while doing worse than others is a loss. Making the argument that people are averse to getting less than others seems easy enough. Making the argument that people positively gain from getting more than others, particularly given the evidence we have seen in section 7.1, is slightly harder to make. We are going to look, therefore, at something called inequality aversion. If a person is **inequality-averse** then she gets disutility from earning less than others and, possibly, gets disutility from earning more than others. This means she may be willing to sacrifice some of her material payoff to obtain a more equitable outcome with others.

What are the consequences of inequality aversion? We shall see that inequality aversion can do a good job of explaining the experimental data summarized in section 7.1. For example, if Alan is averse to earning more than others then he would be willing to propose a positive amount in a dictator or ultimatum game. Similarly, if Anna is averse to earning less than others she would be willing to reject a low offer in the ultimatum game and punish a free rider in a public good game. To formalize this, we need a good model of inequality aversion.

I am going to look at two models. The principal way in which they differ is in the information they assume a person has. The first model I will look at is applicable to situations with **incomplete information**, when a person does not know the actions of others. The second model can be applied only in situations with **complete information**, when a person knows the actions of all others. To explain this distinction, we can look back at Table 7.2 and contrast scenarios D and E. In a game of incomplete information, person 2 gets to know only that her payoff is $26. She can infer that total contributions to the public good must have been 40 and so the average payoff is $26, but she cannot know whether person 1 was a free rider or not. In a game of complete information she gets to know individual contributions, and so can tell whether person 1 contributed $10 or $0.

Some settings may naturally be ones of incomplete information and some of complete information. Intuitively, this is going to make a difference, because person 2 might think differently about scenarios D and E. That means we need a model for both possibilities. It is worth pointing out, however, that a model applicable to situations of incomplete information can also be applied to situations of complete information. The converse need not be true.

7.2.1 Inequality aversion with incomplete information

The **equity, reciprocity and competition** (ERC) model was proposed by Bolton and Ockenfels (2000) and can capture inequality aversion in settings with incomplete

(or complete) information. I shall explain here a simple, special case of the more general model that they propose.

Imagine a group of n people who have played a game and earned material payoffs of $u_1, u_2, \ldots, u_n$. I will focus on one of the people, called Miguel. In calculating the utility of Miguel, the first thing we do is work out his **relative share** of the total payoff using equation

$$s_M = \frac{u_M}{\sum_j u_j}.$$

The utility of Miguel is then assumed to be

$$U_M(u_1, \ldots, u_n) = u_M - \theta_M\left(s_M - \frac{1}{n}\right)^2,$$

where θ_M measures the strength of his inequality aversion.

If Miguel does not care about inequality aversion then $\theta_M = 0$ and his utility is simply his material payoff u_M. If he does care about inequality aversion then he cares about how close his payoff is to the **social reference point** of an equal share. He dislikes earning more or less than the average material payoff of others. The squared term means that, the further he is away from the social reference point, the more he dislikes earning more or less than the average. So, he does not mind being a little away from the average but dislikes a lot being a long way from the average.

To better illustrate the model, we can work through the experimental games of section 7.1 and show that the model does a good job of fitting what we observe. To simplify matters a little I will equate material payoff with money received, but you need to think about whether this is a reasonable thing to do or not.

We can start with the **dictator game**. If the proposer, Miguel, splits the \$10 so that he keeps \$$p$ then his material payoff is p and his utility is

$$U_M = p - \theta_M\left(\frac{p}{10} - \frac{1}{2}\right)^2.$$

The social reference point in this case is \$5, and inequality aversion means he gets disutility from keeping more or less than \$5. To be a bit more specific, we can maximize utility with respect to p and get equation

$$1 = \frac{2\theta_M}{10}\left(\frac{p}{10} - \frac{1}{2}\right),$$

which can be rearranged to

$$p = min\left\{5 + \frac{50}{\theta_M}, 10\right\}.$$

He will, therefore, keep at least \$5, and, if $\theta_M \leq 10$, will keep \$10. If $\theta_M > 10$ then he will give some money to the receiver so as to not be earning relatively too much. For example, if $\theta_M = 25$ then he will keep \$7 and give \$3 to the receiver, consistent with what we see in Figure 7.1.

Extending this to the **ultimatum game** is relatively simple, because the proposers have the same utility function as in the dictator game. We therefore need only to think about what responders will do, and all they can do is accept or reject. Suppose that Anna is the responder. If she accepts, her utility will be

$$U_A = 10 - p - \theta_A\left(\frac{10-p}{10} - \frac{1}{2}\right)^2,$$

and, if she rejects, it will be $U_A = 0$. So, she will accept if

$$10 - p \geq \theta_A\left(\frac{10-p}{10} - \frac{1}{2}\right)^2,$$

which we can rearrange to

$$\frac{100(10-p)}{(5-p)^2} \geq \theta_A.$$

If $\theta_A = 0$, she will accept any positive offer, but, the higher is θ_A, the higher needs to be the offer before she will accept. For example, if p = \$9.50, and so she is offered \$0.50, then she will reject if $\theta_A > 2.5$. This is because the material gain of \$0.50 is not enough to compensate for the disutility of earning less than Miguel. If she is offered the social reference point of \$5 then she will accept for sure.

I will finish looking at the ERC model with the **linear public good game**. Consider first the case without punishment. Going through the equations is a bit tedious. We can argue quite easily, however, that the unique equilibrium will be one in which no person contributes. To see why, suppose that all people are contributing the same amount, say c^*. What would happen if someone such as Miguel contributes a tiny amount less than c^*? He gains in monetary payoff, because he has lowered his contribution. He is now earning more than the average payoff, because he is contributing less than others. Recall, however, that there is hardly any utility loss from earning a little bit more than the average. So, provided he contributes only a little bit less than others, his utility will increase by contributing less.

You can see this in Table 7.3 by comparing scenarios A and F. If Miguel lowers his contribution to \$19, he gains \$0.60 in material payoff and loses $0.75^2\theta_M = 0.56\theta_M$ because of inequality aversion. This is a good deal if $\theta_M \leq 1$. I will leave it to you to work out how high θ_M can be for him to gain from lowering his contribution by \$0.50 or \$0.01. You should find that he will gain from lowering his contribution a bit. If everyone thinks like this then contributions will drift downwards, either for real or in the minds of people, until no one is contributing.

What if we add punishment? Punishment requires complete information about the distribution of payoffs, and so we can no longer assume incomplete information. But the model can still clearly be applied. Doing so, we find that it can be in everyone's

Table 7.3 The ERC model and linear public good game. If he will not be punished, Miguel gains by contributing $19 (scenario F) rather than $20 (scenario A). If he will be punished (scenario G), he does best to contribute $20. Anna does better to punish Miguel for contributing $19 (scenario G) if she sufficiently dislikes earning less than the social reference point.

Scenario		Person				Reference point
		Anna	Alan	Emma	Miguel	
A	Contributions	20	20	20	20	
	Material payoff	32	32	32	32	
	Relative to reference point	0	0	0	0	32.00
F	Contributions	20	20	20	19	
	Material payoff	31.6	31.6	31.6	32.6	
	Relative to reference point	–0.25	–0.25	–0.25	+0.75	31.85
G	Contributions	20	20	20	19	
	Material payoff	31.4	31.6	31.6	30.6	
	Relative to reference point	+0.1	+0.3	+0.3	–0.9	31.3

interests to contribute $20 provided two (or more) people are willing to punish. In explaining why, I will keep things simple by assuming it costs $t = 0.1$ to punish, and we want to deter people from only contributing $19 rather than $20. Suppose that Anna threatens to lower the payoff of anyone who contributes $19 by an amount of $2. This would cost her $0.20. Moreover, suppose that Alan threatens to lower the payoff of Anna by $2 if she contributes $19. Will this work?

It will, as illustrated in Table 7.3. We first need to check that the threat of punishment would deter Miguel from lowering his contribution to $19. Comparing scenarios A and G we see that, if Miguel contributes $19 and is punished, he loses both in material payoff and relative terms. This is enough to deter him from contributing $19. We next need to check that Anna would be willing to punish. This is slightly trickier, but comparing scenarios F and G we see that, by punishing, she loses in material terms and gains in relative terms. Provided the utility gain from being closer to the social reference point exceeds the $0.20 cost of punishing, she will want to punish.

If, therefore, there are some people like Anna and Alan who are relatively inequality-averse, high contributions can be sustained by the threat of punishment. More generally, we have seen that the ERC model can do a good job of explaining a lot of the evidence we saw in section 7.1, even in situations of complete information. Before long, however, we shall see that, when people do have complete information, their utility will likely depend on individual actions in a way that the ERC model cannot capture so well. What I want to do now, therefore, is introduce a model suited to situations of complete information.

7.2.2 Inequality aversion with complete information

Fehr and Schmidt (1999) have developed a model of inequality aversion in which a person compares his payoff against that of others at the individual rather than group level. For want of a better name it is usually called the **Fehr–Schmidt model**. As before, imagine a group of n people who have played a game and earned material

payoffs of $u_1, u_2, \ldots, u_n$. To work out the utility of Miguel, we first distinguish between people who earned a higher and lower material payoff than him. So, let $H_M = \{j: u_j > u_M\}$ be the set of people who earned more than he and $L_M = \{j: u_j < u_M\}$ be the set of people who earned less. The utility of Miguel is then calculated using

$$U_M(u_1, \ldots, u_n) = u_M - \alpha_M \frac{1}{n-1} \sum_{j \in H_M} (u_j - u_M) - \beta_M \frac{1}{n-1} \sum_{j \in L_M} (u_M - u_j),$$

where $\beta_M \leq \alpha_M$ are measures of Miguel's inequality aversion.

To explain how this works, let's concentrate first on the α_M bit. This captures Miguel's disutility from earning less than others. For every person who has a higher material payoff than he, he suffers a utility loss proportional to the difference in payoff. The higher is α_M, the more he dislikes earning less than others. The β_M bit captures his disutility from earning more than others. Again, for every person who has a lower material payoff than he, he suffers a utility loss proportional to the difference in payoff. The assumption that $\beta_M \leq \alpha_M$ means that he gets more disutility from earning less than others than he does from earning more than others, which, hopefully, seems intuitive.

To better illustrate the model we can work through the same examples as for the ERC model, starting with the **dictator game**. If Miguel splits the \$10 so that he keeps \$$p$ and the receiver gets \$$10 - p$ then his utility will be $p - \alpha_M (10 - 2p)$ if $p <$ \$5, and $p - \beta_M (2p - 10)$ if $p \geq$ \$5. Given that $\beta_M \leq \alpha_M$, there is never any incentive for him to give more than \$5, so he should maximize $p - \beta_M (2p - 10)$ with respect to p. Consequently, he should give \$0 if $\beta_M < 0.5$ and give \$5 if $\beta_M \geq 0.5$. This is not exactly what we observe in Figure 7.1, but we do at least see a possible incentive for the proposer to share something.

In the **ultimatum game** proposers again have the same utility function as in the dictator game, so we just need to think about what receivers will do. Assuming that proposers will at best offer an equal share (i.e. $p \geq$ \$5), the utility function of the receiver, Anna, will depend on whether she accepts or rejects. If she accepts, she gets $10 - p - \alpha_A (2p - 10)$, and, if she rejects, she gets zero. Thus, Anna will accept if

$$p < 10\left(\frac{1+\alpha_A}{1+2\alpha_A}\right).$$

This means that, if Anna does not mind earning less than Miguel, $\alpha_A = 0$, she will accept any offer, but, if she is more inequality-averse, α_A increases, she will accept only offers closer to an equal share. For example, if $\alpha_A = 1$, she will reject anything less than a one-third share.

This possibility of rejection may make the proposer think a little. We know from looking at the dictator game that, if $\beta_M \geq 0.5$, Miguel will offer an equal share, and this will be accepted. It is slightly trickier to say what he should do if $\beta_M < 0.5$. He would want to offer the minimum amount that will be accepted. The problem is, he is unlikely to know the inequality aversion of Anna and so will not know what this minimum is. Instead, he will have to form some expectation about what the responder will do and offer an appropriate share. He therefore faces a risky choice. In the next

section we shall look in more detail at how expectations and beliefs can be captured in a model of social preferences.

To further illustrate the Fehr–Schmidt model, suppose we now add competition into the mix. In the **ultimatum game with proposer competition** it is natural to assume that the receiver will accept the best proposal. This means that there will be inequality for sure, because the receiver and one proposer will get a positive material payoff while all other proposers get zero. Given this, the objective of a proposer should be to offer an amount that will be accepted. This way, he will at least be earning more than others, which is better than earning less than others. If all proposers are aiming to get their offer accepted then they should offer almost all of the money. This is consistent with what we saw in Figure 7.10. I leave you to think about what happens with responder competition.

Finally, we can look at the **linear public good game**. Depending on how inequality-averse people are, it may or may not be an equilibrium to contribute to the public good. Table 7.4 will help illustrate why. Suppose that everyone is contributing \$20. Would Anna want to contribute less? Comparing scenarios A and H, we see that, for every \$1 she decreases her contribution, she gains \$0.60 in material payoff, but increases the gap between her payoff and others by \$1. If she dislikes earning more than others, or, more specifically, $\beta_A > 0.6$, then she does best to contribute \$20. If she does not mind earning more than others, $\beta_A < 0.6$, then she would contribute \$0. We could, therefore, see high or low contributions depending on preferences.

Most realistic, given the heterogeneity that we observed in section 7.1, is that some people have a high β and some a low β. This raises the interesting question of how many people it needs to contribute \$0 for all those with a high β to also

Table 7.4 The linear public good game and Fehr–Schmidt model from the perspective of Anna. Anna may not want to contribute less than others (scenarios A and H) but may want to lower her contribution if someone else contributes zero (scenarios I and J). She may also be willing to punish anyone who contributes zero (scenarios I and K).

Scenario		Person			
		Anna	Alan	Emma	Miguel
A	Contributions	20	20	20	20
	Material payoff	32	32	32	32
	Relative to others	—	0	0	0
H	Contributions	19	20	20	20
	Material payoff	32.6	31.6	31.6	31.6
	Relative to others	—	+1	+1	+1
I	Contributions	20	20	20	0
	Material payoff	24	24	24	44
	Relative to others	—	0	0	–20
J	Contributions	19	20	20	0
	Material payoff	24.6	23.6	23.6	43.6
	Relative to others	—	+1	+1	–19
K	Contributions	20	20	20	0
	Material payoff	21.5	24	24	19
	Relative to others	—	–2.5	–2.5	+2.5

contribute \$0. In Table 7.4 we can compare scenarios I and J to see what happens if Anna contributes \$20 or \$19 and Miguel contributes \$0. In this case, when Anna lowers her contribution to \$1, she gains \$0.60 in material payoff, gains by lowering the gap between her and Miguel but loses by increasing the gap with the other two. Overall, she gains if \$0.60 + 0.33 $\alpha_A > 0.67\ \beta_A$. This means that β_A needs to be at least 1.8 to stop Anna also contributing \$0. It takes only one person, therefore, to contribute \$0 for it to be likely the other three will too. This is consistent with the falling contributions we observe in Figure 7.7.

[Extra] For those who like a bit of algebra, models of social preferences are fun to play around with. In a linear public good game the task is made a bit easier by the fact that, if Anna contributes c_A and Miguel contributes c_M, then, irrespective of what else happens, the difference in their material payoffs is $u_A - u_M = c_M - c_A$, the difference in contributions. So, if there are n people each endowed with \$20 and contributions are $c_1, c_2, \ldots, c_n$, the utility of Anna will be

$$U_A = 20 - c_A + m\sum_j c_j - \alpha_A \frac{1}{n-1}\sum_{j\in H_A}(c_A - c_j) - \beta_A \frac{1}{n-1}\sum_{j\in L_A}(c_j - c_A).$$

Now suppose that K people are contributing \$0, and the rest are contributing \$20; what will Anna want to contribute? If she contributes \$20, her utility is

$$U_A = 20m(n-K) - \alpha_A \frac{20K}{n-1}.$$

If she contributes \$0, her utility is

$$U_A = 20 + 20m(n-1-K) - \beta_A \frac{20(n-1-K)}{n-1}.$$

So, she does better to contribute \$0 if

$$(1-m)(n-1) > \beta_A(n-1-K) - \alpha_A K$$

or the proportion of those contributing \$0 is

$$\frac{K}{n-1} > \frac{\beta_A + m - 1}{\alpha_A + \beta_A}.$$

This gives us a measure of how many people will need to not contribute for others to also not contribute.

As you might imagine, adding punishment only complicates the analysis. We shall, therefore, restrict ourselves to showing that it can be an equilibrium for all people to contribute when punishment is allowed. In fact, we shall show that it is enough to have just one person willing to punish. Suppose that $\beta_A > 0.6$, and so we know Anna would contribute \$20 if everyone else does. Suppose also that Anna threatens to lower the

payoff of anyone who contributes \$0 by \$25 and that this will cost her \$2.50, because $t = 0.1$. Will this work? This size of punishment will clearly deter any person from contributing \$0, but will Anna be willing to pay to punish this much? In Table 7.4, by comparing scenarios I and K, we see what happens if she punishes Miguel for contributing \$0. She loses the \$2.50 in material payoff but is now only \$2.50 behind or ahead of the others rather than \$20 behind Miguel. She must be better off if $5\alpha_A - 5/6\beta_A >$ \$2.50, and, because $\alpha_A > \beta_A > 0.6$, we know that this must be the case. She is therefore willing to punish Miguel so as to narrow the gap in their material payoffs.

7.2.3 An evaluation of inequality aversion models

The ERC and Fehr–Schmidt models both do a good job of fitting the experimental data summarized in section 7.1. In particular, they are able to account for cooperation in some contexts, such as a public good game with punishment, and a lack of cooperation in other contexts, such as a public good game without punishment. The models also have the virtue of being relatively simple and easy to apply. The news is not all good, however, because we shall see that, when pushed a bit further, both the ERC and Fehr–Schmidt models start to struggle. The models do, therefore, have their limitations, and so we shall need to think of alternatives to inequality aversion. Before we do that I want to briefly show that in some situations there are different models of inequality aversion that can outperform the ERC and Fehr–Schmidt models.

Engelmann and Strobel (2004) asked subjects to play a variety of **distribution games**. In the different games, three of which are given in Table 7.5, person 2 was able to decide which of the distribution of payoffs A, B or C should be used. Different models of inequality aversion give different predictions. In the taxation game, for example, the Fehr–Schmidt model predicts option A, while the ERC model predicts option C (but I will leave you to check this).

In the final row of Table 7.5 we have the choices actually made to compare with the predictions. It is clear that the ERC model does a relatively poor job at predicting

Table 7.5 Payoffs, predictions and choices in distribution games. Person 2 gets to choose which distribution of payoffs there should be: A, B or C.

Payoffs	**Taxation game**			**Envy game**			**Rich game**		
	A	**B**	**C**	**A**	**B**	**C**	**A**	**B**	**C**
Person 1	8.2	8.8	9.4	16	13	10	11	8	5
Person 2	5.6	5.6	5.6	7	8	9	12	12	12
Person 3	4.6	3.6	2.6	5	3	1	2	3	4
Predictions									
Efficiency	A			A			A		
ERC			C		B	C	A		
Fehr–Schmidt	A					C	A		
Maximin	A			A					C
Choices									
Percentage	84	10	6	77	13	10	27	20	53

Source: Engelmann and Strobel (2004).

the data. This would suggest that subjects did take account of the full distribution of payoffs. What seemed to matter most to subjects was maximizing the minimum payoff that any person would receive, which we can think of as a **maximin** notion of fairness. The Fehr–Schmidt model can pick up this effect to some extent, because it predicts that a person will get disutility from earning more than others. The model also predicts, however, that a person would be more willing to lower the payoff of those earning more than increase the payoff of those earning less; this was not observed in these experiments, particularly the envy game.

In these distribution games, an inequality aversion model based on maximin would do a good job of fitting the data. Maximin would do a lousy job, however, of fitting the data for an ultimatum game or public good game, with or without punishment. This illustrates how different strategic settings likely call for different models. We shall now look at some models that are not based on inequality aversion.

7.3 Intentions and social norms

In models of inequality aversion, people are motivated solely to earn no more and no less than others. To illustrate why this sometimes appears to miss something, I will return to the ultimatum game and look at an **ultimatum game with restricted choice**. When I first explained the ultimatum game I said the proposer could propose to split the $10 in any way that he chose. Offers of an $8, $2 split were often rejected, seemingly because of the unequal payoffs. Suppose now that we constrain the proposer to offer either an $8, $2 split or a $10, $0 split. The offer of an $8, $2 split does not look quite so unfair now. But will it still be rejected?

A study by Falk, Fehr and Fischbacher (2003) allows us to answer this. They considered four different games. In each game the proposer must choose between just two possible offers. As Table 7.6 summarizes, one of the offers is always an $8, $2 split. What we want to know is how the perception of an $8, $2 offer changes relative to the other offer a proposer could have made. If a responder cares only about inequality aversion then the other offer should make no difference, because she still gets $2 and the proposer $8. We clearly see, however, that it does make a difference. Responders are more likely to reject an $8, $2 split if the proposer could have proposed a $5, $5 split compared to if he could have proposed a $10, $0 split.

Table 7.6 The proportion who offered and the proportion who would reject an $8, $2 split in ultimatum games with restricted choice.

Game	Proposer's choices		Interpretation of an $8, $2 offer	Proportion of $8, $2 offers (%)	Proportion who would reject an $8, $2 offer (%)
	A	B			
(5/5)	$8, $2	$5, $5	Unfair	31	44
(2/8)	$8, $2	$2, $8	Relatively fair	73	27
(8/2)	$8, $2	$8, $2	No choice	100	18
(10/0)	$8, $2	$10, $0	Fair	100	9

Source: Falk, Fehr and Fischbacher (2003).

These results are not consistent with inequality aversion. Arguably, what is missing is that inequality aversion does not take into account the intentions of proposers. If Miguel offers \$8, \$2 when he could have offered \$10, \$0 then he is intending to be fair. If he offers \$8, \$2 when he could have offered \$5, \$5 he is intending to be unfair.

We have already seen something similar if you look back to Figure 7.9 and the moonlighting game. You may recall that in the no intentions experiment a computer randomly chose the proposer's offer, and in this case responders did little to reward or punish proposers. Again, if a computer generates the offer then there is no intention on the part of proposers, and this seems to matter. We need, therefore, some way to build intentions into a model of fairness.

7.3.1 A model of fairness based on intentions

I shall now work through a model of fairness that takes account of intentions. The model is based on the work of Rabin (1993) and Dufwenberg and Kirchsteiger (2004). Unfortunately, the model is not as easy as the inequality aversion models, but, if you stick with it, all will – hopefully – become clear.

Consider Miguel and someone called Federica playing a game. The key thing we want to do is think about what Miguel believes Federica will do, and what Federica believes Miguel will do. **Beliefs** are going to give us a means to think about intentions. Here is why.

Suppose that, if Miguel chooses action a_M and Federica chooses action a_F, they get material payoffs $u_M(a_M, a_F)$ and $u_F(a_M, a_F)$, respectively. Now, if Miguel believes that Federica will choose action b_F, he should also believe that if he chooses action a_M they will get payoffs $u_M(a_M, b_F)$ and $u_F(a_M, b_F)$, respectively. This suggests that we can think of Miguel as choosing payoffs for both of them from the set

$$U(b_F) = \{u_M(a_M, b_F), u_F(a_M, b_F) | a_M \in S_M\},$$

where S_M is the set of actions that he could choose.

Ignore from set $U(b_F)$ any payoff pairs that are not Pareto efficient and let $U_F^h(b_F)$ and $U_F^l(b_F)$ be the highest and lowest payoffs that Federica could get from set $U(b_F)$. We then call

$$U_F^e(b_F) = \frac{U_F^h(b_F) + U_F^l(b_F)}{2},$$

which is the payoff halfway between the highest and lowest, the **equitable payoff**. We now come to the crux of the model; we say that Miguel's **kindness** to Federica if he chooses action a_M and believes that Federica will choose action b_F is

$$k_{MF}(a_M, b_F) = u_F(a_M, b_F) - U_F^e(b_F).$$

So, Miguel is **kind** if he chooses an action that he believes will give Federica more than the equitable payoff. He is **unkind** if he chooses an action that he believes will give Federica less than the equitable payoff.

We next assume that Miguel wants to be kind to Federica if Federica is kind to him. The slight problem is that Miguel cannot know whether Federica is being kind to him, because he does not know Federica's beliefs. So, Miguel will need to form a belief about whether Federica is being kind. To do this, let c_M denote Miguel's belief about b_M. That is, Miguel believes that Federica believes that he will choose action c_M. Miguel's beliefs about how kind Federica is being to him are now

$$\lambda_{MF}(c_M, b_F) = u_M(c_M, b_F) - U_M^e(c_M).$$

This expression naturally extends kindness to beliefs about kindness.

The utility function of Miguel can finally be written as

$$U_M(a_M, b_F, c_M) = u_M(a_M, a_F) + \mu_M k_{MF}(a_M, b_F) \cdot \lambda_{MF}(c_M, b_F),$$

where μ_M is Miguel's desire to reciprocate. If Miguel does not care about fairness then $\mu_M = 0$ and his utility is his material payoff u_M. If he does care about fairness and believes that Federica is being kind, so $\lambda_{MF}(c_M, b_F) > 0$, then he can increase his utility by also being kind, $k_{MF}(a_M, b_F) > 0$. If he believes that Federica is being unkind then he would also want to be unkind. A desire for reciprocity is, therefore, built into the utility function.

It is important that beliefs correspond to reality, and so we say that there is a **fairness equilibrium** if $a_M = b_M = c_M$ and $a_F = b_F = c_F$.

This model may be a bit tricky to get one's head around, but the basic idea is relatively simple. Miguel wants to be kind to Federica if Federica is kind to him, and vice versa. Where things get complicated is determining what is kind or not. This is why we needed to think about beliefs, and beliefs about beliefs. To better understand what is going on, I will work through the **ultimatum game**.

We can think of the proposer, Miguel, as proposing an amount $\$a_M$ to give, while the receiver, Federica, simultaneously sets a minimum amount $\$a_F$ that she will accept. If $a_M \geq a_F$ then Federica accepts the offer. If Miguel believes that Federica will accept offers of b_F or more, he can offer $\$10 = U_F^h(b_F)$ or $\$b_F = U_F^l(b_F)$, or anything in between, and the outcome will be Pareto efficient. The equitable payoff is, therefore,

$$U_F^e(b_F) = \frac{10 + b_F}{2}.$$

Miguel's kindness is

$$k_{MF}(a_M, b_F) = a_M - \frac{10 + b_F}{2}.$$

This means that Miguel has to give quite a lot to be considered kind. For example, if $b_F = \$2$ then an offer of \$6 or less is not kind!

If Federica believes Miguel will offer b_M, she can either reject by setting $a_F > b_M$ or accept with $a_F \leq b_M$. Rejection can never, however, be Pareto efficient. This means

that $\$10 - b_M = U_M^l(b_M) = U_M^h(b_M)$. If Federica accepts, her kindness is zero, and, if she rejects, her kindness is $k_{FM} = -(10 - b_M)$. Federica can never be kind in this game!

Given that Federica cannot be kind, Miguel can never gain by being kind to Federica. Miguel does still have to offer enough, however, that Federica will not reject. If Federica accepts an offer of a_M, her utility is a_M, and, if she rejects, it is

$$\mu_F k_{FM}(b_M, a_F)\lambda_{FM}(b_M, c_F) = -\mu_F(10 - b_M)\left(b_M - \frac{10 + c_F}{2}\right).$$

In equilibrium, therefore, where $b_M = c_F = a_M$, Federica will accept if $2a_M \geq \mu_F(10 - a_M)^2$. An offer of zero will, therefore, be rejected if $\mu_F > 0$, but a high enough offer will be accepted.

These results are similar to those obtained in the inequality aversion models. Can this model also capture the importance of intentions? In the ultimatum game with qualified choice, the model predicts that a responder may reject an \$8, \$2 split in the (5/5) and (2/8) games, but should never reject an \$8, \$2 split in the (8/2) or (10/0) games. This is because the offered split is unkind in the first two games and kind or neutral in the last two games. The model does, therefore, take account of intentions. It fails to completely fit the data, however, because offers of an \$8, \$2 split were rejected some of the time in all the games.

7.3.2 What is fair?

The crux of the intentions-based model of fairness that we have just been looking at is the equitable payoff, because kindness is judged by whether or not utility is above or below this equitable payoff. The equitable payoff can be thought of as an example of a **fair reference point**, about which losses or gains are measured. A theme in Chapters 2, 3 and 4 was the potential arbitrariness of reference points; things are no different here. The equitable payoff is that payoff halfway between the highest and lowest possible (Pareto-efficient) payoff, but why should the fair reference point be this and not something else?

The model is flexible to using any fair reference point, so we do not need to stick with the equitable payoff. But what reference point should we be using? We have looked at two candidates so far. In inequality aversion models, the fairness of a person's material payoff is judged relative to the payoff of others. So, the fair reference point is what others are earning, and the fairest outcome is one in which everyone earns the same material payoff. In the intentions-based model we have just looked at, the fairness of a person's payoff is judged relative to the worst and best material payoff that she could have got. The fairest outcome is one in which she gets at least halfway between this worst and best.

In a symmetric game, such as the ultimatum game with unconstrained choice, these two notions of a fair reference point will turn out the same. But in general they are not the same. For example, in the ultimatum game with constrained choice (10/0), a split of \$8, \$2 is unfair by the standards of inequality aversion because the proposer gets more than the receiver, but is fair by the standards of the intentions-based model because the receiver could have got less than she did.

The results that we have for the ultimatum game with constrained choice suggest that both notions of fair reference point are important. Is it possible to combine both elements in the same model? Falk and Fischbacher (2006) show that it is if we use an intentions-based model of fairness with a different definition of kindness.

In this revised model, the kindness of Miguel consists of two parts:

$$k_{MF}(a_M, b_F) = \delta_{MF}(a_M, b_F)\,\vartheta_{MF}(a_M, b_F).$$

The first part, $\delta_{MF}(a_M, b_F)$, captures inequality aversion and is calculated as follows:

$$\delta_{MF}(a_M, b_F) = u_F(a_M, b_F) - u_M(a_M, b_F).$$

This means that Miguel is **kind** if he chooses an action that he expects to give Federica a higher material payoff than it does him. He is **unkind** if he chooses an action that he expects to give Federica a lower material payoff than it does him. The second part, $\vartheta_{MF}(a_M, b_F)$, of the kindness expression captures intentions and is called the **intention factor**. Letting $u_M = u_M(a_M, b_F)$ and $u_F = u_F(a_M, b_F)$, it is calculated using

$$\vartheta_{MF}(a_M, b_F) = \begin{cases} 1 \text{ if } u_F > u_M \text{ and there exists } a \text{ such that } u_F(a, b_F) < u_M \\ \varepsilon \text{ if } u_F > u_M \text{ and there is no } a \text{ such that } u_F(a, b_F) < u_M \\ 1 \text{ if } u_F < u_M \text{ and there exists } a \text{ such that } u_F(a, b_F) > u_M \\ \varepsilon \text{ if } u_F < u_M \text{ and there is no } a \text{ such that } u_F(a, b_F) > u_M. \end{cases}$$

This might look complicated, but the intuition is straightforward. The first two terms consider the case when Federica is earning a higher material payoff than Miguel. This appears kind, but we ask: did Miguel have any alternative but to be kind? If he did then the intention factor is 1 and Miguel is considered **intentionally kind**. If he did not then the intention factor is $\varepsilon < 1$ and so we consider Miguel **unintentionally kind**.

Similarly, the bottom two terms consider the case when Federica is earning less than Miguel. This looks unkind, but maybe Miguel had no choice? If he did, the intention factor is 1, and, if he did not, the intention factor is $\varepsilon < 1$. For example, offering an \$8, \$2 split in the (10/0) game is considered **unintentional unkindness**, so is not as bad as offering an \$8, \$2 split in the (5/5) game, which is **intentional unkindness**.

This model has both inequality aversion and intentions and can do a great job of fitting the data for the ultimatum game with and without restricted choice. It demonstrates, therefore, that intentions-based models are flexible enough to consider any notion of what is fair or kind (see also Research methods 7.5). That still leaves us asking, however, what we should consider kind. As with models of inequality aversion, the Falk and Fischbacher model proposes that fairness is primarily judged by comparing the relative material payoffs of people. Is there any situation in which fairness should be judged relative to the best and worst payoff that a person could have got?

❖ RESEARCH METHODS 7.5 PSYCHOLOGICAL GAME THEORY

A focus on beliefs and intentions is at the heart of something called **psychological game theory**, first introduced by Geanakoplos, Pearce and Stacchetti (1989). In 'traditional' game theory, as discussed in Chapter 1, the payoffs in the game should capture the emotional as well as the material payoff. The problem with this approach is that a person's emotional payoff may depend not only on what happens but also on what the person believed would happen. For instance, emotions such as envy, surprise and disappointment happen relative to some expectation. A psychological game theory approach allows us to explicitly model this and so keep some separation between emotional and material payoffs.

To illustrate how psychological game theory can be applied in settings other than reciprocity, I'll return to the issue of honesty and deception that we first looked at in Chapter 2. It was suggested in Chapter 2 that someone who lied would feel guilty about that. Following an approach due to Charness and Dufwenberg (2006), we can add a somewhat subtle twist to the story. We are now going to say that a person feels guilty about lying only if they were expected to tell the truth.

To fix ideas, consider the following scenario. Federica is deciding whether or not to employ Miguel to clean her windows. Miguel charges \$10 to clean the windows, Federica values clean windows at \$15, and the cost to Miguel of cleaning the windows is \$5. If Federica decides not to employ Miguel, they both have payoff zero. If she employs him and he works hard then both have net material payoff \$5. If she employs him and he slacks on the job then Federica has lost \$10 and Miguel has gained \$10. Miguel clearly has an incentive to slack on the job, because \$10 beats \$5. Suppose, however, that he promises Federica that he will do a good job, and he feels guilty if he breaks his promise.

In this case, Federica must decide whether or not to employ Miguel, a_F, and Miguel must decide whether or not to work hard, a_M. When Federica makes her decision she forms some belief about what Miguel will do. Suppose that she believes he will choose b_M. When Miguel makes his decision he forms some belief about what Federica believes he will do. Suppose that he believes Federica believes he will choose c_M. Miguel will feel guilty if he believes that Federica believes he will work hard (c_M is 'work hard') and yet he slacks on the job (a_M is 'slack'). If this guilt is enough to offset the \$5 gain from slacking, Miguel will work hard and Federica should employ him. If, however, Miguel believes that Federica believes he will slack (c_M is 'slack') then he feels no guilt from breaking his promise to work hard.

The interesting thing about this story is that Miguel is not really seen as deceiving Federica or lying to her if no one believes his promise. This generates multiple equilibria. There is an equilibrium where Miguel keeps his

promises because he is expected to keep his promises, and an equilibrium where Miguel does not keep his promises because no one expects him to keep them. Psychological game theory is a really nice tool for highlighting such possibilities and allowing us to analyze them.

One issue we have not touched on yet is that of property rights. In all the experiments we have discussed so far, subjects were randomly assigned roles in the game. In this case, it seems natural that fairness should be judged by whether people earn the same material payoff. Now, suppose that people earn their material payoffs in some way and so have a stronger sense of ownership or **property rights** over their material payoff. Is it fair that a person who has earned more than another should give this away?

To start to answer this question we can look at a **dictator game with take and earnings**. Figure 7.13 provides some results from a study by List (2007). In the baseline treatment, both proposer and receiver were given $5 (a show-up fee) and the proposer was given an additional $5 that he could split with the receiver. This corresponds to the standard ($5) game of Figure 7.1 (and the results are similar). In the 'take $1' and 'take $5' treatments, the proposer could not only offer some of his amounts of $5 to the receiver but also take up to $1 or $5 from the receiver. In the earned $5 treatment, the proposer and receiver 'earned' their money by completing a task, including sorting and handling mailings for a charitable fundraiser.

In the baseline treatment, most proposers offer money, resulting in more equality between the proposer's and receiver's payoff. In the 'take $1' and 'take $5'

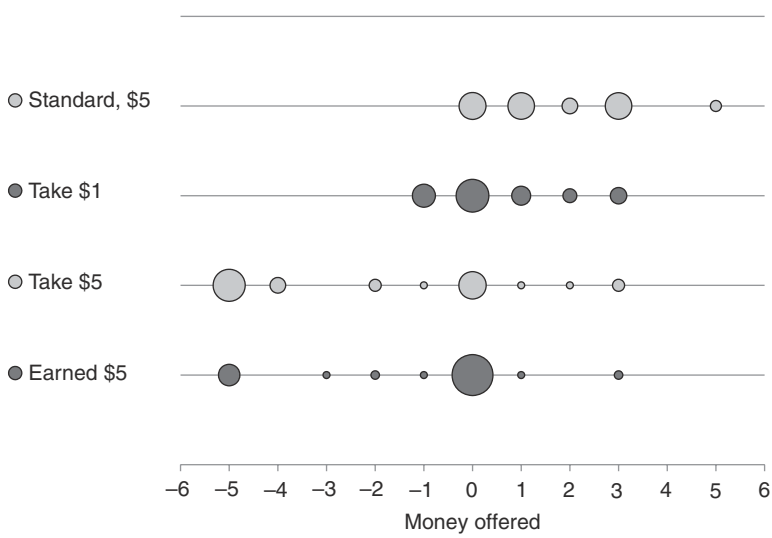

Figure 7.13 The amount shared by dictators in dictator games with take and earnings. Less is shared when there is the chance to take money, or money is earned.
Source: List (2007).

experiments, many fewer proposers give money. The fact that they can take money appears to excuse not giving money. This is consistent with an intentions-based model of fairness (such as that in section 7.3.1), in which the kindness of a proposer is judged by the worst that he could have done to the receiver, and not one judged by how equal proposer and receiver payoffs are. In the 'earned $' experiment, we see that very few proposers take or offer money. Seemingly, money earned is money deserved.

What we learn from this study is that property rights and norms, or institutions, influence what people give, and therefore presumably what is considered fair or kind. Taking money is different from giving money. To have earned money is different from having been randomly given money. Other studies have similarly shown that perceptions of what is fair or kind will depend on how the problem is framed and the social norms that this brings to mind. Sometimes it is natural to judge kindness by equality of monetary payoffs. In others it is natural to judge kindness by whether a person could have done worse or better for another. There is, therefore, no set prescription for what the reference point should be. This should come as no surprise, though, given what we have seen in previous chapters.

7.4 Summary

We started by looking at a series of games and experimental studies that suggest people have social or interdependent preferences. The evidence suggested that many people are reciprocators, in that they want to be kind to those who are kind and unkind to those who are unkind.

In our first attempt to formally model such reciprocation, we looked at two models of inequality aversion. The basic idea behind both models is that a person dislikes earning less than others, and possibly dislikes earning more than others. The difference between the models that we looked at is the information a person is assumed to have. In the ERC model a person compares her payoff to the average payoff of others. In the Fehr–Schmidt model a person compares her payoff to the individual payoffs of others.

A great advantage of the ERC and Fehr–Schmidt models is that they are simple and easy to apply. We also saw, however, that they capture only part of the story of social preferences. To try to fill the gap we considered models that take account of intentions. This meant that we needed to think about what a fair reference point is, what is kind or unkind and how we can model beliefs about kindness. We saw that the answers would depend on context and framing effects.

With intentions-based models we can account for much more of the behavior that we observe, but that is primarily because the models are far more flexible. One symptom of this, which I mentioned only in Research methods 7.5, is that we can very easily end up with ambiguous outcomes: an equilibrium where everyone is unkind because everyone is being unkind, and an equilibrium where everyone is kind because everyone is being kind. Which of these is more plausible? Again, the answer is likely to depend on context and framing effects.

In short, there is no 'perfect' model of social preferences that we can apply without much thought. What we do have, however, is a range of models and ways of thinking about social preferences that allow us to capture behavior on a case-by-case basis.

In some applications the simplicity of an inequality aversion model will be appealing. In other applications it will be more important to take careful account of intentions.

Before we get on to the applications, there is one last thing I want to mention. While we have seen clear evidence that many people voluntarily give to others, we should not overplay the hand. In all the games we have looked at there are always some who behave 'as if selfish', and even those who give appear to give only 'with strings attached'. Furthermore, in Figure 7.13, and a study I will look at in the next section, we see that many appear less likely to give if they have some excuse or wiggle room for not giving. The world is not full of altruists. But it is full of people with very diverse social preferences.

7.5 Giving to charity

A lot of people give time and money to charitable causes. For example, the World Giving Index 2016 (which collected data from 140 countries) reports that around 31 percent of people donated money to charity in 2015 and around 22 percent volunteered time. A survey by the Charities Aid Foundation in 2016 found that over 60 percent of UK adults gave to charity in the previous year, with donations of regular givers averaging £40 per month; an estimated 5 million people (around one in 20 of the population) raised money by competing in a running event! The Giving USA Foundation estimates charitable contributions in the United States in 2016 to be $390 billion. These are big numbers.

To understand why so many people give so much to charity, we surely need to think about social preferences. The story that I have told so far seems incomplete in understanding charitable giving, however. That's because the picture I have drawn is one of people helping and giving to others with strings attached. People help those who they see being kind or are earning less than them. Is there any role for pure altruism, people just helping others unconditionally? This seems plausible because many people give money and time to charity without knowing whether this will help people who are kind or not, or have more money than them.

So, why might a person give money or time to charity? There are three broad explanations. They might just want to help people; this would be **pure altruism**. They might feel good about themselves after giving, because of, say, prestige, respect or social approval; call this a **warm glow** or **joy of giving**. Finally, they might feel bad about themselves if they do not give, because of, say, the social disapproval of others, or guilt. These last two reasons give rise to **impure altruism**, in that the person directly benefits herself from giving, and so her motives for giving are not purely altruistic. But does it matter whether a person gives for pure or impure reasons? Yes, for at least two reasons, which I shall look at in turn.

7.5.1 Crowding out

Suppose that Federica was planning to contribute to a charity but then finds out that someone else has already done so. If Federica is motivated by pure altruism she would no longer want to give, because somebody else has done the good work for her. This is **crowding out**. If Federica is motivated by impure altruism she might still want to

give – indeed, she may feel more incentive to give, in order to compete with the other giver. There will be less crowding out.

To illustrate the issues, I will consider a special case of a **model of warm glow** developed by Andreoni (1990). Consider a group of individuals contributing to a charity or public good. The utility function of Federica is given by

$$U_F(w_F - c_F, C, c_F) = w_F - c_F + \eta_F\sqrt{C} + \rho_F\sqrt{c_F},$$

where w_F is Federica's wealth, c_F is how much she contributes to the public good, C are the total contributions (including hers) and η_F and ρ_F are parameters. If $\rho_F = 0$ then we can say that Federica is a pure altruist, because all she cares about are her wealth after contributing and total contributions to the charity. If $\rho_F > 0$ then we can say that Federica is an impure altruist, because she gets utility from having contributed. If $\eta_F = 0$ then Federica contributes only to get a warm glow.

If we solve for the optimal contribution we get relation

$$1 = \frac{\eta_F}{2\sqrt{C}} + \frac{\rho_F}{2\sqrt{c_F}}.$$

If $\rho_F = 0$ then

$$c_F = \frac{\eta_F^2}{4} - C_{-F},$$

where C_{-F} is total contributions of those other than Federica. The pure altruist, therefore, wants $\eta_F^2 / 4$ to be contributed to the charity, and she does not care who contributes. If C_{-F} increases then she will decrease c_F by the same amount and there will be complete crowding out. If $\eta_F = 0$ then

$$c_F = \frac{\rho_F^2}{4},$$

and so the amount given is independent of the amount given by others. There will be zero crowding out. Whether or not giving is because of altruistic or impure altruistic reasons does, therefore, matter.

Testing whether or not there is crowding out can be done in the lab or with real data on charitable giving. In the lab, the results suggest that there is sizeable but not complete crowding out. For example, Andreoni (1993) and Bolton and Katok (1998) ran two very different experiments to test the extent of crowding out, and find the remarkably similar results of 71.5 percent and 73.7 percent crowding out. That means that a person reduces her giving by \$0.715 or \$0.737 for every \$1 someone else gives. Real data on charitable giving shows far less crowding out. For example, Ribar and Wilhelm (2002) looked at private donations from the United States to 125 international relief and development organizations between 1986 and 1992, and find that giving by the government and other organizations crowded out private donations by

only 18 percent. The evidence consistently shows, therefore, that there is crowding out, but not complete crowding out. What it is hard to agree on is just how much crowding out there is.

Can we reconcile the differences between the experimental and empirical results? Ribar and Wilhelm suggest that we can by taking into account the number of other potential givers. To explain, suppose that Federica is an impure altruist who gets utility from both pure altruism and a warm glow. Because of pure altruism she wants total donations to be sufficiently high – $\eta_F^2/4$ in the example. Because of the joy of giving she wants to give at least a certain amount herself – $\rho_F^2/4$ in the example. If there are few other potential givers, as in the experimental lab, total donations are likely to be lower, so pure altruism will influence her decision more. If there are many other potential givers, as in real life, total donations will likely be high, so the joy of giving will influence her decision more. To see this in the model, suppose that all contributors will contribute the identical amount *c*. Then we can substitute in $C = nc$ and get

$$c = \frac{1}{4}\left(\frac{\eta_F}{\sqrt{n}} + \rho_F\right)^2.$$

If *n* is large then η_F becomes irrelevant and only a warm glow matters. If *n* is small then η_F is relevant and altruism will matter.

It seems, therefore, that we observe a mixture of both warm glow and altruistic reasons for giving. The extent of crowding out that we observe may well depend on the number of other potential givers.

7.5.2 Who is watching?

A pure altruist will give irrespective of whether or not someone will know they have given. Someone motivated by impure altruism may give more if their giving will be observed. This is because giving can serve as a way for someone to signal something good about themselves. For example, giving may signal generosity or wealth. That people may want to give more if being observed offers a potentially important means to increase donations, and the evidence is that some people do indeed give in order to signal. But, before we look at that evidence, I think it is useful to briefly look at a simple model to see how signaling interacts with altruism and a warm glow.

Suppose that giving is seen as a signal of someone's generosity. For this to make sense we need there to actually be some variation in how generous people are. In the model of warm glow, this would mean individuals vary in their values of η and ρ, with a high η and ρ being seen as generous. For simplicity, suppose that there are only two types of individuals: the generous, who have $\eta = 0$ and $\rho = 2$; and the not generous, who have $\eta = \rho = 0$. If giving is not observed then the generous will give one and receive a warm glow, while the not generous will give zero. What happens if giving is observed?

A second ingredient of any model of signaling is that someone wants to appear different from who they are. In this case, a not generous type, such as Federica, may try to look as though she is generous. To be specific, suppose that she receives a prestige bonus of *P* if she is perceived by others as generous, zero if she is perceived

as not generous and $0.5P$ if it is unclear whether she is generous or not. Recall that the generous type would give one if giving is not observed. Federica could, thus, give $c_F = 1$ in order to appear generous. If generous types continue to contribute one then we have a pooling equilibrium in which everyone contributes one. Given that the type is not distinguishable, Federica can expect a prestige bonus of $0.5P$. It is, therefore, optimal for Federica to contribute if $0.5P > 1$.

But what if a generous individual, called Miguel, also gets the prestige bonus? Then Miguel will want to contribute more in order to distinguish himself from Federica. In this case we would obtain a separating equilibrium in which Miguel contributes an amount just above what Federica would be willing to contribute in order to get her prestige bonus. This means that he would contribute a little bit more than $0.5P$. Doing so means that he gets the full bonus of P, rather than $0.5P$, and has to contribute only $0.5P - 1$ more than he intrinsically prefers. In short, Miguel does just enough to make sure Federica would not try to copy him.

To summarize, if only the not generous get a bonus for being perceived as generous then both Federica and Miguel will give one, $c_M = c_F = 1$. If generous individuals also get the bonus then Federica will give zero, $c_F = 0$, and Miguel will give more, $c_M = 0.5P$. Desires for appearing generous can thus increase the giving of either the not generous or the generous. In more general models the implications of signaling become more complicated, but giving is always likely to increase. So, what of the evidence that people do give, at least partially, for signaling reasons?

In short, there is a lot of evidence that signaling is important for some. The most compelling evidence is that very few donors to charity choose to remain anonymous. In public good experiments we also see that giving is less, the more anonymity there is. To illustrate this, and also to make a slightly different point, I want to look at a study by Dana, Weber and Xi Kuang (2007). In the study, subjects played variations of the dictator game in which proposers had reason, or not, to excuse not giving. The results are summarized in Table 7.7.

In all the experiments proposers could, with some variations, offer a fair \$5, \$5 split or an unfair \$6, \$1 split. In the baseline treatment, we see what we would expect, given the evidence in section 7.1, in that most proposers chose the fair option. In all the other treatments, proposers had some plausible reason for not choosing the fair option. We see that the proportion choosing the fair option fell dramatically.

In the hidden information treatment, a proposer knew their own payoff but had to choose to find out the payoff of the receiver. Even though this could be

Table 7.7 Offers in a dictator game with varying amounts of moral anonymity.

Treatment	Proportion choosing fair option (%)	Other details
Baseline	74	
Hidden	38	56% of subjects chose to reveal the receiver's payoff
Two dictators	35	
Possibly random	34	76% of participants chose before the cut-off time

Source: Dana, Weber and Xi Kuang (2007).

done without cost, nearly half the proposers chose not to find out the payoff of the receiver. A majority of those then chose the unfair option. In the two dictator treatments, the offer sent was determined by two dictators simultaneously, either of which could enforce the sending of a fair offer. In the possibly random treatment, the proposer's choice would be made randomly if they did not make their choice before some undisclosed cut-off time. Most did choose before the cut-off, but still chose the unfair option!

What this study nicely shows is how signaling is interrelated with intentions. In the baseline treatment, a proposer's choice made their intentions very clear, and so they could not hide, either from the receiver or from themselves. This means that the proposer's actions would send a clear signal, so they had the incentive to send a good signal and receive a warm glow, avoid guilt, etc. In the other three treatments, a proposer's choice is much less revealing about their intentions. This lessens the incentive to send a good signal, and is enough for the majority of subjects to not choose the fair offer. We see, therefore, that a little bit of wiggle room to not give can be enough to significantly reduce the amount that is given. Signaling is the most plausible explanation for this.

7.5.3 Why do people give?

The picture we get is one of people giving for a variety of reasons. Giving does not appear to be solely due to pure altruism, but is likely a mixture of a warm glow, guilt at not giving, signaling and pure altruism. Understanding this can be useful in getting people to give more.

To illustrate these points, I shall wrap up this section by looking at a study by Alpizar, Carlsson and Johansson-Stenman (2008). They looked at giving by international visitors to Poas National Park in Costa Rica in 2006. A random sample of visitors was interviewed privately about their visit to the National Park, at the end of which they were asked to make a donation. Three things were varied: some were given a small gift of value $3; some were told that 'most common donations have been' $2, $5 or $10; and some made their donations without the interviewer knowing how much they gave.

The results are summarized in Table 7.8. A relatively high proportion did give, but it is noticeable that the proportion who gave is increasing in the lowness of the reference point. The contribution, conditional on giving something, is 25 percent higher in the non-anonymous compared to anonymous treatments. We also see higher contributions, the higher the reference amount. Indeed, many givers matched the reference amount. All these findings are impossible to reconcile with a giver motivated solely by pure altruism or one who cares only about their own payoff. Instead, they suggest the mix of motives for giving that we have looked at.

7.6 Price and wage rigidity

Take a look at the number of unemployed in just about any country and you will see some big numbers. The unemployment rate in the United Kingdom and United States is typically above 5 percent of the workforce, and in a few European countries the

Table 7.8 The proportion who contribute, the average contribution conditional on giving something and the average contribution to Poas National Park.

Treatment	Proportion contributing (%)	Conditional contribution ($)	Average contribution ($)
No gift	48	5.09	2.43
Gift	56	4.56	2.56
Anonymous	51	4.36	2.21
Non-anonymous	53	5.21	2.77
No reference	47	6.00	2.84
$2 reference	61	3.61	2.20
$5 reference	50	3.95	1.98
$10 reference	49	5.97	2.95

Source: Alpizar, Carlsson and Johansson-Stenman (2008).

number can be as high as 20 percent. Some of this could be voluntary unemployment, but, clearly, not all of it. That means that labor markets are not clearing, at least in the short run. If there is an excess supply of labor then the standard story would say that the price of labor, or the wage rate, should fall until the market clears, or supply equals demand. If wages do not fall for some reason then we say there is **wage stickiness**. There are lots of reasons why we might observe wage stickiness, such as regulation and unions, but the reason I want to focus on here is that of reciprocity and fairness.

Why should reciprocity matter? The relationship between that of an employer and a worker is a **principal–agent relationship**. The employer, or principal, hires a worker, or agent, to do a job for her. Given that the employer's profit will depend on how much effort and initiative the worker puts into her job, the employer needs to find a way to motivate the worker to work hard. Crucially, however, in most jobs the employer will not be able to tell for sure how hard the worker is working. At the least, the worker may have ways to make it look as though she is working harder than she is (see section 2.3). The employer, therefore, needs to think carefully about how to motivate the worker.

Now let's think about reciprocity. Positive reciprocity would imply that, if the employer is kind to the worker, the worker will be kind to the employer. Negative reciprocity would imply that, if the employer is unkind to the worker, the worker will be unkind to the employer. What is kind and unkind in this context? That will depend on the reference point. Given what we have seen in Chapters 2 and 4 and this chapter, however, the most obvious reference points would seem to be the wage paid previously or the wage being paid to others. To cut wages or pay a lower wage than to others would then be seen as unkind, while to increase wages or pay a higher wage than to others would be seen as kind.

We might predict, therefore, that employers would pay high wages to motivate workers and appeal to positive reciprocity. They might be reluctant to cut wages and demotivate workers because of negative reciprocity. Numerous surveys of employers have found evidence for both these effects, and particularly the latter. For example, Campbell and Kamlani (1997) asked 184 firms why they did not cut wages as low as they could in a recession. The main concern of employers was that the most productive workers would leave (an adverse selection argument), but a close second was that

Table 7.9 The percentage of employers who thought that a 10 percent wage cut or increase would change worker effort, and change it by more than 10 percent. For example, after a 10 percent wage cut it is expected that 74.7 percent of blue collar workers will decrease effort by 10 percent or more.

Change in wages	Expected change in effort	Proportion of workers who will change effort (%)		
		White-collar workers	Blue-collar workers	Less skilled workers
10% cut	Decrease	85.8	92.0	91.3
	Decrease by 10% or more	61.7	74.7	78.8
10% increase	Increase	55.8	52.1	53.2
	Increase by 10% or more	19.4	18.7	22.5

Source: Campbell and Kamlani (1997).

worker effort would decrease (a reciprocity argument). Table 7.9, for instance, shows the percentage of employers who thought a 10 percent pay cut or increase would have an effect, and a serious effect, on worker effort. A pay rise is predicted to have some benefit, but it is clear that a wage cut is expected to have a significant negative effect.

If employers are reluctant to cut wages because it will cause negative reciprocity then we have a plausible reason for wage stickiness. But we should test this explanation out a bit more thoroughly. I will do so by formalizing the argument a bit more, using a model of fairness, and then looking at some experimental results.

7.6.1 A model of worker reciprocity

To formalize how reciprocity can result in sticky wages I shall work through the intentions-based model of fairness that we looked at in section 7.3.1. Figure 7.14 illustrates. To keep things simple, imagine that a worker, Miguel, can put in either high or low effort $e \in \{H, L\}$, and high effort costs him some amount C. Suppose that, if effort is high, $e = H$, the employer, Federica, gets revenue R and, if effort is low, $e = L$, she gets zero revenue. Federica chooses a wage w and makes a profit, or material payoff, of revenue minus w. Miguel gets a material payoff of w minus effort cost.

Federica does not care about fairness and so simply wants to maximize her profit. Should she pay a high wage?

Miguel can be kind by putting in high effort and unkind by putting in low effort. The revenue of Federica is R in the first case and zero in the latter, so an equitable payoff is $0.5R - w$. The kindness of Miguel is thus $0.5R$ or $-0.5R$ depending on whether he puts in high or low effort. Federica can realistically offer any wage between zero and R, so the equitable payoff is also $0.5R$ and the kindness of Federica is $w - 0.5R$.

This means that, if Miguel puts in high effort, his utility is

$$w - C + \mu_M\left(w - \frac{R}{2}\right)\frac{R}{2}$$

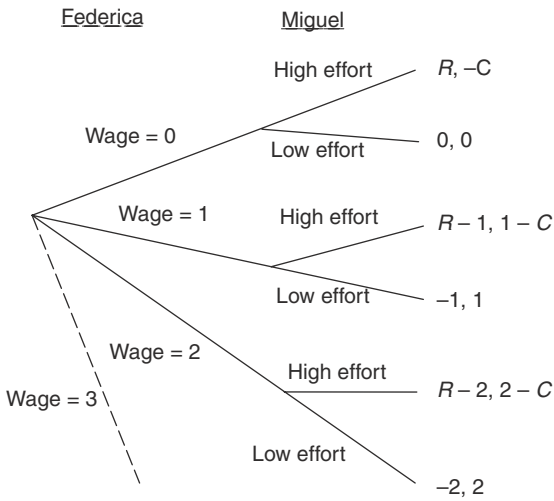

Figure 7.14 A simple model of worker reciprocity. Federica chooses a wage w and Miguel chooses effort. Federica's material payoff is revenue $-$ w and Miguel's is w $-$ effort cost.

and, if he puts in low effort, his utility is

$$w - \mu_M\left(w - \frac{R}{2}\right)\frac{R}{2},$$

where we recall that μ_M is his desire to reciprocate. Miguel will, therefore, put in high effort if

$$w > \frac{R}{2} + \frac{C}{\mu_M R}.$$

We can see that he may be willing to put in high effort if Federica is kind and offers a wage above the equitable payoff of $0.5R$.

Is Federica willing to pay this wage? If Miguel puts in high effort, she gains R in revenue and so she will be willing to pay any wage less than R. Setting $w < R$ gives the condition $2C < \mu_M R^2$. Thus, provided the cost of effort is sufficiently small and/or Miguel has sufficient desire to reciprocate, Federica will be willing to pay a high enough wage that Miguel will be willing to put in high effort.

In this example we see that a worker may be willing to work hard if he believes the employer is being kind and offering a higher-than-equitable wage. We can also see how reciprocity can cause unemployment and wage stickiness. For instance, suppose that R were to subsequently fall, but Miguel still equated kindness with his old wage. We could get wage stickiness because Miguel will think a wage cut is unkind. This could lead to unemployment if Federica can no longer afford to pay the wage Miguel wants.

Do not, however, think of reciprocity as a bad thing that we could do without. Yes, it is the reason why we might get wage stickiness and unemployment, but without it Miguel would never be willing to put in high effort, and Federica would never make any profit! So, reciprocity is beneficial, but could have negative side effects.

7.6.2 Wage stickiness in the lab

The lab provides a controlled environment to try to better understand wage stickiness and has shown the importance of reciprocity. To illustrate, I will look at a study by Fehr and Falk (1999). The basic design of the experiments was similar to that of the gift exchange game that we have already looked at in section 7.1. The main difference was a preliminary stage in which workers and employers were matched through a double auction. In this auction, workers could say what wage they would accept, employers could say what wage they would pay, and, if a worker and employer agreed on a wage, they then played the gift exchange game. This procedure was repeated ten times, but an employer could never know who the worker was, and so there was no opportunity for reputation building. The key to the experiment was that there was an excess supply of workers, 11 workers to seven employers, so unemployment was inevitable.

Figure 7.15 summarizes the main results. The reservation wage shows what wages the seven workers posting the lowest wages were willing to accept. This is, therefore, the wage that employers could have offered and got workers. We can compare this with the wage they did pay. First, look at the treatment in which effort does not count, meaning that the employer's payoff did not depend on worker effort. In this

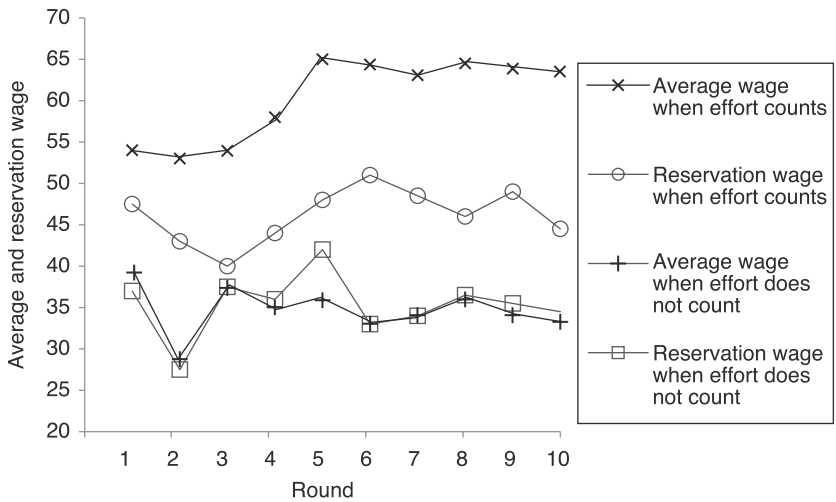

Figure 7.15 The average wage and reservation wage in experimental labor markets. In the treatment without effort the average wage is relatively low. When effort matters, we see that employers pay relatively high wages, well above the minimum that some workers were willing to accept.
Source: Fehr and Falk (1999).

treatment, there was no incentive to pay a high wage, and we consequently observe that the average wage was low and similar to the reservation wage.

Now look at the treatment we are most interested in, in which effort did count. We see that wages are relatively high and, most importantly, above the reservation wage. So, employers could have hired cheaper workers, but chose not to. They chose to hire workers who wanted a high wage! Was this a good decision? Around 34 percent of workers did not make effort conditional on wage, but 66 percent did reciprocate a high wage with high effort. This was enough to mean that paying a high wage did pay back in higher profits.

This experiment demonstrates nicely how positive reciprocity from workers can make it worthwhile to pay high wages. In the second part of their study, Fehr and Falk show that negative reciprocity can also make it worthwhile to pay high wages. They did this by changing the workers' effort costs, so that it was least costly for a worker to put in maximum effort. Workers paid a low wage still on average lowered their effort, but now at a cost to themselves. Workers are therefore willing to sacrifice their own monetary payoff to punish a low-paying employer.

7.6.3 How long to forget a wage change?

It seems clear that reciprocity is important in the labor market and is a causal factor in wage stickiness. But, before we move on, there is one final issue that I want to touch on, that of time. We have already talked in Chapters 2 and 4 about how reference points can change over time. If an employer's wage changes then their reference point should also eventually change. That would suggest that employers need to constantly increase the wage in order to motivate workers through positive reciprocity. This is consistent with the desire for an increasing wage profile that we talked about in Chapter 4. Do we have any idea how long it takes people to adjust to new wages?

A study by Gneezy and List (2006) suggests that it does not take long to forget a wage increase. They recruited people to do routine tasks, such as data entry in the university library and door-to-door fundraising, and some were given an unexpected wage increase. Those given the increase worked harder for the first 90 minutes but after that their effort dropped to the level of those who did not receive an increase. Ultimately, the wage increase did not pay for itself.

In a similar experiment, Kube, Maréchal and Puppe (2006) got a similar result to that of Gneezy and List when looking at a wage increase. They also considered an unexpected cut in the wage, however, and find that this has longer-lasting consequences. Maybe people soon forget a wage increase but never forget a wage cut? This emphasizes the asymmetry that we have already seen in Table 7.9, whereby a decrease in wages was predicted to have stronger effects on worker effort than an increase in wages. This asymmetry is consistent with loss aversion and is key to explaining why reciprocity can lead to wage rigidity. Cutting wages is worse for worker effort than increasing wages is good for effort, so firms should avoid cutting wages.

7.6.4 Firm pricing

It is not only wages that can be sticky. If there is excess demand for a good then firms, according to a standard story, should push up prices until supply equals demand.

Often, however, we observe large excess demand for goods. For example, tickets for major sporting events such as the Superbowl or World Cup final could be sold many times over. Similarly, ski resorts are bursting to capacity at Christmas and New Year. Why do prices not rise more?

There are good arguments to say that it is because of desires for fairness and reciprocity. The explanation is very similar to that for wage stickiness, so I will not spend much time on this issue. But I will talk through some of the results from a study by Kahneman, Knetsch and Thaler (1986) that help illustrate the key points. They conducted a household survey with questions such as:

> A hardware store has been selling snow shovels for $15. The morning after a large snowstorm, they raise the price to $20. Please rate this action as: Completely fair, Acceptable, Unfair, Very unfair.

In times of excess demand, firms have a rationale to increase prices, but 82 percent of respondents in the survey considered this price increase unfair. Reciprocity would suggest that this unfairness might be enough to stop people buying a snow shovel.

As with wages, the most obvious reference price is probably the previous price or the price charged by others. This means that a price increase or relatively high price is unkind while a price decrease or relatively low price is kind. To get price stickiness, we need the asymmetry whereby raising prices is considered worse than lowering prices is considered good. This is what we seem to observe. For example, another survey question was:

> A shortage has developed for a popular model of automobile, and customers must now wait two months for delivery. A dealer has been selling these cars at list price. Now the dealer prices the model at $200 above list price.

Seventy-one percent considered this unfair. Other respondents were told:

> A dealer has been selling these cars at a discount of $200 below list price. Now the dealer sells this model only at list price.

Fifty-eight percent found this acceptable.

So, is it ever acceptable to raise prices? Most people surveyed considered it fair for a firm to pass on increases in wholesale costs to consumers. They seemed to accept that firms need to make an acceptable level of profit, and so are willing to accommodate this. What was considered unfair is for a firm to raise prices because of increased market power. This is clear in the snow shovel question. As we might expect, therefore, the reference point is a little bit more subtle than the previous price.

You may disagree on what is fair or not. The main thing is that firms must take account of customers' likely perceptions of a price rise. This is because we know that people reject in ultimatum games, and so presumably are willing to not buy, a product they consider overpriced. Firms may therefore be reluctant to increase prices, with the consequence that many markets may have excess demand.

Indeed, we can be more specific about which markets will have excess demand. Primarily it will be markets where there are large fluctuations in demand, but little change in costs, and where customer loyalty is important. This is because the large

fluctuations in demand will cause there to be excess demand but the importance of customer loyalty will make firms reluctant to increase prices. This is often the case in sport and entertainment events. Prices do in fact vary for sporting events and in hotels according to demand, but the evidence would suggest that they do not vary enough to stop excess demand at peak times. Kahneman, Knetsch and Thaler quote a skiing industry explanation: 'If you gouge them at Christmas, they won't be back in March.'

7.7 Contract theory

In the preceding section we explored ways in which Federica can incentivize Miguel to exert high effort. The problem that Federica is trying to solve in this setting is one of **moral hazard**. In any principal–agent relationship the incentives of the principal and agent are unlikely to be perfectly aligned. In other words, the agent will have an incentive to do things the principal does not like. For instance, Miguel has an incentive to put in less effort than Federica would like. Similarly, a car mechanic has an incentive to put in less effort fixing a customer's car than the owner would like. Your lecturer has an incentive to put less effort into marking your work than you would like. And CEOs have an incentive to take more risk than shareholders might like.

One way to try and alleviate moral hazard is to design contracts that more closely align the incentives of principal and agent. For example, if CEOs are given shares in the company then they have more incentive to look after the interests of shareholders. Similarly, Miguel could be offered some form of performance-related pay that means his wage more closely correlates with Federica's profit. The design of such employment contracts is part of a general field of game theory called **mechanism design**. Leonid Hurwiz, Eric Maskin and Roger Myerson received the Nobel Prize for Economics in 2007 for having laid the foundations of mechanism design theory.

Like other areas of game theory and economics, mechanism design has taken as its starting point the standard economic model assumptions of rationality and selfishness. So, we know a lot about contracts that would work for *Homo economicus*. There is an increasing recognition, though, that mechanism design needs to take account of the lessons from behavioral economics. There are different strands to this viewpoint.

Most obvious is a recognition that the optimal contract for *Homo economicus* might not work for *Homo sapiens*. We have already seen evidence of this in the chapter. For instance, in section 7.5 we saw that people may want to give to charity because of altruistic motives, and in section 7.6 we saw that reciprocity may be enough to incentivize Miguel to work hard. In both these cases the optimal mechanism would need to take account of intrinsic motives to do good things. Indeed, the last thing we want to do is come up with a mechanism that crowds out intrinsic motives. For instance, performance-related pay might work less well than simply paying Miguel a high wage and hoping for reciprocation.

A more nuanced point that comes out of a behavioral approach to mechanism design is the potential for **exploitative contracts**. Recall the example of gym membership in section 4.7, in which Maria could be exploited because her present bias meant she overestimated how much she would use the gym. In this case, the firm is designing a mechanism to exploit Maria. In Chapter 11 we will briefly look at another

example in which firms use contingent charges, such as a hotel charging for internet access or an airline charging for baggage, to make a good appear cheaper than it really is. Again, the contingent charges are a mechanism to exploit consumers. The interesting thing here is the potential rationale for government involvement to stop exploitation of workers and consumers.

What I want to do here is add to these examples by giving two more examples in which behavioral economics can help inform on optimal contract design. Both examples stick with the setting of Federica employing Miguel. In the first example we will look at what happens if Miguel is loss-averse. In the second example we will look at what happens if Miguel overestimates his ability.

7.7.1 Contracts for loss-averse workers

One of the more difficult issues in optimal contract design is that of designing incentives around risk. Indeed, this is the basic problem with performance-related pay. To illustrate, I will look at an example taken from Kőszegi (2014) that looks at how to eliminate moral hazard in a setting with loss aversion. Suppose that Federica wants to employ Miguel to do a job. As previously (see section 7.6.1), Miguel can exert either high effort or low effort. This time, however, output is stochastic. In particular, output is determined according to the probabilities detailed in Table 7.10. For instance, if Miguel exerts high effort, there is only a 1/3 probability output will be high. Crucially, Federica observes only output and not effort. This makes it difficult for her to appropriately reward Miguel, because Miguel could work hard and yet output be low, or Miguel could not work hard and output be medium. So, what contract is best?

Given that Federica observes only output, the best she can do is to condition reward on output. So, suppose she pays Miguel a basic wage of w and then gives a bonus of b_M if output is medium and a bonus of b_H if output is high. Our task is to work out the optimal combination of wage and bonus. Before we can do that we need a few more details. Suppose that it costs Miguel 20 to exert high effort and 18 to exert low effort. Also suppose that Miguel's utility is given by

$$u(x) = \begin{cases} x + (x - r) & \text{if } x \geq r \\ x - \lambda(r - x) & \text{if } x < r, \end{cases}$$

where x is his wage plus any bonus minus the cost of effort, r is his reference level and λ is a measure of his loss aversion. This form of utility function should be familiar from Chapters 2 and 3. The final thing we need to do is determine the reference level. I will assume that Miguel has a stochastic reference level (see section 3.3.4).

Table 7.10 The probability of low, medium or high output depending on the effort of Miguel.

Miguel's effort	Low output	Medium	High
High	1/3	1/3	1/3
Low	2/3	1/3	0

We can now work out Miguel's expected utility from choosing low or high effort. If he chooses low effort then there is a 2/3 chance output will be low and he will get material payoff $w - 18$, and there is a 1/3 chance output will be medium and he will get payoff $w + b_M - 18$. This means there is a 2/3 chance he will think of the reference level as $w - 18$ and a 1/3 chance he will think of the reference level as $w + b_M - 18$. So, if he exerts low effort and output is low, his reference-dependent utility will be $w - 18 - \lambda b_M/3$. In this instance, Miguel feels loss because he misses out on the bonus b_M. If he exerts low effort and output is medium then his reference-dependent utility will be $w + b_M - 18 + 2b_M/3$. In this instance, he gains from the bonus. His overall expected utility from choosing low effort is, therefore,

$$U(L) = \frac{2}{3}\left(w - 18 - \frac{1}{3}\lambda b_M\right) + \frac{1}{3}\left(w + b_M - 18 + \frac{2}{3}b_M\right) = w + \frac{1}{3}b_M - 18 - \frac{2}{9}(\lambda - 1)b_M.$$

We can go through similar calculations for Miguel's expected payoff if he chooses high effort, though things are slightly more complex, given that there are now three possible outcomes. You can work out that his overall expected utility is

$$U(H) = w + \frac{1}{3}b_M + \frac{1}{3}b_H - 20 - \frac{2}{9}(\lambda - 1)b_H.$$

The objective of Federica is to incentivize Miguel to exert high effort while also minimizing the wage and bonus that she pays. In order to incentivize Miguel, two conditions need to be satisfied. First, there is a **participation constraint** that $U(H) > 0$ or above the level Miguel can achieve with some other job. Second, the **incentive-compatibility constraint** is that $U(H) > U(L)$, and so Miguel will want to work hard.

The simplest case to consider is when $\lambda = 1$, and so Miguel is not loss-averse. In this case, Federica would want to minimize her expected payout, $w + b_M/3 + b_H/3$, subject to the requirement that Miguel takes up the job, $w + b_M/3 + b_H/3 \geq 20$, and works hard, $b_H/3 \geq 20 - 18$. One simple way of doing this is to set $w = 0$, $b_M = 0$ and $b_H = 60$. In other words, she offers a big bonus if output is high. This big bonus clearly provides a disincentive to exert low effort and is just enough to incentivize Miguel to take up the job in the first place. A simple example of performance-related pay! Would this contract work if Miguel is loss-averse?

Suppose that $\lambda = 2$, and so Miguel is loss-averse. The participation constraint now becomes $w + b_M/3 + b_H/9 \geq 20$. It is clear that just offering a high bonus of $b_H = 60$ is no longer going to work. This is because Miguel would feel the loss of not getting the bonus more than he would appreciate the gain if he did get it, and so it would be better for him to simply not take up the job. A different contract is therefore needed. Given Miguel's aversion to loss it is optimal for Federica to increase the wage and limit bonuses so as to minimize any potential losses. The optimal contract does, though, need to satisfy the incentive compatibility constraint of $b_H/9 + 2b_M/9 \geq 2$. This can just be satisfied by setting $b_H = b_M = 6$. In order to satisfy the participation constraint, she can then set $w = 17.33$.

The crucial thing to take from this example is the way in which loss aversion completely changes the type of contract that is going to work for Miguel. Without loss aversion, Federica can incentivize Miguel with the offer of a large bonus if output is high. With loss aversion, Federica does best to offer a high wage and relatively small bonus if output is high. It may not be very surprising that Federica needs to tailor the contract to the preferences of Miguel. If, however, a completely different type of contract is needed then it is going to be difficult for Federica to get this right. And, not only does she need to take account of loss aversion, she also needs to take account of things such as reciprocity.

7.7.2 Exploitation of an overconfident worker

The discussion so far has focused on Federica finding the optimal contract to incentivize Miguel. As of yet there is no sense in which Federica can exploit Miguel to increase her own profit. To give an example of how this may be possible, suppose that Miguel is **overconfident about his ability** (which we shall see in section 8.4 is quite likely). In particular, suppose that he believes that, if he were to put in high effort, there is a 1/3 chance of medium output and a 2/3 chance of high output. In this case, his expected utility from high effort is

$$U(H) = w + \frac{1}{3}b_M + \frac{2}{3}b_H - 20 - \frac{2}{9}(\lambda - 1)(b_H - b_M).$$

If Miguel is not loss-averse – i.e. $\lambda = 1$ – then it is enough to set $w = 0$, $b_M = 0$ and $b_H = 30$. In interpretation, Miguel's overconfidence that he will achieve high output means Federica needs to provide less incentive. Indeed, she can expect to pay $(60 - 30)/3 = 10$ less than if Miguel knew his true ability. Federica can, therefore, exploit Miguel's overconfidence by offering a bonus less than Miguel really needs.

Interestingly, if Miguel is loss-averse – i.e. $\lambda = 2$ – then Federica is less able to exploit his overconfidence. This is because she has to pay a high fixed wage to incentivize him, and so there is little to be gained from changing the bonus. Even so, she can benefit from his overconfidence. For instance, suppose she were to set $b_M = b_H = 20$ and $w = 0$. In Miguel's mind, this is a guaranteed payment of 20 if he exerts high effort, and so this contract will appeal. Federica can expect to pay only $20 \times 2/3 = 13.33$, however. This is a lot less than the $17.33 + 6 \times 2/3 = 21.33$ she would expect to pay if Miguel was not overconfident. Again, Federica can exploit Miguel.

This example further illustrates how the optimal contract can be hugely dependent on Miguel's preferences and beliefs. In this case, Miguel's overconfidence means that Federica can offer him a contract that he perceives as insuring him from any loss. So, Federica gains from taking account of both loss aversion and overconfidence. It is interesting to consider what would happen if Miguel took up the contract, worked hard and then produced low output! This raises a host of other issues, such as how he would react to this new information (see Chapter 5) and whether he will be angry and negatively reciprocate. Designing optimal contracts is clearly not a simple matter.

7.8 Further reading

The review articles by Thaler (1988), Camerer and Thaler (1995) and Fehr and Gächter (2000b) provide a gentle introduction to the economics of social preferences, and the article by Sobel (2005) a more in-depth survey. Much of the book by Camerer (2003) is also relevant. Interest in social preferences was sparked in large part by the ultimatum game experiments reported by Güth, Schmittberger and Schwarze (1982). See also the articles by Charness and Rabin (2002) and Binmore and Shaked (2010) and the review articles by van Damme and co-authors (2014) and Güth and Kocher (2014). For recent advances in psychological game theory, see Battigalli and Dufwenberg (2009) and Attanasi, Battigalli and Manzoni (2015). A survey comparing use of the strategy method with game method is provided by Brandts and Charness (2011). For interesting insights on public goods, a classic book by Olson (1971) is recommended, and I'll also recommend at this point a great book by Frank (1991). The papers by Frank, Gilovich and Regan (1993) and Levitt and List (2007) raise some interesting issues. See Andreoni (2006) and Andreoni and Payne (2013) for more on the economics of charitable giving, and Shang and Croson (2009) for an interesting paper that overviews relevant literature. Bewley (2007) provides a survey on fairness in wage setting. See also Fehr, Goette and Zehnder (2009). For more on behavioral contract theory, see the review article by Kőszegi (2014).

7.9 Online materials

On the companion website you will have access to flashcards, quizzes and answers to the review questions.

7.10 Review questions

7.1 What is the Nash equilibrium in a dictator game? What about the trust game and the ultimatum game? Be careful to distinguish sub-game perfect Nash equilibria from other Nash equilibria.

7.2 Why might payoffs be lower in a public good game with punishment even if punishment increases contributions? Is it good to have the threat of punishment or not?

7.3 In section 7.1.5 I mentioned a study by Andreoni and Miller (2002). They asked subjects to play 11 different dictator games. In each game the proposer was given a number of tokens, from 40 to 100, and asked to divide the tokens between them and a receiver. Each token was worth \$0.10 to \$0.40 to the proposer and \$0.10 to \$0.40 to the receiver. For example, a token might be worth \$0.40 to the proposer and \$0.10 to the receiver, or \$0.20 to the proposer and \$0.30 to the receiver. How do you think the share given to receivers depended on the worth of tokens? How should it depend, to be consistent with utility maximization?

7.4 Carefully explain the differences in incentives for the proposer in a dictator game, ultimatum game and dictator game with third-party punishment.

7.5 Should we equate material payoff with money?

7.6 Find the general conditions such that punishment can work in a linear public good game in the ERC model and the Fehr–Schmidt model. Hint: for the ERC model, find the level of punishment such that, after she has paid for punishing, Anna will have exactly the social reference point; for the Fehr–Schmidt model, find the level of punishment such that the free rider, Miguel, and punisher, Anna, will end up with the same material payoff.

7.7 Why can punishment work in a linear public good game according to the Fehr–Schmidt model if only one person is willing to punish, but in the ERC model only if two people are willing to punish?

7.8 How can an intentions-based model of fairness be adapted to take account of earned versus unearned payoff, and social norms?

7.9 In the intentions-based model of section 7.3.1, the equitable payoff was derived looking only at Pareto-efficient outcomes. Why is this sensible?

7.10 Is strong reciprocity a form of altruism?

7.11 What is the difference between psychological game theory and 'standard' game theory?

7.12 Is it better to be in a firm in which you earn $50,000 and the average salary is $80,000 or in a firm in which you earn $45,000 and the average salary is $30,000?

7.13 Why can loss aversion help explain why a cut in wages is worse than a rise in wages is good? How is this related to wage stickiness and habit formation?

7.14 Should employers fine poorly performing workers or give bonuses to well-performing workers in order to provide more effective incentives?

7.15 What can we learn from the ultimatum game about the interaction between a buyer and seller of a good?

7.16 Explain how we can distinguish warm glow giving from purely altruistic giving.

7.17 Suppose that someone donating to charity is told 'Most people give $10'. Why might this influence how much the person gives?

7.18 Discuss different reasons why performance-related pay may not incentive a worker as much as a flat salary.

7.19 There is some evidence (see Chapter 8) that men are more overconfident of ability than women. What consequences might this have in the labor market if employers offer optimal contracts?

Part III

Origins of behavior



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Chapter 8

Evolution and culture

> It is not the strongest of the species that survives, nor the most intelligent, but rather the one most adaptable to change.
>
> Clarence Darrow

In this chapter, and the next, our objective will shift away from trying to describe behavior to that of trying to understand the reasons for behavior. For example, we want to ask: why are people risk-averse? Why do people value fairness? Why is there reference dependence? Hopefully, you agree that these are fascinating questions. In answering them, we can not only learn why people behave the way they do but can also get additional insight into how they behave. Potentially, therefore, we can learn a lot.

In this chapter I am going to take seriously the fact that everyone's behavior is shaped by our evolutionary history and the culture that we grow up and live in. I will begin by looking at the consequences of evolution and then turn to the role of culture. Note that the section on culture will make a lot more sense if you are familiar with the material in Chapter 7.

8.1 Evolution and economic behavior

I shall not go deeply into the science of evolution, but it will be useful to know the basic idea. Evolution by natural selection relies on three things: **variation**, meaning that different individuals are different in their traits, abilities, characteristics, preferences, and so on; **inheritance**, meaning that offspring inherit similar traits to their parents; and **mutation**, meaning that with random chance the traits of the individual are altered during reproduction and are slightly different from the traits of their parents. Over generations, the consequences of **natural selection** are twofold: (1) traits that result in an individual having somewhat more offspring than others will become more common in the population, because of variation and inheritance; and (2) traits that increase the relative number of offspring of an individual can emerge in the population, because of mutation and inheritance.

To illustrate, consider a very stylized example in which there are two types of people, some having green hair and some blue. Those with green hair find it easier to attack prey and catch food, so survive longer and have more children. On average, greens have four offspring each and blues have two. Suppose that there are ten greens

and ten blues to start with. After one generation there are 40 greens and 20 blues; after two generations there are 160 greens and 40 blues; after three generations there are 640 greens and 80 blues. Clearly, the greens are becoming a larger proportion of the population, from 50 percent to 88 percent in three generations.

Now, suppose that because of some random mutation one person is born with yellow hair and another with red. A person with yellow hair finds it even easier to attack prey, and so has on average six offspring, while someone with red hair is very poor at attacking prey, and so has only one offspring. If you do the math you will find that it is not many generations before those with yellow hair dominate the population.

This simple example avoids many complexities that we need to think about, but it will do for the moment, and I shall provide a more nuanced picture as we proceed. The important thing at this stage is to start explaining what evolution has got to do with modern economic behavior.

8.1.1 Looking for food and finding a utility function

Food and water are essential to survival, and much time would have been spent in our evolutionary history trying to find enough of them. Indeed, finding food is still the main preoccupation of many in the developing world today. Evolution should, therefore, favor those with good strategies for finding food and water. Can this shed light on how humans behave today?

Consider a person called Edward, looking for food. Every day he leaves his cave and can go either east or west in search of food. Suppose that on day t, if he were to go west, he would find f_t^W calories of food, and, if he were to go east, he would find f_t^E. The more calories he finds, the better. The problem is, Edward does not know f_t^W or f_t^E without going and looking, and he can go only one way per day. Which way should he go?

Let's suppose that the amount of food in either direction on any one day is random. The probability of finding f calories by going east is p_f^E and by going west is p_f^W. If

$$\sum_{f=0}^{\infty} p_f^E f > \sum_{f=0}^{\infty} p_f^W f \tag{8.1}$$

then Edward will maximize his calorie intake by going east. Remember, however, that having children is the objective favored by natural selection, not eating a lot. Let $u(f)$ denote the expected number of extra children Edward (and his wife, Anna) will have if Edward gets f calories of food on any one day. We do not need $u(f)$ to be an integer, and, given that Edward will have far fewer children than meals, $u(f)$ will likely be a small number. If

$$\sum_{f=0}^{\infty} p_f^E u(f) > \sum_{f=0}^{\infty} p_f^W u(f) \tag{8.2}$$

then Edward will maximize the expected number of children by going east.

Look at the difference between equations (8.1) and (8.2). In the first we calculate the expected number of calories, in the second the expected number of children. It may well be that going west gives the higher expected calories, but going east the higher expected number of children. For instance, the amount of food to the east may always be a nice consistent amount, which is good for bringing up a family, while that to the west is either very high or very low, which yields a higher average but also periods with too little food. Evolution will favor those who maximize the expected number of children, and so obey equation (8.2) rather than (8.1).

This means that Edward not only needs to work out how much food there is likely to be in each direction, he also needs to work out the relationship between the amount of food he eats and the number of children he will likely have! Given that Edward will have relatively few children, this looks like a hopeless task. Suppose, however, that Edward is born with the function u already hard-wired into his brain. What I mean is that, when he sees and eats foods with a calorie content of f, he gets an enjoyment, satisfaction or utility of $u(f)$. Does Edward's problem become simpler?

Robson (2001b) shows that it does. Edward can now maximize his expected number of children with the following heuristic: 'Today I go east, tomorrow I go west, the next day I go east, and so on; each day I keep a cumulative count of the difference in utility between going east versus west; if this difference reaches some critical threshold then from that moment on I always go east or west, depending on which gives the higher utility.' This is a simple solution to a complex problem, and it works because Edward is born with the utility function $u(f)$ and the heuristic of how to use it.

The utility function has two key benefits here. The first is that all Edward has to think about is food. If he finds enough food to maximize his utility then, without knowing it, he will maximize his expected number of children. This is because evolution will favor those who have the correct utility function – i.e. get satisfaction from eating food proportional to the expected number of children that eating food provides. The utility function $u(f)$ thus encapsulates the **accumulated experience of past generations**.

The second benefit of the utility function is that Edward can adapt to novelty. Over generations it will likely change whether going east or west is better. It would, therefore, not be a good evolutionary strategy to follow the actions of one's ancestors, and, say, go west because a parent did. The utility function allows Edward to follow the strategies of his ancestors, as given by heuristics. Again, evolution will favor those with the best heuristic, meaning that Edward, and his children, will be very **adaptable to a changing environment**.

We have already seen in Chapter 3 the distinction between expected utility, equations (8.2) and (3.1), and expected value, equations (8.1) and (3.2). Now we have an explanation for it, and a justification to assume that people behave as if maximizing an expected utility function. Indeed, we see that evolution may have favored people with a cardinal utility function whereby satisfaction from goods comes in proportion to the expected number of children such goods give. We call such goods **intermediate goods**, because they are anything that increases, or decreases, the expected number of children. Food and shelter are obvious examples of intermediate goods, but status, generosity, esteem and a host of other things may be equally relevant, as we shall see.

Is this useful in understanding modern economic behavior? Yes. It gives us important clues as to what should give us utility and what probably will not. It also gives us clues as to what we should have good heuristics for, search being one of them. What I want to do now is add time and risk to the mix.

8.1.2 Choosing when to have children

I will start by adding time to the story, by questioning when Edward (and Anna) should choose to have children. Let's suppose that Edward can have children up to T years of age. What he needs to do is to decide when to have children during those T years. In our story, this choice is primarily determined by when he chooses to consume the intermediate goods such as food and shelter and not necessarily when he gets together with Anna. We represent his choice by a **reproductive profile** $x = \{x_1, \ldots, x_T\}$ that says how many children, x_t, he expects to have when age t. If all people, including Edward and his children, survive from one year to the next with probability s, what reproductive profile should Edward choose?

Again, this does not look a simple problem. It turns out, however, that, if Edward has an inter-temporal utility function hard-wired into his brain, he can solve this problem relatively easily. As before, he will not need to choose when to have children but, instead, just consume intermediate goods in a way that maximizes his utility (and get together with Anna every so often). But what will the inter-temporal utility function look like? Robson and Samuelson (2009) show that evolution will favor people with the inter-temporal utility function

$$u^T(x_1, \ldots, x_T) = \sum_{t=1}^{T} \delta^{t-1} x_t, \tag{8.3}$$

where

$$\delta = e^{-(\ln \varphi - \ln s)}$$

and $\ln \phi$ is the population growth rate. The term $\ln \phi - \ln s$ is equal to the population growth rate plus the population death rate. You might want to refer back to Chapter 4 to see in more detail how to interpret equation (8.3). But, basically, it means that evolution will favor people who use exponential discounting and have consistent, non-time-varying preferences with a discount factor that depends on the population growth rate and death rate.

This gives us an explanation for exponential discounting and a prediction of what the discount factor should be. But why this specific discount factor? If Edward delays having children one year then he can expect to fall behind others because of population growth and the possibility that he does not survive the year. So, a child next year needs to be discounted by the rate of population growth and death. The higher is population growth and the lower the chance of survival, the more impatient Edward should be to consume intermediate goods and have children.

[Extra] For those curious how to derive equation (8.3), here are some of the details. Starting when Edward is zero years old, we can roll forward time and keep

track of how many descendants he has of each age. Let $N_a(t)$ denote the number of descendants of age a after t years. Things will evolve according to the equation

$$[N_1(t+1), ..., N_T(t+1)] = [N_1(t), ..., N_T(t)] \begin{bmatrix} sx_1 & s & 0 & \dots & 0 \\ sx_2 & 0 & s & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ sx_{T-1} & 0 & 0 & \cdots & s \\ sx_T & 0 & 0 & \cdots & 0 \end{bmatrix}. \quad (8.4)$$

To explain, the number of one-year-olds after $t + 1$ years is equal to the number of surviving offspring born to those who were one in year t – i.e. $N_1(t)\ sx_1$ – plus the number born to those who were two in year t – i.e. $N_2(t)\ sx_2$ – etc. The number of two-year-olds in year $t + 1$ is equal to the number of one-year-olds who survive from year t to $t + 1$ – i.e. $N_1(t)\ s$ – etc. The matrix in equation (8.4) is called the Leslie matrix, and is crucial to finding what reproductive profile evolution will favor. Technically, evolution will favor the profile that maximizes the dominant eigenvalue of the Leslie matrix. What this means in practice is that evolution will favor the utility function in equation (8.3).

8.1.3 Aggregate risk

So far we have described Edward as maximizing expected utility using exponential discounting. In Chapters 3 and 4 we saw that people may not maximize expected utility or use exponential discounting, so how can we reconcile things a little? One way is to add aggregate risk to the story. To explain: an **idiosyncratic risk** is one that affects only one individual in the population, while an **aggregate risk** is one that will likely affect all the population. For example, a cold winter or flood is an aggregate risk, while getting attacked by a predator is more likely to be idiosyncratic. I have focused on idiosyncratic risk so far, but in our evolutionary past aggregate risk was undoubtedly relevant. With the help of four examples, I shall illustrate that it likely had important consequences.

To guide us through the examples, I shall immediately preview the main insights, and also point towards Table 8.1. In the presence of aggregate risk we find that evolution may favor people who (1) diversify, (2) avoid aggregate risk, (3) have a relative high discount rate, and so are impatient, and (4) have a discount rate that decreases with age consistent with decreasing impatience. Clearly, aggregate risk makes a difference. Most interestingly, it offers an explanation for both risk aversion and hyperbolic discounting. So, let's work through the four examples.

In the first example, think about people collecting enough food to survive winter. Winters are either long or short, and people are of two types: 'Long' types, who store enough food to survive any winter, and 'Short' types, who store enough to survive a short winter but not enough to survive a long winter. Because of the extra time spent collecting food, 'Long' types have a 0.5 chance of surviving the summer while 'Short' types have a 0.75 chance. What will happen? Despite being better at surviving summers, the 'Short' types are doomed to extinction because they will all die in the first long winter! Evolution, thus, favors 'Long' types.

Table 8.1 The four examples involving aggregate risk.

Example	Brief description	Implication
Collecting food for winter	The extended family of types who randomly determine how much food to collect will be most able to prosper in good years and survive bad years	Risk aversion and desire to diversify
Living together or separately	It's better they live separately; if they live together there is a small probability of the population growing very large but a large probability that the population will end up small	Should avoid aggregate risk
Surviving the winter	Aggregate risk should not affect how much people discount the future, because it is relative individual differences that matter, and aggregate risk affects everyone equally	The discount rate will be relatively high, suggesting impatience
Something happens to 24-year-olds	Before they are 24 people should be relatively impatient, because the aggregate shock at 24 lowers the expected growth rate of the population	Hyperbolic discounting, and increasing patience with age

Consider, now, a new type of person, who at the start of each summer randomly decides whether to store enough food for a long winter or short winter. The extended family of such 'Mixed' types stand to do well, because in any year some of them will collect enough for a long winter and some will not. That some of them do not collect enough for a long winter means the family can grow at a relatively quick rate. That some of them do collect enough for a long winter means the family will never be wiped out by one long winter. It turns out that evolution will favor people who store enough for a long winter with probability $3p$, where p is the probability of a long winter (see Review question 8.2).

This example shows that, in the presence of aggregate risk, the optimal strategy may involve randomizing. This might look like an argument for risk-loving behavior, but actually it is more of an argument for risk aversion and diversification, because it shows that risks, such as a long winter, should be guarded against.

Another way to interpret this first example is that, in the presence of aggregate risk, it may make sense to have idiosyncratic gambling. The next example takes this point one stage further by showing that idiosyncratic risk is preferred to aggregate risk.

A group of people are now deciding where to live, and have two options. They could all live together in one big village or live separately. Every year each person has two children, and, no matter where they live, there is a 0.5 chance that one child will not survive the winter because, say, there is an attack by a rival tribe. The difference is that, if everyone lives together, this risk is aggregate, so there is a 0.5 chance that everyone in the group will have only one surviving child. If they all live separately, the risk is idiosyncratic, so each person independently has a 0.5 chance of having only one surviving child. Should they live together or separately?

It turns out that they should live separately (see below). The general lesson to take from this is that evolution will favor people who avoid aggregate risk. To give some intuition for why this is the case: after a few years, aggregate risk leads to a small probability of there being a very large population but a large probability of there being a very small population. It is best to avoid such a gamble, and people do that by

avoiding aggregate risk. Interestingly, however, a person by himself may not be able to avoid aggregate risk because it depends on what others do. For example, Edward and Anna cannot live alone if others follow them. This creates interdependence between individuals, which we shall explore a bit more later on.

[Extra] Why should they live separately? For simplicity, suppose that no adult ever survives the winter. In both cases the expected population size by year t is $N(t) = 1.5^t N(0)$, where $N(0)$ is the original population size. In terms of expected utility, therefore, both options are equivalent. In an evolutionary sense, however, they are different. If they live separately, the number of people will be very close to $N(t) = 1.5^t N(0)$, because the risk is idiosyncratic. The growth rate of the population will, therefore, be

$$\frac{\ln N(t) - \ln N(0)}{t} = \ln 1.5.$$

If everyone lives together, the population size will be $N(t) = 2^{b(t)}N(0)$, where $b(t)$ is the number of years in which all children survived. The growth rate now is

$$\frac{\ln N(t) - \ln N(0)}{t} = \frac{b(t)}{t}\ln 2 \to \frac{1}{2}\ln 2 = \ln\sqrt{2}.$$

The growth rate if they live separately exceeds that if they live together.

In the final two examples I want to return to the issue of when Edward should have children. First, imagine that the length of winter can be anything from $l = 100$ to 50 days long. The longer it is, the less chance that a person survives the winter, with the probability of survival $s(l) \in (0,1)$. That the probability of survival, $s(l)$, depends on the random length of the winter implies that there is aggregate risk. That each person survives independently with probability $s(l)$ means that there is also idiosyncratic risk. Does the aggregate risk make any difference to how Edward should discount the future?

Robson and Samuelson (2009) show that Edward should maximize expected utility, discounting exponentially, with a discount rate equal to $\ln \phi - \ln s$, where $\ln \phi$ and $-\ln s$ are the rate of population growth and death rate that there would be without aggregate risk. This sounds like equation (8.3), but there is a difference. Aggregate risk, for the same reasons as in the previous example, will lower the actual population growth rate below what it would be without aggregate risk. Some numbers will help illustrate.

Suppose that in 'normal' conditions a population grows at 2 percent a year, with a death rate of 4 percent a year – numbers plausible in human evolution. Also suppose that every so often a catastrophe occurs that decimates the population and means that, in the long run, population growth is around zero a year, which is also plausible in human evolution. Evolution will favor individuals who ignore the aggregate risk of the catastrophe and have a discount rate of 2 + 4 = 6 percent. If we had gone off equation (8.3), we would have expected a discount rate of 2 + 0 = 2 percent.

The catastrophe proves irrelevant, in an evolutionary sense, because it is an aggregate shock that affects all in the same way. This means that it will not alter the number of children that Edward has relative to another, and so he can ignore it. This

makes people relatively impatient, because they should discount at, say, 6 percent rather than 2 percent a year. It also means that we should not be surprised to see higher discount rates than might seem sensible, given observed population growth and death rates.

In the final example, suppose that, for some reason, the average temperature has an effect on the survival rate of people who are 24 years old. To model this, we say that the probability of survival is $s(l) + d$ (temp.) for anyone aged 24, for some function d(temp.), while the probability of survival for anyone not aged 24 remains $s(l)$. For instance, if $d(5) > 0$ and $d(-5) < 0$, someone aged 24 has a higher probability of survival than others if the temperature is 5°, and a lower probability if it is –5°. This creates an aggregate shock that affects all 24-year-olds.

If Edward is not yet 24, the possibility of an aggregate shock when he is 24 should reduce the value he puts on children he might have after he is 24. This is because aggregate risk (as we have seen) lowers the expected growth rate of the population. Edward should, therefore, be relatively impatient to consume intermediate goods and have children just before he is 24. Intriguingly, this is the case even if the aggregate shock at age 24 makes it is more likely that he will survive when he is 24 than at any other age!

This shows that, with age-specific aggregate risk, the discount rate will change with age. So, evolution need not necessarily favor individuals who use exponential discounting. In fact, provided the age-specific effects are relatively small, the best thing is hyperbolic discounting, whereby the discount rate decreases with age. Thus, evolution can help explain why we see the types of time preferences discussed in Chapter 4, though we have yet to see any explanation for time inconsistency.

8.1.4 Competing with others

The picture I have given so far is one of evolution as a struggle between Edward and the environment. This is only half the story, however, because evolution is also a struggle between different people. A further two examples will help illustrate why this matters. Table 8.2 summarizes.

Suppose that some people are cooperative and some are non-cooperative. People randomly meet each other in pairs and collect food. Two cooperative people go together and share any food they collect. A non-cooperative person

Table 8.2 The two examples with interdependence.

Example	Brief description	Implication
Cooperators and non-cooperators	If cooperators can find some way to interact more with each other than non-cooperators then they can prosper; the result is a signaling game of hide-and-seek, in which cooperators try to distinguish themselves and non-cooperators try to copy them	Seemingly irrelevant things can become important evolutionary signals
Different ability at finding food	Women should prefer to mate with men who are better at finding food, which means that men need to appear good at finding food	People should be good at signaling their desirable attributes

Table 8.3 The payoffs if people are cooperative or non-cooperative.

		Person 2	
		Cooperative	**Non-cooperative**
Person 1	Cooperative	1, 1	0, 2
	Non-cooperative	2, 0	0.5, 0.5

collects food with a cooperative person but then takes all the food. Finally, two non-cooperative people go off and steal the food from cooperative people! Payoffs are summarized in Table 8.3 (and see Research methods 8.1 for how to interpret these payoffs).

❖ RESEARCH METHODS 8.1 EVOLUTIONARY GAME THEORY

Once we recognize the importance of interactions between individuals it is natural to turn to game theory. Game theory has developed tools specific for modeling evolution, which we can illustrate with a hawk–dove game introduced to model animal conflict.

Consider two animals competing for food with a value of *V.* Each animal can either attack or share. If they share, each gets *V*/2. If one attacks and the other shares, the attacker takes all the food. If both attack, there is a conflict that costs 0.5C, and each has an equal chance of winning the food. The payoff matrix is given in Table 8.4. To put this in an evolutionary context, we think of the strategy, attack or share, as a genetic trait or **phenotype** that is constant through the animal's lifetime and inherited by offspring. We also think of the payoffs as representing the net additional offspring that the animal will get from consuming the food (in a similar way to section 8.1.1).

Imagine that the population was full of animals who share. If a mutation produces an animal with the attack strategy, this animal would have relatively many offspring. The strategy 'Attack' would, therefore, spread. Now imagine that the population was full of animals that attack. If $V < C$, an animal that shares would have relatively many offspring, and the strategy share would spread. The population will be stable if the payoff from attacking is the same as sharing. If proportion p^* of animals attack, this requires

$$\left(\frac{V-C}{2}\right)p^* + (1-p^*)V = (1-p^*)\frac{V}{2}.$$

This simplifies to $p^* = V/C$. This is known as the evolutionarily stable strategy, and means that we should expect to see a mix of animals that attack and share.

Table 8.4 Payoffs in the hawk–dove game.

		Animal 2	
		Attack	**Share**
Animal 1	Attack	$(V - C)/2, (V - C)/2$	$V, 0$
	Share	$0, V$	$V/2, V/2$

More generally, an **evolutionary stable strategy** details the proportion of the population that should be of each type if the population is to remain stable over time. Typically, we would expect the population to converge to an evolutionary stable strategy, and so these strategies tell us a lot about the likely consequences of evolution.

I will briefly note that there is an important, and intriguing, connection between an evolutionary stable strategy and symmetric Nash equilibrium. The connection is as follows: if p^* is an evolutionary stable strategy, then 'both animals choose p^*' is a Nash equilibrium; the converse need not hold. The notion of evolutionary stable strategy offers, therefore, a form of equilibrium refinement (see Research methods 6.3). This is intriguing, because the two concepts are subtly different, and have different historical routes. Recall: a Nash equilibrium is such that anyone who deviates gets a lower payoff than if he did not deviate. An evolutionary stable strategy, by contrast, is such that anyone who deviates gets a lower payoff than those who do not deviate.

Unfortunately, evolution will favor non-cooperative people, because they always have a higher payoff than cooperative people. Suppose, however, that cooperative people start to grow longer noses than non-cooperative people, meaning that people can tell who is cooperative and who is not. Now the odds swing back in favor of the people who are cooperative, because they can make sure they interact only with each other. This means that they can get a payoff of one, while non-cooperative people only get a payoff of 0.5. Evolution will now favor cooperative people. Unfortunately, though, the non-cooperative people might fight back by growing noses as long as the cooperative people.

This example illustrates how evolution can result in signaling games of hide-and-seek (not unlike that which we looked at in the last three chapters). Evolution will favor cooperators who can signal that they are cooperators. But it will also favor non-cooperators who can look like cooperators. This game of hide-and-seek can mean that seemingly irrelevant traits, such as the length of noses, can become an important battleground of evolution.

What's going on here is something called assortative matching. We say that there is **assortative matching** if it is more likely that certain types will meet than others. If a cooperator is more likely to interact with a cooperator, a non-cooperator is likely to interact with a cooperator, and then there is assortative matching. This is crucial if cooperators are not to be exploited by non-cooperators, and so assortative matching is crucial if evolution is to favor cooperative people.

We see interesting evidence of assortative matching in human interaction, the 'lie detector' being one example. Experimental economists and psychologists have also shown that individuals have varying abilities to predict what an opponent will do in a game even if they only meet or see that person for a few seconds. For example, Pradel, Euler and Fetchenhauer (2009) did a study in which schoolchildren, aged ten to 19, were asked to predict what their classmates would do in a dictator game (see Chapter 7 for more on this game). As we would expect, there was a split of children offering nothing (8 percent of children), half the amount (49 percent) or something in between.

❖ RESEARCH METHODS 8.2 EXPERIMENTING WITH CHILDREN

The study by Pradel, Euler and Fetchenhauer involved children aged 10 to 19 as subjects, and I am going to talk later about two more studies that involved children. So, psychologists and behavioral economists do sometimes use children as subjects. This can be for convenience, because a school might provide a large group of willing volunteers. For the most part, though, it's because looking at how children behave can provide insight into the causes of behavior.

For instance, if we get children to play the ultimatum game and find they very rarely reject offers, we can say that the desire to reject seems to develop later in life. Indeed, this might suggest that it is a cultural thing (though care is needed, because things can develop later in life for genetic reasons, breasts being a fairly obvious example). We shall come back to such issues in section 8.2.4. Ideally, we would like to experiment with very young children to see how economic behavior develops with age. That poses some challenges, though. For example, young children need an incentive other than money and can more easily lose attention.

The most important thing for us is that the children were good, better than chance, at predicting what others would do, and were particularly good at knowing who would give nothing. This is illustrated in Figure 8.1, where we see that children were much better than chance at predicting the giving of friends and those they disliked. Furthermore, this was not just due to friends being like them, but a genuine ability to predict giving.

If we are able to predict what others will do then assortative matching becomes very relevant in understanding human cooperation. On a related note, we also find that evolution can lead to a desire for Edward to signal his good attributes. To illustrate why, suppose that males differ in their skill at finding food. If mating is random then a male will collect food until the marginal benefit equals the marginal cost. This means that less skilled males will collect less food. Evolution would then favor females who are fussy who they mate with, however, and search out those more skilled at finding food.

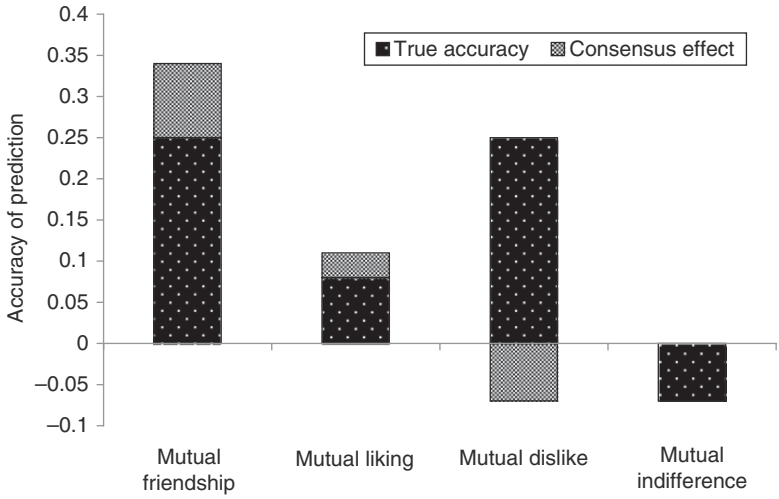

Figure 8.1 Accuracy of children's ability to predict the giving of fellow students in a dictator game, above that of guessing randomly. High accuracy could be because children expected those they like more to behave more like them, a consensus effect. Even taking this into account, however, children were good at predicting the giving of those they knew.
Source: Pradel, Euler and Fetchenhauer (2009).

If females prefer to mate with males who are skilled in finding food then matching will no longer be random. Moreover, males have an incentive to signal that they are skilled at finding food, even if they are not highly skilled. They can do this by collecting lots and lots of food. Indeed, consider the most extreme case, when females will mate only with the male who collects the most food. This turns things into an auction in which the highest bidder wins. We can expect a lot of food!

In the preceding example, evolution will favor those who are good at signaling their ability to find food. More generally, we see another important evolutionary struggle, sometimes called **sexual selection**, between men and women. Evolution will favor people who are fussy who they mate with and look for the partner most likely to have healthy, well-looked-after children. This means that relative traits such as strength, health and faithfulness become important. It also means that a person's ability to signal their desirable attributes becomes important. This provides an explanation as to why status and conspicuous consumption may be valid intermediate goods, because those with higher status, who use conspicuous consumption to signal their status, may look more attractive to the opposite sex.

We have now started to touch on situations in which culture will matter. This is because we have moved away from talking about intermediate goods, such as food, that objectively increase the probability of having children, to goods, such as stylish hair or a new sports car, that increase the probability of having children if and only if they impress others. In different cultures, different things might impress: in one place, stylish hair, and, in another, stylish clothes. So, now seems a good point to start adding culture to the mix.

8.2 Culture and multi-level selection

For the most part, humans are very similar the world over. We all smile, laugh, cry and get angry in similar ways over similar things. This is a consequence of evolution. Behavior is not, however, solely determined by our genes. It is also shaped by the culture that we live in. Elsewhere in the book I have looked in some detail at how people can be influenced by what they see others doing and saying, and so I will not spend too much time here showing why culture matters. Instead, I want to focus on how culture can interact in interesting ways with evolution. To put this into some context, it is useful to start by seeing if there are cross-cultural differences.

8.2.1 Cross-culture comparisons

I am going to look at three cross-culture comparisons. In each we compare how people behave in different cultures to see if there are systematic differences. You may want to look back at the previous chapter for an explanation of the games we shall look at.

I'll start with a fascinating study by a group of economists and anthropologists, in which the **ultimatum game** was played with members of 15 small-scale societies (Henrich and co-authors 2004). This can be coupled with a study by Roth and co-authors (1991), and the hundreds of other ultimatum game experiments performed around the world with university students, to give a picture of how behavior varies across cultures. Figure 8.2 shows proposer offers in a variety of different places.

In all cases we see that proposers give something, and so the picture we gave in Chapter 7 still holds true. There is clearly a large variation, however, in how much people give. To see why, it's useful to look in a bit more detail at some of the societies.

The most giving was observed in Lamalera. This is a village located on an island in Indonesia where whale hunting is integral to village life. Whale hunting can be done only in groups, and so cooperation and sharing are a big part of village life.

At the other extreme is the Machiguenga of the Peruvian Amazon. Here we observe the least giving. The Machiguenga have traditionally lived in single-family units, and so sharing and cooperation outside the family are rare.

The Achuar and Quichua are two ethnic/political groups in Conambo in the Ecuadorian Amazon. To look at the data, it would seem that the Quichua give less than the Achuar. It turns out, however, that the differences can largely be put down to politics. High-status males in Conambo give meat in order to maintain strong alliances, and a strong correlation was observed between meat giving and offers in the ultimatum game. The relatively low status of the Quichua men can, therefore, explain the relatively low offers.

Finally, Au and Gnau are the languages spoken in two villages located near to each other in Papua New Guinea. The results in these villages were noticeable for the large rejection rates. The overall rejection rate was 32.8 percent, and three offers of a 0.7 share were rejected! When gifts are given by the Au and Gnau, it is understood by both the giver and receiver that the acceptor has incurred a debt that must be repaid at some future point. It can, therefore, be wise to not always accept a gift.

This gives only a snapshot of the differences across the different societies, but it does illustrate how behavior was influenced by everyday experience. People more

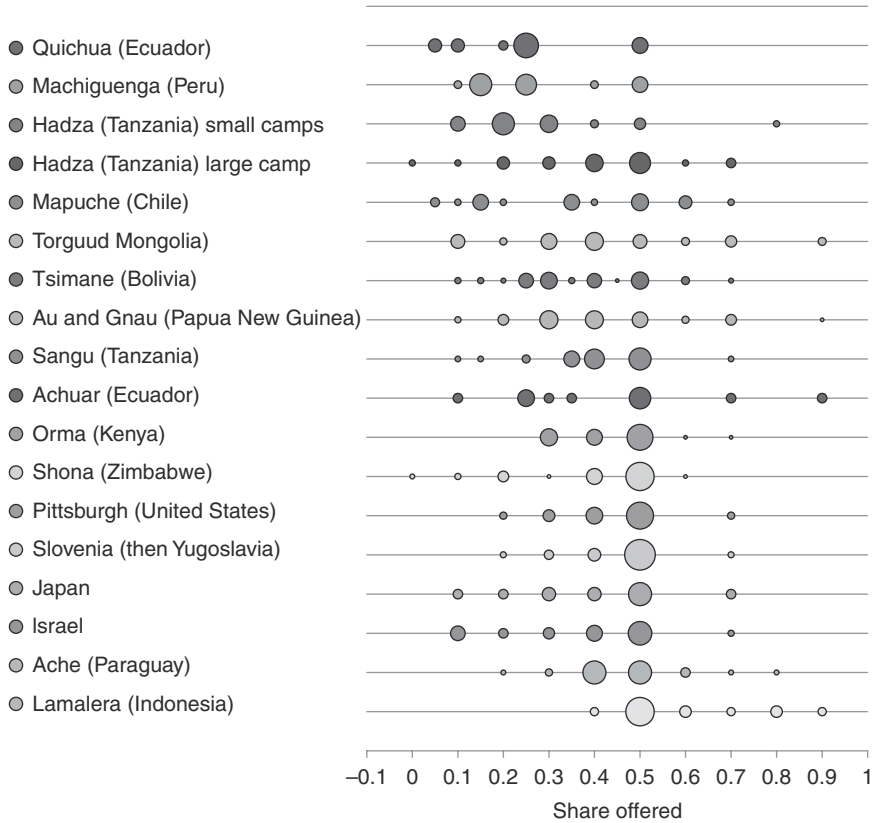

Figure 8.2 Ultimatum game offers across different cultures.
Sources: Henrich *et al.* (2004) and Roth *et al.* (1991).

used to sharing shared more, and people not used to sharing shared less, and so on. Henrich and co-authors take this one stage further by constructing two indices for each of the 15 societies they considered: payoffs to cooperation, which depends on local ecology; and aggregate market integration, which depends on things such as social complexity and settlement size. There was a strong correlation between payoffs to cooperation and aggregate market integration and offers in the ultimatum game. Those more used to cooperating, sharing and interacting with others offered more.

❖ RESEARCH METHODS 8.3 STATISTICS AND LIES: MORE, LESS, OR JUST DIFFERENT?

Comparing behavior in the ultimatum game across countries provides a good opportunity to illustrate what data and statistical tests tell us and what they do not. Figure 8.2 plots the distribution of offers. Why not just plot the mean,

mode or median offers instead? This does not give such a complete picture of the data and can mislead. For example, in the Ache, the modal offer was 0.4, which is less than the 0.5 in most other societies, including the Mapuche, but the mean offer was higher than in many other societies, including the Mapuche.

While it is tempting to want to say that giving is higher in one society than in another, in many cases this is simply impossible to do. For example, were offers higher in the Achuar, or Ache or Sangu? There is no way to answer that question. This also means that care is needed in interpreting statistical tests. Many tests, such as the Mann–Whitney, have a null hypothesis that two samples are drawn from the same distribution and an alternative, in this instance, that one sample stochastically dominates another. The primary thing learnt from a rejected test is that the samples are drawn from different distributions. A little care, therefore, is needed in interpreting the results of such tests.

A more recent study by Gächter, Herrmann and Thöni (2010) provides a cross-cultural comparison of behavior in a **public good game with punishment**.

The most interesting thing to come out of this was differences in the use of punishment. Figure 8.3 plots the frequency of punishment relative to whether the subject contributed more than the average, the same as the average or less than the average by varying amounts.

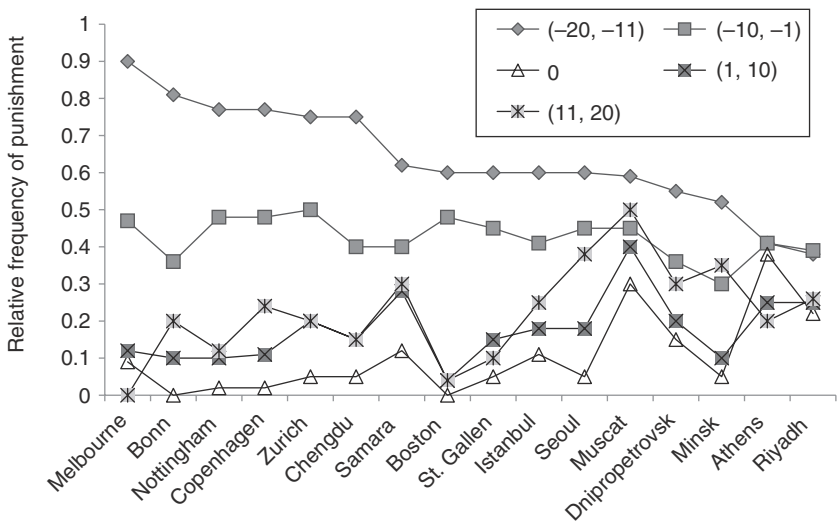

Figure 8.3 Punishment in a public good game in different locations. The proportion of times people were punished depended on whether they contributed a lot less (−20, −11), less (−10, −1), the same (0), more (1, 10) or a lot more (11, 20) than the person punishing.
Source: Herrmann, Thöni and Gächter (2008).

As we might expect, those who contributed less were punished, and the frequency of punishment did not vary too much across the locations. More surprising was the frequency of punishment for those who contributed above average. In places such as Boston, Melbourne and Nottingham, there was little punishment of those contributing above average. In Muscat, Athens and Riyadh, by contrast, those giving above average were punished as much as those giving below average. The best explanation for this is **antisocial punishment**, as a form of revenge. In particular, there was a strong correlation between punishment received in one period and that given out in the next. So, those punished for giving less seemingly tried to punish those who had punished them!

The consequence of antisocial punishment was less cooperation. For example, mean contributions were 18 out of 20 in Boston and 17.7 in Copenhagen, but only 5.7 in Athens and 6.9 in Riyadh. The good news of the previous chapter, therefore, namely that a threat of punishment can increase cooperation, needs to be toned down a little. The threat of punishment does not seem to work everywhere. So, why were there such differences in punishment?

One explanation is the strength of law and norms of civic cooperation in society. The stronger the rule of law, the less people might think that they need to rely on revenge, and, the stronger the norms of cooperation, the less people will want to punish those behaving in a socially desirable way. Consistent with this, Gächter, Herrmann and Thöni show that punishment of those who contributed less is positively correlated with stronger norms of civic cooperation, while punishment of those who contributed more is negatively correlated with the strength of both law and norms of civic cooperation. Again, we see local institutions and norms making a difference.

The final cross-cultural comparison I want to look at is that of **trust**. The World Values Survey is a large-scale survey of attitudes across a very broad selection of countries. The survey was first conducted in 1981 and it has been done several times since. The survey is a fantastic source of information on how values and attitudes change across time and location. One question asked is: 'Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?' The possible answers are 'Most people can be trusted' and 'Can't be too careful'. Figure 8.4 details by country what proportion thought most people can be trusted in the 1981 to 2000 data.

There is, clearly, large variation in what proportion say they can trust. Trust is relatively low in eastern Europe, South America and Africa, and relatively high in Scandinavia. But does this matter? In the previous chapter we saw that trust can be beneficial in simple experimental games, and it appears that this does transfer over to the real economy. Knack and Keefer (1997), for example, using data up to 1991, find that a 10 percent rise in the proportion who say they can trust is associated with a 0.8 percentage point increase in economic growth. This would suggest that the level of trust is potentially very important for the economy.

So, what lessons can we take from these three cross-cultural comparisons? There seem to be two main ones. First, there is significant variation in how people behave across societies and cultures. In some places we see, on average, more giving, more cooperation and more trust than others. Second, this variation appears to come in large part from variations in institutions and norms. This does create a kind of chicken and egg problem of what comes first. Does behavior influence the norms

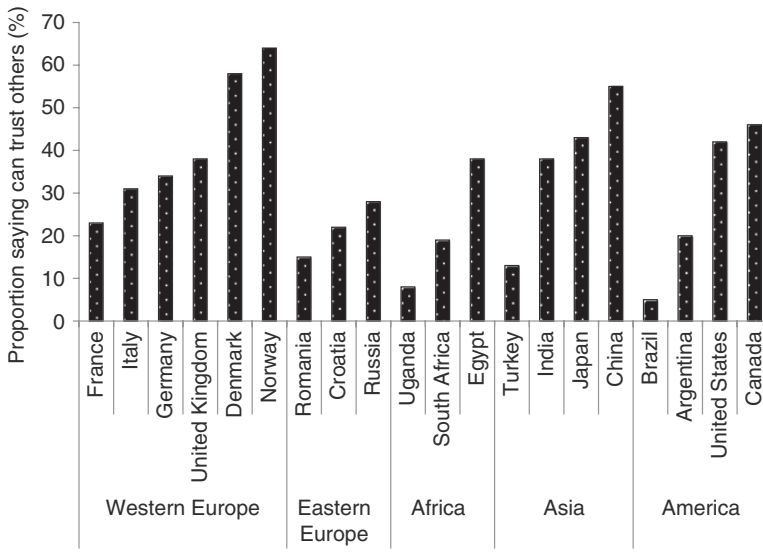

Figure 8.4 The proportion of respondents who said 'Most people can be trusted'.
Source: World Values Survey 1981–2008 database, v.20090901 (www.worldvaluessurvey.org).

and institutions that emerge, or do the institutions and norms that emerge influence behavior?

There are good reasons to think that it is at least partly the latter. The catalyst can be the environment that people live in. Some environments require cooperation for people to survive, such as that of the Lamalera, while some environments allow people to be largely self-sufficient. The environment could, therefore, influence the norms and institutions that emerge, which then influence behavior. The main thing, though, is that we observe cultural variation. What I want to do now is see what implications this variation might have for evolution.

8.2.2 Group selection

That there are cultural differences implies that there must be distinct groups of people for different cultures to emerge. In our evolutionary past people will clearly have been separated into groups based on kinship and geography. This can give rise to something called **group selection**, in which evolution acts on the level of the group rather than on the level of the individual. We have already come across this in section 8.1, where we saw that evolution might favor groups who avoid aggregate risk. I now want to take things one stage further.

To do so I will look at a stylized model of group selection called the **haystack model**. It gets its name from thinking of mice living in haystacks, but I will frame things in terms of people living in fields. So, imagine that nature provides fields with enough food that people can survive happily. The model then consists of three recurring stages, summarized in Table 8.5. In the first stage people randomly decide what field to live in. They then live together for T generations. After T generations all the fields are

Table 8.5 The three recurring stages of the haystack model. Natural selection takes place in stage 2, but group selection takes place overall.

	Stage description
1	People in the population randomly decide what field to live in
2	People who chose the same field live together as a distinct colony for T generations
3	The fields are destroyed and everyone in the population is mixed together

destroyed and everyone in the population is mixed together, before nature provides new fields and the process repeats itself. The key is the periodic separation of people into distinct colonies, and then the mixing together when the fields are destroyed. This can lead to group selection.

To see why, suppose that there are cooperative people and free riders. Cooperative people go out and collect food, which is then shared amongst others. This costs c but benefits everyone in the colony by amount b/N, where N is the number of people in the colony. Free riders do not collect any food for others. So, if $q(t)$ is the proportion of cooperators in a colony at time t, we can assume that cooperators have $a + bq(t) - c$ children and free riders have $a + bq(t)$ children, where a is some positive number.

Suppose that $b/N < c$ but $b > c$. Then we have a linear public good game in which collecting food is the public good. Because $b > c$, the more cooperators a colony has, the more children will be produced in that colony. But, because $b/N < c$, a free rider will always have more children than a cooperator. This latter effect means that, if T is small, free riders will have more children, and free riders will eventually dominate the population. Fortunately, if T is large, colonies with more cooperators will have time to grow relatively large. This means that cooperators could be dominant in the population.

To illustrate, suppose that there are two fields. Field 1 starts with ten cooperators, field 2 with nine cooperators and one defector. Figure 8.5 plots the proportion of each type over time. In the second field the number of free riders exceeds the number of cooperators after 18 generations. Eventually, free riders will dominate this colony. The colony in field 1 grows at a higher rate, however, because it is full of cooperators. This means that after ten generations the proportion of cooperators in the total population is increasing. Thus, cooperation can be sustained in the population if colonies with cooperators have long enough to benefit from being cooperative.

The haystack model shows that cooperative behavior *can* be sustained through group selection. As in section 8.1.3, this is possible because of assortative matching. A cooperator in a colony full of cooperators is more likely to meet a cooperator than is a free rider in a colony full of free riders. Thus, cooperators can do relatively well despite their apparent disadvantage relative to free riders. For this to work, we need the repeating process of reproduction in distinct colonies, which then disperse and intermingle.

Some care is needed, however, not to read too much into group selection. Group selection does not mean that evolution will always favor what is good for the group, and in reality it is typically a relatively weak force compared to individual selection. This is apparent in the haystack model, where it is very unlikely that cooperation

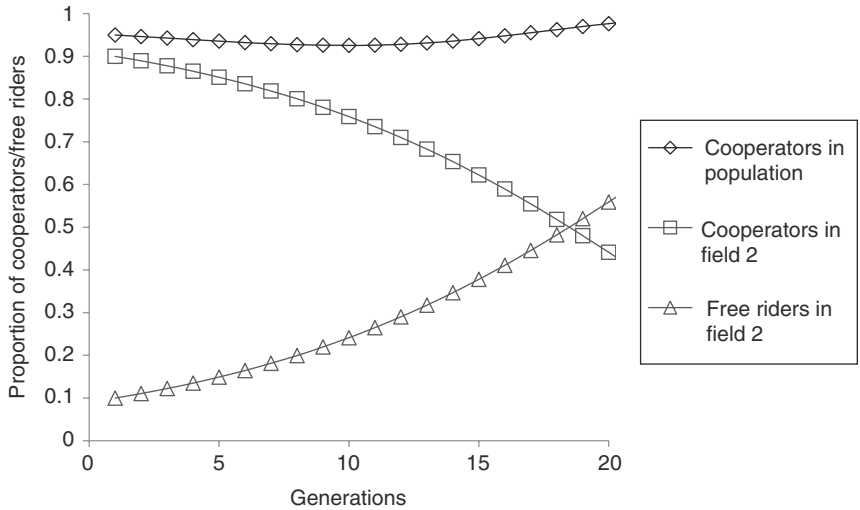

Figure 8.5 An example of the haystack model, where $a = 1$, $b = 3$ and $c = 0.5$; field 1 starts with ten cooperators, and field 2 with nine cooperators and one defector. Field 2 becomes full of free riders, but the overall population is still full of cooperators.

would emerge or spread. The best to hope for is that cooperation is not overtaken by free-riding. Even this happens only in ideal conditions!

The weakness of group selection should not come as much surprise, though, given what we saw in the previous chapter. There we saw that in social dilemmas most people behave either selfishly or reciprocate. We simply do not observe the cooperative, altruistic self-sacrifice that would be best for the group. That does not mean that group selection is not important. It just means that we need to be realistic in what it can achieve. What it might be able to do is help explain reciprocal altruism, as I shall now show.

When we looked at the haystack model we assumed a linear public good game without punishment. In this case, irrespective of $q(t)$, a cooperator always has c fewer offspring than a free rider. Now suppose a third type emerges, willing to punish. A punisher will collect food and also, at cost k, lower the payoff of any free rider by p. Now, if there are $q(t)$ cooperators, $l(t)$ punishers and $f(t)$ free riders, cooperators will have $a + b(q(t) + l(t)) - c$ children, free riders $a + b(q(t) + l(t)) - pl(t)$ and punishers $a + b(q(t) + l(t)) - c - kf(t)$. What difference does punishment make?

We can see that cooperators will have more children than free riders if $c < pl(t)$. Thus, cooperators will have more children than free riders if the punishment given to free riders is sufficiently high. Evolution could favor cooperation. There is a second-order free rider problem that we need to worry about, however. A punisher pays the cost both of collecting food and of punishing, so likely has the fewest offspring. Will there be enough punishers, therefore, to sustain cooperation, or will people free-ride on punishing?

To illustrate, suppose that women punish men who free-ride by not having children with them. In an evolutionary sense, this is an extreme form of punishment that

favors cooperation. But a woman who does not punish has a greater choice of men to mate with than one who does not, so will likely have more children. This means that evolution will favor women who do not punish. One partial solution is to punish women who have children with a free rider – that is, to punish people who do not punish. But then it would be necessary to have someone to punish those who do not punish those who do not punish, and so on! This is the **second-order free rider problem**.

The second-order free rider problem means that punishing behavior could not have arisen by individual selection. Crucially, however, the numbers have now changed in a way that cooperation and punishment is easier to sustain by group selection. To explain: without punishment, a cooperator always has c fewer offspring than a free rider, and so a cooperator is always at a large disadvantage relative to a free rider; group selection has a lot to do to overcome this disadvantage. With punishment, a cooperator has more offspring than a free rider, and a punisher has $kf(t)$ fewer offspring than a cooperator. The crucial thing this means is that a punisher is only at a disadvantage if there are free riders. Group selection has relatively less to do to help punishers overcome any disadvantage. This can tip the balance in a way that group selection can have a realistic role to play.

So, group selection is a plausible explanation for why we observe reciprocity. The story is still a little incomplete, however. That's because, although we now have a good story for why cooperation could be sustained through punishment, we still do not have much of a story for why cooperation and punishment would emerge in the first place. For that we need to add culture.

8.2.3 Gene–culture coevolution

Group selection relies on people being separated into distinct groups or cultures. Behavior is still determined by genetic evolution, however. So, Edward is, say, cooperative because he has a cooperative gene that he inherited from his parents. His payoff is influenced by those in his culture, but his behavior is not. This is consistent with cross-cultural variation but gives only a very minor role to culture itself. What happens if we start to add a bit more cultural influence?

Observing and learning from others can lead to cultural transmission, in which ideas are passed from one person to another. Ideas might be passed from parent to child, **vertical cultural transmission**, or between people of the same generation, **horizontal cultural transmission**. In the same way that good strategies or behaviors might be favored by genetic evolution, good ideas might be favored by **cultural evolution**. For example, Edward might copy or inherit the cooperative, trusting behavior he sees in a friend. Or, he might copy the latest fashion of the day.

The interaction between cultural and genetic evolution raises some fascinating questions. These are encapsulated in the **nature versus nurture debate** about whether a person's behavior is more shaped by their genes or the environment they grow up in. Is Edward more likely to be cooperative because he has a cooperative gene or grew up with cooperative people? I am not going to explore such questions here. But I think it is worth briefly thinking about how cultural evolution interacts with genetic evolution. Does cultural evolution, for example, overwhelm genetic evolution?

To me that seems unlikely. After all, genetic evolution gave us and shapes our ability to learn from others. Genetic evolution would also put a stop to certain types of cultural evolution, particularly by vertical cultural transmission. For example, a family passing from generation to generation the idea 'Use contraceptives and only have one child' is probably not destined to last very long. Neither is a family with the norm 'Smoke, drink and exercise rarely'. Things can work the other way, though, and cultural evolution influences genetic evolution. For example, if there is a cultural norm to mate with tall people then the average height of the population will likely rise over generations.

It seems most apt, therefore, to think of genetic and cultural evolution as part of one big process in which they complement each other rather than dominate each other. This interaction is termed **gene–culture coevolution**. One interesting aspect of gene–culture coevolution that I want pick up on here is how it can increase the possible strength of group selection. Let's see why.

Suppose that there is a **conformist bias**, whereby individuals copy the most prevalent behavior in the group or culture. We saw some evidence in Chapter 5 that this is plausible. I now want to think through what consequence this has in the haystack model. A conformist bias means that colonies will become relatively homogeneous, or, at least, more homogeneous than they would be without a conformist bias. For example, if most people in a colony are cooperative then others will conform to that, and so the colony as a whole becomes more cooperative.

This has two consequences. First, it means that the differences between colonies become starker. We tend to get fewer colonies with a mixture of types, and more colonies dominated by one type, either cooperators or free riders. The second thing it does is make colonies resistant to change. For example, a free rider who joins a colony full of cooperators may conform and change his behavior, rather than having lots of children who subsequently turn the colony into one of free-riding.

Both of these things make the group selection explanation for reciprocation more plausible. Indeed, it may be enough for punishment and cooperation to emerge rather than just survive, because cooperation, once established in one colony, can easily spread from colony to colony through cultural transmission. In addition, in a group that has a cultural norm of cooperation, genetic evolution may favor individuals who do not free-ride. Cooperation in one colony may be enough, therefore, for cooperation to spread through the population.

The exact details of how reciprocation emerged are still hotly debated, and there are certainly a lot of ifs and buts in the story that I have just told. We saw in the last chapter, however, that reciprocation is common in human behavior, and so this must have an evolutionary explanation. Gene–culture coevolution seems to offer the most promising explanation we have.

8.2.4 Reciprocity in children and chimpanzees

In section 8.1 I focused mainly on how evolution can shape individual decision-making. In looking to explain social preferences, such as reciprocation, I have drawn increasingly on the role played by culture. It is natural to take account of culture when looking at social preferences, because the focus is on interpersonal decision-making. There is a sense, however, in which gene–culture coevolution can seem like a bit of a fudge.

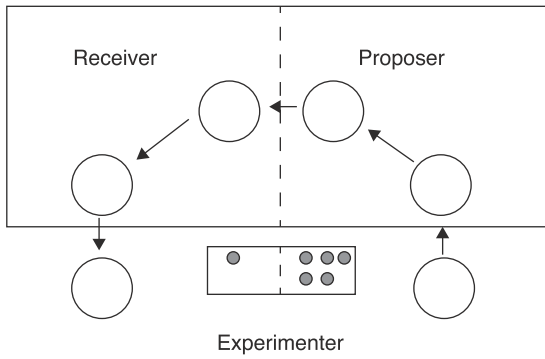

Figure 8.6 Experimental setup for an ultimatum game with chimpanzees.
Source: Proctor *et al.* (2013).

The accusation would be that we are explaining behavior such as reciprocity by adding more and more things to a complex model. Can we not find a cleaner distinction between the role played by evolution and culture?

The answer is almost certainly 'No'. Neither evolution nor culture, on their own, seem enough to explain reciprocity. What we end up with, therefore, is an inevitable interplay between the two. One interesting way to illustrate this is to look at the social preferences of children and chimpanzees. By studying children and chimpanzees we get as close as we possibly can to stripping culture out of the equation (see Research methods 8.2). We shall see, though, that culture still plays its role.

Let me start with a study by Proctor and co-authors (2013). They looked at chimpanzees playing a restricted choice version of the dictator and ultimatum game. The basic setup of an ultimatum game experiment is depicted in Figure 8.6. Two chimpanzees are next to each other. The proposer is offered two tokens and must choose one. One of the tokens equates to an offer of three banana rewards each, and the other token equates to five banana rewards for the proposer and one for the receiver. Having chosen a token, the proposer must hand it to the receiver. The receiver then decides whether to hand it back to the experimenter. If it is handed back, the receiver and proposer get some banana. In the dictator game experiment, the proposer simply chooses a token and returns it straight to the experimenter.

In the dictator game experiment chimpanzees chose the token giving an equal payoff 11 percent of the time. In the ultimatum game experiment the proportion rose to 72 percent. In other words, the chimpanzees were far more likely to offer a fair split than to give a fair split. This is similar to what we see in human subjects. Slightly puzzling, however, was that no chimpanzee ever rejected an unfair offer.

To delve a little more deeply, we can look at a study by Jensen, Call and Tomasello (2007). They also had chimpanzees play a restricted choice ultimatum game. The slight twist was that they considered more games (and used raisins rather than bananas). Specifically, they looked at all of the four games in Table 7.6. Recall that in these games the responder could offer an 8 and 2 split or something else, with the something else ranging from an equally unfair 2 and 8 split to an even more unfair 0 and 10 split. (We can think of the study by Proctor and co-authors as focusing on the 5/5 game.)

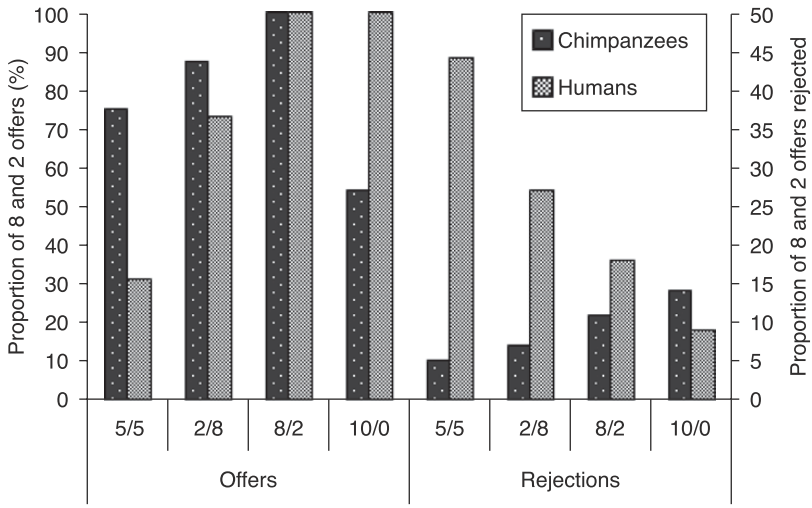

Figure 8.7 The proportion of 8 and 2 offers by game and the proportion of 8 and 2 offers rejected.
Sources: Jensen, Call and Tomasello (2007) and Falk, Fehr and Fischbacher (2003).

The left-hand side of Figure 8.7 looks at the proportion of times that chimpanzees offered the 8 and 2 split. This can be compared with human subjects. If you think carefully about the trade-offs in each game you should come to the conclusion that chimpanzees were less generous than humans. They were, however, willing to share a lot of the time. More interesting is the proportion of offers rejected. The right-hand side of Figure 8.7 shows that chimpanzees hardly ever rejected an 8 and 2 offer. This is in stark contrast to human subjects.

Both studies we have looked at show a similar pattern. Proposer chimpanzees were willing to share, but receiver chimpanzees hardly ever rejected. If offers are never rejected, however, why were the chimpanzees willing to share in the first place? Inequality aversion may be part of the story, but comparison of the dictator and ultimatum game experiments tells us that it is only part of the story. Something else, therefore, is at work. That something else may be social norms. Chimpanzees that do not share food in the wild are more likely to be threatened or subject to a temper tantrum. Moreover, the ultimatum game requires both chimpanzees to work on the task, and so it may have seemed more normal to share the reward in this case than in the dictator game.

Consider next a study by Harbaugh, Krause and Liday (2003) on giving in children. Figure 8.8 details the amount offered by children, aged between seven and 18, in a dictator and ultimatum game. We see that young children gave very little in the dictator game. Most children offered a near 50:50 split in the ultimatum game, however. The behavior of young children looks similar to that of chimpanzees, and things remain similar, albeit a bit more nuanced, when we look at rejection rates. Young children were willing to reject low offers, but they were less willing to reject than older children or adults.

Both chimpanzees and children used a 50:50 sharing rule. Is this due to culture or evolution? Recall from section 8.2.1 that not all cultures use a 50:50 sharing rule. We might, therefore, conclude that 50:50 sharing is a learnt social norm. This would

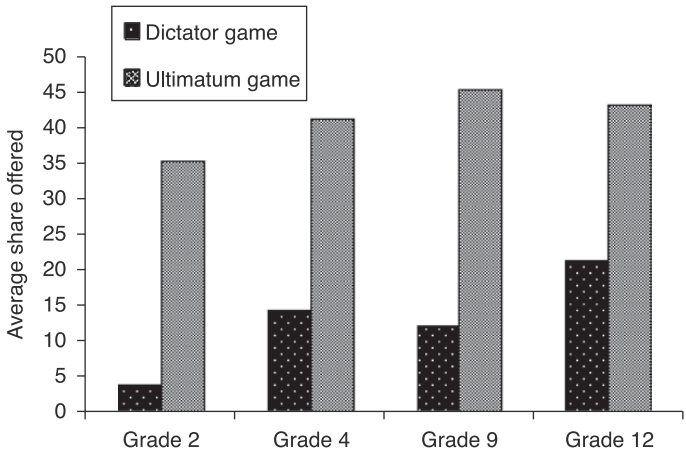

Figure 8.8 The share given in the dictator game and offered in the ultimatum game by school-aged children.
Source: Harbaugh, Krause and Liday (2003).

explain why different cultures end up at different norms. It would also explain why the chimpanzees and children offered to share even though low offers were unlikely to be rejected. This story has to be incomplete, however. Without an evolutionary advantage to sharing, there is no way we would see such a universal willingness to share something (even if not always 50:50). Evolution, therefore, must play its role. Indeed, in the next chapter we shall see that giving and punishment are associated with 'older' or more 'primitive' areas of the brain.

The basic thing I am trying to illustrate here is the complex interplay between evolution and culture. Even in the seemingly simplest case of chimpanzees and children, we see that there is no simple way to pull apart the influence that evolution and culture have. Over time we may find answers to some of the questions raised in this section, and this may give a somewhat clearer distinction between their respective influences. It seems clear, however, that a simple, clean distinction between the two is never going to emerge. Behavior, such as reciprocity, is shaped by a complex mix of both evolution and culture.

8.3 Summary

For most of us, finding food is now a simple task of going to the nearest supermarket, and we live well away from the savannahs of Africa. Hopefully, however, I have shown you that an evolutionary perspective is still useful in understanding modern economic behavior. Here's what we have seen.

We first saw that evolution will favor people who have good heuristics for maximizing expected utility. Intermediate goods that most increase the likelihood of having children will give most utility. We also saw that, when risks are idiosyncratic, evolution will favor people who use exponential discounting.

We also saw that the presence of aggregate risk and interdependence between people can have important implications. Aggregate risk can help explain risk aversion,

Table 8.6 Gender equality data (PPP = purchasing power parity).

Country	Ranking	Labor force participation (%)		Wage for similar work	Estimated earned income (PPP $)		Members of Parliament (%)	
		Men	Women	Men/ Women	Men	Women	Any	Minister
Iceland	1	87	83	0.79	51,143	36,817	41	44
United Kingdom	20	82	71	0.68	52,575	28,237	29	23
United States	45	77	66	0.65	66,338	43,122	19	26
Russia	75	79	68	0.65	29,321	17,496	14	6
India	87	82	28	0.57	9,045	2,103	12	22
China	99	84	70	0.65	16,177	10,049	24	12
Japan	111	85	66	0.66	48,796	25,091	9	22

Source: World Economic Forum (2016).

impatience and decreasing impatience with age. Interdependence between people can help explain why signaling and seemingly irrelevant traits become important; these, in turn, can help explain desires for things such as status and esteem, and a keen ability to distinguish the motives or traits of others.

We next looked at the consequences of culture. Cross-cultural comparisons of behavior in the ultimatum game and public good game, and of trust, provided evidence that important cultural differences do exist. These differences seem to depend a lot on local institutions and norms. We then questioned whether group selection can help explain reciprocation, and argued that it probably can if we take account of cultural transmission. This suggests that social preferences are shaped by a mix of evolution and culture.

In sum, we have seen an evolutionary explanation for a cardinal utility function over intermediate goods, risk aversion, exponential and hyperbolic discounting, signaling and a desire for status and esteem, and reciprocity. That is by no means all the behaviors we have looked at in this book, but it is a pretty impressive list. For me, though, the most exciting thing about an evolutionary perspective is that it can suggest questions that we might have ignored without it. For example, we have seen that an evolutionary perspective predicts a big difference between aggregate risk and idiosyncratic risk. We now need to go away and ask whether we do observe this, and whether it matters. An evolutionary perspective also gives different insight into what is likely to give people high utility and what not. This will prove useful when we look at happiness in Chapter 10.

8.4 The gender gap

An issue that has always attracted much interest is the differences in opportunities and outcomes between men and women. I will not delve too deeply into how large and important these differences are, but Table 8.6 gives some figures from the latest *Global Gender Gap Report* (World Economic Forum, 2016). Even in the highest-ranked country, Iceland, we see that men participate more in the labor market, earn higher

wages and are more represented in parliament. An important question is whether such differences are due to discrimination against women or due to different preferences and choices of women. In the former case the gender gap is a problem we need to solve, but in the latter case it is not.

Maybe we can learn something about the gender gap by looking at whether men and women behave differently when making economic decisions. Let's see. Before I do that, though, one point is worth making. Men and women generally behave in very similar ways. There is certainly far more heterogeneity in behavior within gender than there is variation in behavior between genders. That said, there are some notable differences that we do observe between the behavior of the average woman and the average man. It is these differences in the averages that I want to look at.

8.4.1 Attitudes to risk

One difference we observe is in attitudes to risk. Women, it seems, are more risk-averse than men. For example, several studies have shown that women are less likely than men to take risks in the economics laboratory. Other studies have shown that women are less willing to speed and more likely to wear a seat belt, brush their teeth and have regular blood pressure checks than men. To illustrate, I shall look at a study by Jianakoplos and Bernasek (1998).

They looked at a sample of over 3,000 US households in 1989 to see whether women were less likely to invest in risky financial assets than men. The main findings are summarized in Table 8.7, which contains estimates of the share of risky assets that a household held. Across the board, single females invest less in risky assets than single men or couples. Also noteworthy is the effect of a change in wealth from $20,000 to $100,000. Both men and women are predicted to increase their share of risky assets, but men increase the share more than women. This is a sign of greater risk aversion in women.

In the long term risky assets give a higher rate of return, and so differences in risk attitude can have large implications for investment return. For example, a return of 5 percent rather than 4.5 percent per year on a $1,000 investment will result in a

Table 8.7 Predicted proportion of portfolio invested in risky assets. The baseline case is a 41- to 45-year-old white homeowner, employed, high school degree, $20,000 investment wealth, and two children.

	Single females (%)	Single males (%)	Married couples (%)
Base case	43	51	48
Wealth = $10,000	35	39	39
= $20,000	43	51	48
= $100,000	62	79	69
= $500,000	81	100	89
Age = 26–30	23	63	47
= 31–35	32	42	37
Education = 6 years or fewer	77	79	50
Race = white	43	51	48
= black	58	49	42

Source: Jianakoplos and Bernasek (1998).

$241 difference after 20 years. This can build up to a huge difference in wealth over a lifetime, suggesting that women will on average end up with less wealth than men.

Before moving on, I want to pick out three other notable things in Table 8.7. We see that young women are predicted to invest much less in risky assets than young men, while less educated, and black, women are predicted to invest more in risky assets. This is interesting, because it would suggest that differences in risk attitudes are due to culture as well as evolution. There is, for instance, no genetic reason why black women would invest so much more in risky assets.

8.4.2 Attitudes to competition

A second difference that we observe between men and women is how they respond to competition. Some studies have shown that men improve their performance when faced with competition, while women do not. Other studies have shown that, given the choice, men choose competitive environments, while women do not. To illustrate, I shall look at two studies, the first by Gneezy and Rustichini (2004), the second by Niederle and Vesterlund (2007).

The Gneezy and Rustichini study was performed in a school in Israel and involved nine- to-ten-year-old children running a 40 meter track. First, each child ran by themselves, and times were recorded. Next, pairs of children who recorded similar times were matched together and raced against each other. Finally, some children were just asked to run a second time by themselves. The results are summarized in Figure 8.9. In the first run, and when there was no competition, there are no differences between boys and girls. In a competitive environment, however, boys ran significantly more quickly and girls ran more slowly. Boys thus responded more to competition than girls.

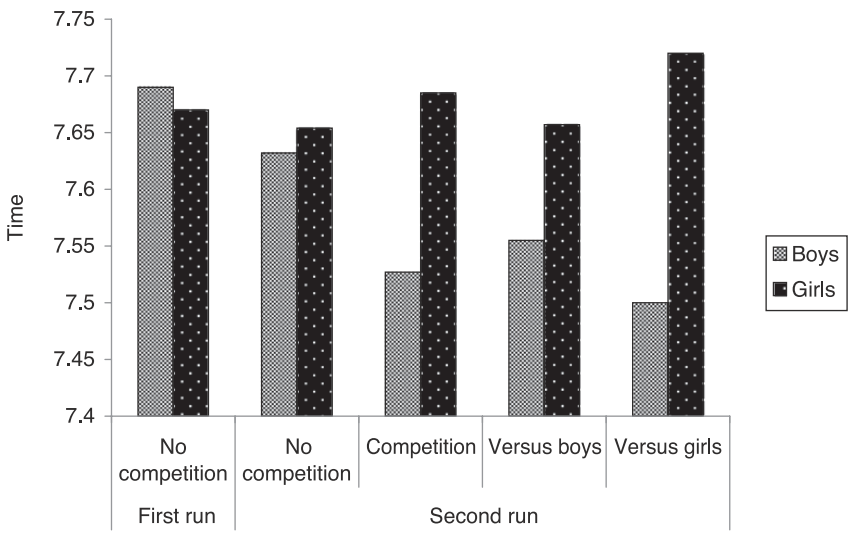

Figure 8.9 Running times of boys and girls. With no competition, there is not much difference in times. With competition, the boys get faster and the girls slower.
Source: Gneezy and Rustichini (2004).

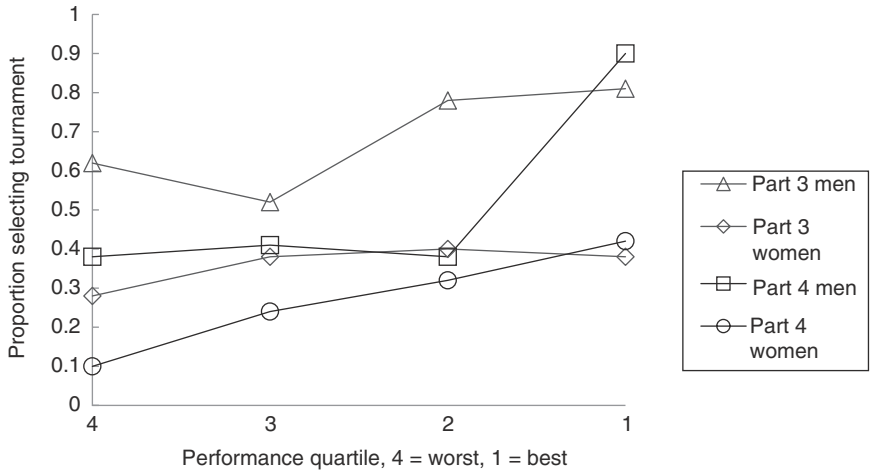

Figure 8.10 The proportion who chose competition conditional on how well they had performed.
Source: Niederle and Vesterlund (2007).

The Niederle and Vesterlund study consisted of four parts. In the first three parts, subjects had five minutes to answer as many simple arithmetic problems as they could. In part 1 they were paid by a piece rate of $0.50 for each correct answer. In part 2 subjects were put in groups of four and payment was by tournament, the subject with the most correct answers getting $2 per correct answer. In part 3 participants could choose whether to be paid by piece rate or tournament. In part 4 subjects were asked to choose whether they wanted their part 1 performance paid by piece rate or tournament.

This sounds a bit complicated, so I will go through it step by step. In the first two parts gender had no effect. Men and women were as good as each other with both the piece rate, when there was no competition, and the tournament, when there was competition. Interesting things start to happen in part 3 of the experiment, when subjects had to decide whether to be paid by piece rate or tournament. Figure 8.10 illustrates what happened by plotting the proportion of subjects who chose the tournament relative to how well they had performed. We see that men were more likely to choose the tournament than women. Indeed, the best-performing women, on average, chose the tournament less often than the worst-performing men! The story is similar, if slightly less dramatic, in part 4, when subjects were asked how they wanted to be paid for the first part of the experiment. Again, men were more likely to choose the tournament payment.

❖ RESEARCH METHODS 8.4 INCENTIVES OVER TIME

If participants play the same game a number of times, issues arise about how to pay them and maintain good incentives. An obvious solution is to pay participants for every decision they make. If, however, a participant is going to make 25 decisions and be paid around $10 in total, there is never going to be

much incentive to make good choices. In addition, if participants have in mind some reservation amount that they want from participating in the experiment then choices may change if they achieve that amount. This is a real concern in the Niederle and Vesterlund study, in which someone who does well in parts 1 and 2 may be more willing to take a risk in parts 3 and 4. To avoid this it is typical to pay participants for one randomly selected round or part of an experiment. This should maintain good incentives but can lead to somewhat lucky and unlucky participants!

We see, therefore, that men appear to react differently to competition from how women do. I shall shortly try to explain why this might be the case. For now, however, I want to point out that much is still unknown about how competition effects interact with gender. For example, the Gneezy and Rustichini study suggests that women are less competitive with other women (see Figure 8.6). A complementary study by the same authors showed an opposite effect. We see similar complications when looking at social preferences.

8.4.3 Social preferences

In terms of social preferences, there appears to be no simple relationship between gender and behavior. For example, in one set of experiments women might seem more altruistic, but then, in another, men might seem more altruistic. In specific contexts, however, gender has been shown to matter. To illustrate these points I shall compare studies by Eckel and Grossman (2001) and Solnick (2001). In both studies, the ultimatum game was played between a receiver and a proposer of known gender. In both cases we see that gender has an effect on receiver behavior.

In the study by Eckel and Grossman, subjects had to divide $5 and interacted in groups of four proposers and four receivers sitting opposite each other. A proposer could not know which of the four people sitting opposite would be the receiver to his or her offer, but one of them would be. Table 8.8 provides a snapshot of the results by looking at whether an offer of $1 is accepted or rejected (according to a model that fits the data). We see that a $1 offer from a female is more likely to be accepted than

Table 8.8 The percentage of $1 offers accepted from a model fitted to the experimental data. For example, if the proposer was female, 36 per cent of males would accept the offer.

		Responder (%)	
		Male	**Female**
Proposer	**Male**	**20**	**39**
	Female	36	78
	Unknown	17	42

Source: Eckel and Grossman (2001).

Table 8.9 Minimal acceptable offers. For example, if the proposer was female, the average amount that a male would accept was $3.39.

		Responder	
		Male	**Female**
Proposer	**Male**	**$2.45**	**$2.82**
	Female	$3.39	$4.15
	Unknown	$2.71	$3.30

Source: Solnick (2001).

a similar offer from a male. We also see that females are more likely to accept than men. Eckel and Grossman call this chivalry and solidarity.

In the study by Solnick, subjects had to split $10 and were not sitting opposite each other. A further difference is that receivers wrote down minimal acceptable offers before seeing the actual offer they had received (see Research methods 7.2). Table 8.9 details the minimal acceptable offers of receivers. This time things look very different. We see that both male and female receivers set a higher minimal acceptable offer when the proposer was a female rather than a male. We also see females setting higher minimal acceptable offers than men. This means that, if a male offered a $7, $3 split to another male, it was likely to be accepted, while between two females it was likely to be rejected. There are no signs of chivalry or solidarity here.

The results of these two studies come to very different conclusions about receiver behavior. These differences could arise from the strategic environment. Female offers were more likely to be accepted when subjects could see each other and when they were responding to an actual offer. Female offers were less likely to be accepted when subjects could not see each other and were deciding before an offer had actually been made. Gender does matter, therefore, but seemingly in ways that are conditional on the strategic environment.

8.4.4 Why are men and women different?

We have seen that gender seems to matter for behavior. Why? There are personality differences between men and women that could help explain the differences in economic behavior. Here are three of them.

First, women appear to experience stronger emotions than men. This could influence their behavior. For example, we saw in Chapters 2 and 3 the importance of loss aversion and risk aversion. Both of these are partly emotional responses connected to the fear of losing money or something else. Stronger emotions may also explain why women were less likely to reject when face to face with an ultimatum game opponent.

Second, men appear to be more overconfident than women, at least in some situations. This could make men more willing to gamble and to prefer competition. For example, Niederle and Vesterlund find that 75 percent of men thought that they would be the best in their group, compared to just 43 percent for women. Remember: there

were four in the group, so the number should have been 25 percent. Both men and women were overconfident, but men were far more overconfident!

Finally, men appear more stimulated by challenge than do women. Again, this could account for less risk aversion and a preference for competition. In the Niederle and Vesterlund study, for instance, there was less difference between men and women in part 4, which was about past behavior, than part 3, which was about a task yet to be done.

But why do we see these personality differences? Is it genetic or cultural? The answer is probably a mix of both. There are good evolutionary reasons why we could see such differences. For example, there is far more possible variation in the number of offspring that a man can have relative to a woman, so evolution could favor risk-taking, competitive men. Moreover, women differ from men in levels of testosterone and other hormones. Culture can still play its part, however, if boys and girls are raised and educated relative to stereotype. For instance, if women are expected to be less assertive then we should not be surprised that, on average, women might become less assertive.

So, what about the gender gap? If men are less risk-averse and more competitive than women, it is not hard to understand why a gender gap could arise. The extent to which the gender gap is a problem is a more difficult question to answer. If genetic differences mean men and women want different things then maybe the gender gap is not such a problem. If the differences are cultural, however, then there is arguably a lot more to be worried about. And the evidence would suggest that culture is at least part of the story.

8.5 The economics of family

As I have already pointed out when looking at team decision-making, it is typical in economics to focus on individuals. Occasionally we might talk about households, but it is only occasionally. Many economic decisions are made as a family, however, or at least with the family in mind. For example, Edward and Anna might want to save money for their children's education, or in order to bequeath something when they die. An evolutionary perspective is crucial to understanding such decisions.

One good reason to take account of evolution is kin selection. As a rough rule of thumb, half a child's genes are inherited from the father and half from the mother. This means that evolution should favor mothers and fathers who look after their children and make sure that 'their genes' are passed on. More precisely, there is a **coefficient of relatedness** between a parent and child of 0.5, between a grandparent and child of 0.25, between brother and sister of 0.5, and so on. **Hamilton's rule** states that an individual who can, at cost C, benefit a relative by B should do so if $Br > C$, where r is the coefficient of relatedness between them. Put in its most extreme form, Anna or Edward may be willing to sacrifice their own life to save two or more of their children.

Kin selection brings the unit of selection down to the level of the gene (compared to the individual and group selection we have seen so far) and is a form of assortative matching, because it means an individual tries to interact with and help family rather than unrelated individuals. This makes it much easier to understand transfers between family members. For instance, it is not really altruistic for Edward and Anna to pay for their children's education, it is just good evolutionary sense.

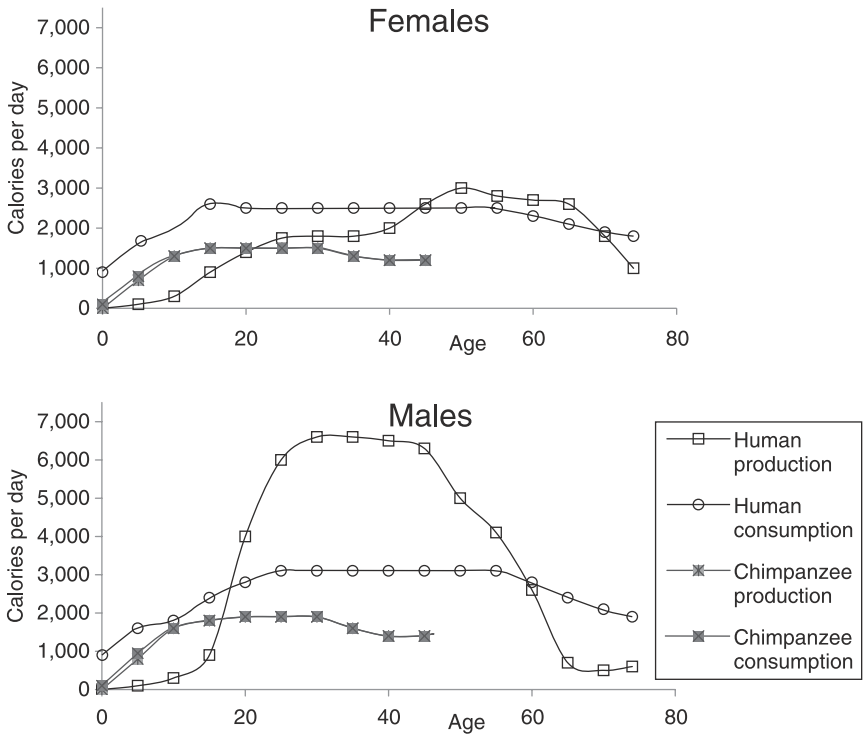

Figure 8.11 Food production and consumption by age in males and females, and in chimpanzees.
Source: Kaplan *et al.* (2000).

Kin selection matters in all species, but humans have taken it to a relatively extreme level. To see why, it's interesting to compare humans with primates, our closest relatives. Humans are distinctive in at least five ways: we (1) have bigger brains, (2) live longer, (3) spend longer as infants being dependent on parents, (4) spend time in a post-reproductive age looking after grandchildren and (5) find men supporting women and their offspring.

To illustrate, Figure 8.11 compares human consumption and production of food in three hunter-gatherer societies (Ache, Hadza and Hiwi) and chimpanzees. In humans we see huge differences over a lifetime between production and consumption, and between men and women. In chimpanzees we see hardly any difference between consumption and production. What this means is that chimpanzees are largely self-sufficient. Humans are far from self-sufficient. Indeed, the basic story with humans is that men aged 20 to 60 acquire enough food to support those under 20 and women aged 20 to 45 who are looking after children. This involves big transfers within the family.

The leading explanation for why we see such transfers is humans' bigger brains. Bigger brains mean we need longer to develop, and that means we need the support of family when we are young. It also means that we need to live longer in order to repay the debt that is accumulated in the early years of life. Luckily, a bigger brain also means

that we have learnt more complicated ways to acquire food, and hence produce large amounts of food. It is a complicated mix, but everything holds together. What is not so clear is the catalyst that started the process. Did evolution favor big brains and we acquired ways to produce them? Or did evolution favor males who produced large amounts of food (see section 8.1.4), and big brains were a consequence?

What we can say for sure is that transfers within the family are an integral part of life. That is not going to stop just because food is more plentiful, for most of us, than in our evolutionary past. We should expect to see transfers between family members. This suggests we should more readily think of the household as most relevant in many economic decisions. It can also help us understand a bit better the transfers we do see, whether they be bequests, paying for education or supporting a poor relative.

To illustrate, one recent trend of family life is the growing importance of **remittances** from migrant workers to their families back in the native country. For example, in 2015 remittances to Nepal were equivalent to 32.2 percent of that country's GDP, 18.5 percent in Honduras, 9.2 percent in Serbia and 3.3 percent in India (figures from the World Bank). In the other direction, over $61 billion was sent from the United States in 2015 in remittances. These are big numbers. Moreover, this possibility of sending back remittances has been a major motivation for individuals to seek work abroad, with the consequent increase in immigrant labor.

From an evolutionary point of view, we should expect transfers to flow downwards from older to younger generations. This would suggest that remittances serve a 'helpers-at-the-nest' purpose, whereby a worker abroad, say, sends money back to his parents so that they can support younger offspring. An alternative possibility is that children are supporting somewhat poorer parents, and so there is a transfer from younger to older generations. This is hard to justify with an evolutionary explanation but easier to support with cultural transmission if it becomes an accepted norm for children to go abroad and support poorer parents. It is not clear yet which way the flows typically go, but a better understanding of the evolutionary and cultural forces at work may help provide the answers.

I began the section by highlighting that many economic decisions are made as a family. It is common in economics to simply assume that family members will cooperate for joint benefit. Sharing of food and remittances would seem consistent with that approach. There is inevitable scope for tension within a family, however, and the question of whether resources should flow up or down generations is one example of this. After all, the coefficient of relatedness is always less than one, and so self should matter most no matter how much people love each other.

Several experimental studies have looked at decision-making within a household. One objective in doing so is to see whether families cooperate in order to maximize joint efficiency. A study by Cochard, Couprie and Hopfensitz (2016) provides a nice illustration. In the study, subjects played a public good game twice, once with their partner and once with a stranger who was the same gender as their partner (see Research methods 8.5). Figure 8.12 summarizes the headline result. We see that both men and women are more likely to cooperate when matched with their partner. (Continuing the discussion from section 8.4.3, you can also see that women are more likely to cooperate with a stranger than men.) Cooperation rates are, however, some way below 100 percent. This suggests that not everyone is maximizing joint efficiency.

Analysis from other parts of the study shows that no subjects appeared to be motivated solely by own payoff. Around a third of subjects did seem partly

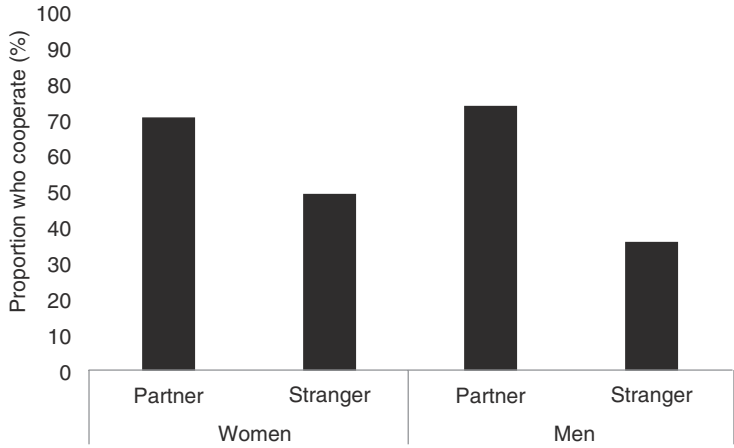

Figure 8.12 The proportion of subjects who cooperate in public good games when playing with their partner or a stranger.
Source: Cochard, Couprie and Hopfensitz (2016).

motivated by inequality aversion, though; they seemingly did not want the difference between own and partner's payoff to be too large. Interestingly, women (not men) with children were significantly less likely to cooperate. This was the same whether they were married or not. One explanation would be that women take on a large share of the childcare tasks and, in so doing, become more focused on own payoff – an issue that may also have something to do with the gender gap! Decision-making within households may not, therefore, be as efficiency-enhancing as some would suggest.

❖ RESEARCH METHODS 8.5 EXPERIMENTS IN THE HOUSEHOLD

Running experiments with couples generates some complicating factors. One issue is what to count as a couple. In the study by Cochard, Couprie and Hopfensitz, couples were required to have been living together for at least one year but did not need to be married. The couples in the experiment had lived together an average of 7.9 years, 44 percent were married and 48 percent had a child in the household.

A trickier issue is how to preserve anonymity of choice given that subjects are explicitly told they will be playing with their partner! Cochard, Couprie and Hopfensitz achieve anonymity by having an experiment with multiple parts and randomly selecting payment from some parts (recall Research methods 8.4). Given that some parts are not with a partner and some differ in the payments available, it would be impossible for a subject to infer what their partner chose. Subjects could, thus, choose confident in the knowledge that their

partner would never know what they chose. The trade-off with this approach is that it necessarily makes payment less transparent than we might like.

Some studies simply do away with anonymity and have a design in which partners are able to infer the choice of their partner. Clearly, this may influence choice. That, though, provides us with additional insight. For instance, if choice is anonymous we gain insight into whether couples wanted to maximize joint efficiency, while, if choice is not anonymous, we gain insight into whether couples wanted to be seen to maximize joint efficiency.

8.6 Development economics

One of the most basic goals of economics is to help alleviate poverty. The ideal would be that sound economic theory can create a world where resources are exploited efficiently and everyone has a comfortable standard of living. Unfortunately, economics does not have a great track record. While people in developed countries can have pretty much whatever they want, many in the rest of the world have to survive on next to nothing. Clearly, economics is not entirely to blame for this! There is, however, a litany of failed economic policies in developing countries. In this section I want to briefly look at how behavioral economics can explain such failure and offer alternatives.

Explaining why many economic policies have failed in developing countries is the easy bit. Development economics and economic policy have almost always been based on the standard economic model. This has led to policies that would work fine for *Homo economicus* but do not always work so well in reality. For instance, according to the standard economic model, people should be the same the world over, and so whatever works well in North America or western Europe should work well anywhere. This mindset has driven policy. We have already seen, however, that this mindset is flawed, because it ignores important, and systematic, cultural differences. Recall, for example, how the possibility of punishment works well to sustain cooperation in North America and western Europe but makes things worse in the Middle East and eastern Europe.

There is now an increasing recognition that development economics needs to take account of behavior, bias, cultural differences, and the like. Consequently, development economics is fertile ground for applying behavioral economics. It is worth highlighting that this is a two-way process, in which behavioral economists can learn a lot from trying to tackle the problems relevant to developing countries. To illustrate, I shall look at two of the more active areas of study: the education production function and microfinance.

8.6.1 The education production function

The Millennium Summit of the United Nations set eight millennium development goals to be reached by 2015. Number two on the list was to ensure that 'children

everywhere, boys and girls alike, will be able to complete a full course of primary schooling'. This goal clearly recognizes the role that education can play in reducing poverty, and narrowing the gender gap. Steady progress is being made in achieving the goal, with enrolment of 91 percent in 2015 compared to 81 percent in 2010. Universal education is not enough in itself, however. We also need the education to be of the highest possible standards. A lot of the early experimental work in development economics focused on this issue.

I shall start with a study by Glewwe, Kremer and Moulin (2009). They questioned whether providing textbooks to rural Kenyan schools led to an increase in test scores. This is a particularly pertinent question, given that policy-makers seem to think it should. Indeed, the objective of the World Bank Free Primary Education Support Project in Kenya, which ran from 2003 to 2007 and cost $55 million, was to ensure 'an adequate supply and better use of instructional materials'. So, do textbooks work? The simple answer seemed to be 'No'. The study looked at an earlier project that began in 1995 and progressively increased the number of textbooks in 100 schools. The textbooks made no apparent difference to test scores, except for the most able students. That the textbooks were in English, the third language of most of the students, presumably did not help.

Spending money on textbooks may not, therefore, be a good idea. This is a surprising negative result. We can look next at a surprising positive result. The Primary School Deworming Project, beginning in 1998, progressively made free deworming treatment available to 75 Kenyan schools. Miguel and Kremer (2004) demonstrate that the project reduced school absenteeism by around a quarter. This was partly due to an externality effect, whereby untreated students benefited from the treatment of others because it slowed the transmission of worms. The success of the program is very good news in terms of improving child health and the amount of schooling. It did not, however, translate into an improvement in test scores.

Can anything improve test scores? A study by Banerjee and co-authors (2007) looked at two interventions in India. In one intervention the weaker students were taken out of the classroom to work with a young woman from the community for two hours per day; the teaching of this balashki (or child's friend) focused on core skills. In the other intervention all students were given access to a computer-assisted learning program adapted to their ability. Figure 8.13 provides data on the short-term effect of the program by comparing performance on a test at the start of the school year with one at the end of the year. Almost across the board, test scores were higher with the balashki and computer-assisted learning.

The improvement in test scores did drop off somewhat after a student left the project. The overall picture, however, is still a very positive one. In many ways this is a surprising finding. In particular, the balashkis were given only minimal training before entering the classroom, and many expressed doubts about the merit of computer-assisted learning. Even so, both interventions seemed to improve upon standard teaching methods. This offers real hope for increasing the quality of education. For instance, the minimal training means that the balashki project is cheap and easy to scale up.

The three studies we have looked at in this section are by no means all that have been done on the education production function. They do, however, nicely illustrate the kind of things we can learn. The main lesson seems to be that one cannot

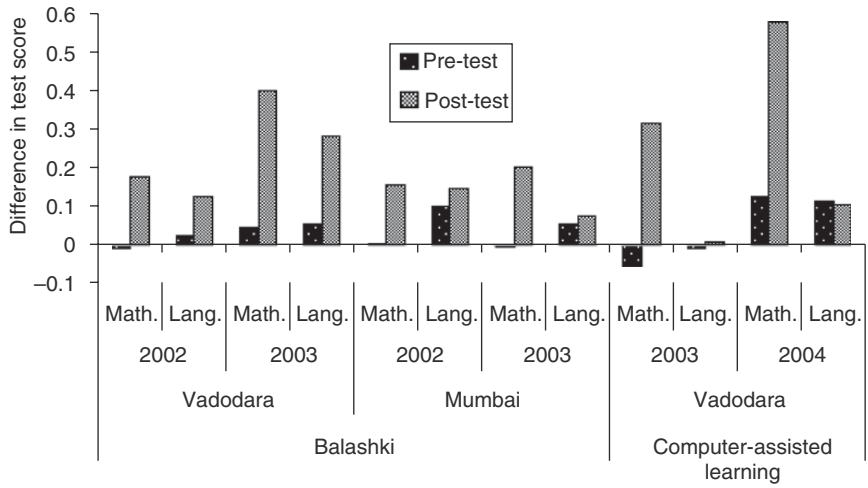

Figure 8.13 The difference in average test score comparing schools involved in the project and those not, in a particular year. Post-test scores are higher in schools involved in the project. Vadodara and Mumbai are the cities where the project took place.
Source: Banerjee *et al.* (2007).

trust intuition (or the standard economic model). Some things do not work when we might have expected them to, and some things work when we might not have expected them to. This emphasizes the need for a lot more investigation of potential policy interventions (see Research methods 8.6) if we are to alleviate poverty and not waste money.

❖ RESEARCH METHODS 8.6 EXPERIMENTING WITH POLICY

In order to properly evaluate the merits of different policy interventions we need well-designed randomized trials. The minimum this requires is an intervention and control treatment. For example, we need some randomly selected schools to implement a policy of deworming or balashki while others do not. Comparison of the schools in the intervention treatment with the control treatment allows us to draw robust conclusions about the consequences of the intervention. Randomized trials of this nature do raise some difficult ethical, methodological and practical problems, however.

Let us think about the ethical problems first. A basic principle in economic experiments is that of **informed consent**, which means that the subject knowingly signs up to take part. The problem we have here is that children are not in a position to make informed consent. In general, this need not be a barrier to running experiments with children. Earlier in the chapter, for instance, we looked at three experiments with children that clearly raised no real ethical

issues. When it comes to an experiment with, say, deworming treatment, the ethics become a bit less clear-cut. In this instance, a child assigned to the control treatment would have more ill health and less education than a child in the intervention treatment. Is that ethical? The difficulty in answering this question makes it hard to do randomized trials with children, whether in development economics or medicine.

Next we can look at the methodological problems. The general rule with experiments is to add one thing at a time in order to not overly complicate matters. This approach is not well suited to a development context, in which money is tight, time is short and people want to see immediate results. Multiple interventions, therefore, can be bundled together with the same population. The difficulty that this creates is in disentangling the various effects. More generally, a policy intervention in a developing country is also almost inevitably going to be quite context-specific. This makes it difficult to draw general conclusions, or, at least, we need care in doing so. For example, if greater access to textbooks failed to improve performance in Kenya because of the specific textbooks being used, what do we learn? The main lesson we learn is the importance of high-quality textbooks, not that textbooks do not matter.

Finally, we come to the practical problems. Many practical problems arise from the ethical and methodological problems already discussed. For example, the ethical issues may mean reluctance from government and aid agencies to take part in studies. There are additional practical problems from working in developing countries, however. Most notably, the experimental environment is likely to be much less stable than in a standard experiment. Attrition rates at school, for instance, could be high. Implementation of the project may also be patchy. For example, in Mumbai, only two-thirds of the schools assigned balashkis actually received any, because of administrative difficulties. Such effects need to be carefully monitored and taken control of.

8.6.2 Microfinance

One of the most exciting things to happen in developing countries over the last few decades has been the growth of microfinance. The basic idea is simple enough. Conventional lending does not work well for relatively poor people because the chance of default is large and there is no collateral. A microfinance scheme 'solves' these problems by lending to a group of people who collectively take responsibility for the loan. The pooling of risk within the group lowers the chance of default. This looks like a win-win solution: lenders can make a profit and poor people can borrow money to buy land or equipment, secure a job or deal with a family emergency.

Most credit for the emergence of microfinance goes to Muhammad Yunus. In 1976 he began a project in Bangladesh lending to the poor. That project resulted in the establishment of Grameen Bank in 1983. The guiding principle was that loans

are better than charity at alleviating poverty. In 2006 Yunus and Grameen Bank won the Nobel Peace Prize for 'their efforts to create economic and social development from below'. Microfinance has now spread to over 50 countries. In 2010 an estimated 137.5 million poor families had a microloan, typically of around $100. In short, microfinance looks like a resounding success story.

Appearances can be deceptive, however. One thing we can say with some certainty is that microfinance is a profitable business. This means that group lending must work, in the sense that repayment rates are high enough. To see why this is an interesting finding I will look at a study by Abbink, Irlenbusch and Renner (2006), who ran a laboratory experiment with a microfinance game.

To explain how the game works, consider a group of four people. Each person is lent $175 for a total of $700. After one year the bank wants a repayment of $840 from the group, equating to an interest rate of 25 percent. Group members independently invest the $175 in their 'business'. With probability 0.83 that the investment is successful, the person ends up with $420. With probability 0.17, the investment is not successful and the person ends up with $0. At the end of the year each group member reports whether his investment was successful, and those with successful investments split the cost of repaying the loan. For example, if three are successful they each pay $280 each.

In order to repay the loan at least two group members must report a successful investment. The interesting word here is 'report'. It is assumed that only the person himself knows whether or not his investment was successful. In order to repay the loan, therefore, it is not enough that two or more investments are successful. We also need group members to honestly report their success. If the loan is repaid then the bank will extend a loan in the subsequent year, and the whole process is repeated. This setup captures two key elements of microfinance schemes. First, it is difficult to monitor borrowers, and so there is much reliance on self-reporting. Second, the main benefit from repaying the loan is the chance to borrow again.

The game is such that no group member has a material incentive to report a successful investment. I will not go through the full explanation, but, to give the basic idea, suppose Rob's investment was successful and he is deciding what to report. If two or three of the other group members report a successful investment then the loan will be repaid; by saying that his investment was not successful, Rob avoids repaying his share of the loan. If no other group members report a successful investment, the loan will not be repaid whatever Rob does; by reporting a successful investment, Rob would have to pay his $420 but not be able to get a loan in the future. It is only if exactly one other group member reports a successful investment that Rob can possibly gain from reporting his successful investment; that is not incentive enough.

If group lending is going to work, therefore, we need group members to report successful investments. We need guilt aversion, social preferences, social norms or something similar. Let us turn to some experimental data. Figure 8.14 shows the repayment rates that would occur given the amount of honest reporting observed in the study. Treatments were run with groups of size two, four and eight. To explore the group element in more detail, a treatment was also run with groups composed of four friends.

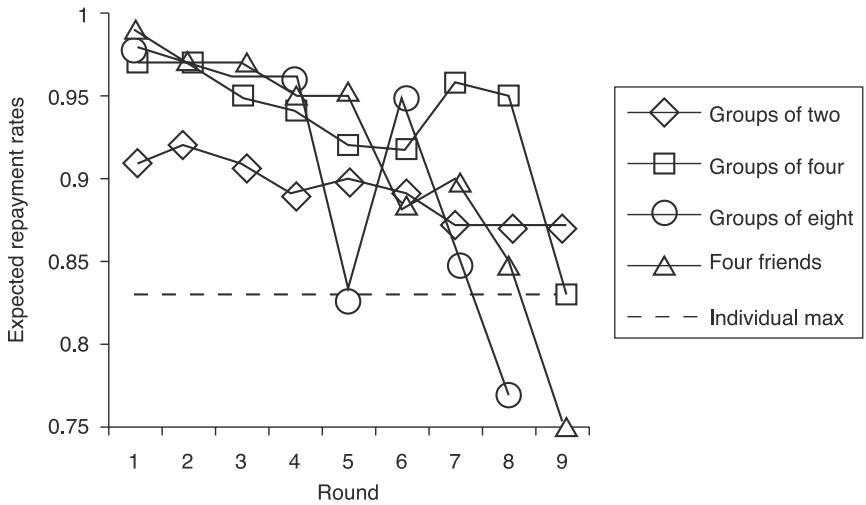

Figure 8.14 The probability of the loan being repaid given the observed levels of honest reporting.
Source: Abbink, Irlenbusch and Renner (2006).

The main thing to note from Figure 8.14 is that enough subjects honestly reported successful investments to generate a high repayment rate. To put things in context, the 'Individual max' line plots the repayment rate that the bank would have expected if it lent on an individual basis and closely monitored every loan. Lending to groups is a significant improvement over this hypothetical ideal. So, lending to groups works, and it seems that the size or composition of the group does not make much difference.

The high repayment rates we see in Figure 8.14 are an indication of why microfinance has been successful for the banks. We need to question next whether microfinance has been beneficial for borrowers. The evidence, so far, is not positive. In particular, a spate of recent studies has found little evidence of the purported benefits of microfinance. To illustrate, I will look at a study by Banerjee and co-authors (2015).

In Hyderabad, India, 52 of 104 poor neighborhoods were randomly selected for the opening of a branch by the microfinance institution Spandana. Operations were progressively rolled out between 2006 and 2007. The study surveyed a large sample of local people three times: in 2005, 2007 and 2009. Comparison of the neighborhoods with and without access to microfinance allows a relatively clean test of its consequences.

There was evidence that people used microfinance to make inter-temporal choices. Specifically, those who borrowed money appeared to spend the money on a durable good. The cost of the loan was offset by reduced consumption on 'temptation goods' and an increase in hours worked. There was no evidence, however, of microfinance lifting people out of poverty. Indeed, if it helped any, it tended to help larger businesses; the average business remained small and not very profitable. Moreover, the take-up of microfinance was not high; fewer than 40 percent of eligible

households took out a microfinance loan, even though most borrowed from some other source.

The results from this study in India are consistent with studies in other countries, including Morocco, Mexico, Mongolia and Bosnia. Microfinance may not, therefore, be the miracle cure for poverty. This is probably due in large part to the lack of good investment opportunities for the poorest in society. Few businesses are able to turn around a 25 percent return in one year! This is not to say that microfinance is a bad thing. If it helps people smooth consumption and pre-commit to not spending money on temptation goods then it can have some benefits. It is not clear, however, that even these benefits are worth much. Would it not be better to find some cheaper way for people to smooth consumption, such as through saving rather than borrowing? In Chapter 11 (see in particular section 11.2.1) we shall look at an alternative to microfinance.

8.7 Further reading

For those eager to learn more: much of the material in section 8.1 was based on the work of Arthur Robson. Robson (2001a, 2001b, 2002) and Robson and Samuelson (2009) are useful references. For more on cultural differences, I would suggest following up the references in the text, particularly the book by Henrich and co-authors (2004). See also Cardenas and Carpenter (2008). Sobel (2002) provides an interesting review of the literature on trust and social capital. For more on group selection, and the evolution of cooperation, Bergstrom (2002) and the book edited by Gintis and co-authors (2005) are worth reading, while the 8 December 2006 issue of *Science*, with articles by Boyd, Bowles and Nowak, provides an interesting snapshot of research in this area. For an interesting earlier contribution to the debate, see Güth and Yaari (1992). For more on the effects of gender, see the excellent review article by Croson and Gneezy (2009), and, for the economics of the family, Bergstrom (1996) is a good starting point. The fascinating article by Kaplan and co-authors (2000) explains the data on food production and consumption, while Robson and Kaplan (2003) offer an economic perspective. See Pascual-Ezama and co-authors (2015) for a cross-country comparison of cheating. Webley (2004) provides interesting insights on children's understanding of economics. On a more general level, the books by Frank (1991), Ridley (1997) and Laland and Brown (2005) are three that I would pick out from the host of excellent, general readership books taking an evolutionary perspective. For more on development behavioral economics and microfinance, see Bertrand, Mullainathan and Shafir (2004), Rodrik (2008), Banerjee and Duflo (2009) and Banerjee (2013). Some recent studies look at household efficiency in developing countries; Kebede and co-authors (2013) is one example.

8.8 Online materials

On the companion website you will have access to flashcards, quizzes and answers to the review questions.

8.9 Review questions

8.1 Give some examples of intermediate goods that may enter the utility function. Is money an intermediate good?

8.2 Let's go back to the first example of long and short winters. Suppose there are 1,600 adults of type L, 1,600 of type S and 1,600 randomizer types. The randomizer types decide at the start of each summer whether to gather enough food for a long winter, and do gather enough with probability 0.4. Suppose we have the sequence summer, short winter, summer, short winter, summer, short winter, summer, long winter; how many of each type will likely survive?

8.3 How can evolution explain reference dependence? How about loss aversion?

8.4 Are present bias and time inconsistency good evolutionary strategies?

8.5 Design an experiment to find out whether people dislike aggregate risk compared to idiosyncratic risk.

8.6 Is natural selection relevant in understanding modern economic behavior or is it only culture that matters?

8.7 Why is it so hard to get cooperation to emerge in the haystack model?

8.8 Discuss reasons why we might observe antisocial punishment.

8.9 Does maximizing utility mean maximizing happiness?

8.10 What do you think is the main cause of the gender gap?

8.11 How can confirmatory bias help explain the gender gap?

8.12 Does it help a person to be overconfident of ability?

8.13 Why do many people save money all their life in order to bequeath it to their children?

8.14 Why does it not make evolutionary sense for a young person to support an older parent? Why could it happen because of cultural influence?

8.15 Why may there be conflict within the household as to how to allocate resources?

8.16 Why should development economics take account of cultural differences?

8.17 Is informed consent necessary to run an experiment?

8.18 In Chapter 6 we briefly looked at the role of networks in facilitating coordination in groups. Discuss the consequences of assortative matching for network structure.

Chapter 9

Neuroeconomics

> Giving up smoking is the easiest thing in the world. I know because I've done it thousands of times.
>
> Mark Twain

We use our brains to think and make decisions, whether consciously or unconsciously, and so, if we understand how the brain works, we should be able to better understand economic behavior. Clearly, things are not quite as simple as this, because the brain is a very complicated thing. Over the last 20 years or so our understanding of the brain, and our ability to measure brain activity, has, however, improved a lot. It has improved so much that neuroscience is now a viable tool for economists to use in interpreting economic behavior. The goal of this chapter is to see what we have learnt so far.

9.1 An introduction to the brain

We will not get very far unless we know some basics of how the brain works. I am not, though, going to go into very much depth on this subject; all we need is a simple model of the brain that is not too far from reality. The model that economists typically work with assumes modularity. The idea behind **modularity** is that the brain can be split into a number of distinct components or **brain modules**, each of which are responsible for specific functions such as breathing, seeing, language, being altruistic or getting angry. We can think of these brain modules as distinct areas of the brain, and so we can split the brain into areas that are responsible for different things. This is schematically shown in Figure 9.1. Controlling all of this is the **interpreter**, which monitors the output of the modules, interprets them and reacts.

Modularity is a very crude model of how the brain works. For instance, we know that modules can work in parallel with other modules, and communicate with each other while doing so. We also know that there is nothing anywhere near as simple as an 'altruism module' or an 'anger module', and any classification is inevitably one of judgment. To a rough approximation, however, modularity does not seem too bad, and it is a very useful model for economists to work with.

It is useful because it potentially allows us to infer what causes a particular behavior. For example, suppose that someone called Rob has the chance to give Federica

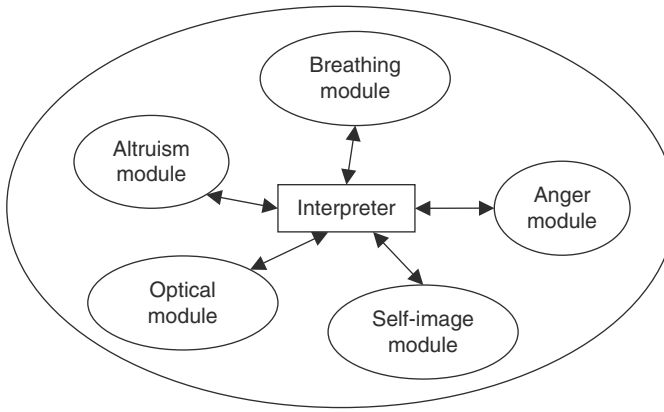

Figure 9.1 A schematic representation of modularity. The brain consists of modules that perform specific functions. An interpreter monitors and interacts with these modules.

money, and does so. Furthermore, suppose that we know there is a brain module associated with altruism and another associated with self-image. If we monitor the brain and find that the altruism module is activated, but not the self-image module, when Rob gives money then we could infer that he gave to be kind to others and not to promote self-image.

For this to be useful we need to measure brain activity. Historically neuroscience has learnt a lot from people who, for one reason or another, usually an accident, have damage in an area of the brain. If the person loses the ability to perform a function, such as get angry, then we can infer that the anger module is where the brain damage is. More recent techniques such as EEG, PET and fMRI allow us to measure brain activity directly (see Research methods 9.1). In principle, these techniques allow us to see what parts of the brain are activated when Rob is, say, angry, and also measure the intensity of the brain activity, so we can say how angry he is.

❖ RESEARCH METHODS 9.1 MONITORING BRAIN ACTIVITY

There are three basic brain-imaging methods: electroencephalography (EEG), positron emission topography (PET) and functional magnetic resonance imaging (fMRI). Each of these methods uses a different technique to measure activity in different parts of the brain. In an experiment, brain activity is measured during some baseline task and during some experimental events, such as the person being told they have some money. It is the difference in brain activity between the baseline and the experimental event that is important, because it gives us a clue as to what parts of the brain are activated, and how much they are activated, by the experimental event.

EEG is the only one of the three methods that directly measures brain activity. It does so via electrodes attached to the scalp, which pick up the electrical current that results when neurons fire.

PET and fMRI make use of the fact that active neurons need energy, and they get this energy from glucose in blood. Neural activity, therefore, increases blood flow. It also, because the active neurons take glucose from the blood but not oxygen, relatively increases the oxygen content of the blood.

A **PET** scanner measures the blood flow in the brain. To do this, a radioactive tracer isotope is injected into the person's blood circulation. As this isotope decays it emits positrons, which, on meeting nearby electrons, produce photons that can be detected. By counting the number of photons emitted it is possible to get an idea of how much blood and blood flow there is in different areas of the brain.

An **fMRI** scanner measures the oxygen content of blood. To do this, the head is put in a magnetic field. It so happens that the way blood interacts with this magnetic field depends on the oxygen content of the blood. This can be detected to show where in the brain the blood contains more or less oxygen, and thus where there is more or less neural activity. This is the most widely used method in neuroeconomics today.

The technicalities of EEG, PET and fMRI are interesting, but not ultimately our main concern. One thing is worth mentioning, however: there is no perfect way, yet, of measuring brain activity. EEG can record very accurately the time of brain activity (down to milliseconds) but is relatively inaccurate in saying where the activity was, or at picking up activity in some parts of the brain. PET and fMRI scanners can record much more accurately where the brain activity takes place (typically to within two to three millimeters) but are much less accurate on the time of activity (seconds rather than milliseconds). In addition, none of the methods is very user-friendly. EEG is relatively unobtrusive and portable compared to PET and fMRI but still a lot more involved than participating in a normal economic experiment.

The idea that we can measure the intensity of brain activity is exciting, because it suggests that we can measure pleasure, pain and utility. Care is needed in pushing this idea too far, though, not least because of the crudeness of the modularity model of the brain. Even so, in principle, as Rob thinks about giving away money, we could measure activity in the altruism, self-image and selfish modules, and see whether we can predict what he will choose from the intensity of activity. For instance, if the selfish module is activated a lot, we might be able to predict him keeping the money.

This already gives us two things to look at while we watch the brain: what areas of the brain are activated, and how much they are activated. Potentially, this can help us understand a bit better the behavior we observe, and we'll see how this works in practice soon enough. Before that, there are a couple of important things that I would like to do: introduce some areas of the brain of most relevance to economists, and talk a bit about different brain processes.

9.1.1 An economist's map of the brain

Neuroscience brings with it a load of terminology that is probably alien to many; at least, it is to me. An understanding of this terminology is useful, however, and so I will try to explain the main terms you are likely to come across. On the basis that this is an economics textbook and not a medical textbook, I shall give only the most basic introduction.

It is traditional to start with the big picture and progressively narrow things down. So, I should start with the nervous system. The **nervous system** transmits signals from one part of the body to another. In doing so it allows an individual to detect changes within themselves and their environment, analyze them and respond to them. The latter two of these are of primary importance for us, meaning that we can focus on the central nervous system. The **central nervous system** consists of seven basic parts: the spinal cord; the cerebellum; the medulla, pons and midbrain, which are collectively called the **brainstem**; plus the **diencephalon** and the **cerebrum**. These are illustrated in Figure 9.2.

Although the cerebellum, midbrain and diencephalon are of interest to economists (Table 9.1 outlines their main functions), most attention is usually given to the cerebrum. This is the big part of the brain responsible for higher-level information processing and decision-making. The cerebrum is more than big enough that we need to distinguish different parts of it, and there are various ways of doing this.

The one you are probably familiar with is the distinction between **left** and **right cerebral hemispheres**, corresponding to the left and right sides of the brain. Equally intuitive, but probably less familiar, is a distinction between the outside and inside of

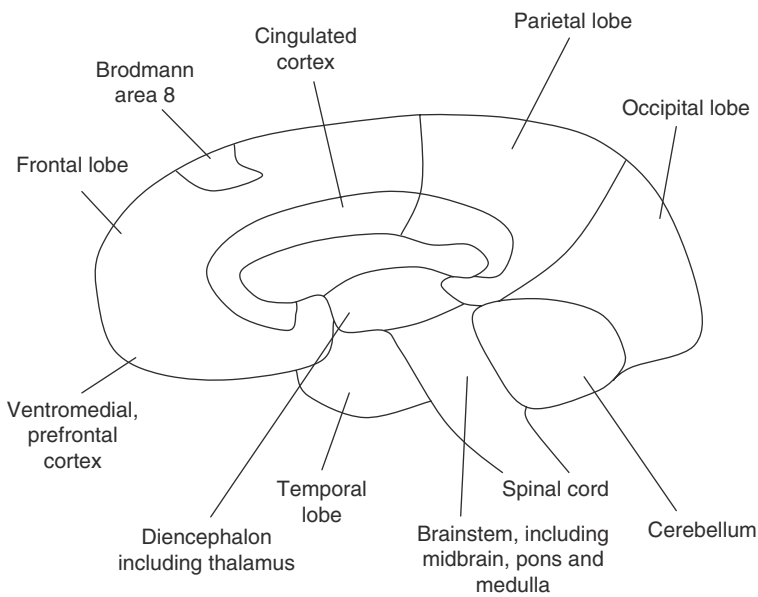

Figure 9.2 Some areas of the brain seen from the medial surface of the brain (i.e. by slicing the brain down the middle).

Table 9.1 Parts of the central nervous system and some of the main functions that go on in that part.

Parts of the central nervous system	Functions include
Brainstem (including pons, medulla and midbrain)	Control of reflexes such as heart rate, breathing and coughing Basic aspects of reward and pleasure Conduit of information in the system
Cerebellum (little brain)	Sensory motor skills necessary for smoothly coordinated movements
Diencephalon (including thalamus)	Control of homeostatic and reproductive functions, such as body temperature, hunger and thirst

the cerebrum. The **cerebral cortex** is the distinctive outside layer of the cerebrum, and is essentially a crumpled sheet of neurons around two to four millimeters thick. Within that we find something called the **basal ganglia**.

As well as distinguishing left and right, and inside and outside, we need ways to pinpoint different areas of the cerebral cortex and basal ganglia. To get some intuition for how this is done, imagine a map of a country. There are natural features, such as mountains and rivers, that we might use to pinpoint a specific location. There are also artificial features, such as county or city boundaries, that we might want to use. Sometimes the natural and artificial features will coincide, but sometimes not. Something similar happens with mapping the cerebrum.

Let's start with the cerebral cortex and the artificial features. One classification splits it into four basic parts: the **frontal**, **parietal**, **temporal** and **occipital lobes**. These can be further distinguished into things such as the **dorsolateral prefrontal cortex** and **ventromedial prefrontal cortex** (see Research methods 9.2). If these names seem a bit unwieldy then you may prefer a classification originally conceived by German neurologist Korbinian Brodmann. In this case, the cerebral cortex is split into 52 areas, known as **Brodmann areas**. Each Brodmann area corresponds to a particular area of the cerebral cortex; area 8, for instance, is illustrated in Figure 9.2. These areas do not necessarily correspond to natural features of the brain, but they do allow us to pinpoint specific locations of the cerebral cortex.

❖ RESEARCH METHODS 9.2 THE TERMINOLOGY OF NEUROANATOMY

Various terms are used to specify locations of the brain, some of which are illustrated in Figure 9.3.

Simplest are: **anterior** and **posterior**, which indicate the front and back of the head; **inferior** and **superior**, which indicate above and below; and **medial** and **lateral**, which indicate towards the center or to the side.

Other terms come from drawing the axis of the body and brain. **Dorsal** refers to the back or upper half, depending on whether the focus is the body or

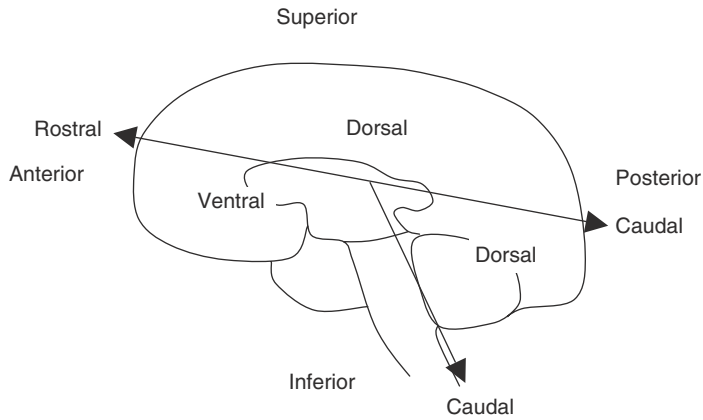

Figure 9.3 Terminology to describe location in the brain.
Source: Purves *et al.* (2008).

brain. **Ventral** refers to the front or lower half. **Rostral** indicates the direction towards the nose and **caudal** towards the back.

This is enough to understand terms such as the posterior parietal cortex and anterior cingulated cortex. The only other thing you should need to know is that the frontal lobe is commonly subdivided into the prefrontal cortex (corresponding to the frontal gyri at the front of the frontal lobe) and the precentral cortex (corresponding to the precentral gyrus at the back of the frontal lobe). Now you can understand terms such as dorsolateral prefrontal cortex and ventromedial prefrontal cortex.

Now for the natural features. Recall that the cerebral cortex is like a crumpled sheet. The bits of the sheet that stick out are called **gyri** (singular: **gyrus**), while the concavities between are called **sulci** (singular: **sulcus)** or **fissures**. These are the mountains and rivers of the brain, and they provide a natural way to pinpoint specific locations of the cerebral cortex. For example, the cingulated cortex including the cingulate gyrus and cingulate sulcus runs along the bottom of the frontal and parietal lobes, taking in Brodmann areas 23, 24, 26, 29, 30, 31 and 32.

Mixing both artificial and natural features gives us a further collection of names, some of which you may be familiar with, including the amygdala, hippocampus and insular. These tend to be towards the middle of the brain and not necessarily in the cerebral cortex. The **amygdala**, for instance, consists of a group of **nuclei** (with a nucleus being a compact collection of neurons) at the anterior pole of the temporal lobe. The **hippocampus** is a cortical structure lying inside the temporal lobe. The **insula** is a portion of the cerebral cortex folded deep between the frontal and temporal lobes.

The final thing we need to do is distinguish parts of the basal ganglia, the area inside the cerebral cortex. Most important are two nuclei called the **caudate** and **putamen**, which are collectively referred to as the **striatum**.

Being able to pinpoint specific locations of the brain is critical in neuroscience, and so understanding the above terminology is useful in following the literature of

neuroeconomics. It is not really necessary, however, for us to know exactly where, say, Brodmann area 41 is; we can leave that to the neuroscientists. More important for us is to know what different areas of the brain do. Tables 9.1 and 9.2 give you some idea as to basic functions. What I want to do now is introduce some functions of more specific interest to economists. In order to do that, it's necessary to talk a bit more about brain processes.

9.1.2 Brain processes

Neuroscience is not just about filling in a map of the brain. It also helps us understand the processes by which decisions are made. To see why, it is useful to think of brain activities as fitting within a two-dimensional framework in which we distinguish between controlled and automatic, and between cognitive and affective, processes. This is summarized in Table 9.3.

We have already briefly come across the distinction between controlled and automatic processes in Chapter 2 when comparing perception, intuition and reasoning. **Controlled processes** involve step-by-step logic, in which someone consciously thinks through some problem or about some event. **Automatic processes** are

Table 9.2 Parts of the cerebrum and some of the main functions that go on in that part.

Parts of the cerebrum	Functions include
Frontal lobe	Controlling movement
	Organizing, planning and decision-making
Parietal lobe	Bodily sensation, including aspects of vision
	Aspects of language
	Attending to stimuli
Temporal lobe	Hearing
	Higher-order visual tasks
	Recognizing stimuli
Occipital lobe	Initial processing of visual information
Amygdala	Aspects of emotional processing
Insula	Visceral and autonomic functions, including taste
Hippocampus	Certain kinds of memory

Table 9.3 Four different brain processes with an example.

	Cognitive processes (free of emotion)	Affective processes (involving emotion)
Controlled processes		
Step-by-step conscious	Thinking through a math problem	An actor recalling being angry to better play his role in a play
Automatic processes		
Parallel subconscious	Playing a golf shot	Jumping when hearing a loud noise

not accessible to consciousness and go on in parallel, allowing many tasks to be performed simultaneously. A person will typically have little conscious idea, during or afterwards, how automatic processes work.

Affective processes involve emotions that typically result in a desire or drive to do something. For example, Rob feels fear, and so runs away, or feels hungry, and so eats. We can simply think of **cognitive processes** as everything else, so that thought processes are free, or almost free, of emotion – for example, Rob driving his car or preparing his dinner.

Common sense tells us that most of what we do must involve automatic processes. As I write this, I am breathing, my heart is pumping, my fingers are pressing the keyboard, my eyes are watching the screen and the fields outside, I'm listening to the cars and birds, and thinking what to write, plus a lot of other things. Clearly, there is no way that this could all be done consciously, and so most of it goes on at a subconscious level.

For this to work, we need to learn how to do things automatically, because, the more we can do automatically, the more we can do. It is possible to do quite complicated things automatically, such as hit a golf ball, or play a piano concerto. More interesting for our purpose is the reverse logic. We cannot possibly consciously decide what to do using controlled processes, and so we must subconsciously decide what is 'important enough' to warrant conscious control. This leads to the idea of executive control.

In the absence of executive control we could imagine that behavior would be determined by a **default mode**. This is basically what would happen if we stopped to think consciously; life would go on with everything happening according to automatic processes. Young children and some adults with brain damage are thought to essentially be like this. Most of us, though, are not like that. Instead, we have **executive control systems** that sometimes cut in and give us control to move away from default mode.

It's natural, therefore, to think of automatic processes as inevitable and the norm, while control processes are rarer and happen only when automatic processes let them happen. This is important, because it goes a long way to justifying the important chain of reasoning that I highlighted in Chapter 2: framing and context influence intuition and perceptions, which influence reasoning. On first encountering an event, automatic processes, both cognitive and affective, get to work, and it's only subsequently that controlled processes are given a role.

So, whether Rob feels a loss or a gain, feels happy or angry, or perceives something as kind or unkind will happen automatically and subconsciously. This must then influence the starting point for any subsequent reasoning that he does about how to behave. Hence, we observe context and framing effects. Stylized though it is, this is a nice way to think about economic behavior, and one good thing about neuroscience is that it allows us to put some details to the idea. In particular, we can hope to understand a bit more about how automatic and controlled processes work and interact. For that, we need to know more about executive control.

9.1.3 Executive control systems

To work well, executive control systems need to perform a few intermediate processes. Most important is when to switch from automatic to control. Following that

there is the need for inhibition and task switching. **Inhibition** is the suppression of automatic behavior no longer deemed appropriate while **task switching** involves starting any new behavior. An important part of this is likely to be **simulation**, in which the possible consequences of different actions are worked through.

To illustrate, one test or means to measure executive control is the **Wisconsin card-sorting task**. (You can find free versions of this test on the Web, and it takes only around 20 minutes, so maybe you might want to have a go before reading on.) The task requires participants to sort a deck of cards according to some rule. The color, shape and number of the symbols on the cards are different and determine the rule. For instance, the rule might be to sort cards according to color, or to sort according to number of symbols. Subjects are not told the rule, but are told whether they are sorting the cards correctly, so they can learn by trial and error what the rule is. The key to the test is that, without the subject being told, the rule periodically changes. For instance, it might go from sorting by color to sorting by shape. Of interest is whether, and how quickly, subjects adapt to the new rule.

This task means that subjects learn something well enough for it to become automatic, and then at some point have to realize that the default mode is no longer working, inhibit the old rule and switch to a new one! Some of this might happen at a conscious level, but executive control need not be conscious. For instance, while driving along the highway, Rob might swerve to avoid something before he knows he's done it.

There are a few areas of the brain thought to be involved in executive control processes, including the dorsolateral prefrontal cortex and ventromedial prefrontal cortex. These areas stand out as prime candidates because they are linked to many other areas of the brain, and so have the ability to control. Brain imaging allows us to confirm their role, as summarized in Table 9.4. That these areas are important in executive control means that they are going to be important areas of interest to economists, and, although the functions might look a bit abstract at first sight, underlying them are things we need to think about. For example, simulating behavioral consequences and learning the connection between stimuli and reward are fundamental parts of making economic decisions.

In following this up there are three things that I now want to look at in turn, as we start to focus more on economic behavior: how outcomes are evaluated; how people learn about the value of outcomes; and, finally, how people make decisions.

Table 9.4 Main components of executive control systems.

Parts of the brain	Functions include
Dorsolateral prefrontal cortex (DLPFC)	Initiating and shifting behavior – e.g. to switch task or change action Inhibition, in the sense of getting rid of irrelevant information – e.g. forgetting the old rule in Wisconsin card-sorting task Simulating behavioral consequences and abstracting from reality
Ventromedial prefrontal cortex (VMPFC)	Inhibition, in the sense of adhering to rules of behavior – e.g. behaving appropriately in social situations Learning when things change by following the link between stimuli and reward
Anterior cingulated cortex	Noticing conflict and deciding whether to run default mode or not
Basal ganglia	Learning the connection between stimuli and reward

9.2 Valuing rewards and learning

In this section we are going to start looking directly at economic behavior but with a focus on automatic processes, primarily the valuation of rewards and learning. This will be a precursor to looking at the more involved controlled processes that are essential in economic decisions. The natural place to start is with the evaluation of rewards.

9.2.1 Reward evaluation

The basic question I want to start with is how people process rewards. For example, if Rob wins $10, or loses $10, what goes on in his brain? The action appears to start in the midbrain. The neurotransmitter **dopamine** is heavily involved in the evaluation of rewards, and there are two important structures in the midbrain that are associated with dopamine. These are called the **substantia nigra** and **ventral tegmental area** (**VTA**). Most important for decision-making is probably the VTA, because this has pathways projecting into other important areas of the brain. In particular, it has pathways to the frontal lobe, and pathways projecting to areas important for emotional and affective processing, such as the amygdala, hippocampus and other areas of basal ganglia. So, you might want to think of the VTA as the hub, sending out messages to other areas of the brain about what has happened.

In principle, this means we should keep track of what happens in the VTA and all the areas that connect to the VTA. But you may well have had enough of terminology by now, and so I am going to make life a bit easier for us. In order that we can focus unencumbered on the economics, I will primarily refer from here onwards to just two areas of the brain, summarized in Table 9.5. First is something called the nucleus accumbens, which I shall call the **NAcc**, and second is the ventromedial prefrontal cortex, which I shall call the **VMPFC**. Why these two areas? Both connect to the VTA; the nucleus accumbens is part of the basal ganglia – i.e. is inside the cerebrum – and is involved with reward evaluation, while the VMPFC is part of the cortex – i.e. on the outside of the cerebrum – and is important in executive control. Keeping watch on these two areas will, therefore, give us a good idea of what is going on more generally.

(Just for the record, I am going to think of the nucleus accumbens and something called the ventral striatum as the same thing, and the VMPFC and something called the orbitofrontal cortex as the same thing. They are not quite the same things, but that should not matter here.)

Let's start looking at some data. I will begin with a study by Breiter and co-authors (2001). In each round of the study a subject was shown a circle split into

Table 9.5 Two areas of the brain that connect with the VTA and are important in economic valuation and decision-making.

Area	Description
Nucleus accumbens (NAcc)	In the caudate nucleus, in the basal ganglia – i.e. on the inside of the cerebrum
Ventromedial prefrontal cortex (VMPFC)	Part of the frontal lobe and part of the cortex – i.e. on the outside of the cerebrum

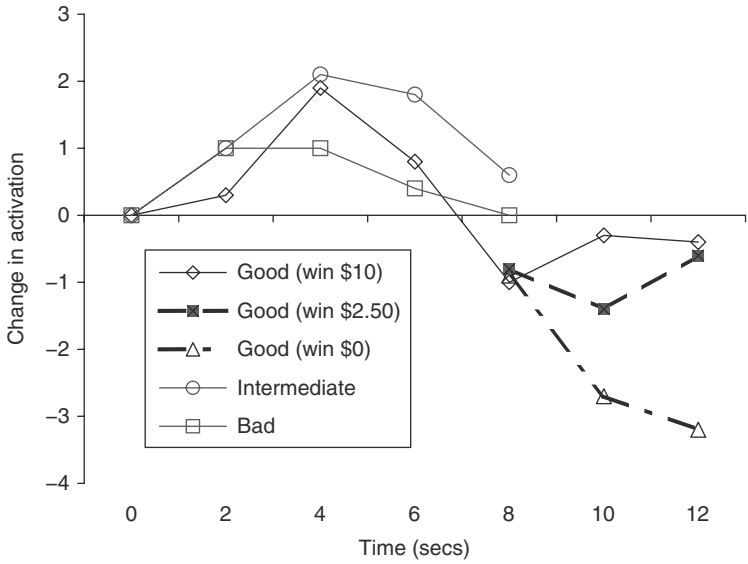

Figure 9.4 Activity in the nucleus accumbens on seeing the circle, around two seconds, and knowing how much is won, around eight seconds.
Source: Breiter *et al.* (2001).

three equally sized sectors labeled with an amount of money. For example, the circle might have one sector labeled $10, one labeled $2.50 and one labeled $0. An arrow would then spin and finally land on one of the sectors, with the subject winning the corresponding amount of money. So, if the arrow landed on the $10 sector, the subject won $10. There were different circles used: a good one (with labels $10, $2.50 and $0), an intermediate one ($2.50, $0 and –$1.50) and a bad one ($0, –$1.50, and –$6). Of interest is to see how subjects reacted to seeing which circle would be used and how they reacted to knowing how much they had won.

Figure 9.4 gives us some idea of what happened in the NAcc, and Figure 9.5 does the same for the VMPFC. We are going to see a few of these types of figures, so let me explain what they show. The vertical axis shows the change in signal from the relevant area of the brain relative to some baseline level. So, a positive suggests that the area is more activated than normal, and a negative that it is less active than normal. The horizontal axis plots time relative to some start point. In this case, the subject saw the circle being used after around two seconds and saw what they had won after eight seconds. We can, therefore, track changes in activation as things happen.

To keep the figures manageable, I have plotted only outcomes for the good wheel, but in both the NAcc and VMPFC we see activation on seeing the circle and then also activation when the outcome is known. Further, the activation appears related to the desirability of the circle used and amount won. The better the circle, the more activation is observed, and, the better the outcome, the more activation is observed. One thing to note for future reference is how the activation in the VMPFC on seeing the outcome is positive, while that in the NAcc is negative.

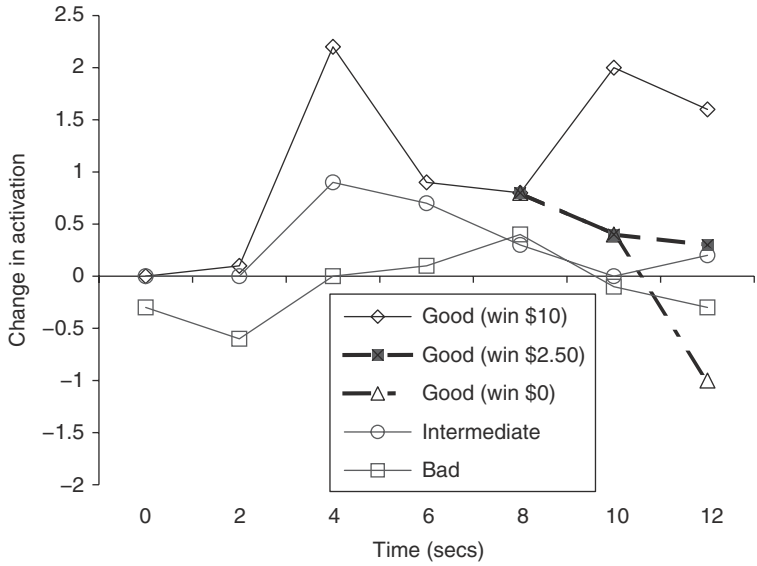

Figure 9.5 Activity in the ventromedial prefrontal cortex on seeing the circle, around two seconds, and knowing how much is won, around eight seconds.
Source: Breiter *et al.* (2001).

We need to be careful of reading too much into one study (see Research methods 9.3), but the results I have given above are representative of what is observed in other studies, namely that the NAcc and VMPFC respond to rewards and might respond by an amount related to the size of the reward. (We shall see some more examples of this as we go on.) Maybe, therefore, we have found the utility function! That would be very premature to say, but it is an exciting possibility. To delve further, we can look at how the brain learns about rewards.

❖ RESEARCH METHODS 9.3 SAMPLE SIZES AND CONFIDENCE INTERVALS

One criticism of neuroeconomics is the smallness of the sample sizes sometimes used. For example, the study by Breiter and co-authors (2001) used 12 subjects. This is fairly typical in neuroscience. Is 12 enough? To answer that, we need to enter the murky world of statistical tests and confidence intervals.

To get us started, think about an experiment, neuro or not, with the objective of finding what proportion of people use one heuristic or another. For example, we might be interested in what proportion of people use level one or level two reasoning (see Chapter 6). We invite subjects to the lab and get them to play games over a number of rounds. From this we need to (1) estimate the heuristic used by individual subjects and (2) estimate what this tells us about the general population. Then we find out what we want to know.

The word 'estimate' here is important, because we cannot know for sure the answer to either of these two questions, and that's why sample size matters. I am not going to go into the technicalities of confidence intervals, but, basically, (1) the more rounds a subject plays, the more confident we can be in our estimate of what heuristic he is using, and, (2) the more subjects we ask to the experiment, the more confident we can be in our estimate of what happens in the general population. So, we need lots of rounds and lots of subjects. Fortunately, the law of large numbers is on our side, and so 'lots' need not mean 'ridiculously many', but, the more, the better, and something like 100 is great, 50 not bad but 10 too few.

In a normal experiment we can easily invite 100 subjects and have them play as much as we want. The cost involved in brain imaging makes that impractical in neuroscience, however. So, although we can confidently estimate what an individual subject is doing, it is harder to confidently generalize any results to the wider population. It's not really clear at the moment whether this is a problem or not. Going on statistical logic it is a problem, but if we think that brains are fairly homogeneous across the wider population it need not be.

You might be wondering how this relates to the more standard economic experiment involving treatments. For instance, we might have one treatment in which subjects play an ultimatum game to split $10, and another in which they play an ultimatum game to split $100, and we are interested if behavior is different in the two treatments. Some might say that we just need to have enough subjects to spot a difference, because, if we observe a difference in, say, the 100 subjects we invite to the lab, there must be a difference in the general population. Economists are prone to become fixated with such statistically significant differences. Of more interest, though, is usually the magnitude of any difference. To estimate this with confidence, the more subjects, the better.

9.2.2 Learning about rewards

To get us started, imagine a sequence of events, such as Rob putting a bet on a roulette wheel, seeing the roulette wheel stop on his number, getting the cash he won, buying a packet of chips with his winnings and then eating the chips. Ultimately, Rob should be happy. What I want us to think about now is when he feels happy. Is it when he bets, wins or eats the chips?

This brings us to the distinction between primary and secondary reinforcers. You can think of a **primary** or **unconditioned reinforcer** as something that has been 'hard-wired' into Rob's brain, by evolution, to be important, such as food, water and sex. You can think of a **secondary reinforcer** as something that Rob has learnt to indirectly associate with a primary or other secondary reinforcer, such as money. In the previous chapter we saw that evolution could have favored lots of possible intermediate goods, and so the distinction between a primary and secondary reinforcer is not necessarily an easy one. The basic principle, however, is simple enough: Rob

does not need to learn about a primary reinforcer, but he does need to learn about a secondary reinforcer.

For Rob, food is a primary reinforcer, while money is a secondary reinforcer, and seeing the roulette wheel stop, or start, are even more removed secondary reinforcers. So, when does he start to get excited at winning?

To illustrate what might happen, I will briefly mention a study by O'Doherty and co-authors (2002). The study involved subjects seeing a visual cue before getting 0.5 milliliters of either a pleasant taste, an unpleasant taste or a neutral taste. Some areas of the brain, including the VTA and NAcc, were activated when the subject saw the cue, rather than felt the taste. Other areas, however, including the VMPFC, reacted more to the taste than to the cue. To take things one stage further, we can look at a similar study by Knutson and co-authors (2001). This study involved money rather than taste, but the results are similar. On seeing the cue of a future payment, areas such as the VTA and NAcc were activated, but areas of the VMPFC reacted more on seeing the reward than the cue.

To put a few more details to this, Figure 9.6 plots the timeline as subjects saw the visual cue at zero and then found out whether they had won money or not. Focus first on what happens if they won money. Then we see activation of the NAcc when they were anticipating winning, but not on knowing they had won. There was no such effect in the VMPFC. On finding out they had not won there was negative activity in both the NAcc and VMPFC. You might now want to check back and see what happens in Figures 9.4 and 9.5.

A stylized story of what is happening is that some areas of the brain, such as the NAcc, learn to treat secondary reinforcers like primary reinforcers, while other

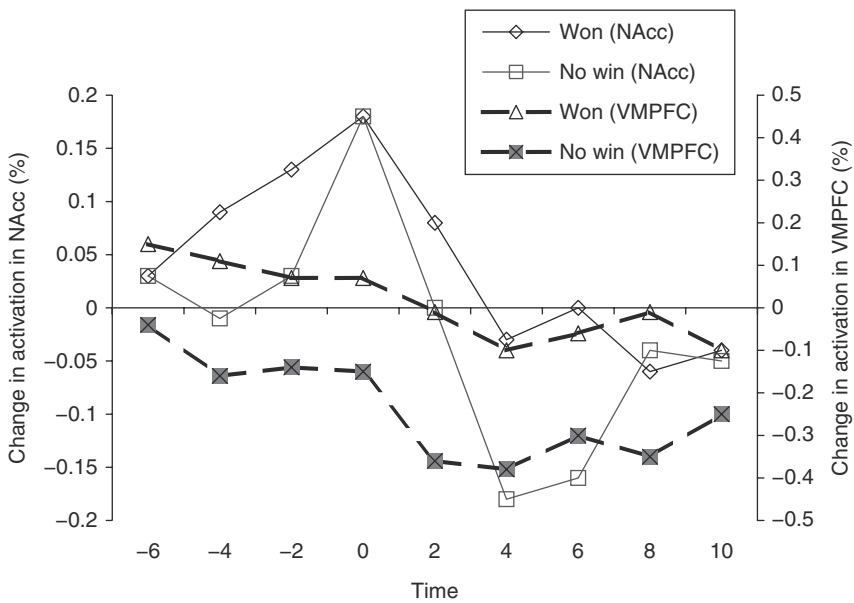

Figure 9.6 The response in the nucleus accumbens and ventromedial prefrontal cortex to anticipating, and then winning, or not winning, money.
Source: Knutson *et al.* (2001).

parts, such as the VMPFC, keep a check on things and so react more to receiving the reward than expecting to receive it. This fits the story of executive control that I summarized in Table 9.4, with the NAcc, in the basal ganglia, and VMPFC updating ongoing predictions about how much reward can be expected in the future from a particular secondary reinforcer.

An interesting example of this is the **Iowa gambling task** (and, again, free versions are available on the Web if you want to try before reading further). This task involves four decks of cards, which subjects are repeatedly asked to choose from. Every card from decks A and B pays a large reward (e.g. $100) but sometimes also entails a large loss (e.g. $1,250). Every card on decks C and D pays a small reward (e.g. $50) and only sometimes entails a small loss. Overall, decks A and B have negative expected value, while decks C and D have positive expected value, and most subjects soon choose only from decks C and D. Those with damage to the VMPFC, however, continue to choose from decks A and B. It's as though the VMPFC is not doing its job of checking whether decks A and B are the good choices that the NAcc thinks they are.

On a related theme, I next want to mention a study by O'Doherty and co-authors (2001). The task that subjects did was partly like the Iowa gambling task, in that they had two choices: option A, which gave large losses and small rewards, or option C, which gave large rewards and small losses. Once subjects had learnt that C was the best choice, however, the reward structure was reversed, similarly to the Wisconsin card-sorting task. Figure 9.7 illustrates what happened in two areas of the

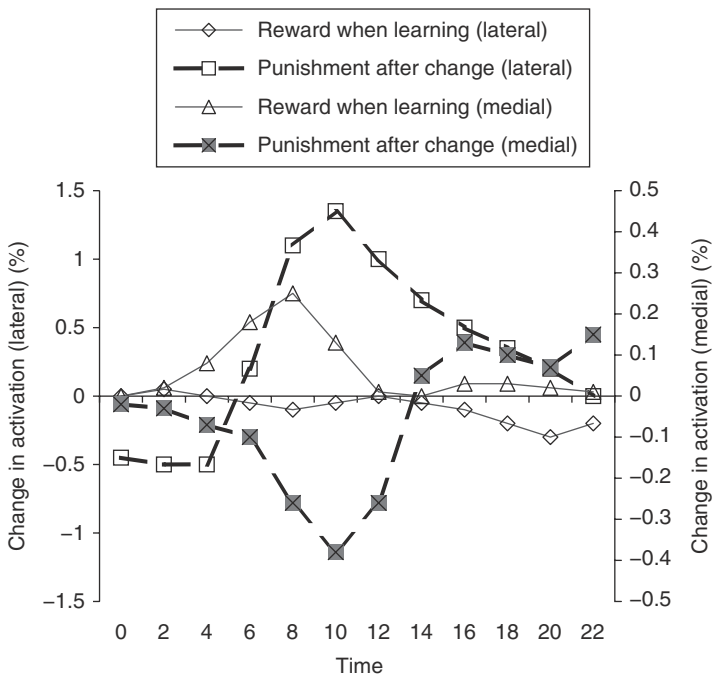

Figure 9.7 The change in activation in two areas of the VMPFC following a reward during a learning stage, or loss in a task reversal stage.
Source: O'Doherty *et al.* (2001).

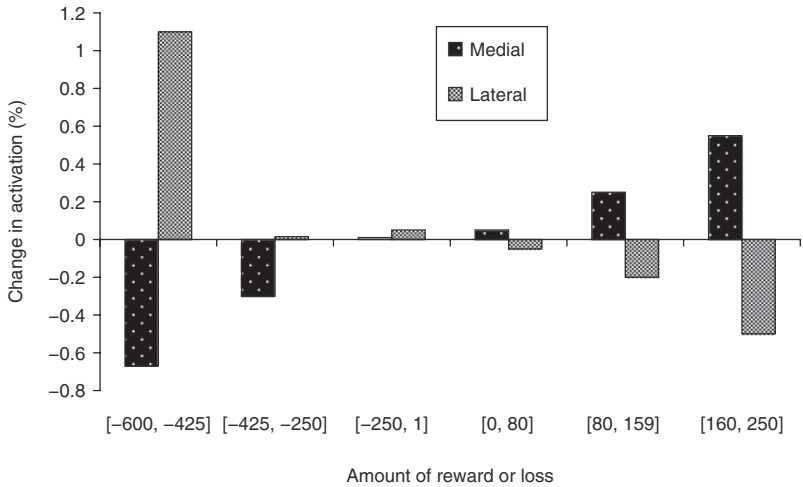

Figure 9.8 The change in activation in two areas of the VMPFC following a reward or loss of varying magnitudes.
Source: O'Doherty *et al.* (2001).

VMPFC while the subjects were learning that option C was best, and then after the change to option A being best.

The interesting thing here is how one area of the VMPFC (which I have labeled 'lateral') was activated following a loss and another (labeled 'medial') was activated following a reward. Figure 9.8 takes this one stage further, by showing how the activation correlates with the magnitude of the reward or loss. Not only is this consistent with the VMPFC being involved in learning but it also suggests that different areas of the VMPFC are activated for rewards and losses. If losses and gains affect different areas of the brain, then it's no surprise that this distinction matters when modeling economic behavior.

9.2.3 Risk and uncertainty

One important issue that we need to address is that of risk. In most of the studies I have mentioned so far there has been risk. For instance, in the study by Breiter and co-authors (2001), subjects saw the arrow spin to decide what reward they would get. This raises the question of whether it's the expected reward that matters, and how people learn when the outcome of an event is unknown. For example, what would be happening in Rob's brain between the roulette wheel starting to spin and stopping?

I will start by looking at a study by Preuschoff, Bossaerts and Quartz (2006). In the study, subjects were asked to bet whether the second of two cards would have a higher or lower number than the first, the numbers ranging from one to ten. They then saw the number on the first card, followed by the number on the second card. The gap between seeing the two numbers was long enough for subjects to realize their chance of winning. For example, if the subject said the second card would be

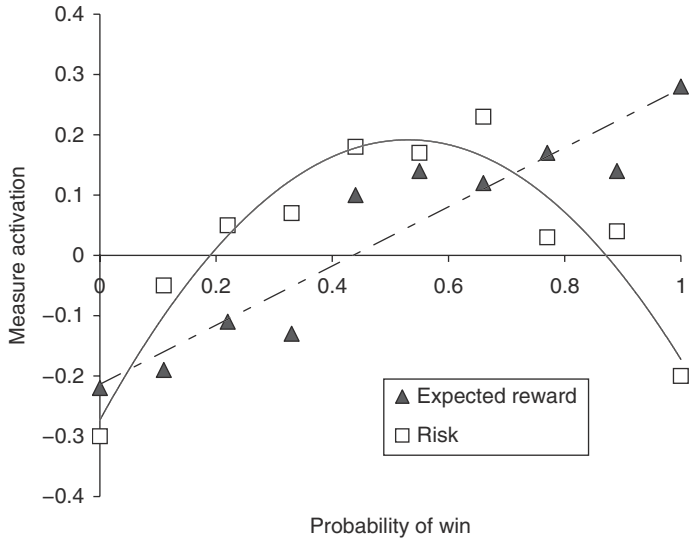

Figure 9.9 Activity in the NAcc (left ventral striatum) increases with expected value and risk.
Source: Preuschoff, Bossaerts and Quartz (2006).

lower, and the number on the first card was nine, then they had a good chance of winning. If the number was five, the chance of winning was less and the risk was higher.

Figure 9.9 plots a measure of activation in the NAcc against the expected reward and the level of risk, where risk is highest if the probability of a win is 0.5. We see that activation is increasing in expected reward and is higher when there is more risk. A similar story has been told for the VMPFC. The brain, it seems, has no problem working with risk and expected values. What about uncertainty and ambiguity?

Recall that, with a risky choice, the probabilities of different outcomes are known, but with uncertainty they are not. For instance, suppose Rob is asked to bet on whether a red or blue card will be randomly drawn from a deck of 20 cards.

If he knows that ten of the cards are red and ten blue then he has a choice with risk. If he does not know how many are red or blue then he has a choice with uncertainty. In a study by Hsu and co-authors (2005), subjects were given both risky and ambiguous choices like this in order to see whether uncertainty made a difference.

Figure 9.10 plots the timeline of activation in the NAcc and an area of the VMPFC. The subject would see the choices at zero seconds, and the average decision what to choose was made at around six to seven seconds. Look first at what happens in the NAcc with a risky choice as compared to a choice with uncertainty. Clearly, once a decision has been made, the NAcc is activated much less when there is uncertainty, suggesting that the (expected or subjective) value put on an uncertain outcome is less than a risky outcome. Subjects, it appears, did not like uncertainty. Now look at what happens in the VMPFC. Here we see greater activation when there is uncertainty and when the subject is deciding what to choose. A similar pattern was observed in the amygdala, suggesting that some of the areas of the brain were activated to 'try to find the missing information'.

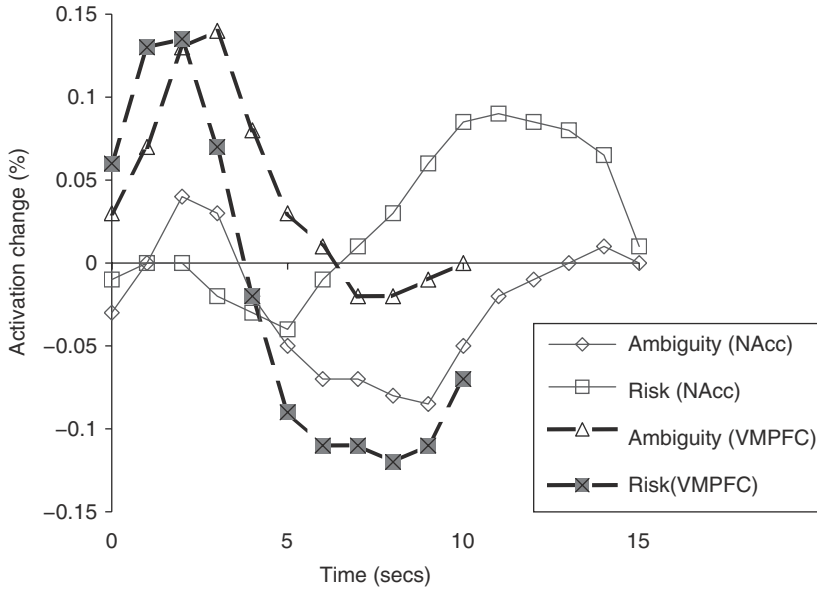

Figure 9.10 Activation in the NAcc (left striatum) and VMPFC (left orbitofrontal cortex) depends on whether the choice is with risk or uncertainty
Source: Hsu *et al.* (2005).

The picture we get, therefore, is that when there is uncertainty the brain is alerted to the fact that information is missing, and because of this a choice with uncertainty is not valued as much as a choice with risk. Indeed, it seems that a choice with uncertainty was given a very low expected reward. That fits with what we looked at in Chapter 5 and the suggestion that people are ambiguity-averse. (We don't know yet what happens with compound lotteries.) I will come back to the differences between risk and uncertainty shortly, but I want to finish this section by looking a bit more at different types of reward.

9.2.4 Different types of reward

We have seen that brain activation can be nicely correlated with expected reward. An important concept in taking this idea further is **reward prediction error**. Basically, this is the difference between the reward expected and what was received. So, if Rob expected to win $5 and wins $0 or $10, his prediction error is –$5 or $5. Reward prediction error is clearly important in determining how excited Rob should be and how he should learn for the future.

In the studies that we have looked at so far, it is difficult to distinguish how the areas of the brain react to reward prediction error as opposed to expected reward. It seems, however, that the NAcc is activated in proportion to the reward prediction error. To illustrate, I will look at a study by Hare and co-authors (2008).

In the study, before going into the scanner, subjects stated a willingness to pay for 50 food items and were given some 'spending money'. Once in the scanner, they

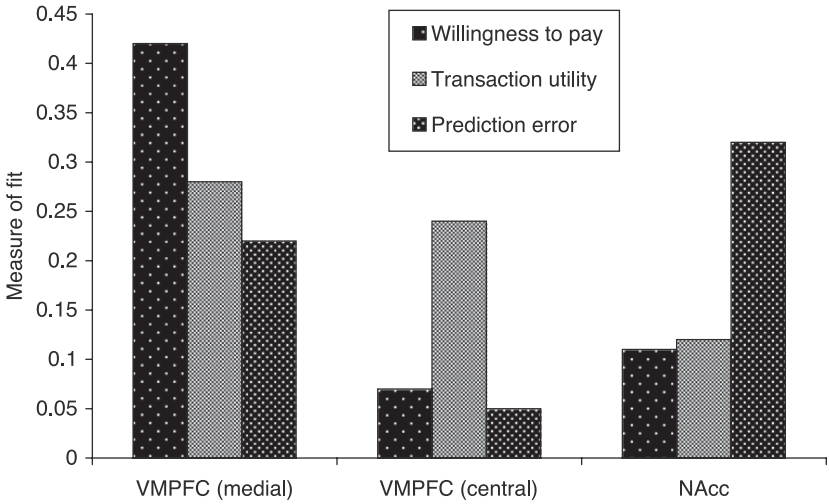

Figure 9.11 The fit between values and activity in three areas of the brain.
Source: Hare *et al.* (2008).

were given the chance to say 'Yes' or 'No' to purchasing food items at a randomly determined price while also learning that their spending money had randomly been increased or decreased by some amount. For example, a subject might be presented with a screen displaying a candy bar, for sale at price $2, and something to say he or she had lost $1 of spending money.

The purpose of doing this was to distinguish three things that could influence the subject's reaction and choice. First, there is the willingness to pay for the good; second, there is the difference between price and willingness to pay; and, third, there is the reward prediction error. You might recall that in Chapter 2 I gave these first two things the names 'acquisition utility' and 'transaction utility'. Figure 9.11 illustrates how much activity in three different areas of the brain correlated with the three different values. The picture we get is, again, consistent with different areas of the brain reacting differently to these values. The NAcc is activated by the prediction error, and areas of the VMPFC by the willingness to pay and transaction utility.

I think we have already seen some exciting results. We have seen that the distinctions made earlier in the book, between losses and gains, risk and uncertainty, and acquisition and transaction utility, really do make sense, not only in trying to model choice but as a description of how people think. The next step is to see if this translates into choice and decision-making.

9.3 Making decisions

So far the focus has been on people passively reacting to events largely out of their control. This firmly puts us in the domain of automatic processes, even if with some executive control. Now it's time to go up a gear and look at what happens when people have some choice, and so use reasoning and control processes.

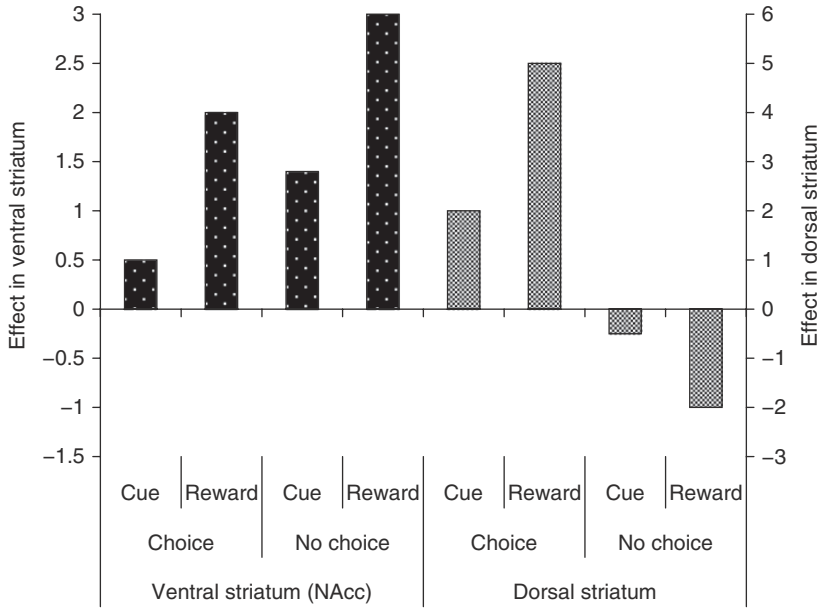

Figure 9.12 Activation in the dorsal and ventral striatum during a learning task on seeing the options or cue and getting the reward. Activation in the dorsal striatum depends on whether there is choice or not.
Source: O'Doherty *et al.* (2004).

To get us started, I want to look at a study by O'Doherty and co-authors (2004) that nicely illustrates the potential differences between passive learning and choice. In the study, subjects were repeatedly exposed to two options before getting either a tasty juice or a neutral reward. One of the options came with a 60 percent chance of getting the juice and the other with a 30 percent chance of getting the juice. So, subjects had to learn the correlation between option and reward. More interesting for our purpose was that sometimes subjects chose an option and sometimes the computer chose for them.

The main results of the study are summarized in Figure 9.12. Focus first on the left-hand side and the ventral striatum. For our purposes the ventral striatum and NAcc are the same, and, as we have now come to expect, this area was activated in proportion to the expected reward and reward prediction error. Now look at what happens in the dorsal striatum. This time there is activation when the subject has the chance to make a choice, but not when the computer will decide. This relationship is consistent with an **actor–critic model**, in which the ventral striatum is cast in the role of critic, evaluating rewards, while the dorsal striatum is cast in the role of actor, selecting what to choose.

9.3.1 Choice and strategy

We have seen that choice does make a difference to what goes on in the brain, with other areas activated to seemingly help make the decision. An intriguing issue is

whether we can use this to predict choice and strategy. To show that this idea may not be too crazy I will look at two studies. In both, subjects had to choose between risky or uncertain gambles.

In a study by Huettel and co-authors (2006), subjects were presented with pairs of gambles or prospects of which some gambles involved risk and some uncertainty. For example, they might have to choose between a risky gamble that pays out \$12 with probability 0.75 and \$20 with probability 0.25, and an uncertain gamble that pays out \$0 or \$35 with unknown probability. We thus move away from asking how subjects respond to seeing risk or uncertainty to how they choose between the two.

For each subject, having seen what they chose, it is possible to construct indices of how much they seem to like (or not) ambiguity and risk. We can then see whether these preferences correlate with the brain activation observed while the subject was making choices. That is, we are trying to find a link between brain activation and choice. Figure 9.13 suggests that we can find such a link, though we are going to have to look at areas of the brain other than the NAcc and VMPFC. Figure 9.13 summarizes how ambiguity led to the greatest activation in the lateral prefrontal cortex for those whose choices suggest they like ambiguity. Similarly, risk led to the greatest activation in the parietal cortex for those whose choices suggest they are less risk-averse.

We see, therefore, that not only do risk and ambiguity activate different areas of the brain, but by seeing how much the brain is activated we can potentially predict choices. Can we go one stage further and predict strategy?

To see if we can, I will look at a study by Venkatraman and co-authors (2009). In the study, subjects were first presented with a risky prospect containing five possible outcomes – for example, to win \$80 with probability 0.25, lose \$75 with probability

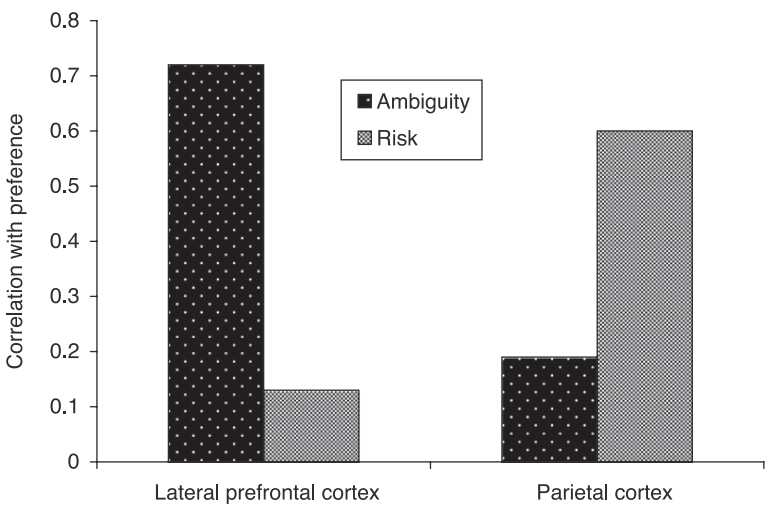

Figure 9.13 Areas of the lateral prefrontal cortex and parietal cortex are more activated in those with a preference for ambiguity and risk.
Source: Huettel *et al.* (2006).

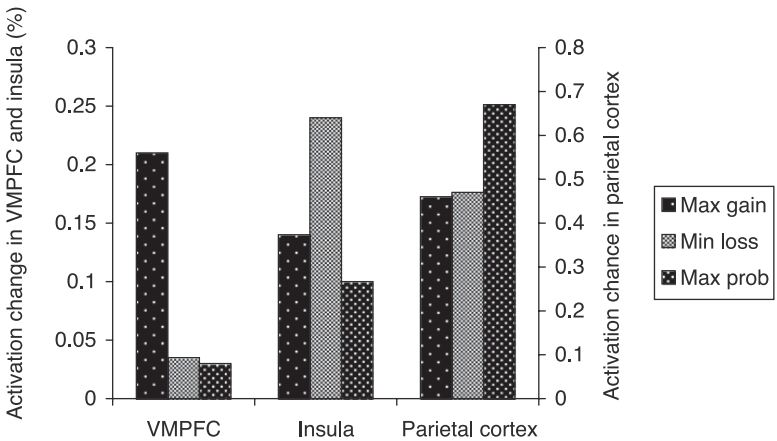

Figure 9.14 Different choices correlate with increased activation in different areas of the brain.
Source: Venkatraman *et al.* (2009).

0.2, and so on. They were then given two alternatives to improve the prospect and had to choose which they wanted. For example, the alternatives might be to make it a $100 win with probability 0.25, or a $55 loss with probability 0.2.

There are two basic types of strategy that one might employ in this scenario. The 'simple strategy' is to maximize the probability of making a difference, and so, for instance, to add $20 to the $80 outcome because this is more likely to matter than adding $20 to the $75 loss. A more 'complex strategy' would be to maximize gains, by adding $20 to the $80, or minimize losses, by adding it to the $75 loss, or maximize expected utility in some other way.

In standard experiments, without brain imaging, most subjects were observed using the simple strategy, but some did use a complex strategy. The main thing, though, was that subjects seemed to stick to a strategy through multiple rounds. We can, therefore, potentially disentangle choice from strategy.

Let's look first at choice. Figure 9.14 shows that choice was correlated with activation in different brain areas. For example, choosing the option that maximized the possible gain was associated with greater activity in the VMPFC, while choosing the option that minimized the possible loss was associated with greater activity in the insula. So far, so good.

Now let's look at strategy. What we want to do here is see whether some areas of the brain are more activated when someone makes a choice not consistent with their usual strategy. That is, what happens when someone who usually uses a simple strategy decides to do something 'more complicated', such as maximize the possible gain, or when someone who usually uses a complex strategy decides to do the 'simpler thing' of maximizing probability? When this happens the dorsomedial prefrontal cortex becomes more activated, and, what's more, as Figure 9.15 shows, we observe greater connectivity between this and the areas associated with choice.

So, greater activation of some areas correlates with choices made, while greater activation in others correlates with strategy, or with the differences between choice and strategy. It looks, therefore, as though we can watch both choice and strategy

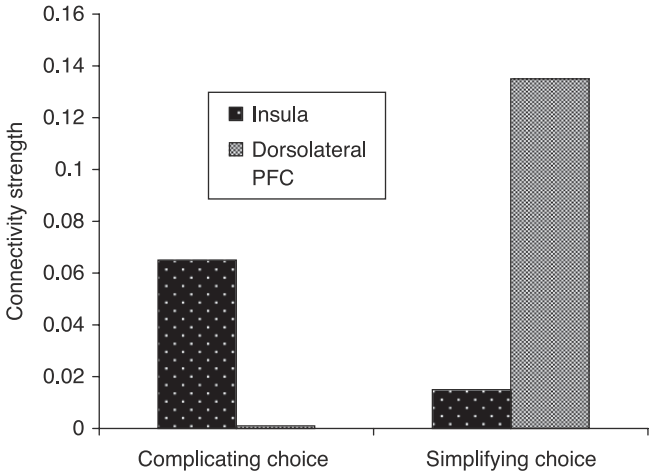

Figure 9.15 Increased connectivity is observed between the dorsomedial prefrontal cortex and the insula and parietal cortex when someone makes a choice inconsistent with their normal strategy.
Source: Venkatraman *et al*. (2009).

being played out in the brain. This is fascinating stuff – fascinating enough that I want to give you a second example.

9.3.2 Framing effects

We have already seen evidence that a gain activates a different area of the VMPFC from a loss. The question I want to look at now is whether this difference can lead to framing effects. Let us look at a study by De Martino and co-authors (2006). Subjects had to choose between a sure amount of money or a risky gamble. For example, they could take $20 or gamble with a 40 percent chance of getting $50 and a 60 percent chance of getting $0. What differed is the way these options were presented. Using a gain frame, the sure thing was framed as a gain of $20. Using a loss frame, the sure thing was framed as a loss of $30.

As we would expect, subjects were more likely to choose the sure option in a gain frame than in a loss frame. Can we pick up this framing effect in the brain? Figure 9.16 shows activation in the amygdala and anterior cingulated cortex. When subjects made the 'simple' choice, the amygdala was positively activated. When subjects made the 'difficult' choice, the amygdala was negatively activated and the anterior cingulated cortex was more activated. Recall (see Table 9.2) that the amygdala is involved with emotional processing, while the anterior cingulated cortex (see Table 9.4) is involved with conflict resolution.

The picture we get, therefore, is of framing and context effects working exactly as discussed in Chapters 2 and 3. Consistent with people being risk-loving for losses (see section 3.1.3), the automatic process is to gamble on the loss frame and not to gamble on the gain frame. This automatic process is potentially inhibited, however, by the executive control system. Such inhibition activates a more reasoned choice, less

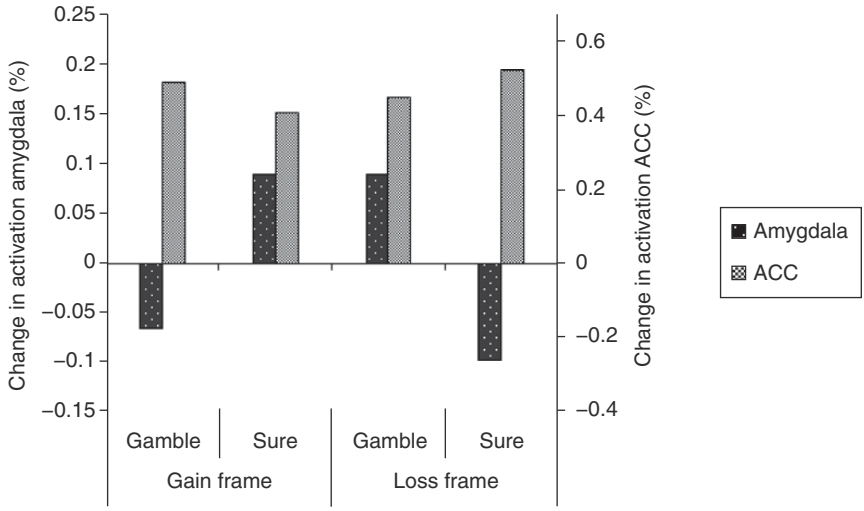

Figure 9.16 Activation in the amygdala and anterior cingulated cortex depending on the frame and choice.
Source: De Martino *et al.* (2006).

subject to the framing effect. Again, we can see choice and strategy being played out in the brain.

What is new is that we can also question the sense of that strategy. Subjects varied a lot in how much they were affected by framing. Some subjects nearly always made the same choices in a gain and loss frame, while other subjects were more affected by the frame. This allowed the authors of the study to construct a 'rationality index', whereby a subject less affected by framing was considered more rational. Interestingly, they found that this rationality index was highly correlated with activity in the orbital and medial prefrontal cortex (OMPFC) and VMPFC. We have not come across the OMPFC yet, but it is connected to the amygdala and known to be important in evaluating the outcomes of decisions. We also know that the VMPFC is involved in inhibition (see Table 9.4).

This suggests that some subjects were better able to recognize the potential emotional bias that the framing effect would create. In other words, the executive control system of 'rational' subjects was more aware that it needed to get involved and inhibit automatic processes. In this case, a subject needed to be aware of their own bias. In strategic interaction a subject needs to think through the strategy and potential bias of others. It is to this that I now briefly turn.

9.3.3 Strategic behavior

Where better to begin looking at strategic interaction than the p-beauty contest (see Chapter 6)? In a study by Coricelli and Nagel (2009), subjects played 13 different p-beauty contests in which p ranged from near zero to near two. Recall (see section 6.2.1) that people playing the p-beauty context differ in their depth of reasoning. A type L-1 expects others to choose randomly, while a type L-2 expects others to be type L-1, and so on.

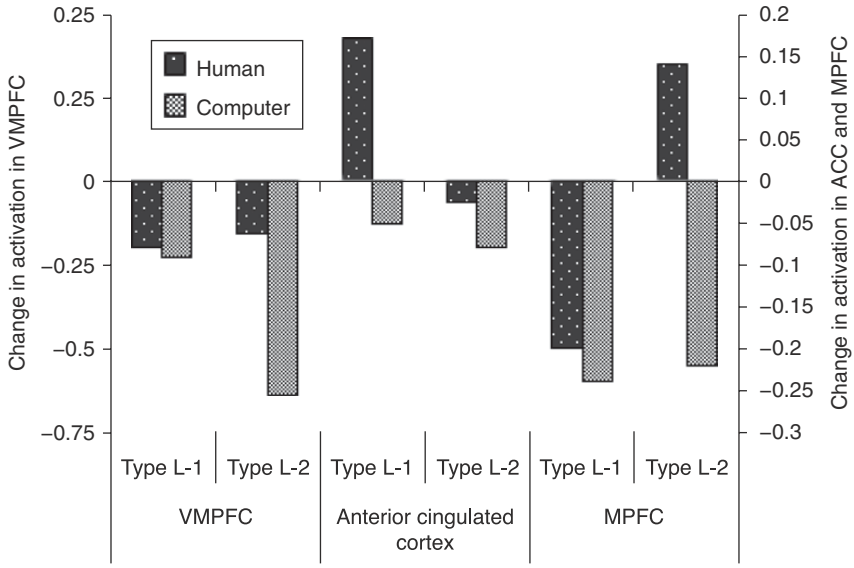

Figure 9.17 Activation in the VMPFC, anterior cingulated cortex and medial prefrontal cortex (MPFC) depending on level of reasoning and whether the opponents are humans or computers.
Source: Coricelli and Nagel (2009).

In order to better evaluate depth of reasoning, Coricelli and Nagel also had subjects play 13 contests with computer opponents, when the computers chose a number randomly. A type L-1 would be expected to behave the same against human opponents as computer opponents. A type L-2, however, would be expected to behave differently against humans as opposed to computers.

Figure 9.17 summarizes some of the ways in which brain activity differed between subjects classified as type L-1 and type L-2 or higher. Let us focus first on the medial prefrontal cortex (MPFC). Notice that this is strongly activated in type L-2 subjects when playing against humans. The MPFC is believed to be involved in **theory of mind** or reasoning through other people's minds. In type L-1 subjects the MPFC is not activated, but the anterior cingulated cortex is. We have focused on the anterior cingulated cortex's role in the executive control system, but it is also associated with thinking about self.

These results provide good evidence that level-k thinking is an appropriate way to model choice in the beauty contest. We see that those subjects classified as type L-2 or above thought through what others would do, while those classified as type L-1 thought more about themselves.

By comparing the distance between a subject's guesses and the guesses of others it is possible to construct an index of 'strategic rationality'. The authors find that this index correlates strongly with activity in the MPFC and VMPFC. This result is very similar to that found for the framing study, looked at in the previous section. The only change is that different areas of the brain are activated to reflect the different task involved. In the framing task, 'rational' subjects had greater activation in the OMPFC,

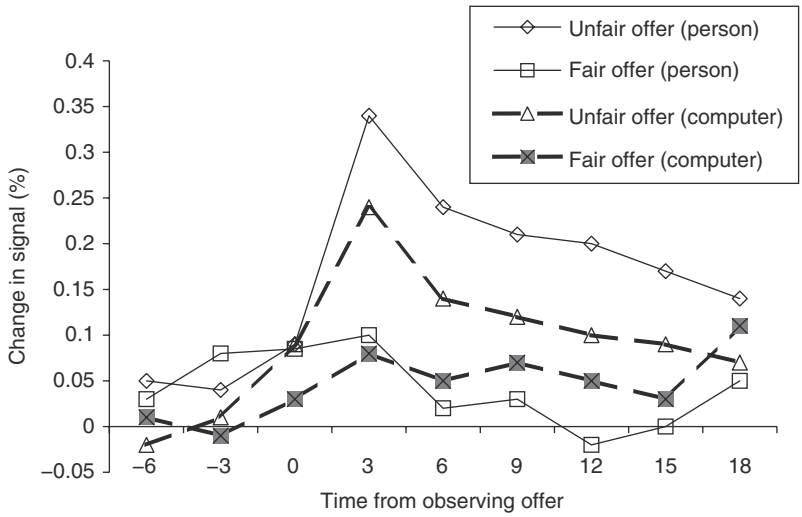

Figure 9.18 Timeline of activation in the (right anterior) insula depending on whether the offer is fair or unfair.
Source: Sanfey *et al.* (2003).

which is associated with evaluating outcomes. In the strategic task, 'rational' subjects had greater activation in the MPFC, which is associated with theory of mind. The VMPFC appears to play a role in both cases.

We have now seen a few examples of how choice and strategy can be picked apart and analyzed. What I want to do now is question how far this ability can take us in understanding social preferences and inter-temporal choice.

9.3.4 Fairness and norms

In getting us started with social preferences, I want to talk about a study by Sanfey and co-authors (2003). Basically, the study involved subjects playing the ultimatum game (see Chapter 7) to see how people react to unfair and fair offers. Subjects, in the role of receiver, were offered proposed splits of $5 and $5, $9 and $1, $8 and $2, and $7 and $3. The main question of interest is how subjects responded to a fair offer, $5 and $5, as compared to unfair offers, such as $9 and $1.

One area of the brain we have yet to talk about is the insula, but now it becomes a lot more relevant. As Figure 9.18 shows, the insula was activated more when a subject received an offer that was unfair rather than fair, and more when the proposer making the unfair offer was known to be a human rather than a computer. There was also greater activation the more unfair the offer. Particularly interesting is Figure 9.19, which compares activation in the insula and dorsolateral prefrontal cortex (DLPFC) for offers that were subsequently rejected relative to those that were accepted. We see that the rejection of an offer comes with relatively large activation of the insula. So, what is going on?

The only mention the insula has had so far is in Table 9.2, where I said that it was involved in visceral and autonomic functions such as taste. In fact, things that are

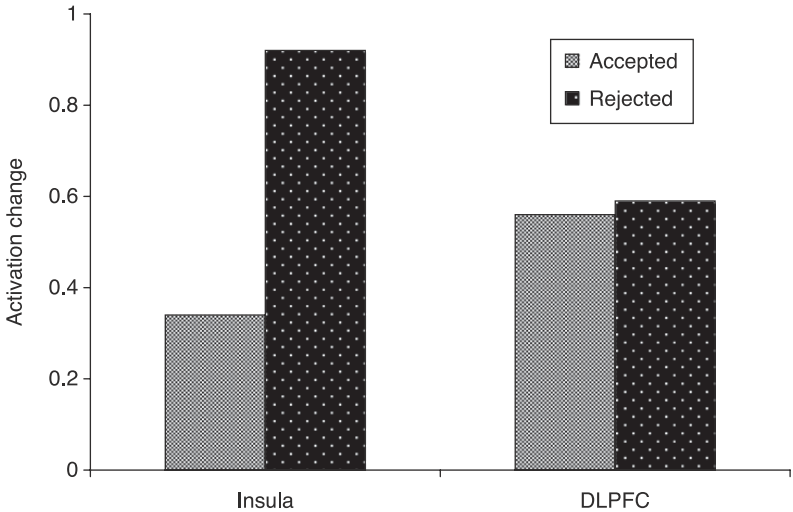

Figure 9.19 Activation in the (right) insula and (right) dorsolateral prefrontal cortex for offers that are subsequently rejected or accepted.
Source: Sanfey *et al.* (2003).

consistently associated with the insula include negative emotional states such as pain and distress, hunger, thirst and disgust from taste or odor. This clearly fits with a story of receivers feeling a negative emotional response on receiving an unfair offer. By contrast, we know the DLPFC is usually associated with cognitive processes and executive control, and so is possibly more focused on getting as much money as possible. This suggests a potential conflict between the insula, that wants to reject, and DLPFC, that wants to accept. Figure 9.19 is certainly consistent with this. Further credence to this idea comes from increased activation of the anterior cingulated cortex for unfair offers. Remember that this is thought to be involved in conflict resolution.

So, we get a neat story. The insula is upset by an unfair offer and wants to reject, while the DLPFC wants the money, so is happy to accept. Sometimes the insula wins out and sometimes the DLPFC. Neat stories often have a catch, though, and this one is no exception. To see why, I want to look at a study by Knoch and co-authors (2006). In this study, low-frequency repetitive transcranial magnetic stimulation was applied to either the right or left DLPFC of subjects. This suppresses activity in that part of the brain. Basically, it switches off either the left or right DLPFC. So, what should happen if we switch off the DLPFC?

If the DLPFC wants to accept any offer, the insula wants to reject an unfair offer and we switch of the DLPFC, we should see more rejection of unfair offers. Figure 9.20 shows that we observe the complete opposite! We see far less rejection of unfair offers when the right DLPFC is switched off than when the left is switched off, or when neither is switched off. This is despite the unfair offer still being perceived as unfair.

The story we get now, therefore, is one in which the DLPFC wants to reject a low offer, because of fairness norms, and other areas of the brain want to accept, and get the money. Consistent with this, stimulation of the right DLPFC had less of an

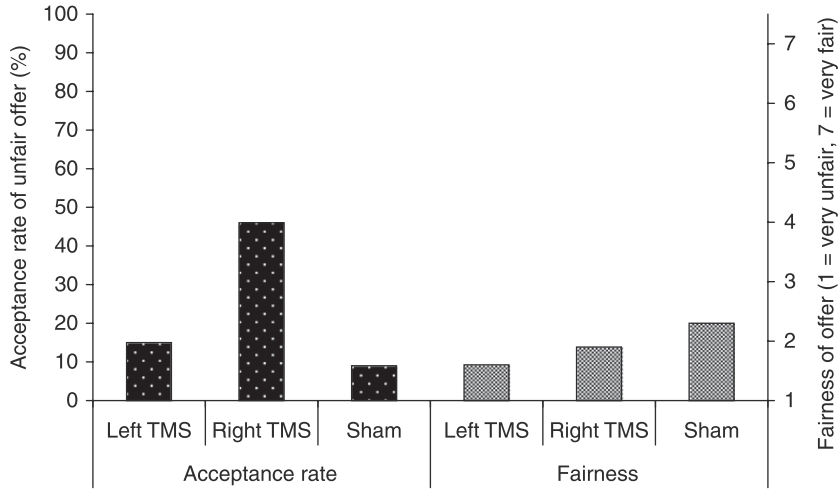

Figure 9.20 Acceptance rate and perceptions of an unfair offer in the ultimatum game. Those with stimulation (TMS) to the right DLPFC are more likely to accept, despite still thinking the offer is unfair. In the sham treatment there is no stimulation. *Source:* Knoch *et al.* (2006).

effect when the low offer was received from a computer rather than human proposer. So, it's not really very clear what's going on, but either way the (right) DLPFC does seem important for reciprocity.

To explore this further, we can look at a study by Spitzer and co-authors (2007). In the study, subjects played two different games several times. One game was a variation on the dictator game, in that one subject, the proposer, was given 125 tokens and a second subject, the receiver, was given 25 tokens; the proposer could give up to 100 of his or her tokens to the receiver. The second game was a variation on the ultimatum game, in that the receiver had the chance to punish the proposer after seeing how many tokens he or she had been given; for every one token that the receiver spent punishing, five tokens would be taken from the proposer. For example, if the proposer gave no tokens to the receiver and the receiver spent 25 tokens on punishing, then both proposer and receiver would end up with zero tokens. Some subjects played the ultimatum game with another subject (and knew it), while some played against a pre-programmed computer (and knew it).

As discussed in Chapter 7, the dictator game gives us an idea of how much a subject would voluntarily give to another, while the ultimatum game gives us an idea of how much he or she would give with the threat of punishment. The focus switches, therefore, more onto the proposer than the receiver. When we compare activation during the ultimatum game to that in the dictator game we see greater activation in several areas, including the DLPFC and VMPFC. Furthermore, activation was higher in these areas, and the insula, when the receiver was a person and not a computer. There was also correlation between activation and the amount offered.

The task in the ultimatum game is already getting quite complex, because the proposer needs to think through both how much they want to give and how much

the receiver is likely to accept. It's no surprise, therefore, that various areas of the brain, including the DLPFC, VMPFC and insula, are more activated. The story we get, though, seems again to be one of the prefrontal cortex acting as a check on other areas that might want to take the money. To progress further, it's interesting to look a bit more at preferences for punishment and equality.

9.3.5 Punishment and inequality aversion

A study by De Quervain and co-authors (2004) gives us interesting insights into preferences for punishment. They looked at a trust game (see Chapter 7) in which an investor can invest zero or ten tokens with a proposer; if he or she invests then the proposer has 50 tokens and must decide whether to return zero or 25 of those to the investor. The variation on this basic trust game is to allow the investor the option of punishing the proposer if the tokens were invested but not repaid.

Focusing solely on those cases in which the investor invested but got no return, Figure 9.21 illustrates the activation in the striatum. There are four different conditions depending on whether the investor chose to punish (intentional) or the computer did, whether it was costly for the investor to punish (costly) and whether the punishment actually lowered the payoff of the proposer (costly and free, versus symbolic). Recall that the striatum is involved in reward evaluation, and so it's particularly interesting that this should be activated more in those conditions when the investor could meaningfully punish the proposer. Activation also correlated with the amount of punishment. So, do people get satisfaction from punishing?

One way to try to answer this is to compare what happens when punishment is free to when it is costly. Several subjects punished the maximum amount when punishment was free, and so we can say that any differences in activation between them were probably due to differences in a desire to punish. We then ask whether

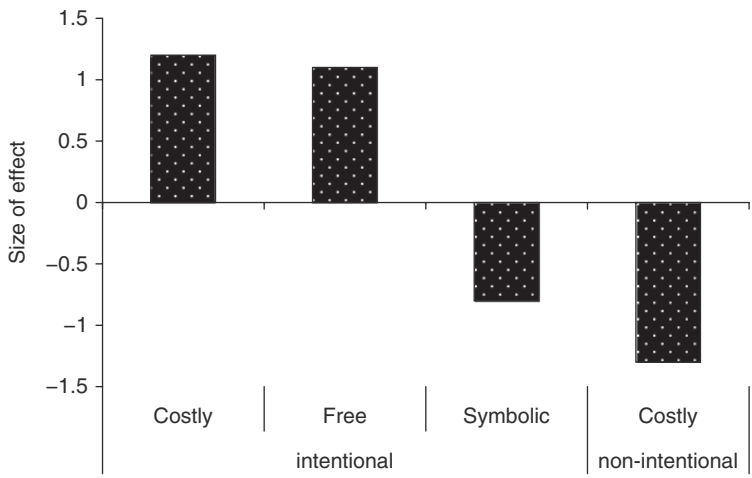

Figure 9.21 Activation in the striatum depends on the type of punishment.
Source: De Quervain *et al.* (2004).

those who had higher activation when punishment was free punished more when it was costly, and they did. This suggests that people did get satisfaction from punishing! One further piece of evidence is that the VMPFC was more activated when punishment was costly, suggesting that it was trading off the desire to punish with the cost of punishing.

If a preference for punishing seems a bit depressing, then a study by Tricomi and co-authors (2010) gives a slightly more upbeat conclusion. In the study, subjects participated in pairs. Having been given an initial $30 each, one subject was randomly given another $50 and the other nothing, creating a 'rich' and a 'poor' subject. Subsequently subjects were asked to evaluate future potential payments to them and the other subject. As we can see from the left-hand side of Figure 9.22, evaluations were consistent with inequality aversion: a rich subject valued a payment to them only slightly above a payment to the poor subject, while poor subjects rated relatively highly payments to them and negatively payments to the rich subject. So, how did this inequality aversion translate into brain activity?

The right-hand side of Figure 9.22 shows a consistency between the valuations and activation in the NAcc and VMPFC. More interesting is the NAcc, which we have seen is activated by expected reward. Rich subjects seemed to be more activated by transfers to the poor subject, while poor subjects seemed to be more activated by transfers to the rich subject. This suggests that inequality aversion really does exist. The role of the VMPFC is, again, probably to trade off the desire to equalize payoffs with the desire to get money.

It's fascinating to see that, potentially, there is pleasure in punishing and pleasure in redressing inequality. You might have noticed, however, that this brings us back full circle in trying to discern the competing incentives in the brain. If punishing and giving provide satisfaction then it's not obvious whether the 'basic emotion' is to take the

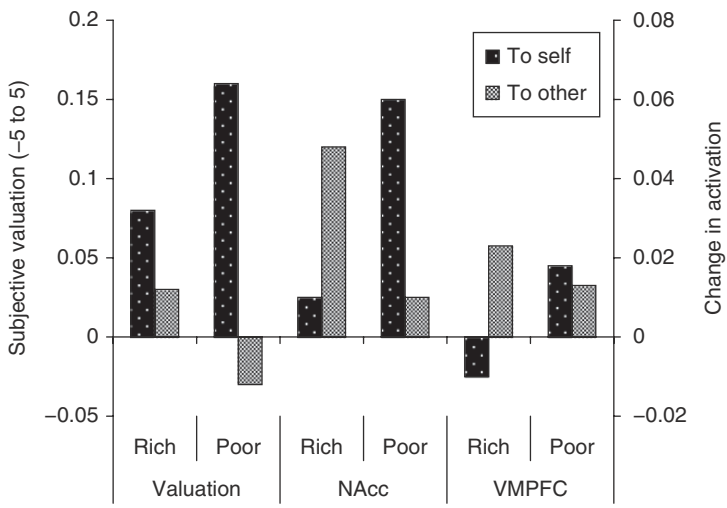

Figure 9.22 The effect of inequality on the valuation of subsequent transfers and activation in the NAcc and VMPFC.
Source: Tricomi *et al.* (2010).

money or to punish. With such competing evidence, we simply do not know yet what conflict the VMPFC and executive control systems have to sort out or inhibit when it comes to social preferences. But there's not much doubt that there are conflicting incentives, and it's this that I will now take up in a bit more detail.

9.3.6 Present bias and a brain in conflict

If parallel processes are going on in the brain all the while, it is not surprising that these sometimes **conflict** in suggesting different behavior. Indeed, if the basic purpose of the executive control systems is to switch from automatic to control, and in so doing inhibit automatic processes, there is a clear potential for conflict. So, if it feels as if one part of your brain is telling you one thing and another part something else, that may not be too far from reality.

We have already seen a potential tension between taking money offered or punishing in the ultimatum and trust game. It's not hard to find other examples. Someone falling asleep at the wheel despite trying to stay awake is a fairly extreme example. Another example you may be familiar with is the **Stroop task**, in which a person is asked to read color words very quickly and some words are in the 'wrong' color (e.g. the word 'blue' might be colored green). People take longer to reply when the word and color are incongruent because of the conflicting signals they get, the anterior cingulated cortex sorting the conflict out.

That someone could be 'of two minds' is potentially very important, particularly in thinking about present bias and time inconsistency (see Chapter 4). We could imagine, say, that Maria's automatic processes suggest she enjoys herself today, while her control processes say she should do her homework. If the automatic processes prove stronger than her control processes then she might want to do her homework, think it was best she did her homework, but still delay it to another day. We have a plausible explanation for time inconsistency.

To see if this story fits, I will look first at a study by McClure and co-authors (2004). In the study, subjects made a series of inter-temporal choices between an early or delayed monetary reward. Sometimes the reward was available immediately, but sometimes only with some delay. This made it possible to distinguish brain areas more activated by choices involving an immediate outcome. To illustrate, Figure 9.23 compares activation in the NAcc and DLPFC. We can see that activation of the DLPFC is not influenced by the delay until getting the reward, but the NAcc is. A reward today activates this area more than a delayed reward.

In Chapter 4 I looked at a quasi-hyperbolic model of (β, δ) preferences in which the δ captured exponential discounting and the β captured present bias. We can now see areas of the brain, such as the NAcc, that possibly give us the β part, while others, such as the DLPFC, give us the δ part. If so, choice should correlate with relative activation of areas associated with each part. Figure 9.24 shows that this is exactly what was seen; when a subject chose the early option, areas of the brain that we might associate with the β part were more activated, and, when they chose the delayed option, areas of the brain that we might associate with the δ part were activated.

To continue the story, I want to briefly mention a study by Diekhof and Gruber (2010). They ran experiments in which, in the first stage, subjects did a task in which they learnt that choosing certain colors gave a reward. In the second stage they were

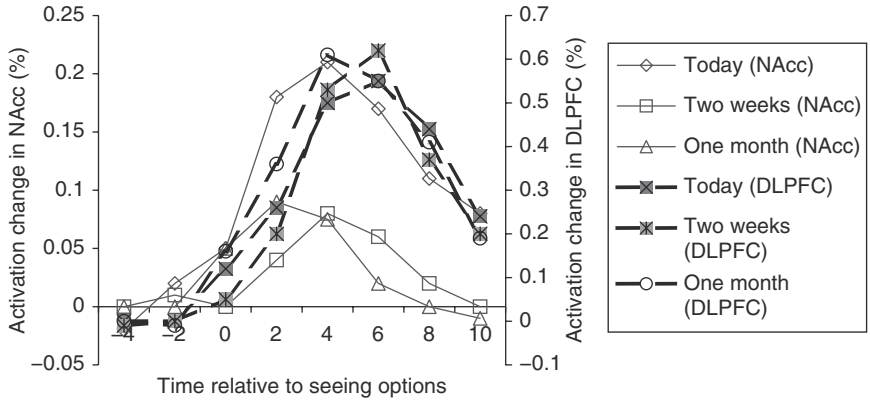

Figure 9.23 Activation in the NAcc and DLPFC depending on whether the reward is immediate or delayed.
Source: McClure *et al.* (2004).

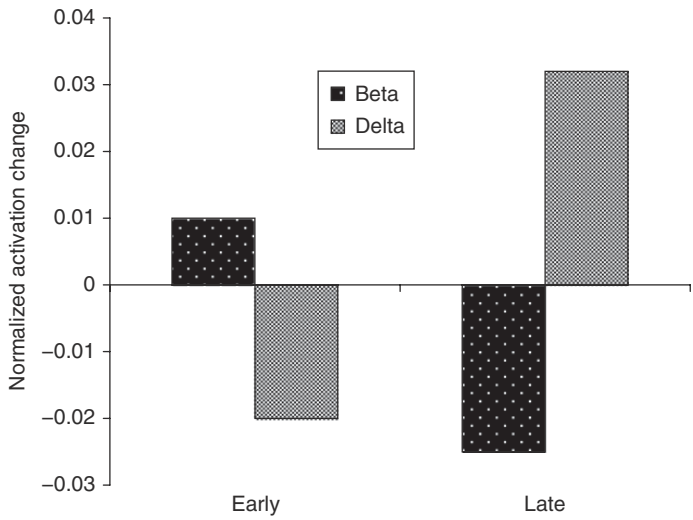

Figure 9.24 Areas associated with immediate reward (beta) are more activated when the subject chooses the early option, and areas associated with delayed reward (delta) are more activated when choosing the later option.
Source: McClure *et al.* (2004).

given a task that involved choosing certain colors a fixed number of times. This created a conflict between the long-term goal of collecting the right number of each color and the impulse to choose a color that gave an immediate reward. As we would expect, a color associated with a reward led to greater activation of the NAcc. The interesting thing came during the second stage, when a negative functional connectivity was observed between the NAcc and areas of the prefrontal cortex. In other

words, the more the prefrontal cortex was activated, the less the NAcc. Furthermore, the stronger this negative connectivity, the more likely the subject performed the task correctly.

I think that these two studies are interesting to contrast in the way that conflict resolution appears to emerge. In the first we have the impression of two areas of the brain conflicting, with the one activated to the greater extent winning. So impulsivity might win over patience, or a desire to punish might win over a desire to take the money (as in Figures 9.19 and 9.20). The second study gives the impression more of areas of the prefrontal cortex overseeing and directing what is going on elsewhere in the brain. So patience subdues impulsivity or the need for money subdues the desire to punish.

Both these ways of looking at things are really just part of the same executive control systems. On the one hand, controlled processes compete with automatic processes and, on the other hand, controlled processes direct automatic processes. One interesting implication of this, which can have important economic consequences, is that behavior can depend on the cognitive load of the executive control systems. Basically, if executive control is looking after one thing, it might let slip something else.

An interesting example of this is provided in a study by Shiv and Fedorikhin (1999). The study involved subjects having to remember either a two-digit or seven-digit number while deciding whether to choose between chocolate cake or fruit salad. Figure 9.25 shows that subjects were significantly more likely to choose the chocolate cake when the subject had to remember a seven-digit number and was shown the snacks for real. This is consistent with automatic processes being more likely to overpower the controlled processes if the controlled processes are busy doing something else.

A particularly thought-provoking example of the potential for cognitive load is provided by Mani and co-authors (2013). The focus of their study was poverty. In

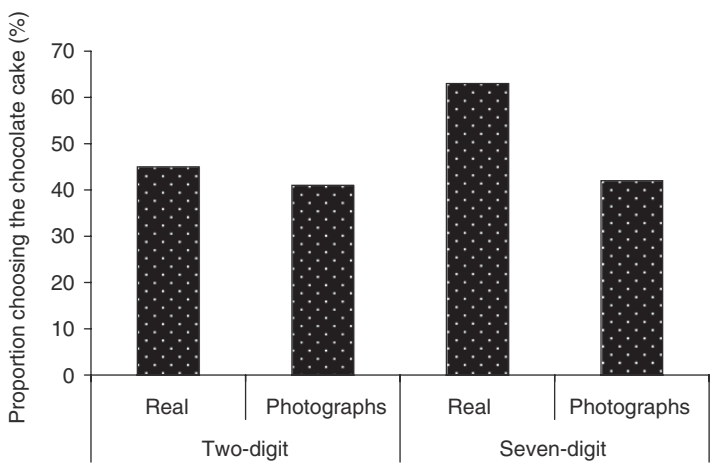

Figure 9.25 The proportion of subjects choosing chocolate cake is higher when the subject has to remember a seven-digit number and is shown the snacks for real.
Source: Shiv and Fedorikhin (1999).

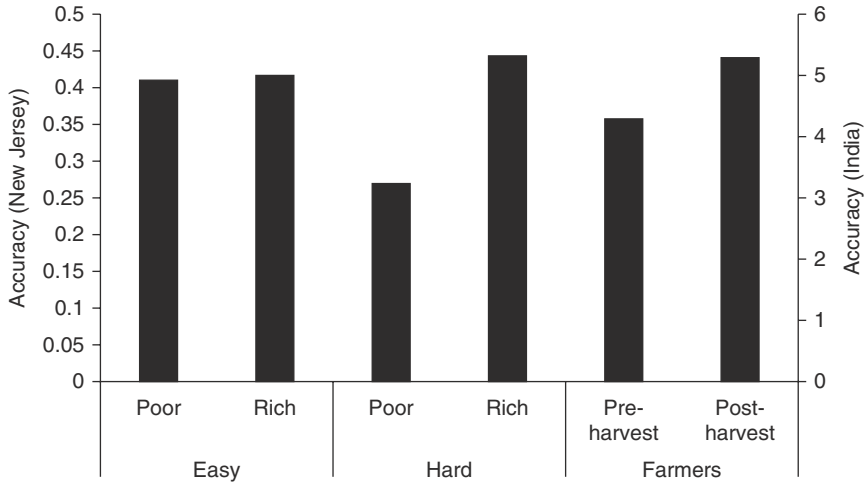

Figure 9.26 Accuracy on Raven's matrices for shoppers in New Jersey and farmers in India depending on whether they are facing a hard or easy financial scenario.
Source: Mani *et al.* (2013).

a first experiment they got shoppers in a New Jersey mall to complete tasks that require cognitive effort. The task I will focus on here is something called a Raven's matrix, in which subjects are shown a sequence of shapes and asked to pick which shape from a set of alternatives best fits in a missing space. The key twist to this experiment was that some subjects were primed, before doing the cognitive tasks, to think about hard financial decisions, such as a needing $1,500 to fix a car, while some were primed with easy decisions, such as needing only $150 to fix a car.

The left-hand side of Figure 9.26 summarizes the main results. The easy task gives us a baseline level of accuracy. In this setting, rich and poor, as determined by income, performed the same. Things are different with a hard priming. Here we see that poor people did worse than rich people and worse than the baseline level. This would suggest that worries about money generate cognitive load. It may be, however, that poor people are just different from rich people. A follow-up experiment got around this problem.

The experiment was run with sugarcane farmers in India. Harvest dates for sugarcane are deliberately staggered over a three- to five-month period. This creates a random variation in harvest dates. For instance, in July there will be some farmers who have already harvested and some who have not. Given that farmers earn most of their income from sugarcane, those who have not harvested will typically be in a harder financial situation than those who have already harvested. As we can see on the right-hand side of Figure 9.26, subjects who had not harvested at the time of the study performed less well.

This result confirms the idea that financial hardship generates cognitive load. Moreover, this cannot be put down to poor people being different from rich people. It is about people who find themselves in financial hardship. If poverty imposes cognitive load and makes people less able to switch on controlled processes in order to make

informed decisions then it may become harder for people to escape from poverty. From another perspective, policy-makers may want to design policies that reduce the cognitive burden imposed on poor people, such as through simplifying forms to apply for poverty relief.

9.3.7 Multiple-self models

A conflict being played out in the brain suggests why people have **self-control problems**. Recall that in Chapter 4 we saw that someone like Maria may want to pre-commit by, say, agreeing to meet with friends tomorrow, which means her homework has to be done today. Furthermore, she might want to pre-commit even if she does her homework today, because she still avoids the cost of overcoming temptation. This kind of thing is very hard for the standard economic model to cope with. Why would someone want to pre-commit and thereby reduce the options available to her? Now we have an explanation.

It's reassuring to have an explanation for self-control problems, but can we go further? One possibility is to explicitly model the conflict going on in the brain. This leads to the idea of **multiple-self** or **dual-self** models. I will not go into too many details here, because I will work through an example of such a model when looking at addiction, but the basic idea is simple. We can think of one self, whom Thaler and Shefrin (1981) suggest calling the **planner**, who wants to maximize lifetime utility, and another self, the **doer**, who wants to maximize utility today.

The easiest way is to think of the doer as being the one who always makes the decisions what to do now. This gives a lot of the power to automatic processes and is potentially quite worrying for the individual, because the doer will always sacrifice the future for today. Fortunately, the planner is given two main ways to constrain the doer. First, it could change the preferences or incentives of the doer. For example, Maria's planner might make it that the doer likes doing homework, or bets with a friend who will get the best homework mark. Second, it could impose rules that change the constraints of the doer. For example, Maria's planner might pre-commit to playing with friends tomorrow.

This type of model is very useful in understanding and modeling pre-commitment and self-control. It gives a big role to automatic processes, which we shall see is apt when thinking about something such as addiction. But should we be using dual-self models a lot more? Maybe it makes sense, given the real conflict between automatic and controlled processes going on in the brain.

As ever, there can be a danger in taking an idea too far. The reason I say this is that conflict is already at the heart of the standard economic model. For example, when I wrote down Anna's utility function over breakfast cereals, $u(x, TQ, HQ)$, in Chapter 2, encapsulated within that is a conflict between her wanting money and wanting the cereal of the highest quality. Anna faces a potentially tough choice in trading off money for a tasty cereal, and so she may spend some time in the grocery store being in two minds what to buy. We do not need a multiple-self model to capture that.

This is not an issue of whether there is or is not a conflict between automatic and controlled processes. Indeed, I have made a point of saying that automatic processes will influence all decisions through perception and intuition. So Anna's choice will inevitably involve both controlled and automatic processes; we just have better

ways to model this. For example, Anna might get angry at a price rise, an automatic process, and then be reluctant to buy despite knowing it's still a price she is willing to pay, a controlled process. In Chapter 2 we saw how this can be captured by distinguishing between transaction utility and acquisition utility.

Dual-self models are probably not so important, therefore, in improving how well we can model economic behavior. We have also seen that, when it comes to things such as fairness, it is very hard to disentangle the automatic 'doer' behavior and the objectives of the controlled 'planner'. Indeed, in terms of modeling economic behavior, it seems more natural to think of the dual self as a distinction between automatic processes that determine the reference point versus controlled processes that work around that reference point. That's the approach we have taken throughout most of the book.

9.4 Summary

This chapter has covered only a small selection of the many results from neuroscience and neuroeconomics, but I think that there is already enough here to be impressed with what we can learn.

We began with the idea of modularity, in which we think of different brain areas as performing different tasks. Modularity suggests that we can learn something useful from observing which areas of the brain are activated, and by how much they are activated, when a person observes an outcome or makes a particular decision.

The second idea I talked about was that of different brain processes. We contrasted automatic processes with controlled processes and gave executive control the job of knowing when to switch from automatic to control. This suggests that there can potentially be conflict in the brain when one area wants to do one thing, another area something different, and executive control has to sort this conflict out.

In looking at economic behavior, the focus to start with was on how people react to observing outcomes. We saw that some areas, such as the nucleus accumbens, are activated by the expectation of a reward and are potentially activated in proportion to the expected reward, or reward prediction error. We saw that other areas, such as the ventromedial prefrontal cortex, are more activated on getting the reward. While it may be a step too far to say that this is like seeing the utility function in the brain, it does allow us to see what types of outcome activate different reward areas, and so learn something about preferences. We saw evidence consistent with ambiguity aversion, with present bias, with the distinction between transaction and acquisition utility and with inequality aversion.

The next step was to add an element of choice to see how the brain is activated when making decisions. We saw that additional areas were activated suggesting the actor–critic distinction between areas of the brain primarily responsible for evaluation and those responsible for action. We saw that brain activation often does correlate with the action chosen and strategy used.

In looking at social preferences we saw that people appear to react negatively to an unfair offer and get a reward from punishing. Harder is to pull apart what are the automatic and controlled processes.

As I briefly mentioned in Chapter 1, neuroeconomics is a controversial topic, and so it seems natural, now that you have seen something of what it has to offer, to try

to make sense of the controversy. The basic controversy is whether neuroeconomics can add anything useful to economics, and this is primarily a manifestation of the debates that we looked at in Chapter 1. For instance, the main charge is that knowing how the brain works is neither necessary nor sufficient to understand, model and predict economic behavior.

I hope to have shown you in this chapter that knowing how the brain works can definitely help us to understand economic behavior a bit better. It's fascinating to see things such as inequality aversion and present bias seemingly being acted out within the brain. Seeing such things should reassure that the models I have looked at earlier in the book, taking account of things such as inequality aversion and present bias, make sense.

I also believe that knowing how the brain works can potentially help us model and predict economic behavior a bit better. For example, if we can distinguish automatic from controlled processes, this could help us predict how behavior will change as the context changes – to understand better, say, when people will want to punish unfair behavior and when they will not.

It is slightly worrying, though, to have to use the phrase 'potentially help us model and predict'. While neuroeconomics is undoubtedly fascinating in understanding economic behavior, it is much harder to see how it can significantly help us to better model and predict economic behavior. It's hard to believe, for instance, that watching people's brains as they play the ultimatum game can make us any more able to predict what they will do than we have learnt from the thousands of standard experiments already run. Given that economic theories are usually judged on their ability to predict behavior, this is where neuroeconomics loses out.

So, while neuroeconomics is clearly informative and worth doing, it's unlikely to revolutionize economics. It's good – but not that good.

Where neuroeconomics can be more revolutionary is in neuroscience. While economists are unlikely to learn too much from watching people's brains as they play the ultimatum game, neuroscientists can potentially learn a lot, precisely because economists know so well how people behave when playing such games. So, while we might debate just how good neuroeconomics is for economics, it is definitely good for neuroscience.

Turning around this logic, I also think economics has a lot more to learn from neuroscience. In particular, the most challenging chapter for me to write in this book was Chapter 6, on interaction and learning. Our ability to predict how people learn and behave in strategic situations is still quite poor. Given that learning, memory and social cognition are important components of neuroscience, maybe we can find some answers there.

9.5 Addiction

Addiction is a big social issue. For instance, the National Survey on Drug Use and Health suggests that, in 2015, over 10 percent of Americans aged 12 years or older had used illicit drugs during the past year. The (now closed) National Drug Intelligence Center (NDIC) estimated the economic cost of drug use at nearly $215 billion, and summarized: 'The damage caused by drug abuse and addiction is reflected in an overburdened justice system, a strained health care system, lost productivity, and

environmental destruction' (NDIC 2010). That's before we start on cigarettes and alcohol!

You might think that addiction is not a domain for economists, but it is also a big economic issue. For instance, cigarette taxes are an important source of government revenue. The US federal government raised around $17.4 billion from cigarette taxation in 2014. Remarkably, the UK government, despite a far smaller tax base, raised nearly as much, at £12 billion! A primary motivation for setting taxes high is to cut demand and lower the costs associated with addiction. The merits of such policies are hotly debated, however. Behavioral economics is ideally placed to inform the debate.

In order to get started, we need to have a definition of addiction. One definition is to say that someone like Rob is **addicted** to a good if (1) the more he consumes the good, the more he wants to consume of the good – call this **reinforcement**; and (2) the more he consumes the good, the lower his utility from consuming the good in the future – call this **tolerance**.

This is a nice definition for economists, because we can easily capture reinforcement and tolerance in simple models. It is probably too broad a definition for many, however. Indeed, according to this definition I am addicted to lots of things, including long-distance running and watching cricket. Given, however, that this is an economics book, we will go with the definition for now, possibly thinking of it more as a definition of habit than addiction. More important is to get on and start modeling addiction.

9.5.1 A model of rational addiction

To prove that this really is an economics book, I will start with a model of rational addiction (credit to Becker and Murphy 1988). It might seem strange to put 'rational' and 'addiction' in the same sentence, but this provides a natural starting point. Here is an example of how the model works.

Rob can choose to consume an addictive good, called cigarettes, or a non-addictive good, called money. Let a_t denote the amount of cigarettes he consumes in period t and let c_t denote the amount of money he spends. We then assume that his utility from smoking will depend on his previous consumption of cigarettes. To do this we need to keep track of how much he has smoked in the past. We do that by letting

$$S_t = d(S_{t-1} + a_{t-1})$$

denote the **stock of past consumption**, where d measures the addictiveness of the good. The higher is d, the more addictive cigarettes appear to be, because the more past consumption is remembered. If $d = 1$, cigarettes are very addictive, because past consumption is never forgotten.

We then assume that Rob's utility in period t will depend on present consumption and the stock of past consumption:

$$u_t = \begin{cases} c_t + 10\sqrt{a_t - S_t} & \text{if } a_t \geq S_t \\ c_t - 10(S_t - a_t) & \text{if } a_t < S_t \end{cases}.$$

We also assume that he uses exponential discounting (see Chapter 4) to find his inter-temporal utility:

$$u^T = u_1 + \delta u_2 + \delta^2 u_2 + \ldots + \delta^{T-1} u_T.$$

What Rob needs to do is plan his addiction. He knows that cigarettes are addictive, but he likes them nonetheless, and so needs to plan how many cigarettes to consume early in his life knowing that he will subsequently become addicted. The problem Rob faces is analogous to that faced by Maria, in section 4.6.1, deciding when to consume goods, knowing she will become used to that level of consumption. One slight difference is that Maria could smooth consumption freely by saving for the future. To make things more interesting we are going to assume that Rob cannot save for the future, either cigarettes or money.

Working out Rob's optimal consumption of cigarettes is a somewhat tedious exercise in math. So, I will just look at some addiction plans and see what Rob would think about them. Table 9.6 looks at three. For each one, we can see what utility Rob would get depending on what happens to the price of cigarettes in period 4. If the price of cigarettes is \$1 in all four periods then Rob does best to consume a lot of cigarettes: he prefers to consume an increasing amount, and in period 4 spends all his money on cigarettes. If the price of cigarettes increases to \$2 in period 4 then Rob does best to be a little restrained: he buys fewer cigarettes in periods 1 to 3, so that he does not become too addicted before cigarettes become expensive. If the price increases to \$2.50 then he does best to abstain altogether.

All this was assuming $d = 1$, so cigarettes are highly addictive. If $d = 0.5$, so cigarettes are less addictive, Rob still does best to consume a lot of cigarettes.

This example illustrates the main predictions that come out of a rational addiction model, namely that Rob should base his present consumption on how addictive cigarettes are, and what will likely happen to the future price of cigarettes. If cigarettes

Table 9.6 The utility of several different addiction plans, assuming Rob has \$10 per period and a pack of cigarettes costs \$1 in periods 1 to 3. In period 4 the amount of cigarettes that Rob can consume is constrained by his income. Unless stated, $d = 1$.

Price in period 4	Cigarettes in period				Utility
	1	2	3	4	
\$1	0	0	0	0	40.0
	1	2	4	10	70.3
	1	1	2	5	51.0
\$2	0	0	0	0	40.0
	1	2	4	5	33.0
	1	1	2	5	46.0
\$2.50	0	0	0	0	40.0
	1	2	4	4	23.0
	1	1	2	4	36
\$2.50 and d = 0.5	0	0	0	0	40.0
	1	2	4	4	73.6
	1	1	2	4	70.5

are addictive and the future price will increase then he should not consume now, so as to not become trapped wanting something he cannot afford. If cigarettes are less addictive and/or he does not expect any increase in the price then he can consume more. We can call this **forward-looking behavior**.

Before we question the plausibility of forward-looking behavior, I want to look at some data. The data is from a study by Gruber and Kőszegi (2001) that looks at whether demand for cigarettes changes when taxes are announced. Given that there is often a delay of two months or more between a government announcing that taxes will increase and taxes actually increasing, it is possible to check for forward-looking behavior. We would expect consumption of cigarettes to fall when the tax increase is announced, not when prices actually rise. We might, however, expect demand for cigarettes to increase when the tax increase is announced, because it is a chance to hoard them before prices rise.

Gruber and Kőszegi looked at some data on both sales and consumption of cigarettes, and find the predicted effects. In the sales data there is evidence of a hoarding effect, with sales increasing after the announcement of a tax increase. For example, an increase in the tax from 10 to 35 cents in California was announced in November 1988 and become effective in January 1989. Cigarette sales jumped from an average of 6.68 packs per person in November to 8.71 packs in December, before falling back to around six packs. In the consumption data there is evidence of decreased consumption once the tax increase is announced. It seems, therefore, that consumers are using forward-looking behavior.

This suggests that Rob could be in control of his addiction, and the only rationale for policy intervention would be if Rob's smoking has a negative externality on others. Given the biases we have seen elsewhere in this book, this might seem a surprising conclusion. It is not the end of the story, however. That's because we observe a lot of other behavior that is not so consistent with a rational model of addiction. To that I now turn.

9.5.2 Biases and addiction

The rational model of addiction suggests two important biases that we might observe. First, the rational model requires forward-looking behavior, and we have seen (in Chapter 4) that people such as Rob may not be able to stick with a forward-looking plan because of a present bias. He might always plan to 'quit tomorrow' but never do so. Second, the rational model says that Rob will smoke more if he thinks cigarettes are less addictive. Chapter 5 provided lots of reasons why he may be biased in underestimating its addictiveness. For example, confirmatory bias and the law of small numbers suggest that he may ignore the scientific evidence on how addictive smoking is because, say, his friend was able to quit smoking.

Both these biases have the same basic prediction: Rob will become more addicted and consume more cigarettes than he planned to do. Evidence of this is not hard to find. To demonstrate, we can look at some data from an Office for National Statistics (ONS) report published in the United Kingdom about smoking-related attitudes and behavior (ONS 2004). Figure 9.27 summarizes how many people surveyed wanted to quit smoking. A lot of them did. Given this, we might expect that they would try to quit. Indeed, well over 60 percent had tried and failed. Figure 9.28 summarizes the

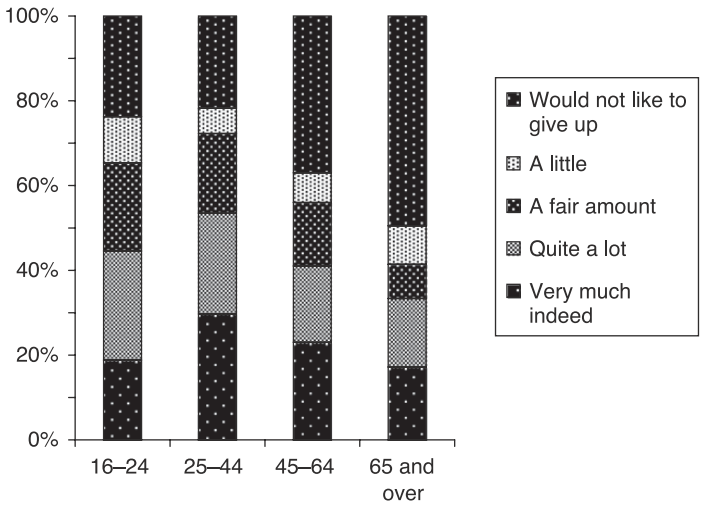

Figure 9.27 The proportion of people surveyed, by age group, who want to quit smoking.
Source: ONS (2004).

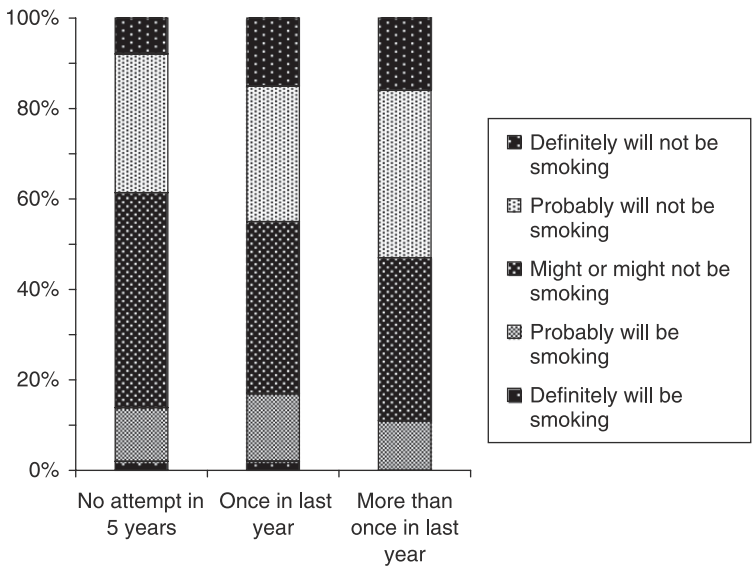

Figure 9.28 What smokers who are planning to quit in the next 12 months expect will happen after 12 months.
Source: ONS (2004).

Table 9.7 The consequences of underestimating the addictiveness of cigarettes, and present bias in the example.

Bias	Cigarettes in period				Expected utility
	1	2	3	4	
None	0	0	0	0	40.0
	1	2	4	4	23.0
Thinks d = 0.5	0	0	0	0	40.0
	1	2	4	4	73.6
	1	3	4	4	73.4
	—	2	4	4	51.0
	—	3	4	4	51.2
Present bias β = 0.8	0	0	0	0	34.0
	1	2	4	4	22.2
Thinks d = 0.5 and has present bias β = 0.8	0	0	0	0	34.0
	1	2	4	4	62.6
	1	4	4	4	61.6
	—	2	4	4	44.4
	—	4	4	4	44.8

thoughts of those who planned to quit in the next 12 months. There was some realism about the difficulties of quitting, but, overall, people were more optimistic than they probably should have been about their chances of quitting.

There is lots more data I could have given, and all of it is hard to reconcile with a rational model of addiction. It looks, therefore, as though people are not forward-planning very well and are underestimating the difficulties of quitting smoking. More complex is to say exactly why this is. Is it because of present bias, or an underestimation of the addictiveness of cigarettes? To try to answer this question, let's look in a bit more detail at how each bias would work in the earlier example.

Suppose that the price of cigarettes is expected to be \$1 in periods 1 to 3 and to increase to \$2.50 in period 4. Also, suppose that $d = 1$, and so cigarettes are very addictive. Table 9.7 illustrates the basic ideas by contrasting four different scenarios. The first scenario considered is that of Rob having no bias. Of the consumption plans I am going to look at, Rob would do best to not smoke (as in Table 9.6). What if he is biased?

In the second scenario Rob thinks cigarettes are less addictive than they are. Specifically, he thinks that $d = 0.5$. In period 1 we can see that he would prefer a plan with only two packs of cigarettes in period 2. When he gets to period 2, however, his preferences change. In particular, because he is more addicted than he thought he would be, he now prefers having three packs of cigarettes rather than two. Rob will consume more cigarettes than he planned, and more than would maximize his utility.

In the third scenario Rob has present bias, with β =0.8. In this scenario not a lot happens. We can see from the relative utilities that consuming cigarettes has become slightly more attractive, but still not enough to warrant smoking. You can play around with the numbers yourself, but it turns out that present bias is not going to make much difference to what Rob will do unless he has a really strong present bias. I'll come back to why this is the case later.

It looks as if underestimating the addictiveness of smoking is the real problem Rob needs to be wary of. Present bias can make a difference, however, if Rob underestimates the addictiveness of smoking. To appreciate this point, consider the final scenario, in which Rob has present bias and thinks $d = 0.5$. When he gets to period 2 he would now prefer to smoke four packs of cigarettes rather than the two he originally planned. Rob will thus consume even more cigarettes and get an even lower utility. The less addictive Rob thinks smoking is, the more present bias is going to matter.

A combination of underestimating the addictiveness of smoking and present bias can, therefore, easily lead to Rob smoking more cigarettes than he planned and getting more addicted than he planned. It is no surprise, therefore, that he might end up wanting to quit but be unable to do so. The data we looked at in Figures 9.27 and 9.28 becomes a lot easier to explain. The role for government policy also changes, because there becomes a clear rationale for more information on the addictiveness of smoking, and it also makes sense to give people help in trying to quit. But the story we have so far is primarily one of why Rob might become addicted. What we need to do now is look more at the potential difficulties of quitting.

9.5.3 Cues and addiction

One important factor in addiction can be **cue-elicited behavior**. The idea here is that some cue, such as a particular sound, smell or situation, makes someone crave something. For example, if Rob typically smokes a cigarette after dinner, or when he goes to the bar with friends, then his craving for a cigarette might be higher than normal at these times. We have seen how secondary reinforcers can have this effect, and such cue-elicited behavior is often cited as a reason for relapse in people trying to quit an addiction. For example, Laibson (2001) gives the story of a smoker who hadn't smoked for years, but on going to the beach one day had an intense craving to smoke. The craving was a consequence of smoking on the beach being common at one time in his life.

We can model cue-elicited behavior without too much change to the rational addiction model. All we need to add are the cues. Let's suppose that each day Rob either eats dinner at home or eats out. This will be the cue. At dinner he then decides to smoke or not. Instead of keeping stock of how many cigarettes Rob smokes, we now distinguish how many cigarettes he smokes during dinners at home and how many at dinners out.

To explain how we can do this, imagine that on day $t-1$ Rob decided to eat out. Then we let

$$S_t^{out} = d(S_{t-1}^{out} + a_{t-1})$$

denote the stock of past consumption of cigarettes when eating out, and keep $S_t^{in} = S_{t-1}^{in}$ as the stock of past consumption of cigarettes when eating in. Similarly, if he had decided to eat in we would let

$$S_t^{in} = d(S_{t-1}^{in} + a_{t-1})$$

denote the stock of past consumption of cigarettes when eating in, and keep $S_t^{out} = S_{t-1}^{out}$ as the stock of past consumption of cigarettes when eating out.

Rob's utility in period t then depends on where he eats and the difference between present consumption and the stock of past consumption. For instance, if he eats out his utility is

$$u_t = \begin{cases} c_t + 10\sqrt{a_t - S_t^{out}} & \text{if } a_t \geq S_t^{out} \\ c_t - 10(S_t^{out} - a_t) & \text{if } a_t < S_t^{out}. \end{cases}$$

If he eats in his utility is

$$u_t = \begin{cases} c_t + 10\sqrt{a_t - S_t^{in}} & \text{if } a_t \geq S_t^{in} \\ c_t - 10(S_t^{in} - a_t) & \text{if } a_t < S_t^{in}. \end{cases}$$

We still assume that he uses exponential discounting to find his inter-temporal utility:

$$u^T = u_1 + \delta u_2 + \delta^2 u_2 + \ldots + \delta^{T-1} u_t.$$

To better understand the model, let's start by assuming that Rob eats out every day. Then the model is exactly the same as the rational model of addiction. But now we need to think more about discounting. Until now I have assumed the discount factor was $\delta = 1$, meaning there was no discounting of the future. Table 9.8 looks at what happens if the discount factor is $\delta = 0.7$ rather than $\delta = 1$. We can see that the decrease in the discount factor makes smoking look somewhat more attractive. The reason for this is that the costs of smoking come primarily in period 4, when Rob is very addicted. The lower the discount factor, the less Rob worries about these costs.

So, the discount factor can matter a lot in Rob's choices. Incidentally, this does not contradict my saying earlier that present bias may not have such a large effect. That's because the benefits of smoking come not just today but over the initial few periods, and so present bias discounts most of the benefits as well as the costs of smoking. The costs of smoking do come after some time, however, and so something that discounts the future the further away it is will make a difference.

Why should we care about the discount factor? To see why the discount factor is important in understanding cues, imagine that Rob eats in six days a week and eats out on a Saturday. As Table 9.9 illustrates, when deciding whether to

Table 9.8 The effect of the discount factor on Rob's desire to smoke during dinner, assuming that $d = 1$ and a pack of cigarettes costs $2 in all periods.

Cigarettes on day				Utility	
1	2	3	4	$\delta = 1$	$\delta = 0.7$
0	0	0	0	40.0	25.33
1	2	3	5	28	27.73

Table 9.9 The discount factor matters if there are two cues.

Day of week	Cue	If smokes, when it will first matter
Monday	Eat in	Tuesday
Tuesday	Eat in	Wednesday
Wednesday	Eat in	Thursday
Thursday	Eat in	Friday
Friday	Eat in	Sunday
Saturday	Eat out	Next Saturday
Sunday	Eat in	Monday

smoke on a Monday he knows that whether or not he smokes will affect his stock of past consumption when eating in, and so will affect his utility on the following day. By contrast, on a Saturday he knows that whether or not he smokes will affect his stock of past consumption when eating out, and so will not affect his utility until one week. This means that the effective discount factor for eating out is a lot lower than that for eating in. For example, if the discount factor per day is $\delta = 0.95$ then the effective discount factor for eating in is 0.95, but for eating out it is $0.95^7 = 0.7$.

We have already seen that a low discount factor makes it optimal for Rob to smoke more cigarettes. A probable outcome, therefore, is that he would not smoke when eating in but would smoke when eating out. This is why cues matter, and why Rob may decide to smoke only when he gets certain cues. If we add present bias or an underestimation of the addictiveness of cigarettes to the story, we can see why addictiveness can become cue-dependent and why Rob may become more addicted to smoking when eating out than he wanted to.

Cue dependence is interesting in thinking about how people can quit addiction. If addiction is cue-dependent then people can avoid addiction by avoiding the cues that cause it. For example, if Rob does become too addicted to smoking when eating out then he can simply stop eating out and solve the problem. This looks like good news for those wanting to quit, because it provides a viable way to avoid addiction. It is not all good news, however, because Rob would have to sacrifice eating out, which he might like, and may remain vulnerable to reigniting his addiction if he ever does eat out. This last point may suggest why many people struggle to quit.

9.5.4 Addiction and neuroscience

Economists love simple models, and also love drawing analogies, but it's sometimes useful to check whether it all makes sense. This seems particularly pertinent in looking at addiction, for which economists are a long way from familiar ground. So, are the models we have just been looking at reasonable models of addiction, and how far can we generalize them in looking at habit and more mundane consumer purchases? Neuroscience can potentially guide us to an answer.

Fortunately, we do not need to delve too deeply into the complex neuroscience of addiction in order to learn some interesting things. In terms of hard drugs, we know that, on initial use, such drugs increase dopamine transmission, which you might

recall is usually associated with a positive reward. Users are, therefore, likely to associate the drug with pleasure, and begin to associate cues that might have preceded using the drug with that pleasure. This is the same story as we looked at above in valuing rewards, and so there is little at this stage to distinguish drugs from anything else – e.g. money.

If the person continues to use the drug, however, then they go through various stages in which the neural circuitry appears to essentially be rewired by the drug use. This leads to addiction. In the final stage of addiction the VMPFC seems to lead the way in demanding the drug, rather than passively responding to consumption of the drug. Potentially, this can become permanent, making it very hard for an addict to ever quit. It also means that addicts will be less responsive to 'normal' rewards, and less able to make decisions.

What do we learn from this? It lends strong support to a cue-based model of addiction, because it is cues that activate the VMPFC to crave the drug. It also lends support to the potential importance of underestimating the addictiveness of drugs, because drugs really are addictive and people may underestimate their effect. Present bias starts to appear less relevant, though, because, if a person really did know the consequences of drug addiction, any present bias would have to be very extreme for them to still use the drug. It's probably not present bias, therefore, that's at work, but a potential uncontrollable craving for the drug that is unrelated to present bias.

Knowing this gives us some idea of what a model of addiction should look like (and actually points towards a model of addiction proposed by Bernheim and Rangel (2004) that I'm not going to look at here, but I suggest you do). It also gives us a better idea how to design policy to deal with the problems of addiction. For instance, it seems clear that conventional economic incentives, such as increasing the price of a drug, are liable to be relatively ineffective. Thus taxes, although they raise a lot of government revenue, are not going to solve the problems of addiction. More apt is probably to increase awareness of the problems of addiction and to help people avoid cues that trigger their addiction.

We can also see that addiction and habit are quite different things. Thus, an analogy gets us only so far, and we probably will need different models to capture drug addiction versus a habit to watch cricket. Drawing the dividing line is not necessarily easy, because drugs come in varying strengths, from heroin to caffeine to chocolate, but a distinction does need to be made. We thus learn how to better model addiction, and also can be reassured that things such as reference dependence and present bias are probably enough to capture habit.

9.6 Further reading

Textbooks on neuroscience and cognitive neuroscience, such as Purves and co-authors (2008) and Glimcher and Fehr (2013), are a natural place to start. The book by Glimcher (2004) gives a more general introduction to the subject. Overlapping survey articles include Camerer (2007) and Camerer, Loewenstein and Prelec (2004, 2005). Fudenberg and Levine (2006) provide a recent dual-self model of self-control. Schelling (1984) and Bernheim and Rangel (2007) provide interesting contrasts on the implications of multiple selves for measuring welfare. On addiction, the paper by

Bernheim and Rangel (2004) is recommended, and Kalivas and Volkow (2005) provide more on the neuroscience of addiction. For more on the debate over whether neuroeconomics is a good or a bad thing for economics, see Caplin and Schotter (2008).

9.7 Online materials

On the companion website there is more details around the example of Rob's addiction. You will also have access to flashcards, quizzes and answers to the review questions.

9.8 Review questions

9.1 Why does the brain interpret money as a reward? Is it a primary or a secondary reinforcer?

9.2 Do you think the brain's reward systems react to expected value or 'expected utility'?

9.3 When does Rob feel the value of winning at roulette: when the wheel stops, when he collects his money or when he eats the burger he buys with the money? Relate this to hedonic editing.

9.4 Why is reward prediction error relevant in interpreting Figures 9.4 to 9.7?

9.5 Relate reward prediction error to reference dependence.

9.6 Given that brain activation is constrained to fall within a fixed range, how can activation be in proportion to the reward? What does this suggest about our differing ability to distinguish relative from absolute magnitudes?

9.7 Discuss ways in which we could try and discern whether the automatic process is to be selfish or 'kind'.

9.8 How might policy-makers take into account the finding that poverty impedes cognitive function when designing policy?

9.9 Why might it be consistent with forward-looking behavior that people want to quit smoking? What if I tell you that only 26 percent of people said they wanted to quit smoking because of financial reasons (ONS 2004), but 86 percent said they wanted to quit because of health reasons?

9.10 What policy should the government adopt to stop people doing drugs or smoking?

9.11 Why is addiction different from habit?

9.12 Is a multiple-self model helpful in modeling economic behavior?

9.13 What can we learn from neuroeconomics about the effect framing has on choice?

9.14 Consider and categorize the different brain processes that might happen for someone who is the proposer in an ultimatum game.

9.15 Does neuroeconomics go against the methodology of positive economics (see section 1.1.4)?



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Part IV

Welfare and policy



Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

Chapter 10

Happiness and utility

> The great source of both the misery and disorders of human life seems to arise from overrating the difference between one permanent situation and another. Avarice overrates the difference between poverty and riches: ambition, that between a private and a public station: vain-glory, that between obscurity and extensive reputation.
>
> Adam Smith, *The Theory of Moral Sentiments*

Having looked at how people behave, and why, we shall now start to look at the welfare and policy implications of that behavior. The basic question I want us to look at in this chapter is whether people make the economic choices that will give them greatest satisfaction and happiness, and, if not, whether they would rather someone else made choices for them. To answer that question we need to know what makes people happy, whether people know what makes them happy and how much people value choice. In the next chapter we shall make the logical next step in asking what role policy can play.

10.1 What makes us happy?

The story throughout the book has been one of people trying to maximize utility. For the most part, we can think of utility as a measure of happiness, so, if people maximize utility, they maximize happiness. The standard economist's answer to the question 'What makes us happy?' is, therefore, simple: 'Whatever we do!' If we observe Emma going to the cinema then we could infer that going to the cinema makes Emma happy; why else would she go? If we observe Alan buying car insurance then we infer that this makes him happy; why else would he buy it? This is called **revealed preference**, because Emma's and Alan's actions reveal their preferences.

Can we be sure, however, that a person's choices truly reveal what makes them happy? To answer that question we need to have an alternative measure of what makes people happy. One alternative is to ask people how happy they are, and see how their answer changes as things happen in their life. For example, if we observe Emma getting a job, a revealed preference argument would suggest that having a job makes her happier. If we know how happy Emma said she was before she got the job, and how happy she says she is after getting the job, then we can check whether she really is happier.

Fortunately, there is a lot of data on how happy people say they are. For example, a question in the World Values Survey asks respondents: 'All things considered, how satisfied are you with your life as a whole these days?' Many other surveys ask similar questions. There are pertinent issues about how much we can rely on this data (see Research methods 10.1 and 10.2), but the data does throw up some consistent results, and so we'll assume that it is reliable. This means we can check whether choices really do reveal what makes people happy. Before we do that, though, it is useful to have a look first at what people say makes them happy.

❖ RESEARCH METHODS 10.1 THE LAW OF LARGE NUMBERS: CAN WE TRUST HAPPINESS DATA?

Two ways we can measure happiness are to ask people questions such as 'How satisfied are you with your life as a whole these days?' or to ask them to look back at an event and try to recall the pleasure, pain or utility they felt during the event. Do people give a reliable answer to such questions?

We know that people's answers to such questions can be affected by current mood and how the question is framed. This was amply demonstrated in an experiment by Norbert Schwarz, where half those surveyed were 'lucky enough to find a dime' before being surveyed. Those who found the dime reported higher 'life satisfaction'. We also know that the same people can give quite different answers to how happy they are if they are asked two or more weeks apart. Clearly, therefore, we need to be a little skeptical of happiness measures.

So, should we be using happiness data at all? It will probably suffice for our purposes. That is because we will aggregate the data of many people. For example, we will be interested whether unemployed people are less happy than employed people. If we ask lots of unemployed and employed people whether they are happy and average out the responses, we should get a reliable measure of whether those who are employed are happier. If the data is a little unreliable or noisy then we may have to ask more people, but we can still get reliable information.

❖ RESEARCH METHODS 10.2 AN ALTERNATIVE MEASURE OF HAPPINESS

The **day reconstruction method** involves participants filling out a diary summarizing episodes that occurred the previous day. They should describe what they did, how long they did it for, etc., and also how they felt during each event of the day. A study by Kahneman and co-authors (2004) used the

Table 10.1 The day reconstruction method. Net affect is a measure of positive minus negative emotions. All responses were on a six-point scale, so 6 is the most positive and 0 the most negative.

Activity	Time spent (hours)	Net affect	U-index
Intimate relations	0.23	4.83	0.040
Socializing after work	1.14	4.15	0.073
Relaxing	2.17	3.96	0.078
Dinner	0.81	3.94	0.074
Lunch	0.57	3.91	0.078
Exercising	0.22	3.85	0.088
Praying/worship	0.45	3.78	0.105
Socializing at work	1.12	3.78	0.100
Watching TV	2.19	3.65	0.095
Phone at home	0.93	3.52	0.126
Napping	0.89	3.35	0.131
Cooking	1.15	3.27	0.138
Shopping	0.41	3.23	0.157
Computer (non-work)	0.51	3.22	0.165
Housework	1.12	2.99	0.161
Childcare	1.10	2.99	0.199
Evening commute	0.61	2.77	0.209
Working	6.89	2.68	0.211
Morning commute	0.47	2.09	0.287

Source: Kahneman *et al.* (2004).

method on a sample of 909 working women in Texas. Some of the data is given in Table 10.1. Net affect measures positive minus negative experience, and there are not too many surprises in what gives relatively high and low net affect.

One way to use the day reconstruction method to measure happiness is the U-index. The **U-index** asks how much of the time a person spends doing unpleasant activities. There is no simple definition of when an activity is unpleasant, but one possibility is to count any event for which the most intense feelings are negative. In Table 10.1 there is the U-index for the various daytime activities. So, for example, 28 percent of the time spent commuting in the morning was unpleasant.

Figure 10.1 compares the U-index with the standard measure of happiness and household income and age, and we see a strong correlation between the two measures. The U-index does have one big advantage over the standard measure, however: interpersonal comparisons are more plausible. Interpersonal comparisons are hard to believe when focusing on satisfaction questions, because how people choose to rate themselves on, say, a ten-point satisfaction scale can vary widely from person to person without any difference in their happiness. So, that Maria rates herself seven and Alan six does

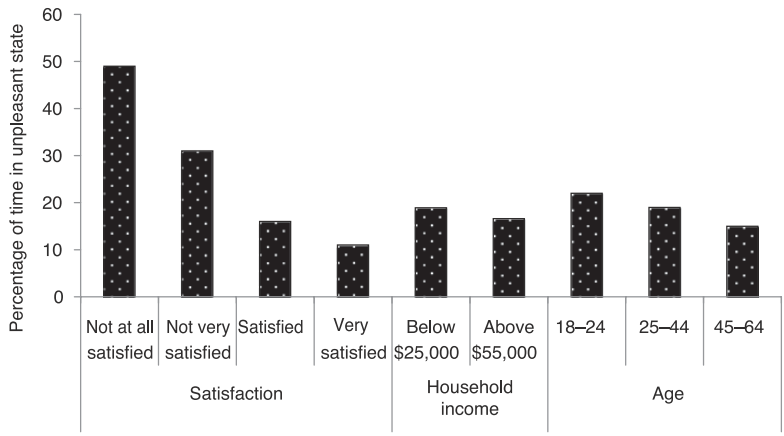

Figure 10.1 The U-index compared to life satisfaction. The lower was life satisfaction, the less time people spent in an unpleasant state.
Source: Kahneman and Krueger (2006).

not convincingly tell you that Maria is happier than Alan. The U-index is intuitively a more objective measure. People may more reliably remember whether or not an event is unpleasant, and the time spent doing unpleasant events is a readily understandable measure.

10.1.1 Happiness is relative

When we look at how happy people say they are, or how satisfied they say they are with their life, we get some surprises. Two studies in the 1970s, one in economics and one in psychology, first highlighted the issues, and it is interesting to look at both.

Brickman, Coates and Janoff-Bulman (1978) compared the happiness of lottery winners, accident victims and a control group of people randomly selected from the phone book. The lottery winners had won between $50,000 and $1 million on the Illinois State Lottery. The accident victims were either paraplegic or quadriplegic as a result of their accidents.

The main results of the study are given in Figure 10.2. The first thing we see is very little difference in reported happiness between the lottery winners and control group. Winning the lottery does not seem to make people any happier. Being in an accident does have a more pronounced effect and, not surprisingly, lowers present happiness and increases the perception of past happiness. What is surprising is that the drop in happiness is not very large, and predictions of future happiness are no different from those of the lottery winners. Winning $1 million or becoming paraplegic is clearly a life-changing event, but this study suggests neither change life satisfaction as much as we might expect.

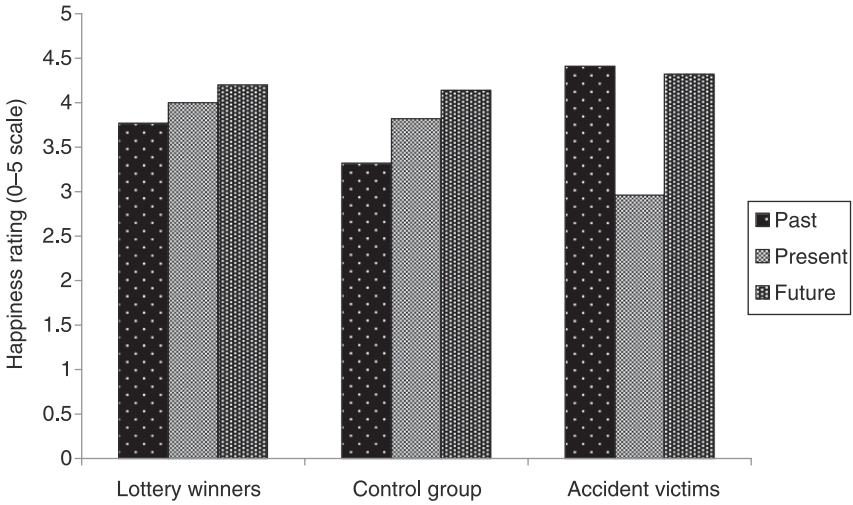

Figure 10.2 The change in reported life satisfaction following a life-changing event. *Source:* Brickman, Coates and Janoff-Bulman (1978).

Easterlin (1974) looked at whether happiness correlates with income. Figure 10.3 reports data from a 1970 survey of the US population. As we might expect, those with higher income reported being happier. No surprises so far. If, however, increased income leads to increased happiness then it seems logical that (1) people should be happier, on average, in countries with higher income, and (2) people should become happier, on average, if incomes increase over time. The curious thing that Easterlin pointed out is that neither of these things happens!

Whether or not growth in average income leads to greater happiness remains a controversial topic. The findings of Easterlin have, however, largely survived the test of time. I say 'largely' because modern data does suggest that people are, on average, happier in countries with higher income. Figure 10.4, for example, provides a recent cross-country comparison of reported happiness relative to per capita GDP. There is a positive correlation. The correlation is not particularly strong, though, and lower GDP tends to correlate with war, despotic governments and lots of other reasons why people might be unhappy. More convincing is to look at changes in income within country. Figure 10.5 shows reported happiness in the United States over time, and, despite real income tripling over this period, reported happiness changed little. The story is the same in other countries. Increases in average income do not seem to cause increases in happiness.

Since the two pioneering studies discussed above there have been many more, some of which we shall look at below, and most confirm these early findings. Happiness, and unhappiness, it seems, can be a little elusive. Before we start to look at the important implications of this, we need to ask how such findings can be explained. The leading explanation is reference dependence. If happiness depends on the relative rather than the absolute level of things such as income and health then changes in the absolute level need not change happiness. The two most likely reference points here are what others have, and what the person used to have. I will primarily focus

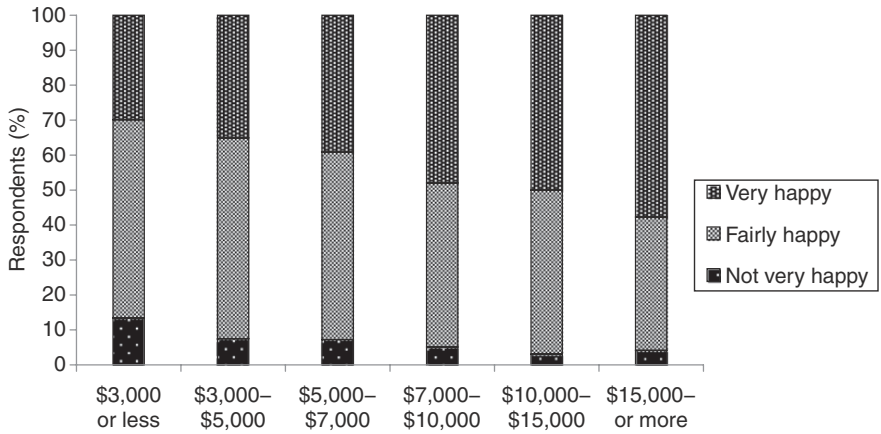

Figure 10.3 Life satisfaction and income. Those with higher income report higher life satisfaction.
Source: Easterlin (1974).

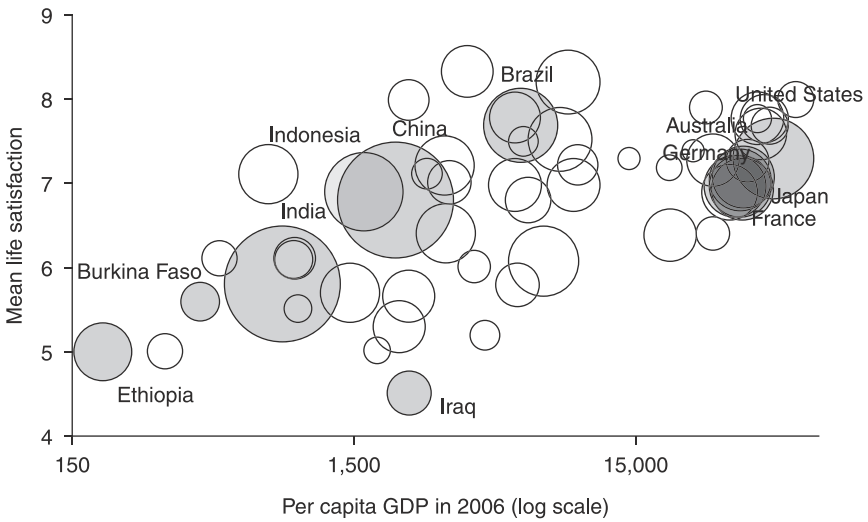

Figure 10.4 Life satisfaction and GDP per capita in different countries in 2006. There is a positive correlation between the average income in a country and satisfaction. Size of bubble reflects population size.
Sources: World Values Survey 1981–2008 database, v.20090901 (www.worldvaluessurvey.org), and World Bank.

on the latter of these two in the rest of the chapter, so let's briefly think about the implications of the first.

If people want to earn as much income as relatives and friends (and we saw evidence in Chapters 2 and 7 that they do) then those earning relatively more will be happier than those earning relatively less. This is consistent with Figure 10.3. Logically,

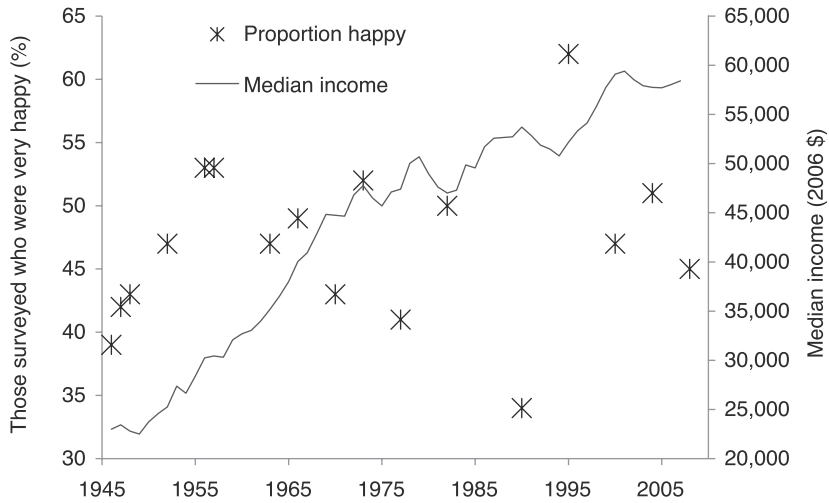

Figure 10.5 Life satisfaction and income over time. Median income increases a lot but reported satisfaction does not.
Sources: Easterlin (1974), Gallup and US Census Bureau.

however, there must always be some in a society that are earning relatively more, and some earning relatively less. This can explain why we see no correlation between average happiness and income over time. That people judge their own income relative to others can, therefore, explain a lot. But it cannot explain everything, as we shall now see.

10.1.2 Adaption and habituation

Looking back at Figure 10.2 we can see that happiness does not depend solely on circumstances relative to others. The lottery winners should be relatively wealthy, but are not happy. The accident victims should have relatively ill health, and were less happy when asked, but did not expect to be less happy in the future. To explain this, we need to introduce the ideas of adaption and habituation.

To illustrate and motivate these concepts, we can look at some data on how people reported happiness or satisfaction changes in response to other 'life-changing' events. Clark and co-authors (2003, 2008) analyzed data from Germany that tracks people's life satisfaction over the period 1984 to 1997. They used the first five or seven years of the panel to obtain a person's baseline level of satisfaction and then looked at how satisfaction changes in the years leading up to and after events such as marriage, divorce, a first child and unemployment. Most of these events are predictable, and so it is also useful to have some data on an event that is more unpredictable. For that, we can look at data by Oswald and Powdthavee (2008) from the United Kingdom on the consequences of disability.

I am not going to look at all of the data, but Figure 10.6 plots a representative sample. In terms of marriage, the birth of the first child and layoff, we appear to see **anticipation**, in that happiness increases or decreases prior to the event.

Disability, as we might expect, causes a more sudden jump in unhappiness, because it is less predictable. After the event, in the case of marriage and the birth of a first child, we see **adaptation**, in that happiness returns to the baseline level. There is some evidence of adaptation in the case of disability, but not for being unemployed.

Adaptation is a fascinating possibility, because it implies that happiness will bounce back to baseline after a good or a bad life event. This has profound implications for policy, and how we look at life. So why does it happen? There are lots of interrelated possible causes, all with different interpretations, of which I will mention four.

We can distinguish two effects that result from a real change in utility. In terms of negative events, there may be **readjustment**, whereby the person finds new ways to fulfill their life. For example, someone who liked to play football but becomes paraplegic might learn to play wheelchair basketball. In this case, the utility function or reference level has not changed, but the person has had to relearn how to maximize utility. **Habituation** is when a person becomes accustomed to their new life. For example, the lottery winner becomes accustomed to buying expensive things, and the accident victim becomes accustomed to being in a wheelchair. This would imply a change in the reference level and/or a change in the utility function (similar to that which we looked at in section 4.6.1).

We can next distinguish effects that focus more on how satisfied people feel they are. One is a **contrast effect**, whereby an extreme good or bad event changes the scale by which all future events are judged. For example, nothing might seem 'as good' as winning the lottery or 'as bad' as being in an accident. Taking this one stage further, we get an **aspiration treadmill**, when a person adjusts aspirations to the

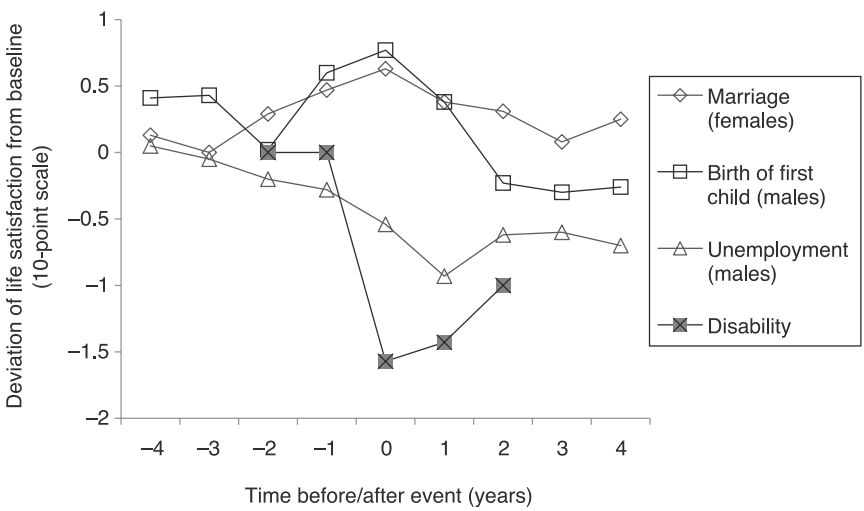

Figure 10.6 Changes in happiness relative to the baseline in, before and after life-changing events.
Sources: Clark *et al.* (2003) and Oswald and Powdthavee (2008).

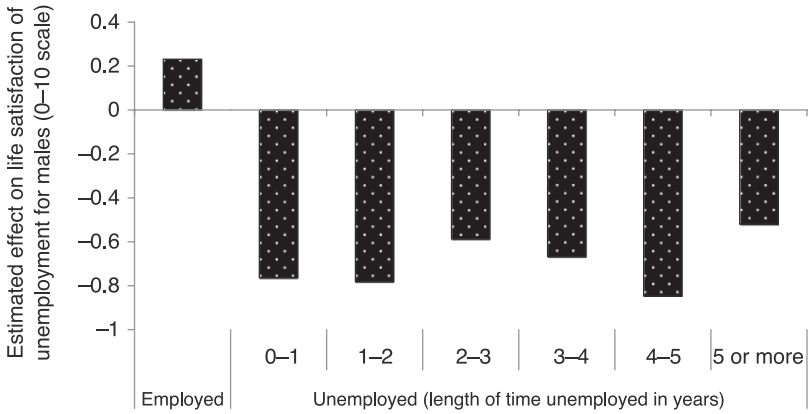

Figure 10.7 Estimated effect on life satisfaction of being employed versus unemployed. There is no evidence of adaptation to being unemployed.
Source: Clark *et al.* (2008).

satisfaction they normally experience. Consequently, the lottery winner is likely to downplay current satisfaction and the accident victim to overplay current satisfaction, given what they experienced previously.

Readjustment and habituation suggest that utility regains its former level. The contrast effect and aspiration treadmill suggest that utility may not regain its former level, but people's answer to the question as to how satisfied they feel does. This is a crucial distinction. Unfortunately, we do not yet have good answers to what best explains adaptation. It does not necessarily follow, therefore, that utility bounces back to the baseline.

What we do know is that adaptation seems to happen only some of the time. For some types of event, adaptation can be incomplete and/or take a very long while. For instance, in Figure 10.6 we see little evidence of adaptation after being unemployed. To illustrate this further, Figure 10.7 plots the estimated effect of unemployment on life satisfaction for males. We see that being unemployed has long-run effects, and so any adaptation is very weak.

We will come back to the consequences of adaptation soon enough, but I now want to return to the question of whether a person's choice reveals what makes them happy.

10.2 Do we know what makes us happy?

To illustrate some of the issues, imagine that someone called Ian has just come back from his annual summer holiday to Italy, and we want to find out how happy he is and why he chose to go.

Thinking about things from a revealed preference perspective, we work back from the fact that Ian went on holiday. We would say that his decision utility for a

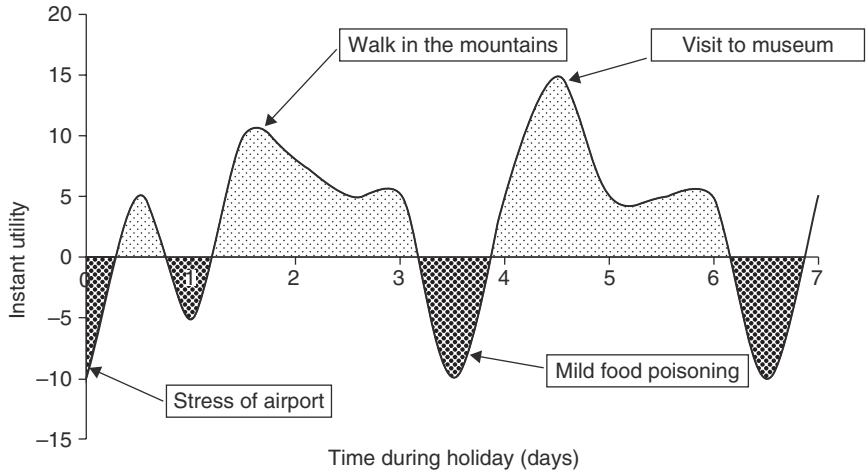

Figure 10.8 The instant utility over time of Ian on a holiday in Italy.

holiday in Italy was greater than that for a holiday in France, or staying at home and not spending the $2,000 he spent on flights and hotels. We know this because we know that he decided to go for a holiday in Italy. More generally, **decision utility** is revealed by a person's behavior. Whenever I have referred to utility in this book up until now I was really talking about decision utility.

Thinking about things from a slightly different perspective, we can start with the concept of instant utility. **Instant utility** measures the pleasure or pain that Ian experienced at specific points in time. Figure 10.8 plots what this might look like over the course of his holiday. Like any good holiday, we see some ups – a nice walk in the mountains – and some downs – mild food poisoning. From this we can calculate **total utility**, which is, basically, the sum of instant utilities. In the example, this would be the area under the instant utility curve, distinguishing between the positive area above the axis and the negative area below the axis. The bigger the positive area and the smaller the negative area, the greater is his total utility.

When we ask people how happy they are, we get a snapshot of their instant utility at that point in time. From this we could estimate total utility. Happiness data gives us, therefore, a means to measure total utility, and this looks like a good measure of happiness. The question we need to think about is whether total utility and decision utility coincide. If they do, then Ian's choice reveals what makes him happy, and so he must know consciously or subconsciously what makes him happy. Everything works nicely. But what if decision utility and total utility are different? Then things become less simple, and there are two good reasons why they can be different. I will look at each in turn.

10.2.1 Remembered utility

When looking back on past experiences, it seems that people typically do not remember total utility. I shall illustrate this with a study by Kahneman and co-authors (1993). The study involved subjects experiencing and evaluating pain.

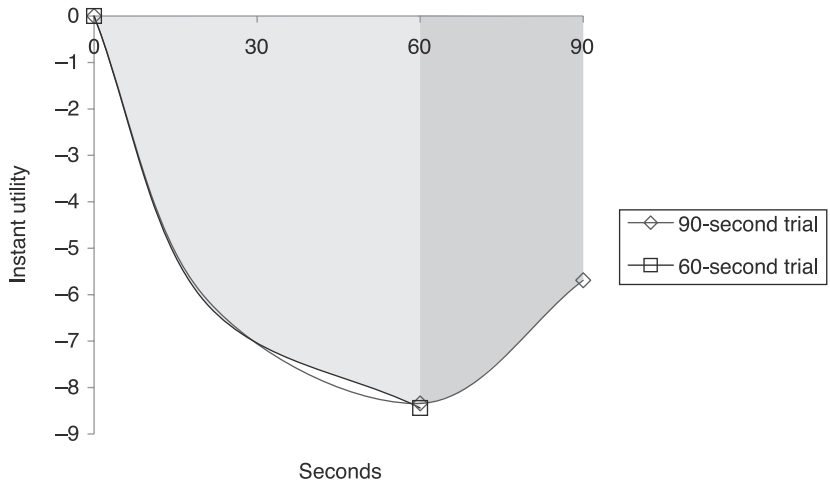

Figure 10.9 Instant utility while hand is immersed in cold water for 60 or 90 seconds.
Source: Kahneman *et al.* (1993).

Subjects were first asked to immerse their hand in cold water for 60 seconds. The water was kept at 14°C. Seven minutes later subjects were asked to put their other hand in cold water for 90 seconds. This time the water was kept at 14°C for 60 seconds and then gradually increased to 15°C. Seven minutes later each subject was told that they would be asked to put their hand in water a third time, but they could choose whether to repeat the first, 60-second, trial or the second, 90-second, trial.

After all this was done, subjects were asked to plot their level of discomfort during the trials. If we convert this to instant utility, we get something like Figure 10.9. The first thing to note is that the reported discomfort was the same in the 60-second trial as the first 60 seconds of the 90-second trial. This means that total utility had to be less for the 90-second than for the 60-second trial, because there was an extra 30 seconds of discomfort. Surely, therefore, people preferred to repeat the 60-second rather than the 90-second trial? In fact, 69 percent of subjects said that they would prefer to repeat the 90-second trial!

Figure 10.10 looks at how subjects rated the 90-second trial relative to the 60-second trial on several criteria. We see that subjects did correctly realize the 90-second trial was longer, but otherwise remembered it as less bad. Why would they do that? The only difference between the two trials, and the only thing that could explain it, is that the 90-second trial ends better. If people put extra weight on the last thing that they remember about an event, we have a plausible explanation for why subjects preferred the 90-second trial.

This study, and many others, suggests that people do not always make decisions based on total utility. It also suggests that we need to take seriously the idea that people may not remember very well how much utility an event gave them. We can use the term **remembered utility** to measure the amount of utility a person thinks an event gave them. Evidence also suggests that the remembered utility of an event

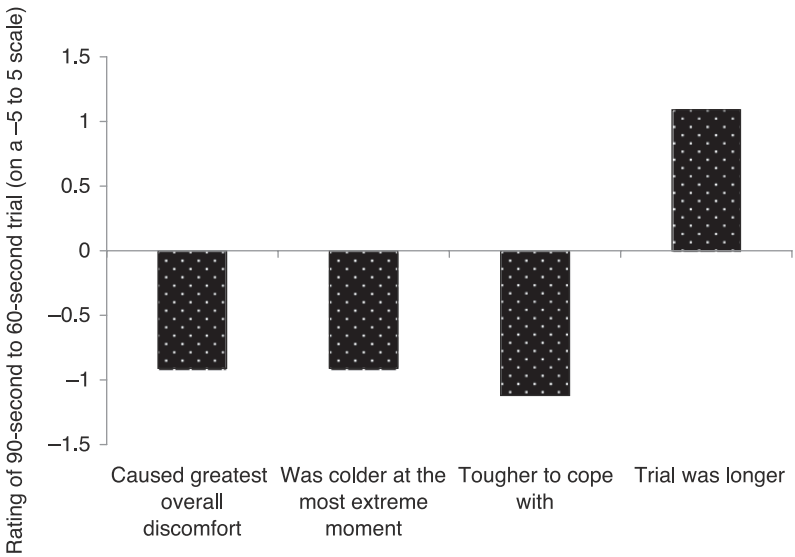

Figure 10.10 How the 90-second trial is rated compared to the 60-second trial. The 90-second trial was rated better on everything except length.
Source: Kahneman *et al.* (1993).

is accurately predicted by averaging the peak level of instant utility and end level of instant utility. This is called **peak–end evaluation**.

Ian's holiday will thus be remembered by the instant utility at the best and/or worst moment, and the instant utility at the end of the holiday. It is clear that this can mean a significant difference between remembered utility and total utility, or between decision utility and total utility. Do not necessarily expect, therefore, that people choose things that will give them highest total utility. Maybe Ian should have stayed at home and saved $2,000?

Thanks to a study by Wirtz and co-authors (2003) we can answer this question a bit more definitively. They tracked 41 students as they went on their spring vacation. The students filled in a questionnaire at various times before, during and after the vacation. Figure 10.11 summarizes the results. Consistent with the earlier discussion, we find a mismatch between total utility and remembered utility. The interesting, new thing is that subjects overplayed both the positive *and* the negative emotions experienced on the vacation. This illustrates how expectations and memories tend to be more intense than reality. Ian will not only remember the holiday as better than it really was, he will also remember it as worse than it really was!

Such differences between total and remembered utility raise some difficult but fascinating questions. For instance, peak–end evaluation implies there will be duration neglect. **Duration neglect** is when the duration of an event does not influence someone's memory of how good or bad the event was. This means that people can forget lots of good, or bad, moments and remember an event based on only a few instances. So, what is better: an event that gives higher total utility or an event that will give higher remembered utility? Was the 60-second trial better because subjects suffered less total discomfort or was the 90-second trial better because subjects

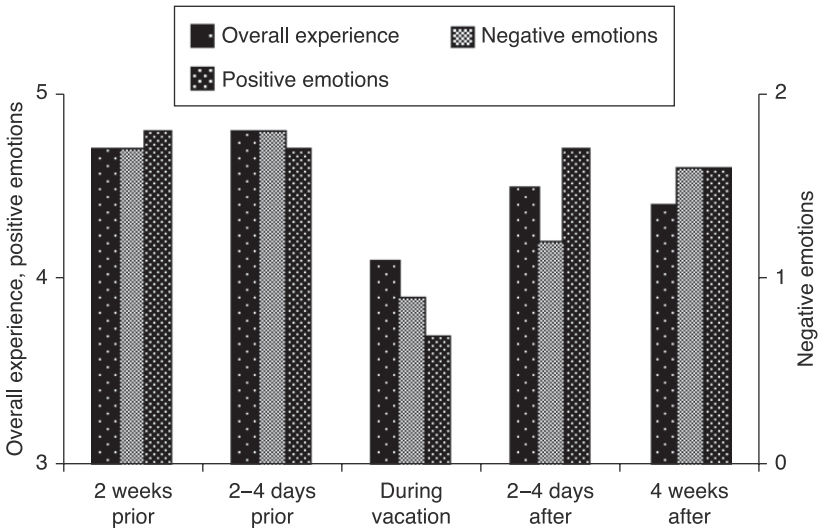

Figure 10.11 Students' expectations before, experience during and recollections after a spring vacation.
Source: Wirtz *et al.* (2003).

remembered less discomfort? What if we are talking about 60 or 90 days, months or years, rather than seconds?

10.2.2 Projection bias

Having seen that people may not be accurate in looking back at their past, we next ask whether they are good at predicting the future. Adaptation raises the possibility that people may not be good at this either. We need to question whether people will predict the extent to which their tastes and preferences adapt over time. The simple answer seems to be that most people correctly predict the direction of change, but not the magnitude. This leads to **projection bias**, in which people overestimate how much future preferences and tastes will resemble current preferences and tastes. Projection bias can show up in various different ways that I want to discuss and distinguish.

There is an **impact bias** if a person expects an event to have a bigger effect than it subsequently does. To illustrate, we can look at a study by Smith and co-authors (2008), who surveyed patients as they were waiting for a renal transplant, and again 12 months after they had had the transplant. Figure 10.12 summarizes the main results. An impact bias is shown by the patients predicting a greater increase in quality of life than actually occurred. We also see that patients remembered their quality of life as worse than it was. This could be because after the transplant we observe remembered utility, while beforehand we observe instant utility.

The opposite of the impact bias is when a person expects an event to have no effect but it subsequently does. One possible cause of this is the endowment effect that we looked at in Chapter 2. Recall that the endowment effect causes people to

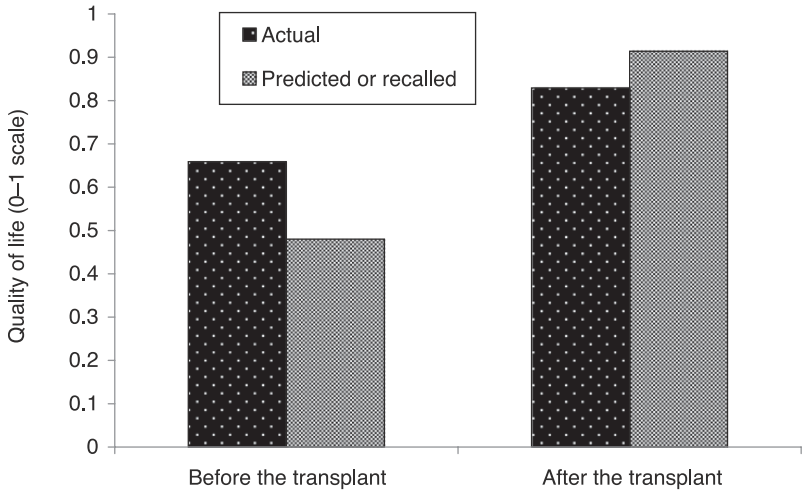

Figure 10.12 Predicted or recalled quality of life compared to actual before and after a renal transplant. Quality of life increases less than expected. People recalled quality of life being worse than it was.
Source: Smith *et al.* (2008).

value an item more when they own it. Do people predict this effect? Loewenstein and Adler (1995) asked this question in a study that involved subjects valuing a mug before and after they were given it. Figure 10.13 summarizes the results. We can see that, before they got the mug, subjects said they would be willing to sell it for around $3.50 to $3.70. Having been given the mug, subjects increased their asking price. The control group, who were not influenced by giving a prediction before receiving the mug, provide a good measure of willingness to sell the mug once owned, and valued it at 50 to 100 percent more. This means that subjects underestimated the effect that ownership of the mug would have on their willingness to sell.

The impact bias and a failure to predict the endowment effect suggest that people are poor at predicting how they will adapt to events and how their reference level may change over time. The accident victim may overestimate the long-run effects of disability, and the lottery winner may underestimate how much his reference level will increase as he becomes accustomed to his new wealth. The final effect I want to look at shows that people can underestimate the effect of short-run mood.

I will illustrate with a study by Read and van Leeuwen (1998). They asked 200 office workers in Amsterdam to choose what, from a range of healthy and unhealthy snacks, they would want to get in one week's time. The workers were asked either after lunch or late in the afternoon and were told that they would be given the snack after lunch or late afternoon. The expectation was that people would be hungry late afternoon but satiated after lunch, and this would affect choice. The results in Figure 10.14 suggest that it did. There is no good reason why current hunger should affect what the person will want in one week's time. For many, however, it seemed that current hunger did affect what the person expected they would want in one

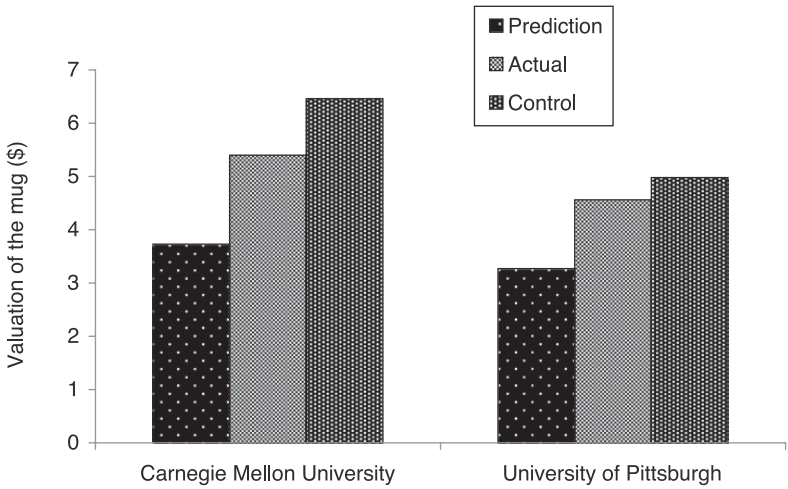

Figure 10.13 The endowment effect. Subjects failed to predict how ownership of a mug would increase their valuation of the mug.
Source: Loewenstein and Adler (1995).

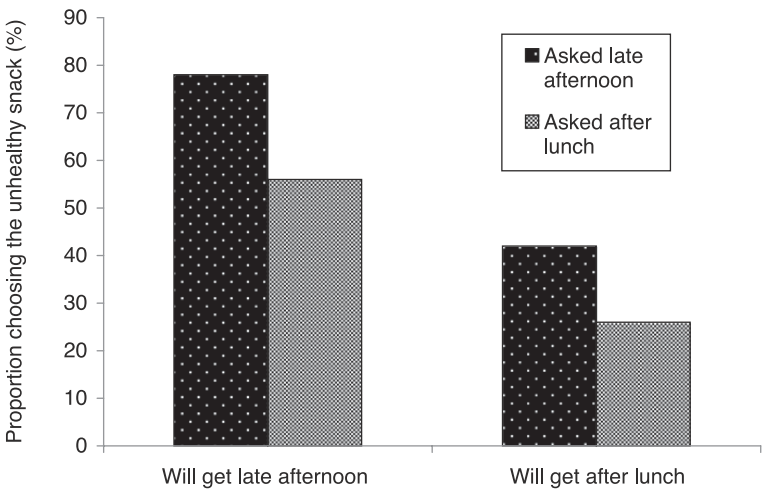

Figure 10.14 The proportion who chose an unhealthy snack for one week's time depends on current hunger.
Source: Read and van Leeuwen (1998).

week's time. Those who were relatively hungry when asked were somewhat more likely to choose the unhealthy snack.

In summary, we observe a projection bias in which people overestimate how much future preferences and tastes will resemble current preferences and tastes, underestimate how much they will adapt to events and are influenced by current mood.

10.3 Choice and commitment

Projection bias and the difference between decision and remembered utility suggest that people may not know what is best for them. We have already seen elsewhere in this book that, even if people do know what is best for them, they may not make the best choices because of, say, present bias. Things do not look, therefore, very promising. But can people overcome projection bias and time inconsistency? One way they might is by pre-commitment, either by pre-committing themselves to a plan of action or by allowing someone else to choose for them. What I want to do in this section is explore how well this might work.

To get us started on this question, let's stick with the subject of food a bit longer. After one week the workers in the study by Read and van Leeuwen were found and reoffered the choice of snacks. This time they got the snack they chose. As we can see in Figure 10.15, a lot of those who said they would want the healthy snack actually chose an unhealthy snack. This looks like the dynamic time inconsistency problem that we saw in Chapter 4. Interestingly, however, it seems that many stuck by the choice of a healthy snack that they knew they had made beforehand. This would explain why the time they were initially asked affected subsequent choice.

We can follow this up with a study by Gilbert, Gill and Wilson (2002) that looked at the effect of hunger on grocery shopping. Shoppers in a grocery store were asked to take part in a survey in which they had to list their planned shopping. Some were given back the list and some were not, and some were given a muffin to try before shopping and some not. The number of unlisted items that the shopper subsequently bought was recorded. Figure 10.16 summarizes the main result. Not surprisingly, shoppers purchased more items than they put on the list. Consistent with projection bias, we see that those who were relatively hungry, because they did not eat the

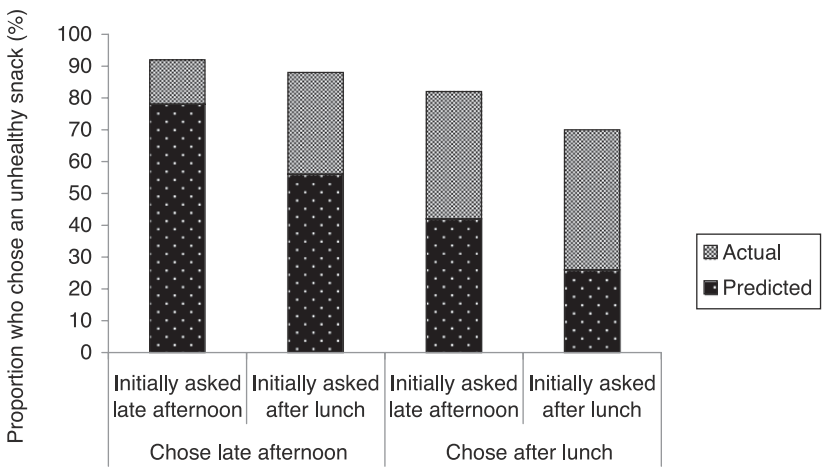

Figure 10.15 More chose an unhealthy snack than predicted they would, but fewer chose the unhealthy snack if they said they would choose a healthy snack.
Source: Read and van Leeuwen (1998).

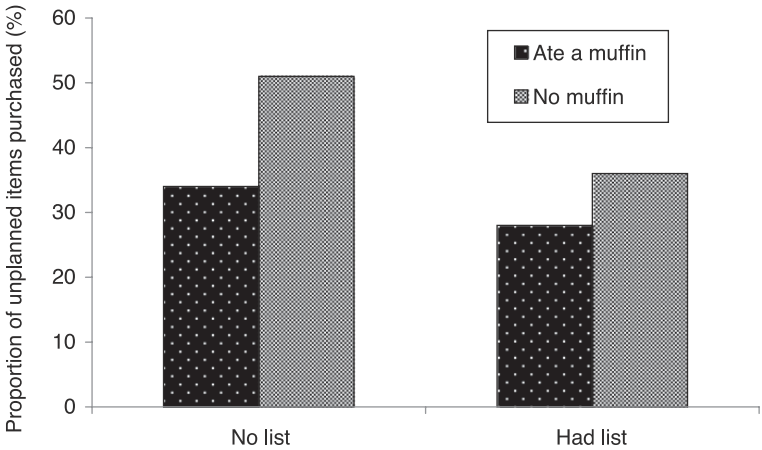

Figure 10.16 How many unplanned items shoppers bought depended on whether or not they ate a muffin and whether or not they had a shopping list to hand.
Source: Gilbert, Gill and Wilson (2002).

muffin, bought more unplanned items. Interestingly, however, those with the shopping list to hand were seemingly much less affected by their current mood.

In both these examples, behavior seemed to be influenced by a loose pre-commitment. In the first case, knowing that they said they would eat a healthy snack seemed enough to persuade some to stick with their choice. In the second case, having what they said written down, in the form of a shopping list, made people stick more closely to what they had planned. Potentially, therefore, pre-commitment is a good thing.

10.3.1 Does present bias matter?

Before we look in more detail at whether pre-commitment can work, it seems pertinent to take a step back and ask whether things such as projection bias and present bias do really lower utility. In section 10.6.1 I will show you that projection bias can seriously lower utility. Here, I want to look at present bias and ask whether present bias lowers utility.

The first problem we have in answering this question is how to evaluate utility. The measure I am going to use takes a long-run perspective in which we ignore any present bias, and say that total utility, if utility in periods one to T is u_1 to u_T, is

$$U^0(u_1, u_2, \ ..., u_T) = u_1 + \delta u_2 + ... + \delta^{T-1} u_T = u_1 + \sum_{\tau=1}^{T} \delta^{\tau-1} u_\tau$$

where δ is the long-run discount factor. This definition has its problems, which you might want to think about, but it's as good as we have.

Now imagine that Ian, fresh from his holiday, is deciding when to finish his dissertation. It's April 1 today, and the work has to be handed in by April 30. It plays on his mind that it is not done, and so his utility is 99 on every day it is not done and then

Table 10.2 Ian's utility depends on when he finishes the dissertation. Total utility ignores any present bias, while decision utility takes into account present bias. Assuming that $\delta = 1$ and $\beta = 0.98$, total utility is highest if he does the dissertation immediately, but present bias means he will constantly put it off until tomorrow and do it on April 30.

Plan	Utility on each day					Total utility	Decision utility	
	1st	2nd	3rd	4th	30th		Today	Tomorrow
Do it today	0	100	100	100	100	2,900	2,842	—
Do it tomorrow	99	0	100	100	100	2,899	2,843	2,744
Do it on 3rd	99	99	0	100	100	2,898	2,842	2,745
Do it on 30th	99	99	99	99	0	2,871	2,816	2,719

100 a day once it is finished. He knows that it would take just one day's work to finish it, and on that day his utility would be zero. Table 10.2 illustrates how things look on April 1. It is clear that Ian gets the highest total utility if he does the dissertation immediately. But what if he has present bias?

Recall (you might want to look back at Chapter 4 to refresh your memory) that we can write decision utility as

$$U^T(u_1, u_2, \ ..., u_T) = u_1 + \beta\delta u_2 + ... + \beta\delta^{T-1}u_T = u_1 + \beta\sum_{\tau=1}^{T}\delta^{\tau-1}u_\tau$$

where β measures the present bias. Let us assume $\beta = 0.98$, meaning Ian has only a very slight present bias. We can see that this present bias is enough that the decision utility is higher for doing the dissertation tomorrow. This, in itself, looks not too bad, because total utility is only one lower if he does it tomorrow. The problem is that tomorrow he will want to put it off until the next day, and so on. This means that, if he is naïve about his present bias, he will actually do the dissertation on April 30 and lose 29 in total utility.

Ian's naïve present bias will reduce his total utility by 1 percent. Nothing can be gained from arguing whether that is big or small; the main point we need to take from this is that a little bit of present bias can accumulate into a large loss in total utility. On any one day Ian is willing to sacrifice only one unit of total utility because of present bias. Unfortunately, he does not realize that he will end up sacrificing one unit every day, over lots of days, and so the total loss will be larger than he thought.

If Ian was sophisticated, and so knew about his present bias, he would do the dissertation straight away. Sophistication does not mean, however, that he always avoids the welfare costs of present bias.

To illustrate, we can think about Emma, who is sophisticated and eager to buy the latest book by her favorite author. As soon as she buys it she will read it and have utility 100. To complicate matters, the bookstore has a deal whereby the price of the book decreases by $1 every day, starting at $29. Table 10.3 summarizes the choice that Emma faces, again assuming $\beta = 0.98$. The best thing that Emma can do is wait until the price has gone down to $0. If she was naïve, that is what she would do, but, because she is sophisticated, she will buy today! That's because she can predict that on

Table 10.3 Emma's utility depends on when she buys the book. Assuming that $\delta = 1$ and $\beta = 0.98$, total utility is highest if she waits to buy the book. Present bias means she knows she will buy it earlier than that, so ultimately she buys it today.

	Utility on each day					Total utility	Decision utility	
	1st	2nd	28th	29th	30th		On 28th	On 29th
Buy it today	71	0	0	0	0	71	—	—
Buy it on 28th	0	0	98	0	0	98	98	—
Buy it on 29th	0	0	0	99	0	99	97	99
Buy it on 30th	0	0	0	0	100	100	98	98

the 29th she will buy, which means that on the 28th she should buy, and so on back to today. Again, small day-to-day losses accumulate to a much bigger total loss.

It turns out that these examples are representative of what happens in general. If the costs are immediate and benefits later, there can be a big loss in total utility from present bias for someone who is naïve. If the benefits are immediate and the costs later then there can be a big loss in total utility from present bias for someone who is sophisticated.

Clearly, therefore, present bias is a potential problem. If someone keeps on putting something off, or bringing something forward, the losses will accumulate. And being sophisticated is not enough to avoid such losses; it can make them worse.

10.3.2 Pre-commitment

Now we know that present bias and projection bias really are a problem, it becomes even more relevant to see whether pre-commitment can work. In principle, it can. Ian might commit to doing the dissertation by, say, making a bet with his friend that he will do it today. Emma might pre-commit by putting her money in a 'safe place' until the end of the month. In Chapter 4 (see, for example, section 4.2.3) I have already argued that people may want to pre-commit. Here, and in the next chapter, I want to take that argument a bit further, and look at some data. I will start with a study by Ariely and Wertenbroch (2002).

In the study, people were recruited to take part in a proofreading task that involved reading three texts of around ten pages. Recruits were paid for every error that they spotted, but had $1 deducted for each day of delay beyond a deadline. Recruits were randomly assigned to one of three conditions: an evenly spaced deadline condition, with deadlines set at seven-day intervals for each of the texts; an end-day condition, in which all three texts were to be handed in at the end of 21 days; or a self-imposed deadline condition, whereby recruits were invited to make their own deadlines for when to hand in the three texts over the next 21 days, with these deadlines subsequently being enforced.

If a person has no present bias then they should prefer the end or the self-imposed condition, as these give maximal flexibility when they do the task. If someone has a present bias and is sophisticated then they should prefer the self-imposed condition, as they can commit to a plan that is optimal for them. Figure 10.17 illustrates what happened. We see that, in terms of performance, those in the evenly spaced

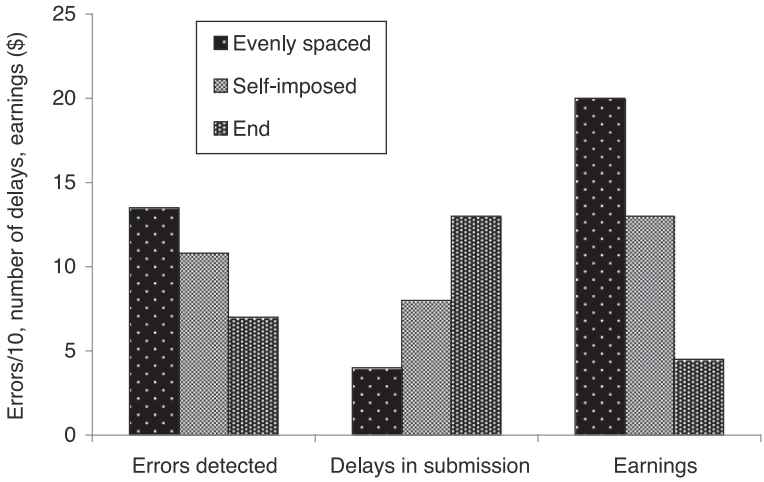

Figure 10.17 Performance in a proofreading task. Performance depended on whether the deadlines were evenly spaced, self-imposed or all at the end.
Source: Ariely and Wertenbroch (2002).

condition did better than those in the self-imposed condition, who did better than those in the end condition.

These results illustrate again that people can benefit from pre-commitment. A desire to pre-commit suggests that people are sophisticated to some extent about their present bias. Maybe, however, they do not choose the best plan when they do commit. The subjects in the previous study who self-imposed a plan did not do as well, for instance, as those with an evenly spaced plan.

To illustrate the consequences of this, recall the gym example in Chapter 4. Buying membership may be a way that Ian could try to overcome his present bias by pre-committing to going to the gym. The problem is, he might underestimate the strength of his present bias and so still not go to the gym. He will have ended up paying a lot of money for a pre-commitment device that does not work! Pre-commitment, therefore, is not a panacea for present bias. It might work for some people in some situations. But sometimes it might make matters worse, such as for Ian not going to the gym.

If people do not always make great choices, and others could potentially make better choices, it seems natural to let those others decide. Indeed, even if people can make great choices, they might still prefer someone else to decide so that they can avoid the psychological costs of temptation and the like. This seems to go partly against ideals of liberty and free choice, but if people were to choose not to choose then the ideal of free choice becomes somewhat confused. To progress on this issue, it's useful to ask whether people like having choice or not.

10.3.3 Do people like having choice?

To illustrate how people respond to choice, I am going to look at a study by Iyengar and Lepper (2000). In the study, subjects were asked to select one chocolate from a selection of either six, a limited choice, or 30, an extensive choice. Half the subjects

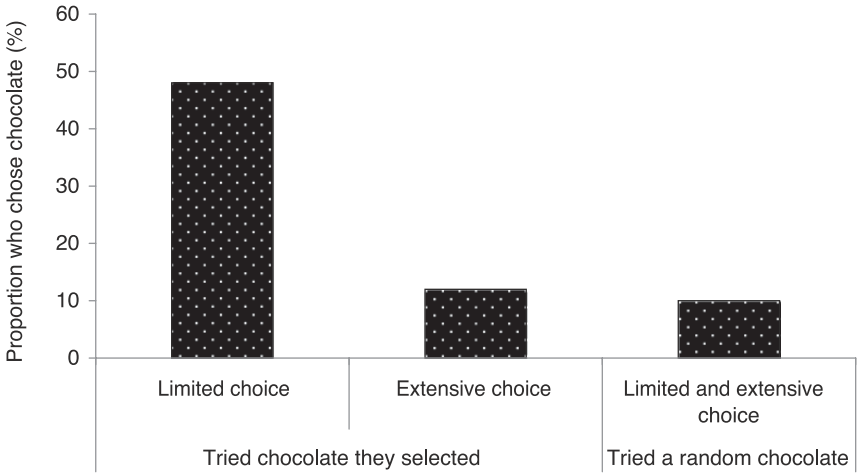

Figure 10.18 The proportion who chose chocolate rather than money depending on whether there was a limited or extensive choice, and whether they got to taste a chocolate they chose or the experimenter chose.
Source: Iyengar and Lepper (2000).

then got to taste the chocolate they chose, and half got to taste a chocolate selected by the experimenter. On leaving the experiment, as payment for participation, subjects were given the choice of receiving $5 cash or a box of chocolates worth $5. The main result is clear from Figure 10.18, where we see that the proportion choosing chocolate over cash was far higher for subjects asked to choose from a limited number of chocolates and allowed to taste the one they selected.

Why is there such a difference? Figure 10.19 offers some clues. There are significant differences in the perceptions of the participants depending on whether they had a limited or extensive choice. Interestingly, participants did appear to enjoy choosing more when there was a more extensive choice. On all other measures, however, participants seemed to benefit from a more limited choice: it was quicker, less difficult and less frustrating and they were more satisfied with their choice.

An experiment about chocolate seems a long way from us making conclusions on the benefits of free choice, but I like this study because it nicely illustrates two results observed much more generally. The first result is that more choice is not always a good thing. More choice can be bad, because it complicates the decision and can lead to temptation and a need for self-control. This is consistent with what we saw in Chapter 2, where many preferred to delay a difficult decision, or seemed overwhelmed when there were too many choices on offer.

Should we, therefore, be taking choice away from people? No, because the second result is that people like having a choice. We see this in the study by Iyengar and Lepper, in which subjects were more likely to want a box of chocolates than cash if they had tasted the chocolate they chose during the experiment – i.e. their choice mattered. More generally, the evidence is strong that people do like to make choices. One explanation is that people have a greater interest and motivation in things that they have **self-determined** through their own actions.

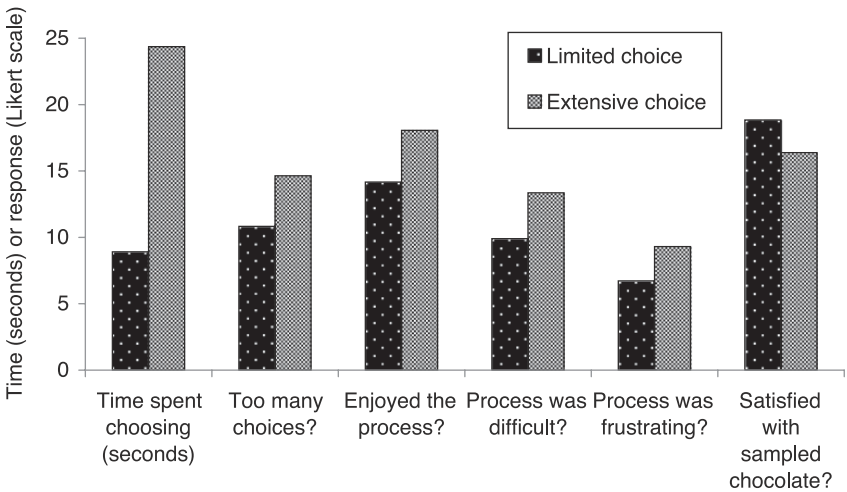

Figure 10.19 How opinions and behavior changed if subjects had a limited choice or extensive choice of chocolates.
Source: Iyengar and Lepper (2000).

So, when faced with an increased choice of options, many people defer making a decision or feel less satisfied with the decision they make. This is called **choice overload**. People do, however, like to have some control and choice.

Putting these things together gives a somewhat bizarre conclusion. People like having choice, but do not like difficult choice. What people should like most, therefore, are easy choices! There's not much evidence to say whether this is a fair conclusion or not, but what we do have is largely supportive. That's because choice seems to lead to the largest increases in motivation in experiments when the choice is largely irrelevant, such as choosing what color pen to use or what name to give a game character. This raises some interesting policy implications that we can come back to in the next chapter.

10.4 Summary

Instant utility is how much utility a person gets at a specific point in time, and can be thought of as a measure of happiness. We saw that a person's instant utility, or level of happiness, appears to depend on their position relative to others and relative to what has happened to them in the past. In particular, we see varying degrees of adaptation to life-changing events.

We then distinguished between decision utility, total utility and remembered utility. Roughly speaking, decision utility is how much utility a person thought they would get from doing something, when they decided to do it. Total utility is how much utility they actually got from doing it. Remembered utility is how much utility they think they got from doing it.

We saw that decision utility, total utility and remembered utility can be different because of things such as duration neglect and projection bias. This suggests that people may not know what things would make them happiest.

We then revisited the issue of present bias and showed that it can cause a big drop in utility. This suggests that people may not choose the thing that would make them happiest even if they know what they should choose.

Given the suggestion that people may not always make good choices, we questioned whether people like having choice. We saw that many people like to precommit and do not appear to like difficult choices. Generally, however, people appear to enjoy some level of choice and self-determination.

We get, therefore, an interesting picture of people not always knowing what makes them happy, or will make them happy. We also see people willing to precommit or delegate choice to others, provided they retain some freedom of choice. This seemingly provides a clear rationale for policy intervention – if, of course, we can trust policy-makers to know any better what makes people happy!

10.5 Health and happiness

Health care is one context in which it can be vital to understand what makes people happy, because the potential benefits or costs of being ill and getting treated can be so large. In this section I shall look at two possible applications of what we have learnt about the nature of happiness and utility.

10.5.1 Measuring the value of treatment

Valuing the costs and benefits of medical treatments is fundamental to allocating health care resources efficiently and most effectively. This is particularly the case in state-financed health care systems, such as the National Health Service in the United Kingdom. Given the limited resources available, priority should be given to treatments that have higher net benefits than others. This notion is often controversial, because it means denying treatment to some patients. In particular, patients are denied treatment if the net benefit of treating other patients is higher. Such economic logic is a hard sell when it comes to denying cancer patients much-needed drugs. The logic behind such rationing is sound enough, however, given the limited resources available. More difficult is accurately measuring the net benefit of treatments, and that's what I want to look at in more detail here.

To calculate the net benefit of a treatment we need some measure of its cost and benefits. The cost is a relatively objective thing to measure in principle. For example, the cost of a drug or the cost of a physician's time can easily be calculated. Much harder is to measure the benefits. There are two possible dimensions to any benefit: a treatment could improve the quality of life or prolong the length of life. How can we possibly measure the value of an improvement in quality of life, or an extra year of life?

One widely used way to do so is the QALY. Each year of perfect health is counted as one **quality-adjusted life year** (**QALY**). A year of less than perfect health is given a QALY value of less than one. The lower the quality of life, the lower the value, with zero being equivalent to death. To get the overall value of a treatment we follow the QALY value of a treatment over the subsequent life of the patient. To illustrate, Figure 10.20 tracks the quality of life of a patient if they have no treatment or two possible treatments, A and B. Without treatment, the quality of life deteriorates. With

treatment A, the patient is stabilized and their life extended. With treatment B, the patient's quality of life initially deteriorates, because of the side effects, but subsequently increases to a relatively high level.

Without treatment, the patient gets 1.5 QALYs because there are 0.5 QALYs in the first year, 0.4 in the second year, and so on, adding up to 1.5. With treatment A, the patient gets 2.4 QALYs, and, with treatment B, they get 2.7. The net benefit of treatment A is, therefore, 2.4 − 1.5 = 0.9 QALYs, and that of treatment B is 1.2 QALYs.

In this example it looks as if treatment B is better than treatment A. We do need to take into account the cost, however. Suppose treatment A costs $10,000 and treatment B costs $20,000. Then, with treatment A, each QALY costs $4,167, but, with treatment B, it costs $7,407. In a publicly funded health service someone has to decide whether these amounts of money are worth paying. In the United Kingdom an organization called the National Institute for Health and Care Excellence (with the sometimes unfortunate acronym NICE) does that, and treatments below $45,000 a QALY are usually funded. In the US system things are less transparent, but a $50,000 per QALY ratio appears to be applied for much publicly funded treatment.

In order to get the data to calculate QALYs, the EQ-5D system is often used. Patients are asked to fill out a questionnaire that asks five questions about mobility, pain, self-care, anxiety and ability to do normal activities. There are three possible answers to each question. A formula is then used to convert these answers into a QALY value. Table 10.4 gives some examples. If patients are tracked and fill in the EQ-5D questionnaire over time, we can obtain the data to draw Figure 10.20. So, everything looks in place to make things work. The whole process hinges on the formula used to derive the numbers in Table 10.4, however, and this is where the things that we have learnt can start proving useful.

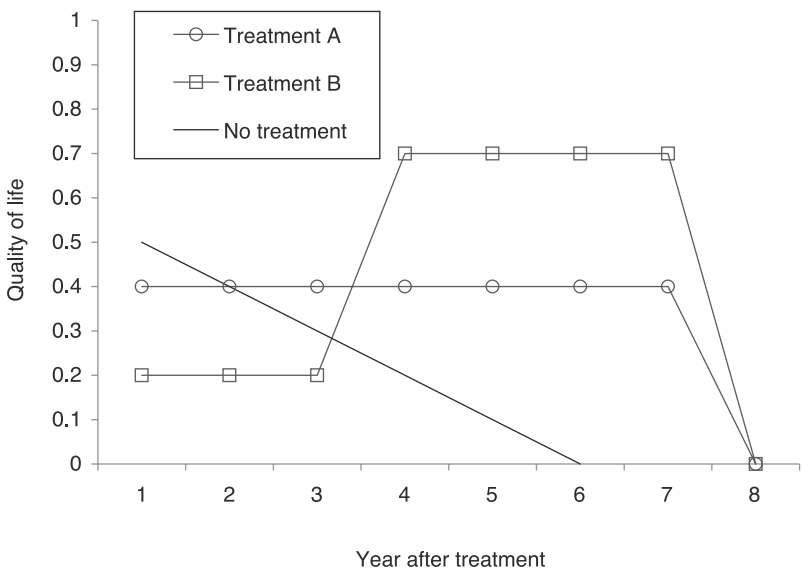

Figure 10.20 Changes in quality of life following two possible treatments, A and B, or no treatment.

Table 10.4 Examples of QALY measures using the EQ-5D system.

Description	QALY
No problems	1.000
No problems walking about; no problems with self-care; some problems with performing usual activities; some pain or discomfort; not anxious or depressed	0.760
No problems walking about; some problems washing or dressing self; unable to perform usual activities; moderate pain or discomfort; extremely anxious or depressed	0.222
Confined to bed; unable to wash or dress self; unable to perform usual activities; extreme pain or discomfort; moderately anxious or depressed	–0.429

Source: Phillips and Thompson (2009).

The formula used to calculate a QALY measure is derived using a **time trade-off valuation technique** or **visual analogue scale**. In both cases, a random sample of people are asked to imagine a particular quality of life, such as those in Table 10.4. The time trade-off technique then asks them to say what length of time in full health they consider to be equivalent to a longer period with this quality of life. The visual analogue scale directly asks them to rate the quality of life on a linear scale. Are people able to make accurate judgments of this sort?

Given what we know about adaptation and habituation, the answer is probably 'No'. People who are relatively healthy are likely to underestimate the quality of life of those in poorer health than they are. Conversely, people who are relatively unhealthy are likely to overestimate the quality of life of those who are healthier than they are. This matters if people find it easier to adapt to some states of ill health than others, and we have seen that adaptation rates can vary. To illustrate the problem, Figure 10.21 plots how a person might adapt to treatment. We see that the person adapts very well to the health state resulting from treatment B and so would eventually report a quality of life near to full health. They adapt less well to the health state resulting from treatment A or no treatment.

If we ask non-patients to think what it would be like to have such treatments, we are likely to get the low numbers that patients really do feel one year after treatment. This ignores the possibility of adaptation and makes treatment A look better. If we ask patients after six or seven years of treatment, they would be more positive and feel particularly positive about treatment B. This ignores the fact that they have adapted. The key thing we need to do is not only rate health states such as those in Table 10.4 but also to think how quickly patients will likely adapt to these health states. This should then enter the calculation for working out QALYs.

Unfortunately, it is not current practice to take into account adaptation when calculating QALYs. One reason is probably that the notion of adaptation can seem a little controversial, because it raises the question: if people are going to adapt to ill health, why do they need treatment? This looks a tricky question until we compare Figures 10.20 and 10.21. If we look at what happens with no treatment, adaptation does indeed improve the quality of life. The crucial thing, however, is that adaptation increases the quality of life from having the treatment by even more! So, yes, people may adapt to their poor health state, but they may also adapt to the potentially less than perfect health state that results from treatment. Adaptation does not, therefore, mean that we are any more or less likely to recommend treatment.

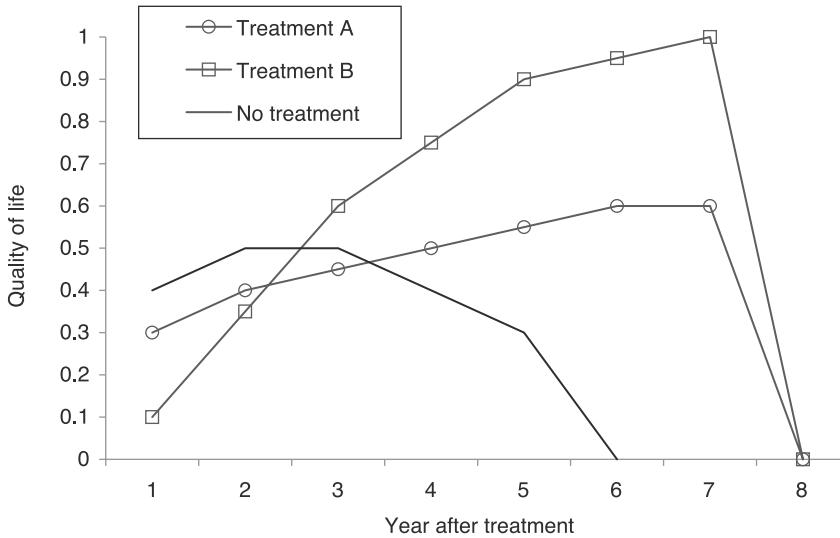

Figure 10.21 Changes in quality of life following two possible treatments, taking into account adaptation.

Valuing improvements in quality of life is controversial and difficult, but if we are to make the best health choices then we need these measures. I have only touched on the issue here but, hopefully, you have seen how the ideas we have looked at in this chapter, particularly adaptation, can help inform the debate.

10.5.2 Improving the remembered utility of treatment

For effective prevention of health problems it is important that people turn up for relevant screening procedures. For example, colonoscopy is a screening procedure for colorectal cancer, one of the leading causes of cancer-related death. Colonoscopy saves lives, but only if patients are willing to undergo the fairly unpleasant procedure. Whether or not they are willing will depend to some extent on their recollections of previous treatments or procedures. Someone who remembers their last colonoscopy as painful and unpleasant may be less likely to turn up for a subsequent colonoscopy, and that could be very bad. This means that anything that can improve the remembered utility of the procedure should be a good thing. Redelmeier, Katz and Kahneman (2003) investigated whether this can be done.

In the study, outpatients who were due for a colonoscopy were asked to take part. During the procedure patients were given a handheld device with which they could indicate their discomfort on a scale of 0 to 10. Afterwards patients and physicians were asked questions about how much pain and discomfort were experienced during the procedure. The procedure could last anything from a few minutes to over an hour. The striking thing, however, was that patients' recollections of pain and discomfort did not depend on how long the procedure lasted. Instead, they depended on the peak and end pain. This is further evidence of duration neglect.

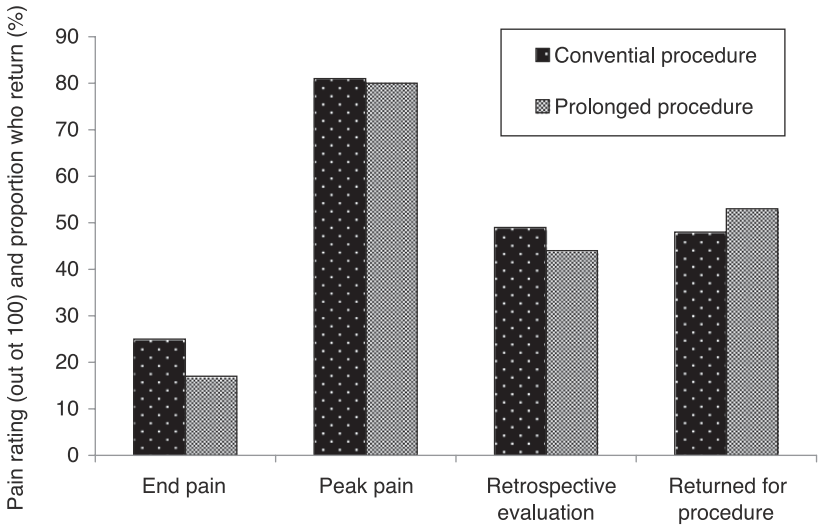

Figure 10.22 Measures of pain and rate of return for a repeat colonoscopy depending on whether the procedure was prolonged or not.
Source: Redelmeier, Katz and Kahneman (2003).

Duration neglect and peak–end evaluation suggest that extending the duration of a procedure while decreasing the end pain could make patients remember it less negatively. To test for this, a randomized trial was done in which the colonoscopy procedure was extended by three minutes for selected patients (see Research methods 10.3). During the extra three minutes there would have been discomfort, but less than during the main part of the procedure. Figure 10.22 summarizes some of the data. Consistent with peak–end evaluation, end pain is less with the prolonged procedure and the retrospective evaluation of pain is less. Patients were subsequently tracked to see whether they returned for a repeat colonoscopy, for which the median follow-up time is around five years. Those who took part in the prolonged procedure were marginally more likely to return.

The effect on the proportion of patients who returned for a repeat colonoscopy was not huge in this study; but, if it comes from waiting just three extra minutes, and if it means one or two people suffer less from cancer, then that is a big gain. We do not know whether other more significant changes could have larger effects, but it is an intriguing possibility. This possibility arose from distinguishing between remembered and total utility.

❖ RESEARCH METHODS 10.3 DOUBLE-BLIND EXPERIMENTS

In an experiment with multiple treatments it is sometimes a concern that the researcher knows what treatment a particular subject is in. If, for instance, the clinicians, in the study by Redelmeier, Katz and Kahneman, knew that a particular patient was in the treatment with a prolonged procedure, then they

might have treated the patient differently during the procedure. We need to avoid this, because it can clearly affect the results.

A **double-blind experiment** is one in which the researcher does not know what treatment is being run during an experimental session. Generally, this is not very simple to achieve, but there are ways to make it happen. In this particular study, the treatment that a subject was assigned to was randomly determined and written down on a piece of paper in a sealed envelope. Only when the colonoscopy procedure was near to completion would the envelope be opened and the clinicians know whether to extend the procedure or not: a neat way to obtain double blindness.

10.6 Saving and retirement

In Chapters 2 and 4 we spent some time looking at the implications of choice and time inconsistency for saving, borrowing and the life cycle. It is natural to pick up these topics again in this chapter. I will start by looking at the implications of projection bias.

10.6.1 Projection bias in saving

In Chapter 4 we looked at a model of habit formation in which Maria would become habituated to past levels of consumption. Such habituation is representative of adaptation, and we showed that it could lead to an increase in saving following an increase in income. At that stage we did not need to assume any present bias. In the last chapter, when looking at addiction, we looked more at how present bias can matter once someone becomes habituated. What I want to do now is add projection bias to the story.

To see why projection bias matters, we can work through an example. Suppose that Ian's utility in year t is

$$u(c_t, h_t) = \begin{cases} \sqrt{c_t - h_t} & \text{if } c_t \geq h_t \\ -\sqrt{h_t - c_t} & \text{if } c_t < h_t \end{cases},$$

where c_t is his consumption and

$$h_t = c_{t-1}$$

is the reference level of consumption to which he is habituated. So, Ian's reference level adapts every year to his consumption of the previous year. I leave you to think about how this relates to the model we looked at in section 4.6.1.

Suppose that, because of projection bias, Ian underestimates how much his reference level will adapt. Specifically, suppose he believes that

$$h_t = \alpha h_{t-1} + (1 - \alpha)c_{t-1},$$

Table 10.5 The consequences of projection bias in a simple example. Ian's income will fall in year 3. If he has no projection bias he will choose a gradually increasing consumption profile and get total utility 3. With projection bias he chooses a flatter consumption profile and gets less utility than he expects.

	Time				Projected utility	Actual utility
	Year 0	Year 1	Year 2	Year 3		
Income	1	4	4	1	—	—
Planned consumption						
$\alpha = 0$	1	2	3	4	3.00	3.00
$\alpha = 1$	1	3	3	3	4.23	1.41
$\alpha = 0.5$ and in year 0	1	2.33	3	3.67	3.47	2.79
$\alpha = 0.5$ and in year 1	1	2.33	3.14	3.53	2.94	2.68

where α is some parameter. If $\alpha = 0$, Ian does not have projection bias, but, if $\alpha > 0$, he does.

Table 10.5 illustrates what consequences a projection bias could have. The example has been set up so that Ian's income is going to increase in the next two years and then fall in year 3. You could think of this as a condensed version of working and then retiring. How can Ian best smooth his consumption? The best he could do would be to increase his consumption slowly. This is what he will do if $\alpha = 0$, and so he has no projection bias. What if he does have a projection bias?

The easiest case to look at is when he has a very strong bias, because $\alpha = 1$. In this case, Ian completely fails to realize that he will adapt to increased consumption. Consequently, he maximizes his projected utility by increasing consumption as much as possible relative to his current reference level. This does not work very well, because his reference level does adapt, and so, while he is happy in year 1, he is not so happy in years 2 and 3.

More realistic is the case when $\alpha = 0.5$, and so Ian knows that he adapts but underestimates how much he adapts. He would now maximize projected utility by choosing a plan intermediate between those he would choose when $\alpha = 0$ and $\alpha = 1$. The consumption profile is increasing, but flatter than optimal. The most interesting thing is what happens after one year. At this point Ian realizes that he underestimated how much he would adapt, and so would want to change his plan in order to maximize projected utility over years 2 and 3. What's interesting is that the new plan that he comes up with is worse than the original one! Basically, Ian knows that he needs to increase his consumption more now than he thought he would have to in order to be above his reference level. This means he brings forward consumption and ends up doing worse in year 3.

Stylized though it is, this example illustrates two key implications of projection bias for saving and retirement. First, Ian saves less for retirement than he should have done. Second, he will be time-inconsistent and save less for retirement than he planned to do. We see this in the 0.14 difference between how much Ian thinks he will save in year 2 and how much he actually will save.

Both these implications are bad, and can accumulate over many years to a real shortfall in savings for retirement. Could we avoid this? In principle the second

problem should, at least, be avoidable, because Ian might be willing to pre-commit to a savings plan that does not allow him to change his mind. In the next chapter we shall see how such plans may work. The first problem could be avoided by 'forcing' Ian to save more, but this brings us back to the issue of choice and autonomy.

10.6.2 Investor autonomy

One way to 'nudge' people such as Ian to save more for their retirement is to automatically enroll them in a savings plan. We are going to look at this in the next chapter, but you could probably guess from what we did in Chapter 4 that people usually stick to such plans. Even if they wanted to withdraw, present bias means they 'put if off until tomorrow' and never do get around to changing. So, let's imagine for now that automatic enrollment is like 'forcing' Ian to save.

This might be a good thing, but we encounter a problem. If we are going to automatically enroll Ian in a savings plan, we have to decide where to invest, how much to invest, and so on. Should we also decide that, or let Ian choose? This is a fascinating situation to question whether choice is good or not. Let's look at a study by Benartzi and Thaler (2002).

In the first part of the study, they collected information about people enrolled in the University of California – Los Angeles (UCLA) retirement plan. There is no default option in this plan, so participants have to choose their own portfolio. Based on this information, they predicted the retirement income a person could expect if invested in (1) the portfolio he or she was invested in, (2) the average portfolio and (3) the median portfolio. (Just to clarify, the median portfolio will be significantly different from the average portfolio, because many invested a lot in cash, and so the average portfolio is skewed towards investing more in cash.) Participants in the study were then asked to rate, on a scale of 1 to 5, how attractive a portfolio appeared, without knowing which one was theirs. As we can see in Figure 10.23, the median portfolio received a higher rating than their own portfolio. Indeed, only 21 percent of participants rated their own portfolio highest!

In the second part of the study, Benartzi and Thaler surveyed employees at SwedishAmerican Health Systems, based in Rockford, Illinois. Employees at SwedishAmerican are automatically provided with an individual portfolio selected by an investment management firm, but have the possibility to opt out and choose their own portfolio. Focusing only on employees who had decided to opt out, they again predicted retirement income and asked participants how they would rank (1) their own portfolio, (2) the average portfolio and (3) the default portfolio suggested by the investment management firm. This time we see in Figure 10.23 that the default portfolio gets the highest ranking. Only 20 percent preferred their own portfolio, and recall that participants had specifically opted out of the default option in order to choose their own portfolio!

These results suggest that investor autonomy is not worth much. In both cases, participants rated higher a portfolio that was not chosen by them. We might be concerned that this is attributable to their having different expectations about what might happen to the stock market, etc. When surveyed, however, most participants expressed no strong beliefs about what would be likely to happen. The more plausible explanation seems to be that participants did not have the skills and information to pick the best portfolio for their preferences.

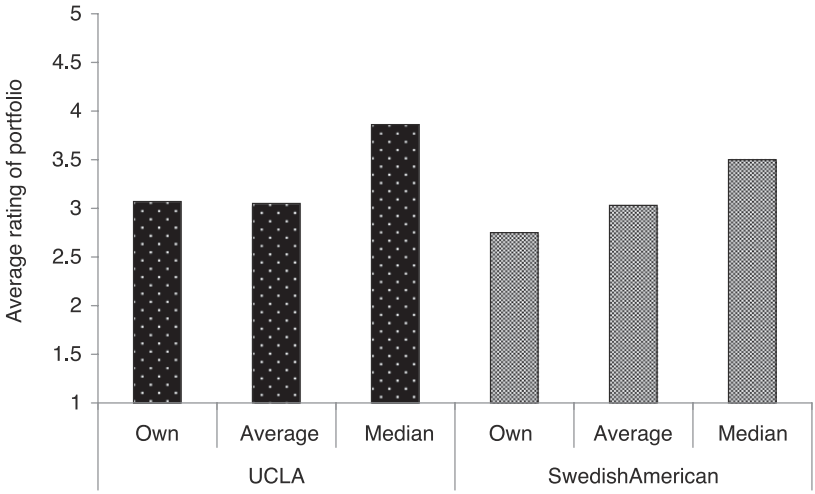

Figure 10.23 The average rating of portfolios at UCLA and SwedishAmerican.
Source: Benartzi and Thaler (2002).

So, should we be letting people choose where they invest? The simple answer given these results would seem to be 'No', but we do need to be careful. First, there is the problem of deciding where best the money should be invested and how. Twenty percent of those in the study were most happy with the portfolio they chose, and can we justify taking choice away from these to benefit the other 80 percent? Second, is the problem that people may like making a choice, even if they make a bad choice? The potential costs of investing unwisely are huge, so people would really have to like making a choice for this to matter much. Even so, it suggests that we might want to think about giving investors some choice.

10.7 Welfare trade-offs

In this final section I want to start talking more explicitly about policy. One fundamental issue in policy-making is to measure the welfare consequences of different policies. After all, the objective is to pursue policies that increase overall welfare, and to do that we need to measure welfare. Here, I shall look at two very different examples of measuring the welfare consequences of policy. There are lots of other examples I could have given, but I think these are two nice examples to illustrate that what we have learnt in the chapter can be useful. The first example shows how we can fill in some of the blanks of the standard economic model. The second example shows how we can open up new issues not considered in the standard model.

10.7.1 The inflation–unemployment trade-off

It is common in macroeconomics to think of a government trying to maximize a social welfare function. Two things usually thrown into the welfare function are inflation and

Table 10.6 Changes in reported life satisfaction (on a 1–4 scale) if there is unemployment or inflation.

	Change in life satisfaction
If unemployed	–0.33
If 1% rise in unemployment rate	–0.02
If 1% fall in inflation	+0.014

Source: Di Tella, MacCulloch and Oswald (2001).

unemployment, with high inflation and unemployment associated with lower social welfare. The textbook story would be that a government wishes to reduce inflation, which will cause unemployment in the short run, and so wants to find the disinflationary path that minimizes the loss in social welfare. For this approach to make sense we need to have some idea of the relative loss in social welfare if there is unemployment or inflation.

A standard approach would be to assume a one-to-one trade-off, so the welfare function would look something like

$$W = -\left(Y - \bar{Y}\right)^2 - \left(\Pi - \bar{\Pi}\right)^2,$$

where $Y - \bar{Y}$ is the deviation of national income from the optimum and $\Pi - \bar{\Pi}$ is the deviation of inflation from the optimum. The reason to assume a one-to-one trade-off seems to be the lack of any better idea! Happiness data can potentially give us a better idea, because we can estimate how much unemployment and inflation decrease reported happiness, and observe the trade-off.

Di Tella, MacCulloch and Oswald (2001) did this, looking at data from 12 European countries over the period 1975 to 1991. Table 10.6 summarizes some of the numbers they come up with. Unemployment causes two negative effects on happiness. There is a direct effect, due to people who are unemployed being on average less happy. There is an additional indirect effect, what we could call 'the fear of unemployment', that lowers the happiness of people still employed. Inflation has only an indirect effect.

With the numbers in Table 10.6 we can calculate the trade-off between inflation and unemployment. A 1 percent increase in the unemployment rate lowers the average satisfaction of everybody by 0.02 and lowers the satisfaction of the 1 percent who become unemployed by 0.33. Overall, therefore, average satisfaction decreases by $0.33 \times 0.01 + 0.02 = 0.0233$. To find the fall in inflation needed to compensate for this decrease in happiness, we calculate $0.0233/0.014 = 1.66$. So, a 1.66 percent fall in inflation increases average satisfaction by as much as a 1 percent increase in unemployment decreases average satisfaction. Thus, the unemployment–inflation trade-off is 1.66, and we might want a social welfare function

$$W = -1.66\left(Y - \bar{Y}\right)^2 - \left(\Pi - \bar{\Pi}\right)^2.$$

This is not the end of the story, because, for example, unemployment and inflation may influence different parts of society differently, and we might want to take account

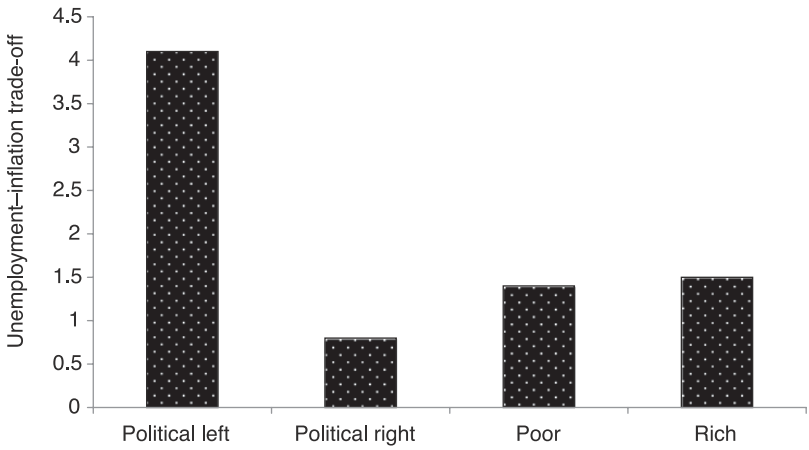

Figure 10.24 The inflation–unemployment trade-off of those on the political left is higher than of those on the political right, and this is not because of differences in income.
Source: Di Tella and MacCulloch (2006).

of that. We also know that being unemployed can have long-term effects on satisfaction, while a change in the inflation rate is unlikely to. The specific number 1.66 should not, therefore, be taken as definitive. But 1.66 does have more justification than the standard assumption of 1. This illustrates how we can use measurements of happiness and life satisfaction to better inform on optimal policy. We have focused on the inflation and unemployment trade-off but a similar idea can be applied to a wide range of other policy questions, such as the optimal level of unemployment benefit and optimal tax rates.

Before moving on I think it is interesting to briefly mention a study by Di Tella and MacCulloch (2005) that asked whether people on the political right had a different inflation–unemployment trade-off from those on the political left. Why would they? Typically we would associate right-wingers with policies that focus on reducing inflation rather than unemployment. By contrast, left-wingers would be associated with policies that focus on reducing unemployment rather than inflation. Figure 10.24 illustrates that those who rated themselves to the political right did have a relatively low trade-off and those to the left a relatively high trade-off. This is not due to them being rich or poor. Recall that these numbers are estimated by looking at how satisfied people say they are. So, these differences do not mean that those on the political right say they care more about inflation than those on the political left. It means that the life satisfaction of those on the political right really does depend more on inflation than it does for those on the political left!

10.7.2 Tax saliency

One other thing that it is important to have a welfare measure of is the deadweight loss caused by taxes. I will finish this chapter with an example of how a seemingly irrelevant change in the framing of a tax can potentially change the associated welfare loss.

In Europe the advertised price of a product includes any relevant taxes. In the United States and Canada it is normal for the price of a product to not include the sales tax. To work out the full price of an item a customer would, therefore, have to calculate and add on the sales tax of 4 to 10 percent. Cognitively, this is not an easy thing to do, even if he carries around a calculator to do the math, because some things are subject to sales tax, some not, etc. A study by Chetty, Looney and Kroft (2009) questioned the consequences of not including the sales tax in the advertised price.

The first part of the study involved an experiment carried out in a large grocery store in California. On some of the products they placed a label, in addition to the main label, giving the price of the good including the tax. This did not change the price of the good but did change the appearance of the price to the customer.

It might be that customers would be more likely to buy such a product because the label allowed them to avoid the cognitive cost of working out the tax. On the other hand, they might be less likely to buy the product because the label made the product appear more expensive. What do you think happened?

By way of comparison, the study looked at the change in demand on goods that did have the amount of sales tax displayed (treated goods in the treatment store) versus the change in demand for similar goods in a similar part of the store (control goods in the treatment store) and those same goods in a different store (control and treated goods in the control store). Figure 10.25 summarizes what happened. The story is fairly clear: demand fell on goods that had the tax included in the price. This fall was also relatively large, at 6 to 8 percent, depending on what was used as the comparator.

To check this result, the second part of the study looked at demand for alcohol. Alcohol is subject to two taxes: an excise tax, which is included in the price, and the sales tax, which is not. This allows for a natural experiment in which we can track changes in the two tax rates and changes in demand. Chetty and co-authors looked at data from 1970 to 2003, and find that changes in the excise tax cause a much larger change in demand than changes in the sales tax. This is consistent with the story that customers focus on the price on the label, which includes excise tax, and largely ignore the sales tax, which is added at the till.



Figure 10.25 Change in demand if sales tax is displayed.
Source: Chetty, Looney and Kroft (2009).

The clear suggestion is that the advertised price is salient. If the price on the label is higher then consumers buy fewer, even though the actual price, accounting for tax, is the same. What are the consequences of this? To answer that question, we can work through a simple model of consumer choice.

Suppose there are two goods, bananas and money. If Ian buys b bananas and has \$$y$ left, his utility is

$$u(b, y) = 2\alpha\sqrt{b} + y,$$

where α is some parameter. Further suppose that he has \$$m$ to start with and the price of a banana is p ignoring the tax and $p(1 + t)$ taking into account the tax.

[Extra] If Ian takes into account the tax, he will maximize his utility subject to the budget constraint $y = m - p(1 + t)b$. The optimum is where the marginal utility from consuming bananas equals the price, and so where

$$\frac{\alpha}{\sqrt{b}} = p(1+t) \quad \text{or} \quad b_T = \left(\frac{\alpha}{p(1+t)}\right)^2 \quad \text{and} \quad y_T = m - \frac{\alpha^2}{p(1+t)}.$$

If Ian does not take into account the sales tax, he will maximize his utility subject to the budget constraint $y = m - pb$. The optimum in this case is where

$$\frac{\alpha}{\sqrt{b}} = p \quad \text{or} \quad b_N = \left(\frac{\alpha}{p}\right)^2 \quad \text{and} \quad y_N = m - (1+t)\frac{\alpha^2}{p}.$$

In both cases, Ian pays the sales tax, as seen when we calculate y. With this we can work out the difference in utility $\Delta = u(b_T, y_T) - u(b_N, y_N)$ whether Ian does or does not take into account the sales tax. It is

$$\Delta = 2\alpha\left(\frac{\alpha}{p(1+t)}\right) + m - \frac{\alpha^2}{p(1+t)} - \left[2\alpha\left(\frac{\alpha}{p}\right) + m - (1+t)\frac{\alpha^2}{p}\right],$$

which we can simplify to

$$\Delta = \frac{\alpha^2}{p(1+t)} - \frac{\alpha^2(1-t)}{p} = \frac{\alpha^2 t^2}{p(1+t)}. \tag{10.1}$$

Whether you followed the algebra or not, the main thing to know is that equation (10.1) tells us the amount Ian will lose if he does not take into account the sales tax. This difference will be small if t is small. For instance, if $t = 0.07375$, which is the relevant figure for California, we get

$$\Delta = 0.0051\frac{\alpha^2}{p}.$$

This is not going to be a big number unless α is big, meaning Ian really likes bananas. For example, if $\alpha = 10$, $m = 200$ and $p = 1$ then, ignoring the sales tax, he will buy 100 bananas and spend \$107.38. If he was to take into account the sales tax, his utility would be \$0.51 higher. That does not look like much of a gain, and so he does not lose much by ignoring the sales tax.

It turns out that this is quite a general result, and not restricted to the functional form for utility that I assumed. Primarily this is because, by maximizing his utility, even if at slightly the wrong price, Ian reaches a point at which he is relatively indifferent to changes in his consumption. So, if he were to take account of the tax, it might change his behavior but would not much change his utility. If ignoring the tax means little welfare loss for the consumer then it is no surprise that we rarely see shoppers with calculators working out the tax on a good.

We can see the implications for the government with the help of Figure 10.26 and Table 10.7. These work through the consequences of a 5 and 10 percent tax. If there is no tax, Ian will buy 100 bananas, have \$100 left and have total utility 300. For a higher tax, we can compare his utility and the tax revenue if he does or does not take account of the tax. Simplest is to measure welfare as utility plus the tax revenue, so it is \$300 with no tax.

If the tax is 5 percent and Ian takes account of it, there is a deadweight welfare loss because some consumer surplus is lost. This is given by area B in Figure 10.26, and comes to \$0.21. If the tax were 10 percent, the deadweight loss would be equal to areas B + C + D, and amount to \$0.79. You might argue that these are not big numbers. If, however, there are 1,000 or 10,000 consumers like this, the overall welfare loss of the tax starts to add up.

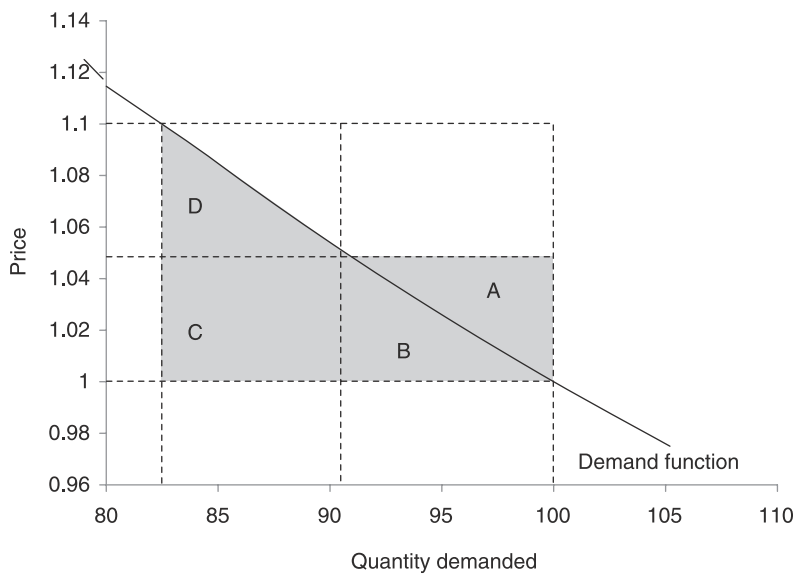

Figure 10.26 The consequences of ignoring the sales tax. If Ian takes account of the tax, there is a deadweight loss of B and D + C + B. If he does not take account of the tax, there is no deadweight loss but A + B is transferred from Ian to the government.

Table 10.7 The consequences of a 5 and 10 percent tax when $\alpha = 10$, $m = 200$ and $p = \$1$.

Tax rate	Is the tax salient?							
	No	Yes	No	Yes	No	Yes	No	Yes
	Bananas bought		Money left		Utility		Tax revenue	
t = 0	100	100	$100.00	$100.00	300	300	$0.00	$0.00
t = 0.05	100	91	$95.00	$104.45	295	295	$5.00	$4.55
t = 0.1	100	83	$90.00	$108.70	290	291	$10.00	$8.30

What if Ian does not take account of the sales tax? He continues to buy 100 bananas. This means that, if the tax is 5 percent, we no longer lose area B. Instead, area A + B gets transferred to the government in tax. We can see this by the fact that consumer utility goes down by five but tax revenue goes up by $5. There is no deadweight loss of the tax, just a redistribution of welfare, in which Ian becomes relatively worse off.

This is important, because it suggests that taxes are not as bad as we might think in terms of efficiency, but are possibly worse in terms of equity. Ultimately, however, it is good for overall welfare that Ian does not take account of the tax. This suggests that care is needed to not push the tax beyond the threshold at which the consumer does take note of it. For example, if Ian ignores a 5 percent tax but takes note of a 10 percent tax, increasing the tax from 5 to 10 percent would cause a big deadweight loss of A + B + C + D.

This example illustrates how it is important to take account of the likely behavior of people in designing policy. You might say that this is obvious, but the example shows that things can be a little subtle. In the next chapter I want to develop a more general framework with which to think through how behavioral economics can effectively inform policy.

10.8 Further reading

The literature on happiness gets ever larger, and so you should not be short of places to find out more. Books include Frey (2008), and survey articles Frey and Stutzer (2002), Di Tella and MacCulloch (2006) and Kahneman and Krueger (2006). See also Easterlin (2013). Two nice papers that did not get a mention when talking about experience utility and projection bias are Kahneman, Wakker and Sarin (1997) and Loewenstein, O'Donoghue and Rabin (2003). The implications for policy are discussed by Loewenstein and Ubel (2008), and a paper by Patall, Cooper and Robinson (2008) provides a comprehensive survey of the psychology literature on whether choice is good or not. I also want to mention a classic paper by Laibson (1997), which looks at savings, commitment and time inconsistency. Finally, a useful source of information for more on how QALYs are measured and used is the website of the National Institute for Health and Care Excellence (www.nice.org.uk). A paper by Dolan and Kahneman (2008) provides a good overview of the issues.

One thing that I did not talk about, but is of interest, is incorporating measures of happiness into measures of economic output. In February 2008 the then French president, Nicolas Sarkozy, asked Joseph Stiglitz to create a commission, subsequently called 'the Commission on the Measurement of Economic Performance and Social Progress', to identify the limits of GDP as an indicator of economic performance and social progress. The commission website (www.stiglitz-sen-fitoussi.fr) has a copy of the committee's report and videos of various speeches made. Unfortunately, not much change has followed the commission, possibly because the financial crisis got in the way, but expect more discussion of such issues in the future.

10.9 Online materials

On the companion website you will find more details on the example of Emma's book, Ian's dissertation and that around tax salience. You will also have access to flashcards, quizzes and answers to the review questions.

10.10 Review questions

10.1 Explain why peak–end evaluation causes duration neglect.
10.2 Can we trust measures of happiness and life satisfaction?
10.3 Should a person choose the thing that maximizes decision utility, total utility or remembered utility?
10.4 If a person has done something before, should we expect that remembered utility equals decision utility?
10.5 Why do you think a high national output does not translate into higher levels of life satisfaction? What are the implications of this for how we measure economic performance?
10.6 If happiness is strongly related to relative income then how should a government think about maximizing total happiness?
10.7 Discuss why current mood can influence the choice a person makes about some event that will not happen until some distant point in the future.
10.8 Consider the distinction between consumption goods, such as a car or fridge, and experience goods, such as a restaurant meal or holiday. For which type of good are the differences between decision utility, total utility and remembered utility likely to be largest?
10.9 Suppose you are taking a course in behavioral economics that involves doing three assignments. Would you rather that the course organizer gives staggered deadlines for handing in the work or that you merely need to hand in the work before the end of the course? Why?
10.10 Discuss the different ways in which pre-commitment can help a person increase their utility.
10.11 How can we take account of adaption when calculating QALYs?
10.12 Given adaptation and projection bias, are patients informed enough to decide what medical treatments are best for them?

10.13 Explain why projection bias is different from present bias. How does this show up in people's decisions to save for retirement?

10.14 Should we take account of adaptation, or the lack of adaptation, to unemployment when calculating the unemployment–inflation trade-off? What effect would this have?

10.15 Should Europe change policy and not require firms to include the sales tax when advertising price?

Chapter 11

Policy and behavior

> A government agency will be set up to paint contour lines on to hills and colour roads the same as on maps. This will help people know where they are.
>
> The Official Monster Raving Loony Party 2010, UK general election manifesto

In this final chapter I want to look at whether behavioral economics can be used by policy-makers to help people make better choices. For instance, Maria might buy gym membership but then never go to the gym. Alan might not save enough for retirement. Can we change things so that they make better choices? This question raises some tricky issues, which I will get to, but it has to be a major objective of behavioral economics. It allows us to apply what we know in trying to make the world a better place where people are a little bit (or a lot) happier.

In looking at how behavioral economics can inform policy, I think it's useful to distinguish two broad categories of economic policy. The first is about designing institutions so that good things happen. The second is about manipulating individual incentives so that good things happen. There is inevitably some overlap between the two. For example, in designing a good institution we need to provide good individual incentives. The distinction is a useful one, however, and so I will look at each category in turn, beginning with the design of institutions.

11.1 Designing good institutions

An economic **institution** is, basically, the means and rules by which people interact. We have already seen plenty of examples of institutions in this book, and seen plenty of examples of why institutions matter. For example, in Chapter 2 we saw that whether the market is a double-auction, posted offer or negotiated price format will likely matter. In Chapter 6 we saw that whether the auction is an English, Dutch or first- or second-price format might matter. We also saw that whether there is leadership, communication and things such as rebates or refunds can help people coordinate in providing, say, a threshold public good. Finally, in Chapter 7 we saw that the threat of punishment can help sustain cooperation.

What I want to do here is look at three very different examples of institution design. As well as further illustrating the importance of institutions, this will take us into the world of practical policy design and intervention. I'll briefly mention that institutions can be of two basic kinds: **formal institutions**, such as markets, and

informal institutions, such as social norms. The first example I want to look at is more about informal institutions and the latter two more about formal institutions.

11.1.1 The tragedy of the commons

To get us started, consider a lake used for fishing. Anyone who wants to go fishing is free to do so, and can fish as much as they want. What's going to happen? What might happen is **overuse** or **overexploitation** of the resource, whereby too many fishermen come along, catch too many fish and leave the stock of fish too low to be sustainable. We don't want that to happen.

A fishing lake is one example of a common pool resource. A **common pool resource** is any good that it is difficult to exclude people from consuming, but one person's consumption lowers the possible consumption of others. It might, for instance, be hard to stop people fishing in the lake, and one fisherman's catch cannot be another's. Other examples include forests, water systems for drinking or irrigation, and the global atmosphere. As this list suggests, common pool resources are immensely important in thinking about environmental issues.

The **tragedy of the commons** is that we should expect overuse of common pool resources. That is, we should expect the worst. To see why, we can look at a specific version of a common pool resource game, used in some experiments.

Suppose that there are eight people, including Anna, who each have ten tokens. They can invest their tokens in something that yields a 'safe return' or spend it on extracting goods from a common pool resource. I will think of people deciding whether to harvest food in their garden, for a safe return, or go fishing in the local lake, where the return will depend on what others do. Tokens could represent money, time or a combination of both, but I will think of time here. Table 11.1 summarizes some of the possible payoff combinations. For example, if Anna spends four hours fishing and the others spend 14 hours fishing, she gets a payoff of 104.

To make some sense of the numbers in Table 11.1, note that Anna gets 50 for sure if she spends all her time in the garden, which is why it's a safe return. If the others do relatively little fishing then her payoff will be higher if she spends time fishing, because the lake is full of fish. By contrast, if the others do rather a lot of fishing then her payoff will be lower if she spends time fishing, because the lake is already empty of fish.

A good outcome is that each person spends four hours fishing, to get a payoff of 90. If the others are going to fish a total of 28 hours, however, we can see that Anna does best to spend ten hours fishing. It looks unlikely, therefore, that people would stick to fishing for only four hours. More plausible is that they each spend eight hours fishing to get a payoff of 66. Why? If the others are going to fish a total of $7 \times 8 = 56$ hours, Anna does best to spend eight hours fishing, and so everything fits a bit better. This is all captured by the fact that 'everyone fish four hours' is the Pareto optimum, while 'everyone fish eight hours' is the Nash equilibrium.

If everyone did fish for eight hours then we would observe a version of the tragedy of the commons. People would be fishing too much and getting a lower payoff than they could do by fishing a bit less. This would not be good, but what happens in reality?

The evidence suggests that the tragedy of the commons can be avoided, sometimes. Indeed, Elinor Ostrom won the Nobel Prize in Economics in 2009 because she '[c]hallenged the conventional wisdom by demonstrating how local property can be successfully managed by local commons without any regulation by central authorities or privatization'. Examples of successful use of common pool resources include grazing land in the Swiss Alps and Japan, irrigation systems in Spain and the Philippines, and use of the groundwater basins around Los Angeles. In each of these cases, users of the resource found some way to maintain use of the resource at the efficient level, or, at least, below the level of overuse.

To give a specific example: in 1995 Alaska's halibut fisherman decided to use something called **individual fishing quotas**. The basic idea is that individuals are given ownership over a share of the total amount that can be fished. The change was dramatic. The length of season went from a dangerous, three-day race to get as many fish as possible to an eight-month, demand-led, sustainable solution.

For every success story, though, there is an example of the news not being so good. For instance, the bluefin tuna has reached dangerously low levels in the Mediterranean, but there is little consensus on how the decline can be halted, especially given that attempts to use individual fishing quotas in international waters have largely failed.

The pertinent question, therefore, is why the tragedy of the commons can be averted in some cases and not others. Ostrom addressed this question by looking at case studies of common pool resources with the objective of seeing what distinguished the successes from the failures. An up-to-date list of the differences is given in Ostrom (2010) and includes things such as: the existence of clear and locally understood boundaries between users and non-users; users of the resource having some say in how the resource is managed; and rules for using the resource designed to include punishment for excessive use.

Such a list, providing as it does a neat dividing line between likely success and failure, is incredibly useful. For instance, it can help inform a policy-maker when intervention is needed to avoid the tragedy of the commons, and when it is not. It can also help policy-makers set things up so that users of the resource can help themselves, such as with individual fishing quotas. Clearly, however, this works only if the list is the right one – i.e. we have the right dividing line between likely success and failure. And case studies necessarily involve some worrying comparisons, such as Alaskan and Spanish fisherman. One way to supplement case studies, and make sure we have things right, is to use lab experiments.

Lab experiments are useful because they allow specific elements to be varied one at a time. To illustrate, I will look at a study by Ostrom, Walker and Gardner (1992). Subjects were involved in a common resource problem, like that in Table 11.1, for 20 rounds. In the first ten rounds, subjects independently had to decide how much to spend extracting goods from the resource, and in most cases there was overuse. Subjects were then given various means to try and avoid overuse in the ten subsequent rounds. Figure 11.1 summarizes the results. It looks a bit messy, so let's go through it slowly.

In the baseline case, subjects were not given any means to avoid overuse, and profits remain well short of those possible. In some treatments, subjects were given the chance to communicate with each other either though a one-off ten-minute

Table 11.1 Some of the possible payoff combinations in a common pool resource game involving Anna and seven others.

Hours spent in extracting goods from the common pool resource						
Anna	**The other seven**					
	0	**14**	**28**	**42**	**56**	**70**
0	50	50	50	50	50	50
2	85	78	71	64	57	50
4	118	104	90	76	62	48
6	149	128	107	86	65	44
8	178	150	122	94	66	38
10	205	170	135	100	65	30

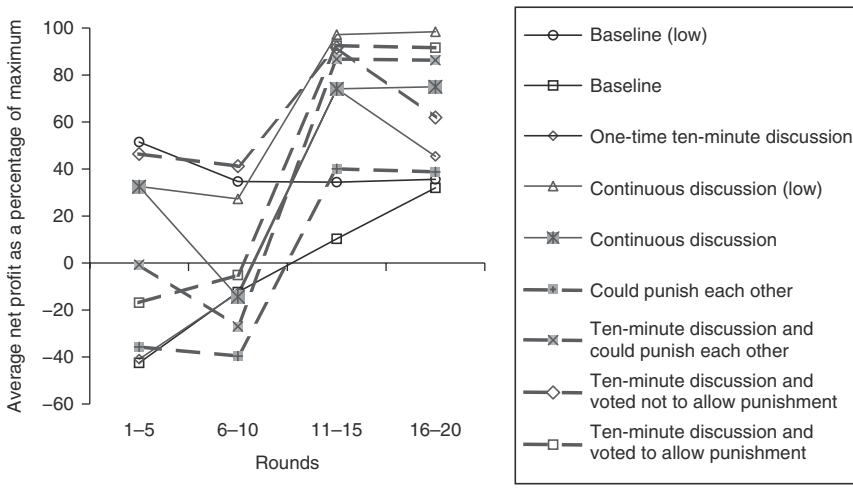

Figure 11.1 The average net profit as a percentage of the maximum over 20 rounds of a common resource problem. A negative percentage indicates that subjects would have done better to invest in the asset with a safe return. 100 percent is the Pareto optimum.
Source: Ostrom, Walker and Gardner (1992).

discussion or through repeated communication before every round (see Research methods 11.1); discussion seemed to have a positive impact, and repeated discussion led to a sustained increase in profits.

In another treatment, subjects were given the chance to punish others (in a similar way to that in the public good game of Chapter 7); this seemed to have less of an impact. A further treatment combined the chance to communicate with the opportunity to punish others; this did appear to have a positive impact. In a final treatment, subjects voted, after discussion, on whether or not to allow punishment. Whatever the outcome of the vote we again see a positive impact.

❖ RESEARCH METHODS 11.1 COMMUNICATION

As a general rule, experimental economists seem averse to letting subjects communicate with each other. There are good theoretical reasons for doing this, because economic activity often takes place between people who simply cannot communicate with each other. There are also good practical reasons for doing this, because it stops the experimenter losing control; for example, one subject with a 'strong' personality could end up biasing a whole experiment. In theory, communication should not matter anyway – at least, according to the standard economic model!

People do often communicate, though, and intuitively it makes a difference, so we do need to take account of it. There are various different ways that subjects can be allowed to communicate in experiments.

The most restrictive way is to let a subject say what action he plans to take. So, before a subject chooses, say, how much to extract from a common pool resource he gets the chance to say to others what he plans to choose. This can be done in a way such that the experimenter loses no control of the experiment, and so is the preferred way of many.

A step up from this is to let subjects chat via computer (see Research methods 7.4). Now the experimenter starts to lose control, because a subject is free to say anything. But, at least all that is said is recorded and relatively easy to analyze.

The final step, as in the experiment by Ostrom, Walker and Gardner (1992), is to let subjects talk to each other face to face. Now the experimenter loses a lot of control, and trying to record what happens is essentially impossible; we might be able to record what was said, but recording body language and the like is clearly very difficult. The latter type of experiment is very rare in economics.

The lesson from this study was that repeated communication in combination with some threat of punishment or sanction helped people avoid the tragedy of the commons. Other studies have shown that the benefits of communication are less if there are 'outsiders' who do not take part in the communication. This is reassuringly similar to what we see in case studies. A combination of case studies and experiments can therefore give us important insights into when and how the tragedy of the commons can be avoided. The crucial things necessary to avoid the tragedy of the commons seem to be communication, a threat of sanctions, an ability to exclude outsiders and a sense of ownership.

11.1.2 Matching markets

In Chapter 2, when we looked at markets, we contrasted double-auction, posted offer and negotiated price institutions. The basic objective of these institutions is to

match buyers and sellers so that they can do mutually beneficial deals. Recall that in a double-auction market there is an auctioneer who accepts bids and asks, and matches a seller to a buyer if the bid is higher than the ask. In a posted offer market, sellers advertise their prices and buyers walk around and find a price they like. In a negotiated price market, sellers and buyers walk around and ask each other what they are willing to sell or buy. What we hope to see is buyers being matched to sellers whenever the seller is willing to supply at a price a buyer is willing to buy.

In Chapter 2 I suggested that double-auction markets are, typically, efficient. What that basically means is that we get the best match of buyers and sellers. Great! Our focus, however, was on a market where each seller had the same thing to sell; the goods were **homogeneous**. Often this seems appropriate; for instance, one share in a company, or one can of a particular brand of cola, is as good as any other, so a buyer should be relatively indifferent as to who they buy from. In many other cases, however, goods are not homogeneous. For example, no two houses, restaurant meals or used cars are exactly alike. This makes it much more difficult to match buyers and sellers efficiently.

To illustrate the problem, we can look at the problem of matching workers to employers. In many professions, newly trained graduates simultaneously try to find entry-level jobs with employers. What we hope to see is the best match between the worker or supplier of labor and the employer or demander of labor. Workers will have different preferences over where they would rather work, however, and employers will have different preferences over who they would rather hire. It is very easy for this to become a bit of a mess, with great candidates getting no offers and great employers finding that no one accepts their offers. Obtaining the best match is far from easy. One profession that has tried hard to tackle this problem is the medical profession.

The problem in the medical profession is to match newly trained doctors with hospitals willing to employ them. To demonstrate the problems there can be, we can look at the experience of the United States. Before 1945 the market for new doctors was decentralized, like a negotiated price market. The outcome was an **unraveling of contract dates**, in which the best students were being hired earlier and earlier as hospitals tried to get the best candidates before anyone else did. In the end, students were being hired two years before graduation. This meant that hospitals were hiring students before they had a chance to see how good they really were, or students had a chance to see what type of medicine they would most want to practice. This is inefficient.

In 1945 medical schools banded together to try to improve matters, but a new problem arose. This time candidates who had offers from one place would wait to see if they got an offer at a preferred place. This might sound reasonable but, if everyone is doing it, then everyone is waiting for everyone else to make a decision. Nothing happens until the deadline for acceptance, and then there is a last-minute rush and decisions are being made with little time to think. This is also inefficient.

In 1952 the National Resident Matching Program was set up as a central clearinghouse for applications. A way had to be found to match doctors with hospitals that would avoid the previous problems. Since 1998 the program has used a **matching algorithm** designed by economists, notably Alvin Roth, and the process is a lot more efficient. Let's look first at the algorithm used.

Table 11.2 An example of possible preferences of junior doctors looking for work and hospitals looking for junior doctors.

		Preference		
		First	**Second**	**Third**
Doctors	Bill	Birmingham	New York	Boston
	Carol	Birmingham	Boston	New York
	Emma	Boston	New York	Birmingham
Hospitals	Birmingham	Bill	Carol	Emma
	New York	Bill	Carol	Emma
	Boston	Bill	Carol	Emma

Table 11.3 The deferred acceptance algorithm allocating doctors to hospitals for the example.

	Doctors assigned			Doctors rejected
	Birmingham	**New York**	**Boston**	
Round 1	Bill and Carol	—	Emma	Carol
Round 2	Bill	—	Emma and Carol	Emma
Round 3	Bill	Emma	Carol	—

After a process of interviews and visits, doctors submit a ranking of their preferred hospitals, and hospitals submit a ranking of their preferred doctors. Table 11.2 gives an example to work through. Something like a **deferred acceptance algorithm** is then used. The algorithm is as follows: each doctor is assigned to his or her first choice of hospital. Table 11.3 keeps track of who is assigned where in the example. The posts at each hospital are then filled with the most preferred doctors assigned to them, and other doctors are rejected. In the example, Carol is rejected because Birmingham prefers to hire Bill. Any doctor rejected at this stage is assigned to his or her second choice of hospital. The posts of each hospital are then refilled with the most preferred doctors assigned to them, and other doctors rejected. This time it is Emma who is rejected, because Boston prefers Carol. This process goes on until no doctors are rejected.

Things are a bit more complicated than this because, for example, there might be couples who want to work in a similar location. The basic idea remains the same, however. The crucial thing about the deferred acceptance algorithm is that it leads to a **stable matching**, in the sense that no candidate or hospital can look back and wish they had done a different ranking.

Not all possible algorithms are stable. For instance, consider a **priority matching algorithm** that works as follows: the priority of a match is measured by the product of the two rankings. For example, if a doctor ranks a hospital first and the hospital ranks the doctor third, the product is $1 \times 3 = 3$; if a doctor ranks a hospital fourth and the hospital ranks the doctor second, the priority is $4 \times 2 = 8$. Posts are filled in order of priority. Thus, priority one matches are filled first, when both doctor and hospital ranked each other first, then priority two matches are filled, and

Table 11.4 Different methods used in the United Kingdom.

Regional market	Algorithm	Stable	Unraveled
Edinburgh	Deferred acceptance	Yes	No
Cardiff	Deferred acceptance	Yes	No
Birmingham	Priority matching	No	Yes
Newcastle	Priority matching	No	Yes
Sheffield	Priority matching	No	Yes
Cambridge	Linear programming	No	No
London Hospital	Linear programming	No	No

Source: Kagel and Roth (2000).

so on, until all vacancies are filled. To see why this algorithm need not be stable, we can look back at the example. The matching between Emma and Boston has priority three while the matching between Carol and Boston has priority four. This means that Emma will end up at Boston and Carol at New York. Once they know this, Boston can ring Carol and ask her to come, and she will accept. The assignment begins to unravel.

If the outcome of an algorithm is likely to not be stable then we can expect unraveling of contract times, because applicants and hospitals such as Carol and Boston will be reluctant to trust the algorithm. The United Kingdom provides a nice field experiment, because different regions have over time used different algorithms. Table 11.4 summarizes what happened. I will not explain what the linear programming algorithm is, but you can see that it provides something of a puzzle, because it is, in principle, unstable yet did not lead to an unraveling of contract times.

One way to try to answer this puzzle is to replicate each matching algorithm in the experimental laboratory and see what happens. That is what Kagel and Roth (2000) and Unver (2005) did. The experiments they ran lasted for 25 rounds. In each round, subjects, playing the role of worker and employer, tried to find a good match and were paid between $14 and $4 depending on who they matched with. To capture the possibility of unraveling, each round was split into three periods, called −2, −1 and 0. If a worker and employer agreed a match in period −2 they each had to pay a fine of $2, and if they agreed in period −1 they had to pay a fine of $1. What we want to see is whether things unraveled, in the sense that subjects agreed matches in periods −2 and −1, despite the cost that this entailed.

Figure 11.2 summarizes what happened in the experiments. Start by focusing on the first ten rounds, when matching was decentralized and so there was no matching algorithm used. We see that subjects paid the fine to get a good match early, and so there was an unraveling of contract time. From round 11 onwards the matching algorithm was used. It was used in period 0 for those workers and employers who had not matched so far in periods −2 and −1. We see that the fine subjects were willing to pay fell quite dramatically, and became almost zero with the deferred acceptance algorithm.

This shows that the deferred acceptance algorithm is working well. The priority and linear programming algorithms work less well, in that the costs remain significant. The suggestion would be, therefore, that it is something of an accident that the linear programming algorithm worked for Cambridge and London Hospital.

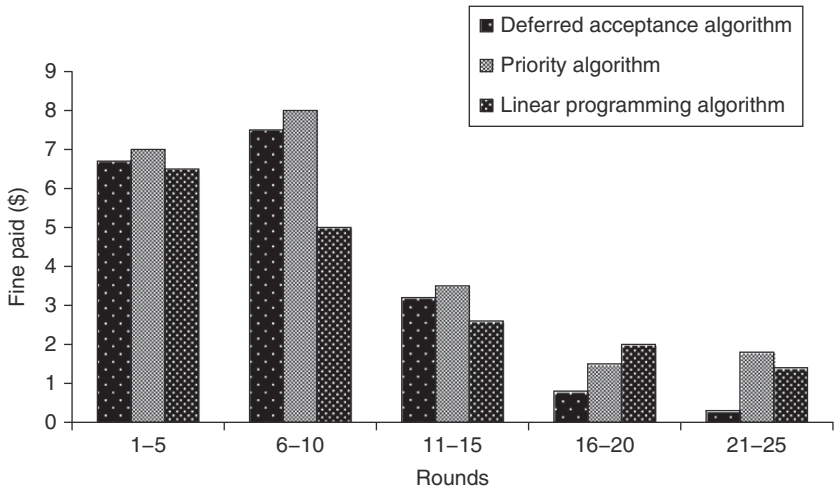

Figure 11.2 Matching algorithms compared.
Source: Unver (2005).

The main thing, however, is that these experiments confirm the advantages of the deferred acceptance algorithm. The deferred acceptance algorithm looks as if it does a good job both in theory and in the experimental laboratory. This has translated into success in the real world. The algorithm has proved successful in matching doctors to hospitals and is now being used in other areas as well, such as matching prospective students with schools. (The biggest mystery is why economists have not used it in their own profession to match junior faculty to departments!) In 2012 Alvin Roth won the Nobel Prize in Economics 'for the theory of stable allocations and the practice of market design'.

11.1.3 Spectrum auctions

The final example I want to look at is auction design. In particular, I want to look at the so-called **spectrum auctions**. Before getting into the economics of these auctions, we need to know a bit about what is up for sale.

The electromagnetic spectrum is the name given to the range of electromagnetic frequencies that exist, from gamma rays, with a relatively high frequency, to long waves, with a relatively low frequency. Towards the low-frequency end of the spectrum are radio waves with a frequency of around 9KHz to 300GHz. Television, radio, mobile phone signals, wireless networking, emergency communications, the signal for the door to open when someone presses a car key fob, and lots of other things are all transmitted via radio waves.

This creates a competition for a portion of the electromagnetic spectrum. For example, if a television channel is transmitting its programs on a certain frequency, it does not want anyone else transmitting on that same frequency, causing interference. Somehow, therefore, the spectrum needs to be split and allocated in such a way that television, radio and everything else can work.

To illustrate, the frequencies between 470MHz and 862MHz are typically set aside for television. These are then split into 49 chunks, or frequency channels, of 8MHz each, giving the frequency channels 21 to 69. A frequency channel is enough space to broadcast a television channel, or six to eight channels using digital television. Potential broadcasters who are allocated a frequency channel can broadcast without fear of interference. But how to decide which potential broadcaster gets allocated a frequency channel?

In many instances, governments directly allocate portions of the spectrum to companies. For example, a government committee might decide who gets what after companies have applied for, or tendered for, the right to use a portion of the spectrum. An alternative, which has become increasingly common, is to auction portions of the spectrum. We saw in Chapters 2 and 6 how auctions can work, and the basic rationale behind auctioning spectrum is that the market, rather than the government, decides who gets what. This could be more efficient, and raise more revenue for the government.

Auctioning spectrum is quite complicated, however, and much more complicated than the auctions I looked at in Chapter 6. That's why designing the spectrum auctions is no simple matter. To get some intuition for why things are so complicated, we can look at something called a chopstick auction.

I will start with a very simple version of the **chopstick auction**. Imagine that someone called Mathan goes to his favorite Chinese restaurant one evening and finds that the restaurant owners have come up with a game. At the start of the evening they are going to simultaneously auction ten sets of chopsticks. They will do so using a first-price sealed bid auction. So, Mathan needs to write down how much he bids for the first set of chopsticks, how much for the second set, and so on. If he bids the highest amount for a set of chopsticks, he wins that set. If Mathan wins a set of chopsticks then he can enjoy his meal, but if he does not he will struggle to eat his meal. He wants, therefore, to win one – and only one – set of chopsticks, and does not care which one he wins.

The bidding for each set of chopsticks in isolation is very similar to the auctions we looked at in Chapter 6. The complication here is that there are ten sets being sold simultaneously. For this reason, it is called a **multi-unit auction**. How much should Mathan bid for each set?

This turns out to be a difficult question to answer. To see why, first imagine that there are ten diners in the restaurant. In principle, this looks promising, because there are ten sets of chopsticks available, and ten diners wanting chopsticks. If Mathan bids $1 for set one and $0 for the others, Emma bids $1 for set two and $0 for the others, and so on, they will all get a set of chopsticks for $1. Somehow, though, the diners need to coordinate. How can Mathan know that he is supposed to bid $1 for set one and not sets two or three? If the restaurant owners do not allow the diners to communicate with each other, there is no way he could know, and so we will likely get coordination failure. This makes it hard to know how Mathan should bid. If there are 20 diners after the ten sets of chopsticks, or some chopsticks are more valuable than others, his task becomes ever more complicated.

We find these complications in auctioning spectrum. To demonstrate, I will look at the spectrum auctions that have received most attention, namely those used to allocate spectrum to mobile phone companies. A mobile phone company

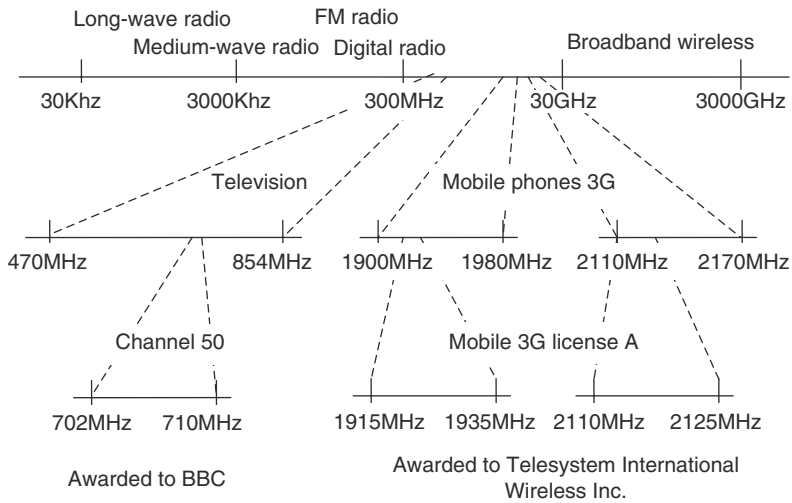

Figure 11.3 The radio wave portion of the electromagnetic spectrum and how it is allocated in the United Kingdom.

needs a portion of spectrum in order to operate. For instance, around the year 2000, with the introduction of 3G or third-generation mobile standards, phone companies needed a new portion of spectrum to start running these services. In most Western countries governments ran auctions to allocate it. Since then auctions have become the routine way to allocate mobile phone licenses across the world.

To be specific, I will look in more detail at the UK 3G auction. The UK government decided to offer five licenses, which you can think of as five portions of the spectrum. Figure 11.3 illustrates that license A consisted of a total of 35MHz. By comparison, license B consisted of a total of 30MHz, and the other three licenses a total of 25MHz. It was decided that license A would be available only to a new entrant in the UK mobile phone market (which excluded BT, One2One, Orange and Vodafone). How does this compare to the chopstick auction?

There were five licenses or five sets of chopsticks up for sale. A company would want one and only one license, but the more spectrum, the better, so license A is better than license B, which is better than licenses C to E. Overall, there were 13 companies that bid, and so demand clearly exceeded supply, but – remember – four of these could not bid for license A. It is very hard to know how a company should bid in such circumstances. This makes it far from obvious whether auctions are a better way to allocate the spectrum than a government committee. In order to work well, auctions need to be well designed.

That is where economists can, hopefully, be of some use. The UK government employed a team of economists including Ken Binmore and Paul Klemperer to advise on the design of the UK 3G spectrum auction. The design eventually used was a simultaneous ascending auction. I will go into the details of what that entails when I look at the US example below. What I want to highlight first is that the auction was a big success. The total raised from the five licenses was £22.5 billion, which was way in

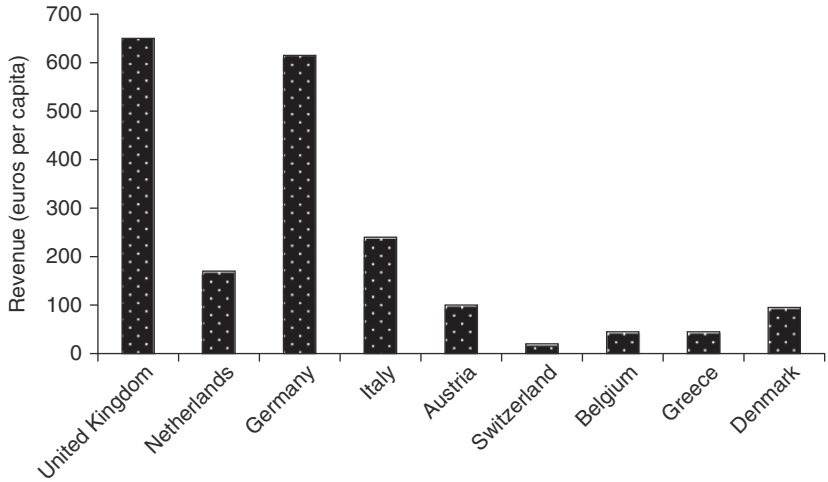

Figure 11.4 The revenue in euros per capita from the 3G spectrum auctions in different European countries in date order. The UK, German and Danish auctions were considered a success.
Source: Klemperer (2004).

excess of what many had expected. It was also far more than was raised from auctions in other European countries, as Figure 11.4 shows.

Figure 11.4 illustrates quite starkly how the success of an auction can depend a lot on the design of that auction. The UK, German and Danish auctions were considered a success, others not so. The Netherlands, for example, used the same auction format as in the United Kingdom but had five incumbents compared to the United Kingdom's four. With five incumbents bidding for five licenses, other potential bidders stayed away, and much less revenue was raised than expected. A similar thing happened in Switzerland. Things can, therefore, go badly wrong. An auction in 1990 in New Zealand of television spectrum is particularly notorious. A sealed bid second-price auction was used, and only NZ$36 million was raised, compared to a predicted NZ$250 million. One lucky company bid NZ$100,000 but had to pay only the second highest bid, of NZ$6, and another bid NZ$7 million but had to pay only NZ$5,000!

If things can go so wrong then it is clearly crucial to design the auction well. The most successful auctions have generally been simultaneous ascending auctions. The first significant use of an ascending auction was by the Federal Communications Commission (FCC) in the United States in 1994, based on the ideas of economists Paul Milgrom, Bob Wilson and Preston McAfee. The basic idea behind the **simultaneous ascending auction** is that there are a number of rounds in which bidders make sealed bids for any items, or portions of spectrum, that they want to bid on. At the end of each round the **standing high bid** for each item is posted as well as the minimum bid for the item in the next round. In order to give bidders an incentive to bid early, an **activity rule** means that bidders can bid only if they have bid in earlier rounds. Table 11.5 illustrates how the bidding progressed in the UK 3G spectrum auction.

Table 11.5 Bidding over the rounds in the UK 3G spectrum auction. The standing high bid for each round is given in millions of pounds.

Round	License				
	A	**B**	**C**	**D**	**E**
1	TIW	Orange	—	Crescent	Epsilon
	£170.0 m	£107.4 m		£89.3 m	£89.3 m
2	NTL Mobile	Telefónica	WorldCom	Crescent	Epsilon
	£178.5 m	£113.1 m	£90.2 m	£89.3 m	£89.3 m
3	NTL Mobile	TIW	WorldCom	3GUK	One2One
	£178.5 m	£176.5 m	£90.2 m	£93.8 m	£93.8 m
4	NTL Mobile	Orange	Telefónica	3GUK	One2One
	£178.5 m	£186.7 m	£95.3 m	£93.8 m	£93.8 m
5	Spectrumco	TIW	Telefónica	WorldCom	Epsilon
	£187.5 m	£196.1 m	£95.3 m	£99.2 m	£100.1 m
150	TIW	Vodafone	BT3G	One2One	Orange
	£4,384.7 m	£5,964.0 m	£4,030.1 m	£4,003.6 m	£4,095.0 m

Source: Ofcom.

The main benefit of a simultaneous ascending auction is that bidders can coordinate over time by seeing how bids progress. This was particularly important in the auctions run in the United States, and those in Germany. To see why, think back to the chopstick auction and imagine that, instead of auctioning ten sets of chopsticks, the restaurant owner auctions 20 chopsticks. Now Mathan wants to win two and only two chopsticks, and his task seems even harder than before. Having the auction run over successive rounds can allow bidders to coordinate a little better. Mathan can be surer of getting the two chopsticks he wants, rather than ending up paying for one that he does not want, or three, which is excessive.

Splitting things up like this, and selling chopsticks rather than sets of chopsticks, might seem like an unnecessary complication. It can, make sense in spectrum auctions, however, because it lets the market decide how many firms there will be. For example, in Germany the 3G spectrum auctions involved 12 portions of spectrum (compared to the five in the United Kingdom) that, in blocks of two or three, were enough to operate a viable network. The number of firms in the market could therefore have ended up as four, five or six. Designing auctions in which bidders may want to win multiple items does, however, pose a new set of challenges. This has led to the use of **package** or **combinatorial auctions**, in which bidders can bid on combinations of items. For example, Mathan might be able to bid $1 for chopstick one, $5 for chopsticks one and two, $3 for chopsticks four, five and six, and so on.

11.1.4 Behavioral economics and institution design

Hopefully, you have seen that institutions matter, and designing good institutions is not easy. For example, avoiding the tragedy of the commons, matching doctors to hospitals or designing an auction that will work efficiently when there are multiple items to sell are all tricky. The tragedy of the commons is avoided in some situations and not in others. Some matching markets and auctions have gone badly wrong and

others have been very successful. The question I want to finish with is: what does institution design have to do with behavioral economics?

At first glance one might say, 'Not too much.' Auction design, for example, draws heavily on game theory and the economics of industrial organization. We would not need behavioral economics to predict the poor performance of the 3G spectrum auctions in the Netherlands and Switzerland; a bit of industrial organization theory would be enough. Look a little deeper, however, and behavioral economics does seem a vital tool in institution design.

This is easiest to see in designing institutions to solve things such as the tragedy of the commons. There, an understanding of social preferences and the ability to coordinate are crucial to understanding what will work and what will not. The case for behavioral economics is weaker in the design of auctions or matching markets, which are more like a combinatorial mathematical problem. Even here, though, behavioral economics can prove crucial. One reason is that the practicalities of institution design far exceed our theoretical ability; accordingly, laboratory experiments can prove useful alongside theory to give an idea of how different institutions will perform. We saw this in looking at designing matching markets. Both the teams designing the UK 3G spectrum auction and the first FCC auctions in the United States also used laboratory experiments to test possible auction designs.

Behavioral economics therefore is, and has proved to be, a useful tool in institutional design, both in informing how people can realistically be expected to behave and as an important test bed for possible institutions. In many ways this seems inevitable, because institution design forces us to step outside the comfort of the standard economic model. It forces us to recognize that it does not always make sense to assume people are like *Homo economicus*. Behavioral economics gives us the tools and techniques to step with confidence beyond the standard economic model.

11.2 Nudge and behavior change

The second category of government policy I distinguished earlier was that of manipulating individual incentives. In recent years there has been much excitement that behavioral economics can change the way in which policy is viewed when it comes to individual incentives. To understand why there is such excitement, it is worth pointing out first something obvious: a policy focus on individual behavior is nothing new. For example, why would central banks or governments increase the interest rate? The answer might be that they want to reduce inflation, but for this to work they need individuals to respond to the change in interest rate. For instance, they need Mathan, who has a mortgage, to reduce his consumption because he knows he will have to pay more interest on his mortgage debt. To give a second example, why would governments give tax relief on saving for retirement? Presumably, they want to increase saving for retirement. Again, for this to work, they need people such as Mathan to change his consumption plans and save more because the tax relief makes saving look a better deal.

So, how can behavioral economics change the way we look at policy? The traditional role for economics and policy has been about changing **incentives for *Homo economicus***. The focus has been on incentives that would matter to someone who maximizes utility without any mistakes or biases. A change in the interest rate or

changes in tax relief are like this. Behavioral economics suggest that **incentives for *Homo sapiens*** can also matter. Account should be taken of incentives that matter to someone who is biased in a predictable way. It seems best to illustrate with some examples. I will look at two, related to saving and time inconsistency.

11.2.1 Savings accounts

If people are time-inconsistent when it comes to saving decisions then properly tailored saving plans may be able to influence behavior. A number of field experiments have now shown this to be the case. I shall illustrate with a study by Ashraf, Karlan and Yin (2006).

The study was carried out in partnership with Green Bank of Caraga, a rural bank in the Philippines. In the first part of the study, a household survey was done on 1,777 existing or former clients of the bank. One part of this survey involved hypothetical time preference questions (similar to those we looked at in section 4.1.2). This allowed respondents to be classified as time-consistent, time-inconsistent or more patient now (meaning their discount factor decreased further into the future). In the second part of the study, half of those surveyed were offered a new SEED (Save, Earn, Enjoy Deposits) account. Of the remaining people, half received no further contact (a control group) and half were encouraged to save more using existing accounts (a marketing group).

The SEED account was specifically designed so that a client had to specify a goal. This could be date-based, such as to save for a birthday, or amount-based, such as to save for a new roof. The client had complete flexibility in setting the goal, but having set the goal they had no access to their money until the goal was reached. This meant very restricted access to funds. In addition, the interest rate of a SEED account was no higher than that of a normal savings account.

For a person who is time-consistent, the SEED account looks like a bad deal, because she has no access to her money and gets no higher interest rate. For someone who is time-inconsistent, the SEED account may look attractive, because it is a strong commitment device to save and overcome the short-term impulse to spend.

In all, 202 SEED accounts were opened, of which 147 were date-based. After 12 months 116 of the accounts reached maturity or met the goal, and all bar one opted to take out a new SEED account. Clearly, the SEED account proved relatively popular. The most fascinating thing is that a reliable predictor of who would take up a SEED account was time-inconsistent preferences. Women who exhibited time-inconsistent preferences were 15.8 percent more likely to take up the SEED product than those with time-consistent preferences. For men, the figure was 4.6 percent. This is consistent with the idea that the account would appeal to those who needed to commit to saving because of present bias.

The final thing we need to check is whether people saved more if they opened a SEED account. The simple answer seems to be 'Yes'. After 12 months savings for those offered the SEED account were estimated to be 82 percent higher than those in the control group, and also much higher than those in the marketing group. This increase in saving seems to be because the SEED product offers consumers incentives to save. In particular, those who took up a SEED account saved a lot more than anyone else. This is clearly apparent in Figure 11.5.

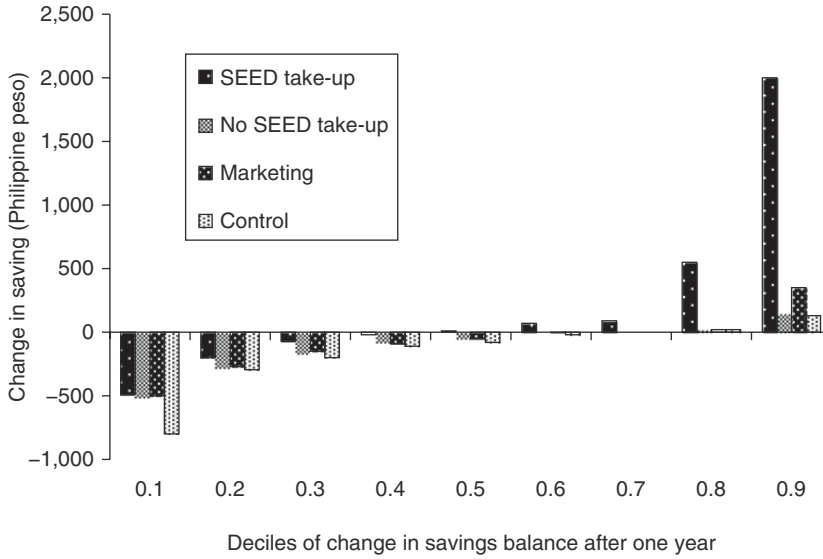

Figure 11.5 Changes in saving behavior after one year. The majority of those taking up the SEED account increased saving. For many, the increase in saving was very large.
Source: Ashraf, Karlan and Yin (2006).

The SEED account thus seemed to be popular and successful. It provided no incentives that would tempt *Homo economicus* because there was no higher interest rate, or anything that should have appealed to someone who was unbiased. The account did provide incentives for *Homo sapiens* to save, however, and seemingly did appeal to those with present bias who wanted a means to pre-commit to more saving.

11.2.2 A default to save

Accounts that restrict access to funds may help some overcome the urge to spend. They should work well for the consumer who is time-inconsistent but sophisticated and so looking for some way to commit for the long term. What about the consumer who is time-inconsistent and naïve? Such a consumer would never choose a restrictive savings account, so another solution is needed. One option is to automatically enroll him in a savings plan and thus put the emphasis on him to opt out of saving rather than opt in. To see how this could work, we can look at a study by Madrian and Shea (2001).

The study looks at data from a large US company that changed enrollment and eligibility for its 401(k) retirement savings plan on April 1, 1998. Before this date employees were eligible to enroll in the plan only after one year of employment, and had to choose to enroll. After the change all employees were immediately eligible to enroll, and new employees were automatically enrolled unless they chose to stay out. Figure 11.6 shows how dramatic the changes proved. The participation rate of employees hired after the change was above that of any subset of current employees.

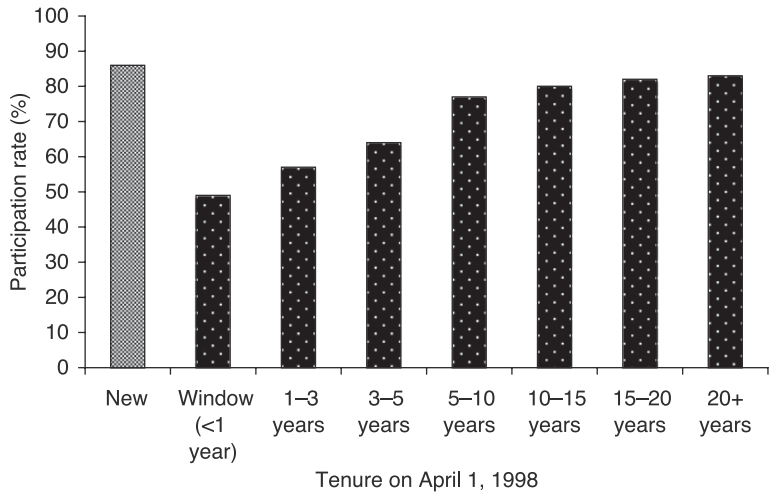

Figure 11.6 Participation rate in a 401(k) retirement savings plan. There was a dramatic increase in the participation rate once the default was to be enrolled in the savings plan.
Source: Madrian and Shea (2001).

The most telling comparison is between new employees and those employed less than a year before the change (the window group). The only substantive difference between these two groups is that enrollment was the default for one and non-enrollment the default for the other.

Clearly, making enrollment the default option increased participation, and we might consider this a good thing. There is, however, one problem. If employees are automatically enrolled then their investment choices also need to be set automatically. In this particular instance there were two choices to make: how much to contribute, and whether to invest in stocks, bonds or some other type of investment. The default chosen was a contribution rate of 3 percent of income invested in a money market fund. Basically, no one chose this combination before the change, but, of those automatically enrolled, 71 percent stayed with this default option. Figure 11.7 illustrates the big contrast between the window group and the new group of employees.

Clearly, new employees tended to stick to the default, despite this probably not being the best thing they could do. This provides something of a conundrum. While the increased participation looks like a good thing, that people follow the default investment choices looks like a bad thing. Indeed, an employee who would have chosen to opt into the savings plan with different investment options may end up being worse off if he is enrolled in the plan by default.

I will come back to this conundrum shortly. For now, I want you to see again the distinction between incentives for *Homo economicus* and those for *Homo sapiens*. Changing the default option should make no difference to *Homo economicus*. Whatever the default, he would have worked out his optimal savings plan and enacted it. For *Homo sapiens* who are time-inconsistent or have other biases, the change can – and,

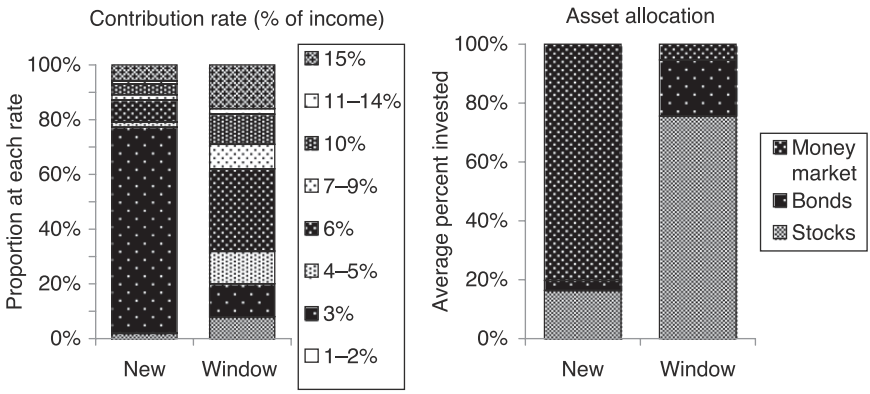

Figure 11.7 Contribution rate and asset allocation in the new and window groups. *Source:* Madrian and Shea (2001).

we see, does – matter. Someone with present bias, for example, may procrastinate and never get around to changing their savings choices despite a desire to do so.

11.2.3 Nudge

In both the previous examples we see changes, namely the introduction of a new account and a default of enrollment in a savings plan, that made a big difference. To calm, calculating *Homo economicus* they should not have made a difference, but it probably came as little surprise that they did make a difference for normal people. This opens up interesting new ways to think of economics and policy, informed by behavioral economics.

The two examples illustrate two different possibilities. In the first example, when a new account is created, something 'real' was changed, in that people had different choices. Ideas from behavioral economics suggested that these different choices may prove popular and worthwhile. In the second example, when the default is changed, nothing really changes, in that everyone still has the same choices as before. The framing of those choices is changed, however, and ideas from behavioral economics suggest why this change in framing may matter.

So, behavioral economics can suggest changes to choices, and changes to the framing of choices, that may make a difference. One of the reasons why this has excited interest in policy-making is that both can come at relatively little cost! Changing the default and administering a new account are relatively cheap to do, particularly when compared with the things governments might normally do, such as change tax relief on savings. Basically, changing the incentives for *Homo sapiens* is a lot cheaper than changing them for *Homo economicus*. But to put behavioral economics to best use we need a framework to work with. The idea of nudge, made popular in a book by Richard Thaler and Cass Sunstein (2008), is one way to go. To explain what a nudge is, we first need to explain the idea of a choice architect. A **choice architect** is anyone responsible for framing a decision that others will take. One example would be a sales assistant explaining to Emma the merits of

different products that she might buy. The assistant has discretion to frame the choices in different ways, by maybe emphasizing the good points of one product and the bad of another. A second example would be the way that share prices are listed when Alan looks up how well his investments are doing. Investments could be listed in alphabetical order or performance order, declining stocks could be highlighted in bright red or left in normal font, and so on.

This book has been full of examples when differences in framing have mattered. For instance, whether the sales tax is included or not in the price and whether the default is to be enrolled or not enrolled in the savings plan are classic examples of how framing matters. We can be pretty confident that how the sales assistant talks to Emma and how Alan's investments are displayed on the screen will also matter.

Now I want to move beyond saying that framing matters to asking whether we can change the framing of a decision in a way that helps people make better choices. This brings us to the idea of **nudge**. To quote from Thaler and Sunstein (2008: 6),

> A nudge, as we will use the term, is any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid.

So, something that changes the incentives of *Homo economicus* cannot be a nudge. Something that changes the incentives of *Homo sapiens* can. Changing the default option to that of being enrolled in a savings plan rather than being not enrolled is a fantastic example of a nudge. The options and economic incentives of the person are exactly the same whatever the default. All we have done is change the choice architecture or the framing of the choices in a seemingly small way. Even so, we get a predictable and big change in behavior.

Can policy-makers harness the power of nudge, and nudge people into making better choices? To answer this question, it's useful to think about two subsidiary questions. When do people need a nudge? And how can people be most effectively given a nudge?

The first question is relatively easy for us to answer, given what we have already covered in this book. We have seen that people are pretty clever, most of the time, and do learn with experience. So, nudges should be needed only sparingly, but there are four good reasons why we may need to nudge someone such as Alan. (1) Alan may face a very complicated choice and, clever or not, can easily get it wrong. (2) He has had very little experience or chance to learn what is best. (3) He is relatively uninformed about what he should choose. (4) He has self-control problems, and so may not make the choices he would have wanted himself to make.

If something ticks all four boxes then it looks a good place for a nudge. Saving for retirement is one clear example: it involves complicated choices; Alan is going to get old only once, so he has no chance to go back and start again having learnt from experience; he may have no real idea what savings he will need in 30 or 40 years' time; present bias may also mean he puts off doing things that he knows he should be doing. In the United States, health care is another thing likely to tick all four boxes.

This gives us a framework to think about when nudges might be needed. The next thing we need to think about is how to most effectively use nudges. This means that it's time for our first mnemonic. Table 11.6 briefly describes the six principles

Table 11.6 Principles for good choice architecture, spelling 'NUDGES'.

Principles	Brief description
iNcentives	People do respond to incentives such as price and cost, but only if these are salient
Understand mappings	People may need help understanding the mapping, from the choices they may make to the outcomes they will get
Defaults	Defaults matter a lot because of present bias and choice overload, so think carefully about them
Give feedback	People do learn, so give feedback on when things are going well or badly
Expect error	People make mistakes, so we need something that is as forgiving as possible to mistakes they may make
Structure complex choices	The more complex the choice, the more problems a person has, and the more likely context effects will matter; so, keep things simple

Source: Thaler and Sunstein (2008).

that Thaler and Sunstein (2008) suggest for good choice architecture. Only the first of these would get much attention in thinking about pushing *Homo economicus*. The basic idea, therefore, is to supplement incentives with things that take account of why the nudge is needed: to lessen the complexity (structure complex choices), create experience (give feedback), improve information (understand mappings) and work around biases (defaults, incentives; and expect error).

This might look fine in principle, but can it work? Things are not necessarily so easy. For example, we have seen that making it the default option for people to be opted into a savings plan may not be enough. We also need to encourage participants to choose the best contribution rate and stock options. Nudging people is not simple, therefore, and can leave some worse off than they might have been. But there are some basic nudges that can move things in the right direction for most people. For example, most would probably agree that making it the default option for people to be opted into a savings plan is a good thing.

Thaler and Sunstein suggest various other nudges, and you can probably come up with some yourself. Before we move on, I want to look at one of these, which I think is pretty neat. It's called **save more tomorrow**. The idea is that a saver, such as Mathan, can commit himself to save more tomorrow, in the sense that his contributions to a savings plan will automatically increase every time his pay increases. Why does this make sense? Mathan might learn that he is saving too little but not want to increase his current savings, because that would lower his income, and we know from Chapters 2 and 4 that people don't like that. If he says 'I will save more tomorrow', we also know that this is probably not going to happen, because of present bias. The save more tomorrow plan means that Mathan commits to saving tomorrow.

If he knows he has a present bias then this looks like a great commitment device that helps Mathan to save and overcome any time inconsistency. So, does it work? Thaler and Benartzi (2004) report on the successful first implementation of the plan; employees chose the save more tomorrow plan and stuck with it, and their savings rate almost quadrupled over the sample period. It's safe to say that this would not have happened without the plan, and many large employers have now adopted the idea. So, this is one nudge that does look as though it works.

11.2.4 Nudge and behavior change

One concern that many express about nudge is that it sounds a bit too much like Big Brother. It sounds a bit too much like the policy-maker deciding on someone's behalf, and can we trust policy-makers to know what is best?

An important aspect of nudge is that it should not involve sacrifice of liberty. This was clear in the quote I gave above from Thaler and Sunstein: '[W]ithout forbidding any options or significantly changing their economic incentives… [T]he intervention must be easy and cheap to avoid.' In principle, therefore, a nudge should not impact on personal liberty. This is a bit different, however, from saying that a nudge is not about changing behavior in the way a policy-maker thinks is best. Basically, a nudge should not change the set of options someone such as Mathan has, but its very objective is to change the option that Mathan will choose.

There is no escaping the fact, therefore, that nudge does require us to think about whether the policy-maker knows best. It is about **behavior change**. That's one reason to think through, as we have done, when people will most need a nudge. Is that enough? Well, we saw in the last chapter that most people might be comfortable with behavior change, and others making difficult decisions on their behalf, provided they retain some choice and self-determination. A nudge retains choice and so looks reasonable enough on this count. Even so, you may be wondering whether we should be trusting that policy-makers know what's best. The crucial point I want to make here is that we have no choice but to do this.

Recall that in Chapters 2 and 9 I argued that framing and context effects are inevitable: context influences perceptions and intuition, which influence reasoning, which influence behavior. What that means is that behavior change is an inevitability of policy. To quote from a discussion document published by the United Kingdom's Institute for Government (Dolan and co-authors 2010):

> We may not agree on how we would like policy-makers to change our behaviour. But whether we like it or not, the actions of policy-makers, public service professionals, markets and our fellow citizens around us have big, and often unintended, impacts on our behaviour. 'Doing nothing' is rarely a neutral option: we are always busy shaping each other's behaviour.

It's the lack of a neutral option that is key here. Is the neutral option that Mathan be not put in a savings plan, put in a savings plan or put in a save more tomorrow plan? Quite simply, there is no neutral option; something has to be chosen. This means that policy-makers have no choice but to change behavior. The key thing is to try to change behavior in the right direction.

That means it is time for our second mnemonic. The Institute for Government has come up with the idea of MINDSPACE. The basic idea is to provide a practical means to think about behavior change and apply insights from behavioral economics and psychology in policy-making. This is being used by policy-makers in the United Kingdom, and so gives one great example of policy-makers using behavioral economics to try to change behavior in the right direction. Table 11.7 gives a brief overview of the main principles underlying MINDSPACE, and most of these are things that we have looked at in this book.

The MINDSPACE principles are designed to help policy-makers better understand how possible policies might change behavior; they are a kind of checklist of

Table 11.7 An overview of MINDSPACE.

Principle	Description
Messenger	We are heavily influenced by who communicates information
Incentives	Our responses to incentives are shaped by predictable mental short cuts, such as strongly avoiding losses
Norms	We are strongly influenced by what others do
Defaults	We 'go with the flow' of pre-set options
Salience	Our attention is drawn to what is novel and seems relevant to us
Priming	Our acts are often influenced by subconscious cues
Affect	Our emotional associations can powerfully shape our actions
Commitments	We seek to be consistent with our public promises, and reciprocate acts
Ego	We act in ways that make us feel better about ourselves

Source: Dolan *et al.* (2010).

things to think about. The interesting thing that I want to finish with is how this checklist can serve different purposes. I will quote again from the report (Dolan and co-authors 2010):

Enhance. MINDSPACE can help policy-makers understand how current attempts to change behaviour could be improved.

Introduce. Some of the elements in MINDSPACE are not used extensively by policy-makers, yet may have considerable impact.

Reassess. Government needs to understand ways it may be changing the behaviour of citizens unintentionally. It is quite possible that the state is producing unintended – and possibly unwanted – changes in behaviour. The insights from MINDSPACE offer a rigorous way of assessing whether and how government is shaping the behaviour of its citizens.

The nice thing about this is that it illustrates the three different ways in which behavioral economics can have a positive impact on policy. It can help us come up with new ideas, such as the save more tomorrow nudge; an example of 'introduce'. It can help us improve current policy, such as changing the default to automatic enrollment; an example of 'enhance'. Finally, it can help us realize the problems with current policy, such as appreciating why a default of no automatic enrollment does not work; an example of 'reassess'.

I would argue, therefore, that behavior change is not only inevitable, but is also something to be embraced, as a positive means by which policy-makers can design policies that work and avoid policies that do not.

11.2.5 Consumer protection, health and the environment

The focus of this section so far has been primarily on encouraging saving. This is no accident. Saving is one area of life in which many people need a nudge, and simple nudges can make a huge difference. If you read the general literature on nudge you could easily get the impression, however, that nudge is only about encouraging saving. All the good examples seem to be saving examples! To dispel such thoughts, I want to finish by highlighting some areas other than saving where nudge can make a difference.

In order to do this, I will point you in the direction of the UK government's Behavioural Insights Team, or nudge unit. Set up in 2010, the team is estimated to have saved the UK government over $1 billion in its first five years. Suitably impressed, other governments, including in the United States, have set up similar teams. The policy papers produced by the United Kingdom's team are an indication of how widely nudge can be applied. The reports include 'Applying behavioural insights to charitable giving', 'Fraud, error and debt', 'Behaviour change and energy use', 'Better choices: better deals', 'Behavioural insights for cities' and 'Applying behavioural insight to health'. Let us delve a little deeper into two of the papers.

I shall start with the 'Better choices: better deals' paper. My plan is not to go through the policy prescriptions in full, but merely illustrate some of the issues. (The full papers are well worth a read, and freely available on the team's website.) The basic aim of 'Better choices: better deals' is to empower consumers to make better choices and get better deals, better value, better customer service and better support when seeking help.

To understand the issues at stake, consider Mathan deciding which bank to use. Clearly, Mathan wants to find the best match for his preferences; he will search around to see what the different banks have to offer in terms of fees, interest rates, accessibility, and so on. In Chapter 2 we looked at some search heuristics that Mathan might want to use when doing this, and one of the key things we learnt is that framing and context matter. The most successful banks are going to be keenly aware of the importance of framing and take this into account when designing their products.

One common trick is the use of hidden charges. These are charges that are 'hidden' but will almost certainly have to be paid at some point, such as a charge for going overdrawn, or for missing a credit card payment. Another trick is to offer a tempting, but short-term deal. For example, the bank might offer a higher interest rate for the first six months after opening a savings account. In this case the bank is hoping for naïve present-day bias; Mathan opens an account to take advantage of the great deal but then leaves his money in the account well after the great deal has expired.

For fairness' sake, I should point out that banks are not the only companies doing this. Indeed, it is hard to find companies that are not doing it! Consider, for example, an airline that offers a great fare, until you add on the fuel surcharge, taxes, fee for hold luggage and charges for food and drink. (One airline recently floated the idea of charging passengers to use the toilet!) Or consider a hotel that offers a great rate, but then charges for breakfast and use of the internet and swimming pool.

Ruses such as hidden charges and joining offers would not fool *Homo economicus*; they do fool us. The idea of 'Better choices: better deals' is to put more power in the hands of consumers. What does this mean in practice?

One key raft of policy proposals is about making sure that consumers have access to all relevant information in an easy-to-understand format. For example, credit card companies are now obliged to issue an annual statement that clearly states all fees and costs incurred on the card during the year; no more hidden charges. To give a second example, energy suppliers have agreed to simplify their tariffs and provide customers with clear information about the lowest available tariff; no more short-term great deals.

A second raft of proposals focuses on making it easier for consumers to learn from each other's experiences. For example, online consumer feedback is to be encouraged, as are comparison websites. In order to maintain the integrity of such

sites, the Office of Fair Trading will monitor for manipulation by companies (and a number of companies have already been prosecuted). Performance and complaints data will also be made more transparent.

The basic rationale behind all these proposals is to give companies less room to hide. That can, hopefully, nudge consumers to make better choices. This is clearly good news for consumers. It is not so obvious that it is good news for companies. Note, however, that the manipulation of consumers can be something of a race to the bottom: if one company starts playing dirty, it is very hard for others to survive unless they go down the same route. The companies that are offering the best deal, rather than making it look as though they are offering the best deal, should welcome any proposals to make consumers more informed. So, this can be a win-win outcome for consumers and good firms.

Let us look next at applying behavioral insight to health. To understand the issues in this case, consider Table 11.8. The table lists the main risks for lost years of healthy life as measured by the World Health Organization (WHO). In high-income countries we see that the top ten risks are full of things that an individual can avoid. You might rejoinder that a person has to die of something, and so why not die of tobacco or alcohol? The risks are ranked using QALYs, however (see Chapter 10), and so the losses are very real; people are missing out on a higher quality of life. Given what we have seen about addiction and habit, it is hard to believe that this is optimal in any sense (except perhaps for the tobacco companies).

This somewhat depressing picture suggests that we need ways to encourage healthy behavior. The policy paper 'Applying behavioural insight to health' of the Behavioural Insights Team sets out some possibilities. They consider ways to help people quit smoking, eat more healthily, exercise more, and so on. One example is the Department of Health's Change for Life campaign, with the slogan 'Eat well, move more, live longer'. The campaign advocates simple but potentially effective things, such as always using a shopping list. We do not know yet what nudges are going to be effective, and so a lot more research is needed, but the hope is that simple nudges can make a big difference.

And nudge can not only help people improve their own life; it might also help them improve the lives of others. To illustrate, consider organ donation. The United

Table 11.8 The leading risks for lost years of healthy life ranked from highest risk, one, to tenth highest risk.

Rank	High-income countries	Low-income countries
1	Tobacco	Childhood underweight
2	Alcohol	Unsafe water and sanitation
3	Overweight	Unsafe sex
4	High blood pressure	Suboptimal breastfeeding
5	High blood glucose	Indoor smoke from fuel
6	Physical inactivity	Vitamin A deficiency
7	High cholesterol	High blood pressure
8	Illicit drugs	Alcohol use
9	Occupational risks	High blood glucose
10	Low fruit and vegetables	Zinc deficiency

Source: WHO (2009).

Kingdom has traditionally used a system of **informed consent** for organ donation; this means that the default option is no organ donation. Other countries, such as Austria, use a system of **presumed consent**; this means that the default option is organ donation. It will probably come as no surprise that organ donation rates appear higher in countries with presumed consent.

The United Kingdom, like many other countries, needs more organ donors. The trend, therefore, is to move away from informed consent. Wales, for instance, adopted a presumed consent policy in 2015, and France in 2017. Presumed consent is controversial, however, and so a more neutral policy such as **prompted choice** may be desirable. In this case, people are required to decide, when filling in something such as a driving license application, whether they want to become an organ donor; most do. Several US states, including Texas and California, have recently introduced such a policy. Thus, the simplest of nudges can save lives.

11.3 Summary

I started by looking at how behavioral economics can help design good institutions, looking at the tragedy of the commons, matching markets and auctions as examples. We saw that insights from behavioral economics can be useful in better predicting behavior. We also saw that experiments can be a useful testing ground to see how institutions will perform.

Next we looked at how behavioral economics can help inform on policies that are directed at changing individual behavior. Two examples showed that small changes could have big, and desirable, consequences. This leads us to the concept of nudge and behavior change. We saw that behavioral economics can help us come up with new ideas, such as the save more tomorrow program, and can also help us reassess why existing policies may not be working.

An important point in looking at both institution design and changing individual incentives is that context and framing will matter. This means that policy-makers cannot be neutral. What they do is going to influence behavior and/or the success of the institution. They need, therefore, to come up with the policies that have the most positive impact, and, for that, they need behavioral economists to come up with the guiding principles, such as those in Tables 11.6 and 11.7. It's testament to the development of behavioral economics that it is now making an impact on policy-making. But, hopefully, this is only the beginning. Behavioral economics has the power to completely change the way we think about all aspects of the economy and economic behavior. What's more, I believe it gives us a much better way to think about and understand economic behavior. That's why my hope is that behavioral economics will soon become something that all economists do some of the time, rather than something some economists do all of the time.

11.4 Further reading

The book by Ostrom (1990) and article (2010) are a good place to start for more on the tragedy of the commons. See also Schmitt, Swope and Walker (2000). For more on matching markets, see Roth (2002, 2008). On auction design, there are

lots of good resources. The books by Klemperer (2004) and Milgrom (2004) are good places to start. The web is also a great source of information. For more on the UK spectrum allocation, see the Ofcom website (www.ofcom.org.uk), and, for more on the US spectrum allocation, see the FCC website (www.fcc.gov). Thaler and Sunstein (2008) is the obvious place to read more about nudge, and there is a lot more about MINDSPACE on the Institute for Government website (www.instituteforgovernment.org.uk), including the full report and interviews with the authors. See also Halpern (2015).

11.5 Online materials

On the companion website you will have access to flashcards, quizzes and answers to the review questions.

11.6 Review questions

11.1 What is the difference between a weak link game, a threshold public good game and a common pool resource game?

11.2 Why is it easier to achieve sustainable fishing for national, inshore fishing than international, offshore fishing?

11.3 Do you think subjects should be allowed to communicate with each other during experiments?

11.4 Carefully explain the difference between a stable and unstable matching algorithm.

11.5 Should the objective of a government running a spectrum auction be to raise as much money as possible?

11.6 The European 4G auctions raised a lot less money than the 3G auctions, particularly in the United Kingdom. Does that mean the auctions were designed badly?

11.7 Why is it useful to distinguish between policies about institutions and about individual behavior? You might want to contrast saving choices with interactive situations, such as a threshold public good game.

11.8 Why might people need a nudge in deciding what insurance to buy, or health care provider to use in the United States? Where else might nudges make sense?

11.9 How could a charity or private firm exploit the save more tomorrow idea?

11.10 Is there any difference between a nudge that works at a conscious level and one that works at a subconscious level?

11.11 What are the advantages and disadvantages of employers being automatically enrolled in a default savings plan?

11.12 Would *Homo economicus* be influenced at all by nudges?

11.13 Should governments ask citizens whether they agree with behavior change, or just get on with it?

11.14 Come up with some nudges of your own.

11.15 Can nudge be useful in avoiding the confirmatory bias and the law of small numbers?

Bibliography

Abbink, K., Irlenbusch, B., and Renner, E. (2006) 'Group size and social ties in microfinance institutions', *Economic Inquiry*, 44: 614–628.

Agarwal, S., Liu, C., and Souleles, N. S. (2007) 'The reaction of consumer spending and debt to tax rebates: Evidence from consumer credit data', *Journal of Political Economy*, 115: 986–1019.

Alberti, F., and Cartwright, E. J. (2015) 'Does the endowment of contributors make a difference in threshold public-good games?', *FinanzArchiv: Public Finance Analysis*, 71: 216–239.

Allcott, H. (2011) 'Social norms and energy conservation', *Journal of Public Economics*, 95: 1082–1095.

Allcott, H., and Mullainathan, S. (2010) 'Behavioral science and energy policy', *Science*, 327: 1204–1205.

Alm, J. (2012) 'Measuring, explaining, and controlling tax evasion: Lessons from theory, experiments, and field studies', *International Tax and Public Finance*, 19: 54–77.

Alm, J., Cherry, T., Jones, M., and McKee, M. (2010) 'Taxpayer information assistance services and tax compliance behavior', *Journal of Economic Psychology*, 31: 577–586.

Alpizar, F., Carlsson, F., and Johansson-Stenman, O. (2008) 'Anonymity, reciprocity, and conformity: Evidence from voluntary contributions to a national park in Costa Rica', *Journal of Public Economics*, 92: 1047–1060.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2008) 'Eliciting risk and time preferences', *Econometrica*, 76: 583–618.

Anderson, L. R., and Holt, C. A. (1997) 'Information cascades in the laboratory', *American Economic Review*, 87: 847–862.

Andreoni, J. (1990) 'Impure altruism and donations to public goods: A theory of warm-glow giving', *Economic Journal*, 100: 464–477.

—— (1993) 'An experimental test of the public-goods crowding-out hypothesis', *American Economic Review*, 83: 1317–1327.

—— (1995) 'Cooperation in public-goods experiments: kindness or confusion?', *American Economic Review*, 85: 891–904.

—— (2006) 'Philanthropy', in Kolm, S.-C., and Ythier, M. (Eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity*, Vol. II, *Applications*: 1201–1269, Amsterdam: Elsevier.

Andreoni, J., and Miller, J. (2002) 'Giving according to GARP: An experimental test of the consistency of preferences for altruism', *Econometrica*, 70: 737–753.

Andreoni, J., and Payne, A. A. (2013) 'Charitable giving', in Auerbach, A. J., Chetty, R., Feldstein, M., and Saez, E. (Eds.), *Handbook of Public Economics*: 1–50, Amsterdam: North Holland.

Andreoni, J., and Sprenger, C. (2012) 'Risk preferences are not time preferences', *American Economic Review*, 102: 3357–3376.

Antonakis, J., and Dalgas, O. (2009) 'Predicting elections: Child's play!', *Science*, 323: 1183.

Ariely, D. (2012) *The Honest Truth about Dishonesty: How We Lie to Everyone – Especially Ourselves*. New York: HarperCollins.

Ariely, D., Loewenstein, G., and Prelec, D. (2003) 'Coherent arbitrariness: Stable demand curves without stable preferences', *Quarterly Journal of Economics*, 118: 73–105.
Ariely, D., and Wertenbroch, K. (2002) 'Procrastination, deadlines, and performance: Self-control by precommitment', *Psychological Science*, 13: 219–224.
Artavanis, N., Morse, A., and Tsoutsoura, M. (2015) 'Tax evasion across industries: Soft credit evidence from Greece', Working Paper no. 21552, Cambridge, MA: NBER.
Asch, S. E. (1956) 'Studies of independence and conformity: I. A minority of one against a unanimous majority', *Psychological Monographs: General and Applied*, 70: 1–70.
Ashraf, M., Karlan, D., and Yin, W. (2006) 'Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines', *Quarterly Journal of Economics*, 121: 635–672.
Attanasi, G., Battigalli, P., and Manzoni, E. (2015) 'Incomplete-information models of guilt aversion in the trust game', *Management Science*, 62: 648–667.
Baicker, K., Buckles, K. S., and Chandra, A. (2006) 'Geographic variation in the appropriate use of cesarean delivery', *Health Affairs*, 25: w355–w367.
Ballew, C. C., and Todorov, A. (2007) 'Predicting political elections from rapid and unreflective face judgments', *Proceedings of the National Academy of Sciences*, 104: 17948–17953.
Banerjee, A. V. (2013) 'Microcredit under the microscope: What have we learned in the past two decades, and what do we need to know?', *Annual Review of Economics*, 5: 487–519.
Banerjee, A. V., Cole, S., Duflo, E., and Linden, L. (2007) 'Remedying education: Evidence from two randomized experiments in India', *Quarterly Journal of Economics*, 122: 1235–1264.
Banerjee, A. V., and Duflo, E. (2009) 'The experimental approach to development economics', *Annual Review of Economics*, 1: 151–178.
Banerjee, A., Duflo, E., Glennerster, R., and Kinnan, C. (2015) 'The miracle of microfinance? Evidence from a randomized evaluation', *American Economic Journal: Applied Economics*, 7: 22–53.
Barber, B., and Odean, T. (2008) 'All that glitters: The effect of attention and news on the buying behaviour of individual and institutional investors', *Review of Financial Studies*, 21: 785–818.
Barberis, N. C. (2013a) 'The psychology of tail events: Progress and challenges', *American Economic Review*, 103: 611–616.
—— (2013b) 'Thirty years of prospect theory in economics: A review and assessment', *Journal of Economic Perspectives*, 27: 173–195.
Barberis, N. C., and Xiong, W. (2009) 'What drives the disposition effect? An analysis of a long-standing preference-based explanation', *Journal of Finance*, 64: 751–784.
Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C., and Sugden, R. (2010) *Experimental Economics: Rethinking the Rules*, Princeton, NJ: Princeton University Press.
Barseghyan, L., Molinari, F., O'Donoghue, T., and Teitelbaum, J. C. (2013) 'The nature of risk preferences: Evidence from insurance choices', *American Economic Review*, 103: 2499–2529.
Bateman, I., Kahneman, D., Munro, A., Starmer, C., and Sugden, R. (2005) 'Testing competing models of loss aversion: An adversarial collaboration', *Journal of Public Economics*, 89: 1561–1580.
Battigalli, P., and Dufwenberg, M. (2009) 'Dynamic psychological games', *Journal of Economic Theory*, 144: 1–35.
Becker, G., and Murphy, K. (1988) 'A theory of rational addiction', *Journal of Political Economy*, 46: 675–700.

Benartzi, S., and Thaler, R. H. (1995) 'Myopic loss aversion and the equity premium puzzle', *Quarterly Journal of Economics*, 110: 73–92.
——— (1999) 'Risk aversion or myopia? Choices in repeated gambles and retirement investments', *Management Science*, 45: 364–381.
——— (2001) 'Naive diversification strategies in defined contribution saving plans', *American Economic Review*, 91: 79–98.
——— (2002) 'How much is investor autonomy worth?', *Journal of Finance*, 57: 1593–1616.
——— (2007) 'Heuristics and biases in retirement savings behavior', *Journal of Economic Perspectives*, 21: 81–104.
Benzion, U., Rapoport, A., and Yagil, J. (1989) 'Discount rates inferred from decisions: An experimental study', *Management Science*, 35: 270–284.
Berg, J., Dickhaut, J., and McCabe, K. (1994) 'Trust, reciprocity, and social history', *Games and Economic Behavior*, 10: 122–142.
Berg, J., Dickhaut, J., and Rietz T. (2010) 'Preference reversals: The impact of truth-revealing monetary incentives', *Games and Economic Behavior*, 68: 443–468.
Berg, N., and Gigerenzer, G. (2010) 'As-if behavioral economics: Neoclassical economics in disguise?', *History of Economic Ideas*, 18: 133–166.
Bergstrom, T. (1996) 'Economics in a family way', *Journal of Economic Literature*, 34: 1903–1934.
——— (2002) 'Evolution of social behaviour: Individual and group selection', *Journal of Economic Literature*, 16: 67–88.
Bernheim, D., and Rangel, A. (2004) 'Addiction and cue-triggered decision processes', *American Economic Review*, 94: 1558–1590.
——— (2007) 'Toward choice-theoretic foundations for behavioral welfare economics', *American Economic Review*, 97: 464–470.
Bertrand, M., Mullainathan, S., and Shafir, E. (2004) 'A behavioral-economics view of poverty', *American Economic Review*, 94: 419–423.
Bewley, T. (2007) 'Fairness, reciprocity, and wage rigidity', in Diamond, P., and Vartiainen, H. (Eds.), *Behavioral Economics and Its Applications*: 157–187, Princeton, NJ: Princeton University Press.
Bikhchandani, S., Hirshleifer, D., and Welch, I. (1998) 'Learning from the behavior of others: Conformity, fads, and informational cascades', *Journal of Economic Perspectives*, 12: 151–170.
Binmore, K. (2008) 'Review of behavioral economics and its applications', *Economic Journal*, 118: F248–F251.
Binmore, K., and Shaked, A. (2010) 'Experimental economics: Where next?', *Journal of Economic Behavior and Organization*, 73: 87–100.
Blinder, A. S., and Morgan, J. (2005) 'Are two heads better than one? Monetary policy by committee', *Journal of Money, Credit and Banking*, 37: 789–811.
Bolton, G., and Katok, E. (1998) 'An experimental test of the crowding out hypothesis: The nature of beneficent behavior', *Journal of Economic Behavior and Organization*, 37: 315–331.
Bolton, G., and Ockenfels, A. (2000) 'ERC: A theory of equity, reciprocity, and competition', *American Economic Review*, 90: 166–193.
Bornstein, G., and Yaniv, I. (1998) 'Individual and group behavior in the ultimatum game: Are groups more "rational" players?', *Experimental Economics*, 1: 101–108.
Bortolotti, S., Devetag, G., and Ortmann, A. (2009) 'Exploring the effects of real effort in a weak-link experiment', DOI: 10.2139/ssrn.1355444.
Bowles, S. (2006) 'Group competition, reproductive leveling and the evolution of human altruism', *Science*, 314: 1555–1556.

Boyd, R. (2006) 'The puzzle of human sociality', *Science*, 314: 1560–1563.
Brandts, J., and Charness, G. (2011) 'The strategy versus the direct-response method: a first survey of experimental comparisons', *Experimental Economics*, 14: 375–398.
Brandts, J., and Cooper, D. (2006) 'A change would do you good... An experimental study on how to overcome coordination failure in organizations', *American Economic Review*, 96: 669–693.
—— (2007) 'It's what you say, not what you pay: An experimental study of manager–employee relationships in overcoming coordination failure', *Journal of the European Economic Association*, 5: 1223–1268.
Brandts, J., and Holt, C. A. (1992) 'An experimental test of equilibrium dominance in signaling games', *American Economic Review*, 82: 1350–1365.
Breiter, H., Aharon, I., Kahneman, D., Dale, A., and Shizgal, P. (2001) 'Functional imaging of neural responses to expectancy and experience of monetary gains and losses', *Neuron*, 30: 619–639.
Breitmoser, Y. (2012) 'Strategic reasoning in p-beauty contests', *Games and Economic Behavior*, 75: 555–569.
Brickman, P., Coates, D., and Janoff-Bulman, R. (1978) 'Lottery winners and accident victims: Is happiness relative?', *Journal of Personality and Social Psychology*, 36: 917–927.
Brown, R. (1995) *Prejudice: It's Social Psychology*, Hoboken, NJ: Wiley-Blackwell.
Browning, M., and Crossley, T. F. (2001) 'The life-cycle model of consumption and saving', *Journal of Economic Perspectives*, 15: 3–22.
Bruni, L., and Sugden, R. (2007) 'The road not taken: How psychology was removed from economics, and how it might be brought back', *Economic Journal*, 117: 146–173.
Brunner, C., Camerer, C., and Jacob, K. (2011) 'Stationary concepts for experimental 2 X 2 games: Comment', *American Economic Review*, 101: 1029–1040.
Buehn, A., and Schneider, F. (2012) 'Shadow economies around the world: Novel insights, accepted knowledge, and new estimates', *International Tax and Public Finance*, 19: 139–171.
Cadsby, C., and Maynes, E. (1999) 'Voluntary provision of threshold public goods with continuous contributions: Experimental evidence', *Journal of Public Economics*, 71: 53–73.
Caginalp, G., Porter, D. P., and Smith, V. L. (2000) 'Overreactions, momentum, liquidity, and price bubbles in laboratory and field asset markets', *Journal of Behavioral Finance*, 1: 24–48.
Camerer, C. (2003) *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton, NJ: Princeton University Press.
—— (2007) 'Neuroeconomics: Using neuroscience to make economic predictions', *Economic Journal*, 117: C26–C42.
Camerer, C., and Ho, T.-H. (1999) 'Experience-weighted attraction learning in normal form games', *Econometrica*, 67: 827–874.
Camerer, C., Ho, T.-H., and Chong, J.-K. (2004) 'A cognitive hierarchy model of games', *Quarterly Journal of Economics*, 119: 861–898.
Camerer, C., and Loewenstein, G. (2004) 'Behavioral economics: Past, present and future', in Camerer, C., Loewenstein, G., and Rabin, M. (Eds.), *Advances in Behavioral Economics*: 3–51, Princeton, NJ: Princeton University Press.
Camerer, C., Loewenstein, G., and Prelec, D. (2004) 'Neuroeconomics: Why economics needs brains', *Scandinavian Journal of Economics*, 106: 555–579.
—— (2005) 'Neuroeconomics: How neuroscience can inform economics', *Journal of Economic Literature*, 43: 9–64.

Camerer, C., and Thaler, R. H. (1995) 'Anomalies: Ultimatums, dictators and manners', *Journal of Economic Perspectives*, 9: 209–219.
Camerer, C., and Weber, M. (1992) 'Recent developments in modeling preferences: Uncertainty and ambiguity', *Journal of Risk and Uncertainty*, 5: 325–370.
Campbell, C., and Kamlani, K. (1997) 'The reasons for wage rigidity: Evidence from a survey of firms', *Quarterly Journal of Economics*, 112: 759–789.
Caplin, A., Dean, M., and Martin, D. (2011) 'Search and satisficing', *American Economic Review*, 101: 2899–2922.
Caplin, A., and Schotter, A. (Eds.) (2008) *The Foundations of Positive and Normative Economics: A Handbook*, Oxford: Oxford University Press.
Cardenas, J. C., and Carpenter, J. (2008) 'Behavioural development economics: Lessons from field labs in the developing world', *Journal of Development Studies*, 44: 311–338.
Carroll, C., Overland, J., and Weil, D. (2000) 'Saving and growth with habit formation', *American Economic Review*, 90: 341–355.
Cartwright, E. J., Gillet, J., and van Vugt, M. (2013) 'Leadership by example in a weak-link game', *Economic Inquiry*, 51: 2028–2043.
Cassar, A., and Friedman, D. (2004) *Economics Lab: An Introduction to Experimental Economics*, London: Routledge.
Chamberlin, E. H. (1948) 'An experimental imperfect market', *Journal of Political Economy*, 56: 95–108.
Charness, G., and Dufwenberg, M. (2006) 'Promises and partnership', *Econometrica*, 74: 1579–1601.
Charness, G., Feri, F., Meléndez-Jiménez, M. A., and Sutter, M. (2014) 'Experimental games on networks: Underpinnings of behavior and equilibrium selection', *Econometrica*, 82: 1615–1670.
Charness, G., and Rabin, M. (2002) 'Understanding social preferences with simple tests', *Quarterly Journal of Economics*, 117: 817–869.
Charness, G., and Sutter, M. (2012) 'Groups make better self-interested decisions', *Journal of Economic Perspectives*, 26: 157–176.
Chaudhuri, A., Schotter, A., and Sopher, B. (2009) 'Talking ourselves to efficiency: Coordination in inter-generational minimum effort games with private, almost common and common knowledge of advice', *Economic Journal*, 119: 91–122.
Chetty, R. (2006) 'A new method of estimating risk aversion', *American Economic Review*, 96: 1821–1834.
Chetty, R., Looney, A., and Kroft, K. (2009) 'Salience and taxation: Theory and evidence', *American Economic Review*, 99: 1145–1177.
Cheung, Y.-W., and Friedman, D. (1997) 'Individual learning in normal form games: Some laboratory results', *Games and Economic Behavior*, 19: 46–76.
Clark, A. E., Diener, E., Georgellis, Y., and Lucas, R. E. (2003) 'Lags and leads in life satisfaction: A test of the baseline hypothesis', Working Paper no. 2003–2014. Paris: DELTA.
Clark, A. E., Diener, E., Georgellis, Y., and Lucas, R. E. (2008) 'Lags and leads in life satisfaction: A test of the baseline hypothesis', *Economic Journal*, 118: F222–F243.
Clotfelter, C., and Cook, P. (1991) 'Lotteries in the real world', *Journal of Risk and Uncertainty*, 4: 227–232.
Coats, J., Gronberg, T., and Grosskopf, B. (2009) 'Simultaneous versus sequential public good provision and the role of refunds: An experimental study', *Journal of Public Economics*, 93: 326–335.
Cochard, F., Couprie, H., and Hopfensitz, A. (2016) 'Do spouses cooperate? An experimental investigation', *Review of Economics of the Household*, 14: 1–26.

Cohen, A., and Einav, L. (2007) 'Estimating risk preferences from deductible choice', *American Economic Review*, 97: 745–788.
Cohen, J. D., Ericson, K. M., Laibson, D., and White, J. M. (2016) 'Measuring time preferences', Working Paper no. 22455, Cambridge, MA: NBER.
Cohn, A., Fehr, E., and Maréchal, M. A. (2014) 'Business culture and dishonesty in the banking industry', *Nature*, 516: 86–89.
Conlisk, J. (1996) 'Why bounded rationality?', *Journal of Economic Literature*, 34: 669–700.
Cooper, D. (2006) 'Are experienced managers expert at overcoming coordination failure?', *Advances in Economic Analysis and Policy*, 6: 1–30.
Cooper, D., and Kagel, J. (2005) 'Are two heads better than one? Team versus individual play in signaling games', *American Economic Review*, 95: 477–509.
Coricelli, G., and Nagel, R. (2009) 'Neural correlates of depth of strategic reasoning in medial prefrontal cortex', *Proceedings of the National Academy of Sciences*, 106: 9163–9168.
Costa-Gomes, M., and Crawford, V. (2006) 'Cognition and behavior in two-person guessing games: An experimental study', *American Economic Review*, 96: 1737–1768.
Costa-Gomes, M., Crawford, V., and Iriberri, N. (2009) 'Comparing models of strategic thinking in van Huyck, Battalio, and Beil's coordination games', *Journal of the European Economic Association*, 7: 365–376.
Crawford, V. (1995) 'Adaptive dynamics in coordination games', *Econometrica*, 63: 103–143.
Crawford, V., and Iriberri, N. (2007) 'Level-k auctions: Can a non-equilibrium model of strategic thinking explain the winner's curse and overbidding in private value auctions?', *Econometrica*, 75: 1721–1770.
Croson, R. T., and Gächter, S. (2010) 'The science of experimental economics', *Journal of Economic Behavior and Organization*, 73: 122–131.
Croson, R. T., and Gneezy, U. (2009) 'Gender differences in preferences', *Journal of Economic Literature*, 47: 448–474.
Croson, R. T., and Marks, M. B. (2000) 'Step returns in threshold public goods: A meta- and experimental analysis', *Experimental Economics*, 2: 239–259.
—— (2001) 'The effect of recommended contributions in the voluntary provision of public goods', *Economic Inquiry*, 39: 238–249.
Croson, R. T., and Treich, N. (2014) 'Behavioral environmental economics: promises and challenges', *Environmental and Resource Economics*, 58: 335–351.
Dana, J., Weber, R. A., and Xi Kuang, J. (2007) 'Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness', *Economic Theory*, 33: 67–80.
Darley, J. M., and Gross, P. H. (1983) 'A hypothesis-confirming bias in labeling effects', *Journal of Personality and Social Psychology*, 44: 20–33.
De Martino, B., Kumaran, D., Seymour, B., and Dolan, R. J. (2006) 'Frames, biases, and rational decision-making in the human brain', *Science*, 313: 684–687.
De Quervain, D. J. F., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., and Fehr, E. (2004) 'The neural basis of altruistic punishment', *Science*, 305: 1254–1258.
Dechenaux, E., Kovenock, D., and Sheremeta, R. M. (2015) 'A survey of experimental research on contests, all-pay auctions and tournaments', *Experimental Economics*, 18: 609–669.
DellaVigna, S. (2009) 'Psychology and economics: Evidence from the field', *Journal of Economic Literature*, 47: 315–372.
DellaVigna, S., and Malmendier, U. (2004) 'Contract design and self-control: Theory and evidence', *Quarterly Journal of Economics*, 119: 353–402.

——— (2006) 'Paying not to go to the gym', *American Economic Review*, 96: 694–719.
Devetag, G., and Ortmann, A. (2007) 'When and why? A critical survey on coordination failure in the laboratory', *Experimental Economics*, 10: 331–344.
Di Tella, R., and MacCulloch, R. J. (2005) 'Partisan social happiness', *Review of Economic Studies*, 72: 367–393.
——— (2006) 'Some uses of happiness data in economics', *Journal of Economic Perspectives*, 20: 25–46.
Di Tella, R., MacCulloch, R. J., and Oswald, A. J. (2001) 'Preferences over inflation and unemployment: Evidence from surveys of happiness', *American Economic Review*, 91: 335–341.
Diekhof, E., and Gruber, O. (2010) 'When desire collides with reason: Functional interactions between anteroventral prefrontal cortex and nucleus accumbens underlie the human ability to resist impulsive desires', *Journal of Neuroscience*, 30: 1488–1493.
Dolan, P., Hallsworth, M., Halpern, D., King, D., and Vlaev, I. (2010) *MINDSPACE: Influencing Behavior through Public Policy*, London: Institute for Government.
Dolan, P., and Kahneman, D. (2008) 'Interpretations of utility and their implications for the valuation of health', *Economic Journal*, 118: 215–234.
Dufwenberg, M., and Kirchsteiger, G. (2004) 'A theory of sequential reciprocity', *Games and Economic Behavior*, 47: 268–298.
Dufwenberg, M., Lindqvist, T., and Moore, E. (2005) 'Bubbles and experience: An experiment', *American Economic Review*, 95: 1731–1737.
Easterlin, R. A. (1974) 'Does economic growth improve the human lot? Some empirical evidence', in David, P. A., and Reder, M. W. (Eds.), *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*: 89–126, New York: Academic Press.
——— (2013) 'Happiness, growth, and public policy', *Economic Inquiry*, 51: 1–15.
Ebert, S., and Strack, P. (2015) 'Until the bitter end: On prospect theory in a dynamic context', *American Economic Review*, 105: 1618–1633.
Eckel, C., and Grossman, P. (2001) 'Chivalry and solidarity in ultimatum games', *Economic Inquiry*, 30: 171–188.
Ellison, G. (2006) 'Bounded rationality in industrial organization', in Blundell, R., Newey, W., and Persson, R. (Eds.), *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Vol. II: 142–174, New York: Cambridge University Press.
Ellsberg, D. (1961) 'Risk, ambiguity, and the savage axioms', *Quarterly Journal of Economics*, 75: 643–669.
Engelmann, D., and Strobel, M. (2004) 'Inequality aversion, efficiency, and maximin preferences in simple distribution experiments', *American Economic Review*, 94: 857–869.
Epper, T., and Fehr-Duda, H. (2015) 'Comment on "Risk preferences are not time preferences": Balancing on a budget line', *American Economic Review*, 105: 2261–2271.
Falk, A., Fehr, E., and Fischbacher, U. (2003) 'On the nature of fair behavior', *Economic Inquiry*, 41: 20–26.
——— (2008) 'Testing theories of fairness: Intentions matter', *Games and Economic Behavior*, 62: 287–303.
Falk, A., and Fischbacher, U. (2006) 'A theory of reciprocity', *Games and Economic Behavior*, 54: 293–315.
Falk, A., Fischbacher, U., and Gächter, S. (2013) Living in two neighborhoods: Social interaction effects in the laboratory. *Economic Inquiry*, 51: 563–578.
Farber, H. S. (2005) 'Is tomorrow another day? The labour supply of New York City cabdrivers', *Journal of Political Economy*, 113: 46–82.

——— (2008) 'Reference-dependent preferences and labor supply: The case of New York City taxi drivers', *American Economic Review*, 98: 1069–1082.
——— (2015) 'Why you can't find a taxi in the rain and other labor supply lessons from cab drivers', *Quarterly Journal of Economics*, 130: 1975–2026.
Fehr, E., and Falk, A. (1999) 'Wage rigidity in a competitive incomplete contract market', *Journal of Political Economy*, 107: 106–134.
Fehr, E., and Fischbacher, U. (2004) 'Third-party punishment and social norms', *Evolution and Human Behavior*, 25: 63–87.
Fehr, E., and Gächter, S. (2000a) 'Cooperation and punishment in public goods experiments', *American Economic Review*, 90: 980–994.
——— (2000b) 'Fairness and retaliation: The economics of reciprocity', *Journal of Economic Perspectives*, 14: 159–181.
Fehr, E., Goette, L., and Zehnder, C. (2009) 'A behavioral account of the labor market: The role of fairness concerns', *Annual Review of Economics*, 1: 355–384.
Fehr, E., and Schmidt, J. (1999) 'A theory of fairness, competition, and cooperation', *Quarterly Journal of Economics*, 114: 817–868.
Fehr-Duda, H., and Epper, T. (2012) 'Probability and risk: Foundations and economic implications of probability-dependent risk preferences', *Annual Review Economics*, 4: 567–593.
Feri, F., Irlenbusch, B., and Sutter, M. (2010) 'Efficiency gains from team-based coordination: Large-scale experimental evidence', *American Economic Review*, 100: 1892–1912.
Fischbacher, U. (2007) 'z-Tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics*, 10: 171–178.
Fobian, C. S., and Christensen-Szalanski, J. J. (1993) 'Ambiguity and liability negotiations: The effects of the negotiators' role and the sensitivity zone', *Organizational Behavior and Human Decision Processes*, 54: 277–298.
Forsythe, R., Horowitz, J., Savin, N., and Sefton, M. (1994) 'Fairness in simple bargaining experiments', *Games and Economic Behavior*, 6: 347–369.
Frank, R. (1991) *Passions within Reason: The Strategic Role of the Emotions*, New York: W. W. Norton.
——— (2007) 'Behavioral economics and health economics', in Diamond, P., and Vartiainen, H. (Eds.), *Behavioral Economics and Its Applications*: 195–234, Princeton, NJ: Princeton University Press.
Frank, R., Gilovich, T., and Regan, D. (1993) 'Does studying economics inhibit cooperation?', *Journal of Economic Perspectives*, 7: 159–171.
Frank, R., and Hutchens, R. (1993) 'Wages, seniority, and the demand for rising consumption profiles', *Journal of Economic Behavior and Organization*, 3: 251–276.
Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002) 'Time discounting and time preference: A critical review', *Journal of Economic Literature*, 40: 351–401.
Frey, B. (2008) *Happiness: A Revolution in Economics*, Cambridge, MA: MIT Press.
Frey, B., and Stutzer, A. (2002) 'What can economists learn from happiness research?', *Journal of Economic Literature*, 40: 402–435.
Friedman, M. (1953) *Essays in Positive Economics*, Chicago: University of Chicago Press.
Fudenberg, D., and Levine, D. K. (2006) 'A dual-self model of impulse control', *American Economic Review*, 96: 1449–1476.
Fudenberg, D., Levine, D. K., and Maniadis, Z. (2012) 'On the robustness of anchoring effects in WTP and WTA experiments', *American Economic Journal: Microeconomics*, 4: 131–145.
Gabaix, X., Laibson, D., Moloche, G., and Weinberg, S. (2006) 'Costly information acquisition: Experimental analysis of a boundedly rational model', *American Economic Review*, 96: 1043–1068.

Gächter, S., and Falk, A. (2002) 'Reputation and reciprocity: Consequences for the labour relation', *Scandinavian Journal of Economics*, 104: 1–26.
Gächter, S., Herrmann, B., and Thöni, C. (2010) 'Culture and cooperation', *Philosophical Transactions of the Royal Society B – Biological Sciences*, 365: 2651–2661.
Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989) 'Psychological games and sequential rationality', *Games and Economic Behavior*, 1: 60–79.
Genesove, D., and Mayer, C. (2001) 'Loss aversion and seller behaviour: Evidence from the housing market', *Quarterly Journal of Economics*, 116: 1233–1260.
Gigerenzer, G., Todd, P. M., and ABC Research Group (1999) *Simple Heuristics that Make Us Smart*. New York: Oxford University Press.
Gilbert, D. T., Gill, M. J., and Wilson, T. D. (2002) 'The future is now: Temporal correction in affective forecasting', *Organizational Behavior and Human Decision Processes*, 88: 430–444.
Gilboa, I., and Schmeidler, D. (1995) 'Case-based decision theory', *Quarterly Journal of Economics*, 110: 605–639.
Gillet, J., Cartwright, E. J., and van-Vugt, M. (2010) 'Leadership in the weak link game', School of Economics, University of Kent working paper.
Gilovich, T., Vallone, R., and Tversky, A. (1985) 'The hot hand in basketball: On the misperception of random sequences', *Cognitive Psychology*, 17: 295–314.
Gintis, H., Bowles, S., Boyd, R., and Fehr, E. (Eds.) (2005) *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*, Cambridge, MA: MIT Press.
Glewwe, P., Kremer, M., and Moulin, S. (2009) 'Many children left behind? Textbooks and test scores in Kenya', *American Economic Journal: Applied Economics*, 1: 112–135.
Glimcher, P. W. (2004) *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics*, Cambridge, MA: MIT Press.
Glimcher, P. W., and Fehr, E. (Eds.) (2013) *Neuroeconomics: Decision Making and the Brain*, 2nd Edn., London: Academic Press.
Gneezy, U. (2005) 'Deception: The role of consequences', *American Economic Review*, 95: 384–394.
Gneezy, U., and List, J. (2006) 'Putting behavioral economics to work: Field evidence on gift exchange', *Econometrica*, 74: 1365–1384.
Gneezy, U., and Rustichini, A. (2004) 'Gender and competition at a young age', *American Economic Review*, 94: 377–381.
Gode, D. K., and Sunder, S. (1993) 'Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality', *Journal of Political Economy*, 101: 119–137.
Goeree, J. K., and Yariv, L. (2015) 'Conformity in the lab', *Journal of the Economic Science Association*, 1: 15–28.
Goldfarb, A., and Xiao, M. (2011) 'Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets', *American Economic Review*, 101: 3130–3161.
Grosskopf, B. (2003) 'Reinforcement and directional learning in the ultimatum game with responder competition', *Experimental Economics*, 6: 141–158.
Gruber, J., and Köszegi, B. (2001) 'Is addiction "rational"? Theory and evidence', *Quarterly Journal of Economics*, 116: 1261–1303.
Gul, F., and Pesendorfer, W. (2001) 'Temptation and self-control', *Econometrica*, 69: 1403–1435.
Güth, W., and Kocher, M. G. (2014) 'More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature', *Journal of Economic Behavior and Organization*, 108: 396–409.

Güth, W., Schmittberger, R., and Schwarze, B. (1982) 'An experimental analysis of ultimatum bargaining', *Journal of Economic Behavior and Organization*, 3: 367–388.
Güth, W., and Yaari, M. (1992) 'An evolutionary approach to explain reciprocal behavior in a simple strategic game', in Witt, U. (Ed.), *Explaining Process and Change: Approaches to Evolutionary Economics*: 23–34, Ann Arbor, MI: University of Michigan Press.
Hafenbrädl, S., Waeger, D., Marewski, J. N., and Gigerenzer, G. (2016) 'Applied decision making with fast-and-frugal heuristics', *Journal of Applied Research in Memory and Cognition*, 5: 215–231.
Halevy, Y. (2007) 'Ellsberg revisited: An experimental study', *Econometrica*, 75: 503–536.
Halpern, D. (2015) *Inside the Nudge Unit: How Small Changes Can Make a Big Difference*, London: W H Allen.
Harbaugh, W. T., Krause, K., and Liday, S. G. (2003) 'Bargaining by children', Economics Working Paper no. 2002–40, Eugene, OR: University of Oregon.
Hare, T., O'Doherty, J., Camerer, C., Schultz, W., and Rangel, A. (2008) 'Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors', *Journal of Neuroscience*, 28: 5623–5630.
Harrison, G. W., List, J., and Towe, C. (2007) 'Naturally occurring preferences and exogenous laboratory experiments: A case study of risk aversion', *Econometrica*, 75: 433–458.
Harrison, G. W., and Rutström, E. E. (2009) Expected utility theory and prospect theory: One wedding and a decent funeral. *Experimental Economics*, 12: 133–158.
Hashimzade, N., Myles, G. D., and Tran-Nam, B. (2013) 'Applications of behavioural economics to tax evasion', *Journal of Economic Surveys*, 27: 941–977.
Heath, C., and Soll, J. B. (1996) 'Mental budgeting and consumer decisions', *Journal of Consumer Research*, 23: 40–52.
Henker, J., and Owen, S. (2008) 'Bursting bubbles: Linking experimental financial market results to field market data', *Journal of Behavioral Finance*, 9: 5–14.
Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., and Gintis, H. (2004) *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, Oxford: Oxford University Press.
Herrmann, B., Thöni, C., and Gächter, S. (2008) 'Antisocial punishment across societies', *Science*, 319: 1362–1367.
Ho, T.-H., Camerer, C., and Weigelt, K. (1998) 'Iterated dominance and iterated best response in experimental "p-beauty contests"', *American Economic Review*, 88: 947–969.
Hoffman, E., McCabe, K., Shachat, K., and Smith, V. L. (1994) 'Preferences, property rights and anonymity in bargaining games', *Games and Economic Behavior*, 7: 346–380.
Hoffman, E., McCabe, K., and Smith, V. L. (1996) 'Social distance and other-regarding behavior', *American Economic Review*, 86: 653–660.
Hogarth, R. M., and Reder, M. W. (1987) *Rational Choice: The Contrast between Economics and Psychology*, Chicago: University of Chicago Press.
Holt, C., and Laury, S. (2002) 'Risk aversion and incentive effects', *American Economic Review*, 92: 1644–1655.
Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., and Camerer, C. (2005) 'Neural systems responding to degrees of uncertainty in human decision-making', *Science*, 310: 1680–1683.
Huber, J., and Kirchler, M. (2012) 'The impact of instructions and procedure on reducing confusion and bubbles in experimental asset markets', *Experimental Economics*, 15: 89–105.

Huberman, G., and Jiang, W. (2006) 'Offering versus choice in 401(k) plans: Equity exposure and number of funds', *Journal of Finance*, 61: 763–780.
Huck, S., Müller, W., and Normann, H. T. (2001) 'Stackelberg beats Cournot: On collusion and efficiency in experimental markets', *Economic Journal*, 111: 749–765.
——— (2002) 'To commit or not to commit: Endogenous timing in experimental duopoly markets', *Games and Economic Behavior*, 38: 240–264.
Huettel, S., Stowe, J., Gordon, E., Warner, B., and Platt, M. (2006) 'Neural signatures of economic preferences for risk and ambiguity', *Neuron*, 49: 765–775.
Hussam, R. N., Porter, D., and Smith, V. L. (2008) 'Thar she blows: Can bubbles be rekindled with experienced subjects?', *American Economic Review*, 93: 924–937.
Institute of Medicine (2001) *Crossing the Quality Chasm: A New Health System for the 21st Century*, Washington, DC: NAS Press.
Iyengar, S. S., and Lepper, M. R. (2000) 'When choice is demotivating: Can one desire too much of a good thing?', *Journal of Personality and Social Psychology*, 79: 995–1005.
Jehiel, P. (2005) 'Analogy-based expectation equilibrium', *Journal of Economic Theory*, 123: 81–104.
Jensen, K., Call, J., and Tomasello, M. (2007) 'Chimpanzees are rational maximizers in an ultimatum game', *Science*, 318: 107–109.
Jianakoplos, N., and Bernasek, A. (1998) 'Are women more risk averse?', *Economic Inquiry*, 36: 620–630.
Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006) 'Household expenditure and the income tax rebates of 2001', *American Economic Review*, 96: 1589–1610.
Kagel, J. H., and Levin, D. (2016) 'Auctions: A survey of experimental research', in Kagel, J. H., and Roth, A. (Eds.), *Handbook of Experimental Economics*, Vol. II: 563–637, Princeton, NJ: Princeton University Press.
Kagel, J. H., and Roth, A. E. (2000) 'The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment', *Quarterly Journal of Economics*, 115: 201–235.
Kahneman, D. (2003) 'Maps of bounded rationality: Psychology for behavioural economics', *American Economic Review*, 93: 1449–1475.
——— (2011) *Thinking, Fast and Slow*, London: Allen Lane.
Kahneman, D., and Dolan, P. (2008) 'Interpretations of utility and their implications for the valuation of health', *Economic Journal*, 118: 215–234.
Kahneman, D., Fredrickson, D. L., Schreiber, C. A., and Redelmeier, D. A. (1993) 'When more pain is preferred to less: Adding a better end', *Psychological Science*, 4: 401–405.
Kahneman, D., Knetsch, J., and Thaler, R. H. (1986) 'Fairness as a constraint on profit seeking: Entitlements in the market', *American Economic Review*, 76: 728–741.
Kahneman, D., and Krueger, A. B. (2006) 'Developments in the measurement of subjective well-being', *Journal of Economic Perspectives*, 20: 3–24.
Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., and Stone, A. A. (2004) 'A survey method for characterizing daily life experience: The day reconstruction method', *Science*, 306: 1776–1780.
Kahneman, D., and Tversky, A. (1973) 'On the psychology of prediction', *Psychological Review*, 80: 237–251.
——— (1979) 'Prospect theory: An analysis of decision under risk', *Econometrica*, 47: 263–291.
——— (1983) 'Choices, values, and frames', *American Psychologist*, 39: 341–350.

Kahneman, D., Wakker, P., and Sarin, R. (1997) 'Back to Bentham? Explorations of experienced utility', *Quarterly Journal of Economics*, 112: 375–405.
Kaiser Family Foundation (2008) '2008 update on consumers' views of patient safety and quality information', Report no. 7819, Menlo Park, CA: Kaiser Family Foundation.
Kalivas, P., and Volkow, N. (2005) 'The neural basis of addiction: A pathology of motivation and choice', *American Journal of Psychiatry*, 162: 1403–1413.
Kaplan, H., Hill, K., Lancaster, J., and Hurtado, A. M. (2000) 'A theory of human life history evolution: Diet, intelligence, and longevity', *Evolutionary Anthropology*, 9: 156–185.
Karlsson, N., Loewenstein, G., and Seppi, D. (2009) 'The "ostrich effect": Selective attention to information about investments', *Journal of Risk and Uncertainty*, 38: 95–115.
Kebede, B., Tarazona, M., Munro, A., and Verschoor, A. (2013) 'Intra-household efficiency: An experimental study from Ethiopia', *Journal of African Economies*, 23: 105–150.
Ketcham, J., Smith, V. L., and Williams, A. W. (1984) 'A comparison of posted-offer and double-auction pricing institutions', *Review of Economic Studies*, 51: 595–614.
Kilbanoff, P., Marinacci, M., and Mukerji, S. (2005) 'A smooth model of decision making under ambiguity', *Econometrica*, 73: 1849–1892.
Kirchler, M., Huber, J., and Stöckl, T. (2012) 'Thar she bursts: Reducing confusion reduces bubbles', *American Economic Review*, 102: 865–883.
Klemperer, P. (2004) *Auctions: Theory and Practice*, Princeton, NJ: Princeton University Press.
Knack, S., and Keefer, P. (1997) 'Does social capital have an economic payoff? A cross-country investigation', *Quarterly Journal of Economics*, 112: 1251–1288.
Knoch, D., Pascual-Leone, A., Meyer, K., Treyer, K., and Fehr, E. (2006) 'Diminishing reciprocal fairness by disrupting the right prefrontal cortex', *Science*, 314: 829–832.
Knutson, B., Fong, G., Adams, C., Varner, J., and Hommer, D. (2001) 'Dissociation of reward anticipation and outcome with event-related fMRI', *NeuroReport*, 12: 3683–3687.
Kocher, M. G., and Sutter, M. (2005) 'The decision maker matters: Individual versus group behaviour in experimental beauty-contest games', *Economic Journal*, 115: 200–223.
Kooreman, P. (2000) 'The labeling effect of a child benefit system', *American Economic Review*, 90: 571–583.
Kőszegi, B. (2014) 'Behavioral contract theory', *Journal of Economic Literature*, 52: 1075–1118.
Kőszegi, B., and Rabin, M. (2006) 'A model of reference-dependent preferences', *Quarterly Journal of Economics*, 121: 1133–1165.
—— (2007) 'Reference-dependent risk attitudes', *American Economic Review*, 97: 1047–1073.
Kube, S., Maréchal, M., and Puppe, C. (2006) 'Putting reciprocity to work: Positive versus negative responses in the field', Discussion Paper no. 2006–27, St. Gallen: Department of Economics, University of St. Gallen.
Kugler, T., Bornstein, G., Kocher, M. G., and Sutter, M. (2007) 'Trust between individuals and groups: Groups are less trusting than individuals but just as trustworthy', *Journal of Economic Psychology*, 28: 646–657.
Laibson, D. (1997) 'Golden eggs and hyperbolic discounting', *Quarterly Journal of Economics*, 112: 443–447.

— (2001) 'A cue-theory of consumption', *Quarterly Journal of Economics*, 116: 81–119.
Laibson, D., Repetto, A., and Tobacman, J. (2007) 'Estimating discount functions with consumption choice over the lifecycle', Working Paper no. 13314, Cambridge, MA: NBER.
Laland, K., and Brown, G. (2005) *Sense and Nonsense: Evolutionary Perspectives on Human Behaviour*, Oxford: Oxford University Press.
Landsberger, M. (1966) 'Windfall income and consumption: Comment', *American Economic Review*, 56: 534–539.
Larrick, R. P., and Blount, S. (1997) 'The claiming effect: Why players are more generous in social dilemmas than in ultimatum games', *Journal of Personality and Social Psychology*, 72: 810–825.
Larrick, R. P., and Soll, J. B. (2008) 'The MPG illusion', *Science*, 320: 1593–1594.
Levitt, S., and List, J. (2007) 'What do laboratory experiments measuring social preferences reveal about the real world?', *Journal of Economic Perspectives*, 21: 153–174.
Lichtenstein, S., and Slovic, P. (1971) 'Reversals of preference between bids and choices in gambling decisions', *Journal of Experimental Psychology*, 89: 46–55.
Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., and Combs, B. (1978) 'Judged frequency of lethal events', *Journal of Experimental Psychology: Human Learning and Memory*, 4: 551–578.
List, J. (2004) 'Neoclassical theory versus prospect theory: Evidence from the marketplace', *Econometrica*, 72: 615–625.
—— (2007) 'On the interpretation of giving in dictator games', *Journal of Political Economy*, 115: 482–493.
Loewenstein, G., and Adler, D. (1995) 'A bias in the prediction of tastes', *Economic Journal*, 105: 929–937.
Loewenstein, G., O'Donoghue, T., and Rabin, M. (2003) 'Projection bias in predicting future utility', *Quarterly Journal of Economics*, 118: 1209–1248.
Loewenstein, G., and Prelec, D. (1992) 'Anomalies in intertemporal choice: Evidence and an interpretation', *Quarterly Journal of Economics*, 107: 573–597.
—— (1993) 'Preferences over sequences of outcomes', *Psychological Review*, 100: 91–108.
Loewenstein, G., and Ubel, P. A. (2008) 'Hedonic adaption and the role of decision and experience utility in public policy', *Journal of Public Economics*, 92: 1795–1810.
Loomes, G., Starmer, C., and Sugden, R. (1991) 'Observing violations of transitivity by experimental methods', *Econometrica*, 59: 425–439.
—— (2003) 'Do anomalies disappear in repeated markets?', *Economic Journal*, 113: C153–C166.
Loomes, G., and Sugden, R. (1982) 'An alternative theory of choice under uncertainty', *Economic Journal*, 92: 805–824.
—— (1983) 'A rationale for preference reversal', *American Economic Review*, 73: 428–432.
—— (1986) 'Disappointment and dynamic consistency in choice under uncertainty', *Review of Economic Studies*, 53: 271–282.
Lord, C. G., Ross, L., and Lepper, M. R. (1979) 'Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence', *Journal of Personality and Social Psychology*, 37: 2098–2109.
Lucking-Reiley, D. (1999) 'Using field experiments to test equivalence between auction formats: Magic on the internet', *American Economic Review*, 89: 1063–1080.
Luhan, W. J., Kocher, M. G., and Sutter, M. (2009) 'Group polarization in the team dictator game reconsidered', *Experimental Economics*, 12: 26–41.

McCall, T. B. (1996) *Examining Your Doctor: A Patient's Guide to Avoiding Harmful Medical Care*, New York: Carol Publishing.
McClure, S., Laibson, D., Loewenstein, G., and Cohen, J. (2004) 'Separate neural systems value immediate and delayed monetary rewards', *Science*, 306: 503–507.
McKelvey, R., and Palfrey, T. (1995) 'Quantal response equilibria for normal form games', *Games and Economic Behavior*, 10: 6–38.
McKelvey, R., Palfrey, T., and Weber, R. A. (2000) 'The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria', *Journal of Economic Behavior and Organization*, 42: 523–548.
Madrian, B. C., and Shea, D. F. (2001) 'The power of suggestion: Inertia in 401(k) participation and savings behavior', *Quarterly Journal of Economics*, 116: 1149–1187.
Mani, A., Mullainathan, S., Shafir, E., and Zhao, J. (2013) 'Poverty impedes cognitive function', *Science*, 341: 976–980.
Marks, M. B., and Croson, R. T. (1998) 'Alternative rebate rules in the provision of a threshold public good: An experimental investigation', *Journal of Public Economics*, 67: 195–220.
Mazar, N., Amir, O., and Ariely, D. (2008) 'The dishonesty of honest people: A theory of self-concept maintenance', *Journal of Marketing Research*, 45: 633–644.
Mehra, R., and Prescott, E. (2003) 'The equity premium puzzle in retrospect', in Constantinides, G. M., Harris, M., and Stulz, R. (Eds.), *Handbook of the Economics of Finance*: 889–938, Amsterdam: North-Holland.
Mehta, J., Starmer, C., and Sugden, R. (1994) 'The nature of salience: An experimental investigation of pure coordination games', *American Economic Review*, 84: 658–673.
Miguel, E., and Kremer, M. (2004) 'Worms: Identifying impacts on education and health in the presence of treatment externalities', *Econometrica*, 72: 159–217.
Milgrom, P. (2004) *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.
Miravete, E. (2003) 'Choosing the wrong calling plan? Ignorance and learning', *American Economic Review*, 93: 297–310.
Mullainathan, S., and Thaler, R. H. (2000) 'Behavioral economics', in *International Encyclopedia of the Social and Behavioral Sciences*: 1094–1100, Cambridge, MA: MIT Press.
Nagel, R. (1995) 'Unraveling in guessing games: An experimental study', *American Economic Review*, 85: 1313–1326.
National Drug Intelligence Center (2010) *National Drug Threat Assessment 2010*, Johnstown, PA: NDIC.
Neumark, D., and Postlewaite, A. (1998) 'Relative income concerns and the rise in married women's employment', *Journal of Public Economics*, 70: 157–183.
Niederle, M., and Vesterlund, L. (2007) 'Do women shy away from competition? Do men compete too much?', *Quarterly Journal of Economics*, 122: 1067–1101.
Northcraft, G. B., and Neale, M. A. (1987) 'Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions', *Organizational Behavior and Human Decision Processes*, 39: 84–97.
Nowak, M. (2006) 'Five rules for the evolution of cooperation', *Science*, 314: 1569–1572.
Odean, T. (1998) 'Are investors reluctant to realize their losses?', *Journal of Finance*, 53: 1775–1798.
O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, R., and Dolan, R. J. (2004) 'Dissociable roles of ventral and dorsal striatum in instrumental conditioning', *Science*, 304: 452–454.

O'Doherty, J., Deichmann, R., Critchley, H., and Dolan, R. J. (2002) 'Neural responses during anticipation of a primary taste reward', *Neuron*, 33: 815–826.
O'Doherty, J., Kringelbach, M., Rolls, E., Hornak, J., and Andrews, C. (2001) 'Abstract reward and punishment representations in the human orbitofrontal cortex', *Nature Neuroscience*, 4: 95–102.
O'Donoghue, T., and Rabin, M. (1999) 'Doing it now or later', *American Economic Review*, 89: 103–124.
—— (2000) 'The economics of immediate gratification', *Journal of Behavioral Decision Making*, 13: 233–250.
—— (2006) 'Incentives and self-control', in Blundell, R., Newey, W., and Persson, R. (Eds.), *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Vol. II: 215–245, New York: Cambridge University Press.
Office for National Statistics (2004) *Smoking-Related Behaviour and Attitudes, 2003*, London: HMSO.
Ok, E. A., Ortoleva, P., and Riella, G. (2014) 'Revealed (p)reference theory', *American Economic Review*, 105: 299–321.
Olson, M. (1971) *The Logic of Collective Action: Public Goods and the Theory of Groups*, Cambridge, MA: Harvard University Press.
Oskarsson, A. T., Van Boven, L., McClelland, G., and Hastie, R. (2009) 'What's next? Judging sequences of binary events', *Psychological Bulletin*, 135: 262–285.
Ostrom, E. (1990) *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge: Cambridge University Press.
—— (2006) 'The value-added of laboratory experiments for the study of institutions and common-pool resources', *Journal of Economic Behavior and Organization*, 61: 149–163.
—— (2010) 'Beyond markets and states: Polycentric governance of complex economic systems', *American Economic Review*, 100: 641–672.
Ostrom, E., Walker, J., and Gardner, R. (1992) 'Covenants with and without a sword: Self-governance is possible', *American Political Science Review*, 86: 404–417.
Oswald, A., and Powdthavee, N. (2008) 'Does happiness adapt? A longitudinal study of disability with implications for economists and judges', *Journal of Public Economics*, 92: 1061–1077.
Pascual-Ezama, D., Fosgaard, T. R., Cardenas, J. C., Kujal, P., Veszteg, R., de Liaño, B. G. G., Gunia, B., Weichselbaumer, D., Hilken, K., Antinyan, A., and Delnoij, J. (2015) 'Context-dependent cheating: Experimental evidence from 16 countries', *Journal of Economic Behavior and Organization*, 116: 379–386.
Patall, E. A., Cooper, H., and Robinson, J. C. (2008) 'The effects of choice on intrinsic motivation and related outcomes: A meta-analysis of research findings', *Psychological Bulletin*, 134: 270–300.
Phelps, C. E., and Mooney, C. (1993) 'Variations in medical practice: Causes and consequences', in Arnould, R., Rich, R., and White, W. (Eds.), *Competitive Approaches to Health Care Reform*: 139–178, Washington, DC: Urban Institute Press.
Phillips, C., and Thompson, G. (2009) 'What is a QALY?', Newmarket, UK: Hayward Medical Communications.
Porter, D. P., and Smith, V. L. (2003) 'Stock market bubbles in the laboratory', *Journal of Behavioral Finance*, 4: 7–20.
Post, T., van den Assem, M. J., Baltussen, G., and Thaler, R. H. (2008) 'Deal or no deal? Decision making under risk in a large-payoff game show', *American Economic Review*, 98: 38–71.

Pradel, J., Euler, H., and Fetchenhauer, D. (2009) 'Spotting altruistic dictator game players and mingling with them: The elective assortation of classmates', *Evolution and Human Behavior*, 30: 103–113.
Pratt, J. W., Wise, D. A., and Zeckhauser, R. (1979) 'Price differences in almost competitive markets', *Quarterly Journal of Economics*, 93: 189–211.
Preuschoff, K., Bossaerts, P., and Quartz, S. (2006) 'Neural differentiation of expected reward and risk in human subcortical structures', *Neuron*, 51: 381–390.
Proctor, D., Williamson, R. A., de Waal, F. B., and Brosnan, S. F. (2013) 'Chimpanzees play the ultimatum game', *Proceedings of the National Academy of Sciences*, 110: 2070–2075.
Purves, D., Brannon, E., Cabeza, R., Huettel, S., LaBar, K., Platt, M., and Woldorff, M. (2008) *Principles of Cognitive Neuroscience*, Sunderland, MA: Sinauer Associates.
Rabin, M. (1993) 'Incorporating fairness into game theory and economics', *American Economic Review*, 83: 1281–1302.
——— (1998) 'Psychology and economics', *Journal of Economic Literature*, 36: 11–46.
——— (2000) 'Risk aversion and expected-utility theory: A calibration theorem', *Econometrica*, 68: 1281–1292.
——— (2002) 'Inferences by believers in the law of small numbers', *Quarterly Journal of Economics*, 117: 775–816.
Rabin, M., and Schrag, J. (1999) 'First impressions matter: A model of confirmatory bias', *Quarterly Journal of Economics*, 114: 37–82.
Rapoport, A., and Budescu, D. V. (1997) 'Randomization in individual choice behavior', *Psychological Review*, 104: 603–617.
Rapoport, A., Seale, D. A., and Winter, E. (2002) 'Coordination and learning behavior in large groups with asymmetric players', *Games and Economic Behavior*, 39: 111–136.
Read, D., and van Leeuwen, B. (1998) 'Predicting hunger: The effects of appetite and delay on choice', *Organizational Behaviour and Human Decision Processes*, 76: 189–205.
Redelmeier, D. and Kahneman, D. (1996) 'Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures', *Pain*, 66: 3–8.
Redelmeier, D., Katz, J., and Kahneman, D., (2003) 'Memories of colonoscopy: A randomized trial', *Pain*, 104: 187–194.
Ribar, D., and Wilhelm, M. (2002) 'Altruistic and joy-of-giving motivations in charitable behavior', *Journal of Political Economy*, 110: 425–457.
Ridley, M. (1997) *The Origins of Virtue*, London: Penguin Books.
Robson, A. (2001a) 'The biological basis of economic behavior', *Journal of Economic Literature*, 39: 11–33.
——— (2001b) 'Why would nature give individuals utility functions?', *Journal of Political Economy*, 109: 900–914.
——— (2002) 'Evolution and human nature', *Journal of Economic Perspectives*, 16: 89–106.
Robson, A., and Kaplan, H. (2003) 'The evolution of human life expectancy and intelligence in hunter-gatherer economics', *American Economic Review*, 93: 150–169.
Robson, A., and Samuelson, L. (2009) 'The evolution of time preference with aggregate uncertainty', *American Economic Review*, 99: 1925–1953.
Rodrik, D. (2008) 'The new development economics: We shall experiment, but how shall we learn?', Research Working Paper no. 08–055, Cambridge, MA: John F. Kennedy School of Government, Harvard University.

Rosenbaum, S. M., Billinger, S., and Stieglitz, N. (2014) 'Let's be honest: A review of experimental evidence of honesty and truth-telling', *Journal of Economic Psychology*, 45: 181–196.
Roth, A. (1995) 'Introduction to experimental economics', in Kagel, J., and Roth, A. (Eds.), *Handbook of Experimental Economics*: 3–20, Princeton, NJ: Princeton University Press.
—— (2002) 'The economist as engineer: Game theory, experimentation, and computation as tools for design economics', *Econometrica*, 70: 1341–1378.
—— (2008) 'Deferred acceptance algorithms: History, theory, practice, and open questions', *International Journal of Game Theory*, 36: 537–569.
Roth, A., and Erev, I. (1995) 'Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term', *Games and Economic Behavior*, 8: 164–212.
Roth, A., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991) 'Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study', *American Economic Review*, 81: 1068–1095.
Samuelson, L. (2006) 'Economic theory and experimental economics', *Journal of Economic Literature*, 43, 65–107.
Sanfey, A., Rilling, J., Aronson, J., Nystrom, L., and Cohen, J. (2003) 'The neural basis of economic decision-making in the ultimatum game', *Science*, 300: 1755–1759.
Schelling, T. (1984) 'Self-command in practice, in policy, and in a theory of rational choice', *American Economic Review*, 74: 1–11.
—— (1990) *The Strategy of Conflict*, Cambridge, MA: Harvard University Press.
Schmidt, U., Starmer, C., and Sugden, R. (2008) 'Third generation prospect theory', *Journal of Risk and Uncertainty*, 36: 203–223.
Schmitt, P., Swope, K., and Walker, J. (2000) 'Collective action with incomplete commitment: Experimental evidence', *Southern Economic Journal*, 66: 829–854.
Schunk, D. (2009) 'Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence', *Journal of Dynamics and Control*, 33: 1719–1738.
Selten, R. (1994) 'Reinhard Selten: Biographical', Nobel Prize Foundation, www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1994/selten-bio.html.
Selten, R., Abbink, K., and Cox, R. (2005) 'Learning direction theory and the winner's curse', *Experimental Economics*, 8: 5–20.
Selten, R., and Chmura, T. (2008) 'Stationary concepts for experimental 2x2 games', *American Economic Review*, 98: 938–966.
Sent, E.-M. (2004) 'Behavioral economics: How psychology made its (limited) way back into economics', *History of Political Economy*, 36: 735–760.
Shafir, E., Simonson, I., and Tversky, A. (1993) 'Reason-based choice', *Cognition*, 49: 11–36.
Shafir, E., and Thaler, R. H. (2006) 'Invest now, drink later, spend never: On the mental accounting of delayed consumption', *Journal of Economic Psychology*, 27: 694–712.
Shang, J., and Croson, R. T. (2009) 'A field experiment in charitable contribution: The impact of social information on the voluntary provision of public goods', *Economic Journal*, 119: 1422–1439.
Shapiro, M. D., and Slemrod, J. (2003) 'Consumer response to tax rebates', *American Economic Review*, 93: 381–396.
Shefrin, H. M., and Thaler, R. H. (1988) 'The behavioural life-cycle hypothesis', *Economic Inquiry*, 26: 609–643.
Shiller, R. J. (2003) 'From efficient markets theory to behavioural finance', *Journal of Economic Perspectives*, 17: 83–104.

Shiv, B., and Fedorikhin, A. (1999) 'Heart and mind in conflict: The interplay of affect and cognition in consumer decision making', *Journal of Consumer Research*, 26: 278–292.
Shogren, J. F., and Taylor, L. O. (2008) 'On behavioral-environmental economics', *Review of Environmental Economics and Policy*, 2: 26–44.
Siemens, J. C. (2007) 'When consumption benefits precede costs: Towards an understanding of "buy now, pay later" transactions', *Journal of Behavioral Decision Making*, 20: 521–531.
Simon, H. (1955) 'A behavioural model of rational choice', *Quarterly Journal of Economics*, 69: 99–118.
—— (1991) *Models of My Life*, London: Basic Books.
Simonsohn, U., and Loewenstein, G. (2006) 'Mistake #37: The effect of previously encountered prices on current housing demand', *Economic Journal*, 116: 175–199.
Simonson, I., and Tversky, A. (1992) 'Choice in context: Tradeoff contrast and extremeness aversion', *Journal of Marketing Research*, 29: 281–295.
Slonim, R., and Roth, A. (1998) 'Learning in high stakes ultimatum games: An experiment in the Slovak Republic', *Econometrica*, 66: 569–596.
Smith, D., Loewenstein, G., Jepson, C., Jankovich, A., Feldman, H., and Ubel, P. (2008) 'Mispredicting and misremembering: Patients with renal failure overestimate improvements in quality of life after a kidney transplant', *Health Psychology*, 27: 653–658.
Smith, V. L. (1962) 'An experimental study of competitive market behavior', *Journal of Political Economy*, 70: 111–137.
—— (1991) 'Rational choice: The contrast between economics and psychology', *Journal of Political Economy*, 99: 877–897.
—— (1998) 'Experimental economics: Behavioral lessons for microeconomics theory and policy', in Jacobs, D., Kalai, E., Kamien, M., and Schwartz, N. (Eds.), *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*: 104–122, Cambridge: Cambridge University Press.
—— (2002) 'Markets, institutions and experiments', in Nadel, L. (Ed.), *Encyclopedia of Cognitive Science*: 991–998, London: John Wiley.
—— (2003) 'Constructivist and ecological rationality in economics', *American Economic Review*, 93: 465–508.
—— (2007) *Rationality in Economics: Constructivist and Ecological Forms*, Cambridge: Cambridge University Press.
Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988) 'Bubbles, crashes, and endogenous expectations in experimental spot asset markets', *Econometrica*, 56: 1119–1151.
Sobel, J. (2002) 'Can we trust social capital?', *Journal of Economic Literature*, 40: 139–154.
—— (2005) 'Interdependent preferences and reciprocity', *Journal of Economic Literature*, 43: 392–436.
Solnick, S. (2001) 'Gender differences in the ultimatum game', *Economic Inquiry*, 39: 189–200.
Spiegler, R. (2014) *Bounded Rationality and Industrial Organization*, New York: Oxford University Press.
Spitzer, M., Fischbacher, U., Herrnberger, B., Grön, G., and Fehr, E. (2007) 'The neural signature of social norm compliance', *Neuron*, 56: 185–196.
Sprenger, C. (2015) 'An endowment effect for risk: Experimental tests of stochastic reference points', *Journal of Political Economy*, 123: 1456–1499.

Starmer, C. (2000) 'Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk', *Journal of Economic Literature*, 38: 332–382.
Stern, N. (2006) *Review on the Economics of Climate Change: The Stern Review*, Cambridge: Cambridge University Press.
Strotz, R. H. (1956) 'Myopia and inconsistency in dynamic utility maximization', *Review of Economic Studies*, 23: 165–180.
Sugden, R. (1995) 'A theory of focal points', *Economic Journal*, 105: 533–550.
Sutter, M. (2005) 'Are four heads better than two? An experimental beauty-contest game with teams of different size', *Economics Letters*, 88: 41–46.
——— (2009) 'Deception through telling the truth?! Experimental evidence from individuals and teams', *Economic Journal*, 119: 47–60.
Sutter, M., Czermak, S., and Feri, F. (2013) 'Strategic sophistication of individuals and teams: Experimental evidence', *European Economic Review*, 64: 395–410.
Svorenčík, A., and Maas, H. (Eds.) (2015) *The Making of Experimental Economics: Witness Seminar on the Emergence of a Field*, Heidelberg: Springer.
Sydnor, J. (2010) '(Over)insuring modest risks', *American Economic Journal: Applied Economics*, 2: 177–199.
Terrell, D. (1994) 'A test of the gambler's fallacy: Evidence from pari-mutuel games', *Journal of Risk and Uncertainty*, 8: 309–317.
Thaler, R. H. (1981) 'Some empirical evidence on dynamic inconsistency', *Economics Letters*, 8: 201–207.
——— (1988) 'Anomalies: The ultimatum game', *Journal of Economic Perspectives*, 2: 195–206.
——— (1990) 'Anomalies: Saving, fungibility, and mental accounts', *Journal of Economic Perspectives*, 4: 193–205.
——— (2008a) 'Mental accounting and consumer choice', *Marketing Science*, 27: 15–25.
——— (2008b) 'Mental accounting and consumer choice: Anatomy of a failure', *Marketing Science*, 27: 12–14.
Thaler, R. H., and Benartzi, S. (2004) 'Save more tomorrow™: Using behavioral economics to increase employee saving', *Journal of Political Economy*, 112: S164–S187.
Thaler, R. H., and Shefrin, H. M. (1981) 'An economic theory of self-control', *Journal of Political Economy*, 89: 392–406.
Thaler, R. H., and Sunstein, C. R. (2008) *Nudge: Improving Decisions about Health, Wealth, and Happiness*, New Haven, CT: Yale University Press.
Todorov, A., Mandisodza, A. N., Goren, A., and Hall, C. C. (2005) 'Inferences of competence from faces predict election outcomes', *Science*, 308: 1623–1626.
Tricomi, E., Rangel, A., Camerer, C., and O'Doherty, J. (2010) 'Neural evidence for inequality-averse social preferences', *Nature*, 463: 1089–1092.
Tversky, A., and Kahneman, D. (1974) 'Judgment under uncertainty: Heuristics and biases', *Science*, 185: 1124–1131.
——— (1981) 'The framing of decisions and the psychology of choice', *Science*, 211: 453–458.
——— (1983) 'Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment', *Psychological Review*, 90: 293–315.
——— (1986) 'Rational choice and the framing of decisions', *Journal of Business*, 59: S251–S278.
——— (1992) 'Advances in prospect theory: Cumulative representation of uncertainty', *Journal of Risk and Uncertainty*, 5: 297–323.
Tversky, A., Kahneman, D., and Choice, R. (1981) 'The framing of decisions', *Science*, 211: 453–458.

Tversky, A., and Shafir, E. (1992) 'Choice under conflict: The dynamics of deferred decision', *Psychological Science*, 3: 358–361.
Tversky, A., Slovic, P., and Kahneman, D. (1990) 'The causes of preference reversal', *American Economic Review*, 80: 204–217.
Unver, U. (2005) 'On the survival of some unstable two-sided matching mechanisms', *International Journal of Game Theory*, 33: 239–254.
Van Damme, E., Binmore, K., Roth, A., Samuelson, L., Winter, E., Bolton, G. E., Ockenfels, A., Dufwenberg, M., Kirchsteiger, G., Gneezy, U., and Kocher, M. G. (2014) 'How Werner Güth's ultimatum game shaped our understanding of social behavior', *Journal of Economic Behavior and Organization*, 108: 292–318.
Van Huyck, J., Battalio, R., and Beil, R. (1990) 'Tacit coordination games, strategic uncertainty, and coordination failure', *American Economic Review*, 80: 234–248.
Venkatraman, V., Payne, J., Bettman, J., Frances Luce, M., and Huettel, S. (2009) 'Separate neural mechanisms underlie choices and strategic preferences in risky decision making', *Neuron*, 62: 593–602.
Walker, M., and Wooders, J. (2001) 'Minimax play at Wimbledon', *American Economic Review*, 91: 1521–1538.
Wang, J. T. Y., Spezio, M., and Camerer, C. (2010) 'Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender–receiver games', *American Economic Review*, 100: 984–1007.
Warner, J., and Pleeter, S. (2001) 'The personal discount rate: Evidence from military downsizing programs', *American Economic Review*, 91: 33–53.
Weber, R. A. (2003) '"Learning" with no feedback in a competitive guessing game', *Games and Economic Behavior*, 44: 134–144.
Webley, P. (2004) 'Children's understanding of economics', in Barrett, M., and Buchanan-Barrow, E. (Eds.), *Children's Understanding of Society*: 43–68, Hove, UK: Psychology Press.
Weizsäcker, G. (2010) 'Do we follow others when we should? A simple test of rational expectations', *American Economic Review*, 100: 2340–2360.
Wirtz, D., Kruger, J., Scollon, C. N., and Diener, E. (2003) 'What to do on spring break? The role of predicted, on-line, and remembered experience in future choice', *Psychological Science*, 14: 520–524.
World Economic Forum (2016) *The Global Gender Gap Report 2016*, Geneva: World Economic Forum.
World Health Organization (2009) *Global Health Risks: Mortality and Burden of Disease Attributable to Selected Major Risks*, Geneva: WHO Press.

Index

absolute magnitude effect 166, 179
absolute risk aversion 100
acquisition utility 52, 433, 450
action sampling 280–281
actor–critic model 434
adaption 471–473, 477–478, 490–492
adaptive learning 273–275
addiction 449, 451–460
adversarial collaboration 50–51
affective processes 421–422, 424
aggregate risk 377–380, 389
Allais Paradox 102–105, 108–112, 114–116, 136
altruism 335, 354–358, 391
ambiguity 158–159, 204, 207–212, 231–232, 247, 431–433, 435
amygdala 420–421, 424, 431, 437–438
anchoring effect 42–44
anticipated utility 182–183
anti-social punishment 388
artefactual field experiments 302–303
aspiration level 33–34, 273–274
aspiration treadmill 472
asset account 87
assortative matching 382–383, 390, 403
attitudes to competition 399–401
attitudes to risk 98–100, 105–106, 184–189, 398–399
attractiveness of a strategy 272
auctions 69–72, 290–296, 319, 512–517
automatic processes 421–422, 438, 445
availability heuristic 233–234, 244
axioms 177–178

backward induction 267–268
Bayes updating 206–208, 214–216, 221, 223–234, 231
beauty contest 250–262, 265, 269, 273–279, 285–287, 319, 438–439
behavior change 202, 524–526
behavioral finance 12–13, 26
belief based learning 274–278, 281, 287
best reply 274
beta-delta preferences 172–176, 445–448
borrowing 84–86, 190–192, 410–413
brainstem 418
Brodmann area 418–419
bubble 236–244
buy–sell imbalance 244–245

calling plan 197–198
capital punishment 212–213
caudate 420, 424
central nervous system 418–419
cerebrum 418–421, 424
charity 354–359
children 85–86, 89–90, 198–201, 224, 246–247, 383–384, 393–396, 399–400, 403–410, 413
chimpanzee 393–396
choice arbitrariness 33, 38–44, 52, 54, 190, 290
choice architect 521–522
choice bracketing 55–57
choice overload 486–487, 523
choose–reject discrepancy 42
chopstick auction 513–516
climate policy ramp 199–201
cognitive bias 11–12, 212–221, 316
cognitive hierarchy 258–260
collectively rational recommendation (CRR) 263–264
combinatorial auction 516
commitment 176–179, 449, 480–485, 518–525
common pool resource 505–506
common value auction 293
communication 302–303, 504, 507–508
compound lottery 209–211
confirmatory bias 212–216, 221, 230–231, 232–234, 244, 247, 454
conflict of interest 303
conflicting choice 38–39
conformist bias 393
conformity 226–227, 230, 234, 247, 286
constant absolute risk aversion (CARA) 100
constant relative risk aversion (CRRA) 100, 144

context effect 10, 44–47, 65–66, 95, 131, 135–136, 164–166, 190, 203, 437, 523–524
contingent valuation 49
contract theory 365–368
contrast effect 472
controlled processes 421–422, 447–450
coordination game 261–265, 296–309, 314–315, 319
Cournot 274–276, 316–318
crash 236–245
cross-culture comparison 385–389
crowding out 354–356
cue-elicited behavior 457–459
cultural transmission 392–393, 397, 405
current income account 87–88

day reconstruction method 466–467
deception 21, 59–64, 95, 227, 351
decision utility 474–476, 482–484, 486
decision weights 113–114, 151, 157, 189
deductible 144–146
default 193, 422–423, 494, 519–528
default mode 422–423
deferred acceptance algorithm 510–512
delay-speed up asymmetry 166–167, 179–181, 190
demand 3, 68–75, 78, 82–83, 236–245, 359–365, 452–454, 498
dictator game 60, 322–325, 327–330, 335–336, 339–340, 342, 352, 357, 383–384, 395–396, 442
directed cognition 33–36
disability 471
disappointment 110–113, 118–119, 136, 150, 351
discovered preference hypothesis 15
disposition effect 139–142
distribution games 345
diversification 92–96, 122, 187, 378
doer 449–450
dominates 252
dopamine 424, 459
dorsolateral prefrontal cortex (DLPFC) 423, 440–443, 445–446
double auction 69–73, 237–241, 505, 508–509
double blind 323–324, 337, 491–492
dual-self model 449–450
duration neglect 477, 486, 490–491
Dutch auction 291–293

economic psychology 12, 14
editing phase 116, 130
education production function 407–410
efficient market hypothesis 244
elation 111–112, 118
electro-encephalogram (EEG) 416–417
elimination by aspects 33–37
Ellsberg paradox 210–212
empirical optimum 224–225
endogenous timing 317–318
endowment effect 47–51, 65–66, 134, 154, 247, 477–479
English auction 291–293
equilibrium refinement 265–269, 270, 280, 319, 382
equilibrium selection 267–268
equity premium puzzle 136–139, 144
equity, reciprocity and competition (ERC) 338–345, 353
evaluation period 137–144
evaluation phase 116, 130
evolutionary game theory 381–382
evolutionary stable strategy (ESS) 382
executive control system 422–424, 429, 437–439, 441, 445–447, 450
expected utility 97–108, 112–114, 116–124, 135–136, 144, 148–152, 184–188, 208–211, 375
expected value 99–100, 375
experience weighted learning (EWA) 276, 278–279, 287
experienced subjects 241–242
experimental economics 12, 17
exponential discounting 162–171, 184, 189–190, 376–380, 396, 445
external validity 17, 154–155, 299
extremeness aversion 40–44
eye tracking 61, 258

fair reference point 349–353
fairness 52, 334, 336, 346–353, 359–361, 364–365, 369, 440–443, 450
fairness equilibrium 348
family 80–81, 232–233, 403–407, 413
fanning out 108–414
feedback theory 244
Fehr-Schmidt model 341–346
fictitious play 275–277
field experiments 17–18, 293, 408, 518
first fundamental theorem of welfare economics 69–71
first price auction 292–294

focal point 260–265
food 52–53, 373–384, 396, 404–405, 427, 431–432, 480–481
foregone gain 50
forward induction 267–268, 290
forward looking behavior 454–455
fourfold pattern of risk attitudes 155–160
fragile cascades 223
framing effect 45–49, 64, 84, 116, 120–123, 128, 130, 152–154, 159, 202–203, 240, 353–354, 422, 437–489, 497–498, 521–526
free-riding 390–392
functional magnetic resonance imaging (fMRI) 416–417
fungibility 86–90
future income account 87–88

gain–loss asymmetry 166, 179
gambler's fallacy 216–219, 247, 296
game method 328
game theory 8–9, 12, 22–23, 23–26, 251–253, 351, 365, 381, 517
gender gap 397–403, 406
gene–culture co-evolution 392
gift exchange game 322, 325–326, 331, 335, 362–363
group selection 389–392
growth 191–192
gym membership 193–197

habit formation 190–192, 492–494
habituation 471–473, 489, 492
happiness 10, 22–23, 32, 397, 463–480, 487–490, 497
hawk–dove game 277, 381–382
haystack model 389–392
health care 231–235, 487–492, 522
hedonic editing 55–57, 66, 83, 123
Homo economicus 3–14, 44–45, 57–58, 289, 321, 517–526
hot hand fallacy 217–219
household bargaining 283–285, 405–407, 413
housing 81–84, 236
hyperbolic discounting 170–178, 190, 377–379, 445–448

imitation 326
impact bias 477
impact effect 143
impulse balance 281–282, 294
impure altruism 354–356
independence 108–110, 126–127, 135, 178
indifference curve 109–113
individual fishing quotas 506
inequality aversion 338–347, 350–354, 395, 443–445
infinite regress problem 11, 35
inflation 309–310, 495–497
information cascades 26, 223–226
inhibition 423, 437–438
instant utility 474–478, 486
institutions 3, 8, 12, 69–75, 239, 303–305, 353, 388–389, 397, 504–517
insula 420–421, 436–437, 440–443
insurance 144–146
intellective task 286
intentional kindness 350
intentions 333, 346–354, 358, 360–361
intergenerational discount factor 199–201
interpreter 415–416
inter-temporal utility 163–181, 184, 201, 376
intuition 35, 45–47, 50, 52–54, 57, 65–66, 245, 421–422, 449, 524
intuitive equilibrium 229
investment funds 91–95, 139
Iowa gambling task 429
iterated deletion of dominated strategies 252–254, 288

judgmental task 286

Kahneman, D. 7–8, 11, 15, 51, 116, 263
kindness 347–350, 353, 360–361

labor market 75–81, 228, 358–363, 397–398
latent class approach 276–277
law of one price 73–75
law of small numbers 216–221, 230, 232–234, 244, 454
learning direction theory 295
level-k thinking 256–262, 270, 286, 290, 294–295, 439
life satisfaction 465–473, 495–496
life cycle hypothesis 84–91, 190–193
limit pricing 311–313
linear public good game 322, 330–333, 335, 340–344, 390–392
litigation 155–159
loss aversion 47–51, 65–66, 82–83, 116–118, 126, 137–140, 144–145, 160, 166, 179–181, 281–282, 363, 366–368, 402
lottery numbers 216–217

marginal propensity to consume 87–90
market entry 314–316
market equilibrium 68–75
markets 8–9, 26, 67–75, 236–245, 508–512
marriage 471–472
matching algorithm 510–512
matching markets 509–512, 516–517
maximin 345–346
maximum likelihood estimation 276–277
mental accounting 45–57, 86–91, 123, 130, 166
meta-data analysis 226
microfinance 410–413
MINDSPACE 524–525, 529
minimum effort game 297–303
modularity 415–417, 450
monetary policy 309–310, 495–496
moonlighting game 322, 333, 335, 347
moral hazard 365
MouseLab 35, 258
multi-level selection 389–393
multi-unit auction 513
multiple-self models 449–450

naïve 174–176, 190, 194–197, 482–483, 519, 526
narrow framing 52–56, 86, 123
Nash equilibrium 23–24, 224–225, 229, 251–259, 261, 265–271, 280, 288–290, 296–299, 331–332, 381–382
natural selection 373–374
negotiated price 69, 73–75, 504, 508–509
news 142–144, 233–34, 244–245
no loss in buying hypothesis 50–51, 119, 145
nucleus accumbens (NAcc) 424, 425–427, 429–435, 444–447
nudge 494, 517, 521–528

one over n heuristic 92–95
optimism 114, 119, 136
opt-out versus opt-in 193, 494, 519, 527
ostrich effect 142–144
Ostrom, E. 506
Overconfidence 214, 221, 368, 402–403

Pareto dominant 265–266, 268, 290, 297–279
Pareto efficient 23–24, 69, 330–331, 349, 505
Pareto, V. 5–6
payoff function 22–23
payoff sampling 280–282
peak-end evaluation 476–477, 490
perception 8, 45–47, 50–57, 65–66, 121, 346, 421, 449
personal equilibrium 121
pessimism 114, 119, 136
planner 449–450
polarization of attitudes 213, 221
pooling 229, 312
positive economics 13, 16
positive self-concept 64
positron emission topography (PET) 416–417
posted offer 69, 72–73, 504, 508–509
pre-commitment 176–178, 311, 449, 480–484, 518, 523
preference reversal 126–136
prefrontal cortex 418–419, 423–428, 432–443, 446, 447, 450
preproperate 174–175, 190, 482
present bias 172–179, 189–190, 194–198, 365, 445–451, 454–460, 480–484, 492, 518–523
pricing 194–198, 311–313, 363–365
primary reinforcer 427–428
principal-agent relationship 359, 365
priority matching algorithm 510–512
private value auction 291
procedural invariance 49, 127–128, 130–131, 135
procrastination 174–175, 198, 203, 482
projection bias 477–483, 492–494
prominence effect 131
property rights 352–353
proposer 322–330, 333–336, 339–343, 346–349, 352–353, 357–358, 385–387, 394–395, 401–402, 440–443
prospect 98
prospect theory 115–126, 130, 133–137, 140–145, 155–157, 160
psychological game theory 351–352
punishment 5, 322, 328–333, 335, 340–341, 344–346, 387–388, 391–393, 429, 442–444, 504, 506–508
pure coordination game 261–264

quality adjusted life year (QALY) 487–490, 501
quantal response equilibrium (QRE) 269–271, 282, 300
quasi-hyperbolic discounting 172–176, 190, 192, 445–446

random sequences 220–221
rank dependent expected utility 112, 116–119, 136
readjustment 472

reasoning 45–46, 247–478, 252, 257, 260, 421–422, 433, 438–439, 524
rebate 304
reciprocity 330–335, 338, 348, 359–365, 391–394, 392–396, 441
recognition heuristic 264
recollected utility 182–184
reference-dependent utility 46–47, 118–121, 124, 136, 150, 156–157, 367
reference point 46–56, 66, 74–83, 115–126, 133–136, 143, 150–153, 157, 179–181, 190, 338–341, 349–353, 358–359, 363–364, 450, 469
reference point updating effect 143
reflection effect 107
refund 303, 505
regret theory 131–136
regularity condition 39, 41
reinforcement 452
reinforcement learning 273–274, 277–282, 315
relative risk aversion 100
remembered utility 474–480, 488, 492
responder 322, 328–334, 340–343, 345–349, 394, 401–402
restricted choice 346–350, 394
retirement 10, 85–86, 91–92, 492–495, 519–522
revealed preference 200–201, 465–466
revenue equivalence 290–293
reward prediction error 432–434, 450
risk averse 98–110, 122–126, 137, 142–144, 150, 156–159, 185, 209–211, 398, 403
risk dominant 265–266, 297
risk loving 99–100, 106, 117, 125–126, 136, 145, 156–159, 378, 437

salience 261–263, 525
satisficing 33–37
save more tomorrow 523
saving 26, 84–87, 91–95, 190–193, 492–495, 518–526
scale compatibility 131
Schelling salience 261–262
search 33–38, 66–67, 73–75, 95, 374–376, 526
second price auction 292–294, 515
secondary reinforcer 427
second-order free-rider problem 392
second-order lottery 209–211
self-control 176–179, 449–450, 460, 483–487, 522
self-determination 487–488
Selten, R. 8–9, 12
sequences 168–170, 179–184, 218
sequential game 266–267
session 18
sexual selection 384
shading the bid 291
short selling 239
signaling 228–231, 253, 312, 356–359, 379–382, 397
Simon, H. 7, 11
simulation 423
single blind 21, 323
small area variation 235–237
Smith, Adam 5, 465
Smith, Vernon 8, 12, 14, 69
smoking 452–459, 527
social dilemma 331
social distance 323
social norms 346, 353, 395, 411, 505
social reference point 339–40
sophisticated 174–176, 190, 196–197, 482–484, 519
sophisticated beliefs 257–260
sophisticated truth-teller 61–63
spectrum auctions 512–517
stable matching 510–512
Stackelberg equilibrium 316–319
standard economic model 3–10, 14–16, 24, 43–44, 66, 135, 231, 289, 365, 407, 449, 495, 517
stochastic reference point 123–126
strategy 22–24
strategy method 328, 369
streaks 217–220
striatum 420, 424, 431–434, 443
Stroop task 445
sub-game perfection 9, 268

task switching 423
tax evasion 146–155
tax rebate 88–90
tax saliency 497–501
teams 23, 62–63, 283–290, 299–300, 309–313, 335–338
temptation 176–179, 190, 412–413, 449, 484–485
Thaler, R. 10, 521–524
third party punishment 322, 328, 330, 335
threat 266–268, 341, 388, 442, 508
threshold public good 18–24, 303–305
time trade-off valuation technique 489

time-inconsistency 173–176, 190, 192–198, 204, 380, 445, 480, 493, 501, 518, 523
tolerance 452
total utility 474–476, 486, 491
tradeoff contrast hypothesis 40
tragedy of the commons 505–508, 516–517
transaction utility 51–54, 65–66, 433, 450
transitivity 107–108, 126–128, 178
treatment 20
turnaround game 301–303
Tversky, A. 8, 11, 15, 51, 116, 263
TV game show 121–122
two-part tariff 194–196

ultimatum game 322, 327–330, 334–335, 337–338, 340–343, 346–350, 369, 385–387, 394–396, 401–402, 440–442
uncertainty 97, 104, 153, 201, 204–212, 247, 430–435
unemployment 9, 11, 309–310, 358–362, 471–473, 495–497
update beliefs 207–208
updating rule 272–278
utility 22–23, 31–33, 44, 46, 66, 319, 335, 374–376, 426, 449–450, 465–467

value function 46–47, 55, 116–118
ventral tegmental area (VTA) 424, 428
ventromedial prefrontal cortex (VMPFC) 423–426, 428–439, 442–445, 460
visual analogue scale 489
voting 245–247

wage stickiness 362–365
warm glow 354–358
weakest link game 200, 297–303
weighting probabilities 117
willingness to accept (WTA) 49–51, 127, 129, 211
willingness to pay (WTP) 49–51, 127
windfall income 89
winners curse 293–296
Wisconsin card sorting task 423
World Values Survey 389, 466

zero-intelligence traders 71



Taylor & Francis eBooks

Helping you to choose the right eBooks for your Library

Add Routledge titles to your library's digital collection today. Taylor and Francis ebooks contains over 50,000 titles in the Humanities, Social Sciences, Behavioural Sciences, Built Environment and Law.

Choose from a range of subject packages or create your own!

Benefits for you

- Free MARC records
- COUNTER-compliant usage statistics
- Flexible purchase and pricing options
- All titles DRM-free.

Benefits for your user

- Off-site, anytime access via Athens or referring URL
- Print or copy pages or chapters
- Full content search
- Bookmark, highlight and annotate text
- Access to thousands of pages of quality research at the click of a button.

REQUEST YOUR **FREE** INSTITUTIONAL TRIAL TODAY

Free Trials Available
We offer free trials to qualifying academic, corporate and government customers.

eCollections – Choose from over 30 subject eCollections, including:

- Archaeology
- Architecture
- Asian Studies
- Business & Management
- Classical Studies
- Construction
- Creative & Media Arts
- Criminology & Criminal Justice
- Economics
- Education
- Energy
- Engineering
- English Language & Linguistics
- Environment & Sustainability
- Geography
- Health Studies
- History
- Language Learning
- Law
- Literature
- Media & Communication
- Middle East Studies
- Music
- Philosophy
- Planning
- Politics
- Psychology & Mental Health
- Religion
- Security
- Social Work
- Sociology
- Sport
- Theatre & Performance
- Tourism, Hospitality & Events

For more information, pricing enquiries or to order a free trial, please contact your local sales team:
www.tandfebooks.com/page/sales

Routledge
Taylor & Francis Group

The home of Routledge books

www.tandfebooks.com