

TEMPORAL AND SOCIAL DISCOUNTING OF PAIN AND ILLNESS

A thesis submitted to Imperial College London
for the degree of Doctor of Philosophy (Ph.D.)

Giles William Story

Centre for Health Policy,
Department of Surgery and Cancer, Imperial College London

2015

(In collaboration with Wellcome Trust Centre for Neuroimaging,
University College London)

ABSTRACT

Unhealthy behaviour often entails short-term indulgence at the expense of long-term health. This thesis examines an hypothesis that temporal discounting, a measure of the extent to which a person devalues delayed benefit, predicts unhealthy behaviour. The work also evaluates temporal discounting as a psychological model for unhealthy behaviour, in particular unhealthy behaviour that is enacted in spite of healthy goals, and behaviour with painful consequences. Studies examining relationships between temporal discounting and health behaviour are systematically reviewed, with the finding that discounting of reward correlates with many forms of appetitive unhealthy behaviour. It is proposed that while steep discounting predisposes to unhealthy behaviour, goal-incongruent behaviour is better explained by the interfering effect of prelearned or innate values. Also, conventional discounting fails to account for the fact that many people prefer to expedite inevitable pain or illness. An explanation is that people dislike waiting for pain, termed ‘dread’. The empirical work of this thesis establishes how dread depends on delay, by asking participants to titrate the timing and severity of their own pain or that of others. For the average participant, the dread of pain accumulated at a decreasing rate as pain was delayed. Dread was found to be less marked when deciding on behalf of socially close others. Unexpectedly a tendency to dread future pain in one-off choices did not predict saving a budget of pain relief in sequential choices. Further experiments examined social discounting for pain, finding that participants appear more averse to causing pain in others than in themselves, a tendency that is discounted with social distance. Conclusions are that temporal discounting of reward is a promising marker of appetitive unhealthy behaviour, with a considerable evidence base, while dread offers a candidate marker for engagement in health-promoting behaviour with painful consequences, a possibility which demands further investigation.

CONTENTS

Declaration of Originality	9
Declaration of Copyright	11
Acknowledgements	12
List of Tables.....	13
List of Figures	14
1. Introduction	17
1.1 Health Behaviour.....	18
1.2 Predicting Health Behaviour	19
1.3 Behavioural Economics.....	21
1.3.1 Choices, Preferences and Utilities	21
1.3.2 Irrationality	22
1.4 Behavioural Health Economics	23
1.5 Temporal Discounting as a Marker of Unhealthy Choice.....	24
1.5.1 Exponential Discounting	25
1.5.2 Hyperbolic and Quasi-hyperbolic Discounting	25
1.6 Social Discounting and Social Health Choices	28
1.7 Thesis Outline.....	29
2. Health Discounting	30
2.1 Systematic Review Methods	31

2.2 Positive Time Preference for Health	31
2.3 Zero and Negative Time Preference for Health	32
2.4 Correlations between Health and Monetary Discounting	34
2.5 Conclusions	35
2.5.1 Experiential Outcomes and Zero Time Preference	35
2.5.3 Dread and Negative Time Preference	36
2.5.2 Discounting as a Trait.....	36
3. Relationships between Discounting and Health Behaviour	42
3.1 Systematic Review Methods	42
3.2 Hypothetical Health Discounting and Health Behaviour	43
3.3 Monetary or Appetitive Discounting and Health Behaviour.....	44
3.3.1 Discount Rates and Self-Reported Health.....	44
3.3.2 Discount Rates and Cigarette Smoking.....	45
3.3.3 Discount Rates and Alcohol Use.....	49
3.3.4 Discounting and Illicit Substance Misuse	52
3.3.5 Discounting, Obesity and Eating Behaviour	56
3.3.6 Discount Rates and Preventive Health Behaviour	58
3.3.7 Discount Rates and Risky Sexual Behaviour or Drug-Taking Practices ..	59
3.3.8 Discount Rates and Multiple Health Behaviours	61
3.4 Conclusions	61
3.4.1 Trait and State Variability In Discounting	62
3.4.2 Concluding Remarks	63
4. Discounting and Self-Control	65
4.2 Hyperbolic Discounting and Dynamic Inconsistency	65
4.3 Evidence for and against Hyperbolic Discounting.....	66
4.4 Goal-Incongruent Actions Often Result From State Changes	68

4.5 A Reinforcement Learning Approach to Cue-Triggered Preference Reversal	69
4.5.1 Environmental Cues can Trigger Goal-Incongruent Habits	71
4.6 Discussion	76
4.6.1 Innate (Pavlovian) Influences	76
4.6.2 Future Directions	77
5. General Experimental Methods	79
5.1 Estimating Utility from Discrete Choice Experiments.....	79
5.2 Bayesian Model Comparison	82
5.2.1 Maximum Likelihood Estimation	84
5.2.2 Hierarchical Bayesian Model-Fitting	86
5.3 Continuous Utility Estimates.....	89
5.3.1 Utility Matching Procedures	89
5.3.2 Willingness-to-Pay	89
5.4 Experimental Painful Cutaneous Shock Stimuli	90
5.4.1 Threshold Procedure for DS7 Stimulator	91
5.4.2 Threshold Procedure for DS5 Stimulator	92
5.5 Ethics Statement	94
5.6 Participant Recruitment	94
6. Time Preference for Pain and Illness	96
6.1 Experiment 1A: Time Preference for Painful Shocks	98
6.1.1 Methods	98
6.1.2 Results	103
6.1.3 Discussion	123
6.2 Experiment 1B: Time Preference for Hypothetical Dental Appointments	124
6.2.1 Methods	124
6.2.2 Results	125

6.2.3 Discussion	129
6.3 Experiment 1C: Time Preference for Hypothetical Illness	130
6.3.1 Methods	130
6.3.2 Results and Discussion	134
6.4 Dreading Future Pain: General Discussion	137
7. Allocating Pain across Time	140
7.1 Do People Really Perform Dynamic Utility Maximisation?.....	141
7.2 Methods	144
7.2.1 Experiment 2: Dynamic Consumption of Pain Relief	144
7.2.2 Dynamic Programming	147
7.2.3 Model Fitting Procedures	151
7.3 Results	151
7.3.1 Simulating Consumption Behaviour	151
7.3.2 Observed Consumption of Relief	154
7.3.3 Predicting Consumption from One-Off Choices between Delayed Pains	158
7.3.4 Modelling Relief Consumption Using Heuristics	162
7.4 Discussion	167
8. Social Preference for Pain	172
8.1 Social discounting	172
8.2 Dreading the pain of others?	174
8.3 Experiment 3A: Preferences for Others' Delayed Pain	177
8.3.1 Methods	177
8.3.2 Results	183
8.3.3 Summary of Results from Experiment 3A	195
8.4 Experiment 3B: Social Discounting of an Hypothetical Painful Treatment	195
8.4.1 Methods	196

8.4.2 Results	198
8.4.3 Summary of Results from Experiment 3A	204
8.5 Experiment 3C: Interaction of Social and Temporal Discounting for Pain	205
8.5.1 Methods	205
8.5.2 Results	205
8.6 Social Preferences for Pain: General Discussion	208
9. Social Allocation of Pain	212
9.1 Inequity Aversion	213
9.2 Are People Inequity Averse for Pain?	214
9.3 Does Unfair Treatment Exacerbate Pain?	216
9.4 Experiment 4: Responses to Pain Inequality	217
9.4.1 Methods	217
9.5 Results	228
9.5.1 Distribution of Mean Dictator Game Offers	228
9.5.2 Inequality Aversion for Pain and Money	230
9.5.3 Self-Oriented Preferences for Money but not Pain	232
9.5.4 Status Quo Bias	235
9.5.5 Modelling Dictator Behaviour.....	236
9.5.6 Evaluation of Experienced Pain Allocations.....	240
9.5.7 Ratings of Fairness, Responsibility and Inclination to Punish	242
9.6 Discussion	243
10. General Discussion	247
10.1 Delay Discounting as a Viable Behavioural Marker of Unhealthy Choice	248
10.2 Environmental Influences on Discounting	249
10.2 Future Directions for Studying Self-control in health choice	253
10.2.1 ‘Model-based’ and ‘Model-Free’ Interactions	254

10.2.2 Relationships Between ‘Model-Based’ Control and Discounting	255
10.2.3 Pre-commitment, Willpower and Sophistication	256
10.3 Dread and Avoidance Behaviour	258
10.3.1 Avoidable Pain	258
10.3.2 What Does Dread Represent?.....	260
10.4 Social Discounting of Pain	262
10.5 Spreading and Inequality Aversion for Pain	263
10.6 Limitations.....	264
10.6.1 Face Validity of Decision Tasks	264
10.6.2 Participants and Sampling Biases.....	265
10.6.3 On Computational Modelling.....	266
10.7 Fin.....	267
Bibliography	268
Appendix: Binary Choice Options	300

DECLARATION OF ORIGINALITY

Unless otherwise referenced this thesis is the work of Giles William Story. The thesis contains work that has been previously published in or submitted to peer-reviewed journals. At the time of writing the systematic literature review described in Chapter 2, the ideas in Chapter 4, Experiments 1A and 1B described in Chapter 6 and Experiment 2 described in Chapter 7 have been published:

- Story, G. W., Vlaev, I., Seymour, B., Darzi, A., & Dolan, R. J. (2014). Does temporal discounting explain unhealthy behaviour? A systematic review and reinforcement learning perspective. *Frontiers in Behavioural Neuroscience*, 8: 76
- Story, G. W., Vlaev, I., Seymour, B., Winston, J. S., Darzi, A., & Dolan, R. J. (2013). Dread and the disvalue of future pain. *PLoS Computational Biology*, 9(11), e1003335
- Story, G. W., Vlaev, I., Dayan, P., Seymour, B., Darzi, A., & Dolan, R. J. (2015). Anticipation and choice heuristics in the dynamic consumption of pain relief. *PLOS Computational Biology*, 11(3), e1004030

Experiments 3 and 4, described in Chapters 8 and 9 respectively have been submitted for publication under the following titles:

- Story, G. W., Chowdhury, R., Vlaev, I., Crockett, M., Darzi, A., & Dolan, R. J. (2015). Dreading the pain of others? Negative social and temporal discounting for pain.
- Story, G. W., Vlaev, I., Metcalfe, R., Crockett, M., Darzi, A., & Dolan, R. J. (2015). Social redistribution of pain and money.

The papers are not simply reproduced here in their entirety, but the constituent work appears in a form suitable to support the narrative of the thesis.

For all the studies submitted for publication, Giles Story wrote both the published manuscripts and the updated exposition submitted here, designed the experiments, wrote the necessary computer code for running the protocols, performed all systematic reviews included in this thesis, analysed all the experimental data and produced all figures and tables. Giles Story also recruited participants and carried out the experiments (for Experiment 3A with the assistance of Rony Chowdhury). Co-authors, namely Ivo Vlaev, Ben Seymour, Rob Metcalfe, Zeb Kurth-Nelson, Molly Crockett, Ray Dolan, Peter Dayan and Ara Darzi provided comments and suggestions on previously submitted manuscripts, and circumscribed editorial input. Ivo Vlaev and Ben Seymour provided advice on the use of statistical methods. Peter Dayan and Zeb Kurth-Nelson provided advice in the use of computational methods. Peter Dayan also reviewed the equations presented in Chapters 6 and 7 and checked the computer code used for the dynamic programming analyses. Aspects of the experimental and analytical methods that have been used in previously published studies are identified and referenced.

DECLARATION OF COPYRIGHT

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Previously published work included in the thesis is reproduced under a Creative Commons Attribution (CC BY) licence. Under the CC BY licence, authors retain ownership of the copyright for their article. Provided the work is appropriately cited, no permission is required from the authors or the publishers to reprint.

ACKNOWLEDGEMENTS

I would like to thank several people for their help with this work:

- Professor the Lord Ara Darzi for giving me the opportunity to carry out this research, for his continued enthusiasm for the work and for being willing to incorporate a psychiatry trainee within the Department of Surgery and Cancer.
- Professor Ray Dolan for taking me on as a member of his research group, and for encouraging me to pursue a Ph.D. He is a great source of expertise on the neuroscience of decision-making and has advised me superbly on all aspects of the work.
- Professor Ivo Vlaev for his effervescent energy, for boosting my morale and for his excellent advice on the behavioural economic aspects of the work.
- Professor Peter Dayan for some most interesting discussions on intertemporal choice and for his extremely thorough and patient supervision on the computational aspects of this work.
- Dr Zeb Kurth-Nelson, Dr Michael Moutoussis and Dr Molly Crockett for some stimulating discussions on intertemporal and interpersonal choice.
- Professor Roger Carpenter for inspiring my interest in decision-making.
- My fiancée Emma Grant, for her great love and support.
- My parents for giving me the chance to study and for supporting me throughout my life.

LIST OF TABLES

Table 2.1	Studies Measuring Discount Rates for Health Outcomes
Table 3.1	Studies Comparing Discount Rates and Smoking Behaviour
Table 3.2	Studies Comparing Discount Rates and Alcohol Use
Table 3.3	Studies Comparing Discount Rates and Illicit Substance Misuse
Table 3.4	Studies Comparing Discount Rates and Obesity or Eating Behaviour
Table 3.5	Studies Comparing Discount Rates and Self-Reported Health, Preventive Health Behaviours or Multiple Health Behaviours
Table 7.1	Putative Mechanisms Underlying <i>Save-Now-Spend-Later</i> , <i>Spend-Now-Suffer-Later</i> and <i>Spread-Spending</i> Heuristics
Table A1	Binary choice options for Experiment 1A
Table A2	Binary choice options for Experiment 1B
Table A3	Binary choice options for <i>self-now-self-later</i> and <i>other-now-other-later</i> conditions in Experiment 3A
Table A4	Binary choice options for <i>self-now-other-now</i> choices in Experiment 3A
Table A5	Binary choice options for <i>self-now-other-now</i> (social discounting) choices in Experiment 3B

LIST OF FIGURES

Figure 1.1	Hyperbolic and exponential discounting
Figure 4.1	Interactions between model-based and model-free decision-making
Figure 5.1	Softmax (logistic) action selection function for binary choice
Figure 6.1	Trial structure of the task in Experiment 1A
Figure 6.2	Observed time preference: Experiment 1A
Figure 6.3	Choice data of individual participants in Experiment 1A categorised by time preference
Figure 6.4	Temporal value functions predicted by alternative dread-discounting models
Figure 6.5	Model predictions on the task: Experiment 1A
Figure 6.6	Parameterisation of the general form Exponential Dread model
Figure 6.7	Parameterisation of the Undiscounted Exponential Dread model
Figure 6.8	Model comparison and framing effects in Experiment 1A
Figure 6.9	Time preference of sample participants in Experiment 1A and fits of the (discounted) exponential dread model
Figure 6.10	Time preference for a hypothetical painful dental appointment in Experiment 1B
Figure 6.11	Model comparison for Experiment 1B
Figure 6.12	Time preference for illness in Experiment 1C predicted by an Undiscounted Exponential Dread model
Figure 6.13	Time preference for illness in Experiment 1C
Figure 6.14	Model comparison for Experiment 1C
Figure 7.1	Trial structure of Experiment 2
Figure 7.2	Anticipation-discounting functions
Figure 7.3	Effects of changes in the instantaneous utility function on dynamic consumption
Figure 7.4	Anticipation-discounting and dynamic utility maximisation

Figure 7.5	Observed distribution of relief consumption across time in Experiment 2
Figure 7.6	Consumption paths from six sample participants in Experiment 2
Figure 7.7	Histograms of relief consumption over time in Experiment 2
Figure 7.8	Relationships between one-off (binary) and dynamic intertemporal choices
Figure 7.9	Fits of an anticipation-discounting model with variable utility and choice randomness to consumption choices in Experiment 2
Figure 7.10	Heuristic model fits to consumption choices in Experiment 2
Figure 7.10	Policy weightings of heuristic models in Experiment 2
Figure 8.1	Protocol for Experiment 3A
Figure 8.2	Negative discounting for pain in the <i>self-now-self-later</i> condition in Experiment 3A
Figure 8.3	Time preference for pain for self and other in Experiment 3A
Figure 8.4	Social modulation of dread and social discounting in Experiment 3A
Figure 8.5	Distribution of social discounting across participants in Experiment 3A
Figure 8.6	Skin conductance responses to delayed pain for self or other in Experiment 3A
Figure 8.7	Experiment 3B: Dread-discounting of hypothetical painful treatments for self and other
Figure 8.8	Social discounting of hypothetical painful treatment in Experiment 3B
Figure 8.9	Social discounting of pain in individual participants in Experiment 3B
Figure 8.10	Social modulation of dread in Experiment 3C
Figure 9.1	Experiment 4: Modified pain dictator game protocol, <i>Give</i> and <i>Take</i> frames
Figure 9.2	Experiment 4: Modified pain dictator game protocol, <i>Give or Take</i> frame
Figure 9.3	Experiment 4: Modified pain dictator game, choice options within each frame
Figure 9.4	Experiential phase of Experiment 4: Trial structure for receivers
Figure 9.5	Distribution of mean resulting dictator game offers in Experiment 4
Figure 9.6	Dictator behaviour for pain as a function of default allocations in Experiment 4

- Figure 9.7 Dictator behaviour for money as a function of default allocations in Experiment 4
- Figure 9.8 A status quo bias in dictator behaviour
- Figure 9.9 Social utility functions based on a Social Discounting model
- Figure 9.10 Model fits to pain dictator game choices in Experiment 4
- Figure 9.11 Model fits to money dictator game choices in Experiment 4
- Figure 9.12 Rated pain intensity as a function of social context in Experiment 4
- Figure 10.1 Discounted dread and avoidance

INTRODUCTION

“The mind has great influence over the body, and maladies often have their origin there.”

Molière – *l’Amour Medecin* (1665)

In the United States in 1900 pneumonia, tuberculosis and gastrointestinal infection were the chief causes of mortality, together responsible for one third of all deaths (Department of Commerce and Labor, 1906); by 1997 these were replaced by heart disease, cancer and stroke, together responsible for almost two-thirds of all deaths (Hoyert, Kochanek, & Murphy, 1999). The fraction of total mortality attributable to non-communicable disease is projected to increase further over the first quarter of the twenty-first century, placing unsustainable demands on existing healthcare systems (R. Smith, Corrigan, & Exeter, 2012). A substantial proportion of non-communicable disease results from particular behaviour patterns, for example it is estimated that close to a fifth of all deaths in the US in the year 2005 were attributable to tobacco smoking, while physical inactivity and obesity were each responsible for nearly a tenth of all deaths over the same period (Danaei et al., 2009), and reducing the prevalence of these behaviours is therefore crucial to reducing the burden of non-communicable disease. However, the determinants of unhealthy behaviour remain poorly understood, and have only become a focus of research in recent decades (Conner & Norman, 2005a).

The current dissertation proposes that understanding the reasons why people engage in unhealthy behaviour will assist the preventive effort. Existing research in health psychology has also adopted this view (Adler & Matthews, 1994; Armitage & Conner, 2000; Baum & Posluszny, 1999; Conner & Norman, 2005a), though the prevailing models of health behaviour require updating in light of recent advances in the behavioural sciences. In particular, formal descriptions of how human decision-

making diverges from classical economic notions of rationality, collectively referred to as behavioural economics (Camerer et al., 2005; Camerer, Loewenstein, & Rabin, 2004; Glimcher & Rustichini, 2004; Sanfey et al., 2006), offer a promising means to understand why people often make unhealthy choices despite knowing that such choices are unwise and planning to behave more healthily. The current work aims to characterise individual differences in health-related decision-making using tools derived from behavioural economics.

The primary focus of the work is on personal health decisions, such as deciding whether to smoke a cigarette, whether to attend a potentially unpleasant cancer screening appointment, or how to consume a budget for spending on one's health. By examining radically simplified versions of these choices presented to experimental participants, I construct models of the underlying psychological processes, in the hope that these models will yield parameters capable of predicting a person's real-world health-related choices. A secondary focus concerns whether people behave differently when making decisions regarding other peoples' health, a task facing health care professionals and policy-makers, in the hope that understanding such differences might assist in improving shared decision making in clinical practice. This aspect of the work also addresses a theoretical parallel between psychological processes involved in evaluating our own long-term wellbeing and those involved in evaluating the wellbeing of others.

1.1 HEALTH BEHAVIOUR

All behaviour might be said to influence health in some way. However, for the sake of parsimony, researchers have tended to focus on a set of easily defined behaviours that are known to be beneficial or deleterious to health. Frequently studied health-promoting behaviours include exercise participation, healthy eating, attendance for screening procedures or immunisations, and concordance with recommended medication. Frequently studied harmful behaviours include tobacco smoking, excessive alcohol intake, unhealthy eating, illicit substance misuse, and risky sexual practices (Ogden, 2012). Behaviours in either category could easily belong to the opposite category by defining them in terms of their absence; therefore I will use the term 'health behaviour' to the behaviours above, irrespective of valence. This usage

differs from existing definitions, which refer specifically to behaviour beneficial to health. For example, Conner and Norman define health behaviour broadly as “Any activity undertaken for the purpose of preventing or detecting disease or for improving health and well-being” (Conner & Norman, 2005b, p2).

1.2 PREDICTING HEALTH BEHAVIOUR

Demographic factors are known to influence engagement in health behaviours. For example, young age, high levels of education, and high levels of social support are associated with health-promoting behaviours (Ford et al., 1991; C. E. Ross & Wu, 1995), while lower socioeconomic status is associated with harmful behaviours such as smoking, alcohol excess and substance abuse (Allgöwer, Wardle, & Steptoe, 2001; Baum, Garofalo, & Yali, 1999; Lynch, Kaplan, & Salonen, 1997; Steptoe, Wardle, Pollard, Canaan, & Davies, 1996; Van Oers, Bongers, Van de Goor, & Garretsen, 1999). Unsurprisingly, however demographic data cannot explain all variability in health behaviour; that is individuals from similar backgrounds nevertheless differ in their behaviour (Cummings et al., 1980). As a result over the last three to four decades researchers have attempted to measure the more proximate antecedents of behaviour, in the form of individual cognition, with the aim of improving our ability to predict health behaviour (Conner & Norman, 2005b), and thereby allowing interventions to be directed towards at-risk groups.

The most widely used models for understanding the psychological determinants of health behaviour are the social cognition models (SCMs) (Fiske & Taylor, 2013), which have been widely influential in health psychology and have influenced the design of behaviour change interventions (Conner & Norman, 2005b). Earlier SCMs, such as the Health Belief Model (HBM) (Rosenstock, 1974) and Protection Motivation Theory (PMT) (R. W. Rogers, 1975) are primarily concerned with the individual as a thinker; thoughts are seen to follow on from perception and to precede behaviour, and are accessed by asking people to report their beliefs explicitly. The Health Belief Model for example, proposes that the key determinants of health behaviour are beliefs about the threat posed by illness and about the relative benefits of and barriers to performing the behaviour, and originated from attempts by US public health researchers to improve the effectiveness of health education

programmes. In an early demonstration of this Hochbaum (1958) found that attendance for chest X-ray screening to detect tuberculosis showed a relationship with perceived susceptibility to tuberculosis, and a belief that the disease could be asymptomatic. Haefner (1974) went on to show that providing health education which aimed to increase perceived susceptibility, perceived severity and anticipated benefits, resulted in a greater number of routine visits to the doctor over the following months in the intervention group compared with the control group.

The SCMs as a whole are reasonable predictors of the *intention* to engage in health behaviours. For example Milne, Sheeran, and Orbell (2000) report small to medium associations between constructs of PMT and health intentions ($|r_+|=0.10-0.34$; r_+ represents the average Pearson correlation coefficient, r , across studies weighted by sample size). However, meta-analysis has shown that the intention to perform health behaviour explains less than half of the variance in behaviour itself (Sheeran, 2002; Webb & Sheeran, 2006). For example, in an analysis of 47 studies Webb and Sheeran (2006) showed that behaviour change interventions which had a medium-to-large effect on intention ($d_+ = 0.66$, 95% confidence interval [0.51 to 0.82]; d_+ represents the weighted average Cohen's d , a measure of effect size, see Cohen, 1992) had only a small-to-medium effect on behaviour ($d_+ = 0.36$, 95% confidence interval [0.22 to 0.50]). Consistent with this the SCMs show weaker correlations with health behaviour itself, and are particularly challenged to predict behaviour prospectively. For example, in meta-analysis average correlations between HBM constructs and health *behaviour* are weak (Pearson's $r = 0.15, 0.08, 0.13$ and 0.21 for susceptibility, severity, benefits and barriers respectively (Harrison, Mullen, & Green, 1992); and associations between PMT constructs (R. W. Rogers, 1975) and future engagement in health behaviours such as condom use, attendance at cancer screening or smoking cessation are either low or non-significant (J. Cohen, 1992; Milne et al., 2000).

Consistent with the explanatory gap between health intentions and behaviour, the theoretical structure of the earlier SCMs does little to incorporate the fact that our behaviour does not always correspond to our declared goals (Ainslie, 2001; Allan, Johnston, & Campbell, 2010; Bargh, 1992; Nosek, Hawkins, & Frazier, 2012; Read & Van Leeuwen, 1998). Notably more recent social-cognition models, such as the Theory of Planned Behaviour (Ajzen, 1991) and Implementation Intentions

(Gollwitzer & Sheeran, 2006) have attempted to address this. A further limitation of the SCMs is that they tend to lack a specification for the mechanisms linking the constructs, which are treated as independent variables (Conner & Norman, 2005a, 2005b). Finally the constructs of the SCMs are measured through self-report, making them susceptible to self-presentation biases, or peoples' motivation to conform to the implicit requirements of the survey (Sheeran & Orbell, 1996). Recent advances in decision-making psychology and neuroscience provide alternative or complementary approaches with the potential to address some of these limitations.

1.3 BEHAVIOURAL ECONOMICS

In recent years there has been a growing interest in understanding human decision-making from an information processing perspective (Camerer, Loewenstein, & Prelec, 2005; Dayan, 2012; Dayan & Balleine, 2002; R. J. Dolan & Dayan, 2013; Glimcher, 2004; Glimcher, Dorris, & Bayer, 2005; Glimcher & Rustichini, 2004; Montague, Dolan, Friston, & Dayan, 2012; Rangel, Camerer, & Montague, 2008; Sanfey, Loewenstein, McClure, & Cohen, 2006). This endeavour has seen a confluence between the fields of economics, psychology, and behavioural neuroscience, giving rise to the fields of behavioural economics and neuroeconomics (Camerer et al., 2005; Camerer, Loewenstein, & Rabin, 2004; Glimcher & Rustichini, 2004; Sanfey et al., 2006).

1.3.1 CHOICES, PREFERENCES AND UTILITIES

While the social cognitive approach described above aims to access mental processes directly through self-report, economists have historically tended to take *choice* behaviour as their starting point. In the first half of the twentieth century a group of economists referred to as the neo-classical school began to formalise mathematical notions of rational choice (see for example Debreu, 1991; Samuelson, 1937), laying down axioms which if followed would lead to internally consistent choices. Rather than seeking to describe how people actually behave, the resulting models of behaviour focussed on how individuals could make the best possible

decisions (Baron, 2000). A key principle of this approach is that choosing an alternative is said to *reveal preference* for that alternative (Houthakker, 1950; Sen, 1971). To take a health-related example, if a consumer of health-care opts for elective surgery over management with Medicine A, they are said to *prefer* surgery to Medicine A.

Neo-classical economists also derived sets of axioms for how agents should best choose between outcomes that are uncertain, or delayed in time (Von Neumann & Morgenstern, 1944; Samuelson, 1937). These theories explain preferences by assuming that decision-makers choose *as if* to maximise an unobservable internal quantity termed utility, related to the objective value of outcomes by a continuous function, the utility function. For neo-classical economists, preferences and utilities were unobservable quantities, existing only in theory. However the subsequent emergence of behavioural economics helped to lay the foundations for using the concepts of utility and preference to study the psychological antecedents of choice in their own right (Glimcher, Camerer, Fehr, & Poldrack, 2009).

1.3.2 IRRATIONALITY

The field of behavioural economics arose out of a series of observations over the latter half of the twentieth century showing that human decision-makers violate the axiomatic assumptions of the neo-classical models (Camerer & Loewenstein, 2004; Camerer et al., 2005; Camerer et al., 2004; Glimcher et al., 2009; Kahneman & Tversky, 2000). In each case, to account for these ‘irrational’ deviations from normative theory, behavioural economists introduced new *descriptive* theories. For example, expected utility theory (EUT) states that rational decision-makers should compute the value of a risky gamble by weighting the utility of each outcome by its probability of occurrence (Von Neumann & Morgenstern, 1944). However, by offering participants choices between structured gambles, Kahneman and Tversky illustrated that humans deviate from EUT, displaying systematic biases in the interpretation of stated probabilities, which were taken as the basis for their Prospect Theory, a descriptive theory of risk preference (Kahneman & Tversky, 1979).

Behavioural economic research has also revealed biases in how people select actions (Kahneman & Tversky, 2000). For example people show a tendency to choose

pre-selected ‘default’ options (Choi, Laibson, Madrian, & Metrick, 2004; Kahneman, Knetsch, & Thaler, 1991), and often use convenient rules of thumb, termed heuristics, to make potentially complex decisions, such as choosing how much money to save for retirement (Benartzi & Thaler, 2007; Gigerenzer, 2008; Gigerenzer, Hertwig, & Pachur, 2011). Similarly, studies have demonstrated that relatively subtle changes in how an option is described, for example as a gain in reward from a low baseline versus a loss from a high baseline, can sometimes have a marked effect on preference, leading to a set of effects collectively referred to as framing (Kahneman & Tversky, 2000; Tversky & Kahneman, 1981). These deviations from neo-classical notions of rationality provided behavioural economists with the building blocks for a new psychologically grounded theory of decision-making. The leap of faith required was to go beyond the *as if*, and to consider the behavioural economic theories as hypotheses for how humans process information to arrive to a decision (Camerer et al., 2005; Glimcher et al., 2009; Glimcher & Rustichini, 2004).

Several behavioural and neuro- economic approaches evoke multiple psychological (or neural) systems for decision-making as a means to account for observed deviations from optimality (e.g. Angeletos, Laibson, Repetto, Tobacman, & Weinberg, 2001; Kahneman, 2011; Loewenstein, 1996; Loewenstein & O'Donoghue, 2004; McClure, Laibson, Loewenstein, & Cohen, 2004). These models reflect a long tradition in psychology, which pre-dates behavioural economics, of distinguishing between fast, automatic and slower, controlled or deliberative information processing (Bargh, 1992; Hofmann, Friese, & Strack, 2009; Schneider & Chein, 2003; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977; E. R. Smith & DeCoster, 2000; Tiffany, 1990), or between emotionally laden, so-called ‘hot’ states and more seemingly rational, ‘cold’ states (Chapman, 2005; Metcalfe & Mischel, 1999). Such models are useful in so far as they acknowledge, unlike some of the social cognitive models discussed above, that human behaviour is not exclusively preceded by a deliberative weighing-up of information.

1.4 BEHAVIOURAL HEALTH ECONOMICS

Within the behavioural economic tradition described there is now a growing endeavour to identify decision-making phenotypes which correlate with maladaptive

behaviour (Montague et al., 2012; Bickel et al., 2007; Burks, Carpenter, Götte, & Rustichini, 2012). In particular, behavioural economic measures offer new tools with which to attempt to predict health behaviour (Vlaev & Darzi, 2012). For example, in a sample of 1094 adults, L. R. Anderson and Mellor (2008) find significant negative relationships between individual risk aversion and health behaviours (including cigarette smoking, heavy drinking and being overweight or obese), after controlling for demographic characteristics. Furthermore, a new interventional science based on behavioural economic principles attempts to change (or ‘Nudge’) behaviour by exploiting commonly exhibited biases and heuristics (Thaler & Benartzi, 2004; Thaler & Sunstein, 2008). A recent example is the MINDSPACE framework, a policy tool commissioned by the UK government to assist civil servants in increasing the behavioural impact of their policy interventions (P. Dolan et al., 2012; P. Dolan, Hallsworth, Halpern, King, & Vlaev, 2009; King, Greaves, Vlaev, & Darzi, 2013).

Behavioural economic models have the potential to address some of the shortcomings of the social cognitive models of health behaviour described above; for example they emphasise choice-based methods of eliciting value, which are less vulnerable to self-presentation effects than subjective report. Also, behavioural economic approaches acknowledge that humans often make inconsistent choices, providing a starting point for addressing intention-action discrepancies. Finally, the economic approach brings with it mathematical formalism, which requires researchers to be more explicit about their underlying assumptions.

1.5 TEMPORAL DISCOUNTING AS A MARKER OF UNHEALTHY CHOICE

The current dissertation seeks primarily to critically appraise one behavioural economic tool, the individual measurement of ‘temporal discounting’, as a model for predicting and understanding apparently shortsighted personal health choice. Given a choice, people tend to prefer immediate over delayed rewards, for example preferring to receive \$10 today rather than \$15 a month from today. This implies that people assign lower utility to delayed rewards relative to more immediate ones, a phenomenon referred to as temporal discounting, for the reason that delayed reward is discounted (Samuelson, 1937, for reviews see Frederick, Loewenstein, & O'Donoghue, 2002; Kalenscher & Pennartz, 2008). By analysing how a person

chooses between immediate and delayed rewards (Mazur, 1987), for example by titrating upwards the magnitude of a small immediate reward until it is deemed subjectively equivalent to receiving a larger delayed reward, it is possible to estimate the extent of temporal discounting, summarised by a single parameter, the temporal discount *rate* (e.g. Green & Myerson, 2004; Myerson & Green, 1995). Unhealthy behaviour frequently has a delayed effect on health, leading to the hypothesis that an individual's tendency to make unhealthy choices is positively correlated with their temporal discount rate (Bickel, Jarmolowicz, Mueller, Koffarnus, & Gatchalian, 2012; MacKillop et al., 2011).

1.5.1 EXPONENTIAL DISCOUNTING

Neo-classical economic theory posits that goods ought to be discounted exponentially with delay (Samuelson, 1937). Formally, an outcome with utility A if received straightaway ($t = 0$) is worth $A\delta^t$ if delayed t periods into the future. The present-time value, V , of receiving A at time t is thus given by:

$$V(A, t) = A\delta^t \quad (1.5.1)$$

Here the discount rate, δ , embodies a constant proportional decrease in value with each time period of delay. The assumption of a constant discount rate leads to the prediction that if two rewards are separated by a given delay, the ratio of their value remains constant no matter how far into the future the first reward occurs. This property, known as the stationarity axiom, ensures that decision-makers who discount exponentially make consistent choices between alternative streams of future reward (Strotz, 1957) (Figure 1.1a).

1.5.2 HYPERBOLIC AND QUASI-HYPERBOLIC DISCOUNTING

Exponential discounting is a normative, rather than descriptively accurate theory. In particular, both humans and animals violate the assumption of a constant proportional discount factor; instead net discounting is lower when measured over longer delays than over shorter ones, implying that the discount factor for each added unit of delay decreases as delays grow longer (Ainslie, 1974, 1975, 2001; Ben Zion,

Rapoport, & Yagil, 1989; Chung & Herrnstein, 1967; Green et al., 1994; Mazur, 1987; Strotz et al., 1957). In other words delaying a reward a short time into the future markedly decreases its subjective value, but each additional increment of delay has a smaller proportional effect. This pattern is better accounted for by a hyperbolic function, written in its simplest form as follows, where k represents the hyperbolic discount rate:

$$V(A, t) = \frac{A}{1 + kt} \quad (1.5.2)$$

An alternative is the “quasi-hyperbolic” approximation to hyperbolic discounting (Angeletos et al., 2001; Laibson, 1997; McClure, Ericson, Laibson, Loewenstein, & Cohen, 2007; McClure et al., 2004; Phelps & Pollak, 1968), which is formalised as exponential discounting, combined with a preference for immediate reward, expressed as a second discount factor, β , applied to all time-periods except the first:

$$V(A, t) = \begin{cases} A & \text{if } t = 0 \\ A\beta\delta^t & \text{if } t > 0 \end{cases} \quad (1.5.3)$$

The hyperbolic and quasi-hyperbolic forms qualitatively capture the observation that individuals make far-sighted plans when outcomes are distant, but reverse their choices in favor of short-term rewards when the future is reached (Figure 1.1b and c; Green et al., 1994; Kirby & Herrnstein, 1995; Strotz et al., 1957; see Kalenscher & Pennartz, 2008 for a review). To extend the example above, people might archetypally prefer to receive \$10 today over \$15 in a month, yet prefer \$15 in thirteen months over \$10 in twelve months (though see Chapter 4 for a more detailed discussion of the evidence for this). In economic terms this phenomenon is referred to as myopic preference reversal (Kirby & Herrnstein, 1995).




Figure 1.1 Hyperbolic and exponential discounting. Discounted value, $V(A, t, \tau)$ under three discount functions is plotted as a function of the decision maker's position in time, τ . Here A is the instantaneous utility of the outcome and t the time at which it is due to be delivered. A larger-later reward, LL, of utility l is due to be received at t_3 and a smaller-sooner reward, SS, of utility s is due to be received at t_2 . **A** Exponential discounting. The ratio of the discounted value of the two rewards is constant irrespective of how far away the earlier reward is in time, leading to consistent choices. In this case the decision-maker always prefers LL (i.e., $V(l, t_3, \tau) > V(s, t_2, \tau)$ for all $\tau < t_3$). **B** Simple hyperbolic discounting with a low discount rate. The ratio of the value of two rewards is no longer constant as a function of τ , but SS becomes relatively more valuable as both rewards are approached. The hyperbolic discount rate, k , governs the steepness of the curvature. Here, where k is low LL is still preferred throughout, even when the smaller sooner reward is immediately available. **C** Hyperbolic discounting with a high discount rate. At t_1 , when both rewards are distant, LL is preferred, i.e., $V(l, t_3, t_1) > V(s, t_2, t_1)$, however SS becomes increasingly desirable as it approaches in time, such that at t_2 , SS is preferred, i.e., $V(l, t_3, t_2) < V(s, t_2, t_2)$.

Hyperbolic or quasi-hyperbolic discounting therefore putatively imparts an explanation for why people intend to perform behaviours which they fail to carry out (though see Read, 2001; Read, Frederick, & Airoldi, 2012). Thus, not only does temporal discounting offer a potential model for unhealthy actions with deleterious long-term consequences, but also hyperbolic discounting purports to explain the marked intention-action gap in health behaviour which the social cognitive models fail to address. Furthermore discounting is easy to measure and as such seems well placed as a predictor of individual differences in health behaviour. Evaluating discounting as a viable model for unhealthy behaviour forms the fulcrum of this thesis.

1.6 SOCIAL DISCOUNTING AND SOCIAL HEALTH CHOICES

A secondary aim of this thesis is to examine how people make choices regarding others' health, pertinent to the perspective of health-care professionals or policy-makers. For example, if temporal discounting were related to an individual's propensity to make unhealthy choices, it would be interesting to ask whether people discount the future differently when choosing on behalf of others. Were systematic differences in intertemporal preference for self and other to emerge, this might have implications for shared decision-making in health care settings.

It is also of interest that similar models have been used to understand both social and intertemporal decision-making. While temporal discounting refers to the fact that the subjective value of reward decreases with delay, in a related sense, the subjective value of giving reward to others appears to decrease with *social distance*, termed social discounting (B. Jones & Rachlin, 2006; B. A. Jones & Rachlin, 2009; Rachlin & Jones, 2008). Social distance can be formalised by asking people to create an imaginary rank ordering of the people closest to them, starting with their dearest friend or relative (B. Jones & Rachlin, 2006). Remarkably, previous work has shown that social discount functions for donations of money are also well described by a decreasing hyperbolic function (B. Jones & Rachlin, 2006; B. A. Jones & Rachlin, 2009; Rachlin & Jones, 2008). Furthermore, temporal and social discount rates for money are correlated, with more patient individuals tending also to be more altruistic (Curry, Price, & Price, 2008; Rachlin & Jones, 2008), suggesting that

temporal and social valuations of reward engage similar computational processes. These findings hint at the possibility of a ‘patient-altruistic’ phenotype (see also Bartels & Rips, 2010; Ersner-Hershfield, Wimmer, & Knutson, 2009; Harris & Madden, 2002), whereby people who make sensible long-term choices are also more likely to be prosocial. If this were to hold true for health-related choices, *social* discounting might even form a useful predictor of *personal* health behaviour.

1.7 THESIS OUTLINE

This thesis first systematically reviews studies measuring temporal discounting for hypothetical health outcomes (Chapter 2). Studies comparing temporal discount measures with any form of health outcome are reviewed subsequently in Chapter 3. Chapter 4 critically appraises the notion that hyperbolic discounting accounts for failures of self-control, and proposes an alternative account in terms of interplay between the goal-directed and habitual control of action. Chapter 5 describes the general experimental methods used in subsequent chapters, including Bayesian model comparison.

The vast majority of existing studies comparing discounting with health behaviour examine temporal preferences for reward. People exhibit rather more unusual preferences, not well described by temporal discounting, when choosing the timing of pain or illness (Berns et al., 2006; Loewenstein & Prelec, 1991; Van Der Pol & Cairns, 2000). To address this in Chapter 6 I test models of time preference for pain and illness. Chapter 7 explores the implications of these models for the use of a personal health budget for the relief of pain in real-time.

In Chapter 8 I test the same models in a setting where individuals must choose the timing of others’ pain, pertinent to the perspective of health-care professionals or policy-makers. In Chapters 8 and 9 I also address the theoretical and empirical overlap between temporal and social decision-making, including social discounting (Jones & Rachlin, 2006; Rachlin & Jones, 2008).

HEALTH DISCOUNTING

“In order to arrive at the measurement of utility... it is necessary to place the individual (*homo economicus*) ... under certain ideal circumstances where his behaviour will render open to unambiguous inference the form of the function which he is conceived of maximising.”

Paul A. Samuelson – A note on the measurement of utility (1937)

If an individual's engagement in healthy behaviour is considered equivalent to their investing in their long-term health, then their discount rate ought to predict such engagement (Fuchs, 1982; Grossman, 1972). A key question in testing this hypothesis is whether to elicit intertemporal choices between monetary outcomes, as is most conventional, or between health states. Clearly in real-world health choices the ultimate delayed outcome is a health state. However, it is pertinent to ask whether measured discounting is equivalent for health and money outcomes.

Neo-classical economic theory dictates that discount rates should be the same across all commodities (Samuelson, 1937), for the reason that if all commodities can be monetised then adopting a different rate of discount for different commodities leads to inconsistent choices (see Brouwer, Niessen, Postma, & Rutten, 2005; Cairns, 1992; Chapman & Elstein, 1995; Severens & Milne, 2004). However, health states are not always directly fungible, for example health cannot be directly invested (Chapman, 1996b), nor exchanged for other commodities, suggesting that discounting for health outcomes might differ from monetary discounting. Furthermore there is no good reason to believe that people should behave in line with normative economic theory.

Studies measuring health discounting require subjects to imagine a hypothetical illness and to trade-off the duration or severity of the illness against when they would prefer to experience it (e.g. Van Der Pol & Cairns, 2002). Individuals who will accept more severe delayed illness over less severe immediate

illness are said to discount future illness. Alternatively, the health state may be described as an improvement in health (for example Ganiats et al., 2000), where individuals who prefer health improvement sooner rather than later, and are willing to accept a smaller-sooner improvement in health over a larger-later improvement are said to discount future health. At the time of writing, few previous authors have given an up-to-date synthesis of the findings of these studies (though see Asenso-Boadi, Peters, & Coast, 2008; Chapman, 2003; Chapman, 2005).

2.1 SYSTEMATIC REVIEW METHODS

I reviewed studies of hypothetical health discounting by systematically searching the PubMed Medline database in January 2014 for full text articles published up to January 2014 whose title contained any of the following keywords: ‘Time Preference’, ‘Discounting’, or ‘Intertemporal’ together with either of the words “Health” or “Illness”. This search yielded 87 studies. The abstracts of these were reviewed for suitability: studies were included if they used any delay discounting paradigm with a non-fatal change in health as the outcome. Studies solely measuring discounting of life years were not included, since a large proportion of these actually measure social preferences, and the focus of this review is on personal preferences regarding qualitative future health states. On the basis of these criteria 20 suitable studies were identified. Reference and citation lists of studies meeting the inclusion criteria were also reviewed to identify further studies suitable for inclusion, yielding a further 15 studies. (The 35 studies reviewed are summarised in Table 2.1).

2.2 POSITIVE TIME PREFERENCE FOR HEALTH

Several established properties of monetary discounting are also exhibited in the health domain. Many people discount future health improvements in the conventional sense, termed positive discounting, or positive time preference for health (Bleichrodt & Johannesson, 2001; Lipscomb, 1989; Olsen, 1993). For example, an early study (Lipscomb, 1989) obtained preference scores from 52 undergraduates for a set of health scenarios, covering an hypothetical 50 year period beginning at age 25.

A period of a specified illness, of defined length (e.g. one year), could occur at a variable time into the 50 year period (e.g. at age 26 or at age 28), followed either by excellent health until aged 75 or by immediate death. Delaying onset to the period of illness had a positive effect on the preference scores, such that an illness delayed further into the future tended to be preferred to those where illness occurred sooner, implying a positive rate of discount for illness.

Where positive temporal discounting occurs for health outcomes, existing findings support hyperbolic or quasi-hyperbolic over an exponential discounting, illustrating a consistency with monetary discounting (Bleichrodt & Johannesson, 2001; Van Der Pol & Cairns, 2002; Van Der Pol & Cairns, 2011). For example, Van Der Pol and Cairns (2002) asked participants to make choices relating to a generic illness, which could be delayed into the future by means of an imaginary treatment, showing that hyperbolic discount functions fitted choices better than exponential ones. The same authors (Van Der Pol & Cairns, 2011) later extended their analysis to directly test the stationarity axiom of exponential discounting, as well as axioms predicted from quasi-hyperbolic and hyperbolic models. They found violations of stationarity extending beyond the first time period, favouring a hyperbolic model over quasi-hyperbolic or exponential alternatives.

2.3 ZERO AND NEGATIVE TIME PREFERENCE FOR HEALTH

The observation that many people positively discount hypothetical future health states is at least congruent with an hypothesis that discounting explains unhealthy behaviour. Problematically however positive temporal discounting for health (unlike monetary) outcomes is far from universal. Some individuals rather choose to advance the timing of illness (Cairns, 1992; Chapman, 1996b; Chapman et al., 2001; Chapman & Coups, 1999; Redelmeier & Heller, 1993; Van Der Pol & Cairns, 2002; Van Der Pol & Cairns, 2011; Van Der Pol & Cairns, 2000) or to defer health improvement (Chapman, 1996b; Chapman & Elstein, 1995; P. Dolan & Gudex, 1995; Olsen, 1993). This is the opposite pattern to conventional (positive) discounting, and is termed negative time preference or negative discounting (Loewenstein, 1987; Loewenstein & Prelec, 1991). In addition, some people do not discount hypothetical future health outcomes at all, preferring to experience better

health, irrespective of delay, which is termed zero time preference (Cairns, 1992; Chapman, 1996b; Chapman et al., 2001; Chapman & Coups, 1999; Chapman & Elstein, 1995; P. Dolan & Gudex, 1995; Olsen, 1993; Redelmeier & Heller, 1993; Van Der Pol & Cairns, 2002; Van Der Pol & Cairns, 2011; Van Der Pol & Cairns, 2000).

Across the majority of studies the proportion of individuals exhibiting negative time preference ranges from 3-10%. The study of Olsen (1993), referred to previously, found that 3.6-5.2% of participants had implied negative discount rates. Redelmeier and Heller (1993) using example outcomes of blindness, depression and colostomy, found that 62% of their sample had discount rates of zero and 10% exhibited negative discount rates. Cairns (1992) performed a health discounting task on a small sample of students, in which the health state was described in terms of an episode of depression. Of 18 participants with consistent health discounting results, three had negative discount rates, six had zero discounting and nine had positive discount rates. Chapman and colleagues (Chapman, 1996a; Chapman et al., 2001; Chapman & Coups, 1999; Chapman & Elstein, 1995) found that up to 10% of their samples display negative discount rates for health. Additionally, across repeated studies Van der Pol and Cairns (2000, 2002, 2011) found that 6.7 – 10% of their respondents displayed negative discount rates for illness. A notable outlier is a study performed by P. Dolan and Gudex (1995), who found much higher proportions of negative discounting. Averaged across all health states, 39% of participants displayed negative time preference, preferring sooner illness, and 36% showed zero time preference. The higher rates of negative discounting in this study may be attributable to the fact that health and illness were explicitly presented in sequence, which is known to increase negative time preference (Chapman, 2000; Loewenstein & Prelec, 1993).

An observation related to that of negative time preference, is that monetary losses tend to be (positively) discounted at a lower rate than monetary gains, referred to as the sign effect (e.g. Benzion et al., 1989; Thaler, 1981). There is also evidence for a sign effect in the health domain; for example Chapman (1996b) found that illness was discounted less than health improvement, to the extent that discount rates for illness were uncorrelated with discount rates for health improvement. MacKeigan, Larson, Draugalis, Bootman, and Burns (1993) report a similar result for preferences

over hypothetical improvements or decrements in health framed as a scenario of arthritis, finding positive discounting for health improvement and for long periods of illness, but negative discounting for fleeting illness.

The substantial proportion of individuals displaying negative or zero discounting in choices between hypothetical future health states presents a theoretical stumbling block for a proposed link between discounting and health behaviour. Negative discounting denotes a preference for expediting hardship ('getting things over and done with'), quite the reverse of the procrastinatory mindset hypothesised to be associated with unhealthy behaviour.

2.4 CORRELATIONS BETWEEN HEALTH AND MONETARY DISCOUNTING

In keeping with the high prevalence of negative and zero time preference for health, discount rates for health and money are poorly correlated across individuals (Cairns, 1992; Chapman, 1996b, 2002; Chapman & Elstein, 1995; Lazaro, Barberan, & Rubio, 2001; Petry, 2003). In an early study addressing this question, Cairns (1992) compared discounting for a health state described in terms of depression with monetary discounting. Sub-samples of discount rates within each domain were correlated with each other, however there was no significant correlation between health and monetary discounting. Chapman and Elstein (1995), who also included a third domain of vacations of varying duration, report an overall mean Spearman coefficient of 0.25 for the correlation between discount rates of different domains, indicating a low to medium correlation. In a later study, Petry (2003) studied discount rates for hypothetical health, money and freedom (from time spent in jail), finding that health and freedom discount rates were correlated, but neither domain was significantly correlated with monetary discount rates.

Chapman and Elstein (1995) suggested three possible confounding explanations for the low cross-domain correlations in discounting in these studies. Firstly, the authors proposed that the difference in discount rates across domains might be accounted for by the sign-effect, if health and monetary outcomes were not matched for sign. Secondly, they suggested that domain independence might be attributable to a phenomenon known as the magnitude effect, an observed violation of the discounted utility model whereby larger amounts are discounted less than smaller

amounts (Benzion et al., 1989; Green et al., 1994; Loewenstein, 1988; Thaler, 1981). This could lead to higher discount rates for health if health outcomes were deemed less valuable relative to monetary outcomes. Finally, Chapman and Elstein noted that apparently different discounting could result if the instantaneous utility functions for health and money differed in shape. Chapman (1996b) addressed these possible confounds, demonstrating that low correlations between health and monetary domains persist even when the outcomes are matched for sign as well as utility. Correlations between the two domains overall showed a trend towards being higher when matched for sign, although all correlations remained weak (Spearman r 0.06-0.28). In other words discount rates for illness showed marginally stronger correlations with monetary losses than with monetary gains, and discount rates for health improvement showed marginally stronger correlations with monetary gains than with monetary losses, though all such correlations are weak and most do not reach conventional levels of significance.

2.5 CONCLUSIONS

To sum up, discounting of hypothetical health is at best weakly correlated with monetary discounting (e.g. Chapman 1996, Spearman r 0.06-0.28) and in such studies a high proportion of individuals either appear not to discount future health at all, or prefer to delay good health. The latter finding appears to challenge the hypothesis that unhealthy choices in some sense *result from* the discounting of delayed health. It remains possible however that people represent health outcomes differently when making real health choices as opposed to responding in hypothetical health discounting tasks.

2.5.1 EXPERIENTIAL OUTCOMES AND ZERO TIME PREFERENCE

The high proportion of zero time preference for hypothetical health may derive from the fact that the health outcomes are described in experiential terms. As mentioned above, unlike money, health cannot be invested, removing a key normative motive for discounting (see Chapter 10: General Discussion). Interestingly Chapman

(2002) demonstrated that higher correlations (Pearson correlation coefficient $r = 0.41$) between health and monetary discount rates result if choices are embedded in a scenario in which health and money are framed as tradable, compared with their being framed as non-tradable (Pearson correlation coefficient $r = 0.07$), suggesting that thinking of health as a commodity leads to health being discounted in a more conventional manner.

2.5.3 DREAD AND NEGATIVE TIME PREFERENCE

Negative discounting for illness outcomes is not unexpected. In choices between genuine delayed painful events, many people prefer to expedite inevitable pain (negative time preference) (e.g. Berns et al., 2006; Cook & Barnes, 1964; Loewenstein, 1987; Loewenstein & Prelec, 1991). Negative time preference for pain has been attributed to the fact that anticipating delayed pain is itself aversive, termed dread (Berns et al., 2006; Loewenstein, 1987). This unpleasant anticipation can be minimised by choosing to “get the pain out of the way”. Consistent with respondents exhibiting dread for delayed illness, in an illness discounting task, individuals who perceived the illness as more severe were more likely to have negative discount rates (Van Der Pol & Cairns, 2000). Despite authors having previously observed that negative discounting of health is consistent with dread (Van Der Pol & Cairns, 2000), no previous studies have examined the functional form of negative time preference. I address this in Chapter 6 of this thesis.

2.5.2 DISCOUNTING AS A TRAIT

The high rates of negative and zero discounting for health call into question the notion of discounting as a unitary construct, pertinent to a trait interpretation of discounting. In a recent analysis, Odum (2011) evaluates the extent to which discounting fulfills criteria for a personality trait. Odum adopts a commonly used definition of a trait as ‘a relatively enduring pattern of thoughts, feelings, and behaviours that reflects the tendency to respond in certain ways under certain circumstances’ (Roberts, 2009). Notably, discounting for monetary outcomes

robustly fulfills the first criterion, namely that a trait should be ‘relatively enduring’, since the test-retest reliability of monetary discounting is high at intervals of up to one year (Pearson $r = 0.71$; Kirby, 2009) and across different elicitation methods (Odum, 2011). Furthermore, Odum demonstrates that discounting for several forms of outcome, excluding health, shows consistency across individuals. Discount rates are not *identical* across commodities: people tend to discount primary reinforcers such as food and water more steeply than money (Odum, 2011) and a number of studies have shown that people with substance dependence discount their drug of abuse more steeply than money (e.g. Bickel, Odum, & Madden, 1999; Madden, Petry, Badger, & Bickel, 1997; Petry, 2001). However, discount rates for money are strongly and significantly correlated with other forms of appetitive outcome, such as the discounting of cigarettes for cigarette smokers, the discounting of heroin for opioid-dependent outpatients and the discounting of food amongst college students (Odum, 2011; Pearson $r = 0.93$; $p = 0.0007$ for money versus the mean of all other outcomes). There is also evidence that discounting is heritable (for a review see MacKillop, 2013). A recent longitudinal twin study estimated the heritability of delay discounting in adolescence at up to 50% (Anokhin, Golosheykin, Grant, & Heath, 2011), rats and mice can be bred to exhibit greater degrees of delay discounting (e.g. K. G. Anderson & Woolverton, 2005; Belin, Mar, Dalley, Robbins, & Everitt, 2008), and steeper discounting in humans is associated with specific polymorphisms related to dopamine signaling (Eisenberg et al., 2007). Thus, for appetitive outcomes at least, discounting has characteristics consistent with a trait variable. However, the low correlations between health and monetary discounting seem to suggest a limitation to the scope of the discounting-trait, and suggest that hypothetical health discounting tasks engage different processes from those with appetitive outcomes.

Table 2.1 *Studies Measuring Discount Rates for Health Outcomes*

Study	Outcome health state	Baseline	Frame	Discounting analysis	N	Discount rate measure	Central tendency measure	Average health discount rate	% Negative	% Zero	Stationarity
Lipscomb (1989)	Generic illness (social)	Full health	Loss ^a	Preference score	52	N/A	N/A	N/A	N/A	N/A	N/A
Cairns (1992)	Depression	Full health	Loss	Open ended (days of illness)	29	Exp: annual r	Mean	-0.15 – +0.13	17% ^b	33% ^c	N/A
Redelmeier & Heller (1993)	Colostomy, blindness or depression	Usual health	Loss	Standard gamble and categorical scaling valuations	121	Exp: annual r	Mean	0.033	10%	62%	r decreased with delay
Dolan & Gudex (1995)	One year generic illness	Full health	Loss	Time trade-off valuations	39	Exp: annual r	Mean	-0.035 – +0.02	39%	36%	N/A
Chapman (1996) (2a)	One year generic illness	Full health	Loss	Open ended (years of illness)	77	Exp: annual r	Geometric mean	0.50	10%	n/s	r decreased with delay
Chapman & Coups (1999)	Influenza illness	Usual health	Loss	Closed ended (days of illness)	412	N/A	Mean	1.00	N/A	85%	n/s
Cairns & Van der Pol (1999)	Generic illness	Usual health	Loss	Open ended (days of illness)	150	n/s	Median	0.06	9% ^d	13% ^e	r decreased with delay
Van der Pol & Cairns (2000)	Generic illness (social and private)	Usual health	Loss	Open ended (days of illness)	308	N/A	N/A	N/A	12.8%	6.8%	n/s

Van der Pol & Cairns (2001)	Generic illness	Usual health	Loss	Discrete choice (days of illness)	787	Exp: annual r	Mean	0.04 -0.10	$>24\%^f$	$>7.5\%^g$	n/s
Bleichrodt & Johannesson (2001)	Back pain	Full health	Loss	Preference scores	172	N/A	N/A	N/A	N/A	N/A	Violation of stationarity
Chapman et al (2001) (1)	Influenza illness	Usual health	Loss	Closed ended (days of illness)	679	N/A	N/A	Negative	n/s	n/s	n/s
Chapman et al (2001) (2)	Intermittent chest pain	Illness	Loss	Closed ended (frequency of pain)	195	N/A	N/A	n/s	n/s	n/s	n/s
Odum, Madden & Bickel (2002) (1)	Sexually transmitted illness	Full health	Loss	Closed ended (days to years of illness)	66	Hyp: k	Median	0.01	n/s	n/s	Hyp fitted better than exp
Van der Pol & Cairns (2002)	Generic illness	Usual health	Loss	Open ended (days of illness)	159	Exp: annual r	Mean	-0.02 - +0.08	6.7%	n/s	Hyp fitted than exp
Petry (2003)	Generic illness	Full health	Loss	Open ended (continuous duration)	141	Hyp: k	Median	0.00 ^b	n/s	n/s	Hyp only fitted
Robberstad (2005)	Malarial illness	Usual health	Loss	Open ended (days of illness)	450	Exp: annual r	Mean	0.07	n/s	n/s	r decreased with delay
Robberstad & Cairns (2007)	Non-fatal illness	Usual health	Loss	Open ended (days of illness)	450	Exp: annual r	Median	0.05 - 0.06	n/s	n/s	Hyp fitted better than exp
Van der Pol & Cairns (2008)	Generic illness	Usual health	Loss	Open ended (days of illness)	1064	Exp: annual r	Mean	0.04 - 0.11	n/s	n/s	r decreased with delay
Van der Pol & Cairns (2008)	Generic illness	Usual health	Loss	Discrete choice (days of illness)	1064	Exp: annual r	Mean	0.02 -0.03	n/s	n/s	r decreased with delay

Van der Pol & Cairns (2011)	Generic illness	Usual health	Loss	Open ended (days of illness)	203	N/A	N/A	10%	n/s	Violation of stationarity
Olsen (1993)	Number of people's health improved from public program (social)	Usual health	Gain	Opened ended (number of people)	250	Exp: continuous <i>r</i>	Mean	0.07 - 0.15 3.6 - 5.2%	7.2-10%	<i>r</i> decreased with delay
Chapman & Elstein (1995) (1)	Treatment providing full health from generic illness	Illness	Gain	Opened ended (years of health)	70	Exp: annual <i>r</i>	Mean	8.93 n/s	n/s	<i>r</i> decreased with delay
Chapman & Elstein (1995) (2)	Treatment providing full health from generic illness	Illness	Gain	Opened ended (years of health)	31	Exp: annual <i>r</i>	Mean	1.22 3% overall	n/s	<i>r</i> decreased with delay
Chapman (1996) (1)	Treatment providing full health from generic illness	Illness	Gain	Opened ended (years of health)	40	Exp: annual <i>r</i>	Mean	0.75 3% overall	n/s	<i>r</i> decreased with delay
Chapman, Nelson & Hier (1999)	Treatment for migraine or Crohn's disease	Illness	Gain	Closed ended (frequency of headache or	74	Exp: annual <i>r</i>	Mean	1.00 - 190.00 n/s	n/s	<i>r</i> decreased with delay
Ganiats et al (2000)	Treatment providing total relief from migraine	Illness	Gain	Discrete choice (months of relief)	141	Exp: annual <i>r</i>	Median	1.16 n/s	n/s	n/s
Lazaro et al (2001)	Treatment providing full health from generic illness	Illness	Gain	Open ended (years of health)	203	Exp: annual <i>r</i>	Mean	0.15 - 0.395 <1% overall	n/s	<i>r</i> decreased with delay
Khawaja, Silverman & Sloan (2007)	Improvement to full health	Usual health	Gain	Open ended (days of health)	431	Exp: continuous <i>r</i>	Mean	0.05 -0.50 n/s	n/s	<i>r</i> decreased markedly with delay
Chapman (1996) (2b)	Treatment providing full health from generic illness	Illness	Gain	Opened ended (years of health)	77	Exp: annual <i>r</i>	Geometric mean	1.50 10%	n/s	<i>r</i> decreased with delay
Chapman (2002) (1)	Treatment providing total relief from migraine (social)	Illness	Gain	Open ended (months of relief)	205	Exp: annual <i>r</i>	Mean	0.37-0.41 n/s	n/s	n/s

Chapman (2002) (2)	Treatment providing total relief from migraine (social)	Illness	Gain	Open ended (months of relief)	317	Exp: annual r	Mean	0.43	n/s	n/s	n/s
Odum, Madden & Bickel (2002) (2)	Treatment providing full health from sexually transmitted illness	Illness	Gain	Closed ended (days to years)	66	Hyp: k	Median	0.01	n/s	n/s	Hyp fitted better than exp
Chapman (1996) (3)	Gains: treatment providing full health from generic illness; losses: generic illness	Gains: Illness Losses: Full health	Gain and loss	Open ended (years of health)	38	Exp: annual r	Geometric mean	0.15	8% overall	n/s	r decreased with delay
Mackeigan et al (1993)	Arthritis	N/A	Gain and loss	Preference scores	108	N/A	N/A	N/A	N/A	N/A	N/A
Baker, Johnson & Bickel (2003)	Gains: 10% improvement in health; losses: 10% decline in health	Usual health	Gain and loss	Closed ended (days of improved or reduced health)	60	Hyp: k	Median	<0.01	n/s	n/s	Hyp only fitted

Notes

Abbreviations: N/A=full text of study not available; n/s=not stated in the paper; Hyp=hyperbolic discount function, Exp=exponential discount function.

‘Baseline’ refers to the hypothetical health state when not receiving the outcome. Unless otherwise stated discount rates are those for private preferences (own health). In studies involving comparator groups, rates stated are for the control group.

- Studies are grouped according to loss or gain frame.
- 3/18 (17%) participants with consistent time preferences had negative discount rates
- 6/18 (33%) participants with consistent time preferences had zero discount rates
- 14/150 (9%) participants showed >3 out of 6 responses implying negative discount rates and no responses implying positive discount rates
- 19/150 (14%) participants showed >4 out of 6 responses implying zero discount rates
- 92/385 (24%) participants had dominant preferences for shortest delay illness
- 29/385 (7.5%) participants had dominant preferences for shortest duration illness
- 0.003

RELATIONSHIPS BETWEEN DISCOUNTING AND HEALTH BEHAVIOUR

“Live fast, die young.”

Willard Motley – Knock on Any Door (1947)

The low correlations between health and monetary discounting, and the high prevalence of negative and zero time preference for health described in the preceding chapter would appear to cast doubt on an hypothesis that health behaviour results from discounting of delayed health in a strictly mechanistic sense. Nevertheless, either measure might prove capable of predicting real-world health behaviour. In fact, authors addressing the hypothesis that discounting predicts health behaviour have tended to simply use conventional monetary discounting, with promising findings. In this chapter I perform systematic review of studies which examine relationships between temporal discount rates and health behaviour or health status.

3.1 SYSTEMATIC REVIEW METHODS

The PubMed database was searched in January 2014 for full-text articles published up to January 2014, containing any of the keywords ‘Discounting’, ‘Intertemporal’, ‘Time Preference’, together with the words ‘Health’, ‘Illness’, ‘BMI’, ‘Obesity’, ‘Alcohol’, ‘Drinking’, ‘Smoking’, ‘Drug’ or ‘Behaviour’. This initial search yielded 104 studies. The abstracts, and where necessary full texts, of these studies were then examined to determine suitability for inclusion, defined as any study which compared the results of a delay discounting paradigm to an observed health-related measure or behaviour. Altogether 34 suitable full text articles were identified by this method. The references and citation lists of these studies were then

reviewed for inclusion, yielding a further 78 suitable studies, making a total of 112 suitable studies (summarised in Tables 3.1 to 3.5). Several of these studies have been included in previously published reviews (Bickel et al., 2012; Bickel, Koffarnus, Moody, & Wilson, 2014; Bickel & Marsch, 2001; Koffarnus, Jarmolowicz, Mueller, & Bickel, 2013; MacKillop et al., 2011; Reynolds, 2006b). However, these have been chiefly in the context of addiction – here a broader range of health outcomes are reviewed. Below, studies are organised first by the modality of discounting, hypothetical health versus money or other appetitive rewards. Within each modality studies are organised by the nature of the health outcome(s), namely tobacco smoking, alcohol use, illicit substance misuse, obesity and eating behaviour, preventive health behaviour, risky sexual behaviour and drug-taking practices or multiple health behaviours.

3.2 HYPOTHETICAL HEALTH DISCOUNTING AND HEALTH BEHAVIOUR

Four studies comparing discount rates for hypothetical health outcomes with observed health behaviour were identified. Only one of these four studies found a statistically significant correlation. Two of the studies examined relationships between health discounting and cigarette smoking. Baker, Johnson, and Bickel (2003) measured health discounting in current and never-before cigarette smokers, finding that current smokers discounted health at a marginally higher rate than the never-before smokers, but this difference did not reach statistical significance. By contrast current smokers discounted money at a significantly higher rate than the never before smokers. Similarly, Khwaja, Silverman, and Sloan (2007) found no significant differences between smokers and non-smokers in health discounting for health outcomes (however neither did this study find a relationship between smoking status and monetary discounting).

A single study in the sample (Petry, 2003) examined discount rates for health outcomes in substance misusers, finding significantly higher discount rates for hypothetical health, money and freedom (from time spent in jail) in a group of current or previous heroin and/or cocaine users compared with a group of controls with no history of substance misuse in their lifetime. Finally, Chapman and Coups (1999) asked whether discount rates for monetary losses, described in terms of a parking fine,

and for a flu-like illness, could explain uptake of free-of-charge influenza vaccinations. The respondents were future oriented, with 85% not significantly discounting the flu-like illness (i.e. showing zero time preference) and 83% not significantly discounting parking fines. Monetary time preferences were related to vaccine uptake: 45% of those with no discounting accepted the vaccine, compared with 29% of those who discounted money in the conventional manner. However, health discount rates were unrelated to vaccine acceptance.

The observed higher health discount rates in substance misusers compared with controls (Petry, 2003) and the trend towards higher health discounting in current cigarette smokers (Baker et al., 2003), taken together with the larger between-group differences in monetary discounting, suggest that health discounting does exhibit a weak relationship with health behaviour, but is a less sensitive predictor than monetary discounting.

3.3 MONETARY OR APPETITIVE DISCOUNTING AND HEALTH BEHAVIOUR

Although discounting for hypothetical health appears to be a weak predictor of observed health behaviour, many studies in the search sample demonstrate that discount rates for money or other appetitive outcomes such as food or drug rewards correlate more substantially with health behaviour or health status. The key findings are discussed below, grouped by health outcome.

3.3.1 DISCOUNT RATES AND SELF-REPORTED HEALTH

Self-reported health is perhaps the most general health outcome measure and correlates with life-expectancy in the developed world (Idler & Benyamini, 1997). A recent household survey in the Netherlands ($N = 2300$) (Van Der Pol, 2011) found that, in a multiple linear regression, higher monetary discount rates were weakly associated with poorer self-reported health status. A further study in a South African population (Chao, Szrek, Pereira, & Pauly, 2009) found evidence for a U-shaped relationship between subjective health and monetary discount rates, whereby those who reported “average” health had lower discount rates than those who are either very

healthy or very sick. The authors suggested that this might be due to the fact that those with very poor health were in more urgent need of money to fund medical care, whereas those in excellent health were able to enjoy the benefits of immediate economic consumption, highlighting the considerable difficulties in establishing a casual pathway between discounting and health.

3.3.2 DISCOUNT RATES AND CIGARETTE SMOKING

Monetary discount rates have consistently been shown to be higher in people who currently smoke tobacco than in non-smokers (Bickel et al., 1999; S. Fields, Leraas, Collins, & Reynolds, 2009; Kobiella et al., 2014; Rezvanfard et al., 2010; S. H. Mitchell, 1999; Odum, Madden, & Bickel, 2002; Reynolds, 2006a; Reynolds, Leraas, Collins, & Melanko, 2009; Reynolds, Patak, & Shroff, 2007; Reynolds, Penfold, & Patak, 2008; Stillwell & Tunney, 2012; Wing, Moss, Rabin, & George, 2012). A recent meta-analysis of studies comparing discount rates with addictive behaviours, (MacKillop et al., 2011) estimated a moderate and highly significant effect (Cohen's $d = 0.57$ $p < 0.0001$) across all studies comparing discount rates in smokers versus non-smokers.

Monetary discount rates also correlate with smoking frequency (L. H. Epstein et al., 2003; Fuchs, 1982; Kang & Ikeda, 2013; Ohmura, Takahashi, & Kitamura, 2005). In keeping with this infrequent smokers exhibit monetary discount rates intermediate between heavy smokers and non-smokers (Heyman & Gibb, 2006; Reynolds & Fields, 2012; Stillwell & Tunney, 2012 $N = 9454$) (however see M. W. Johnson, Bickel, & Baker, 2007 for negative findings; Reynolds, Karraker, Horn, & Richards, 2003) and both smoking frequency and monetary discount rates were found to be higher in a group of young-adult smokers than in a group of adolescent smokers (Reynolds, 2004). The relationship between smoking frequency and discounting does not appear to be mediated by the acute effects of nicotine, since acute nicotine administration to non-smokers has recently been shown to have no effect on intertemporal choice behaviour (Kobiella et al., 2014). However the relationship may be related to the level of nicotine dependence (Sweitzer, Donny, Dierker, Flory, & Manuck, 2008), consistent with discounting being a state-based marker of addiction severity.

Interestingly those who have previously smoked and those who have never smoked do not significantly differ in their monetary discounting (Bickel et al., 1999). Furthermore in a prospective study of smoking cessation, participants were separated into a group who received an intervention program directed at reducing smoking and a control group who continued to smoke as usual. The two groups did not differ in their discounting behaviour at baseline. Whilst the control group showed no changes in discounting over time, the intervention group (who reduced their smoking frequency) showed a significant decrease in discounting for both money and cigarettes after only five days into the program (Yi & Landes, 2012). These results strongly suggest that the state of nicotine addiction acts reversibly to increase discount rates. However there is also evidence to support the idea that discounting operates as an antecedent vulnerability marker, since monetary discount rates in smokers predict rates of smoking adoption (Audrain-McGovern et al., 2009). Monetary discount rates in smokers also predict rates of relapse within smoking cessation programs (Brown & Adams, 2013; Krishnan-Sarin et al., 2007; MacKillop & Kahler, 2009; Sheffer et al., 2012; Yoon et al., 2007), and the ability to abstain from smoking under laboratory conditions (Dallery & Raiff, 2007; Mueller et al., 2009). It seems reasonable to conclude therefore that relationships between discounting and smoking behaviour are subject to both state- and trait-based influences (de Wit, 2009; Odum, 2011).

Table 3.1

*Studies Comparing Discount Rates and Smoking Behaviour (N = 39)**

Study	Health outcome	Modality	Real/ Hypothetical	N	DF
Mitchell (1999)	Smokers vs non smokers	Money	Real (IC)	40	HYP, IP
Bickel, Odum & Madden (1999)	Current smokers vs ex-smokers vs never smokers	Money, cigarettes	Hypothetical	66	HYP
Odum, Madden & Bickel (2002)	Current smokers vs ex-smokers vs never smokers	Health	Hypothetical	66	EXP, HYP
Reynolds et al (2003)	Never smokers vs “triers” vs current smokers	Money	IC	55	HYP
Epstein et al (2003)	Smokers by amount smoked	Money	Hypothetical	78	HYP
Baker, Johnson & Bickel (2003)	Smokers vs never smokers	Health, money, cigarettes	Both	60	HYP
Mitchell (2004)	Nicotine deprivation in smokers	Money, cigarettes	Real (IC)	11	HYP, IP
Reynolds et al (2004)	Smokers vs never smokers	Money	Real	54	HYP
Reynolds (2004)	Adolescent and young adult smokers vs young adult non- smokers	Money	Hypothetical	73	HYP
Ohmura, Takahashi, Kitamura (2005)	Smokers vs never smokers; amount smoked	Money	Hypothetical	50	AUC, EXP, HYP
Reynolds (2006)	Smokers vs non smokers	Money	Real (EDT)	30	HYP
Heyman & Gibb (2006)	Smokers stratified by amount smoked	Money	Both	71	HYP
Field et al (2006)	Smokers vs abstinent smokers	Money, cigarettes	Hypothetical	30	AUC
Johnson, Bickel & Baker (2007)	Smokers vs non smokers; amount smoked	Money	Hypothetical	90	HYP
Reynolds et al (2007a)	Smokers vs non smokers (adolescents)	Money	Hypothetical	51	AUC
Reynolds et al (2007b)	Smokers vs non smokers (adolescents)	Money	Hypothetical	80	AUC
Khwaja, Silverman & Sloan (2007)	Smokers vs ex-smokers vs never smokers	Money, health	Hypothetical	431	EXP
Yoon et al (2007)	Pregnant women and postpartum smoking relapse	Money	Hypothetical	48	HYP
Dallery & Raiff (2007)	Laboratory model of smoking abstinence reinforcement	Money	Hypothetical	30	HYP, AUC
Krishnan-Sarin et al (2007)	Smoking cessation in adolescents	Money	Both (EDT)	30	HYP
Sweitzer et al (2008)	Never smokers vs “tried it” smokers vs regular smokers	Money	Hypothetical	710	HYP
Yi et al (2008)	Effects of smoking cessation on discounting	Money, cigarettes	Hypothetical	56	HYP
Mueller et al (2009)	Smoking abstinence in the laboratory	Money, cigarettes	Hypothetical	19	HYP

Finally short-term abstinence from cigarettes increases discounting in addicted smokers (Field, Santarcangelo, Sumnall, Goudie, & Cole, 2006; S. H. Mitchell, 2004; Yi & Landes, 2012). For example, Field et al. (2006) measured discounting for hypothetical gains of money or cigarettes in a group of 30 smokers: one randomised group performed the procedures following their usual cigarette consumption, the other following a minimum of 13 hours of abstinence from cigarettes. Implied discount rates for both money and cigarettes were significantly higher in the abstinence group.

In summary, monetary discount rates are higher in current smokers, correlate with smoking frequency and prospectively predict the adoption of smoking and abstinence from smoking. The upward effect of nicotine cravings on discount rates and the decrease in discounting concomitant with reductions in smoking indicate that discounting is influenced by state-based environmental and motivational factors. Taken together these results indicate that relationships between discounting and smoking have both state- and trait- based components, a consideration which most likely also applies to relationships between discounting and other addictive behaviours (de Wit, 2009; Odum, 2011).

3.3.3 DISCOUNT RATES AND ALCOHOL USE

Monetary discount rates exhibit consistent relationships with alcohol use. A recent meta-analysis of studies comparing discount rates in persons meeting clinical criteria of an alcohol dependence syndrome with controls (MacKillop et al., 2011) demonstrated a moderate, highly significant effect (Cohen's $d = 0.50$ $p < 0.0001$). Monetary discount rates are higher in currently abstinent alcohol dependent individuals compared with non-dependent controls (Bjork, Hommer, Grant, & Danube, 2004; Boettiger et al., 2007; J. M. Mitchell, Fields, D'Esposito, & Boettiger, 2005), in early-onset as opposed to late-onset alcohol dependence (Dom, De Wilde, Hulstijn, Van Den Brink, & Sabbe, 2006), and correlate with the severity of alcohol dependence (J. M. Mitchell et al., 2005), as well as symptoms of an alcohol abuse disorder (MacKillop et al., 2010). Monetary discount rates have also been shown to be higher in a group with a previous lifetime diagnosis of an alcohol abuse disorder as compared with those without a lifetime history of alcohol abuse (Bobova, Finn, Rickert, & Lucas, 2009).

Table 3.2

Studies Comparing Discount Rates and Alcohol Use (N=17)

Study	Health outcome	Modality	Real/ hypothetical	N	DF
Vuchinich & Simpson (1998)	Light social drinkers vs heavy social drinkers vs problem drinkers	Money	Hypothetical	48	HYP
Petry (2001)	Alcoholics vs abstinent alcoholics vs controls	Money, alcohol	Hypothetical	46	HYP
Bjork et al (2004)	Abstinent alcohol dependent patients vs controls	Money	IC	160	HYP
Dom et al (2006)	Early vs late onset alcoholics vs controls	Money	Hypothetical	142	HYP
Mitchell et al (2005)	Abstinent alcoholics vs controls	Money	Hypothetical	31	E:L
Boettiger et al (2007)	Alcoholics vs controls	Money	Hypothetical	19	E:L
Mackillop et al (2007)	Social drinking vs hazardous drinking (students)	Money	Hypothetical	93	HYP
Field et al (2007)	Light vs heavy drinkers (adolescents)	Money, alcohol	Hypothetical	66	AUC
Rosow (2008)	Frequency of alcohol use, adverse consequences (adolescents)	Money	Hypothetical	17413	EXP
Bobova et al (2009)	Alcohol dependence, +/- previous childhood conduct disorder vs controls (young adults)	Money	Hypothetical	426	HYP
Mackillop et al (2010)	Heavy drinkers	Money	Hypothetical	61	MCQ
Moore & Cusens (2010)	Blood alcohol levels (social drinkers)	Money	Hypothetical	46	HYP
Fernie et al (2010)	Measures of alcohol use (social drinkers)	Money	Hypothetical	75	AUC
Dennhardt & Murphy (2011)	Frequency of alcohol use, adverse consequences (undergraduates)	Money	Hypothetical	206	HYP
Yankelevitz, Mitchell & Zhang (2011)	Alcohol consumption (college students)	Money, alcohol	Hypothetical	65	AUC
Moallem & Ray (2012)	Heavy drinkers, smokers and heavy drinking smokers	Money	Hypothetical	387	HYP
Andrade et al. (2013)	History of alcohol problems or smoking in problem gamblers	Money	Hypothetical	333	HYP

Note Studies are listed in chronological order. Abbreviations: EXP = Exponential function; HYP = Hyperbolic function; AUC = Area under curve; DF = Discount function; MCQ = Monetary Choice Questionnaire, a discounting method (Kirby et al., 1999); IC = Incentive compatible (a design in which a proportion of choices have real outcomes); E:L = early:late choice ratio

Several studies have further linked higher discounting with relatively moderate levels of alcohol consumption. For example, Vuchinich & Simpson, 1998 demonstrated higher monetary discount rates in “problem drinkers” and also heavy social drinkers, compared with light social drinkers, suggesting a relationship between alcohol intake and discount rates even amongst those designated as social drinkers. Similarly, Field, Christiansen, Cole, and Goudie (2007) found that delay discounting for alcohol positively correlated with weekly alcohol consumption (Pearson $r = 0.31$) amongst adolescents, where those in highest tertile of alcohol use had a mean weekly consumption of 23 units, while those in the lowest tertile had a mean of 3 units. Finally, Yankelevitz, Mitchell, and Zhang (2012) found that implied discounting for both money and alcohol was moderately correlated with levels of alcohol use in female students (Pearson $r = 0.43$ for money and 0.41 for alcohol discounting), though no correlation was found in male students.

Amongst students, monetary discount rates appear related to adverse consequences of alcohol use (Dennhardt & Murphy, 2011; Kollins, 2003; Rossow, 2008). Kollins (2003) observed that monetary discount rates were negatively correlated with age at first using alcohol and showed a strong positive correlation with the number of times that students had “passed out” as a result of alcohol use (Pearson $r = 0.73$, $P < 0.01$) and Rossow (2008), studying a sample of 17,413 adolescents in Norway, demonstrated that high monetary discounters became intoxicated more frequently and were more likely to vomit or “pass out” as a result of drinking. Paralleling findings in previous smokers of cigarettes, previously addicted users of alcohol who have achieved long-term abstinence have discount rates intermediate between current users and controls (Petry, 2001). Finally, in an elegant field study, the discount rates of male social-drinkers on their entry to a bar prospectively predicted increases in blood alcohol level on their exit (Moore & Cusens, 2010), such that those that had steeper discounting on entry showed greater increases in alcohol level. Discount rates were not confounded by baseline intoxication, since blood alcohol level at entry to the bar did not significantly predict the baseline discount rates.

In summary, despite a minority of studies reporting negative findings (Ferne, Cole, Goudie, & Field, 2010; Kirby & Petry, 2004; MacKillop, Mattson, Anderson Mackillop, Castelda, & Donovan, 2007), monetary discount rates (and in some studies discount rates for alcohol) show robust relationships with alcohol intake over

a wide range of usage, being higher in currently dependent individuals, where they correlate with the degree of dependence, and predicting use in non-dependent individuals.

3.3.4 DISCOUNTING AND ILLICIT SUBSTANCE MISUSE

Amongst health behaviours, illicit substance misuse exhibits the most consistent relationships with discount rates. In an early study, heroin dependent individuals exhibited monetary discount rates twice those of non-drug-using controls (Kirby, Petry, & Bickel, 1999). Several other studies have also demonstrated significantly higher monetary discount rates in opioid-dependent individuals compared with controls (Kirby & Petry, 2004; Madden et al., 1997). Monetary discount rates are also higher in users of stimulant drugs such as cocaine and methamphetamine than in non-drug-using controls (Coffey, Gudleski, Saladin, & Brady, 2003; Heil, Johnson, Higgins, & Bickel, 2006; Hoffman et al., 2006; M. W. Johnson, 2012; Kirby & Petry, 2004; Moeller & Dougherty, 2002; Monterosso et al., 2007), with one study finding significantly higher monetary discount rates among individuals primarily using crack cocaine than among those primarily using heroin (Bornovalova, Daughters, Hernandez, Richards, & Lejuez, 2005). Indeed, Mackillop et al. (2011) estimated a large and highly significant aggregate effect (Cohen's $d = 0.87$, $p < 0.0001$) across studies comparing discount rates in dependent users of stimulant drugs versus controls and a moderate highly significant effect across studies comparing discount rates in opiate dependent individuals versus controls (Cohen's $d = 0.76$, $p < 0.0001$).

Consistent with nicotine abstinence studies, mild opioid deprivation in opioid dependent individuals increases discounting for money as well as heroin (Giordano et al., 2002), whereas those who have achieved longer term abstinence from heroin have lower discount rates than those currently addicted (Bretteville-Jensen, 1999; Kirby & Petry, 2004). By contrast abstinent formerly dependent cocaine users do not differ in discounting behaviour from current users (Kirby & Petry, 2004; Heil et al., 2006), and Hoffman et al. (2006) found no relationship between length of abstinence and monetary discount rates in amphetamine dependent individuals. These findings suggest that either discount rates do not predict abstinence from stimulants, or that

addiction to stimulants can have an irreversible effect to increase discount rates. Evidence against the former suggestion is that baseline discounting has been shown to predict the duration of abstinence from cocaine under a contingency management intervention (with low-incentives but not with high-incentives) (Washio et al., 2011).

Monetary discounting has not been consistently associated with concurrent cannabis use. M. W. Johnson et al. (2010) found that discount rates for hypothetical money in a group of marijuana dependent individuals did not differ from non-drug using controls, despite their study being adequately powered to detect any such difference (see also Stea, Hodgins, & Lambert, 2011). Similarly Heinz, Peters, Boden, and Bonn-Miller (2013) found that monetary discounting did not correlate significantly with frequency of cannabis use over a 90-day period, although higher discounting was associated with younger age at first cannabis use. A recent study has shown that discount rates for hypothetical large monetary amounts (\$1000) prospectively predicted abstinence outcomes amongst adolescents undergoing treatment for marijuana dependence (Stanger et al., 2012), more recent studies have (Heinz et al., 2013; E. N. Peters, Petry, LaPaglia, Reynolds, & Carroll, 2013) found that discount rates did not predict response to a similar intervention in adults.

Finally studies have demonstrated an additive effect of smoking and alcohol use on discounting (Moallem & Ray, 2012; see also Andrade & Petry, 2012) but not of smoking and other forms of substance misuse (Businelle, McVay, Kendzor, & Copeland, 2010), and the combination of gambling problems and substance misuse appears highly predictive of impulsive choice (however see Ledgerwood, Alessi, Phoenix, & Petry, 2009; Petry, 2001; Petry & Casarella, 1999). In summary, with the exception of cannabis use, monetary discount rates consistently show strong correlations with the use of illicit substances.

Table 3.3

Studies Comparing Discount Rates and Illicit Substance Misuse (N = 27)

Study	Health outcome	Modality	Real/ hypothetical	N	DF
Madden et al (1997)	Opiate dependent patients vs controls	Money, heroin	Hypothetical	56	HYP
Kirby, Petry & Bickel (1999)	Heroin users vs controls	Money	Real (IC)	116	HYP
Bretteville-Jensen (1999)	Active heroin users vs former heroin users vs controls	Money	Hypothetical	270	EXP
Petry & Casarella (1999)	Substance abusing gamblers vs substance abusing non gamblers vs controls	Money	Hypothetical	81	HYP
Petry (2001)	Gamblers with substance misuse vs gamblers without substance misuse	Money	Hypothetical	50	HYP
Giordano et al (2002)	Opiate deprivation in opioid dependent patients	Money, heroin	Hypothetical	13	HYP
Moeller et al (2002)	Cocaine dependent subjects vs controls	Money	Real	74	SCF
Petry (2003)	Current or former substance users vs controls	Health, money, freedom	Hypothetical	141	HYP
Kollins (2003)	Substance abuse (students)	Money	Hypothetical	47	HYP
Coffey et al (2003)	Cocaine users vs controls	Money, cocaine	Hypothetical	25	HYP
Kirby & Petry (2004)	Heroin & cocaine users vs alcoholics vs controls ; abstinence >14days	Money	Real (IC)	145	HYP
Bornovalova et al (2005)	Users of primary crack cocaine vs users of primary heroin	Money	Hypothetical	27	HYP
Heil et al (2006)	Cocaine dependent current users vs dependent abstinent for 30 days vs controls	Money	Hypothetical	63	HYP
Hoffman et al (2006)	Dependent metamphetamine users (2 week abstinent) vs controls	Money	Hypothetical	82	HYP
Monterosso et al (2007)	Methamphetamine users vs controls	Money	Hypothetical	29	HYP
Ledgerwood et al (2009)	Gamblers with substance misuse vs gamblers without substance misuse vs controls	Money	Hypothetical	101	AUC
Businelle et al (2010)	Smokers with SUD (including alcohol) vs smokers without SUD vs never smokers with SUD vs controls (never smoked, no SUD)	Money	Hypothetical	115	HYP
Johnson et al (2010)	Current vs former marijuana users vs controls	Money, marijuana	Hypothetical	88	HYP
Washio et al (2011)	Prospective study of treatment for cocaine-dependent individuals	Money	Hypothetical	36	HYP
Stea, Hodgkins & Lambert (2011)	Cannabis, gambling and alcohol (students)	Money	Hypothetical	218	AUC
Camchong et al (2011)	Cocaine dependent individuals vs controls	Money	Not specified	51	HYP
Bickel et al (2011)	Working memory training in stimulant users	Money	Both	27	EXP-P

Johnson (2012)	Cocaine dependent individuals vs controls	Money	Both (QODT, EDT)	40	HYP
Stanger et al (2012)	Prospective study of treatment for marijuana abuse in adolescents	Money, marijuana	Hypothetical	165	HYP
Landes, Christensen & Bickel (2012)	Prospective study of treatment in opioid dependent individuals	Money	Hypothetical	159	HYP
Peters et al (2013)	Prospective study of treatment for marijuana abuse in adults	Money	Real (EDT)	93/61*	HYP
Heinz et al (2013)	Prospective study of treatment for marijuana abuse in military veterans	Money	Hypothetical	72	HYP

Note Studies are listed in chronological order. Abbreviations: SUD = substance use disorder; EXP = Exponential function; HYP = Hyperbolic function; AUC = Area under curve; IC = Incentive compatible (a design in which a proportion of choices have real outcomes); QODT = Quick Discounting Operant Task (Johnson, 2012); EDT = Experiential Discounting Task (a task with real monetary rewards); SCF = sooner choice frequency; EXP-P = Exponential power function. * $N = 93$ pre-treatment, 61 post-treatment.

3.3.5 DISCOUNTING, OBESITY AND EATING BEHAVIOUR

Researchers have examined relationships between obesity and discounting for both food and money outcomes, citing similarities between eating behaviour and addiction. Obese children have been shown to choose immediate over delayed edible rewards more often than normal weight children, though this effect was not found for non-food rewards (Bonato & Boland, 1983; W. G. Johnson, Parry, & Drabman, 1978). Notably the ability to delay gratification for food rewards at aged four predicts the likelihood of being overweight at aged 11 (Seeyave et al., 2009).

Correlational cross-sectional studies have examined links between measures of monetary discount and Body Mass Index (BMI) in adults, with mixed findings (Borghans & Golsteyn, 2006; L. H. Epstein et al., 2003; Ikeda, Kang, & Ohtake, 2010; Nederkoorn, Smulders, Havermans, Roefs, & Jansen, 2006; Reimers, Maylor, Stewart, & Chater, 2009). In a large sample from the Netherlands financial proxies for the discount rate, for example reported under-saving or excessive expenditure, were significantly correlated with BMI, however measured discount rates themselves were not (Borghans and Golsteyn, 2006). Similarly, Ikeda et al. (2010) found that BMI was positively correlated with a survey measure of procrastination, but showed no correlation with monetary discount rates in a sample of 2987 Japanese adults.

A handful of studies have compared discounting for overweight and normal weight individuals. Weller, Cook, Avsar, and Cox (2008) found that obese women exhibited significantly higher discount rates than healthy weight women, and a further study has shown that people who smoke cigarettes who are also obese to exhibit higher rates than non-obese smokers (S. Fields, Sabet, Peal, & Reynolds, 2011). C. Davis, Patte, Curtis, and Reid (2010) found that obese women with a binge-eating disorder, but not obese women without binge-eating disorder, had significantly higher monetary discount rates than normal weight women. More recently Kulendran and colleagues found significantly higher monetary discount rates in obese adolescents compared with normal-weight adolescents (Kulendran et al., 2013), and demonstrated that monetary discount rates in obese adolescents decreased over the course of a residential obesity intervention (Kulendran et al., 2013).

Table 3.4

Studies Comparing Discount Rates and Obesity or Eating Behaviour (N = 12)

Study	Health outcome	Modality	Real/ hypothetical	N	DF
Bonato & Bolland (1983)	Obese vs healthy weight children	Food/non-food	Real	40	SCF
Johnson, Parry & Drabman (1978)	Obese vs healthy weight children	Food/noon-food	Real	142	SCF
Epstein et al (2003)	BMI in smokers	Money	Hypothetical	78	HYP
Borghans & Golsteyn (2005)	BMI	Money	Hypothetical	2059	SCF
Nederkorn et al (2006)	Obese vs healthy weight women	Money	Hypothetical	51	HYP
Weller et al (2008)	Obese vs healthy weight	Money	Hypothetical	112	AUC
Ikeda, Kang & Ohtake (2010)	BMI	Money	Hypothetical	2987	EXP
Rollins et al (2010)	Healthy weight women, <i>ad libitum</i> eating	Money	Hypothetical	24	HYP
Davis et al (2010)	Obese women +/- binge eating disorder vs healthy weight women	Money	Hypothetical	209	IP
Fields et al (2011)	Obese vs healthy weight adolescent smokers	Money	Real (IC)	35	AUC
Appelhans (2011)	Food reward sensitivity in obese women	Money	Hypothetical	62	AUC
Kulendran et al (2013)	Obese vs healthy weight adolescents	Money	Real	53	HYP

Note Studies are listed in chronological order. Abbreviations: EXP = Exponential function; HYP = Hyperbolic function; AUC = Area under curve; SCF = sooner choice frequency; IP = indifference point; IC = Incentive compatible; BMI = Body Mass Index

Finally, it has been suggested that sensitivity to food rewards interacts with delay discounting, in support of which high discount rates predict palatable food intake amongst normal weight women who find palatable foods highly rewarding (Rollins, Dearing, & Epstein, 2010), an effect which has been replicated in obese and overweight women (Appelhans et al., 2011). Taken together these studies suggest an emerging relationship between discounting and weight status, although further work is clearly required to establish whether particular aspects of eating behaviour, such as caloric intake, or eating frequency show relationships with discounting.

3.3.6 DISCOUNT RATES AND PREVENTIVE HEALTH BEHAVIOUR

Some studies have demonstrated relationships between discounting and preventive health behaviours, however the findings are less consistent than for addictive behaviours. As described above, Chapman and Coups (1999) asked whether discount rates for monetary losses, as well as for a flu-like illness, could explain uptake of influenza vaccinations, with the finding that time preferences for money, but not illness, were related to vaccine uptake. In a later study (Chapman et al., 2001) monetary discounting showed an absent or very weak correlation with compliance with anti-hypertensive or cholesterol lowering medication. Similarly, a meta-analysis (Chapman, 2005) of 16 existing studies, including those described above, found no significant correlation between discounting and preventive health behaviour (Mean Pearson $r = 0.04$, 95% CI = - 0.01, 0.09).

The studies of Chapman and colleagues above suggest that in the population as whole preventive health behaviours show little or no relationship with discounting. However, two subsequent studies suggest that a subset of the highest discounters diverge from the rest of the population in their patterns of preventive health behaviour. Firstly, Axon, Bradford, and Egan (2009), studying 422 hypertensive adults, found that those in the highest quintile of monetary discount rates reported that they would be less likely to alter their diet and exercise plans to improve their future health. The highest discounters were not however significantly less likely to check their blood pressure or to follow their doctors' plans, as assessed by self-report. Secondly, Bradford (2010), analyzing discounting in 978 adults, found that for high discounting women the implied probability of attending mammography was reduced

by 15.30% over the preceding two years and high discounting men had significantly lower rates of prostate examination (probability reduction 8.31%). The influence of discounting on attendance for cervical cancer screening was marginally significant. Across gender, high discounters were significantly less likely to have attended the dentist (probability reduction of 24.8%) or to have had any cholesterol testing (probability reduction 12.38%) or any influenza vaccination (probability reduction 11.05%) over the preceding two years. Additionally, high discounters were significantly less likely to be non-smokers or to have undertaken weekly vigorous activity. These studies suggest that monetary discount rates might be a useful tool for identifying groups at risk of failing to engage in preventive health practices.

3.3.7 DISCOUNT RATES AND RISKY SEXUAL BEHAVIOUR OR DRUG-TAKING PRACTICES

Convergent evidence associates high monetary discount rates with behaviours that increase the risk of contracting sexually transmitted or blood-borne viral infections. Individuals infected with hepatitis C exhibit higher rates of discounting than controls (Huckans et al., 2011), although the direction of causality cannot be established from this study. Higher discount rates are associated with needle sharing amongst heroin users (Odum, Madden, Badger, & Bickel, 2000). Dierst-Davies et al. (2011) found that a sample of homeless, men who abused substances and had sex with men, had higher discount rates than a control sample of men, deemed to be at lower risk of HIV, who had sex with men however had stable housing and did not abuse substances. Finally, Chesson et al. (2006) found relationships between monetary discounting and a range of sexual behaviours and outcomes in a combined sample of university students and adolescents attending clinics (N = 1042). For example, adolescents with higher discount rates were more likely to have had sexual intercourse before age 16 years, to have contracted gonorrhea or chlamydia, or to have become pregnant.

Table 3.5

Studies Comparing Discount Rates and Self-reported health, Preventative Health Behaviour or Multiple Behaviours (N = 17)

Study	Health outcome	Modality	Real/ Hypothetical	N	DF
Fuchs (1982)	Self-rated health, smoking, dentist attendance, seatbelt wearing	Money	Hypothetical	500	EXP
Chapman & Coups (1999)	Influenza vaccination uptake	Monetary loss, flu illness	Hypothetical	412	EXP
Odum et al (2000)	Needle-sharing amongst heroin users	Money, heroin	Hypothetical	32	HYP
Chapman et al (2001) (1)	Influenza vaccination uptake	Monetary loss, flu illness	Hypothetical	679	IP
Chapman et al (2001) (2)	Hypertension medication adherence	Money, chest pain	Hypothetical	195	IP
Chapman et al (2001) (3)	Cholesterol medication adherence	Money, chest pain	Hypothetical	137	IP
Chesson et al (2006)	Risky sexual behaviour in teenagers	Money	Hypothetical	1042	EXP
Chao et al (2009)	Self-assessed morbidity and subjective survival probability	Money	Hypothetical	175	HYP
Axon et al (2009)	Hypertensive adults: BP checks, diet and exercise, compliance with doctors' plans	Money	Hypothetical	422	EXP
Reimers et al (2009)	Age at first sexual activity, smoking, BMI, substance misuse	Money	Hypothetical	42863	SCF
Bradford (2010)	Cancer and cholesterol screening, flu vaccinations, dental visits, non-smoking, exercise	Money	Hypothetical	987	Utility
Daugherty & Brase (2010)	Tobacco, alcohol, drug use, exercise, breakfast, seatbelt, estimated longevity, health concerns, sociosexual orientation	Money	Hypothetical	467	MCQ
Huckans et al (2010)	Hepatitis C status and substance misuse	Money	Hypothetical	83	HYP
Van der Pol (2011)	Self-assessed health	Money	Hypothetical	2300	EXP
Dierst-Davis et al (2011)	Homeless, substance-dependent men who have sex with men vs non-substance dependent men with stable housing	Money	Hypothetical	40	HYP
Moallem & Ray (2012)	Heavy drinkers, smokers and heavy drinking smokers	Money	Hypothetical	387	HYP
Melanko & Larkin (2012)	Healthy lifestyle questionnaire	Money	Both (EDT)	72	AUC

Note Studies are listed in chronological order. Abbreviations: EXP = Exponential function; HYP = Hyperbolic function; AUC = Area under curve; IC = Incentive compatible (a design in which a proportion of choices have real outcomes); EDT = Experiential Discounting Task (a task with real monetary rewards).

3.3.8 DISCOUNT RATES AND MULTIPLE HEALTH BEHAVIOURS

Two studies in the search sample compared discount rates with a broad range of health behaviour. Firstly Daugherty and Brase (2010), collected data from 467 undergraduates on an inventory of health behaviours, namely tobacco, alcohol and drug use, number of visits to a doctor or dentist in the past year, exercise frequency, eating breakfast, seat-belt use when in a vehicle, motorbike or bicycle helmet use, and the use of sunscreen. They found that, in a two-step hierarchical regression analysis, a combination of delay discounting for hypothetical money and survey measures of time perspective explained a significant proportion of the overall variance in health behaviour over and above the combination of the respondents' gender and their personality type (Costa & McCrae, 1990). At the level of predicting individual behaviours, the improvement in model fit achieved by adding the time preference measures at the second step was small (the largest improvement in R^2 was 0.05) but significant for all the behaviours above except helmet wearing. Notably Daugherty and Brase (2010) did not separate individuals by their level of cigarette, alcohol or drug use. As a result the observed relationships between discounting and other behaviours may have been confounded by the effects of these addictive behaviours to increase impulsivity in other domains.

Secondly, Melanko and Larkin (2012) analyzed data from 72 young adults who performed both a discounting task with real monetary rewards and a hypothetical monetary discounting task as well as completing a questionnaire assessing a variety of health behaviours. Unfortunately however their published article has since been retracted.

3.4 CONCLUSIONS

Taken together, the studies reviewed here support an hypothesis that high discount rates for money, and in specific instances food or drug rewards, are correlated with many unhealthy behaviours. Furthermore the effect sizes reported compare favorably to existing social cognitive models of health behaviour (Armitage & Conner, 2000; Conner & Norman, 2005b), establishing high discounting as a

reliable correlate of unhealthy choice. Several of the studies above also identify state-based influences on discounting, and as such it is timely to re-examine the extent to which discounting can be considered as either a personality trait or a state variable (de Wit, 2008; Odum, 2011; Bickel et al., 2012). Personality traits are defined as stable and enduring characteristics, reflecting a general tendency to respond in a given manner under given circumstances and can be seen to represent persistent patterns of internal states (see Costa & McCrae, 1990). State variables by contrast vary over a shorter time scale, and may be altered with changes in the motivational state of the respondents and/or the elicitation conditions (Kraemer, Gullion, Rush, Frank, & Kupfer, 1994).

3.4.1 TRAIT AND STATE VARIABILITY IN DISCOUNTING

In the preceding chapter I reviewed evidence to suggest that discounting for appetitive outcomes has features of a trait marker, exhibiting high test-retest reliability and being correlated across different forms of appetitive reward. If discounting indeed indexes an impulsive personality trait, then high discounting ought to precede unhealthy choice. The majority of studies reviewed above are cross-sectional and are therefore indeterminate as to whether high discounting antecedes unhealthy behaviour, or vice versa. However one study reviewed here demonstrates that monetary discounting can prospectively predict onset of unhealthy behaviour (Audrain-McGovern et al., 2009). Furthermore a growing number of studies have shown that monetary discounting predicts response to behaviour-change interventions (Yoon et al., 2007; Dallery and Raiff, 2007; Krishnan-Sarin et al., 2007; MacKillop and Kahler, 2009; Mueller et al., 2009; Washio et al., 2011; Sheffer et al., 2012; Stanger et al., 2012; Brown and Adams, 2013). Thus discounting has predictive utility and may allow health-behaviour change interventions to be tailored to benefit at-risk groups. In addition, interventions may be targeted at modifying the cognitive mechanisms associated with discounting, which under a trait hypothesis are assumed to contribute to unhealthy behaviour, for example working memory training has been shown to both reduce discount rates and modify addictive behaviour (Bickel, Yi, Landes, Hill, & Baxter, 2011).

Nevertheless, there is also considerable evidence from the studies reviewed that discounting is influenced by state-based factors. Discount rates amongst substance misusers are greater in a state of drug-craving than in a drug-sated state, and discount rates have been shown to decrease following behaviour change (e.g. Kulendran et al., 2013; Landes, Christensen, & Bickel, 2012). These findings are supported by a wealth of additional evidence showing that discounting can be manipulated through contextual framing (see Koffarnus et al. 2013 for a review). The observation that discounting exhibits state-dependence does not necessarily detract from its trait status, provided that the rank ordering of discount rates across individuals is preserved over these changes in state (for example, if addicted individuals who have the highest discount rates in a drug-sated state also have the highest rates when craving the drug). To our knowledge no studies have directly examined this. Until this question is addressed, a conservative conclusion would be to view discounting as a concurrent marker of the extent of unhealthy behaviour, rather than exclusively as an antecedent risk factor.

3.4.2 CONCLUDING REMARKS

This chapter has addressed correlations between discounting and health behaviour, without explicit discussion of the psychological processes putatively linking the two. Hyperbolic discounting has been envisaged by some authors as a mechanistic model, in the sense of being an algorithm by which humans or animals evaluate delayed outcomes (Ainslie, 1992; Ainslie, 2001). This algorithmic interpretation of discounting suggests that unhealthy choices in part *result from* hyperbolic discounting of their distant health consequences. For example, an oft-cited hypothesis is that the dynamic inconsistency theoretically engendered by hyperbolic discounting causes people renege on their best-laid plans. The next chapter, Chapter 4, scrutinises this proposal. Also, whilst hyperbolic discounting is widely accepted as a descriptive model for choice between immediate and delayed rewards (though see Kable & Glimcher, 2010; Luhmann, 2013; Read, 2001; Read et al., 2012; Read & Roelofsma, 2003), few previous authors have mapped out the functional form which best explains intertemporal choice for aversive outcomes. Many health-relevant

choices involve aversive outcomes, for example in the form of a painful preventive procedure, unaccustomed exercise, or even illness itself, and a model for aversive time preferences might assist in predicting choices concerned these kinds of outcome. To address this lacuna, Chapter 5 aims to test and extend existing theoretical accounts of negative time preference for aversive events.

DISCOUNTING AND SELF-CONTROL

“People sometimes behave as if they had two selves, one who wants clean lungs and a long life, an another who adores tobacco.”

Thomas C. Schelling – Choice and Consequence (1984)

Most people are familiar with having made well-intentioned plans for the future that they fail to carry out. For example early in the week one might plan to go to the gym on Sunday morning, only to find that when Sunday morning arrives, staying in bed appears more attractive. The attempt to enact one’s goals despite competing motives is referred to as ‘willpower’ (Ainslie, 2001) or ‘self-control’ (Tuk, 2015; O’Donoghue & Rabin, 2003; Prelec & Bodner, 2003; Thaler & Shefrin, 1981), and is often characterised as a conflict between larger-later and smaller-sooner rewards, which Read, Loewenstein, and Kalyanaraman (1999) have termed ‘virtues’ and ‘vices’ respectively. Such conflicts are a likely source of the marked intention-action discrepancy in health behaviour (Sheeran, 2002).

4.2 HYPERBOLIC DISCOUNTING AND DYNAMIC INCONSISTENCY

A classically rational agent would have no need for willpower, instead adhering effortlessly to previously made plans (Baron, 2000; Samuelson, 1937; Strotz et al., 1957). In axiomatic terms, a self-consistent agent who prefers a smaller reward, r available at time t , over a larger reward, R , available at time $t + d$, denoted by $(r, t) \succ (R, t + d)$, should still prefer the smaller reward when both rewards are delayed into the future by an added delay, ϕ , that is should also satisfy $(r, t + \phi) \succ$

$(R, t + d + \phi)$. This is referred to as the stationarity axiom and is satisfied if the agent discounts delayed reward exponentially.

However, there is substantial evidence that, in choices between immediate and delayed rewards, both humans and animals exhibit relatively lower discount rates over longer delays than over shorter delays (e.g. Ainslie, 1974, 1975, 2001; Benzion et al., 1989; Chung & Herrnstein, 1967; Green et al., 1994; Mazur, 1987), a pattern more consistent with a hyperbolic, rather than an exponential discount function. A decision maker with a hyperbolic discount function violates the stationarity axiom. In other words, hyperbolic discounting predicts that a preference for smaller-sooner over larger-later reward can be reversed simply by adding sufficient delay before both options are available. Hyperbolic discounting is therefore superficially consistent with a tendency for people to choose virtues over vices when viewed from a distance, but to reverse their choice when the situation is at hand (see Figure 1.1). Furthermore, since hyperbolic discount rates correlate with many forms of unhealthy behaviour (see Chapter 3), hyperbolic discounting has been proposed as an algorithmic explanation for impulsive reward-seeking at the expense of long-term plans (e.g. Ainslie, 2001; Angeletos et al., 2001; Bickel et al., 2012; Laibson, 1997; for a discussion of algorithmic models see Farrell & Lewandowsky, 2010; Lewandowsky & Farrell, 2010; Marr, 1982)

4.3 EVIDENCE FOR AND AGAINST HYPERBOLIC DISCOUNTING

Despite the elegance of the above account, direct empirical evidence for it is mixed. The majority of studies providing support for hyperbolic over exponential discounting elicit choices between *immediate* and delayed rewards, where lower discount rates over long delays (e.g. Benzion et al., 1989; Bleichrodt & Johannesson, 2001; Raineri & Rachlin, 1993), or an improved fit with a (quasi-)hyperbolic rather than an exponential model (Kirby & Maraković, 1995; Van Der Pol & Cairns, 2002), are taken to indicate violation of stationarity. However such studies do not in fact demonstrate preference reversals in process. As Read (2001) identifies, such studies confound the length of delay with the interval separating the two rewards, since in studies where the sooner reward is immediate these are identical. Thus although in

such studies implied discount rates are seen to decrease with increasing delay, this could equally be attributable to decreasing discount rates over longer *intervals*.

A more rigorous approach therefore tests for preference reversal directly by adding a ‘front-end’ delay to both sooner and later options and testing whether this indeed causes people to shift their preference from smaller-sooner to larger-later options as hyperbolic discounting predicts. Several studies have demonstrated that adding a front-end delay indeed causes the predicted preference shift (e.g. Green et al., 1994; Kirby & Herrnstein, 1995; Van Der Pol & Cairns, 2011). However, Read (2001) found that when a delay was divided into three intervals and discounting measured separately for each interval, discounting over the final interval was not significantly greater than that over the first interval. In other words adding a front-end delay did not produce the decrease in discounting predicted by hyperbolic discounting. Similarly, Kable & Glimcher (2010) found that discount functions based on a choice set in which all options were delayed by a fixed amount had the same hyperbolic curvature as those based on the same choice set in which the sooner option always occurred immediately. They concluded that their participants evaluated all delayed rewards relative to the earliest available reward, in contrast to conventional hyperbolic discounting, in which all outcomes are evaluated relative to the present time. This “as-soon-as-possible” function does not predict impulsive preference reversal.

Finally, an even more rigorous approach to testing the premises of hyperbolic discounting is to look for the predicted preference reversals longitudinally: if hyperbolic discounting is veridical, preference should shift from larger-later to smaller-sooner rewards with the passage of time. Such longitudinal preference reversals have been successfully demonstrated in choices with health-relevant outcomes. For example, Read and Van Leeuwen (1998) asked people whether they would prefer to receive in one week’s time either a healthy snack (such as a piece of fruit) or an unhealthy snack (such as a chocolate bar). The same individuals were followed up and one week later they were offered an immediate choice between a healthy and an unhealthy snack. Respondents chose healthy snacks more frequently when the choice was made in advance compared to when the snacks were immediately available. However, studies testing the predictions of hyperbolic discounting in real-time using monetary outcomes have had mixed findings (Ainslie

& Haendel, 1983; Read et al., 2012; Sayman & Öncüler, 2009). The earliest of these studies found support for the preference reversals predicted by hyperbolic discounting. In this study Ainslie and Haendel (1983) asked participants on a Monday to choose between smaller amount of (hypothetical) money on to be received on Friday and larger amount to be received the following Monday. Participants were offered the choice again on the Friday, this time between receiving the smaller amount (for real) immediately or the delayed amount on the coming Monday. Consistent with hyperbolic discounting, the most common pattern was a preference for the larger-later amount when choices were made in advance, but for the smaller sooner amount when this was immediate. However subsequent studies have not replicated this finding. Sayman and Öncüler (2009) found the opposite result using a design similar to that of Ainslie and Haendel (1983). Furthermore a study performed over several weeks using real monetary rewards showed that preference reversals towards choosing smaller-sooner amounts (that is, in the direction predicted by hyperbolic discounting) were not significantly more common than those in the opposite direction (Read et al., 2012). Importantly this was the case despite the participants displaying hyperbolic discounting in conventional “cross-sectional” choices.

In summary, the preference reversals of the form predicted by conventional hyperbolic discounting have hitherto not been *consistently* demonstrated, at least with monetary outcomes. Nevertheless, intention-action discrepancies in health behaviour do exist, and failures of self-control accord with everyday experience. This suggests that the preference reversals underlying health-related choices (such as those in Read and van Leeuwen, 1998) do not result directly from hyperbolic discounting in an algorithmic sense. Instead, real-world preference reversals may result from peoples’ inability to predict in advance the impact of motivational and environmental states on their future decision-making.

4.4 GOAL-INCONGRUENT ACTIONS OFTEN RESULT FROM STATE CHANGES

Everyday experience suggests that people often abandon long-term plans in favour of immediate reward in response to environmental cues or changes in internal

motivational state; for example, one might plan to abstain from eating dessert as part of a diet plan, but find it harder to resist when presented with a piece of cake (see for example Allan et al., 2010). This idea is supported by evidence; for example, relapses in drug-taking behaviour following abstinence commonly occur after exposure to a previous drug-taking environment (O'Brien, Childress, Ehrman, & Robbins, 1998). Indeed, hyperbolic discounting often does not appear necessary to explain these state-dependent influences. For example, in a study of analgesic preferences for childbirth (Christensen-Szalanski, 1984), women asked roughly one month in advance of labor preferred to avoid invasive spinal anesthesia in favor of less invasive but less effective pain relief methods, however during active labor women frequently reversed preference and opted for anesthesia. These findings are easily explained by an increase in the marginal utility for anaesthesia during the painful state, which was not accurately predicted in advance, without reference to hyperbolic discounting.

In keeping with the suggestions above, Loewenstein (1996) and Read (2001) have proposed that motivational drives and the cues that elicit them, rather than hyperbolic discounting, are responsible for impulsive preference reversal. However, existing accounts of these 'visceral influences' (Loewenstein, 1996) have tended to lack an explicit framework either for incorporating the motivational information provided by environmental cues, or for how this information becomes associated with cues through learning. The following discussion advances a suggested quantitative framework based on the principles of reinforcement learning for understanding the effects of environmental cues on intertemporal health choice. Key to this account is the notion that cues previously associated with rewarding actions can trigger goal-incongruent habits, leading to preference reversal even in the absence of hyperbolic discounting. A full exploration of learning is beyond the scope of this discussion, which is restricted to the effects of cues after learning has taken place.

4.5 A REINFORCEMENT LEARNING APPROACH TO CUE-TRIGGERED PREFERENCE REVERSAL

Models of reinforcement learning posit that action control proceeds by estimating the expected value of ensuing reward over series of temporally connected future states, encapsulated in a state-action value function (Sutton & Barto, 1998);

such models are therefore well placed to incorporate the effects of state-changes on choice behaviour. Attempting to optimise value in changing environments can be considered a trade-off between flexibility in rapidly incorporating new information and the efficient use of past experience (Daw, Niv, & Dayan, 2005). This trade-off is embodied by two methods of learning: a rather rigid, but computationally lean method, referred to as model-free, and a flexible, planning method capable of simulating future possible outcomes, often referred to as model-based (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; R. J. Dolan & Dayan, 2013; Gläscher, Daw, Dayan, & O'Doherty, 2010; McDannald, Lucantonio, Burke, Niv, & Schoenbaum, 2011; Simon & Daw, 2011; Wunderlich, Smittenaar, & Dolan, 2012). These systems reflect an established distinction in psychology between deliberative and automatic processes (Evans & Stanovich, 2013), but endow this with a normative and explicitly computational basis (Daw et al., 2005; Dayan, Niv, Seymour, & Daw, 2006).

A model-based decision-maker is generally assumed to search through the possible future states consequent on each action. Model-based decision-making corresponds to the definition of “goal-directed” behaviour in animal learning experiments as rapidly sensitive to changes in outcome value or the contingency between response and outcome (Balleine & Dickinson, 1998; Colwill & Rescorla, 1986; Dickinson & Balleine, 1994; Domjan, 2003). A model-free decision-maker, by contrast, through a gradual integration of outcome values encountered through experience, assigns a scalar estimate of long-run future value to taking an action in a particular state, without explicitly representing the corresponding future state of the world. The resulting ‘cached’ action values are relatively insensitive to immediate changes in the outcomes. Instrumental behaviour is initially goal-directed (model-based), but becomes increasingly model-free with learning, such that actions eventually become insensitive to changes in the value of the outcome, acquiring the characteristic of habits (Dickinson, Balleine, Watt, Gonzalez, & Boakes, 1995; Neal, 2006; Ouellette, 1998). Say for example a rat is trained to press a lever to obtain a food reward. Early in learning, lever pressing can be abolished by devaluing the food reward, for example by pairing it with sickness outside of the response box. By implication lever pressing is ‘goal-directed’: the rat presses because it ‘wants’ the food, a want that is abolished by pairing the food with sickness. By contrast, after

extensive training, lever pressing continues to some extent even after the food is devalued (Dickinson, Balleine, Watt, Gonzalez, & Boakes, 1995). This perserverative behaviour bears direct analogy to economic models of habit formation, which modify the instantaneous utility function to depend on past consumption (G. S. Becker & Murphy, 1988).

The differential engagement of these two systems has the potential to explain the environmental dependence of the preference reversals which underlie many forms of unhealthy behaviour. While steep temporal discounting over the model-based valuation of future health would be expected to encourage the initiation of unhealthy behaviour, with repetition, unhealthy behaviour is likely to become increasingly model-free, or habitual. At this stage, even if the decision-maker re-evaluates their goals in favor of making healthy choices, cached action values will continue to encourage unhealthy choice in response to relevant environmental cues, leading to apparently impulsive preference reversals (intention-action discrepancies).

4.5.1 ENVIRONMENTAL CUES CAN TRIGGER GOAL-INCONGRUENT HABITS

As outlined below, the goal-incongruent influence of cached (habitual) action values can produce preference reversal, without invoking hyperbolic discounting. Furthermore, preference reversal can result even if each system in isolation exhibits exponential discounting and discounts the future at the same rate, a crucial distinction from dual-process models of quasi-hyperbolic discounting (Angeletos et al., 2001; Koffarnus et al., 2013; Laibson, 1997; McClure et al., 2007; McClure et al., 2004). To demonstrate this formally, I consider a decision-making agent for whom overall action value is a weighted average of the value from each controller, where both systems discount the future exponentially with a per period rate, γ (γ is the conventional symbol for the discount rate in reinforcement learning approaches; its meaning is equivalent to that of δ in Equation 1).

Say, for example, the agent is a person who has just commenced a new diet plan and is choosing whether or not to consume a calorific biscuit when faced with a cue, the biscuit tin. A simplified (semi-Markov) state space for this decision is depicted in Figure 4.1. State B represents the presence of the biscuit tin. Consuming biscuits leads after a short delay, d_c , to state C, which carries a reward associated with

the tasty biscuits, R_c , and after a longer delay, d_h , to maintaining one's current weight, for simplicity here assigned a reward value of zero. By contrast, abstaining from biscuits leads, via the unrewarded state, A , to a small incremental health benefit in the form of weight loss, R_h , after delay d_h .

Say for example that before starting the diet plan, the goal-directed value of the delayed health benefit was zero ($R_h = 0$), and on starting the diet increases to R_h . The model-based system is capable of flexibly incorporating this new goal-value, however the model-free system has learned the cached value of consuming biscuits (termed a ' Q -value') based solely on the reward previously provided from consumption, and not the health consequences (Figure 2B). It is assumed that the model-based system has some difficulty in forecasting these cached values. Consider then that the agent, is asked to make their decision when situated in state P , at some time delay, p in advance of state B , where cached values have no influence, and that here they are indifferent between indulging and abstaining, that is to say that the model-based value, Q_{MB} , of consumption is equal to that of abstention:

$$Q_{MB}(Consume, P) = Q_{MB}(Abstain, P) \quad (4.5.1)$$

Given by:

$$R_h \cdot \gamma^{p+c+h} = R_c \cdot \gamma^{p+c} \quad (4.5.2)$$

Which simplifies to:

$$R_h \cdot \gamma^{c+h} = R_c \cdot \gamma^c \quad (4.5.3)$$

On reaching B , the presentation of the biscuit tin, cached action values are also "brought online", incrementing the benefit of indulging, such that:

$$Q_{COMBINED} = Q_{MB} \cdot \omega + Q_{MF} \cdot (1 - \omega) \quad (4.5.4)$$

Given by:

$$Q_{COMBINED}(Consume, B) = [R_c \cdot \gamma^c \cdot \omega] + [R_c \cdot \gamma^c \cdot (1 - \omega)] \quad (4.5.5)$$

$$Q_{COMBINED}(Abstain, B) = R_h \cdot \gamma^{c+h} \cdot \omega \quad (4.5.6)$$

And therefore, by Equation 4.5.4:


$$R_h \cdot \gamma^{c+h} \cdot \omega \leq [R_c \cdot \gamma^c \cdot \omega] + [R_c \cdot \gamma^c \cdot (1 - \omega)] \quad (4.5.7)$$

Predicting a preference for indulging for $\omega < 1$. Therefore the presentation of the biscuit tin brings about a preference for sooner consumption. In economic terms, environmental cues such as the biscuit tin can be viewed as updating the utility of the immediately available option, by providing (previously inaccessible) information from prior experience. This is simply a way of formalising the idea that previously formed habits can be difficult to overcome.


The interplay between model-based and model-free systems in the account above bears some similarity to existing dual-systems models of intertemporal choice, which posit a deliberative planning system in opposition with an impulsive system. However, while the former are often mapped onto quasi-hyperbolic models of discounting (McClure et al., 2007; McClure et al., 2004), which combine two exponential discount functions with differing rates, here the two systems may share the same discount rate. Dynamic inconsistency can then result from the different sources of information available to either controller (also see Dayan et al., 2006). In particular, the state-dependent valuations of the cached system can explain why real-world preference reversals occur in response to learned cues and, unlike existing quasi-hyperbolic accounts, why these preference reversals become more prominent with the formation of habits. In addition, unlike existing dual-process accounts,

reinforcement learning models can explicitly model the learning process generating incongruent preferences.

a) Model-based System



b) Model-free System



c) Combined Values

Indifference at P:

$$\omega * Q_{MB}(\text{Consume}, P) = \omega * Q_{MB}(\text{Abstain}, P)$$

Prefer Consumption at B:

$$\omega * Q_{MB}(\text{Consume}, B) + (1 - \omega) * Q_{MF}(\text{Consume}, B) > \omega * Q_{MB}(\text{Abstain}, B) + (1 - \omega) * Q_{MF}(\text{Abstain}, B)$$

Figure 4.1 Interactions between model-based and model-free decision-making. An agent deciding whether or not to consume biscuits given a cue, the biscuit tin. The agent's choice combines model-based and model-free value. **A** Model-based decision-tree from state, P at a time, d_p , in advance of encountering the biscuit tin, denoted by state B . Actions at B are evaluated by searching through the tree. *Consume* is followed after a short delay, d_c , with food reward, R_c , associated with consumption, state C , followed after a longer delay, d_h , by the maintenance of current body weight, the unrewarded state, U . *Abstain* is followed after delay, d_c , by the unrewarded state A , followed after delay, d_h , by weight loss with reward, R_h . The agent is naïve to model-free values. Model-based values, Q_{MB} , are given by the total reward following each action, discounted according to, $D(t)$, here exponential and identical across both controllers. **B** Model-free values: neither the outcomes, nor the transitions between them, are represented. Distant health consequences do not influence these values, Q_{MF} . Model-based and model-free values combine according to a weighted average, governed by ω . At P , where model-free values have no influence, the agent is indifferent between consuming and abstaining. In the presence of the biscuit tin at B however the additional influence of model-free (cached) values induces preference for consumption.

4.6 DISCUSSION

The foregoing evidence and argument challenge hyperbolic discounting as an explanation for intention-action discrepancies. The preference reversal predicted by hyperbolic discounting is not consistently evidenced in longitudinal studies of choices between delayed sums of money, and preference reversal can formally result from failure to forecast the effects of environmental cues and motivational states on behaviour, without evoking hyperbolic discounting. Nevertheless, independent the precise form of the discount function, steep discounting would putatively favor the initiation of unhealthy behaviour at a goal-directed level, following which the acquisition of habits might effectively bind impulsive individuals to their initial choices. In addition, repeated choice of immediately available rewards by individuals with high discount rates would be expected to lead to these individuals acquiring habits more rapidly (by more reinforced choices).

The above process appears particularly relevant to addiction (Everitt et al., 2008; Everitt & Robbins, 2005; Lucantonio, Caprioli, & Schoenbaum, 2014) and may in part account for the strong relationships between discount rates and addictive states. In support of this, animal studies of addiction demonstrate that rats bred to exhibit steeper delay discounting more rapidly acquire compulsive self-administration of cocaine than their low discounting counterparts (Belin et al., 2008). Finally, the shift towards model-free responding might be made all the more irrevocable in chronic addiction (Keramati, Dezfouli, & Piray, 2012), through damage to frontal cortical areas on which *model-based* valuations are thought to depend (Camchong et al., 2011; Gläscher et al., 2010; R. D. Rogers & Robbins, 2001; Smittenaar, FitzGerald, Romei, Wright, & Dolan, 2013), with an attendant decrease in the capacity to exert model-based control over goal-incongruent habits.

4.6.1 INNATE (PAVLOVIAN) INFLUENCES

Notably, while the above account focuses on goal-directed valuations and learned habits, it is also likely that innately signaled sources of cached value, and their interaction with motivational states, have an important behavioural effect. For example, if I am offered a tasty looking snack it I am likely to find it harder to resist

when hungry than when sated. This idea has been formalised by proposing that, in addition to goal-directed and habitual control, an innate controller mediates motivational state changes and initiates stereotyped evolutionarily appropriate actions in response to stimulus cues (Dayan et al., 2006; Domjan, 2003; Pavlov, 2003). This controller can be seen as reporting which stimuli and attendant actions have proved advantageous in our evolutionary past. The innate system is often referred to as ‘Pavlovian’, with reference to Pavlovian conditioning, in which a stimulus capable of eliciting an innate response is temporally paired with an arbitrary environmental cue, which subsequently becomes capable of eliciting the response independently (Domjan, 2003; Pavlov, 2003). Such innate affordances correspond to the ‘visceral’ influences (Loewenstein, 1996) and emotionally ‘hot’ states (Chapman, 2005; Metcalfe & Mischel, 1999) which have been previously associated with self-control failure, but place these influences within a quantitative framework.

The precise contributions of instrumental and Pavlovian effects to real-world choices are difficult to distinguish, for example there is also an ongoing debate as to role of classical (Pavlovian) conditioning in addictive disorders, and its interaction with the aforementioned instrumental processes (Everitt and Robbins, 2005). Nevertheless the mechanism of choice inconsistency proposed above for the case of model-based and model-free interactions would remain largely equivalent for the case of interactions between model-based and Pavlovian values (Dayan et al., 2006). An advantage of the reinforcement learning approach is its ability to generate simulations of these interactions over the course of learning and such models may yield parameters capable of explaining further variance in health behaviour.

4.6.2 FUTURE DIRECTIONS

Novel interventions might be directed at specific constructs within the above framework, and indeed several existing health behaviour interventions can be viewed in this manner. For example strategies aimed at making *healthy* choices habitual are already known to be effective (Lally, Chipperfield, & Wardle, 2008). Similarly, the Implementation Intentions approach (Gollwitzer & Sheeran, 2006; Sheeran, Milne, Webb, & Gollwitzer, 2009) can be seen as encouraging people to set up cached values

in advance, by specifying ‘if-then’ plans for how to enact a particular goal when a situation calling for self-control is encountered. These approaches are discussed in more detail in the General Discussion (Chapter 10). There is an urgent requirement for novel interventions capable of reducing goal-incongruent unhealthy choice, since the increasing burden of disease attributable to unhealthy behaviour is placing unsustainable demands on existing healthcare systems (R. Smith et al., 2012).

This chapter has focused on preference reversals for appetitive outcomes. By contrast, goal-incongruent choices regarding unpleasant, or ‘dreaded’ events have been less extensively studied from an economic perspective (for a discussion of the theoretical implications see Loewenstein, 1987), and no previous studies have set out to establish precisely how the disvalue of future pain depends on its timing. Much of the experimental work described in subsequent chapters of this dissertation addresses this issue.

GENERAL EXPERIMENTAL METHODS

“Would the art of measuring be the saving principle, or would the power of appearance?”

- Plato, Protagoras

This chapter details the generic procedures and analytical methods used in the experiments described in Chapters 6 to 9. Treatment is first given to methods for estimating models of utility from discrete choice experiments, since this forms the bulk of the experimental work contained in this thesis. I then give a brief description of methods that yield continuous estimates of subjective utility, namely matching procedures and incentive-compatible auctions. The final sections of the chapter concern the procedures used to induce pain in the laboratory, physiological measurements and the design of the social decision-making experiments described in Chapters 8 and 9.

5.1 ESTIMATING UTILITY FROM DISCRETE CHOICE EXPERIMENTS

By the principle of revealed preference, choosing a particular option over an alternative is taken to indicate that the chosen option carries higher subjective value (utility) for the decision-maker (see Sen, 1971). Of course, an experimenter might wish to disambiguate alternative functions for how utility depends on factors such as magnitude or delay. A straightforward approach to this is to design a set of choice options that titrate the independent variable of interest. For example, choices used to estimate discounting titrate upwards the size of a small immediate reward, while holding constant the size of a larger delayed reward and the delay (Kirby et al., 1999; Mazur, 1987). The point at which the subject switches their choice from preferring the large delayed reward to preferring the smaller immediate reward provides an estimate

of the *indifference point* at each delay (the point at which immediate and delayed options carry equal subjective value) and hence the discount rate. The Bayesian methods described below are based on analysis of a large number of such binary choices, presented in a randomised order. These methods incorporate stochasticity in choices, enabling a decision maker's preference for A over B to be expressed as a probability that they will choose A over B, and thereby allowing utility to be placed on a continuous scale (Glimcher et al., 2009; Glimcher et al., 2005).

The above approach requires a linking function to transform the utilities associated with alternative actions to the probability of choosing each, conventionally a logit, or softmax, function of the following form:

$$\pi(\text{choose } a) = \frac{\exp(\beta \cdot U(s|a))}{\sum_{a'} \exp(\beta \cdot U(s|a'))} \quad (5.1.1)$$

Thus the probability, π , of choosing action, a , is an exponential function of the utility of the state, s , obtained after taking action a , divided by the sum of the exponentiated value of all alternative actions, a' . (The utility of a state conditional on performing a given action is referred to as an *action value*). The parameter β , referred to as the inverse temperature, governs the steepness of the sigmoid dependence between values and choice probabilities, as illustrated in Figure 5.1, and can be thought of as governing the agent's sensitivity to the underlying utilities. When $\beta = 0$ all actions are chosen with uniform probability. As β increases there is an increasing probability of choosing the option with higher value on any given choice. β is usually fitted as a free parameter in modelling applications. The probability function π is referred to as the 'policy'. For the special case of two alternative actions the above simplifies to:

$$\pi(\text{choose } a) = \frac{1}{1 + \exp[\beta(U(s|a') - U(s|a))]} \quad (5.1.2)$$

This illustrates that the probability of choosing an action a over an alternative a' increases as a function of the value difference between a and a' (see Figure 5.1). More precisely the log odds of choosing a' over a increase as a linear function of the value difference (obtained by rearranging the above):

$$\log \left(\frac{\pi(\text{choose } a')}{\pi(\text{choose } a)} \right) = \beta [U(s|a') - U(s|a)] \quad (5.1.3)$$

This formulation has been shown to provide a good fit to behaviour in decision-making tasks (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; for a theoretical review see J. D. Cohen, McClure, & Yu, 2007).




Figure 5.1 Softmax (logistic) action selection function for binary choice. The probability of choosing an action, a , over an alternative, a' , increases as a sigmoid function of the increase in utility obtained by choosing a rather than a' , $U(s|a) - U(s|a')$. The inverse temperature parameter, β , governs the steepness of this relationship, shown here for three settings of β .

5.2 BAYESIAN MODEL COMPARISON

Bayesian model comparison (Daw, 2011; Raftery, 1995; Stephan, Penny, Daunizeau, Moran, & Friston, 2009) offers a principled means to disambiguate alternative models for how utility itself depends on the independent variables of interest. For example, a researcher might be interested in establishing a function by which the delay preceding a reward and reward magnitude interact to influence the utility of the reward. To do so the researcher could compare alternative models (functional forms) for the effect of these two factors. The Bayesian approach prioritises models that provide a good fit to the data, whilst penalising more complex models, which is important because a more complex model will always fit the data better, but may not make good predictions outside of the current experimental context (Lewandowsky & Farrell, 2010). The approach is based on Bayes' theorem, a fundamental principle of conditional probability:

$$p(H|E) = \frac{p(H) p(E|H)}{p(E)} \quad (5.2.1)$$

Thus the *posterior* probability of an hypothesis, H , being true given some evidence, E , is proportional to the *prior* probability of the hypothesis being true (before seeing the evidence), $p(H)$, multiplied by the probability of observing E under a state of the world in which H is true, $p(E|H)$, (termed the *likelihood* of the hypothesis). To turn the resulting estimate into a probability measure, this product is divided by a normalising constant, given by the probability of observing E under all possible hypotheses, $p(E)$.

For the purpose of model comparison, the hypothesis is the model under test, M , and the evidence is the experimental data, $\mathbf{y}$, giving:

$$p(M|\mathbf{y}) \propto p(M) p(\mathbf{y}|M) \quad (5.2.2)$$

That is, the posterior probability, given the data, of a model being true is proportional to the prior probability of the model being true, $p(M)$, multiplied by the consistency between the model and the data, formally the probability of the data given the model, $p(\mathbf{y}|M)$. The latter is referred to as the *evidence* for model M . If consideration is given only to models that are equally plausible on a priori grounds, $p(M)$ can be safely ignored, and models compared on the basis of the model evidence.

To compare models the best fitting parameters, θ , of each model must first be estimated. Expressed in terms of Bayes' theorem, doing so seeks to maximise the posterior probability of the parameters, given the data and the model (Daw, 2011; Lewandowsky & Farrell, 2010):

$$p(\theta|\mathbf{y}, M) = \frac{p(\mathbf{y}|\theta, M)p(\theta|M)}{p(\mathbf{y}|M)} \quad (5.2.3)$$

In the numerator, $p(\mathbf{y}|\theta, M)$ is the probability of seeing the data given the model and the current parameter settings, and indicates the extent to which the parameterised model provides a close fit to the data; expressed as a function of θ it is referred to as the likelihood function. $p(\theta|M)$ represents the prior belief in the current setting of the parameters themselves. Note that the denominator, $p(\mathbf{y}|M)$, is the model evidence as described above. This can be computed by marginalising over the model parameters, θ , and integrating:

$$p(\mathbf{y}|M) = \int p(\mathbf{y}|\theta, M)p(\theta|M)d\theta \quad (5.2.4)$$

This term takes account of model complexity, penalising models whose parameters deviate further from their prior values. In practice this integral is often difficult to calculate analytically, leading to a number of approximate solutions.

5.2.1 MAXIMUM LIKELIHOOD ESTIMATION

Given uniform priors on the model parameters, a straightforward approach to parameter estimation is to maximising the likelihood term, $p(\mathbf{y}|\theta, M)$, in other words finding the parameters, θ^{ML} , that make the observed pattern of data most probable (Daw, 2011; Lewandowsky & Farrell, 2011):

$$\theta^{ML} = \underset{\theta}{\operatorname{argmax}} p(\mathbf{y}|\theta, M) \quad (5.2.5)$$

The data here consist of a vector of n choices made by each subject, such that:

$$\mathbf{y} = \{c_1, c_2, c_3 \dots c_n\} \quad (5.2.6)$$

The likelihood function for each data point, c_i , is given by the policy, π , which is in turn determined by the utility and softmax functions. Assuming that each choice is an independent observation, the likelihood for all choices can be expressed as the product of the likelihood for each individual choice:

$$\theta^{ML} = \underset{\theta}{\operatorname{argmax}} \prod_i p(c_i|\theta, M) \quad (5.2.7)$$

Taking logs:

$$\log \theta^{ML} = \underset{\theta}{\operatorname{argmax}} \sum_i \log p(c_i|\theta, M) \quad (5.2.8)$$

In practice the maximum can be calculated by minimising the negative log likelihood. To do so here I use non-linear optimisation routines in Matlab (Mathworks, USA), either a simplex optimisation (fminsearch: Nelder-Meald algorithm, Lagarias, Reeds, Wright, & Wright, 1998), or constrained optimisation (fmincon: Active-Set algorithm). Having obtained the maximum likelihood parameter

estimates for each model, an approximation to the (log) Bayesian model evidence is given by the Bayesian Information Criterion (Schwarz, 1978):

$$BIC = -2\log \theta^{ML} + k\ln(n) \quad (5.2.9)$$

where n is the number of fitted choices for each subject, and k the number of model parameters. The second term penalises more complex models, according to their number of parameters.

If each subject's choices are treated as independent, *BIC* estimates for each subject can simply be summed to obtain a group-level estimate of the model evidence (Daw, 2011). This is sometimes referred to as fixed effects analysis, since the relationship of each subject's data to the group (population) level data is ignored. Since the parameters for each subject are estimated independently, frequentist statistical tests can be performed on the parameter estimates, permitting conclusions to be drawn regarding particular effects at the population level (Daw, 2011; Friston, Stephan, Lund, Morcom, & Kiebel, 2005; Holmes & Friston, 1998). In this thesis I chiefly use this 'summary statistic' approach.

A disadvantage of maximum likelihood fitting is that parameter estimates are sometimes poorly defined for particular ranges of the data, allowing parameters to take on extreme values. A solution is to impose sensible bounds on the parameters, which in practice means that extreme-valued parameters take on the lower or upper bound (Daw, 2011). This irregularity often means that parameter estimates are not normally distributed across subjects; for this reason I use non-parametric tests where applying this method. A further drawback of maximum-likelihood methods is that they tend to encounter local minima in the likelihood surface, and therefore often return estimates that do not reflect the global maximum likelihood. To circumvent this I use a random-multistarted overlay (RMSearch), which calls the optimiser 1000 times at starting points drawn from a uniform distribution between the bounds.

5.2.2 HIERARCHICAL BAYESIAN MODEL-FITTING

An alternative to the summary statistic approach is to perform Bayesian model comparison at the group level. The approach described below allows the distribution of parameters across subjects to inform an empirical prior for parameter estimation at the subject level (Daw, 2011; Huys et al., 2011; Huys et al., 2012). Model fitting at the subject level for each subject, j , then aims to find the maximum *a posteriori* estimate of θ_j , given a vector of each subject's choices, $\mathbf{y}_j$:

$$\theta_j^{MAP} = \operatorname{argmax}_{\theta} p(\mathbf{y}_j|\theta_j)p(\theta_j|\vartheta) \quad (5.2.10)$$

where ϑ are the parameters (hyperparameters) of an empirical prior distribution on θ . The use of an empirical prior constrains unreliable parameter estimates to take on sensible values.

Model-fitting in this case proceeds by estimating the maximum-likelihood *hyperparameters*, given the data from all N subjects:

$$\hat{\vartheta}^{ML} = \operatorname{argmax}_{\vartheta} p(\mathbf{y}_1 \dots \mathbf{y}_N|\vartheta) = \operatorname{argmax}_{\vartheta} \prod_i p(\mathbf{y}_i|\vartheta) \quad (5.2.11)$$

Note that this is analogous to the equation for maximum-likelihood at the parameter level, only now is applied to estimation of the group-level priors. (Dependence on the model, M , is omitted for simplicity). The likelihood of the hyperparameters is obtained by integrating over the subject-level parameters, weighted by the empirical prior:

$$p(\mathbf{y}_j|\vartheta) = \int d\theta_j p(\mathbf{y}_j|\theta_j)p(\theta_j|\vartheta) \quad (5.2.12)$$

Since both sides of the equation depend on ϑ , this integral is intractable. I used a procedure described previously (Huys et al., 2011; Huys et al., 2012), which uses Expectation-Maximisation (E-M) to find a solution. The equations below were derived by Quentin Huys and colleagues, though the notation has been adapted to follow on from the preceding discussion. I apply this previously published method in Chapter 8, described in detail below.

In the ‘E-step’ at the k th iteration the maximum *a posteriori* (MAP) parameter estimates for each subject are estimated, given an estimate of the empirical prior from the preceding iteration, $\vartheta^{(k-1)}$:

$$\theta_j^{(k)} = \operatorname{argmax}_{\theta} p(\mathbf{y}_j|\theta_j)p(\theta_j|\vartheta^{(k-1)}) \quad (5.2.13)$$

This maximisation is achieved by minimizing the negative log posterior probabilities by unconstrained nonlinear optimisation in Matlab (Mathworks, MA, USA). In the ‘M-step’, the estimated hyperparameters $\vartheta^{(k)}$ of the empirical prior distribution, mean μ , and factorised variance, σ^2 , are then updated based on the group-level distribution of the estimated MAP parameters. The mean update is straightforward, and is simply given by the mean of the subject-level parameters:

$$\mu^{(k)} = \frac{1}{N} \sum_j \theta_j^{(k)} \quad (5.2.14)$$

The variance update is somewhat more complex, and incorporates both error in the estimation of each subject’s parameters, and between-subject variability. The estimate of the variance on the individual-level MAP parameters, uses a Laplace approximation, which assumes that the posterior is Gaussian around the MAP estimate:

$$p(\theta_j|\mathbf{y}_j) \approx N(\theta_j^{(k)}, \Sigma_j^{(k)}) \quad (5.2.15)$$

where $\sum_j^{(k)}$ is the second moment around $\theta_j^{(k)}$, which approximates the variance. The factorised prior variance, σ^2 , is then updated as follows:

$$(\sigma^{(k)})^2 = \frac{1}{N} \sum_j \left[(\theta_j^{(k)})^2 + \sum_j^{(k)} \right] - (\mu^{(k)})^2 \quad (5.2.16)$$

The Bayesian model evidence, $p(\mathbf{y}_1 \dots \mathbf{y}_N | M)$, can be approximated as BIC_{int} (Huys et al., 2011; Huys et al., 2012):

$$-\frac{1}{2} BIC_{int} = \log p(\mathbf{y}_1 \dots \mathbf{y}_N | \hat{\vartheta}^{ML}) - \frac{1}{2} |M| \log (|\mathbf{y}_1 \dots \mathbf{y}_N|) \quad (5.2.17)$$

where $|\mathbf{y}_1 \dots \mathbf{y}_N|$ is the total number of choices made by all subjects, and $|M|$ denotes number of hyperparameters fitted (a mean and a variance for each parameter). By distinction from conventional BIC, $\log p(\mathbf{y}_1 \dots \mathbf{y}_N | \hat{\vartheta}^{ML})$ is a sum over the model evidence at the subject level by integrating out the subject-level parameters:

$$\log p(\mathbf{y}_1 \dots \mathbf{y}_N | \hat{\vartheta}^{ML}) = \sum_j \log \int d\theta \, p(\mathbf{y}_j | \theta) p(\theta | \hat{\vartheta}^{ML}) \approx \sum_j \log \frac{1}{K} \sum_{k=1}^K p(\mathbf{y}_j | \theta^k) \quad (5.2.18)$$

The right hand expression approximates the integral by summing over K samples, drawn from the empirical prior, $p(\theta | \hat{\vartheta}^{ML})$. Thus the individual-level parameters intervene between the data and the group-level inference, but are averaged out when comparing models. Since the above approach assumes that parameters are distributed normally in the population (or log normally if parameters are log transformed prior to inference), when using this approach I also analyse choice data directly with model-free approaches and non-parametric statistical tests, which do not rely on such distributional assumptions.

5.3 CONTINUOUS UTILITY ESTIMATES

5.3.1 UTILITY MATCHING PROCEDURES

An alternative to the forced choice methods described above is to simply ask participants to directly report their indifference points, for example by asking them to adjust the size of an immediate reward until they feel it just outweighs the subjective utility of a larger delayed reward, referred to as a ‘matching’ procedure (e.g. Read & Roelofsma, 2003; Van Der Pol & Cairns, 2008). Utility functions can then be estimated directly, for example by least squares curve fitting. I use a matching procedure in Experiment 1C to estimate time preferences for delayed illness, following a design previously published by Van der Pol and colleagues (Van Der Pol & Cairns, 2002) (see Chapter 6 for details). To compare models with different numbers of parameters, conventional BIC was estimated from the least squares model fits, by assuming that errors were normally distributed around the model estimate and using the normal probability density as a likelihood function.

5.3.2 WILLINGNESS-TO-PAY

A further means of obtaining a continuous utility estimate is to elicit subjects’ willingness to pay money to obtain a particular outcome. A standard means of doing so in an ‘incentive-compatible’ fashion is to hold a simulated auction using a method proposed by Becker, deGroot and Marschak, abbreviated to ‘BDM’ (G. M. Becker, DeGroot, & Marschak, 1964) (for an application see Vlaev, Seymour, Dolan, & Chater, 2009). Here the participant is asked how much of an endowment they would be willing to pay to receive a given outcome, referred to as their ‘bid’. Subsequently a ‘market price’ for the outcome is selected at random from a uniform distribution between zero and the size of the monetary endowment. If the participant’s bid exceeds the market price, they obtain the outcome and pay the market price (second-price auction); otherwise they forgo the outcome and pay nothing, keeping the entire endowment. This method ensures that the optimum behaviour is to bid an amount that is exactly equivalent in utility terms to the subjective utility of the outcome. Where

this method was used (Experiment 4), participants were fully briefed on the method. Following previous studies from the same laboratory, we also used an animated ‘roulette wheel’ effect to indicate that a market price was being selected randomly (see Chapter 9 for details).

5.4 EXPERIMENTAL PAINFUL CUTANEOUS SHOCK STIMULI

The experiments presented in this thesis examine temporal and social preferences for aversive events. To implement real aversive outcomes I used electrical stimulation applied to the skin on the dorsum of the hand or wrist. Such stimuli are extensively used in studies of aversive processing (e.g. Berns et al., 2006; Crockett, Kurth-Nelson, Siegel, Dayan, & Dolan, 2014; Feldman-Hall et al., 2012; Hein, Lamm, Brodbeck, & Singer, 2011; Seymour et al., 2005; Singer et al., 2004; Vlaev et al., 2014; Vlaev et al., 2009; Winston, Vlaev, Seymour, Chater, & Dolan, 2014) and are considered to be harmless. Participants were fully briefed on the nature of the stimuli, and were free to discontinue the study at any time they wished. They were reassured that all stimuli would be below an intensity level that they had indicated they were willing to tolerate. In fact all participants completed the experiments without undue distress.

The current experimental setup was developed at the Wellcome Trust Centre for Neuroimaging by Ben Seymour, Ivo Vlaev and Joel Winston (e.g. Seymour et al., 2005; Vlaev et al., 2014; Winston et al., 2014), and adapted for use in the current experiments by Giles Story. For each experiment a constant current stimulator (Digitimer, Letchworth, UK) was used to deliver shocks via two silver chloride electrodes affixed to the dorsum of the hand or wrist. Experiments 1A and 2 used a DS7 constant current stimulator, Experiments 3A and 4 used a DS5 constant current stimulator. In each case the stimulator receives an analogue voltage input and converts this into a current, applying the necessary voltage to generate a current level, in milliamperes (mA) specified by the experimenter.

After providing consent, participants underwent a threshold procedure, which allowed selection of current levels corresponding to an equivalent subjective level of discomfort for each participant, controlling for between participant variations in pain

perception. Notably that the pain scales used confound pain and unpleasantness (see Duncan, Bushnell, & Lavigne, 1989; Miron, Duncan, & Bushnell, 1989; Price, McGrath, Rafii, & Buckingham, 1983), however throughout I was primarily interested creating a generic noxious stimulus, rather than in isolating specific sensory components of pain. Owing to the differences between the DS5 and DS7 stimulators, different threshold procedures were used in each case, detailed below. The general principles of these procedures, and necessary Matlab code were developed by Joel Winston and Ben Seymour; the specific stimuli used were generated by Giles Story.

5.4.1 THRESHOLD PROCEDURE FOR DS7 STIMULATOR

The DS7 stimulator receives input from the parallel port of the Stimulus PC, which controls the timing of individual $200\mu\text{s}$ square-wave pulses. The current level is set manually by adjusting a dial on the stimulator itself. In view of this constraint, to generate stimuli of equal duration, the severity of which could be controlled from the computer during the task, I manipulated the frequency of pulses within a 5 second stimulus train, while holding the current constant. (The intensity of each discrete shock did not vary). The independent variable was the probability of receiving a shock at each sampled time interval during the 5 seconds, such that the number of shocks within the stimulus train followed a Poisson distribution. To achieve this, each 5s stimulus was partitioned into 50 sampling intervals of 0.1s duration. There was a fixed probability of receiving a shock within each of these sampled 0.1s windows. For example to achieve an expected shock rate of 10 shocks/5s, the probability of receiving a shock in a given 0.1s interval was set to 10/50 (i.e. 1/5). The mean number of shocks per five seconds varied between 2 and 14, corresponding to mean shock rates of 0.4 and 2.8 shocks/s respectively. These stimuli were used for Experiments 1A and 2, which were performed in the same session with the same participants.

The threshold procedure in this case was designed to set a current level such that each participant rated the five-second stimulus at the maximum shock rate (2.8 shocks/s) as moderately severe pain. To achieve this the expected shock rate was set to 2.8 shocks/s whilst the current amplitude was varied. Participants provided a visual analog pain rating of the entire five-second stimulus on a continuous 10-point visual

analogue scale (VAS), where 0=no sensation at all, 1-3=mild pain or discomfort, 4-6=moderate pain, 7-9=moderate to severe pain and 10=intolerable pain. I spent some time explaining the concept of this scale to participants, in order to set a realistic expectation of the level of shock. For example, the usage of such a scale in a controlled experiment clearly differs from its use in clinical practice, where 10/10 pain would represent excruciating pain. In this case a 10/10 stimulus was set to the maximum pain a participant would be willing to tolerate for the purposes of the experiment (i.e. a level which the participant felt unwilling to endure). Current level was increased in small increments until the participant rated the stimulus as 6/10. The staircase procedure was then repeated, allowing participants time to adapt to initial anxiety about the shocks. I encouraged each participant to give a qualitative description of the resulting sensation as the current was increased, to check his or her interpretation of the rating scale. If their interpretation differed from the initial briefing, thresholding was restarted.

The above procedure determined a single current level corresponding to moderately severe (6/10) pain for each participant. During the Experiments 1A and 2 this current level was held fixed, whilst the mean *rate* of shocks varied. At the end of Experiment 2 I also verified that increasing the mean shock rate indeed corresponded to monotonic increases in rated aversiveness, by asking participants to provide visual analog ratings of stimulus trains with different mean shock rates but constant current amplitude. To do so, shock rate was increased in increments of 2 shocks/5s, starting from the baseline mean rate of 2 shocks/5s up to the maximum rate of 14 shocks/5s at a constant current level equal to that used during the choice phase, followed by a symmetrical decreasing staircase. (See the results section 6.2.2 for data on the extent to which participants adapted to shocks over time).

5.4.2 THRESHOLD PROCEDURE FOR DS5 STIMULATOR

The DS5 stimulator receives a voltage input (-10 to +10V) via a USB data acquisition board (USB-6211, National Instruments, Austin, TX, USA), which in turn receives stimulus commands via USB from a stimulus PC running Matlab (Mathworks, MA, USA) with the National Instruments Data Acquisition (NI-DAQ) toolbox installed. The DS5 was set up such that +10mV inputs mapped to +10mA

outputs. The specific stimuli used for each experiment using the DS5 stimulator (Experiments 3A and 4) are described in the relevant methods sections of Chapters 8 and 9. The threshold procedure used was equivalent in each case. Firstly participants were asked verbally to rate the pain associated with shocks of varying current levels using a 10-point visual analogue scale (VAS), identical to that described above. The experimenter specified current levels manually, by typing into the Matlab command line. Suggested current values were generated automatically from a custom script designed to produce an increasing stochastic staircase of current levels. Stochastic increments were used to limit participants' expectation that subsequent shocks would be of increasing intensity. Two iterations of this staircase procedure were used to identify the maximum intensity (rated 10/10) that the participant felt they could tolerate for the purposes of the experiment. Participants were reassured that during the experiment all stimuli would be below their maximum tolerance. Secondly participants rated, using an on-screen VAS, a randomised sequence of sub-maximal shock intensities generated automatically. An evenly-spaced sequence of current levels between zero and the maximum tolerable level was generated by the computer programme, and the presentation order of these current levels was subsequently randomised to prevent expectation effects. The least-squares fit of a sigmoid psychophysical function (a three-parameter Weibull distribution) describing the subjective ratings at each current level was determined automatically for each participant using a customised curve-fitting routine in Matlab (Mathworks, USA), and inverted to determine current levels (in mA) for use during the experiment, corresponding to specific VAS ratings (the precise subjective rating levels used varied by experiment: see the relevant methods sections of each). In addition, the programme plotted the resulting curve fit, the quality of which was checked visually by the experimenter. If a participant's rating scores were particularly variable, the curve fit was disrupted by outlying values, or significant adaptation to the shocks had taken place since the initial staircase procedure, the entire thresholding procedure was repeated.

5.5 ETHICS STATEMENT

All procedures described here received approval from a local research ethics committee. All participants gave full informed consent before taking part in the experiments. Experiments 1 and 2 were approved by the National Health Service National Research Ethics Service, Central London Research Ethics Committee 3 (Ethics number 08/H0716/6, Amendment AM1). Experiments 3 and 4 were approved by the UCL Research Ethics Committee (4418/001 for Experiment 3A and 4418/002 for Experiments 3B and 3C, 3953/001 for Experiment 4).

5.6 PARTICIPANT RECRUITMENT

For Experiments 1A and Experiment 2 healthy adult participants were recruited from the University College London Division of Psychology and Language Sciences Subject Pool (currently available at <http://uclpsychology.sona-systems.com>). For Experiments 1B, 3A and 4 healthy adult participants were recruited from the University College London Institute of Cognitive Neuroscience Subject Database (currently administered at <http://groupspaces.com/ICNSubjectDatabase/>). Both these subject pools offer a facility for researchers to post advertisements regarding their study for participants to browse. Participants must complete a pre-screen survey before signing up to confirm that they do not suffer from any major neurological or psychiatric conditions and are not currently taking any prescribed psychoactive medicines. Participants are given guidelines when they sign up regarding acceptable conduct for taking part in studies, for example advising them that it is prohibited to use mobile phones or to eat during studies. No data are currently available regarding the demographics of these populations as a whole, however healthy adult volunteers aged 18-55 were recruited for all the above experiments. Specific demographic details of participants are given in the relevant methods section of each experiment.

For Experiments 3B and 3C participants were recruited using Amazon Mechanical Turk (AMT). Healthy adult participants who had worker accounts on Amazon Mechanical Turk (AMT) were recruited via an advertisement on the AMT account of one of the authors (MC). AMT is an online marketplace for work, now

widely used as a method of data collection for psychology experiments (Mason & Suri, 2012). Responses on AMT have been shown to be reliable (Rand, 2012) and replicate well-established findings in the cognitive psychology literature (Crump, McDonnell, & Gureckis, 2013). Demographic data relating to these participants is described in the relevant methods section.

TIME PREFERENCE FOR PAIN AND ILLNESS

“The only thing we have to fear is fear itself.”

Franklin D Roosevelt, 1933

Imagine that you are faced with having to schedule a time for an appointment at the dentist to have a tooth removed. You know that the experience will be painful, but will only last for a short time. Would you prefer an appointment this week or in three weeks' time? In experiments people frequently seek to expedite unavoidable pain in order to “get it of the way” (Badia, McBane, & Suter, 1966; Berns et al., 2006; Cook & Barnes, 1964; M. D'Amato & Gumenik, 1960; Franzini, 1970; Hare, 1966a, 1966c; Hare, Krebs, Creighton, & Petrusic, 1966). In other words, people often appear impatient to receive pain as well as reward. As described in Chapter 2, since this behaviour implies a negative discount rate it is termed ‘negative temporal discounting’, or ‘negative time preference’ (Elster & Loewenstein, 1992; Loewenstein, 1987; Loewenstein & Prelec, 1991). Preventive health behaviours, such as attendance at cancer screening, or physical exercise, are readily framed as entailing short-term discomfort for the sake of long-term relief from illness, suggesting that time preferences for pain or illness might prove capable of predicting such behaviours.

An influential explanation for the tendency to expedite inevitable pain that has been that anticipating pain is itself aversive, aptly termed ‘dread’ (Loewenstein, 1987; Berns et al. 2006). The longer a person has to wait for pain, the longer they have to anticipate it, suggesting that the total disutility from dread ought to increase with the delay (see for example Franzini, 1970), and a person can putatively mitigate this prospective dread by choosing sooner pain. The concept of dread receives empirical support from the finding that brain responses to the anticipation of pain are anatomically contiguous with responses to the direct experience of pain (Koyama, Tanaka, & Mikami, 1998; Ploghaus et al., 1999). Further, using a cue-delayed shock

paradigm and measuring human brain blood-oxygenation-level-dependent (BOLD) activity with functional magnetic resonance imaging (fMRI), Berns et al. (2006) demonstrated that individuals with stronger preference for sooner pain had significantly higher anticipatory brain activity during the cue-shock interval, suggesting that stronger preference for sooner shocks in these individuals was indeed driven by their greater prospective dread. However, despite empirical support for the existence of dread, no previous studies have compared alternative models for how the disvalue of pain might grow with delay. Hitherto unresolved issues are the form of the anticipation, and the way that this anticipation affects present value as dread.

Under existing theoretical accounts, the decision-maker is assumed to sum the moment-by-moment anticipation they can expect to experience whilst waiting for pain, to form a prospective estimate of dread (Berns et al., 2006; Loewenstein, 1987). A simple account proposes that, when doing so, people treat each unit of waiting time as equally aversive. Here the total dread of pain is directly proportional to the delay. Under an alternative account, the moment-by-moment anticipation of pain increases as pain draws nearer. Here the total dread of pain again increases with increasing delay, though at a decreasing rate (for the reason that distant pain generates relatively little immediate anticipation). A further possibility is that the moment-by-moment increase in anticipation is itself prospectively discounted in time (Loewenstein, 1987). In other words, although people putatively expect to be experiencing intense anticipation immediately before the pain, if the pain is distant, this future anticipation might have little influence in the present. In particular, this predicts that prospective pain has a future point at which it is maximally aversive, being preferred both sooner or later. Thus we might prefer to have a painful dental procedure now as opposed to next week, but also next year as opposed to next week. These differing accounts make testable predictions for how the disvalue of prospective pain depends on its timing.

Another unknown property of dread is its stability in the face of biases such as framing (Tversky & Kahneman, 1981). In particular, if dread can be re-framed as relief from a putatively higher amount of pain, it might be possible to reduce dread or even eliminate it. This would be consistent with a ‘sign effect’ shown for monetary outcomes (Frederick et al., 2002; Loewenstein & Prelec, 1992; Tanaka, Yamada, Yoneda, & Ohtake, 2014; Thaler, 1981), whereby losses are discounted less than gains, and also for health outcomes, whereby decrements in health (illnesses) are

discounted less than health improvements (MacKeigan et al., 1993; Odum et al., 2002). By contrast, if people interpret a relief frame in a rational manner, they ought to dread the unrelieved pain, and so exhibit no framing effect. To test the alternative models described above and the influence of framing effects, in two experiments I examined intertemporal choices between painful outcomes at different delays. In Experiment 1A the delays ranged from seconds to around 15 minutes, and outcomes consisted of moderately painful cutaneous electric shock stimuli delivered to the dorsum of the hand. In Experiment 1B outcomes consisted of hypothetical dental appointments with varying degrees of dental pain at delays ranging from today to approximately 8 months.

As reviewed in Chapter 2, negative time preference is frequently observed for hypothetical illness (e.g. Van Der Pol & Cairns, 2000). Whilst it has been suggested that this effect might be explained in terms of the dread of future illness, no previous studies have formally tested this. I do so here in Experiment 1C. This experiment directly follows the design of a previously published study of illness discounting, Van der Pol and Cairns (2002), which asks participants to choose whether to delay an illness by means of an imaginary treatment.

6.1 EXPERIMENT 1A: TIME PREFERENCE FOR PAINFUL SHOCKS

6.1.1 METHODS

Thirty-five healthy adult participants (18 females, mean age 28.5 years, s.d. 5.1 years) took part in the experiment. Participants were recruited by advertisement on the website of the University College London Psychology Subject Pool (see General Experimental Methods 5.6). All experiments were carried out at the Wellcome Trust Centre for Neuroimaging, University College London and each session lasted around 2 hours. All participants gave full informed consent prior to the experiment. Participants were briefed that they would be making choices between different numbers of moderately painful electric shocks that would be delivered at different points in time. Verbal briefing of participants suggested that several had

previously taken part in studies involving similar painful shock stimuli, though unfortunately I did not collect data on this to enter into the analyses.

Shocks were delivered using a Digitimer (Letchworth Garden City, England) DS7 constant current stimulator through silver chloride surface electrodes placed approximately 3 cm apart on the dorsum of the left hand. Each individual shock consisted of a single 200 μ s square-wave pulse. Throughout the experiment the participant sat in front of a computer monitor; where trials were presented on-screen, and decisions were indicated using two keys on the keyboard. The software package COGENT 2000 (University College London) was used for stimulus presentation and response acquisition.

The painful shocks occurred within a five second stimulus train, and the intensity of each discrete shock did not vary. The independent variable governing the severity of pain was the number shocks within the five second stimulus. Since the duration of the stimulus was constant, increasing number of shocks was equivalent to an increasing shock *rate*. The number of shocks within the stimulus train followed a Poisson distribution with uniform probability of receiving a shock at each sampled time interval. To achieve this, each 5s stimulus was partitioned into 50 sampling intervals of 0.1s duration. There was a fixed probability of receiving a shock within each of these sampled 0.1s windows. For example to achieve an expected shock rate of 10 shocks/5s, the probability of receiving a shock in a given 0.1s interval was set to 10/50 (i.e. 1/5). (see also General Experimental Methods 5.4). Participants were briefed on the probabilistic nature of the outcomes and were informed of the number of shocks they could on average expect to receive for a particular outcome, that is, the mean of the distribution; it was assumed that participants made their choices based on this number. Probabilistic outcomes were used to embed the experiment into a more naturalistic context.




Figure 6.1. Trial structure of the task in Experiment 1A. 1. A sequence of two Choice Trials, demonstrating the display of outcome options and outcome phases. The dotted arrow denotes how choices on previous trials determine expected shock rates on the future trials referred to by those choices. 2. A No Choice Trial.

After providing consent, participants underwent a standardised thresholding procedure (Seymour et al., 2005; Vlaev et al., 2014; Winston et al., 2014) (see General Experimental Methods 5.4.1) to control for between participant variations in pain perception. Prior to making choices participants also received six samples of five second stimulus trains at two different shock rates, corresponding to the minimum and maximum rates used during the experiment, of 2 shocks/5s (0.4 shocks/s) and 14 shocks/5s (2.8shocks/s). As detailed in General Experimental Methods 5.4.1, for this experiment, two threshold procedures were performed, firstly to determine a current level to use throughout the experiment, corresponding to 6/10 pain at the maximum shock rate. Secondly, at the end of the experiment, participants' ratings were obtained

of the unpleasantness of varying shock rates. This second procedure had two purposes: firstly to confirm that increasing shock rates were associated with increasing levels of rated aversiveness, and secondly to check that participants had not substantially adapted to the shock stimuli over the course of the experiment (see Results 6.1.2 for data).

Prior to the intertemporal choice phase, participants were briefed with on-screen instructions that embedded the task in a naturalistic health-related scenario:

Imagine that you are suffering from pains as a result of damage to the nerves in your arm, which feel as if you are having brief electric shocks delivered to your hand. The pains come in clusters of ‘shocks’, which are roughly evenly spaced in time, with each episode lasting five seconds. Within each episode the shocks come randomly. You experience a usual baseline level of pain, of on average two shocks per episode. All the shocks are of the same intensity.

On each trial the default outcome was a five second shock train with mean 2 shocks/5s (0.4 shocks/s), which was referred to as a “Baseline Episode”. Participants’ choices determined outcomes with higher shock rates, referred to as “Severe Episodes”. Participants were informed that their choices would not change the total number of pain episodes, only their timing and severity. Intertemporal choice data were collected in two blocks, the order of which was counterbalanced across participants: a block in which outcomes were framed as an increase in shock rate, referred to as the *pain* frame and an otherwise identical experimental block in which outcomes were framed as a decrease in shock rate, referred to as the *relief* frame. The timing of outcomes was expressed as the number of trials (episodes) in the future. The description in the pain frame was as follows:

You are about to have an attack of more severe episodes, lasting about half an hour in total. These severe episodes consist of more shocks than your usual baseline level, sometimes up to around 12 extra shocks per episode. You can choose when you would like to experience each severe episode: for each episode you will be offered two different points in time, and you will be given the choice as to which you would prefer. The severity of the episode will often be different at the two different time

points but, to help you make your choice you will be told the average number of extra shocks, relative to your usual baseline, that you could expect to receive at each time.

The description in the relief frame was as follows:

You are about to have an attack of more severe episodes, lasting about half an hour in total. These severe episodes consist of 12 more shocks per episode than your usual baseline level. You can choose when you would like to experience each severe episode: for each episode you will be offered two different points in time, and you will be given the choice as to which you would prefer. The severity of the episode will be 12 shocks above the baseline, but you will be offered a painkiller to reduce the severity. The effectiveness of the painkiller might depend on when it is used. To help you to choose you will be told how many shocks you can expect the painkiller to relieve at each time-point.

In each of 2 sessions participants made 95 choices between two options involving between 3 to 12 moderately painful electric shocks, delivered at between 4 to 51 trials in the future. Choice blocks proceeded according to a trial-based design in which the unit of time was a single trial and participants' choices determined outcomes on future trials. The sequence of events across a series of trials is shown in Figure 6.1. All choices were genuine, with shock delivered faithfully according to participants' choices. Making no *a priori* assumptions about the direction of participants' time preference, there were an equal number of choices in which the delayed outcome had a higher expected shock rate as choices in which the sooner outcome had a higher expected shock rate. The set of choices used is detailed in the Appendix (Table A1). The presentation of outcome options was counterbalanced and randomised such that sooner outcome options appeared on the left-hand side of the screen on half of trials, and on the right-hand side in the other half of trials.

Finally, there were two types of trial: Choice Trials and No Choice Trials. In each run there were 95 Choice Trials and an approximately equal number of No Choice Trials. The latter were necessary to absorb the outcomes of all the choice trials, such that all choices faithfully led to their outcomes. On Choice Trials participants were first presented with a choice between two options for a Severe

Episode, where each detailed its timing and expected shock rate. After a choice had been made, there followed the painful episode (five-second outcome stimulus) for that trial, whose shock rate was determined by previous choices. On No Choice Trials participants were presented with a screen saying “No Choice This Time”, which was displayed for a constant delay of 1s, and followed directly by the painful episode for that trial.

Choices and outcomes were interleaved: for example if a participant chose on trial one to receive “a Severe Episode with nine extra shocks, five trials in the future”, then on trial six they would experience an outcome with a mean shock rate of nine shocks above the baseline. The outcome for trial one would then be a Baseline Episode, as was the case for all trials not referenced by a previous choice. Prior to each experimental run we generated a novel trial order using a random permutation that was bounded such that no two Choice Trials referred to an outcome on the same trial, and participants were informed of this fact. Although Choice Trials and No Choice Trials were randomly interspersed, the frequency of the latter necessarily increased towards the end of the experimental run, in order to ensure that choice delays did not extend beyond the end of the experiment.

6.1.2 RESULTS

6.1.2.1 Adaptation to Shock Stimuli and Participant Inclusion Criteria

Ratings of stimuli with varying shock rates obtained at the end of the experiment revealed that rated unpleasantness was an increasing concave or sigmoid function of shock rate (fitted using a three-parameter Weibull function). Of the 35 participants, 2 participants were excluded from the analysis after they reported during the experiment that they did not find the shock stimuli aversive. With these 2 participants excluded, inversion of the psychophysical curve fit to the ratings data revealed that the mean participant rated the maximum shock rate used during the experiment at 5.81/10 on the VAS (s.d. 1.22 points, 95% CI [5.38 6.25]), not significantly different from the 6/10 rating obtained at the start of the experiment

(one-tailed t -test, $p > 0.25$), indicating that no significant adaptation had taken place at the group level.

A further 8 participants deterministically chose sooner pain, irrespective of the shock rate: these “maximum dreaders” were excluded from the modelling analysis, since the shape of their preferences could not be reliably assessed using the experimental choices offered. Notably this is relatively high exclusion rate. However it is not unreasonable that some participants fell outside the range of preferences assessed by the choice options, given that preferences for delayed pain had received little empirical exploration previous to this study.

6.1.2.2 Group Level Time Preference

At the group level, participants showed a strong preference for sooner pain, at the expense of an increased number of shocks. Overall time preference in the experiment is given by the mean probability across all choices of choosing later shocks (S2) over sooner shocks (S1), referred to as $p(\text{Choose S2})$. Since there are equal numbers of trials in which $S1 > S2$ as in which $S2 > S1$, overall negative time preference is indicated by $p(\text{Choose S2}) < 0.5$. Group mean $p(\text{Choose S2})$ averaged across both frames and all delay lengths was significantly less than 0.5, [mean $p(\text{choose S2}) = 0.29$, $SE = 0.04$, $N = 33$, One sample $t(32) = -5.23$, $p < 0.001$], this was confirmed with non-parametric testing [median $p(\text{choose S2}) = 0.34$, One sample Wilcoxon Signed Rank test $p < 0.001$], indicating overall negative time preference. As a result, participants chose the larger pain on 32.6% ($SE = 3.21$) of choices overall.




Figure 6.2. Observed time preference: Experiment 1A. Mean $p(\text{Choose S2})$ across participants included in the modelling analysis ($N = 25$) as a function of difference in delay between delivery of S2 and S1 ($D2 - D1$), expressed in units of trials. Delay difference ($D2 - D1$) is binned into tertiles, corresponding to short (1-10 trials), medium (11-20 trials) and long (>20 trials) delay differences. **A:** choice probabilities for all choices. Since there are equal numbers of trials in which $S1 > S2$ as in which $S2 > S1$, and option presentation is counterbalanced, $p(\text{Choose S2})$ at a delay difference of zero is theoretically bounded at 0.5. Blue circles represent choice probabilities for the relief frame, red circles choice probabilities for the pain frame. Error bars represent one standard error from the between subject mean. **B:** choice probabilities for choices in which the more delayed option was a smaller number of shocks. At delay difference of zero, S1 and S2 would occur at the same time-point, under which circumstance it might be assumed that participants would show preference for the smaller number of shocks, denoted by the blue and red square.

Figure 6.2 shows mean $p(\text{Choose S2})$ on the two frames across the 25 subjects included in the modelling analysis, as a function of delay length, expressed as the difference in delay between the two choice options. A 2-way repeated measures analysis of variance (ANOVA) revealed a significant main effect of frame [$F(1,24) = 9.5$ $p = 0.005$], whereby participants chose sooner shocks less frequently in the *relief* frame. There was also a significant main effect of delay [$F(3,72) = 8.2$; $p = 0.002$], as well as a delay by frame interaction [$F(3,72) = 4.2$; $p = 0.023$]. Non-parametric pairwise comparisons between zero, short, medium and long delay differences (across both frames combined) revealed a significant decrease in $p(\text{Choose S2})$ between zero and short delay differences [$p(\text{Choose S2}, \text{Short}) < 0.5$, Wilcoxon Signed Rank test, $p = 0.004$] and between short and medium delay differences (Wilcoxon Signed Rank test, $p = 0.014$), but no significant change between medium and long delay differences

(Wilcoxon Signed Rank test, $p = 0.398$). These pairwise tests are *post-hoc* and therefore exploratory. Based on these results, an hypothesis is that negative time preference decreases at longer delay differences, rather than being constant.

In particular the *slope* of the dependence of $p(\text{Choose S2})$ on delay provides a proxy for the rate of time preference. I therefore tested the hypothesis of decreasing negative time preference as noted above by performing a further 2-way repeated measures ANOVA, entering the slope of $p(\text{Choose S2})$ between each category of delay difference as the dependent variable. This second ANOVA demonstrated a significant main effect of frame on the rate of time preference [$F(1,24) = 15.5$; $p = 0.001$], as well as a significant main effect of delay [$F(2,48) = 4.4$; $p = 0.033$], thus rejecting the null hypothesis of constant negative time preference. There was no significant delay by frame interaction in this analysis [$F(2,48) = 1.6$; $p = 0.205$]. Non-parametric pairwise comparisons revealed that the effect of delay was driven by a significantly more negative slope between short and medium delays than between medium and long delays (Wilcoxon Signed Rank test, $p = 0.013$), consistent with diminishing negative time preference. This suggests the rate of accumulation of dread diminishes with increasing delay, suggesting that on average instantaneous dread increases in time, rather than being constant.

6.1.2.3 Classification of Participants by Time Preference

The group level data displayed in Figure 6.2 conceal substantial heterogeneity in the response patterns of individual participants. I therefore categorised the 25 participants whose data contributes to Figure 6.2 according to their individual pattern of time preference. Four mutually exclusive categories were identified: *zero* time preference, *positive* time preference, *negative* time preference and *reversing* time preference (Figure 6.3).




Figure 6.3 Choice data of participants in Experiment 1A categorised by time preference. $p(\text{Choose S2})$ as a function of delay difference, expressed in units of trials, for all 25 participants included in the modelling analysis. Choice probabilities shown are the mean of those on the two frames. Delay difference scaling is identical that in Figure 2. Time preference is approximated by the slope of the choice probability lines. **A:** participants with no significant time preference at any delay. **B:** participants who show positive time preference, but no significant negative time preference at any delay. **C:** participants who show negative time preference, but no significant positive time preference at any delay. **D:** participants with initial negative time preference followed by significant positive time preference at longer delays. Data are plotted as solid lines to assist visualisation of the choice patterns. Each gray line represents data from a single participant. The bold purple lines represent the between-subject means in each category. At delay difference of zero, S1 and S2 would occur at the same time-point; since there are equal numbers of trials in which $S1 > S2$ as in which $S2 > S1$, the plots are theoretically bounded to cross the probability axis at $p(\text{Choose S2}) = 0.5$.

Participants were classified as having *zero* time preference if $p(\text{choose S2})$ showed no significant deflections from 0.5 at any delay difference (Binomial test, $p > 0.05$) (7/25, Figure 6.3A). The zero time preference group chose the option with the smaller shock rate on 88% of choices [mean $p(\text{choose smaller}) = 0.88$, $SE = 2.8$], demonstrating that this group did not simply respond randomly, but tended to choose the less painful stimulus, irrespective of delay. Participants were classified as having *positive* time preference if they displayed significant increases in $p(\text{choose S2})$, but no significant decreases (Fisher Exact test, $\alpha = 0.05$) (4/25, Figure 6.3B), as having *negative* time preference if they displayed significant decreases in $p(\text{choose S2})$, but no significant increases (12/25, Figure 6.3C), and as having *reversing* time preference if they displayed significant increases in $p(\text{choose S2})$, as well as significant decreases (2/25, Figure 6.3D). The two participants with *reversing* time preference both displayed initial negative, followed by positive time preference, a pattern that would be consistent with prospective dread being itself discounted in time.

6.1.2.4 Group Level Modelling Analysis

I compared alternative accounts for the observed patterns of time preference by fitting models of dread-discounting. The starting point for these models was the model of anticipatory utility introduced by Loewenstein (1987). Loewenstein used the term ‘dread’ to refer to the anticipation of events with negative utility, and ‘savouring’ to refer to anticipation of positively valued events. Applied to the evaluation of future pain, Loewenstein’s model describes the utility for a decision-maker at time $t = 0$ of anticipated pain starting at time $t = d$. Pain itself is assumed to produce a constant stream of disutility $U(x)$, beginning at time d and lasting for duration L such that:

$$U^P(x, t, d, L) = \begin{cases} U(x) & \text{if } d \leq t \leq d + L \\ 0 & \text{otherwise} \end{cases}$$

(6.1.1)

U^P denotes instantaneous utility derived from pain itself. In applications of this model (and related versions of the model) in this thesis, $U(x)$ is implemented as a power function of the subjective intensity, x , of pain on a ten-point VAS:

$$U(x) = -x^b$$

(6.1.2)

where b is a parameter governing the curvature of utility function; $0 < b < 1$ produces concave (diminishing marginal) disutility of pain, $b = 1$ denotes linear utility of pain and $b > 1$ produces convex (increasing marginal) disutility of pain.

In addition, whilst waiting for pain, the decision-maker is assumed to derive instantaneous disutility from anticipation, termed dread (U^D), which in Loewenstein’s formulation is given by the (exponentially discounted) expectation of pain at each moment, $t \leq d$:

$$\begin{aligned}
U^D(x, t, d, L) &= \alpha \int_d^{d+L} U(x) e^{-\delta(\tau-t)} d\tau \\
&= \frac{\alpha}{\delta} U(x) e^{-\delta(d-t)} (1 - e^{-\delta L})
\end{aligned}
\tag{6.1.3}$$

Here δ is an exponential discount rate governing the rate of increase in momentary anticipation as the pain draws nearer, and is distinct from the main discount rate. When evaluating delayed pain from $t = 0$, the decision-maker is then assumed to derive disutility from two sources, the prospective disutility of pain itself, and a prospective estimate of the total moment-by-moment dread they can expect to experience whilst waiting for pain. Under Loewenstein's formulation both sources of utility are simply added together, and each is also assumed to be exponentially discounted with rate, r , giving:

$$\begin{aligned}
U(x, d, L) &= \underbrace{\frac{\alpha}{\delta} \int_0^d U(x) e^{-\delta(d-t)} (1 - e^{-\delta L}) e^{-rt} dt}_{\text{[Utility from dread]}} \quad \underbrace{\int_d^{d+L} U(x) e^{-rt} dt}_{\text{[Utility from pain]}} \\
&\tag{6.1.4}
\end{aligned}$$

The first integral denotes prospective utility from dread, the second prospective utility from pain itself. Integrating both terms gives:

$$U(x, d, L) = U(x) \left[\frac{\alpha}{\delta(\delta - r)} (e^{-rd} - e^{-\delta d}) (1 - e^{-\delta L}) + \frac{1}{r} e^{-rd} (1 - e^{-rL}) \right]
\tag{6.1.5}$$

The contribution of dread is governed by the parameters α and δ . α can be viewed as a regression weight on dread, representing an individual's propensity to dread future pain, while δ determines the change in dread as pain is delayed further into the future.

Loewenstein's model applies to outcomes occurring in continuous time. For the current experiment however, delays were described in discrete time, in units of trials. I therefore adapted the above model for the discrete-time case. Since each

painful outcome was considered to last for one time-period, dependence on the duration of pain, L , was omitted. Loewenstein's formulation uses exponential discounting. Although different forms of discounting function, such as hyperbolic and quasi-hyperbolic, are of importance in standard models of financial discounting, they have a relatively subtle effect here, by comparison to the much coarser changes in the shape of the functions resulting from the addition of dread, and I therefore retained exponential discounting for computational expediency. To reflect notation used in reinforcement learning frameworks, exponential discounting was implemented by multiplying the utilities by a constant discount factor, $0 < \gamma < 1$ for each added unit of delay. I sought to compare alternative restrictions to Loewenstein's formulation above.

Each model shares the same general form, but the models differ in the manner in which dread varies as a function of time and outcome magnitude (see Figure 6.4). Figure 6.5 displays choice probabilities on the paradigm itself predicted by representative parameterisations of each model (blue lines), as well as the mean choice probabilities from each of the four participant sub-groups. (Notably the shape is also dependent on the parameters of each model and the parameter dependence of the more complex models is illustrated in Figures 6.6 and 6.7).

All models assume that the total aversive value of a prospective option, $U(x, d)$ is equal to the aversive value associated with accumulated sum of future dread up until the time, $d - 1$, immediately preceding the actual pain plus the aversive value related to the pain itself:

$$U(x, d) = \alpha \sum_{t=0}^{d-1} U^D(x, t, d) + U^P(x, d) \quad (6.1.6)$$

As previously, the overall weighting given to dread is governed by α .

All models in this case assumed linear instantaneous disutility of pain, such that:

$$U(x) = -x \quad (6.1.7)$$

where x represents the expected shock rate ($b = 1$ in Equation 6.1.2). For all models γ_P represents the discount factor applied to the future pain, and x represents the shock rate. Thus:

$$U^P(x, d) = U(x) \gamma_P^d \quad (6.1.8)$$

What differs between the models is the time course of dread, $U^D(x, t, d)$.

Model 1: The first model, the Null model, assumes that the prospective disvalue of pain depends only on the stated number of shocks, and not on the delay, i.e. $\alpha = 0$ and $\gamma_P = 1$ such that:

$$U(x, d) = U(x) = -x \quad (6.1.9)$$

As shown in Figure 6.5, this model predicts a net 50% probability of choosing later shocks on the current choice paradigm. This arises from the fact that on half the presented options the later outcome carries a larger number of shocks and on half the options the later outcome carries a smaller number.

Model 2: The second model, termed the Exponential Discount model, extends the Null model to include exponential discounting of pain, γ_P , but again assumes $\alpha = 0$, such that:

$$U(x, d) = U(x) \gamma_P^d \quad (6.1.10)$$

The Exponential Discount model predicts a preference for later shock (Figure 6.5).

Model 3: The third model, which is the simplest model to incorporate dread, assumes a constant benefit or cost from anticipation accruing from any delayed outcome, which does not scale with the size of the outcome or the delay to its delivery. In other words, outcomes accrue a fixed cost from anticipation.

$$U(x, d) = A + U(x) \gamma_P^d \quad (6.1.11)$$

Where A is the fixed constant of dread, bounded to be negative.

Model 4: The fourth model, Constant Dread, assumes that dread is constant over time, such that:

$$U(x, d) = U(x) \gamma_P^d + \alpha \sum_{t=0}^{d-1} U(x) \quad (6.1.12)$$

Hence dread accumulates linearly with proportionality constant α when pain is viewed from the perspective of an increasing delay:

$$U(x, d) = U(x)(\alpha d + \gamma_P^d) \quad (6.1.13)$$

Model 5: The fifth model, Exponential Dread, following Loewenstein, assumes an increasing time-course of instantaneous dread, such that dread increases exponentially until the actual time of pain. Parsimoniously, the model assumes that the exponential rate governing the increase in dread is identical to the rate by which pain is discounted, γ_P , such that dread becomes simply equivalent to the predicted (discounted) value of future pain. This differs from Loewenstein's original formulation, where the increase in dread was separate from the main discount factor. The rationale for this difference was to separately parameterise discounting of dread itself, whilst keeping the total number of parameters within sensible limits. Thus the full Exponential Dread model allows the exponential rise of dread, governed by γ_P , to be itself discounted in time by a further discount factor, γ_D , such that prospective dread each time point t during the delay, d , is given by:

$$U^D(x, t, d) = U(x) \gamma_P^{d-t} \gamma_D^t \quad (6.1.14)$$

Such that:

$$U(x, d) = U(x) \left[\gamma_P^d + \alpha \sum_{t=0}^{d-1} \gamma_P^{d-t} \gamma_D^t \right] \quad (6.1.15)$$

Figure 6.6 outlines the parameter dependence of the general form Exponential Dread model. I tested this general form, as well as two nested variants. The first variant, termed Undiscounted Exponential Dread assumes that dread is not itself prospectively discounted, such that $\gamma_D = 1$, giving:

$$U(x, d) = U(x) \left[\gamma_P^d + \alpha \sum_{t=0}^{d-1} \gamma_P^{d-t} \right] \quad (6.1.16)$$

Figure 6.7 outlines the parameter dependence of this Undiscounted Exponential Dread model, which predicts that the cost of waiting for pain increases with increasing delay, though at a decreasing rate. The second variant, Restricted Discounted Exponential Dread, assumes that the rate of exponential rise of dread is the same as the rate of exponential discounting of dread, such that $\gamma_D = \gamma_P = \gamma$, giving:

$$U(x, d) = U(x) \left[\gamma^d + \alpha \sum_{t=0}^{d-1} \gamma^{d-t} \gamma^t \right] \quad (6.1.17)$$

Which further simplifies as follows:

$$U(x, d) = U(x) \gamma^d [\alpha d + 1] \quad (6.1.18)$$

If dread itself is prospectively discounted ($\gamma_D < 1$), reversals of time preference can occur, such that prospective pain is maximally aversive at intermediate delays (Figures 6.4, 6.5 and 6.7).




Figure 6.4 Temporal value functions predicted by alternative dread-discounting models.

For each panel the value of a shock is plotted against delay to its delivery. The value of immediate shock is given by the intersection of the curves with the vertical axis; the scales of the vertical axes are arbitrary and differ between the plots. Parameters of the function are displayed next to each. The top left panel depicts simple exponential discounting with positive rate, with the result that the prospective utility of shock becomes less negative the further it is delayed into the future. The top right panel depicts a model in which all values of delayed shock carry a fixed subtractive cost, A (here set arbitrarily to a value of 5), with the discount factor γ set to 0.95. The remaining panels depict dread-discounting models as labelled: Constant Dread, Undiscounted Exponential Dread (denoted by the prefix U) and Discounted Exponential Dread (denoted by the prefix D). Two versions of the Discounted Exponential Dread model are depicted, a restricted version in which the discount rate used to determine dread, γ_D is equal to the discount rate applied to consumption of shock, γ_P (bottom left) and an unrestricted version in which these parameters are independent (bottom right)

Utilities predicted by each model were transformed into predicted probabilities of choosing either option according to a softmax activation function, and models were fitted by maximum likelihood, using simplex optimisation in Matlab (Mathworks, MA, USA), with bound constraints by transformation (see General Experimental Methods). In addition, for models with smaller numbers of parameters, grid search was performed over the entire parameter space as a check against local minima. Models with differing architecture and complexity were compared using the Bayesian Information Criterion (*BIC*) according to a fixed-effects scheme, by summing the *BIC* values for model fits to individual data. According to conventional criteria, a *BIC* difference of greater than 6 between two models corresponds to strong evidence in favour of one particular model (Kass and Wasserman, 1995).

Consistent with the observed significant negative time preference at the group level, Null model did not adequately capture the group level data (*BIC* = 5009). The Exponential Discounting model, incorporating discounting of pain but not dread, improved the group-level likelihood by comparison with the Null model, consistent with some participants demonstrating positive time preference; however the improvement in likelihood was not sufficient to compensate for an increase in complexity over the null model (*BIC* = 5133). Despite the additional parameters, the Fixed Delay Cost model substantially outperformed both the Null and Exponential Discount models (*BIC* improvements of 368 and 492 respectively). In turn, the time dependence of total dread embodied by the Constant Dread model substantially improved the fit compared with Fixed Delay Cost (*BIC* improvement of 321, see Figure 6.8A). However, consistent with the observation of diminishing negative time preference, Exponential Dread models were the best performing at the group level (*BIC* improvement of 149 over Constant Dread for general form Exponential Dread). The restricted variants were favoured over the general form, with Undiscounted Exponential Dread providing the most parsimonious fit out of all models tested at the group level (*BIC* improvement of 150 over the general form), followed by Restricted Discounted Exponential Dread (*BIC* improvement of 57 over the general form).




Figure 6.5 Model predictions on the task: Experiment 1A $p(\text{Choose S2})$ as a function of delay difference according to alternative models of dread. Choice probabilities shown are the mean of those on the two frames. Delay difference scaling is identical that in Figure 6.3. The fine gray lines represent mean $p(\text{choose S2})$ for the four participant subgroups shown in Figure 6.3. Data points marked by blue squares, joined with lines for illustrative purposes, represent model data simulated at the parameter values denoted in each panel. These do not represent the results of model fitting, but serve to illustrate the basic form of the alternative model predictions. Notably different parameterisations of the more complex models can produce diverse shapes of choice frequency plot. Error bars represent one standard deviation of the binomial distribution. In each case the softmax inverse temperature parameter, β , is set to 0.25, a representative value.




Figure 6.6 Parameterisation of the general form Exponential Dread model. A range of temporal value functions predictable by an Exponential Dread model with separate γ_D and γ_P , at different values of the three parameters γ_D, γ_P and α . The model allows for points of maximal aversion at intermediate values of delay. **A:** Effects of increasing α (left to right) at lower (top row) and higher (second row) values of γ_P with a high value of γ_D . At small values of α the model approaches positive exponential discounting (top left panel). Simple exponential discounting is produced when $\alpha = 0$. **B:** Effects of decreasing γ_D (i.e. increasing the discounting of dread; left to right) at lower (top row) and higher (second row) values of α .




Figure 6.7 Parameterisation of the Undiscounted Exponential Dread model. A range of temporal value functions predictable by an Exponential Dread model with $\gamma_D = \gamma_P$, where dread itself is not subject to discounting, at different values of the two free parameters γ and α . **A:** Effects of increasing α (left to right) at lower (top row) and higher (second row) values of γ . It is evident that at small values of α the model approaches positive exponential discounting (top left panel). Simple exponential discounting results when $\alpha = 0$. At positive α aversiveness (negative value) increases at a decreasing rate with delay, where both α and γ influence the asymptotic boundary. **B:** Effects of decreasing γ (left to right) at lower (top row) and higher (second row) values of α .

6.1.2.5 Group Level Framing Effects

Across participants, the mean frequency of choosing sooner pain in the *pain frame* was 73.7% (SE = 3.95) and 67.3% (SE = 4.22) in the *relief frame* (Figure 6.8B). A two-tailed paired t-test showed that this difference was significant: $t(32) = 2.84$, $p = 0.0077$. The result was confirmed with non-parametric testing for differences between paired samples using the Wilcoxon Signed Rank test, which revealed significant differences between the median choice frequencies ($N = 33$, $Z = -2.6$, $p = 0.0093$).

Analysis of the framing effect was extended using the general form Exponential Dread model to test which parameters of the model best accounted for the differences in intertemporal choice between *pain* and *relief* frames. Thus, I performed a second model comparison whilst restricting which parameters were allowed to vary between frames. The general form model was chosen for this analysis, in order that predictions regarding the basis of the framing effects were not dependent on accepting a specific version of the model. The results of model comparison across frames are shown in Figure 6.8C.

The two models with the lowest *BIC* were the model in which α is allowed to be free (i.e. the discount factors γ_P and γ_D and the choice temperature β are fixed; $BIC = 4142$), and the model in which γ_P is allowed to be free (i.e. the dread discount factors γ_D , the dread weighting parameter, α , and the choice temperature are fixed; $BIC = 4115$), with a *BIC* difference of 27 between these two models, indicating substantial support for the γ_P -framing model. These results suggest that the framing effect was most parsimoniously accounted for by changes in the discount rate governing the rate of increase in the dread of future pain, γ_P . Likelihood ratio tests rejected the (null) No-Framing model in favour of the γ_P -Framing model at both the group (fixed effects) (LR = 10^{108} :1, $\chi^2 = 497.3$, $p < 0.001$, $d.f. = 25$) and individual (Mean individual LR = 2088:1, $\chi^2 = 19.9$, $p < 0.001$, $d.f. = 1$) levels.




Figure 6.8 Model comparison and framing effects in Experiment 1A. **A** Bayesian Information Criterion (BIC), summed across participants ($N = 25$) for the alternative models. Lower values of BIC indicate better fits of the model. Exponential Dread outperformed other models, with Undiscounted Exponential Dread providing the most parsimonious fit at the group level, indicated by the red circle. **B** Mean frequency of choosing sooner pain across all choices by all participants in either frame. Error bars are one standard error. Two-tailed paired t -test showed significant difference between the two frames $t(32) = 2.84$, $p = 0.0077$. **C** Results of fitting the (Discounted) Exponential Dread model, whilst restricting which parameters were allowed to vary between frames. In the unrestricted model (All-Framing) all four parameters, the inverse softmax temperature, β , the discount parameters, γ_P and γ_D , and the anticipation parameter, a , were applied separately to each frame. In the fully restricted model (No-Framing) all parameters were constrained to be equal across frames. The best fit (red circle) was provided by a model in which β , γ_D and γ were fixed across frames, leaving between-frame differences explained by differences in γ_P (γ_P -Framing). Likelihood ratios are displayed at both the group level and the individual level, favoring γ_P -Framing over No-Framing at both the group (fixed effects) ($LR = 10^{108}:1$, $\chi^2 = 497.3$, $p < 0.001$, $d.f. = 25$) and individual levels (Mean individual $LR = 2088:1$, $\chi^2 = 19.9$, $p < 0.001$, $d.f. = 1$).

6.1.2.6 Modelling Analysis of Sub-Groups

At the sub-group level, I hypothesised that the Null model would perform best in the *zero* time preference group, that Exponential Discounting would perform best in the *positive* time preference group, Undiscounted Exponential Dread would perform best in the *negative* time preference group and either general form or Restricted Discounted Exponential Dread would perform best in the *reversal* group. In the *zero* time preference group, the two models with the lowest *BIC* estimates were indeed the Null model ($BIC = 755$) and Undiscounted Exponential Dread ($BIC = 793$), suggesting evidence in favor of the Null model in this group. In the *positive* time preference group the two models with the lowest *BIC* estimates were Fixed Delay Cost ($BIC = 961$) and Exponential Discounting ($BIC = 964$), with a *BIC* difference of 3.5, unexpectedly providing weak evidence in favor of the Fixed Delay Cost model in this group. Closer inspection revealed that the improved fit of this model was driven by a single participant's data. As expected, in the *negative* time preference group the best performing model was Undiscounted Exponential Dread ($BIC = 1868$), followed by Restricted Discounted Exponential dread ($BIC = 1941$). As expected, in the *reversal* group, the best performing model was Restricted Discounted Exponential dread ($BIC = 370$), followed by the general form Discounted Exponential Dread ($BIC = 372$) (followed by Undiscounted Exponential Dread, $BIC = 374$), indicating evidence in favour of dread being discounted in this sub-group. Notably, alternative parameterisations of the general form Exponential Dread model are able to account for the patterns of time preference displayed by each of the four sub-groups, as shown in Figure 6.9, which presents data from a single participant within each group, along with the maximum likelihood estimates predicted by the general form Exponential Dread model.




Figure 6.9 Time preference of sample participants in Experiment 1A and fits of the (Discounted) Exponential Dread model. Observed $p(\text{Choose Later})$, combined across both frames, as a function of delay difference, expressed in units of trials, is displayed for a single participant from each of the four subgroups shown in Figure 6.3, indicated by the purple circles. Delay difference scaling is identical that in Figure 6.2. Data simulated from the general form Exponential Dread model at the maximum likelihood parameter estimates for each participant, subsequently combined across frames, are plotted as cyan squares. Error bars represent one standard deviation of the binomial distribution. **A** Participant with zero time preference. **B** Participant with positive time preference (left hand column). **C** Participant with negative time preference. **D** Participant with reversing time preference: showing initial negative time preference reverting to positive time preference at longer delay differences. The general form of the Exponential Dread model adequately captures all four patterns of time preference.

6.1.3 DISCUSSION

Experiment 1A was designed to compare alternative accounts for how the prospective aversiveness of pain depends on its timing, using an experimental paradigm in which participants made choices between painful outcomes occurring at different delays in the future. In accordance with previous studies most participants (26 out of 33) exhibited negative time preference for pain. These participants preferred to experience the same pain sooner rather than later and were willing to accept more pain in order to hasten its occurrence. At the group-level, negative time preference initially increased with increasing delay, but saturated at long delay. This pattern was best accounted for by a dread-discounting model, termed Exponential Dread, in which dread increases exponentially as pain is approached in time. The total utility from dread is then given by the prospective sum of dread, where the extent to which an individual incorporates dread can be described in terms of the weighting parameter, α . I also found that dread is modulated by relief framing, an effect which was captured by modulation in the rate of instantaneous dread increase. These findings extend an existing study of dread by Berns and colleagues (Berns et al., 2006), who compared a Constant Dread model and an Exponential Discounting model *without dread* in the context of choice between delayed shocks predicted by a cue, finding that the Constant Dread model provided better fits to both the behavioural data and the BOLD response in several regions of interest (right primary and secondary somatosensory cortices, caudal anterior cingulate cortex and right posterior insula) than the model without dread.

Whilst attempt was made to embed Experiment 1A within a naturalistic scenario, the outcomes are nevertheless brief and delays short by comparison to decisions regarding aversive treatments or preventive health behaviours in real life. Furthermore the above conclusions are based on only a relatively small sample of participants. To address this I implemented Experiment 1B, with outcomes of a hypothetical painful dental treatment, and delays ranging up to several months.

6.2 EXPERIMENT 1B: TIME PREFERENCE FOR HYPOTHETICAL DENTAL APPOINTMENTS

In Experiment 1B participants made 70 choices between two possible timings for a hypothetical dental appointment. They were informed that the appointment would last for 15 minutes and that the experience would be painful. Participants were also informed that the appointment was routine and that the timing would not affect their dental health.

6.2.1 METHODS

Thirty healthy adult participants aged 18-55 took part in the experiment. Participants were recruited by advertisement on the website of the University College London Psychology Subject Pool (see General Experimental Methods 5.6). Choices were administered online using Qualtrics software (Qualtrics.com; Provo, UT). Unfortunately specific demographic data of participant characteristics were not obtained for this experiment.

Participants were introduced to the following scenario:

You are due to have a routine dental appointment. The appointment will last for 15 minutes. The appointment is non-urgent but must be booked now to occur sometime in the next year. Importantly the timing will not affect your dental health: having the appointment sooner will provide no added health benefits. However the experience will be very uncomfortable, and at times painful. Exactly how painful will depend on when you choose to have the appointment. The pain will only last for the length of the appointment: you do not experience dental pain at any other time. The dental surgery is very close to where you live, so you will be able to attend an appointment almost immediately if you choose to. As far as your diary is concerned, any of the appointment times are equally possible for you.

Participants were offered binary choices between different timings for the dental appointment, in units of days. The sooner appointment was always designated

as occurring “today”. The later appointment occurred at delays of 1, 5, 13, 32, 89 or 237 days. In each case participants were told how painful they could expect the appointment to be on a scale of 0% to 100%, where 100% represents the worse imaginable dental pain. The outcome magnitudes were 60, 55, 51 46, 37 and 16% dental pain. At each possible delay, each possible magnitude was paired with an outcome of 60% dental pain. Options were counterbalanced, so that there were an equal number of choices in which the larger magnitude pain was the delayed option as choices in which the smaller magnitude pain was the delayed option. So that the timings appeared plausible, participants were asked to imagine that it was currently a weekday morning. Data analysis followed the same methodology as described above for Experiment 1A. Since outcome magnitude was described in experiential terms, a linear utility function was assumed over percentage pain intensities.

After completing the questionnaire, participants were asked debriefing questions, requiring them to indicate when they last attended the dentist, and to what extent (on a 5-point scale) they were afraid of going to the dentist.

6.2.2 RESULTS

6.2.2.1 Group Level Time Preference

All 30 participants were included in the analysis. As for Experiment 1A, there were an equal number of choices for which the larger magnitude pain was the sooner option as choices in which the (identical) larger magnitude pain was the later option, therefore the probability of choosing the later option, in this case $p(\text{Choose A2})$ where A2 refers to Appointment 2, reflects time preference in the same manner as for Experiment 1A.

Consistent with Experiment 1A, at the group level, participants showed a strong preference for sooner dental appointments, at the expense of more severe dental pain: group mean $p(\text{Choose Later Appointment})$ averaged across all delay lengths was significantly less than 0.5, [mean $p(\text{choose S2}) = 0.38$, $N = 30$, $SE = 0.025$, One sample $t(29) = -4.56$, $p < 0.001$], a result which was confirmed with non-parametric testing [median $p(\text{choose S2}) = 0.39$, One sample Wilcoxon Signed Rank test $p < 0.001$], indicating overall negative time preference.




Figure 6.10 Time preference for a hypothetical painful dental appointment in Experiment 1B. Observed $p(\text{Choose A2})$ is plotted as a function of the delay to the later appointment; the sooner appointment was always at 0 days, i.e. “today”. Error bars represent one standard error from the group mean. **A** Group mean $p(\text{Choose A2})$ for all participants ($N = 30$). **B** Mean $p(\text{Choose A2})$ in the zero time preference group ($N = 12$). **C** Mean $p(\text{Choose A2})$ in the positive time preference group ($N = 3$). **D** Mean $p(\text{Choose A2})$ in the negative time preference group ($N = 15$).

To facilitate comparison with Experiment 1, delays to the later dental appointment were grouped into short (1-5 days), medium (13-32 days) and long (89-237 days) categories. A one-way analysis of variance (ANOVA) revealed a significant main effect of delay [$F(3,116) = 8.1$ $p < 0.001$]. Non-parametric pairwise comparisons between zero, short, medium and long delays revealed a significant decrease in $p(\text{Choose A2})$ between zero and short delay differences [$p(\text{Choose S2, Later Appointment}) < 0.5$, Wilcoxon Signed Rank test, $p < 0.001$], between short and medium delay differences (Wilcoxon Signed Rank test, $p = 0.0025$), and between medium and long delay differences (Wilcoxon Signed Rank test, $p = 0.022$),

suggesting consistent negative time preference at the group level. Figure 6.10A displays the group mean $p(\text{Choose A2})$ at all delay lengths offered. The data reflect the finding of consistent negative time preference. Diminishing negative time preference with increasing delay can be clearly appreciated from the shape of the plot.

6.2.2.2 Classification of Participants by Time Preference

I categorised the 30 participants according to their individual pattern of time preference, using an identical method as for Experiment 1A: 12 out of 30 participants showed *zero* time preference, 3 out of 30 showed *positive* time preference, and the remaining 15 out of 30 showed *negative* time preference. In this experiment, no participants displayed *reversing* time preference, i.e. none displayed significant increases in $p(\text{choose S2})$, as well as significant decreases. Figures 6.10 B-D display group mean $p(\text{Choose A2})$ at each delay length for the *zero*, *positive* and *negative* time preference groups respectively.

6.2.2.3 Group Level Modelling Analysis

The results of group level model comparison are shown in Figure 6.11. Consistent with the shape of the group level choices (Figure 6.10), Undiscounted Exponential Dread was the best performing model at the group level ($BIC = 1036$), followed by Constant Dread ($BIC = 1115$). These two models outperformed both the restricted variant of Discounted Exponential Dread ($BIC = 1162$), general form Discounted Exponential Dread Constant Dread ($BIC = 1179$), Fixed Dread Cost ($BIC = 1418$). Exponential Discounting ($BIC = 1784$) and the Null model ($BIC = 1791$).




Figure 6.11 Model comparison for Experiment 1B. Bayesian Information Criterion (*BIC*), summed across participants ($N = 30$) for the alternative models. Lower values of *BIC* indicate better fits of the model. Undiscounted Exponential Dread provided the most parsimonious fit at the group level, indicated by the red circle.

6.2.2.4 Modelling Analysis of Sub-groups

At the sub-group level, a plausible hypothesis was that the Null model would perform best in the *zero* time preference group, that Exponential Discounting would perform best in the *positive* time preference group and that Undiscounted Exponential Dread would perform best in the *negative* time preference group. In the *zero* time preference group, the best performing model was in fact Undiscounted Exponential Dread ($BIC = 218$), followed by the restricted form of Discounted Exponential Dread ($BIC = 229$), followed by the Null model ($BIC = 232$). This unexpected result is consistent with the observation that participants in the *zero* time preference group displayed a small but consistent degree of negative time preference, as shown in Figure 6.10B. In the *positive* time preference group the two models with the lowest BIC estimates were Exponential Discounting ($BIC = 176$) and Undiscounted

Exponential Dread ($BIC = 179$), a BIC difference of 3.1, providing weak evidence in favor of Exponential Discounting in this group. Similarly, as predicted, in the *negative* time preference group the best performing model was Undiscounted Exponential Dread ($BIC = 633$), followed by Constant Dread ($BIC = 674$).

6.2.2.5 Debriefing on Dental Attendance and Dental Fear

The mean time since last attending the dentist was 24.6 months (s.d. 51.1 months). The data were however considerably skewed, and median time since last attendance was 8.0 months (interquartile range 4.5-18.0). The mean level of dental fear was 2.73/5 (s.d. 1.08). Although I did not obtain direct data on whether participants had previously experienced dental pain, taken together these responses suggest that participants were familiar with attending the dentist and reported a degree of apprehension about attending.

6.2.3 DISCUSSION

A multitude of factors, which may interact with the effect of delay, are likely to influence the valuation of future pain in real-world contexts. Nevertheless Experiments 1A and 1B show the functional form of dread appears conserved across two very different experimental contexts: in the context of real painful outcomes experienced at delays of up to approximately 15 minutes, and in the context of an imagined painful experience at delays of up to approximately 8 months.

Notably, although debriefing data suggest that many participants had attended the dentist in recent months, and that the mean participant exhibited some level of dental fear, the extent to which participants were able to adequately imagine dental pain is unclear. A further key limitation of this experimental scenario is that participants may have held a prior belief that seeking sooner treatment is beneficial, either to ameliorate ongoing pain or to improve their dental health. I endeavoured to reduce this possibility by emphasising in the instructions that the appointment was routine and that the timing of the appointment would not affect dental health. Experiments described later in this thesis (Chapter 8) use a scenario designed to further minimise these possible confounds. The latter scenario requires participants to

imagine receiving a painful injection, said to prevent the onset of dementia later in life, and said to be 100% effective provided it is received sometime in the next ten years. In the dementia scenario, the time of onset of illness is distant relative to the treatment, hopefully circumventing a prior belief that earlier treatment might be more beneficial.

A further relevant observation here is that the form of the dread function appears to demonstrate scale invariance, as evidenced by a similar shape when making choices over delays expressed in different units of time (trials or days). As a result the magnitude of prospective dread at a given delay is likely to depend upon the psychological construal of the time scale. Scale invariance is a feature of many psychometric functions, including temporal discounting with rewarding outcomes, and the scale invariance of dread presents a possible target for future study.

6.3 EXPERIMENT 1C: TIME PREFERENCE FOR HYPOTHETICAL ILLNESS

Experiment 1B examined how people might dread a future painful treatment, motivated by the idea that measures of dread might ultimately predict engagement in unpleasant preventive health procedures. However, people also exhibit negative time preference for future illness (e.g. Van Der Pol & Cairns, 2000), suggesting that attendance for preventive procedures might represent a complicated interplay between the dread of the procedure itself and the dread of delayed illness. (In many instances the procedure itself might be further dreaded for the reason that it might reveal that one is likely to become ill). To go some way towards examining this, Experiment 1C compares the ability of models of dread to account for peoples' preferences for the timing of a transient illness.

6.3.1 METHODS

Thirty healthy adult participants (18 female, mean age 22.85 years, s.d. 5.45 years) took part in the experiment. Participants were recruited by advertisement on the website of the University College London Psychology Subject Pool (see General

Experimental Methods 5.6). Choices were administered online using Qualtrics software (Qualtrics.com; Provo, UT).

I used a method for assessing illness discounting introduced by Van der Pol and Cairns (2002). Following Van der Pol and Cairns, participants received the following brief:

Imagine an illness that affects you in the following ways:

“You have some problems with performing your usual activities (e.g. work, study, housework, family or leisure activities) and you have moderate pain or discomfort. You have no problems in walking about, nor with washing or dressing yourself and you are not anxious or depressed.”

All uses of 'ill' and 'ill-health' in the questionnaire refer to the above imagined illness. You will be asked some brief questions relating to this illness.

In each case a treatment will be offered that will postpone the ill-health into the future, however this treatment may increase or decrease the number of days of ill-health experienced. You are asked to state the maximum number of days of the future (postponed) ill-health which would still make it worthwhile for you to accept the treatment.

Questions then assumed the following format:

Imagine that you will be ill starting 2 years from now for 20 days. There is a minor one-off treatment available that will postpone this spell of ill-health to a point further into the future. If the ill-health would then start d years from now what is the maximum number of days of future ill-health (x) that would still make the treatment worthwhile?

Each participant answered 9 of these questions, with delays, d_2 , of 4, 7, 10, 12, 15, 21, 33, 45 and 53 years. The experimental data consisted of the stated number of days of future illness, x , termed x_2 to refer to the delayed illness option. Positive time preference (discounting) of illness is indicated by $x_2 > 20$, corresponding to

taking on extra days of illness so as to postpone it; by contrast $x_2 < 20$ indicates negative time preference, and corresponds to taking on extra days of illness so as to avoid postponing it.

6.3.1.1 Models of Dread for Future Illness

The same models of dread-discounting introduced in Experiments 1A and 1B were tested here. To achieve continuity with the models above, I retained the discrete-time framework and assumed that illness simply lasted for a single-time period, in this case one year. The disutility of illness was assumed to scale linearly with the number of days indicated by participants. These are radical assumptions, and further work is required to test whether they are justified, however they carry the advantage of allowing the severity of the illness, x , to be subsumed within the instantaneous utility function, rather than being time-extended, and thus greatly simplify the analysis. In each case the sooner illness option was given by an illness occurring in 2 years' time ($d_1 = 2$), and lasting for 20 days ($x_1 = 20$), and participants were asked to state how many days of illness, x_2 , at a specified delay, $4 \leq d_2 \leq 52$, would make them indifferent between the two options. Therefore to generate predicted values of x_2 for each model I solved the following:

$$U(x_2, d_2) = U(20, 2) \quad (6.3.1)$$

Recall that the equation for the full Exponential Dread model (Equation 6.1.15 above) can be decomposed into the instantaneous utility function, $U(x)$, and the dread-discounting function (the terms within the square bracket), which depends only on d and the parameters. Denoting the latter by $\Delta(d)$ gives:

$$U(x_2, d_2) = U(x_2)\Delta(d_2) = U(20)\Delta(2) \quad (6.3.2)$$

with linear disutility of illness this gives:

$$x_2 = 20 \frac{\Delta(2)}{\Delta(d_2)} \quad (6.3.3)$$

To implement a Fixed Dread Cost, I assumed a cost, A , associated with choosing the more delayed illness, such that:

$$U(x_2, d_2) - A = U(20)\Delta(2) \quad (6.3.4)$$

Giving:

$$x_2 = 20 \frac{\Delta(2) + A}{\Delta(d_2)} \quad (6.3.5)$$

To compare models with different numbers of parameters, conventional BIC was estimated from the least squares model fits, by assuming that errors were normally distributed around the model estimate and using the normal probability density as a likelihood function.




Figure 6.12 Time preference for illness in Experiment 1C predicted by an Undiscounted Exponential Dread model. **A** Discounted utility of 20 days of illness, $U(x = 20, d)$ as a function of delay, d , in years, to the onset of illness, for $\gamma_p = 0.95$ at increasing values of α . **B** Predicted behaviour on the illness discounting questionnaire at the same parameter settings as shown in **A**, given by the duration of delayed illness, x_2 , at each delay, d , that is subjectively equivalent to 20 days of illness delayed by 2 years.

6.3.2 RESULTS AND DISCUSSION

Figure 6.12A shows the dread-discounting functions predicted by an Undiscounted Exponential Dread model, for $x = 20$, $\gamma_p = 0.95$ at varying settings of α . Figure 6.12B shows the consequent predicted values of x_2 on the illness-discounting task. Exponential discounting alone ($\alpha = 0$) predicts exponentially increasing x_2 as a function of delay. Low values of α combined with discounting (top

right panels) create concave increasing dependence of x_2 on delay. Higher α gives rise to negative time preference, whereby x_2 declines as a function of delay.




Figure 6.13 Time preference for illness in Experiment 1C. **A** Observed group level time preference (red circles) indicated by the mean duration of delayed illness, x_2 , at each delay, d , that is subjectively equivalent to 20 days of illness delayed by 2 years. Error bars represent one standard error above and below the group mean. Note that the first data point, $x_2 = 20$ at $d = 2$ is not part of the observed choices, but is imposed in theory as an intercept. Hollow squares joined by a dashed line represent the mean least squares fit of a Discounted Exponential Dread model with Fixed Delay Cost (formed by taking the mean over individual least squares estimates). **B** Time preferences of four sample subjects displaying (from left to right), predominant discounting with low dread, predominant discounting, predominant dread, and predominant discounting with an initial cost.

One participant was excluded from the analysis since they answered $x_2 = 0$ for all questions. Observed group mean x_2 for the remaining 29 participants is plotted as a function of delay in Figure 6.13A; x_2 is an increasing concave function of delay, consistent with positive discounting combined with low α . However, the intercept of the curve lies below 20: mean x_2 for the shortest delay (4 years) was significantly less than 20 (95% CI [13.41- 19.27], one-sample $t(28) = 2.56$; $p=0.01$). A possible explanation is that participants judged there to be a cost associated with treatment,

such that a treatment which did not delay illness was only deemed worthwhile if it reduced the number of days of illness. This appears plausible in view of the fact that the wording of the scenario did not especially emphasise that the treatment was trivial. Since the questionnaire did not incorporate a question in which treatment did not delay illness ($d_2 = d_1 = 2$), this explanation is indistinguishable from a fixed cost associated with delaying illness.




Figure 6.14 Model comparison for Experiment 1C **A** Bayesian model comparison: the best fitting model was Discounted Exponential Dread with a Fixed Delay Cost, indicated by the red circle. **B** Observed x_2 (vertical axis) is plotted against that predicted by the best-fitting model (horizontal axis) for all choices made by all participants, yielding $R^2=0.971$.

Participants showed heterogeneous responses, with four example patterns shown in Figure 6.13B. Notably several participants displayed a pattern consistent with negative time preference for illness at all lengths of delay, in keeping with the results of Experiments 1A and 1B. Also in keeping with the previous two experiments, Exponential Dread models outperformed Fixed Delay cost and Constant Dread models. In this case the full Discounted Exponential Dread model outperformed its Undiscounted counterpart (Figure 6.14A, $\Delta BIC=75$). To model the observation of an intercept less than 20, I also fitted *post hoc* a model which combined Discounted Exponential Dread with a Fixed Delay Cost (equivalent to a fixed cost associated with treatment). This model outperformed Discounted

Exponential Dread alone ($\Delta BIC=59$) and provided the most parsimonious fit at the group level ($R^2=0.971$; Figure 6.14B).

In summary, Experiment 1C aimed to test the dread-discounting models introduced in preceding experiments against intertemporal choices for future illness, using a previously published paradigm. The sample size for this study was relatively small, given that the study was conducted online ($N=30$). The above data can therefore be seen as an exploratory attempt to characterise time preference for illness in terms of dread discounting. The current data support the hypothesis that models of dread-discounting can be used to delineate quite different patterns of time preference exhibited for hypothetical illness.

6.4 DREADING FUTURE PAIN: GENERAL DISCUSSION

The experiments described above address the phenomenon of negative time preference for aversive events, by offering participants choices between different intensities of pain at different delays. Analysis of such choices permitted mapping out the form of time preference for individual participants, and at the group-level. Model comparison indicated that choices across all three modalities were consistent with a model of dread in which the total prospective cost of waiting for pain accumulates with increasing delay (but at a decreasing rate), providing empirical support for the exponential model of anticipatory utility proposed by Loewenstein.

Under the Exponential Dread model implemented here, as for Loewenstein's original formulation (Loewenstein, 1987), both moment-by-moment dread and the temporally discounted value of pain itself increase as pain is approached in time. An increasing aversiveness by time function for the anticipation of pain bears similarity to observations in studies of fear conditioning. For example, the ability of fear to potentiate the startle response is specific to the learned time interval between conditioned stimulus (CS) and unconditioned stimulus (UCS) (M. Davis, Schlesinger, & Sorenson, 1989; Grillon, Ameli, Merikangas, Woods, & Davis, 1993; L. E. Ross, 1961). Thus, following the CS, fear behaviours increase to reach a maximum at the predicted time of UCS onset. Similarly, in human subjects instructed to expect shock after a stated delay, physiological measures of fear such as galvanic skin response

(GSR) and heart rate both increase roughly exponentially in the period immediately preceding the predicted time of shock delivery (Boucsein & Wendt-Suhl, 1976; Breznitz, 1968; Estes & Skinner, 1941; Folkins, 1970). Consequently, the anticipation of pain can be considered as resembling a temporally discounted value signal, assuming a low level when pain is distant and increasing as pain is approached. The findings here suggest that pain-related decision-making incorporates, to a varying degree, a prospective sum of this anticipation, comprising the dread term of the dread-discounting model.

The current work on dread was motivated by the attempt to predict avoidance of unpleasant preventive health behaviour. A paradox arises with respect to the relationship between dread and avoidance. On the one hand dread increases the aversiveness of future pain, which ought to motivate avoidance; on the other hand if the pain is in fact *unavoidable* then dread motivates speeding up the onset of pain. It is therefore unclear on *a priori* grounds whether higher dread would predict greater avoidance of unpleasant preventive behaviours, or a tendency to get them over and done with. An important variable would appear to be whether a person believes (implicitly or explicitly) that the dreaded event can be avoided through delaying it (see General Discussion).

Experiment 1A finds that choices to expedite pain were more frequent when the same outcomes were framed as an increase in pain than when framed as a decrease in pain, a demonstration that framing biases exert an effect in situations associated with dread. If dread can be shown to influence real-world behaviour, such as decisions to attend for preventive health screening, such framing effects might have implications for strategies designed to encourage attendance. For example, if dread were to be positively associated with the avoidance of unpleasant screening procedures, messages aimed at diminishing dread might prove successful. Such ideas are of course speculative at this stage and require further investigation.

The observation of framing may be similar to that which underlies the well-known sign effect, in where discount rates for rewards are typically different from those for punishments (Loewenstein & Prelec, 1992; MacKeigan et al., 1993; Odum et al., 2002; Ohmura et al., 2005; T. Takahashi, Ohmura, Oono, & Radford, 2009). Indeed, a model in which between-frame differences in temporal value functions were determined by changes in the rate of accumulation of dread, here equivalent to the

discount rate for pain, provided the most parsimonious account for these effects of framing, suggesting that differential anticipation is a sufficient explanation for the sign-effect in this context. This observation is however bound to the framework of the dread-discounting model. It is possible for example that framing induced changes in the instantaneous utility function for pain (Loewenstein & Prelec, 1992).

Why dread is a consistent feature of pain related decision-making is unclear. One possibility is that cognitive and emotional mechanisms associated with preparation for pain interfere with other behavioural processes, such as those involved in reward seeking. It is known for example that non-contingent prediction of shock, signaled by a conditioned stimulus, can reduce the vigor of instrumental responding, an effect referred to as conditioned suppression (Estes, 1948; Estes & Skinner, 1941). Dread, as the prospective sum of anticipated punishment, may therefore signal the likely degree of behavioural suppression during the delay. Another possibility is that dread represents a form of ‘stimulus substitution’ – the observation that cues associated with the prediction of aversive events evoke some of the core properties of the aversive events they predict themselves (Ploghaus et al., 1999; Waschulewski-Floruss, Miltner, Brody, & Braun, 1994). An alternative explanation would be that people have an increasing uncertainty with time that they can engage in an adequate physical or psychological response to deal with pain. Further research is required to uncover the constitutive mechanisms of dread. The insights gained might carry importance for clinicians and health policy makers, since knowledge about the shape of pain value functions and their modulation by framing may be useful when presenting options regarding potentially painful investigations and treatments.

ALLOCATING PAIN ACROSS TIME

“In the analysis of the supply of savings, it is extremely doubtful whether we can learn much from considering such an economic man, whose tastes remain unchanged, who seeks to maximise some functional of consumption alone, in a perfect world, where all things are certain and synchronised.”

Paul A. Samuelson – A note on the measurement of utility
(1937)

The preceding chapter of this thesis described how choices to advance the timing of unavoidable pain could be conceptualised as arising from the ‘dread’ of future pain. This work was motivated by the attempt to predict avoidance of unpleasant preventive health behaviour by measuring time preferences for pain or illness. As an intermediate step in this direction, the current chapter presents an experiment designed to test whether the derived dread functions are capable of predicting behaviour in a more naturalistic experimental context, modelled on the use of a personal health budget. Personal budgets ring-fenced for spending on one’s health have been recently introduced to the National Health Service in the UK, forming part of a wider trend towards giving patients more control over their own care (Darzi, 2008; Forder et al., 2012). Peoples’ time preferences for health or illness might influence how they choose to distribute such a budget over time, suggesting for example that people who tend to dread future illness might show a tendency to save a health-related budget, so as to mitigate the dread of future pain or illness.

Conventionally, intertemporal preferences are estimated from ‘one-off’ binary choices between outcomes at varying delays. By contrast, budgeting choices, in keeping with many real-world decisions, are made in a dynamic context, where the results of current choices influence the possibilities available for subsequent choices. Spending money on a holiday for example inevitably diminishes the amount one has available to save for retirement. The classical economic account of intertemporal

choice predicts that one-off and dynamic choices ought to be consistent, since in both cases a rational agent ought to simply allocate their resources over time so as to maximise total (exponentially) discounted utility (Samuelson, 1937; Strotz et al., 1957). However, several features of the classical model are inconsistent with what is known about how people approach dynamic choices.

7.1 DO PEOPLE REALLY PERFORM DYNAMIC UTILITY MAXIMISATION?

A key simplifying assumption of the classical economic model is that utility is independent of that in other time periods, allowing the effect of delay to be entirely described by the temporal discount function (for critiques see Frederick et al., 2002; Loewenstein & Prelec, 1992). In reality however the assumption of additive utility is violated. For instance eating a meal reduces the utility of food for some time afterwards. Similar violations occur prospectively too. For example although, as noted, people overwhelmingly prefer sooner one-off rewards to delayed rewards of equivalent magnitude, when the same rewards are framed as sequences people tend to prefer sequences which improve over time—behaviour which cannot be reconciled with a single discount function whilst also preserving additive utility (Chapman, 2000; Frederick & Loewenstein, 2008; Loewenstein & Prelec, 1993).

The conventional economic model also assumes that humans have the necessary cognitive capacity to optimise their discounted utility. However, when deciding how to allocate reward over several time steps, the number of possible allocation plans grows exponentially as outcomes further into the future are considered, generating decision-problems of considerable complexity (Daw et al., 2005; S. Johnson, Kotlikoff, & Samuelson, 1987). In response to this people apparently adopt simplifying strategies. For instance, transfers into retirement savings plans cluster around the minimum and maximum allowable contributions, as well as around multiples of five dollars, suggesting that investors choose these as convenient ‘rules-of-thumb’ (Benartzi & Thaler, 2007). Such strategies are examples of ‘heuristics’, which are approximate, generic solutions to classes of problems (Baron, 2000; Gigerenzer, 2008; Gigerenzer et al., 2011). Notably the use of heuristics can generate behaviour in a dynamic context that is not consistent with the preferences

that the decision-maker exhibits in simpler, one-shot contexts (S. Johnson et al., 1987; Lettau & Uhlig, 1999).

In addition the form of temporal discount function interacts with the ability to execute one's best-laid plans. As discussed in previous chapters, a decision-maker with an exponential discount function (and an increasing concave utility function over outcome magnitude) has time-consistent preferences—i.e. will make the same decision between options with different temporal profiles no matter how close or far in time these are (Strotz et al., 1957). Such a decision-maker would naturally adhere to their plans, however frequently they were re-evaluated. By contrast if the discount function is hyperbolic then the decision-maker is expected to exhibit dynamically inconsistent behaviour: by seeking immediate reward, they would tend to undo previous long-sighted plans (Ainslie, 2001; Green et al., 1994; Kirby & Herrnstein, 1995; Strotz et al., 1957; however see Chapter 4 of this thesis for alternative accounts). Temporally inconsistent preferences theoretically compound the complexity of planning resource allocations in real-time, since they necessitate a dynamic model of the behaviour of future selves (Ainslie, 1992; Ainslie, 2001).

All the challenges and violations of the conventional economic wisdom described above make it unlikely that individuals exhibit fully optimal intertemporal allocations. However very few studies have directly examined budgeting decisions in real-time, tested the extent to which these are consistent with discount functions derived from one-off choices, or indeed found a parsimonious description that accounts well for actual choices. Furthermore no previous studies have tested this in a health-related context, where anticipatory utility is likely to be prominent.

To address this I designed a task, referred to as Experiment 2, that involved allocating a limited budget in real-time. The task required participants to perform a series of 60 trials over 14 minutes wherein they were scheduled to receive a number of identical, moderately painful, cutaneous electric shock stimuli on each trial. At the outset, each participant was endowed with a fixed budget of computerised pain relief, an amount insufficient to relieve all shocks in the session. On each trial they were allowed to choose how much relief they wished to use, up to a maximum allowable “dose”. The scenario was embedded within a hypothetical health-related context, and pain relief was described in units of milligrams. Figure 7.1 illustrates the experimental protocol. On each trial, participants received a number of shocks drawn from a

Poisson distribution. Without pain relief, the mean of this distribution was 14 shocks; for every 1mg of relief the participants spent on a trial, the mean decreased by 0.1 shocks. Participants were allowed to spend a maximum of 120mg of relief on a trial; this reduced the mean number of shocks to two, a level termed the ‘baseline pain’. Participants had to spend within a total budget of 2400mg. Before making their choice, participants were informed of the total relief capital remaining, the number of trials remaining and the mean remaining relief per trial. The participants had, on the same day, also taken part in Experiment 1A, where they made binary choices between different numbers of, and delays to, painful shock stimuli, which were identical to those used for the consumption-savings experiment. This permitted testing for consistency between one-off and dynamic choices and examining evidence for possible simplifying heuristics.




Figure 7.1 Trial Structure of Experiment 2. Participants entered an experimental run of 60 trials, on which they could expect to receive mildly painful electric shock stimuli on each trial, referred to as painful episodes. Each episode consisted of a five second stimulus in which there was a uniform probability of receiving a brief shock at each sampled time interval. By default participants could expect to receive a five second stimulus with 14 brief shocks on each trial, however they were provided with a budget of relief at the outset of the experiment, 2400 “milligrammes” in total. At the start of each trial a screen indicated the number of remaining trials, the remaining supply of relief in milligrammes (mg) and the mean relief remaining per trial. Participants then had the opportunity to indicate how much relief they wished to consume on that trial. The relief was then effective on the immediately subsequent painful episode. Each 10mg of relief consumed reduced the expected number of shocks in the stimulus train by one. Participants could consume up to a maximum of 120mg of on each trial, sufficient to reduce the expected number of shocks in the upcoming stimulus to two, referred to as the “baseline pain”. Since the budget of relief was insufficient to relieve all the shocks in the session, participants faced an intertemporal resource allocation problem.

Experiment 1A illustrated a range of intertemporal choice behaviour, with some participants exhibiting positive discounting, others negative discounting and others discounting very little or not at all. These three patterns would be expected to give rise in the dynamic task to spending relief early, saving relief for the end and spreading relief evenly over time, respectively (the latter assumes a concave utility function for relief). I therefore predicted that individuals who positively discounted one-off pains would tend to spend their relief early, those who displayed greater negative discounting (dread) for one-off pains would be more likely to save their relief (to mitigate dread, akin to savouring of relief), and those who did not discount pain at all would be more likely to spread their relief across time. I also compared behaviour with the optimal predictions of an anticipation-discounting model fitted to the one-off choices.

The anticipation-discounting functions introduced in the previous chapter can generate non-adherence to past plans of a different form to that entailed by hyperbolic discounting (Figure 7.2) (see also Loewenstein, 1987). Since I did not elicit participants' plans prior to the experiment, I could not directly test for this. However, to explore the theoretical implications of the model in more detail I simulated optimal consumption choices under various parameterisations of the utility and anticipation-discounting functions, allowing for the possibility that participants might have different degrees of insight into their future tendencies, being either inaccurate (*naïve*) or accurate (*sophisticated*) (O'Donoghue & Rabin, 2003).

7.2 METHODS

7.2.1 EXPERIMENT 2: DYNAMIC CONSUMPTION OF PAIN RELIEF

7.2.1.1 Participants

Participant details, recruitment and informed consent are identical to Experiment 1A. The 30 healthy adult participants whose data were included in the study were those taking part in Experiment 1A. Anticipation-discounting parameters were estimable in 23 participants from these thirty. The remaining 7 participants

always choose sooner pain on the binary choice experiment, precluding reliable model fitting.

7.2.1.2 Procedure

Participants were briefed with on-screen instructions that embedded the task in a naturalistic health-related scenario as for Experiment 1A. These were as follows:

Imagine that you are suffering from pains as a result of damage to the nerves in your arm, which feel as if you are having brief electric shocks delivered to your hand. The pains come in clusters of 'shocks', which are roughly evenly spaced in time, with each episode lasting five seconds. Within each episode the shocks come randomly. You experience a usual baseline level of pain, of on average two shocks per episode. All the shocks are of the same intensity.

You are about to have a severe attack of the pains, lasting about fifteen minutes in total. These severe episodes consist of 12 more shocks per episode than your usual baseline level.

Before the attack you are given a supply of 2400mg of painkiller. On average 10mg provides relief from one shock per episode. Before each episode you will be asked how much of your supply of relief you would like to use. There will be 60 episodes in total, so you have on average 40mg to use per episode. You may wish to use more or less than this amount on each episode.

You can use up to 120mg of painkiller for each episode. This will provide relief down to your baseline level of pain.

As described above, at the start of the session participants were endowed with a fixed budget of computerised pain relief, described in units of milligrams, 2400mg in total. The budget was not sufficient to relieve all the shocks in the session, and participants were informed of this fact, and therefore the possibility that they might expend all their relief before the end of the session. Before each trial (of a total of 60), participants were informed of the total number of trials remaining, the number of units of relief remaining and the calculated mean relief remaining per trial in mg. They were then given the opportunity to indicate how much relief they wished to

consume on that trial, by moving a pointer along a visual scale using the keyboard. There followed a painful shock stimulus, the severity of which was determined by the amount of relief consumed.

As for Experiment 1A, the painful shocks were drawn from a Poisson distribution within a five second stimulus train, where the intensity of each discrete shock, which consisted of a single 200 μ s square-wave pulse, did not vary. Each 5s stimulus was partitioned into 50 sampling intervals of 0.1s duration. There was a fixed probability of receiving a shock within each of these sampled 0.1s windows. For example, to achieve an expected shock rate of 10 shocks/5s, the probability of receiving a shock in a given 0.1s interval was set to 10/50 (i.e. 1/5). By default participants were due to receive a series of 60 stimuli at the maximum expected shock rate of 14 shocks/5s (2.8 shocks/s). Consuming 10mg of relief reduced the expected number of shocks in the immediately following stimulus train by one. Due to the probabilistic nature of the stimuli, this meant that the effectiveness of relief could vary somewhat from one trial to the next. Participants were fully briefed that the pain relief was probabilistic, chosen so as to achieve a more naturalistic context. To assist participants' understanding of this fact, an analogy was drawn to the use of analgesia in real life, whose subjective effectiveness depends on factors other than simply the dose. The maximum allowable consumption of relief on each trial was 120mg, sufficient to reduce the expected shock rate to 2 shocks/5s (0.4 shocks/s), which was referred to as the "Baseline Pain".

Before the experiment, participants underwent a standardised procedure, to control for individual variability in pain perception (see Experiment 1A and General Experimental Methods 5.1.1). Also, prior to entering into the session, participants were given three samples of the maximum (default) and minimum (baseline) shock rates which they could expect to experience with using no relief or using maximum relief respectively. The choice phase was limited to 6 seconds, and each trial lasted 14 seconds in total, the experimental session therefore lasted 14 minutes.

7.2.1.3 Data Processing

To reduce the computational complexity of the modelling and simulation analysis, relief consumption was rounded to the nearest 10mg, creating 13 possible

spending choices on each trial (0 to 12). This procedure produced occasional rounding errors such that the cumulative total rounded consumption exceeded the budget constraint. These errors were corrected by disallowing rounded consumption to exceed the remaining total relief, resulting in fictitious observations on the final trial for some participants. These discrepancies from the true observed consumption profiles were small by comparison to predominant patterns of consumption.

7.2.2 DYNAMIC PROGRAMMING

To simulate consumption paths predicted by the dread-discounting functions derived from one-off choices I implemented a dynamic program (Bellman, 1956; Bertsekas, 1976; Rust, 1989) over all possible states of capital at each time point. A deterministic transition function, $T(c_t, s_t)$ described how actions in the current state mapped to subsequent states, such that:

$$s_{t+1} = s_t - c_t \quad (7.1.1)$$

Where s_t , denotes capital at time t , and c_t consumption at t . Borrowing is not allowed, therefore $s_t \geq 0$ and $c_t \leq s_t$.

Consuming a quantity of relief, c_t , was associated with utility $U(c_t)$ at the current state, where $U(c_t)$ is the utility function for relief. Since $c_t \leq s_t$, by the budget constraint, the function U also depends on current capital s_t . The overall value of consuming relief, c_t , when situated at time, t , with a state of capital, s_t , termed a Q – value, was then described recursively as a function of the resulting relief utility at the current state, $U(c_t, s_t)$, followed by the expected utility of relief at all future states, given a future action policy, π (see Chapter 5, General Methods) and a discount function, $\Delta(d)$, giving rise to:

$$Q^\pi(s_t, c_t) = U(c_t, s_t) + E \left[\sum_{d=1}^{T-t} \Delta(d) \cdot U(c_{t+d}, s_{t+d})_{c_{t+d} \sim \pi} \right] \quad (7.1.2)$$

The action policy, π , dictates the probability of consuming an amount of relief, c_t , given that the agent is currently situated at t and has capital, s_t , here represented by a softmax policy for action selection, such that:

$$p(c_t|s_t, t) = \pi(s_t, t) \quad (7.1.3)$$

Where:

$$\pi : s_t, t \rightarrow \frac{\exp(\beta \cdot Q^\pi(s_t, c_t, t))}{\sum_{c_t} \exp(\beta \cdot Q^\pi(s_t, c_t, t))} \quad (7.1.4)$$

Higher values of the inverse temperature parameter, β , increase the probability that the agent will choose the option with higher utility. $\Delta(d)$ was represented by the anticipation-discounting function derived from one-choices, which assumed the following form:

$$\Delta(d) = \gamma_C^d + \alpha \left[\sum_{\tau=0}^{d-1} \gamma_C^{d-\tau} \cdot \gamma_A^\tau \right] \quad (7.1.5)$$

In this case γ_C represents the discount rate on future consumption, which also determines savouring of future relief, and γ_A represents the discounting on savouring itself. Parameters from the Exponential Dread model fitted to choices in Experiment 1A (γ_D –Framing; both Pain and Relief frames separately), namely α , γ_P and γ_D were carried forward to generate resulting optimal consumption paths on the dynamic experiment. To do so, discounting of relief consumption was set to be equivalent to discounting of pain ($\gamma_C = \gamma_P$) and discounting of dread was set to be equivalent to the discounting of savouring ($\gamma_A = \gamma_D$). Figure 7.2 illustrates typical forms for an anticipation-discounting function for a positively-valenced outcome, where $\gamma_A = 1$. Where anticipation dominates, the overall value is an increasing function of delay. Linear utility function for pain and relief was assumed. Optimal policies were simulated using a high value of the softmax inverse temperature, $\beta = 10000$.

7.2.2.1 Sophistication vs Naivety

The value function $Q^\pi(s_t, c_t)$ expresses the notion that the value of consuming an amount, c , at the current capital level and time period depends on the immediate utility of consuming c plus the expected value of (discounted) future consumption, given accurate knowledge of one's likely future policy for action. This model therefore entails complete sophistication. Put simply, the sophisticated agent knows how they are likely to behave in future and takes this into account. In other words the model assumes for example that the agent at $t = 1$ knows that a future decision-maker at $t = 11$ will apply the same degree of discounting to periods $t = 11, t = 12, t = 13$ and so on as the agent currently applies to periods $t = 1, t = 2, t = 3$ and so on. Naivety by contrast would entail for example that the agent at $t = 1$ assumes that the decision-maker at $t = 11$ will apply the same discount factors to periods $t = 11, t = 12, t = 13$ and so on as the agent currently applies to those time periods. Given dynamic inconsistency in the discounting function, a naïve agent would be expected to change their plans at each time step.

To simulate this, the dynamic program was iterated once for each trial of the simulation, with the following recursive value function:

$$Q_{naive}^\pi(s_t, c_t, t, i) = \Delta(t - i) \cdot U(c_t, s_t) + E \left[\sum_{\tau=t+1}^T \Delta(\tau - i) \cdot U(c_{t+d}, s_{t+d})_{c_{t+d} \sim \pi} \right] \quad (7.1.6)$$

Where each iteration is represented by i , which ranges between 1 and T . Naïve consumption plans at trial i were sampled from a policy based on $Q_{naive}^\pi(s_t, c_t, i)$ over the remaining trials $t = i, i + 1, i + 2 \dots T$. Naïve consumption paths themselves were simulated by sampling from $Q_{naive}^\pi(s_t, c_t, t)$ over all trials $t = 1, 2, 3 \dots T$. In simple terms, the naïve agent continually unravels their previous plans, and must form a new plan at each time step.




Figure 7.2 Anticipation-discounting functions. Anticipation-discounting functions are constructed from a linear combination of the conventionally discounted value of an outcome, i.e. its instantaneous anticipation, and the prospective sum of anticipation whilst waiting for the outcome, displayed here for an outcome with positive utility. **A** Where prospective anticipation (savouring) dominates, the overall value of the outcome decreases as it draws nearer, due to decreasing prospective anticipation. **B** These functions give rise to dynamic inconsistency. Where prospective savouring dominates, preference reverses towards deferral of consumption. Here r_1 is a larger sooner reward and r_2 is a smaller later reward, where both generate a large degree of savouring. When both rewards are distant, the sooner reward is preferred, however as the rewards approach, the delayed reward becomes preferable.

7.2.3 MODEL FITTING PROCEDURES

As for Experiment 1A, model fitting followed a maximum likelihood framework (see Chapter 5, General Methods), using the softmax policy to generate the probability of observing each possible (rounded) level of relief consumption, given a particular set of model parameters. Simplex optimisation was performed using the Matlab (Mathworks, MA, USA) `fminsearch` optimisation tool (Nelder-Mead search algorithm, Lagarias et al., 1998) with the addition of bound constraints by transformation. For each participant 10 iterations of the optimisation were performed, and the maximum likelihood estimate across all iterations was selected. On each iteration the optimiser was called within a random multi-started overlay (RMsearch), with 100 starting points selected from a uniform distribution between the parameter bounds, in order to reduce convergence on local minima.

7.3 RESULTS

7.3.1 SIMULATING CONSUMPTION BEHAVIOUR

To illustrate the effects of changes to the instantaneous utility and anticipation-discounting functions, I used dynamic programming to simulate optimal behaviour on a reduced version of the task lasting 10 time periods (with a budget of 400mg).

7.3.1.1 Effects of the Instantaneous Utility Function

Within the standard economic model, the instantaneous utility function can affect the optimal consumption path, even for a decision-maker who treats the same outcome as equally valuable regardless of its timing (Figure 7.3). For the special case of linear utility, where $U(c_t) = c_t$, provided there is no discounting or interest rate, all possible consumption paths in which total consumption corresponds to spending the entire capital are equally valued. As a result, at each time step, and each state of capital, all possible consumption levels have equal value. The result is that consumption in the first time period, c_1 , could be selected at random from a uniform

distribution, in which case, the expected consumption level, c_1 , is close to 60 units. Single-period consumption, c_t , then continues in this manner until the capital is entirely consumed (The left panel of Figure 7.2A shows an example linear utility function, and the left panel of Figure 7.2B shows two example simulated consumption paths under this utility function with no discounting). A simple explanation for this is that, for the agent with linear utility and no time preference, increases in relief spending in one time period exactly outweigh the inevitable decreases in later time periods – so the agent is indifferent as to when they spend their relief.




Figure 7.3 Effects of changes in the instantaneous utility function on dynamic consumption. **A** Two alternative utility functions, linear (left panel) and concave (right panel). **B** Simulated optimal consumption paths with no discounting or anticipation (zero time preference) under the two forms of utility function. In each case two sample simulated paths are displayed, to illustrate that, with linear utility (left panel) there is more than one optimal path. Left panel: under linear utility with no discounting or anticipation all paths which consume the entire budget are equally valued. Consumption is therefore chosen at random from a uniform distribution until the budget is expended. Right panel: concave utility motivates spreading consumption over time.

With a concave utility function, here illustrated with $U(c_t) = c_t^k$, where $0 < k < 1$, low levels of consumption are relatively more valuable than would be the

case under a linear function. Here, the optimal path is to spread consumption evenly across time (The right panel of Figure 7.3A, shows an example concave utility function, and the right panel of Figure 7.3B shows two sample consumption paths under this utility function with no discounting). This can be expressed as a desire not to be left with little or no relief in any given time period.

7.3.1.2 Effects of Anticipation-Discounting Functions

Figure 7.4 plots predicted consumption paths under four possible parameterisations of the anticipation-discounting function (Figure 7.4A), under both full naivety or full sophistication, for a concave utility function: $U(c_t) = c_t^{0.75}$. It can be appreciated from Figure 7.4B that, where discounting dominates (first column), optimal consumption is decreasing, corresponding to spending relief early. With no discounting (second column), optimal consumption is even over time, owing to concave utility. Where anticipation dominates (third column), the predicted consumption path is increasing, corresponding to saving relief. Where anticipation is itself discounted ($\gamma_A < 1$; fourth column) non-monotonic consumption profiles result. Where there is a degree of savouring ($\alpha > 0$), the consumption paths for naïve and sophisticated consumers diverge, albeit subtly in some cases. The underlying dynamic inconsistency is illustrated in Figure 7.4C, which plots consumption plans made at the first three time periods for fully naïve agents. Rather than consumption itself, these plots depict the naïve plans for future consumption from the current time-period onwards. Where discounting dominates (left column), inconsistency similar to that implied by hyperbolic discounting results: consumption at the next period turns out to be greater than planned. Where savouring dominates (right hand two panels), the naïve decision-maker consumes less than planned. A sophisticated agent takes these future discrepancies into account and adjusts their plan accordingly.

To summarise, normative considerations justify at least three obvious qualitative classes of relief spending – increasing, decreasing and flat (or spreading). They also make strong predictions about the relationship between single- and multi-period decisions, and potentially the effect of degrees of game-theoretic sophistication.




Figure 7.4 Anticipation-discounting and dynamic utility maximisation. **A** Four anticipation-discounting functions. From left to right: predominant discounting, no discounting, predominant savouring, discounted savouring. The parameters of each function are displayed on the plot. **B** Simulated optimal consumption paths under the same four discount functions, with concave utility, $U(c) = c^{0.75}$. Green circles represent simulated consumption paths for a fully naïve decision-maker (See Main Text). Red circles represent consumption for a fully sophisticated decision-maker. **C** Plans for future consumption made in the first three time periods for a naïve decision-maker. The red circles indicate *planned* consumption from the perspective of $t = 1$, the blue circles from the perspective of $t = 2$ and the green circles from the perspective of $t = 3$. Where discounting dominates (left panel), the naïve decision-maker consumes more than planned, where savouring dominates (right hand two panels), the naïve decision-maker consumes less than planned.

7.3.2 OBSERVED CONSUMPTION OF RELIEF

The experimental data consisted of the number of units of relief consumed on each trial by each participant. Figure 7.5A plots the median consumption of relief on each trial at the group level ($N = 30$, bars indicate the interquartile ranges). Across participants, the profile of consumption is increasing over time, showing the tendency

for relief to be saved for towards the end of the experimental session. Robust linear regression on all choices made by all participants ($N = 1980$), using iteratively reweighted least squares with a bi-square weighting function, demonstrated a significantly positive effect of time on relief consumption ($\beta = 0.47$, $p < 0.001$).




Figure 7.5 Observed distribution of relief consumption across time in Experiment 2. **A** Median relief consumption on each trial at the group level is indicated by the solid black circles. Error bars indicate the upper and lower quartiles of consumption. **B** The distribution of group level consumption. The intensity of the grey bars represents the proportion of the 30 participants included in the analysis choosing to consume each amount of relief on a particular trial. Relief is expressed in units, produced by rounding the raw consumption choices to the nearest 10 units of relief. A tendency to conserve relief is evident from below-average consumption over the first 40 trials and above-average consumption over the last 20 trials. A tendency to spread relief across time is evident from the high proportion of choices to spend 4 units of relief, the mean rounded relief per trial over the whole experimental run.

However, the group-level presentation of the data conceals the complexity of subject-specific choices. To examine this I calculated the proportion of participants choosing a particular level of consumption on each trial. The observed distribution of rounded relief-consumption at the group level is displayed in Figure 7.5B. Darker bars indicate a higher proportion of participants choosing a given consumption level on each trial. There were very few choices to consume close to the maximum quota of relief early in the experimental session. Rather, higher intensities corresponding to spending close to zero relief in the first 40 trials, and above-average consumption across the final 20 trials, demonstrated that participants tended to conserve relief for the final portion of the session, which would be consistent with savouring. Since there was a budget of 240 rounded units of relief, to be allocated across 60 trials, even

spreading of relief would entail spending 4 units per trial. Notably, high intensities corresponding to spending close to 4 units of relief indicate that participants also demonstrated a tendency to spread relief across time, which would be consistent with participants having concave utility for relief. There is also a weak tendency to sample the maximum allowable quota of relief throughout the experimental run. An additional interesting feature is that participants were more likely to consume close to the mean relief remaining early in the experiment, tending to switch to consuming zero relief during the middle of the experiment.




Figure 7.6 Consumption paths from six sample participants in Experiment 2. Hollow circles denote the consumption choices on each trial. The bold dashed line denotes the mean relief remaining at the start of each trial. These six participants are selected as representative of the key patterns observed. The first two participants (i and ii) on spread relief over time. The subsequent three participants (iii-v) predominantly conserve relief, as evidenced by an increase in the mean relief remaining per trial over time. Participant vi) consumes above the mean relief remaining towards the start of the session and subsequently adjusts consumption downward.




Figure 7.7 Histograms of relief consumption over time in Experiment 2. A Relief consumption expressed as mg. Each plot represents the distribution of relief consumption over a period of 10 trials. Multiples of 10mg are over-represented. **B** Relief consumption rounded to the nearest 10mg, expressed as ‘units’. Each plot represents the distribution of relief consumption over a period of 10 trials.

At the individual level, participants appeared to display one or more of the above three tendencies, though strikingly, no participant systematically consumed close to the maximum available relief at the outset of the experiment. To illustrate this, consumption profiles from six sample participants are displayed in Figure 7.6, overlaid with the mean relief remaining per trial (dashed lines), termed ρ_t . This quantity (displayed to participants on-screen before each choice) is given by the total remaining relief on that trial, s_t , divided by the number of trials remaining:

$$\rho_t = s_t/[60 - (t - 1)]$$

For any trial during the experiment, consuming exactly ρ_t units of relief on every remaining trial would entail even consumption of relief over the remainder of the experiment.

Figure 7.7A plots raw consumption choices (in mg) as a series of histograms over time, illustrating that multiples of 10mg are over-represented. This suggests that participants used strategies to reduce the dimensionality of the task, rather than performing optimisation at the native resolution. When rounded consumption (in units) is also plotted in this manner (Figure 7.7B), choices to consume zero relief or 4 units of relief are prominent.

7.3.3 PREDICTING CONSUMPTION FROM ONE-OFF CHOICES BETWEEN DELAYED PAINS

Although the participants exhibited the same qualitative patterns of behaviour as expected from the normative accounts (Figure 7.4), this does not mean that each participant's own choices were consistent with their one-off preferences. To compare one-off and dynamic behaviour, I first derived summary measures of behaviour on both tasks. In the one-off choice task, the frequency of choosing sooner pain is a correlate of dread. As described previously, in Experiment 1A one-off choices between delayed pains were elicited under two descriptive 'frames', a *pain* frame, and a *relief* frame. The latter corresponds to that used in the relief consumption experiment. Nevertheless I examined the relationships between dynamic relief consumption behaviour and sooner choice frequency on both frames. The signed slope of the dynamic consumption path (fitted with least-squares linear regression) is

a measure of the overall tendency to conserve relief, while the absolute magnitude of the slope is a measure of the deviation, in either direction, from even spreading of relief.

Contrary to a normative account, there was no significant positive relationship between the tendency to dread (on either frame) and the slope of the consumption path (Figure 7.8A; $p > 0.25$, $N = 30$), although there was a trend in this direction for the relief frame choices (Pearson $r = 0.2$). Neither was there a significant relationship between dread and the tendency to spread relief over time (Figure 7.8B; $p > 0.25$, $N = 30$).




Figure 7.8 Relationships between one-off (binary) and dynamic intertemporal choices. The frequency of choosing sooner pain in the binary intertemporal choice experiment which provides a summary behavioural measure of dread, in both pain (i) and relief (ii) frames (see main text), is plotted against: **A** the slope of the spending profile in the dynamic consumption task (positive slope indicates saving relief) and **B** the absolute slope of the spending profile (a measure of deviation from even spreading of relief). Solid lines indicate a linear least-squares fit through the data. There are no significant relationships between the behavioural metrics on the two tasks.

It is possible that variability in the utility function and softmax temperature parameters could account for some of the observed differences between one-off and dynamic choices. To explore this I implemented a model in which the softmax β , and the exponent governing the utility function, k , were fitted freely, whilst holding the previously derived anticipation-discounting parameters constant (for those participants for whom estimates of anticipation-discounting functions were available from one-off choices, $N = 23$, see Methods). The group-level distribution of consumption predicted by this model is displayed in Figure 7.9B, (Figure 7.9Bi, pain frame preferences; Figure 7.9Bii, relief frame preferences), alongside the observed group level distribution in Figure 7.9A. Although the optimal preferences predict saving of relief at the group level, they underestimate the tendency to spread relief over time, even allowing for concave utility, and the fitted policies are relatively imprecise. To estimate the proportion of variance in the observed data accounted for by the model, I found the mean consumption level for each participant across each 10 trials of the experiment, before calculating the same measure by simulating 10000 consumption paths resulting from the maximum likelihood parameterisation of the model. As shown in Figures 7.9Ci (pain frame preferences) and 7.9Cii (relief frame preferences), there was a significant positive relationship between predicted and observed consumption paths (robust regression, pain frame: $\beta_1 = 0.22, p < 0.001$; relief frame: $\beta_1 = 0.44, p < 0.001$). However least squares fits indicated that the model accounted for only a relatively small proportion of the observed variance (pain frame $R^2 = 0.03$, relief frame $R^2 = 0.07$).




Figure 7.9. Fits of an anticipation-discounting model with variable utility and choice randomness to consumption choices in Experiment 2. **A** Observed consumption by participants for whom anticipation-discounting functions derived one-off choice tasks were available ($N = 23$). Warmer colours indicate that a higher proportion of participants chose to consume that amount of relief on a given trial. **B** Consumption predicted by the anticipation-discounting functions. These plots denote the mean predicted probability across participants of consuming an amount on each trial given the total remaining relief (at the maximum likelihood parameters of each model). i) Anticipation-discounting functions derived from one-off pain frame choices, with softmax temperature and utility parameters freely fitted. ii) Anticipation-discounting derived from one-off relief frame choices, with softmax temperature and utility parameters freely fitted. **C** Mean predicted consumption levels simulated from the maximum likelihood parameterisations of each model over each 10 trials of the experiment for each participant are plotted against the same metric derived from the observed data.

7.3.4 MODELLING RELIEF CONSUMPTION USING HEURISTICS

Given that consumption behaviour showed only weak correspondence with the predictions of anticipation-discounting as derived from one-off choices, I tested alternative generative accounts. This analysis was performed *post hoc*, and focussed on characterising simple computations that might feasibly have produced the observed consumption choices.

The first model, termed the *Direct Action Heuristic* model, proposed that participants implemented the three observed behavioural tendencies directly as heuristics, with some propensity to choose between them. The three are termed *spend-now-suffer-later* (with propensity M_{spend}), *spread-spending* (with propensity M_{spread}), and *save-now-spend-later* (with propensity M_{save}). The extent to which observed relief consumption, c_t , fell below the mean relief remaining on each trial, ρ_t , is given by:

$$d_t = \rho_t - c_t$$

Positive d_t entails using less than the mean relief remaining per trial, while $|d_t|$ indicates the extent of deviation from spreading. Formally, the three heuristics were defined as::

$$M_{\text{spend}}(s_t, c_t, t) = c_t$$

$$M_{\text{spread}}(s_t, c_t, t) = -|d_t|$$

$$M_{\text{save}}(s_t, c_t, t) = \begin{cases} d_t & \text{if } \rho_t < 12 \\ 0 & \text{otherwise} \end{cases}$$

M_{spend} formalises a *spend-now-suffer-later* heuristic, by assuming linear utility for relief consumption, and thus a propensity to consume the maximum allowable relief. M_{spread} formalises a *spread-spending* heuristic, by penalizing deviations from the mean relief remaining, and so generates a propensity to spread relief over time. M_{save} formalises a *save-now-spend-later* heuristic, by assigning higher value to consuming

less relief, provided that the mean remaining relief per trial is less than the maximum possible consumption level. M_{save} therefore generates a propensity to consume as little relief as possible until there is sufficient remaining relief to reduce pain to the baseline level for the remainder of the experiment, at which point the remaining heuristics encourage spending this quantity.

The three action propensities were implemented as separate policies, each with a unique softmax inverse temperature parameter. The final probability of consuming each level of relief was assumed to arise from a weighted average across these policies with weight for a policy determined by its inverse variance, such that:

$$\omega_i = \left[\frac{1}{\text{var}(\pi_i)} \right] / \sum_{j=1}^3 \left[\frac{1}{\text{var}(\pi_j)} \right]$$

where ω_i is the weighting on policy, π_i . This procedure ensured that the *Direct Action* model had only three parameters: the inverse temperatures of each softmax function, respectively termed β_{spend} , β_{spread} and β_{save} . In each case higher β_i had the effect of increasing the weighting on that particular policy, by reducing the variance of π_i .

The group-level distribution of observed consumption choices is reproduced in Figure 7.10A, for comparison with the model fits. The distribution predicted by the *Direct Action* heuristic model is displayed in the left-hand panel of Figure 7.10B. The model provided a parsimonious summary of observed consumption choices, albeit not convincingly capturing the observation that some participants were more likely to consume close to the mean relief remaining per trial (ρ_t) near the start of the experimental run, before switching to conserve relief.




Figure 7.10. Heuristic model fits to consumption choices in Experiment 2 **A** Observed consumption by all 30 participants. Warmer colours indicate that a higher proportion of participants consumed that amount on a particular trial. Black arrows indicate spending zero relief, which becomes more prominent during the middle of the experiment. **B** Relief consumption predicted by alternative heuristic models. These plots show the mean probability across all participants of consuming an amount on each trial given the total remaining relief (at the maximum likelihood parameters of each model). The Direct Action model combines the three observed behavioural tendencies as heuristics: save-now-spend-later, spread-spending or spend-now-suffer-later. The Income Maximisation model implements the saving tendency as an attempt to maximise the mean remaining relief per trial, over a limited future horizon. This model captures the relatively greater tendency to save relief during the middle of the experiment (black arrows). **C** Mean predicted consumption levels simulated from the maximum likelihood parameterisations of each model over each 10 trials of the experiment for each participant are plotted against the same metric derived from the observed data.

The above pattern might have several different explanations. One simplification in the model is that the explicit relative weightings of the heuristics are assumed to be constant. However, participants may have adopted the *spread-spending* heuristic at the outset, before learning the extent that they were able to tolerate pain as the experiment progressed then switching to *save-now-spend-later*. Similarly they may have consumed the mean relief at the outset as a default option, until they learned to trust the experimental setup. A further possibility is that participants, rather than using a *save-now-spend-later* heuristic directly as defined above, may have sought to maximise the mean relief remaining per trial (ρ_t) over the near future: since saving relief would have more immediate effect on ρ_t later in the experiment compared with at the start, the propensity to save would be expected to increase as the experiment continued.

The data do not admit a direct distinction between the above hypotheses. However in order to illustrate one of the possibilities I fitted a modified version of the above model in which the *save-now-spend-later* heuristic described above is replaced with a heuristic to maximise ρ_t over a limited future horizon, which I term an *income maximisation* heuristic (and eponymous model). Thus M_{save} in this model was replaced by a action-value function, which described the value of consuming an amount, c_t , at the current capital level, s_t , and time period, t , given knowledge of the future policy for action, π . In other words this model assumed that participants were in part attempting to maximise the expected mean relief remaining per trial, akin to maximizing their expected income:

$$Q_{\rho-\max}^{\pi}(s_t, c_t, t, \gamma) = \rho_t(s_t, t) + E \left[\sum_{d=1}^{T-t} \gamma^d \cdot \rho_{t+d}(s_{t+d}, t+d)_{c_{t+d} \sim \pi} \right]$$

To account for limited computational resources, the above formula incorporates a probability, $1 - \gamma$, that the decision-maker terminates their search at every level deeper into the tree (the γ parameter is mathematically equivalent to an exponential discount rate). I fitted this part of the model using dynamic programming. The remaining two action propensities, M_{spend} and M_{spread} were implemented in the same manner as the *Direct Action* model, and policies were combined using the same weighting method.

The distribution of consumption at the group level predicted by the *Income Maximisation* model is shown in the right-hand panel of Figure 7.10B. It can be seen that this model accounts for the tendency to save relief being higher during the middle part of the experiment. As expected, the *Income Maximisation* model produced an improvement in BIC of 78 at the group level over the *Direct Action* model. (Notably however the *Income Maximisation* model was optimised *post hoc* to account for a particular feature of the observed data). Least squares fits indicate $R^2 = 0.56$ for the *Direct Action* model and 0.80 for the *Income Maximisation* model (Figure 7.10C).

To illustrate the contribution of each of the three heuristics, the policy weightings of the two heuristic models are displayed in Figure 7.11. The *Income Maximisation* model results in a larger relative weight being placed on saving during the middle half of the experiment (Figure 11A). Also throughout the experiment saving (*save-now-spend-later* and *income maximisation*) and *spread-spending* receive considerably higher weightings (Figures 7.11A and 7.11B) than *spend-now-suffer-later*.




Figure 7.10 Policy weightings of heuristic models in Experiment 2 **A** Mean weighting across participants on each of the three policies (see Methods) on each trial under the maximum likelihood fits of the Direct Action (left) Income Maximisation (right) models. Spend-now-suffer-later has low weighting early in the experiment. Spread-spending has high weightings throughout. For the income maximisation model, saving is weighted most highly in the middle part of the experiment. **B** Box plot showing distribution over participants of policy weightings, averaged over all trials of the experiment. Saving and Spread-Spending heuristics dominate.

7.4 DISCUSSION

People are often required to choose between improving their immediate or future wellbeing. In reality such choices are often made sequentially, such that today's choices influence the possibilities available tomorrow for example, when distributing a limited budget, early spending limits the capital available in later periods. Economic theory proposes that budgeting choices ought to be predictable from choices between equivalent one-off delayed outcomes (Rust, 1989; Strotz et al., 1957), however the theory faces considerable theoretical challenges. Experiment 2 tested this in an experiment analogous to the use of personal health budget, in a group

of participants whose intertemporal preferences for one-off future pains of the same nature had been elicited previously.

Interestingly, no participant systematically consumed close to the maximum available relief at the outset of the experiment, as conventional temporal discounting would predict. Whilst two out of the thirty participants analysed did generate declining profiles of relief, these two participants also showed trial-to-trial variability in consumption, suggesting that they may have chosen consumption levels largely at random (with the decline resulting from exhausting the budget). Tendencies to consume the minimum allowable relief early on, thus saving for the end, and to consume close to the mean remaining relief were prominent, with several participants alternating between these two tendencies. Consistent with ‘round number’ heuristics in retirement-savings decisions (Benartzi & Thaler, 2007), choices to spend multiples of 10mg of relief were over-represented in the data.

Perhaps surprisingly there was no significant correlation between a preference for sooner pain in one-off choices and the tendency to save relief in the dynamic task (although there was a trend towards a positive relationship). This null finding appears to cast some doubt on the idea that dread-discounting might be capable of predicting pain-related choices in the field, and appears particularly surprising given a near exact correspondence between the two sets of choices in the timing and nature of the outcomes. The data suggest either that additional processes are at play when choices are made sequentially, that are not well specified by the existing analysis of one-off choices (Frederick & Loewenstein, 2008), or that faced with the complexity of budgeting choices, people choose relatively arbitrary heuristics.

A factor that is poorly specified in the previous analysis of one-off choices is the instantaneous utility function for pain/relief. This is important, since the form of utility function chosen when estimating dread-discounting will influence the resulting estimates. Also, as shown here through simulation, the instantaneous utility function has marked effects on budgeting choices. This suggests that researchers interested in predicting sequential choice might do well to estimate the instantaneous utility function independently from time preferences.

A further effect that is poorly specified in one-off choices is the degree of sophistication – i.e. whether people can predict their own tendency to renege on their plans. As demonstrated here (as well as in existing studies), the anticipation-

discounting functions described previously for one-off choices predict a novel form of inconsistent choice, distinct from that of hyperbolic discounting, which entails the perpetual deferral of consumption against one's long term interests (Figure 7.2; Loewenstein, 1987). Notably such inconsistent behaviour is the reverse of that conventionally associated with 'failure of self control', and is referred to by Loewenstein (1987) as 'reverse time inconsistency'. It is unclear whether such behaviour is manifest in real-time, or indeed influences the kind of consumption choices demonstrated here.

Finally, relative valuation processes might be expected to have played a greater role in the sequential task. Relative valuation processes involve comparison of outcomes against an assumed baseline, or reference-point (Elliott, Agnew, & Deakin, 2008; Kahneman & Tversky, 1979). For example, due to psychological adaptation to the current pain level, a moderate intensity pain can appear more severe when following a low intensity pain than when following a high intensity pain (Vlaev et al., 2009; Winston et al., 2014). Thus if consumption levels are compared with those that precede them, people might be led to choose deliberate privation in order to increase the hedonic impact of subsequent consumption, in this case predicting saving of relief (Dayan, 2009; Elster & Loewenstein, 1992; Frederick & Loewenstein, 2008). The opposite effect may also occur, namely sensitisation to repeated high-level pain, leading participants to occasionally consume the maximum relief as 'respite'. A further possibility is that decreases in consumption from one time period to the next are valued as more negative than equivalent increases are valued positively, i.e. loss aversion (Kahneman et al., 1991; Tom, Fox, Trepel, & Poldrack, 2007; Usher & McClelland, 2004). Loss aversion would be expected to further penalise deviations from either even spreading or saving, for the reason that any increases in consumption above even spreading inevitably lead to future decreases (Loewenstein & Prelec, 1993).

In summary this plurality of possible mechanisms contributing to dynamic behaviour (summarised in Table 7.1), in combination with the challenging complexity of dynamic choices, might explain the low correlation between one-off and dynamic tasks. Perhaps disappointingly, the current data do not permit firm conclusions to be drawn regarding the valuation processes underlying the observed behaviour, however they do suggest a wealth of interesting directions for future investigation. The current

findings also suggest that researchers interested in predicting real-world health-related behaviour might do well to consider a broader range of metrics, besides measures of time preference alone.

Here, I conceptualised deviations from optimality in terms of heuristics, rule-of-thumb strategies designed to ease computational demands (Gigerenzer, 2008; Gigerenzer et al., 2011). I generated putative heuristics *post hoc*, in light of the three observed behavioural tendencies, finding that consumption behaviour was well-described by a combination of three corresponding simple rules, namely save-now-spend-later, spread-spending and, to a much lesser extent, spend-now-suffer-later. Since the heuristic models were developed *post hoc*, they require independent validation in related experimental contexts to establish their generalizability. However, in their support, preferences for spreading rewards evenly across time and for improvement over time are evident in choices between predetermined sequences of outcomes, including wages (Loewenstein & Sicherman, 1991), health (Chapman, 1996a, 2000) and other desirable or undesirable events such as dining at a favorite restaurant or scheduling a visit from a troublesome relative (Loewenstein & Prelec, 1993). Loewenstein and Prelec (1993) propose a model for classifying these preferences, which resembles the Direct Action heuristic model used here, albeit not in the context of whole sequences of choices over time, as here. These data would suggest that people might be likely to show similar tendencies when allocating a health-related budget over time, and also that individuals might differ considerably in their preferred budget allocations over time. From a policy perspective, such individual differences will be interesting to examine as more data on the use of personal health budgets emerge (Forder et al., 2012).

Finally, the behavioural tendencies identified here might be relevant to how patients use pain-relieving medication in clinical practice. Under-use of analgesia relative to clinicians' recommendations is prevalent in both acute and chronic pain (Broekmans, Dobbels, Milisen, Morlion, & Vanderschueren, 2009; Jacobsen, Møldrup, Christrup, & Sjøgren, 2009; Yates et al., 2002), and is a barrier to achieving effective pain relief (Broekmans et al., 2009; Chang, Wray, Sessanna, & Peng, 2011; Gunnarsdottir, Serlin, & Ward, 2005; Miaskowski et al., 2001; Yates et al., 2002). Fears of addiction, tolerance and side effects are frequently cited attitudinal barriers to the use of analgesics (Levin, Cleeland, & Dar, 1985; Lewis, Combs, & Trafton, 2010;

Ward et al., 1993). However, patients also commonly report saving pain relief, to use only in the event that pain becomes intolerable. For example Yates et al. (2002) examined attitudinal barriers to effective pain management in hospitalised patients with cancer, finding that almost one-third (32.5%) agreed that pain-relieving medication should be saved in case pain gets worse. Whether underuse of analgesia might be predicted from a tendency to save relief in simulated relief-scheduling tasks such as this one forms an interesting subject for future enquiry.

	Valuation Mechanism	
	<i>Relative</i>	<i>Absolute</i>
<i>Save-Now-Spend-Later</i>	Adaptation	Anticipation Risk aversion
<i>Spend-Now-Suffer-Later</i>	Sensitisation	Temporal discounting
<i>Spread-Spending</i>	Loss aversion	Risk aversion (concave utility)

Table 7.1. Putative Mechanisms Underlying *Save-Now-Spend-Later*, *Spend-Now-Suffer-Later* and *Spread-Spending* Heuristics. The psychological processes motivating the choice of heuristics might be classified as both relative (reference-dependent) and absolute valuation mechanisms. Relative valuation mechanisms include adaptation to current consumption levels, sensitisation to repeated punishment and loss aversion. Absolute valuation processes include anticipatory utility, temporal discounting and risk aversion.

In summary this chapter examines how people allocate resources for the mitigation of pain, showing that behaviour is not consistent with conventional economic models of intertemporal preference, but is consistent with a simple set of heuristics that encapsulates saving in the present to spend in the future, spreading consumption out evenly over time and (less prominently) spending in the present at the expense of the future. Applied measures of choice over time have tended to focus exclusively on one-off choice paradigms (e.g. Bickel et al., 2012; Bickel et al., 2014), and the modelling of dynamic decision-making tasks suggests a novel and quantitatively rich behavioural predictor.

SOCIAL PREFERENCE FOR PAIN

Primum est ut non nocere
(First, do no harm)

-arr. Thomas Sydenham

The work described in previous chapters focussed on how people make decisions regarding their own future pain or illness, and sought to validate models of such decisions, with a view to predicting of an individual's engagement in unpleasant preventive health behaviour. Here I test whether decisions regarding *other peoples'* future pain systematically differ from decisions regarding one's own future pain. Healthcare professionals and policy-makers are by definition required to make decisions regarding other peoples' health, and such putative preference gaps might thereby have implications for shared decision-making in clinical settings. Furthermore, existing studies have shown that social and temporal preferences for reward are correlated across individuals, with more patient individuals also tending to be more altruistic (Curry, Price, & Price, 2008; Rachlin & Jones, 2008). Were a similar relationship to hold for aversive outcomes also, measuring an individual's preferences for others' pain might even contribute to predicting their own health behaviour.

8.1 SOCIAL DISCOUNTING

A reader of previous chapters of this thesis needs no reminder of the fact that humans and animals often accept smaller immediate rewards over larger delayed rewards (Ainslie, 1974, 1975; Kalenscher & Pennartz, 2008), indicating that the subjective value of reward decreases with delay, termed temporal discounting

(Frederick et al., 2002; Samuelson, 1937). In a related sense, the subjective value of giving reward to others decreases with social distance, termed social discounting (Jones & Rachlin, 2006; Jones & Rachlin, 2009; Rachlin & Jones, 2008). Social distance can be formalised by asking people to create an imaginary rank ordering of the people closest to them, starting with their dearest friend or relative (Jones & Rachlin, 2006). Remarkably, both temporal and social (Jones & Rachlin, 2006; Jones & Rachlin, 2009; Rachlin & Jones, 2008) discount functions for reward are well approximated by a decreasing hyperbolic function (Ainslie, 1974, 1975; Chung & Herrnstein, 1967; Green et al., 1994; Mazur, 1987; Myerson & Green, 1995), in its simplest form given by:

$$U(x, d) = \frac{x}{1 + Kd} \quad (8.1.1)$$

For the case of temporal discounting $U(x, d)$ denotes the subjective utility of a reward with instantaneous utility, x , received after delay, d . In the case of social discounting d denotes the social distance of the recipient from the decision-maker. In each case, the rate at which value declines with increasing d is governed by K , a subject specific hyperbolic discount rate.

Interestingly, temporal and social discount rates for money are correlated, with more patient individuals tending also to be more altruistic (Curry, Price, & Price, 2008; Rachlin & Jones, 2008), and also cooperating more in economic games such as the Prisoner's Dilemma (Harris & Madden, 2002), suggesting that temporal and social valuations of reward engage similar computational processes. This observation dovetails with a theory that choosing delayed outcomes can be seen as analogous to choosing for a future self (Ainslie, 2001; Bartels & Rips, 2010; Ersner-Hershfield, Wimmer, & Knutson, 2009; Jones & Rachlin, 2006; Jones & Rachlin, 2009; Parfit, 1971; Thaler & Shefrin, 1981; Pronin, Olivola & Kennedy, 2008). Indeed, interpersonal choice can also be couched in intertemporal terms, such that choosing to cooperate with another individual carries deferred benefits if the other returns the kindness at a later date. However, the overwhelming majority of studies in this area have examined rewarding outcomes. How these two forms of discounting operate and interact for painful events is less well understood.

Positive social discounting for pain would imply that people ought to find the prospect of causing pain in themselves more aversive than the prospect of causing pain in others. However, in experimental settings people often choose to take on pain to relieve the pain of others (Batson, O'Quin, Fultz, Vanderplas, & Isen, 1983; Hein et al., 2011). Furthermore a recent study has found that people will pay more money to relieve others' pain than they will pay to relieve their own pain of equivalent subjective severity (Crockett et al., 2014; for a related methodology of paying to relieve others' pain see Feldman-Hall et al., 2012). This 'hyperaltruistic' behaviour implies a greater aversion to harming others than to harming oneself, and might be branded 'negative social discounting of pain'. However, no previous studies have examined how this tendency depends on the social proximity of the other person. In other words, the generic form of social discounting for pain has been hitherto unknown. A plausible prior prediction is that people exhibit negative social discounting for close others (i.e. will take on pain to relieve even a milder pain for someone close to them), but that this effect diminishes with increasing social distance. I refer to this behavioural hypothesis as 'discounted harm aversion', for the reason that it implies an aversion to harming others that is positively discounted with social distance.

8.2 DREADING THE PAIN OF OTHERS?

A further question is whether and how people incorporate temporal discounting and/or anticipatory utility when choosing the timing of events on behalf of others. This question is relevant to situations in which healthcare professionals must advise patients on the timing of medical or surgical procedures, or control how long patients must wait for potentially painful procedures. In previous chapters of this thesis I have described how people frequently choose to expedite, rather than defer, inevitable pain (Badia et al., 1966; Berns et al., 2006; Cook & Barnes, 1964; Loewenstein, 1987), termed negative temporal discounting (Loewenstein & Prelec, 1991; Van Der Pol & Cairns, 2000), with an explanation being that people find waiting for pain intrinsically aversive, aptly termed 'dread' (Berns et al., 2006; Loewenstein, 1987). However no previous studies have formally examined whether people take dread into account when evaluating others' pain.

People are known to show empathic responses to others' pain (Batson et al., 1988; Batson et al., 1983; Jackson, Meltzoff, & Decety, 2005; Lloyd, Di Pellegrino, & Roberts, 2004; Singer et al., 2004), suggesting that people might empathise with others' dread. An illustrative real-world scenario would be that of a nurse or doctor responsible for overseeing a painful procedure, such as taking a blood sample. A nurse who is empathic to the patient's dread might attempt to reduce the salient waiting time before the procedure, for example by preparing their equipment before bringing the patient into the clinic room. However, given heterogeneity in preferences for delayed pain (seen in Chapter 6 of this thesis), people are also likely to be uncertain about others' tendency to dread pain. Faced with this uncertainty, people might be expected to place more emphasis on the intensity of pain itself (a more reliable indicator of aversiveness). This hypothesis would predict diminished dread, and thus diminished emphasis placed on waiting time, when choosing on behalf of others. I term this the 'diminished social dread' hypothesis.

To test these two hypotheses I implemented three experiments, (Experiments 3A, 3B and 3C) in which participants were asked to decide between different painful outcomes, whilst varying the subjective intensity, timing and recipient of the pain. Experiment 3A was implemented in the laboratory, with outcomes of moderately painful brief cutaneous electric shock stimuli of different intensities across different stated delays, of up to 29 seconds (there was a fixed inter-trial interval of 30 seconds, such that choosing sooner shocks did not hasten the end of the experiment). In condition one (presented in a counterbalanced order), termed *self-now-self-later*, all participants ($N=60$ complete datasets) made binary choices between immediate and delayed shocks for themselves. In condition two, termed *other-now-other-later*, Deciders ($N=27$ complete datasets) chose between immediate and delayed shocks on behalf of another participant, the Receiver. In conditions three and four, respectively termed *self-now-other-later* and *other-now-self-later*, Deciders chose between immediate shocks for themselves and delayed shocks for the other participant, or *vice versa*. This design permitted measurement dread for one's own delayed pain, and that of others, as well as allowing an estimate of altruism. Figure 8.1 illustrates the experimental protocol for Experiment 1.

Experiments 3B and 3C ($N=67$ and 70 respectively) were carried out online, with outcomes of a hypothetical painful medical treatment. Experiment 3B consisted

of a replication of the first two conditions of Experiment 3A, and a set of social discounting questions, in which participants chose between immediate pain for themselves and another person at varying degrees of social separation from themselves. Finally, to test for an interaction between delay and social distance, in Experiment 3C I also replicated the *self-now-self-later* and *other-now-other-later* conditions, whilst manipulating whether the other person was either a close friend/relative, or a more distant friend/acquaintance.




Figure 8.1 Protocol for Experiment 3A. **A** In each of four possible conditions (order counterbalanced across participants) the Decider was asked to choose between two possible combinations of intensity, delay and recipient of shock. **B** Choices were selected to be realised with a 1/10 probability, in which event participants saw a yellow warning screen followed by a countdown, **C**, to shock delivery, **D**. All participants performed *self-now-self-later* choices ($N=60$). Only Deciders ($N=27$ included) performed the remaining conditions. For these conditions Deciders were informed that, in the event that an outcome was chosen to count for real, the Receiver would see the Decider's choice, followed by an identical countdown to the shock outcome.

8.3 EXPERIMENT 3A: PREFERENCES FOR OTHERS' DELAYED PAIN

8.3.1 METHODS

8.3.1.1 Participants

63 healthy participants (23 males 40 females; mean age 23.6 years, sd =4.6 years) were recruited from the UCL Institute of Cognitive Neuroscience subject database. The experiments took place at the Wellcome Trust Centre for Neuroimaging, UCL. Each experimental session lasted for 2 hours and participants were compensated at a rate of £10 per hour. Participants were recruited in pairs. The two participants in each pair did not meet each other, but were informed that they would be interacting through the computer via an intranet link. The reason for this was to maintain anonymity and to ensure that choices were not influenced by characteristics of the other participant, such as their age or gender.

8.3.1.2 Briefing and Role Allocation

Before the experiment, the participants were randomly allocated to the role of either 'Decider' or 'Receiver', by a method designed to reassure participants that no deception was involved (Crockett et al., 2014), whilst maintaining their anonymity. While hidden from each other, and remaining silent, both participants were simultaneously asked to reach into an opaque box and take one of two plastic balls concealed therein. The balls were labeled 'D' and 'R', designating their respective roles of Decider and Receiver. Once allocated, the roles did not reverse, and participants were given information on this fact in advance. All participants were fully briefed before the study that the experimental setup did not involve deception. All participants gave full informed consent before taking part in the study, and were free to withdraw their consent at any time.

8.3.1.3 Participant Flow and Exclusions

Of the 63 participants recruited, 58 (29 pairs) completed all parts of the study. In the remaining five cases, either the second participant in the pair did not arrive or time constraints forced the experiment to end before all blocks had been completed. Thus, these five participants completed the *self-now-self-later* block alone. Data from four participants were lost due to experimenter error. As a result, choice data from all four conditions were available from 27 participants designated as Deciders, while *self-now-self-later* choices were available from 60 participants (the decrease from the 60 participants recruited to 27 Deciders is largely due to the fact that they performed the task in pairs, with one Decider in each pair).

8.3.1.3 Pain Stimuli and Thresholding

Throughout the experiment each participant sat in front of a computer screen. The package COGENT 2000 (University College London) was used for presentation of choices and response acquisition. Cutaneous electrical stimuli were delivered through two silver chloride surface electrodes placed approximately 3cm apart on the dorsum of the hand, 1cm distal to the wrist, via a DS5 Digitimer (Letchworth Garden City, London) constant current stimulator. A single ‘shock’ was composed of five 10 millisecond square-wave pulses at 49 millisecond intervals. After providing consent, participants underwent a standardised thresholding procedure, which allowed selection of shock intensities corresponding to an equivalent subjective level of discomfort on a 10-point Visual Analogue Scale (VAS) for each participant (see General Methods).

8.3.1.3 Passive Trials and Physiological Measurements

After thresholding participants were first presented with a sequence of 10 ‘passive trials’ in which the outcome was a shock of 7/10 intensity on the VAS. These trials gave participants an opportunity to become familiar with the delays and the experimental setup. In addition, participants were fully briefed on the nature of the choices and their understanding was checked verbally. Delays of 0, 2, 4, 8, 16 and 29

seconds were used for the main experiment. The 0s delay was not sampled for the passive trials; all other delays were sampled once in a randomised order. In each of two blocks, in randomised order, shocks were delivered to either the Decider or the Receiver. Each passive trial commenced with a screen detailing the intensity of the shock, recipient and delay, for example: “7/10 shock for the other participant, delay 8 secs”. After this screen a countdown timer began, displayed on the screen as a pie chart, with the segment of time remaining decreasing each second, up until the time of the shock outcome (see Figure 8.1). Participants were informed that the other participant would see an identical countdown timer.

During passive trials, the skin conductance response (SCR) and heart rate of the Decider were recorded. Skin conductance was recorded from two silver chloride surface electrodes attached to the tips of the first and middle fingers, on the same hand to which the shock stimuli were being delivered. Heart rate was recorded from a pulse oximeter attached to the ring finger of the same hand.

Conductance recordings were received by a constant voltage coupler DC-powered by a 12 V battery, the output of which was converted into an optical pulse frequency and transmitted using fiber optics to a PC running the data acquisition package Spike2 (Cambridge Electronic Design, Cambridge). The latter PC also received pulses of 10ms length via the parallel port of the stimulus PC, marking the timing of the two events of interest, namely the start of the countdown timer and administration of shock. The analysis software, SCRalyze (SCRalyze b2.1.8, scralyze.sourceforge.net), run from within Matlab (Mathworks, MA, USA), was used to analyze physiological responses. Prior to analysis, conductance data were converted back to a waveform signal with 100 Hz time resolution, bandpass filtered using a first order Butterworth filter (cut-off frequencies of 5 Hz, and 0.0159 Hz; corresponding to a time constant of 10 s), and down-sampled to 10 Hz sampling rate. The time-series was then z-transformed to account for between-participants variance in SCR amplitude.

To analyze skin conductance responses I generated a General Linear Convolution Model (GLM) with regressors corresponding to the pre-shock delay period and shock onset. These were convolved with a canonical skin conductance response (no temporal or dispersion derivatives) for analysis at the within-subject and group levels (Bach, Flandin, Friston, & Dolan, 2010). For each event responses of

Deciders to shocks for themselves were distinguished from responses to shocks for the other participant, creating four events in total: *anticipation-self*, *anticipation-other*, *shock-self* and *shock-other*. Events were compared with the mean baseline SCR signal and with each other using one-sample and paired two-sample *t*-tests respectively.

8.3.1.4 Choice Trials

Immediately after the passive trials the main experiment commenced (Figure 8.1). Four experimental conditions were presented block-wise in a counterbalanced order. In the first, termed *self-now-self-later*, participants chose between immediate and delayed shocks for themselves. Both Deciders and Receivers completed these choices. The remaining three conditions were completed by Deciders only. In the second, termed *other-now-other-later*, the Decider chose between immediate and delayed shocks for the Receiver. In the third, termed *other-now-self-later*, the Decider chose between immediate shocks for themselves and delayed shocks for the Receiver. Finally in the fourth, *self-now-other-later*, condition the Decider chose between immediate shocks for the Receiver and delayed shocks for themselves. There was no opportunity for reciprocity and participants were informed of this fact at the outset.

Within each condition, choices followed a symmetrical adjusting-sooner-amount and adjusting-later-amount design (see Appendix, Table A2 for the full set of choices). For each choice there were two options: an immediate and a delayed pain. In adjusting-sooner pain choices the delayed shock was always of intensity 5/10, while the intensity of the immediate shock varied from one choice to the next, between 1/10 and 9/10. In adjusting-later pain choices, the immediate shock was always of intensity 5/10, while the intensity of the delayed shock varied from one choice to the next, between 1/10 and 9/10. Choices in which the immediate shock was the larger shock were designed to assess for negative temporal discounting, while choices in which the delayed shock was the larger were designed to assess for positive temporal discounting. The symmetrical procedure ensured that the number of choices in which the delayed shock was the higher intensity was equal to the number in which the immediate shock was the higher intensity. The *self-now-other-later*, *other-now-self-later* conditions also incorporated otherwise identical choices in which both

shock options were immediate (0s), allowing an estimate of the intensity of pain which the Decider was willing to endure to avoid a given intensity of pain for the Receiver, to be calculated independently from the effect of delay.

After each choice, the chosen option was highlighted. In the social conditions the Receiver was shown the options presented to the Decider and their subsequent selection. In each condition, each choice had a 0.1 probability of being realised. In this event a yellow warning screen was first displayed to both participants for 2 seconds, after which the countdown timer appeared, followed by the occurrence of a shock at the relevant delay. This form of ‘incentive-compatible’ design in which selected choices count for real is widespread in studies of reward-guided choice. The timing of outcomes for the two participants was synchronised via an intranet link between the two stimulus PCs. Deciders were informed that, in the event that an outcome was chosen to count for real, the Receiver would see the Decider’s choice, followed by an identical countdown to the shock outcome.

8.3.1.5 Model Fitting Procedures

To fit alternative models of temporal and social discounting, I used mixed effects model-fitting, whereby the group-level distribution of parameters forms an empirical prior for individual-level parameter estimation (see General Experimental Methods, Chapter 5; Huys et al., 2011; Huys et al., 2012), constraining unreliable parameter estimates at the individual level. For each model, the probability of choosing an option over the alternative depended on their relative utilities according to a softmax function. Model comparison was performed at the group level using the integrated Bayesian Information Criterion (BIC_{int}), which unlike conventional BIC, approximates the Bayesian model evidence by integrating over subject-level parameters (Huys et al., 2011; Huys et al., 2012, see Chapter 5). Since the above approach assumes that parameters are distributed log normally in the population, throughout I firstly analysed choice data directly with no such distributional assumptions.

8.3.1.6 Subjective Questionnaire Data

After the experiment both participants completed three behavioural questionnaires in a secure online format (Qualtrics, Provo, US) namely: the Personal Altruism Level (PAL) Scale (Rushton, Chrisjohn, & Fekken, 1981), the Behavioural Approach - Behavioural Inhibition System scale (BIS/BAS) (Carver & White, 1994) and the Interpersonal Reactivity Index (IRI) (M. H. Davis, 1983). I analyzed information from the following sub-scales of each survey: the ‘willingness to help strangers’ subscale of the PAL, the BIS subscale of the BIS/BAS scale, and the ‘empathic concern’ subscale of the IRI, intended to assess “other orientated feelings of sympathy and concern for unfortunate others”. I hypothesised that dread would correlate positively with BIS scores and that social discounting would correlate negatively with PAL-willingness to help strangers and IRI-empathic concern scores.

8.3.2 RESULTS

8.3.2.1 Model-free Analysis of *Self now-self later* Choices

To test statistical hypotheses based on choice proportions without making assumptions about the distribution of the data, sampling distributions for one-sample tests were generated by randomly resampling the data with replacement (for paired tests by randomly reassigning data to either condition) with 100,000 iterations. The analysis follows that described in Chapter 6, whereby a 50% probability of choosing the delayed shock, $p(\text{choose later}) = 0.5$ indicates indifference between immediate and delayed shocks, $p(\text{choose later}) < 0.5$ indicates a preference for sooner pain and $p(\text{choose later}) > 0.5$ a preference for delayed pain.

To test for negative discounting in this study I calculated the area between the curve of $p(\text{choose later})$ vs delay and the 50% indifference line, using the trapezium method (Myerson, Green, & Warusawitharana, 2001). I termed this the area-under-the-curve (AUC), although its meaning here differs from conventional usage for positive temporal discounting. Here a negative area indicates negative discounting (preference for sooner pain), and a positive area indicates positive time preference (preference for delayed pain). Mean $p(\text{choose later})$ in *self now-self later* choices ($N=60$ participants) was below 50% (Figure 8.2A, solid grey circles) and decreased as delay increased. Across participants mean AUC (Figure 2B) was indeed significantly negative (mean AUC = -1.16, $p < 0.001$, permutation test) consistent with negative discounting.




Figure 8.2 Negative discounting for pain in the *self-now-self-later* condition in Experiment 3A. **A** Observed $p(\text{Choose later})$ (solid green circles) vs delay in seconds for all participants ($N=60$). Overlaid are the equivalent probabilities derived from the maximum a posteriori (MAP) policies of alternative models of dread: Exponential Dread (hollow circles), Hyperbolic Dread (hollow squares) and Non-Scaled Hyperbolic Dread (solid gray triangles). Error bars are one standard error. **B** The area between a line joining the observed data points a horizontal line at $p = 0.5$ was significantly negative, indicating negative discounting at the group level (permutation test; *** denotes $p < 0.001$). The error bar represents the 95% confidence interval of the mean (resampling method). **C** $p(\text{Choose later})$ (solid blue circles) in adjusting-sooner (i) and adjusting-later (ii) pain choices vs delay, overlaid with the MAP policies of the alternative models. Models in which the cost of waiting scaled with the intensity of the pain tended to underestimate dread at longer delays in adjusting-later pain choices. This was not so for the Non-Scaled model (see main text). Error bars represent one standard error. **D** Integrated Bayesian Information Criteria (BICint) of the alternative models. ‘Exp Dread (R)’ denotes a restricted Exponential Dread model (see main text). The red circle indicates the model with the highest evidence (lowest BICint), Non-Scaled Hyperbolic Dread.

8.3.2.2 Model-based Analysis of *Self now-self later* Choices

I first tested the model of dread proposed by Loewenstein (Loewenstein, 1987), here referred to as ‘Exponential Dread’. The delays in Experiment 3A were experienced in continuous time; therefore in this instance Loewenstein’s original formulation was fitted, rather than the discrete time version described in Chapters 6 and 7. For a description of both versions the reader is referred to Chapter 6. In this case the instantaneous utility of pain, $U(x)$, was implemented as a power function of the subjective intensity, x , of pain on a ten-point VAS (see Chapter 6, Equation 6.2.1), where b is a parameter governing the curvature of utility function. The resulting utilities were normalised by dividing by the utility of the most severe pain option, in this case intensity $x = 9$. In total the Exponential Dread model therefore has five free parameters at the subject level: α, δ, r, b and the softmax inverse temperature β . In Chapter 6 I proposed a modification of Loewenstein’s model, which posits that dread itself is not discounted, and that the rate of increase in dread is governed by the same discount factor as applied to the pain itself (in the current notation $r = \delta$). I also tested this restricted version here. In this case the full model marginally outperformed the restricted model, though there was insufficient evidence to conclude in favour of either version ($\Delta BIC_{int} = 1.45$).

Here I also tested a dread-discounting model based on hyperbolic discounting, termed ‘Hyperbolic Dread’. This model makes some heuristic assumptions. The first is that painful events are not temporally extended, but simply yield momentary utility, $U(x)$, an assumption which appears reasonable, since in this experiment the duration of pain itself was small by comparison to the delay. The second is that total prospective dread is simply proportional to delay, under the assumption that time perception is logarithmic (i.e. obeys the Weber-Fechner law, see for example Han & Takahashi, 2012; Takahashi, Oono, & Radford, 2008), giving:

$$U(x, d) = U(x) \left[\alpha \log(1 + \rho d) + \frac{1}{1 + Kd} \right] \quad (8.3.1)$$

The first term represents the contribution of dread, where ρ is a parameter governing the slope of logarithmic time perception, while the second term denotes hyperbolic

discounting with rate K . Notably a similar dread term would result from the forward integral of hyperbolically discounted utility. This model has five free parameters at the subject level: α , ρ , K , b and β . For both Exponential and Hyperbolic Dread models all parameters were bounded to be greater than zero by an exponential transform.

Figure 8.2A shows the mean maximum *a posteriori* (MAP) policies of the (unrestricted) Exponential Dread and Hyperbolic Dread models, overlaid with the observed data. Whilst Hyperbolic Dread provided a better fit to the data than Exponential Dread ($\Delta BIC_{int} = 33$), both models tended to underestimate the extent of dread at longer delays. Closer examination of the data revealed that this discrepancy was driven by greater-than-expected dread in choices where the immediate pain intensity was 5/10, and delayed pain intensity varied (adjusting later pain choices), compared with choices in which the delayed pain intensity was 5/10 and the immediate pain intensity varied (adjusting sooner pain choices), as shown in Figure 8.2C. For example, in adjusting later choices, the average participant chose an immediate 5/10 pain to avoid waiting for a 3/10 pain on 68% of choices (mean $p(\text{Choose sooner})=0.68$). However in adjusting sooner choices participants only chose a 7/10 pain to avoid waiting for a 5/10 pain on 24% of choices (mean $p(\text{Choose sooner})=0.24$; two-sample permutation test $p<0.001$). This behaviour is not accounted for by non-linear disutility over pain intensity, since the models already incorporate this. A possible explanation is that the cost of waiting for pain increases with increasing delay, but not with increasing pain severity, inflating the relative cost of waiting for minor pain. To test this, I implemented a ‘Non-scaled Hyperbolic Dread’ model where:

$$U(x, d) = \alpha \log(1 + \rho d) + \frac{U(x)}{1 + Kd} \quad (8.3.2)$$

As shown in Figures 8.2A and C, this model provided a close correspondence to the observed data, successfully accounting for the above effects and outperforming the scaled Hyperbolic Dread model ($\Delta BIC_{int} = 62$).

To estimate of the proportion of variance in participants’ choices explained by the Non-Scaled Hyperbolic Dread model, I regressed the observed probability of choosing delayed pain, $p(\text{Choose later})$ for each participant at each delay against the

equivalent probability predicted by the model; this procedure returned $R^2=0.89$. I also tested the parsimony of a restricted version of the Non-scaled Hyperbolic Dread model, in which b was constrained to equal one, giving rise to linear disutility of pain ($U(x) = -x$). The assumption of linear utility marginally improved the model fit ($\Delta BIC_{int} = 3$).

8.3.2.3 Model-free Analysis of Social Choices

Figure 8.3 displays mean $p(\text{choose later})$ in Deciders ($N = 27$) in the four conditions as a function of delay. I tested the ‘diminished social dread’ hypothesis by comparing $p(\text{choose later})$ across the four experimental conditions. Mean AUC was significantly less than zero in both the *self-now-self-later* (Figures 8.3A and 8.4A; mean AUC = -2.02; two-tailed permutation test uncorrected $p=0.002$) and *other-now-other-later* conditions (mean AUC = -1.16; two-tailed permutation test uncorrected $p=0.039$), and significantly more negative in the *self-now-self-later* condition than in the *other-now-other-later* condition (two-tailed permutation test uncorrected $p = 0.002$). These results indicate that participants indeed displayed negative discounting when choosing pain on behalf of others, though to a lesser extent than when choosing their own pain, consistent with the ‘diminished social dread’ hypothesis.

Secondly I compared AUC in *other-now-self-later* and *self-now-other-later* conditions, where Deciders were required to trade-off their own pain with that of the Receiver (Figures 8.3B and 8.4A). Here AUC was calculated relative to the point at which these curves cross the vertical axis, at $p(\text{choose other}, \text{delay}=0)$, rather than the 50% indifference line, to correct for social-discounting. AUC was significantly negative in the *other-now-self-later* condition (mean = -1.14, two-tailed uncorrected $p = 0.014$) and marginally significant in the *self-now-other-later* condition (mean = -0.70, two-tailed uncorrected $p=0.084$). However the difference in AUC between the two conditions was not significant (two-tailed permutation test $p > 0.250$). A possible explanation is that the increased complexity of the trade-off conditions lead to greater noise in participants’ value estimates, tending to diminish the effect of dread and obscuring self-other differences.




Figure 8.3 Time preference for pain for self and other in Experiment 3A. **A** Mean probability in Deciders ($N=27$) of choosing the delayed shock option at each delay in *self-now-self-later* (hollow gray squares) and *other-now-other-later* (hollow blue squares) conditions. **B** The same is plotted for *other-now-self-later* (hollow gray squares) and *self-now-other-later* (hollow orange squares) conditions. In both **A** and **B**, error bars are one standard error above and below the mean. Overlaid are maximum *a posteriori* policies of a Restricted Dread Noise model (hollow circles joined by dashed lines), fitted jointly to all four conditions.

In some choices, termed *self-now-other-now* choices (36 in total per participant), the shock for both participants was immediate (delay 0 in Figure 8.3B). Since there were an equal number of these choices in which the Decider's shock was the more intense as in which the Receiver's shock was the more intense, the proportion of these choices on which Deciders chose a shock for the Receiver, termed $p(\text{choose other})$, provides an approximate measure of their social discounting. While $p(\text{choose other})=0.5$ suggests no social discounting (i.e. that participants were indifferent as to the recipient of the shocks), $p(\text{choose other})<0.5$ would imply negative social discounting (hyperaltruism) and $p(\text{choose other})>0.5$ would imply positive social discounting (self-oriented behaviour). A one-tailed permutation test with the alternative hypothesis mean $p(\text{choose other})<0.5$ returned $p=0.12$, indicating a non-significant trend towards negative social discounting (mean $p(\text{choose other}) = 0.45$; Figure 8.4B).




Figure 8.4 Social modulation of dread and social discounting in Experiment 3A **A** Mean AUC in the four conditions, abbreviated by their initials. Error bars represent 95% confidence intervals (resampling method). AUC was significantly more negative in the *self-now-self-later* condition than in the *other-now-other-later* condition (two-tailed permutation test $p=0.002$), but did not differ significantly between *other-now-self-later* and *self-now-other-later* ($p>0.250$). **B** In *self-now-other-now* choices there was a non-significant trend towards negative social discounting (mean $p(\text{choose other})=0.45$; $p=0.12$, permutation test). Error bars represent a 95% confidence interval (resampling method). **C** Distribution of posterior (log transformed) parameter estimates of an unrestricted Dread Noise model. ω_r is the proportional change in the dread parameter, r , when evaluating others' pain, ω_β the proportional change in the softmax gain, β , in choices involving others' pain, and k is a social discount factor applied to others' pain. Cyan bars show the group mean parameter estimates, green error bars the 95% confidence interval. Black dots represent individual parameter estimates, with error bars of the square root of the second moment around the individual estimates. Mean $\log \omega_r < 0 < 0$ at the group level, consistent with diminished social dread. Mean $\log k > 0$ (mean $k > 1$), consistent with negative social discounting. **D** Integrated Bayesian Information Criteria (BICint) of alternative models. The red circle indicates the model with the highest model evidence (lowest BICint), a Restricted Dread Noise Model.

The distribution of $p(\text{choose other})$ across participants (shown in Figure 8.4) did not appear to be unimodal, but appeared to consist of up to three peaks – a central peak with a mode marginally below 0.5 and two smaller peaks at the extremes. I tested whether there was sufficient evidence to conclude in favor of a multimodal distribution by fitting a Gaussian mixture models to the data. In this case, the most parsimonious model was a unimodal Gaussian with parameters $\mu = 0.45$ and $\rho = 0.05$ (conventional BIC = 3.15), with next best fit provided by a mixture of three Gaussians (BIC=8.65), indicating insufficient evidence to support a trimodal form in this instance.




Figure 8.5 Distribution of social discounting across participants in Experiment 3A. Histogram of the probability of choosing a shock for the other participant, $p(\text{Choose other})$, in *self-now-other-now* choices in Experiment 1.

8.3.2.4 Model-based Analysis of Social Choices

To investigate how the findings above might be instantiated in terms of a model of dread-discounting, I fitted a Non-scaled Hyperbolic Dread model jointly to all four conditions, in which the effect of delay interacted with the recipient of the pain such that:

$$U_{self}(x, d) = \alpha \log(1 + \rho d) + \frac{U(x)}{1 + Kd} \quad (8.3.3)$$

$$U_{other}(x, d) = \kappa \left[\alpha \log(1 + \omega_\rho \rho d) + \frac{U(x)}{1 + \omega_K K d} \right] \quad (8.3.4)$$

$U_{self}(x, d)$ indicates disutility which the Decider derives from their own shocks, and $U_{other}(x, d)$ disutility which they derive from the Receiver's shocks. κ is a social discount factor representing the extent to which shocks for the Receiver create disutility for the Decider, independent of the effect of delay. $\kappa = 1$ indicates that the Decider disvalues (immediate) shocks for the Receiver equally to their own shocks, i.e. no social discounting. $\kappa < 1$ indicates that the Decider derives more disutility from their own shocks than from the Receiver's shocks, i.e. positive social discounting. $\kappa > 1$ indicates that the Decider derives more disutility from the Receiver's shocks than from their own shocks. The weighting factors ω_ρ and ω_K represent the extent to which the effect of delay is altered when evaluating others' pain, relative to one's own pain. Under the Null hypothesis of no interaction between delay and recipient, $\omega_\rho = \omega_K = 1$.

Three forms of the model were compared against the Null model, firstly a general modulation of the effect of delay when evaluating others' pain, termed the Delay model ($\omega_\rho = \omega_K$), secondly a specific modulation of discounting, termed the Discounting model ($\omega_K; \omega_\rho = 1$) and thirdly a specific modulation of dread, termed the Dread model ($\omega_\rho; \omega_K = 1$). Each of these models had seven free parameters at the subject level (to fit 396 choices per subject across four conditions): $\alpha, \rho, K, b, \omega, \kappa$ and β , corresponding to seven group level priors. A previous study (Crockett et al. 2014) found that participants exhibited noisier choices when also required to evaluate others' pain. To test for this here, I fitted additional versions of the models above in which the softmax inverse temperature, β , was diminished by a factor $0 < \omega_\beta < 1$ in all conditions except *self-now-self-later*. These models are referred to as Null-Noise, Discounting-Noise, Delay-Noise and Dread-Noise respectively. Each of these Noise models therefore had eight free parameters at the subject level.

Finally, for all models except the Null and Null-Noise models, I tested the parsimony of two forms of population prior on the social weighting factors ω_ρ or ω_K . The first form of prior (termed Unrestricted) represents a modulation of dread in either direction, and allows all ω_ρ or $\omega_K > 0$, achieved by exponentially transforming the fitted parameters. The second enforces diminished social dread, or increased

social discounting, by restricting ω such that $0 < \omega < 1$ for Delay and Dread models, and $\omega > 1$ for Discounting models (Restricted), achieved by sigmoid and exponential transformation respectively. In particular the Restricted Dread model proposes that *all* participants exhibit either no social modulation of dread, or diminished social dread. This is a strong form of the diminished social dread hypothesis, and the comparative performance of this model therefore provides a conservative test for diminished social dread.

The most parsimonious fit was provided by the (Restricted) Dread-Noise model (Figure 8.3E; ΔBIC_{int} compared with Null model 102). The mean MAP policy of this model, formed by taking the mean of the subject-level policies, is overlaid with the observed choice probabilities in Figure 8.3A. It can be appreciated that the model reproduces the gross pattern of the data adequately. To estimate of the proportion of variance explained, I regressed $p(\text{Choose later})$ for each participant at each delay against the equivalent probability predicted by the model; this procedure returned $R^2=0.91$. That this restricted model, which enforces diminished social dread, outperformed its unrestricted counterpart ($\Delta BIC_{int}=43$) supports the diminished social dread hypothesis. The distribution of (log transformed) MAP estimates of ω_ρ , ω_β and κ for the *unrestricted* Dread-Noise model are shown in Figure 8.3D; 24 out 27 ω_ρ estimates lie below 0 (corresponding to $\omega_\rho < 1$), consistent with the superior performance of the restricted model.

Finally, to test for an effect of social discounting, I tested the parsimony of a version of the best fitting (Restricted Dread-Noise) model in which $\kappa = 1$, i.e. no social discounting. This restriction substantially decreased the quality of model fit ($\Delta BIC_{int} > 1000$), indicating that social discounting was an important source of variance in the data. The posterior estimate of mean $\log(\kappa)$ was greater than zero (mean $\log(\kappa) = 0.36$), indicating $\kappa > 1$, i.e. negative social discounting. However there was substantial between-subject variability (s.d. $\log(\kappa) = 0.91$), consistent with non-significant negative social discounting at the group level in direct analysis of choice data above.

8.3.2.5 No Relationship between Social and Time Preferences for Pain

There was no direct relationship between dread of one's own pain and social discounting for pain. To test for this AUC in the *self-now-self-later* condition was regressed against $p(\text{choose other, delay}=0)$ using robust linear regression with weighted least-squares, yielding no significant relationship between the two ($\beta_1 = 0.19$; $p > 0.25$; ordinary least squares $R^2 = 0.014$).

8.3.2.6 Skin Conductance to Others' Shock Correlates with Altruism

Of the 27 participants designated as Deciders, skin conductance and heart rate data were collected from 21 out of the 27. Physiological data from the remaining 6 Deciders were not obtained due to time constraints. Skin conductance data from two of the 21 were unsuitable for analysis, in one case due to movement artefacts and in another due to hyperhidrosis. Heart rate data from 6 out of the 21 were unsuitable for analysis, due to poor signal from the pulse oximeter. Due to this unexpectedly high rate of data loss, resulting in a small number of complete datasets for this analysis ($N=19$) the findings below should be interpreted as exploratory.

The mean skin conductance response (SCR) at each delay across all participants whose data were suitable for analysis is displayed in Figure 8.6. There was a significant group-level responses to receiving a shock oneself (*shock-self* $t(18) = 5.68$, $p < 0.001$). However, surprisingly there was no significant conductance response at the group level to either the anticipation of shocks for oneself (*anticipation-self* $t(18) = 0.41$, $p > 0.250$), anticipation of shocks for the other participant (*anticipation-other* $t(18) = -0.76$, $p > 0.250$) or administration of shocks to the other participant (*shock-other* $t(18) = 0.69$, $p > 0.250$).

There was no significant relationship between GLM β values for anticipation-self and dread parameters α or ρ from the Non-Scaled Hyperbolic Dread model above (Robust linear regression with bisquare weighting function, $p > 0.250$). GLM β values for the event-related conductance responses to the other participant's shock were also regressed against the social discount factor, κ . As predicted, this yielded a significant positive relationship (β -shock-other vs $\log(\kappa)$: $\beta_1 = 0.45$, $p < 0.001$), indicating

that individuals with a greater sudomotor response to the other participants' shock tended to be more altruistic.

Finally I compared heart rate responses to the anticipation of a shock for oneself, with those to the anticipation of a shock for the other participant. To do so I extracted the mean heart rate across the longer delay periods (8, 16 and 29s). A 2x3, self-other vs delay repeated measures ANOVA (in the 15 participants whose data was suitable for analysis) yielded no significant main effect of either delay ($F(2,28) = 0.80$, $p > 0.250$) or recipient, ($F(1,14) = 1.37$, $p > 0.250$) and no significant interaction ($F(2,28) = 1.10$, $p > 0.250$).

A




Figure 8.6 Skin conductance responses to delayed pain for self or other in Experiment 3A. The mean skin conductance response across Deciders with suitable data ($N=19$) is displayed for each delay to shock, time-locked at time zero to the onset of the countdown timer. The blue vertical line in each plot indicates the onset of shock. The solid trace depicts the mean response when the shock was for the participant themselves, the dashed trace the response when the shock was for the other participant. We found a significant group-level responses to receiving a shock oneself (*shock-self* $t(18) = 5.68$, $p < 0.001$). However, surprisingly there was no significant conductance response at the group level to either the anticipation of shocks for oneself (*anticipation-self* $t(18) = 0.41$, $p > 0.250$), anticipation of shocks for the other participant (*anticipation-other* $t(18) = -0.76$, $p > 0.250$) or administration of shocks to the other participant (*shock-other* $t(18) = 0.69$, $p > 0.250$).

8.3.2.7 Questionnaire Data

Complete questionnaire data were available from 55 participants who made *self-now-self-later* choices. Robust regression of dread parameters (a and ρ) against Behavioural Inhibition System (BIS) score (Carver & White, 1994) yielded no significant relationship ($p > 0.25$ for both β_a and β_ρ). Of the 27 Deciders, complete questionnaire data were available from 25. For these participants there was no significant relationship between the social discount factor (κ) and the willingness to help strangers subscale of the Personal Altruism Level scale (Rushton et al., 1981) ($\beta_\kappa = 1.08$; $p=0.193$) nor between the social discount factor and the empathic concern subscale of the Interpersonal Reactivity Index (M. H. Davis, 1983) ($p>0.250$).

8.3.3 SUMMARY OF RESULTS FROM EXPERIMENT 3A

In summary Experiment 3A replicated previous findings that participants are averse to waiting for pain, showing that this tendency was best accounted for by a Hyperbolic Dread-Discounting model in which the cost of waiting did not scale with the intensity of pain. As predicted, there was evidence for diminished social dread when evaluating others' pain. Model comparison indicated that participants valued others' pain differently from their own, with a trend towards negative social discounting for pain at the group level. Experiments 2 and 3 test whether these effects replicate in a different context.

8.4 EXPERIMENT 3B: SOCIAL DISCOUNTING OF AN HYPOTHETICAL PAINFUL TREATMENT

Experiment 3A was conducted in the laboratory, with delays in the order of seconds. By contrast, the real-world examples evoked above entail longer delays. Experiment 3B was designed to test whether the self-other difference in dread reported in Experiment 3A was robust to changes in context and in the timescale of choices, as well as to test for effects of social distance.

8.4.1 METHODS

8.4.1.1 Participants

Healthy participants who had worker accounts on Amazon Mechanical Turk (AMT) were recruited via an advertisement on the AMT account of one of the authors (MC). AMT is an online marketplace for work, now widely used as a method of data collection for psychology experiments (Mason & Suri, 2012). Responses on AMT have been shown to be reliable (Rand, 2012) and replicate well-established findings in the cognitive psychology literature (Crump, McDonnell, & Gureckis, 2013).

67 participants completed Experiment 3B (34 females; mean age 34.0 years, *s.d.* =9.3 years; median yearly income category 25-30,000 USD; 48/67 university educated). All workers had a record of a high consistency of responding in similar tasks. The mean time to completion was 21 minutes. Participants were compensated \$3, the current market rate for AMT. All participants gave full informed consent before taking part in the study, and were free to withdraw their consent at any time by simply discontinuing the online questions.

8.4.1.2 Experimental Choices

Experiment 3B required participants to consider hypothetical painful outcomes whilst both the timing and intensity of the hypothetical pain and the social distance of the recipient were manipulated. In the first of two sets of choices, participants ($N=67$) made choices between two possible timings for a hypothetical medical treatment, termed Treatment A, which was said to prevent the onset of dementia later in life, and was said to be completely safe, free of charge and 100% effective. Participants were told that the treatment would involve an injection into the bone marrow causing pain lasting for half an hour. The intensity of the pain was said to depend on which nurse was available to perform the injection. Participants were explicitly instructed that neither the level of pain, nor identity of the nurse, nor the timing of the appointment would alter the effectiveness of the treatment. In each case, participants were told the timing of the appointment, and how painful the treatment would be on a 10-point scale.

This first set of choices recapitulated *self-now-self-later* and *other-now-other-later* conditions from Experiment 3A, following the same form of symmetrical adjusting design; participants made 72 choices between two possible treatment appointments either for themselves, or on behalf of another person (in this case the 50th person on an imagined list of their 100 people closest to them, see below). For each choice, one appointment option was always ‘this week’ (delay 0 weeks), and the other was delayed 4, 17, 52 or 416 weeks. To avoid habitual responding, half of the time the sooner appointment appeared on the left of the screen and half of the time on the right.

A second set of choices examined effects of social distance on choices between pain for oneself and pain for others. Here a second treatment, Treatment B, was said to be available, involving a donation of immune cells from another person, and requiring a painful injection into the bone marrow for the donor, but no pain for the person receiving the treatment. Participants chose between two different treatment options for another person: either Treatment A, entailing a painful injection for the other person, or Treatment B, entailing a painful injection for them but no pain for the other person, whilst the social distance of the other person was varied, using the following instruction introduced by Jones and Rachlin (Jones & Rachlin, 2006; Jones & Rachlin, 2009; Rachlin & Jones, 2008):

This experiment asks you to imagine that you have made a list of the 100 people closest to you in the world ranging from your dearest friend or relative at position #1 to a mere acquaintance at #100. The person at number one would be someone you know well and is your closest friend or relative. The person at #100 might be someone you recognise and encounter but perhaps you may not even know their name. You do not have to physically create the list—just imagine that you have done so.

As in previous experiments, choice options followed a symmetrical adjusting amount design ($N=90$ choices per participant), with social distances of #1, #5, #14, #45 and #97 (see Appendix, Table A4 for a full list of choice options). All choices were administered via the online survey software, Qualtrics (www.qualtrics.com; Provo, UT), which participants accessed via an anonymous link from their AMT account, and participants made choices by clicking on their preferred option using the

computer mouse. The order of conditions and trials within conditions was randomised between participants. Pain intensity was described using the following ten-point pain scale: 0 = not at all painful, 2 = mild pain, 4 = annoying, unpleasant pain, 6 = discomforting, troublesome pain, 8 = intense, distressing pain and 10 = excruciating pain. The structure of choices was identical to that of the *self-now-self-later* condition in Experiment 1, with pain levels of 1, 3, 5, 7 or 9 on the above scale.

8.4.2 RESULTS

8.4.2.1 *Self now-self-later* and *other-now-other-later* Choices

In keeping with Experiment 3A, and previous findings with hypothetical dental pain, participants showed highly significant negative temporal discounting for the delayed appointments, for both self (Figure 8.7A; mean AUC = -87.3, one-sample, two-tailed permutation test, $p < 0.0001$) and other (mean AUC = -86.4, $p < 0.0001$). However, here there was no significant difference between negative temporal discounting for self and other (permutation test by randomly reassigning conditions, $p > 0.250$). I explore this null finding in more detail below.

Consistent with Experiment 3A, Hyperbolic Dread models outperformed their Exponential counterparts in accounting for *self-now-self-later* choices, with Non-Scaled Hyperbolic Dread the most parsimonious model (Figures 8.7A and B; $\Delta BIC_{int} = 255$ compared with Scale Hyperbolic Dread). As previously, I also tested the parsimony of constraining $\alpha = 1$ (such that ρ alone governs the effect of dread) or $b = 1$ (linear disutility of pain). The former did not improve model fit ($\Delta BIC_{int} = -98$), however in this case linear utility was found to be a parsimonious assumption ($\Delta BIC_{int} = 60$). Regressing $p(\text{choose later})$ for each participant at each delay against the equivalent probability predicted by the model returned $R^2 = 0.95$. Given the lack of a significant group-level difference in negative discounting for self versus other, in this experiment I did not perform model comparison of social effects on dread.




Figure 8.7 Experiment 3B: Dread-discounting of hypothetical painful treatments for self and other. **A** Mean probability for all participants in Experiment 3B ($N=67$) of choosing a delayed treatment at each delay, when choosing on behalf of oneself (*self-now-self-later*, dark blue circles) or another person at social distance #50 (*other-now-other-later*, cyan circles). AUC did not significantly differ between the two conditions ($p>0.1$). Overlaid are the maximum *a posteriori* fits of Exponential Dread (hollow circles), Hyperbolic Dread (hollow squares) and Non-scaled Hyperbolic Dread (solid gray triangles) models to the *self-now-self-later* data. Error bars represent one standard error above and below the group mean. **B** Integrated Bayesian Information Criteria (BICint) of alternative dread models fitted to *self-now-self-later* choices; the red circle indicates the model with the highest model evidence (lowest BICint), Non-Scaled Hyperbolic Dread. ‘Exp Dread (R)’ denotes a restricted Exponential Dread model (see main text).

8.4.2.2 *Self-now-other-now* Choices: Social Discounting for Pain

A model-free metric of social discounting is provided by $p(\text{choose other})$ with direct equivalence to $p(\text{choose later})$ in delay discounting analyses. Mean $p(\text{choose other})$ across participants ($N=67$) is plotted as a function of social distance in Figure 8.8A. As predicted, participants displayed negative social discounting for close others, i.e. chose pain for themselves to relieve a close other of a less severe pain (mean $p(\text{choose other \#1}) = 0.4$, 95% CI by resampling method [0.33 – 0.47], one-sample two-tailed $p=0.003$), however this tendency was steeply discounted with social distance, such that by social distance #5 participants showed neither negative nor positive discounting (mean $p(\text{choose other \#5}) = 0.5$, 95% CI [0.44-0.56]) and exhibited significant positive social discounting by social distance #45 (mean $p(\text{choose other \#45}) = 0.57$, 95% CI [0.50-0.64], $p=0.04$).




Figure 8.8 Social discounting of hypothetical painful treatment in Experiment 3B **A** Mean probability ($N=67$) of choosing a treatment which entailed pain for the other person, rather than pain for oneself ($p[\text{choose other}]$, solid green squares), as a function of social distance, ranging from the participant's closest friend or relative at distance #1, to a mere acquaintance at #100. Error bars represent one standard error. $P=0.5$ indicates parity between the disvalue of pain for self and other (solid horizontal line). Participants displayed negative social discounting for close others, which decayed with social distance. Overlaid are fits of Subtractive Harm Aversion models. **B** BICint of alternative models of social discounting: hyperbolic discounting (Hyp), Multiplicative Harm Aversion (M-HA), Subtractive Harm Aversion (S-HA), Precautionary Cost (PC) and Subtractive Harm Aversion with Precautionary Cost (S-HA PC). The red circle indicates the model with the highest model evidence, in this case S-HA PC (see main text).

An intuitive explanation for this pattern is that people are generally averse to harming others (Crockett, Clark, Hauser, & Robbins, 2010; Crockett et al., 2014; Cushman, Gray, Gaffey, & Mendes, 2012), but that this tendency decays with social distance (Jones & Rachlin, 2006). I tested this possibility formally through model comparison. Since all treatment options were due ‘this week’ (0 delay), no dread or temporal discounting terms were required here. Given the parsimony of linear utility for the same form of outcome above, all models assumed linear utility:

$$U_{self}(x) \equiv U_{other}(x) \equiv -x \quad (8.4.1)$$

Here the null model was taken to be simple Hyperbolic Social Discounting, of the form shown in Equation 8.1.1. I first tested an alternative model, termed Multiplicative Harm Aversion, in which harming others incurs a proportional cost, κ , (relative to harming oneself):

$$U_{other}(x, d) = \frac{\kappa U(x)}{1 + K_{soc}d} \quad (8.4.2)$$

where d indicates social distance and K_{soc} a subject-specific social discount rate. This model had three free subject-level parameters: K_{soc} , κ and β (to fit 90 choices per subject). I also tested a model, termed Subtractive Harm Aversion, in which harming others incurs a subtractive cost, θ , which is itself hyperbolically discounted with rate K_θ such that:

$$U_{other}(x, \delta) = \frac{U(x)}{1 + K_{soc}d} - \frac{\theta}{1 + K_\theta d} \quad (8.4.3)$$

This model had four free subject-level parameters: K_{soc} , K_θ , θ and β . As shown in Figure 8.8B, both Multiplicative and Subtractive Harm Aversion models conclusively outperformed Hyperbolic Social Discounting, with Subtractive Harm Aversion the better performing of the two ($\Delta BIC_{int} = 1099$ and 1135 respectively). To estimate of the proportion of variance in participants’ choices explained by the Subtractive Harm Aversion model, I regressed $p(\text{choose other})$ for each participant at each delay against the equivalent probability predicted by the model; this procedure returned $R^2=0.92$.

Closer examination of the data revealed heterogeneity in subject-level responses. Example patterns are shown in Figure 8.9. The first three patterns represent no significant degree of social discounting (8.9A), solely positive social discounting (8.9B) and the archetypal pattern of initial negative social discounting followed by positive discounting (8.9C). The participant whose data is shown in Figure 8.9D exhibits initial negative social discounting (hyperaltruism for close others), which is discounted back to the 50% indifference line, but not beyond ($K_\theta > 0, K_{soc} \approx 0$).




Figure 8.9 Social discounting of pain in individual participants in Experiment 3B. p[Choose other] as a function of social discounting for five example subjects, illustrating heterogeneity in responses: **A** non-significant social discounting, **B** positive social discounting **C** and **D** initial negative social discounting (discounted harm aversion) and **E** consistent negative social discounting. Overlaid are the best fits of a S-HA PC model (gray circles joined by solid gray lines). Error bars represent one standard error of the binomial distribution ($N=18$ choices per data point, with the exception of the intercept, which is set to 0.5 on *a priori* grounds).

Figure 8.9E shows a pattern that is accounted for by neither of the models above, namely steep negative social discounting that appears to increase with social distance, in a manner resembling the dread functions seen for delayed pain. This pattern would be consistent with growing uncertainty regarding the other person's response to inflicted pain with increasing social distance, generating a precautionary motive to avoid harming distant others. To examine whether this pattern of persistent or increasing negative social discounting appeared to have support at the group level,

I also fitted *post-hoc* a model with an additional term in which an aversion to harming others increased with delay, in a manner directly equivalent to dread:

$$U_{other}(x, \delta) = \frac{U(x)}{1 + K_{soc}\delta} - \frac{\theta}{1 + K_{\theta}\delta} - \alpha_{soc}\log(1 + \rho_{soc}d) \quad (8.4.4)$$

This model had six free parameters at the subject level: K_{soc} , K_{θ} , θ , α_{soc} , ρ_{soc} and β . Despite the additional parameters, this ‘Subtractive Harm Aversion with Precautionary Cost’ model improved fit quality compared with Subtractive Harm Aversion alone ($\Delta BIC_{int} = 17$; $R^2 = 0.95$). However notably this more complex model was developed after seeing the data, and so may not be considered as plausible *a priori* as the two simpler models. A restricted model with Precautionary Cost alone ($\theta = 0$) performed relatively poorly ($\Delta BIC_{int} = -187$ compared with Subtractive Harm Aversion). Mean MAP policies of Subtractive Harm Aversion and Subtractive Harm Aversion with Precautionary Cost models are overlaid with the observed data in Figure 8.8A.

Unexpectedly there was no significant relationship between social and temporal discounting for pain (robust regression of AUC_{social} vs AUC_{delay} , $\beta_1 = -0.04$, $p > 0.250$), suggesting that dread of pain and hyperaltruism for others’ pain may be distinct phenomena.

8.4.3 SUMMARY OF RESULTS FROM EXPERIMENT 3A

In summary, Experiment 3A revealed significant negative temporal discounting for both self and other, where in this case the ‘other’ was a friend or acquaintance (social distance #50). However, unlike in Experiment 3A, there was no significant self-other difference in negative temporal discounting. There was however robust evidence for negative social discounting (hyperaltruism) for close others, and this tendency appeared to be hyperbolically discounted with social distance.

8.5 EXPERIMENT 3C: INTERACTION OF SOCIAL AND TEMPORAL DISCOUNTING FOR PAIN

In Experiment 3B, *other-now-other-later* choices were on behalf of a person at social distance #50. In fact, the mean participant actually exhibited a small degree of positive social discounting at social distance #50 (see Figure 8.8A); this is in contrast to Experiment 3A where the mean participant exhibited *negative* social discounting for the other participant. An explanation for the lack of significant social modulation of dread at the group level in Experiment 3B might therefore be that the diminished social dread is dependent on negative social discounting. To test this hypothesis directly, in Experiment 3C, participants made *other-now-other-later* choices for people at two different social distances, #1 (their closest friend or relative) and #50 (a friend or acquaintance). Based on the considerations above I predicted diminished social dread for #1 but not for #50.

8.5.1 METHODS

Participants were recruited via AMT and completed the study online. 70 participants completed Experiment 3C (32 females; mean age 36.8 years, s.d. =10.5 years; median yearly income 30,000 USD; 31/70 university educated). and 16 minutes for Experiment 3; for each experiment participants were compensated \$3, the current market rate for AMT. All participants gave full informed consent before taking part in the study, and were free to withdraw their consent at any time by simply discontinuing the online questions. The experimental scenario and design was identical to Experiment 3B, with the exception that the 416-week delay was omitted, to reduce the overall duration of the questionnaire. The order of the three conditions and trials within conditions was randomised between participants.

8.5.2 RESULTS

As predicted, participants exhibited significantly diminished negative temporal discounting when choosing the timing and intensity of pain on behalf of

their closest friend or relative compared with when choosing their own pain (Figures 8.10A and 8.10B; mean AUC *self-now-self-later* = -6.28, mean AUC *other-now-other-later* #1 = -4.47; two-sample permutation test $p=0.022$). Mean AUC when deciding on behalf of the person at social distance #50 was intermediate between these two conditions and significantly different from neither (mean AUC *other-now-other-later* #50 = -5.24; two-sample permutation tests $p>0.10$).

Given the reduced range of delay choices in Experiment 3C, power to resolve differences between alternative models was limited. Nevertheless, by way of demonstration that this finding was consistent with the models introduced above I fitted the same categories of model here. In this instance κ and the social weighting on the softmax temperature, ω_β , have identical effects, both multiplying with U_{other} ; I therefore simply set $\kappa = 1$ and fitted ω_β . Since in this case *other-now-other-later* choices were made at two levels of social distance, the social weighting factors on the softmax temperature (ω_β) and dread (ω_ρ) were fitted separately for each social distance. Based on the findings of Experiment 3B, linear utility was imposed. Each model, besides the null model, therefore had 8 free parameters at the subject level (to fit 162 choices per subject over 3 conditions): $\alpha, \rho, K, \omega_{\rho\#50}, \omega_{\rho\#1}$ (or $\omega_{K\#50}, \omega_{K\#1}$), $\omega_{\beta\#50}, \omega_{\beta\#1}$ and β . For the null model fitted parameters were $\alpha, \rho, K, \omega_{\beta\#50}, \omega_{\beta\#1}$ and β , with $\omega_{\rho\#50} = \omega_{\rho\#1} = 1$ (or $\omega_{K\#50} = \omega_{K\#1} = 1$).




Figure 8.10 Social modulation of dread in Experiment 3C. **A** Mean $p(\text{choose later})$ across subjects ($N=70$) for *self-now-self-later* (solid blue squares), *other-now-other-later* #1 (solid cyan squares) and *other-now-other-later* #50 (solid green squares) conditions. Error bars are one standard error. Overlaid are maximum *a posteriori* policies of a Delay Noise model (hollow circles joined by dashed lines), fitted jointly to all three conditions. **B** Mean AUC in the three conditions, abbreviated by their initials. Error bars represent 95% confidence intervals by resampling. AUC was more negative in *self-now-self-later* than in the *other-now-other-later* #1 choices (two-tailed permutation test $p = 0.02$), but did not differ between the remaining conditions ($p > 0.250$). **C** Integrated Bayesian Information Criteria (BICint) of alternative models. Delay Noise had the highest model evidence (red circle, lowest BICint). Delay-Noise, Dread-Noise and Restricted Delay-Noise all outperformed the Null model. **D** Distribution of posterior (log transformed) parameter estimates of an unrestricted Delay-Noise model. ω_r #50 is the proportional change in the dread parameter, r , when evaluating the pain of a person at social distance #50 (a friend or acquaintance) pain, ω_r #1 the equivalent for distance #1 (closest friend or relative). Bars represent the group mean parameter estimates, green error bars the 95% confidence interval. Black dots are individual parameter estimates, with error bars representing the square root of the variance on the individual estimates. Mean log ω_r #1 was negative at the group level, consistent with diminished social dread for close others.

The results of model comparison are shown in Figure 8.10C. In this instance an unrestricted Delay-Noise model provided the most parsimonious account of the data (although an unrestricted Dread-Noise model also outperformed the Null-Noise model). For this model $R^2 = 0.93$ (calculated as previously described). Log transformed posterior estimates of $\omega_{\rho\#1}$ and $\omega_{\rho\#50}$ are displayed in Figure 8.9D; posterior mean $\omega_{\rho\#1}$ was lower than the posterior mean $\omega_{\rho\#50}$ (mean $\log \omega_{\rho\#1} = -0.69$, s.d. = 1.30; mean $\log \omega_{\rho\#50} = -0.40$, s.d. = 1.42), consistent with the behavioural data. In summary Experiment 3C indicates diminished negative temporal discounting when deciding on behalf of close, but not socially distant others, an effect which is parsimoniously accounted for by a general down-weighting of delay when deciding for close others.

8.6 SOCIAL PREFERENCES FOR PAIN: GENERAL DISCUSSION

The three experiments described above examine the interaction between social and temporal preferences for pain. A prior prediction was that people would exhibit negative temporal discounting of pain for both themselves and others. However it was hypothesised that negative discounting would be less marked when evaluating others' pain, termed the 'diminished social dread' hypothesis. A further prediction was that people would choose to endure pain to relieve an even milder pain for another person, termed negative social discounting, and that this effect would be most pronounced when the other person was a close friend or relative, referred to as the 'discounted harm aversion' hypothesis.

Consistent with previous findings (Chapter 6; Berns et al., 2006; Loewenstein, 1987), participants indeed displayed negative temporal discounting for their own pain, indicated by choices to expedite pain. Also consistent with previous findings (Chapter 6), across all three experiments the cost of waiting for pain increased with increasing delay, though the rate of increase declined at longer delays. Models of dread based on hyperbolic discounting outperformed the exponential models tested previously (Chapter 6). Under the best performing model, hyperbolic discounting of pain was combined with a waiting cost ('dread'), which increased logarithmically with increasing delay, but did not scale with the intensity of pain.

Participants also preferred sooner pain when deciding on behalf of another person. In Experiment 3A, where the other person was another experimental participant in an adjacent room, the extent of this negative temporal discounting was significantly less than when choosing their own pain. However this effect was not significant when participants were asked to trade-off their delayed pain with the other participant's immediate pain or *vice versa*. A possible explanation is that participants exhibited greater choice stochasticity in these more complex choices, in support of which models incorporating greater choice noise for the social conditions, as well as diminished social dread, provided a parsimonious account of the observed data.

I had predicted that people would give less emphasis to waiting costs, relative to the intensity of pain itself, when choosing on behalf of others, as a precaution against misjudging the other's preferences. As described above the data indeed support this diminished social dread hypothesis. However, the results suggest a different explanation for the effect than originally thought. Uncertainty regarding others' preferences for delayed pain would be expected to grow with increasing social distance from the other person, predicting greater diminution in dread for distant others than for close others. In fact the opposite pattern was seen: in Experiments 3B and 3C, with outcomes of a hypothetical painful medical treatment, at delays of up to 8 years, participants exhibited diminished social dread when choosing on behalf of a close friend or relative, but not when choosing on behalf of a more distant friend or acquaintance. A possible explanation would be that although uncertainty regarding the other's preferences ought to be lower when the other person is a close friend or relative, a putative precautionary motive is heightened, although this possibility is of course speculative.

In a subset of participants with suitable data, there were significant skin conductance responses (SCRs) to the receipt of shock itself, but not to the anticipation of shock, nor to the delivery of shock to another person. The latter null findings are surprising, given that such responses have been previously documented (Hein et al., 2011). A possible explanation for the lack of significant anticipatory responses would be that in this setup participants received reliable information as to the timing of shock, putatively tending to diminish anticipatory anxiety, which is known to be sensitive to uncertainty regarding the properties of the noxious stimulus (S. Epstein, Epstein, & Roupinian, 1970). A possible explanation for the lack of response to

others' shocks would be that the SCR data for the response to others' shocks were collected block-wise, and participants may have stopped attending to events on the screen for these trials. Despite the lack of significant group-level response to others' shock, the magnitude of response to others' shock correlated with the tendency to take on pain in lieu of the other person, a finding that has also been reported previously (Hein et al., 2011). Since these particular analyses were relatively underpowered ($N=19$) compared with the remainder of the analyses, owing to the number of participants with incomplete or un-usable data, I interpret them as preliminary data worthy of further investigation.

Participants' choices implied that they valued others' pain differently from their own, consistent with social discounting for pain. In Experiment 3A model fitting indicated a mean tendency towards negative social discounting for pain at the group level (17/27 showed negative social discounting, as indicated by fitted $\kappa > 1$), consistent with previous findings (Crockett et al., 2014). In Experiment 3B, with outcomes of a hypothetical painful medical treatment, participants exhibited negative social discounting for close others, choosing to endure pain to relieve even a milder pain for someone close to them. This hyperaltruistic tendency was steeply discounted with social distance, and reverted to positive social discounting for distant others, supporting a 'discounted harm aversion' hypothesis. Interestingly, this group-level tendency was also seen to a lesser extent in a previously published study of social discounting choices with money, where participants on average gave up close to \$80 in order to win \$75 for to a close friend or relative (Jones & Rachlin, 2006). An interesting question for future work would be to directly compare social discounting for monetary and non-monetary outcomes within participants. The apparently greater degree of hyperaltruism seen here for pain would be in keeping with the observation that people are more generous when choosing how to share non-monetary rewards, as compared to money, with others (A. L. Davis, Miller, & Weber, 2011). Why participants consistently display hyperaltruism for pain is unclear. One possibility is that people overestimate their own ability to deal with pain (relative to that of others). This could be easily tested, by asking participants to rate their perceived ability to cope with pain, relative to other people, and regressing this against hyperaltruism.

Based on previous studies showing that temporal and social discounting for rewards are correlated (Curry et al., 2008; Rachlin & Jones, 2008), I had expected to

find a correlation between social and temporal discounting for pain. Such a correlation would support the idea of a ‘patient-altruistic’ phenotype across domains, such that social preferences for pain might even be used to predict personal pain-related decisions. Surprisingly however there was no significant correlation between the two measures. This might be taken as evidence that discounting for rewards and for aversive stimuli engage distinct processes (see General Discussion).

Finally I would like to draw attention to the finding that a dread-discounting model in which the cost of waiting for pain did not scale with pain intensity provided the most parsimonious fit to the intertemporal choice data, raising the question as to whether the tendency to expedite pain represents a more generic form of impatience, i.e. a simple desire to ‘have things sooner’. In normative terms, this might result from a cognitive (computational) cost associated with keeping track of future events. Since this explanation is subtly different from the notion of anticipatory disutility, one must question whether the term ‘dread’ remains valid. However the non-scaled model was developed *post hoc*, and therefore warrants further investigation with a design intended to directly disambiguate scaled and non-scaled waiting costs.

A possible implication of the effect of diminished social dread is that policy-makers and caregivers, though acting with precaution and with the aim of minimizing harm to patients, might nevertheless tend to underestimate the extent to which patients dislike waiting for unpleasant treatments or procedures. To illustrate this hypothesis with an example, say that a nurse is preparing to take a blood sample from a patient, and to diminish the pain itself applies a topical anaesthetic cream, which takes a few minutes to reach maximum effect. The nurse then faces a trade-off between getting the procedure over and done with (to minimise the patient’s dread) and waiting for the cream to take full effect (to minimise the pain itself). According to an hypothesis of ‘diminished social dread’, the nurse would be archetypally predicted to wait for the cream to take effect, to diminish the intensity of pain, while the patient would prefer to get the procedure over and done with, to minimise their dread. Such scenarios are interesting subjects for a future applied study. Taken together, the findings of the above experiments suggest that people hold a powerful aversion to harming others that is focused on concrete markers of harm such as the intensity of pain, rather than more abstract effects such as dread, and decays steeply with social distance.

SOCIAL ALLOCATION OF PAIN

“To each according to his needs.”

- Louis Blanc, Plus de Girondins, (1851)

This chapter extends the study of social preferences for pain introduced in the preceding chapter, to ask how people might apportion a ‘budget’ of pain between themselves and others. Just as Chapter 7 considers issues of intertemporal equity in apportioning health, the current chapter considers interpersonal equity in health. From an applied perspective, the current chapter seeks to explore effects of health inequality on personal wellbeing and as such diverges from the prevailing theme of this thesis concerning prediction of health behaviour.

In 1980 the Department of Health and Social Security in the UK published an influential report overseen by Sir Douglas Black, which showed that health inequality had increased, rather than decreased, since the introduction of the NHS in 1948 (see A. Gray, 1982). Since the time of this eponymous ‘Black Report’ marked health inequalities have endured, for example between 2001 and 2003 for men aged 25-64, mortality from ischaemic heart disease was 2.9 times higher in social class 7 (routine occupations) than in social class 1 (higher managerial and professional occupations) (White, Edgar, & Siegler, 2007). Inequality tends to be viewed as unjust and undesirable, for example in a recent study of US citizens, 90% indicated that their society was more unequal than they would like it to be (Norton & Ariely, 2011). Here I explore possible hedonic and motivational effects of health inequality at an individual level, by observing how people respond to, and choose to reapportion, unequal allocations of pain between themselves and others.

9.1 INEQUITY AVERSION

Psychologists have studied the effects of social comparison on individual wellbeing for over half a century. In an early account, Adams (1965) proposed a theory of inequity, arguing that deprivation relative to others provokes feelings of injustice, which in turn lead to dissatisfaction. A corresponding thought process in a health context might be: “She and I have both been waiting to see the doctor for over two hours. She arrived after me but was seen first. That seems unfair (unjust). I am dissatisfied with my treatment here.” Under this view, unfairness is a subjective cost associated with having worse prospects than others, which engenders dissatisfaction.

More recently, economists have studied the value of fairness through behavioural experiments, introducing quantitative theories describing the cost of inequity (Fehr & Schmidt, 1999; Fehr & Schmidt, 2006; Forsythe, Horowitz, Savin, & Sefton, 1994; Kahneman, Knetsch, & Thaler, 1986). These theories have emerged from analysis of formal economic games, in which the rules of social exchange are tightly constrained and responses can be quantified in monetary terms. An archetypal example is a scenario referred to as the ‘Dictator Game’, in which one participant, the Dictator, must choose how to allocate an endowment (usually of money) between themselves and another participant, who is anonymous to them (Kahneman et al., 1986). Classically the game is played as a ‘one-shot, one-way’ interaction, with no reciprocity between the two players. Meta-analysis of many hundreds of such games reveals that on average Dictators donate close to a third of the endowment to the other player, despite no financial return for themselves (Engel, 2011). This behaviour has been taken to indicate that people have a basic preference for fairness (Fehr & Schmidt, 1999; Kahneman et al., 1986). In economic models, the latter is often approximated as cost associated with having either more or fewer resources than others, termed inequality aversion (Fehr & Schmidt, 1999; Fehr & Schmidt, 2006).

Psychological and economic methodologies offer complementary means to understand individual responses to health inequality. A specific question of interest from an economic perspective is whether people disvalue inequalities in health in a similar manner to monetary inequality, as revealed through their choices to distribute both health and money. For example, are people more sensitive to health inequality

than to wealth inequality? A route to studying this is to examine how people respond to allocations of either pain or money between themselves and others.

9.2 ARE PEOPLE INEQUITY AVERSE FOR PAIN?

Although many experimental studies have examined how humans redress monetary inequality, how humans respond to unequal allocations of pain between themselves and others is little studied. Commensurate with documented empathic responses to pain in others (Batson et al., 1988; Batson et al., 1983; Jackson et al., 2005; Singer et al., 2004), people will pay money to relieve others' pain (FeldmanHall et al., 2012) and will take on painful stimuli to relieve the apparent suffering of a confederate (Batson et al., 1983; Hein et al., 2011). Furthermore, a recent study implementing a dictator game with painful outcomes found that participants were particularly egalitarian, allocating on average 48% of the pain (time spent immersing one's hand in ice water) to themselves, while in a monetary dictator game the same participants were significantly less charitable, donating only 30% of the endowment on average (A. L. Davis et al., 2011). These findings might suggest that people are more inequality-averse for pain than money, in other words that people particularly dislike having either more or less pain than others. However an alternative possibility is that people are simply more altruistic for pain than for money, in the sense that relieving others' of their physical suffering carries greater motivational weight than does increasing others' wealth.

In support of the idea that pain evokes particularly altruistic responses, people will often take on pain to relieve even a less severe pain for another person (see Experiments 3A and 3B, Chapter 8), and will even pay more money to reduce others' pain than to reduce their own pain by an equivalent subjective amount (Crockett et al., 2014). The latter behaviour, referred to as 'hyperaltruism' (or in Chapter 8, negative social discounting), can be taken to imply that people in fact sometimes even assign greater disvalue to harming others than to harming themselves. In a pain dictator game, hyperaltruistic participants ought to take on more pain than they allocate the other player; indeed in the dictator game study described above, 30% of participants retained more than half the pain, while only 2% gave away more than half the money (A. L. Davis et al., 2011).

A key question emerges, namely how can a putative preference for equal allocations of pain be reconciled with hyperaltruism? Models of social utility which solely implement an heuristic penalty for inequality (Fehr & Schmidt, 1999; Fehr & Schmidt, 2006) cannot account for hyperaltruism. These models assume that people derive utility from the size of their own payoff, but incur a utility cost to having either more or fewer resources than others. Under the simplest of such models (Fehr & Schmidt, 1999), both sources of utility are assumed to be linear. Such a model predicts that in dictator games with an endowment of reward people should either retain all of an endowment for themselves (offering nothing to the other player) or split the endowment equally. This model cannot account for the finding that in dictator games many people choose other allocations that are neither maximally selfish nor maximally equal (Engel, 2011), of which hyperaltruistic allocations for pain are an example (A. L. Davis et al., 2011). A more nuanced account implements a non-linear penalty for inequality (Bolton & Ockenfels, 2000). This model allows for a preference in dictator games to offer the other player something less than half of the endowment, but is nevertheless unable to account for hyperaltruistic offers. Here I show that both hyperaltruism and equality-seeking can be accounted for within a social utility model in which the utility of one's own and others' payoffs are combined in a weighted sum (e.g. Charness & Rabin, 2002), but where each are non-linear in the size of the payoff (obey diminishing marginal utility for rewards and increasing marginal disutility for pain). Under this model, given that people have some degree of concern for others, inequality aversion derives from the more fundamental principle of diminishing marginal utility over individual payoffs.

The conventional dictator game setup is not ideally suited to disambiguating alternative models of social utility, since it provides information only about a person's preferred allocation, and does not allow for an estimate of how a person values different degrees of pain inequality (pertinent for example to whether inequality carries a linear or non-linear cost). To address this, I implemented a modified dictator game with mildly painful cutaneous electric shocks as the outcomes, where participants were given an opportunity to leave unchanged or adjust default allocations of 24 shocks between themselves and an anonymous recipient. Adjustments could be made by giving or taking fixed numbers of shocks. Three 'action frames' were implemented (see also List, 2007): a *Give* frame, a *Take* frame

and a *Give or Take* frame, with names corresponding to the permissible actions (Figures 9.1 and 9.2). By observing how frequently participants adjusted the default allocations at differing levels of inequality (see Figure 9.3), I mapped out a pattern of inequality aversion for pain, albeit coarsely.

I also aimed to test an hypothesis, supported by previous studies, that people are more altruistic when allocating pain, compared with money. To do so, in a subset of participants, I implemented an equivalent set of choices with monetary outcomes where, in two sub-conditions, these choices were framed as either gains or losses of money, where losses were described as debts to be repaid from an endowment. Allocations of monetary losses (debt) were compared with those of pain, since these conditions were matched for valence. It was predicted that participants would adjust the default allocations in the direction of equality for both pain and debt. Based on greater altruism for pain than money in standard dictator games, it was predicted that participants would allocate more debt than pain to the other participant. Finally, based on existing findings of hyperaltruism for pain it was also predicted that hyperaltruism would be more prevalent for pain than for debt.

9.3 DOES UNFAIR TREATMENT EXACERBATE PAIN?

A more ‘psychological’ question is: does perceived unfair treatment negatively impact upon a patient’s experience of pain or illness? Yoda and Kumakura (2007) examined the impact of perceived unfairness on satisfaction with care, characterizing subjective unfairness for medical encounters by asking respondents ($N=162$) to rate their agreement with statements such as “The nurses were more caring and attentive towards other patients than towards me”, with the finding that the effect of unfairness on overall patient satisfaction was comparable to that of perceived symptom improvement or physician competence. However it is also possible that perceived unfairness interacts with the perception of health outcomes themselves. In support of this idea, social rejection and physical pain have overlapping representations (DeWall & Baumeister, 2006; DeWall et al., 2010; Eisenberger, 2012; Eisenberger, Lieberman, & Williams, 2003; MacDonald & Leary, 2005), and perceived injustice following injury is associated with greater subsequent pain and disability (McParland & Eccleston, 2013; Sullivan et al., 2008) (Furthermore social

comparison is known to influence neural responses to reward, e.g. Fließbach et al., 2007; Takahashi et al., 2009; Tricomi, Rangel, Camerer, & O'Doherty, 2010). However prior to the current work no previous studies had examined how pain inequality influences pain perception.

An intuitive prior prediction is that favorable inequality should reduce pain perception, while unfavorable inequality should increase it. A further prediction is that pain resulting from the benevolent actions of another person is perceived as less severe than an equivalent unintentional pain, and *vice versa*. Two previous studies support this prediction, though these studies did not directly test an effect of inequality (K. Gray, 2012; K. Gray & Wegner, 2008). To test these hypotheses I assessed the impact of pain inequality, and perceived intentionality, on the perception of pain itself, by asking participants to rate identical pain outcomes received under differing social conditions. Selected allocations of pain from the dictator game were implemented for real, and attributed either to the play of chance or to the decisions of the dictator (shown in Figure 9.4), whilst participants were asked to rate the intensity of the painful stimuli.

9.4 EXPERIMENT 4: RESPONSES TO PAIN INEQUALITY

9.4.1 METHODS

9.4.1.1 Participants


78 participants (46 female) were recruited from the University College London Psychology Subject Pool. A power calculation based on a Student's *t*-distribution indicated that a sample size of 34 would be required to detect a medium sized effect (Cohen's $d=0.5$), while a sample size of 15 would be required to detect a large effect (Cohen's $d=0.8$) at 80% power (J. Cohen, 1992). I therefore aimed to elicit dictator game choices from at least 35 participants. The final sample size easily exceeded this total ($N=47$; 29 female, mean age 27.7 years, s.d. 8.5 years). A monetary questionnaire was issued at a later date to these 47 participants; here the sample size was determined by the number responding to this request ($N=25$; 16 female, mean age 29.6 years, s.d. 11.6 years). Of the 47 participants completing pain

dictator game games, 31 took part in the experiential component of the study. For this component, 31 participants were separately recruited to act as ‘responders’. Due to experimenter error, datasets from 6 of these 31 pairs were incomplete, resulting in full experiential data from 25 pairs ($N=50$; 30 female, mean age 27.8, s.d. 10.2). Dictators and responders did not actually meet each other, but each was briefed that they would be interacting with another participant in an adjacent testing room through an intranet link; they were informed that there was no opportunity for reciprocity. All participants gave full informed consent prior to the experiment. After the experiment participants were debriefed and given the opportunity to provide feedback. Experiments were carried out at the Wellcome Trust Centre for Neuroimaging, University College London. Participants were compensated for their participation at a rate of £10 per hour.

9.4.1.2 Painful Stimuli

Throughout the experiment each participant sat in front of a computer monitor while choice options were presented on-screen, and indicated their decisions using the keyboard. The presentation package COGENT 2000 (University College London) was used for stimulus presentation and response acquisition. Electrical stimuli were delivered using a DS5 Digitimer (Letchworth Garden City, England) constant current stimulator. A single “shock” consisted of a train of seven 10ms square-wave pulses delivered at 40ms intervals. When multiple shocks were delivered these were spaced 250ms apart. After providing consent, participants underwent a standardised thresholding procedure, which allowed selection of current levels (in mA) for use during the experiment, which corresponded to visual analogue scale pain ratings of 5/10 and 8/10 respectively (see General Experimental Methods, Chapter 5).

Shocks: *Give* Frame



Shocks: *Take* Frame




Figure 9.1 Experiment 4: Modified pain dictator game protocol, *Give* and *Take* frames. In each of three possible action frames (shown here for *Give* and *Take* frames) the Dictator was given the opportunity to alter or leave unchanged default allocations of 24 brief, moderately painful, cutaneous electric shocks between themselves and another participant, termed the Responder. In a *Give* frame the Dictator could choose to either give six of their allocated shocks to the Responder, or leave the allocation unchanged. In a *Take* frame the Dictator could choose to either take on six of the Responder's shocks, or leave the allocation unchanged.




Figure 9.2 Experiment 4: Modified pain dictator game protocol, *Give or Take* frame. In each of three possible action frames (shown here for the *Give or Take* frame) the Dictator was given the opportunity to alter or leave unchanged default allocations of 24 brief, moderately painful, cutaneous electric shocks between themselves and another participant, termed the Responder. In a *Give or Take* frame both giving and taking options were available, as well as the option to leave the default unchanged.

9.4.1.3 Dictator Game Choices

Dictator game choices were elicited through an online interface (Qualtrics.com; Provo, UT); on each the computer suggested a default allocation of 24 shocks between dictator and receiver, and the dictator chose whether to alter this allocation, or leave it unchanged (Figures 9.1, 9.2 and 9.3). Dictators were informed that the subjective intensity level of the shocks would vary between 4/10 and 8/10 on a VAS, but that the subjective level for themselves and the other player would always be the same on a given trial. In a *Give or Take* condition default allocations (number of shocks for the dictator listed first) were [18-6], [12-12] or [6-18]. The dictator chose whether to give 6 shocks to the responder from their own allocation, to take 6 shocks from the responder's allocation (which would be added to their own allocation) or to leave the allocation unchanged. In a *Give* condition default allocations were [24-0], [18-6], [12-12] or [6-18]. The dictator chose whether to give 6 shocks of their shocks to the responder or to leave the default allocation unchanged.

In a *Take* condition default allocations were [18-6], [12-12], [6-18] or [0-24]. The dictator chose whether to take 6 shocks from the responder's allocation, or to leave the allocation unchanged. Notably equivalent outcomes were attainable in all frames. Each choice option was presented twice, creating 22 choices in total. Dictators were informed that a selection of their choices would be played for real in the experiential part of the experiment, should they consent to take part. The order of conditions and trials within conditions was randomised.

In a follow-up experiment I elicited equivalent choices with money in a subset of the same participants, where the numbers of shocks were replaced with quantities of money in pounds sterling. In a Money Gain condition participants were told that they had won a prize of £24 between themselves and the other participant. On each choice the computer suggested a default allocation for the money and the dictator chose to either alter this allocation, or leave it unchanged. In a Money Loss condition participants were informed that they and the other participant had *each* won a prize of £24, however that they must also pay a debt of £24 between them both. Net allocations in the Money Loss condition were equivalent to the Money Gain condition. Participants were informed that some choices from the entire set of choices made by all participants would be randomly selected to be paid out for real.

Default	Give 6	Default	Take 6
[6- 18]* -> [0- 24]		[6- 18] <- [0- 24]*	
[12- 12]* -> [6- 18]		[12- 12] <- [6- 18]*	
[18- 6]* -> [12- 12]		[18- 6] <- [12- 12]*	
[24- 0]* -> [18- 6]		[24- 0] <- [18- 6]*	

Take 6	Default	Give 6
	[12- 12] <- [6- 18]* -> [0- 24]	
	[18- 6] <- [12- 12]* -> [6- 18]	
	[24- 0] <- [18- 6]* -> [12- 12]	

Figure 9.3 Experiment 4: Modified pain dictator game, choice options within each frame. Allocations are in the format [shocks for Dictator – shocks for Responder]. The number of shocks allocated to the Responder, highlighted in bold, was termed the ‘offer’. Asterisks indicate the default allocation (which would result if participants chose ‘leave unchanged’). We also implemented equivalent choices for money, in which shocks were simply replaced with sums of money in pounds sterling (money gain condition) or with amounts of debt in pounds sterling (money loss condition). Since debt in the money loss condition was subtracted from an endowment of £24, the net allocations under the two conditions were equivalent. In both forms of the task a selection of choices were implemented.

9.4.1.4 Experiential Phase

To examine the effects of inequality and intentionality on the subjective experience of pain, selected shock allocations were played for real in the laboratory. This phase consisted of three types of trial (Figure 9.4). On ‘Shocks for You: Chance’ (the *Non-Social-Chance* condition) trials participants were informed that they would receive either 6, 12 or 18 shocks, which would be randomly assigned to them by the computer. On ‘Split Shocks: Chance’ trials (the *Social-Chance* condition) participants were informed that 24 shocks would be split between themselves and the other

participant, according to one of three possible allocations [6-18], [12-12] or [18-6], which would be randomly selected by the computer. Finally a third type of trial (forming the *Social-Intentioned* condition) was referred to as ‘Split Shocks: You Decide’ for the dictator and ‘Split Shocks: Other Player Decides’ for the responder. The responder was informed that the dictator had chosen allocations of shocks in advance, and that on each trial of this type one of the dictator’s chosen allocations would be implemented. Dictators were informed that on each trial of this type one of their previously chosen allocations would be implemented, and that each allocation may be implemented more than once.

Both players were told that the intensity of the shocks could vary from trial to trial and that whilst they would not be told the intensity, the subjective level of the shocks for themselves and their partner would always be the same (both players would be told the *number* of shocks they were due to receive). All dictators made several choices resulting in the three allocations 6-18, 12-12 and 18-6 shocks. In fact, so as to control for context effects in the perception of pain, these three outcomes were selected with equal probability in both the *Social-Intentioned* and *Social-Chance* conditions. Similarly in the *Non-Social* condition, participants received 6, 12 or 18 shock with equal probability. Although participants were not informed of this sampling process, they were otherwise given faithful information regarding the setup. The three trial types (conditions) were interleaved in 6 short blocks of 12 trials each. Within each condition, the three possible allocations of shocks were repeated twice at each of two different intensity levels (5/10 and 8/10). The order of conditions and of trials within each condition was randomly counterbalanced. The outcome allocation on each trial was indicated by a pie chart. In *Non-Social* and *Social-Chance* conditions, this was preceded by a “roulette-wheel” effect, emphasizing that the computer was randomly selecting the shocks. After the outcome was displayed shocks were administered to both players accordingly.




Figure 9.4. Experiential phase of Experiment 4: Trial structure for receivers. **A** On *Shocks for you: Chance* trials participants were informed that the outcomes would only be relevant to them, and did not involve the other participant. The computer then selected an allocation of shocks at random. The participant received the stated number of shocks and was then asked to rate the intensity of the shocks on a visual analogue scale. We subsequently elicited participants' willingness to pay to avoid three further repeats of the same outcome (not shown in this figure). **B** *Split Shocks: Chance* trials proceeded in the same manner, though in this case the computer 'randomly' chose how to allocate 24 shocks between the two participants. **C** The trial structure was identical in *Split Shocks: Partner Decides* trials, though on these trials the receiver was informed that an allocation chosen by their partner would be selected and played for real.

In the first session, after shock delivery each player rated the intensity of the shocks on a visual analogue scale. Immediately afterwards participants' willingness-to-pay to avoid the shocks was obtained by means of a simulated "auction" (not shown in Figure 9.4) (G. M. Becker et al., 1964). Participants were endowed with 40 pence on each trial and were asked to indicate their willingness to pay from this endowment to avoid three repeats of their allocated shocks. On one in ten randomly selected trials their bid was compared to a market price, sampled from a uniform distribution in the range of 0 – 40 pence. If the participant's bid exceeded the market price, they would avoid repeated instances of the allocated shocks, and would pay the amount they bid; otherwise they would receive three repeats of their allocation of shocks and would pay nothing. Participants took home as a "bonus" any money which they did not spend on these auction trials. Participants were informed that their bids only affected their own shocks, and did not alter the other participant's shocks. In a second session, participants rated the fairness of the outcomes, the extent to which they perceived the other player to be responsible for the allocation and the extent to which they felt inclined to punish the other player for the allocation, each on a VAS ranging from 0 to 100.

9.4.1.5 Model Specification and Fitting

I tested alternative social valuation models for the choices made by the dictators. In particular I sought to test a simple form of social welfare preference that is capable of accounting for hyperaltruism. Under this model, termed Social Discounting, a person's social utility is given by a weighted sum of their own utility and that of others, shown here for the two-person case:

$$U^{SD}(s_i, s_j) = u_i(s_i) + \kappa \cdot u_j(s_j) \quad (9.1.1)$$

Here $U^{SD}(s_i, s_j)$ is a social utility function governing the utility to player i of an allocation with payoff s_i to player i and s_j to player j , $u(s)$ is a utility function over individual payoffs and κ is a social discounting parameter which governs the contribution of player j 's payoff. Note that this function is equivalent to the models of

social discounting proposed in Chapter 8 of this thesis. If $u(s)$ is concave for both gains and losses (decreasing marginal utility for gains and increasing marginal utility for losses/harms) then, for $\kappa > 0$ the above function predicts a tendency to spread the payoff across individuals. (Note also the direct analogy here to concave utility producing spreading of resources over time in Chapter 7). This has an intuitive interpretation as a motive to give to those most in need. Owing to concave utility, individuals with lower payoffs will value the same increase in their payoff more than individuals with higher payoffs. To fit the model, the utility functions over individual payoffs were assumed to be identical and concave, such that $u_i(s) \equiv u_j(s)$, both given by the quadratic function $u(s) = s - bs^2$, where $b = 0.01$.

I tested this model against three existing social utility models: Fehr-Schmidt utility (Fehr & Schmidt, 1999), Charness-Rabin social preferences (Charness & Rabin, 2002), and the Equity, Reciprocity and Competition model (ERC) (Bolton & Ockenfels, 2000). The Fehr-Schmidt model penalises unequal allocations, with separate penalties for advantageous inequality and disadvantageous inequality:

$$U^{FS}(s_i, s_j) = s_i - \alpha \max\{s_j - s_i, 0\} - \beta \max\{s_i - s_j, 0\} \quad (9.1.2)$$

$U^{FS}(s_i, s_j)$ is a social utility function governing the utility to player i of an allocation with payoff s_i to player i and s_j to player j , and α parameterises the linear effect of disadvantageous inequity, β the effect of advantageous inequity. As is conventional, parameter bounds were set such that $\alpha \geq 0$ and $0 \leq \beta \leq 1$.

The ERC model also implements a penalty for unequal payoffs, though in a non-linear manner, achieved by penalizing squared deviations from an even distribution of payoffs:

$$U^{ERC}(s_i, s_j) = s_i \cdot a + \frac{b}{2} \cdot (\sigma_i - 0.5)^2 \quad (9.1.3)$$

where σ_i is player i 's share of the total payoff:

$$\sigma_i = \frac{s_i}{s_i + s_j}$$

(9.1.4)

The parameter a weights self-oriented preferences, while b determines the contribution of inequality aversion; both parameters were bounded to be zero or positive.

Charness-Rabin social preferences assume that people behave as if to maximise a flexible weighted average of the payoffs of all players, where the weightings reflect the existence of advantageous or disadvantageous inequity as well as the effect of intentionality. For the two-person case:

$$U^{CR}(s_i, s_j) = (\rho r + \sigma s + \theta q) \cdot s_j + (1 - \rho r - \sigma s - \theta q) \cdot s_i \quad (9.1.5)$$

where

$$\begin{aligned} r &= \begin{cases} 1 & \text{if } s_i > s_j \\ 0 & \text{otherwise} \end{cases} \\ s &= \begin{cases} 1 & \text{if } s_j > s_i \\ 0 & \text{otherwise} \end{cases} \\ q &= \begin{cases} -1 & \text{if player } j \text{ has misbehaved} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (9.1.6)$$

In this case q was set to zero, making θ redundant; the remaining weighting factors ρ and σ were bounded between zero and one.

Since the current experimental choices required participants to either accept a suggested allocation, or alter it, I expected to observe a *status quo* bias – i.e. a tendency to accept the suggested allocation (Kahneman et al., 1991). All four models above were therefore modified for the current choices by adding a *status quo* bias, simply a fixed cost, x , associated with changing the allocation, such that for each model:

$$U(s_i, s_j|a) = \begin{cases} U(s_i, s_j) - x & \text{if } a \neq 0 \\ U(s_i, s_j) & \text{otherwise} \end{cases} \quad (9.1.7)$$

where $a = 0$ denotes accepting the default.

Model fitting was by maximum likelihood estimation; a standard logistic sigmoid (softmax function) was used to transform the social utilities of the allocations

resulting from different actions into corresponding probability of choosing a particular action. Constrained non-linear optimisation was used in Matlab (Mathworks, MA, USA), to find model parameters which minimised the negative log likelihood of participants' choices. To avoid convergence on local minima, the optimiser was run within a random multi-started overlay with 1000 starting points drawn from a uniform distribution within the parameter bounds (RMsearch). Models were compared using the Bayesian Information Criterion (BIC) (see General Experimental Methods, Chapter 5).

9.4.1.6 Subjective Questionnaire Measures

Participants also completed a 20-item Paranoia Scale (Fenigstein & Venable, 1992), a 40-item Empathy Quotient questionnaire (Baron-Cohen & Wheelwright, 2004), a Justice Sensitivity Scale (Victim and Perpetrator sensitivity) (Schmitt, Gollwitzer, Maes, & Arbach, 2005) and a Vengeance Scale (Stuckless & Goranson, 1992). We hypothesised that the tendency amongst Receivers to attribute responsibility for their shocks to the choices of the Dictator would correlate positively with their paranoia score, and that their rated inclination to punish the Dictator for their actions would correlate with justice sensitivity and vengeance scores. We further hypothesised that more altruistic offers amongst Dictators would correlate positively with the empathy quotient.

9.5 RESULTS

9.5.1 DISTRIBUTION OF MEAN DICTATOR GAME OFFERS

The distribution of offers across Dictators (taking the mean resulting offer over all three action frames for each participant) is displayed in Figure 9.5 for pain (Figure 9.5a) and money (Figure 9.5b for losses, 9.5c for gains). Across participants the mean offer of shocks was not significantly different from 12 shocks ($N = 47$; mean = 12.16, permutation test, $p > 0.25$, see Methods), indicating a preference for equal allocations of pain at the group level. By contrast the mean offer of money

(averaged across gain and loss frames) was significantly below £12 ($N = 25$; mean = £11.39, $p=0.011$, two-tailed permutation test), indicating a degree of self-oriented behaviour for monetary outcomes. To compare the degree of self-oriented behaviour for money and shocks I focussed on comparing offers in pain and money loss conditions, since these were matched for valence. As predicted, offers of shocks were significantly lower than offers of debt (mean shocks offered in 25 subjects performing money choices = 11.98, mean debt offered = £12.98; $p=0.021$, two-tailed permutation test, uncorrected; $p=0.042$ Bonferroni corrected for two possible comparisons). Surprisingly there was no significant difference between the mean offers in pain and money gain conditions (mean money retained by dictator = £12.24; $p>0.250$), nor was there a difference between money gain and money loss conditions ($p>0.250$). Notably, given the restricted nature of the choices, this analysis based on mean offers might not be the most sensitive for detecting changes in altruism between modalities.




Figure 9.5 Distribution of mean resulting dictator game offers in Experiment 4. Histograms summarizing the distribution of mean offers across all choices made by each participant. In each case the horizontal axis represents an increasing measure of self-oriented behaviour. **a** Pain dictator game, offers of shocks by all 47 participants, **b** offers of shocks made by 25 participants who also completed monetary choices, **c** money dictator game, offers of debt in the loss frame, **d**, money dictator game, money retained by the dictator in the gain frame. Offers are out of a total endowment of 24 shocks (**a**, **b**) or £24 (**c**, **d**).

Nevertheless, taken together these results indicate that participants were most altruistic when apportioning pain, and least altruistic when apportioning debt. I perform additional analysis on the raw choice frequencies below, by comparing responses to advantageous and disadvantageous inequality for pain and debt.

For pain allocations, 15 out of 47 participants (32%) retained more than half the shocks across all choices, indicating hyperaltruistic behaviour. By contrast, in the money loss frame, only 3 out of 25 participants (12%) assigned more than half of the money to the receiver (one-sided Fisher's exact test $p=0.054$).

9.5.2 INEQUALITY AVERSION FOR PAIN AND MONEY

To examine the pattern of social preferences in more detail, I analyzed the frequency with which participants altered the default allocations as inequality changed. In the *Take* frame participants were increasingly likely to take a fixed number of shocks from the other participant as the default number of shocks allocated to the other participant increased (Figure 9.6A). Similarly in the *Give* frame, participants became less likely to give shocks to the other participant as the baseline number of shocks allocated to the other participant increased (Figure 9.6B). An equivalent effect was seen in the *Give or Take* frame (Figures 9.6c and d). In the monetary domain a similar effect was also observed (Figure 9.7). Importantly, this particular pattern is in keeping with an increasing marginal effect of inequality, since the propensity to correct inequality by the same degree, 6 shocks, increased as the baseline level of inequality increased (from 6-18 to 0-24 in the *Take* frame, and from 18-6 to 24-0 in the *Give* frame).




Figure 9.6 Dictator behaviour for pain as a function of default allocations in Experiment 4. For each plot a single data point represents a mean choice probability across 47 participants; error bars are 95% confidence intervals around the mean, generated by re-sampling the data 100,000 times with replacement. **A** Probability, in the *Take* frame, that dictators chose to take 6 shocks from the receiver's allocation (decreasing the resulting offer by 6 shocks), as opposed to leaving the allocation unchanged, is plotted as a function of the initial allocation of shocks to the receiver (the default offer). Dictators were increasingly likely to take shocks from the receiver as the initial allocation to the receiver increased, consistent with an increasing marginal cost to having fewer shocks than the receiver. **B** Probability, in the *Give* frame, that dictators chose to give 6 of their allocated shocks to the receiver (increasing the resulting offer by 6 shocks) is plotted as a function of the initial allocation of shocks to the receiver (the default offer). Dictators were increasingly likely to give shocks to the receiver as the initial allocation to the receiver decreased, consistent with there being an increasing marginal cost to having more shocks than the receiver. Similarly, in the *Give or Take* frame, the probability that dictators chose to take 6 shocks from the receiver's allocation increased as the initial allocation to the receiver increased (**C**), and the probability that dictators chose to give 6 of their shocks to the receiver decreased as the initial allocation to the receiver increased (**D**).

9.5.3 SELF-ORIENTED PREFERENCES FOR MONEY BUT NOT PAIN

To further compare social preferences for debt and pain at the group level, I separated choices (in the *Give* and *Take* action frames) that entailed disadvantageous inequality (24-0 vs 18-6, 18-6 vs 12-12) from those that entailed advantageous inequality (6-18 vs 0-24, 12-12 vs 6-18). Here, choosing the more equal option in ‘disadvantageous’ choices more frequently than in the symmetrical ‘advantageous’ choices indicates a degree of self-oriented behaviour. This analysis creates a three-way factorial design with two outcome modalities (debt vs pain), two action frames (*Give* vs *Take*) and two directions of inequality (advantageous vs disadvantageous). A main effect of modality would indicate greater or lesser inequality aversion for pain compared with debt irrespective of the direction of inequality. By contrast a modality x inequality-direction interaction would indicate more or less altruistic behaviour for pain.

Three-way repeated measures ANOVA on the probability of choosing the more equal option (termed equality-seeking) for the 25 participants who made choices for both modalities revealed a significant main effect of inequality-direction ($F(1,24)=4.616$, $p=0.042$), driven by some degree of self-oriented behaviour irrespective of modality or action frame (estimated marginal mean equality-seeking = 0.745 for advantageous choices, 0.872 for disadvantageous choices). There was neither a significant main effect of modality ($F(1,24)=0.183$, $p>0.250$), nor a significant main effect of action ($F(1,24)=2.348$, $p=0.139$). Importantly, as predicted, there was a significant interaction between modality and inequality-direction ($F(1,24)=6.233$, $p=0.020$). Follow-up non-parametric tests (by randomly reshuffling condition assignment) indicated that this was driven by greater self-oriented behaviour for debt than for pain. For example, for debt participants were on average 26% more likely to choose the more equal option in disadvantageous choices than in advantageous choices (mean difference = 0.26, two-tailed $p<0.001$, $N=25$). By contrast, for pain there was no significant difference in equality seeking between disadvantageous and advantageous choices (mean difference <0.01 , two-tailed $p>0.250$, $N=25$). Furthermore, the observed asymmetry for debt was significantly greater than that for shocks (two-tailed $p=0.018$, $N=25$).

The above analysis also revealed a significant interaction between action frame and inequality direction ($F(1,24)=14.605$, $p=0.001$). Follow-up permutation testing indicated that this was driven by more self-oriented behaviour in the *Take* frame than in the *Give* frame (difference in mean inequality-direction effect = 0.27, two-tailed $p<0.001$). Since default offers of debt or pain to the other participant were higher in the *Take* frame, this effect is consistent with a *status quo* bias in dictator behaviour. As described below a *status quo* bias was also evident in the mean offers across all three outcome modalities. The remaining interactions in the three-way ANOVA (between modality and action frame, and the three-way interaction) were not significant ($F(1,24)=0.660$, $p=0.424$ and $F(1,24)=0.604$, $p=0.444$ respectively).




Figure 9.7 Dictator behaviour for money as a function of default allocations in Experiment 4. Each data point represents a mean choice probability across 25 participants; error bars are 95% confidence intervals, generated by re-sampling with replacement. **A** Probability (in the *Take* frame) that dictators chose to *take* £6 from the receiver, as opposed to leaving the offer unchanged, plotted as a function of the initial net (default) offer, indicated by the symbol $\blacklozenge$. Also plotted is the probability (in the *Give* frame) that dictators chose to *give* £6 of debt to the receiver, indicated by the symbol $\blacktriangle$. In both cases, dictators were increasingly likely to reduce the offer as the initial offer increased, consistent with there being an increasing marginal cost to having less money than the receiver. **B** Probability (in the *Give* frame) that dictators chose to give £6 to the receiver, as opposed to leaving the offer unchanged, plotted as a function of the net default offer, indicated by the symbol $\blacksquare$. Also plotted is the probability (in the *Take* frame) that dictators chose to take £6 of debt from the receiver, indicated by the symbol $\bullet$. In both cases, dictators were increasingly likely to increase the offer as the initial offer decreased, consistent with an increasing marginal cost to having more money than the receiver. Similarly, in the *Give or Take* frame, the probability that dictators chose to decrease the net offer by £6, either by taking money ($\blacklozenge$) or giving debt ($\blacktriangle$) increased as the initial offer increased (**C**), and the probability that dictators chose to increase the net offer by £6, either by giving money ($\blacksquare$) or taking debt ($\bullet$) decreased as the initial offer increased (**D**). A tendency towards self-oriented choices is evident from the greater propensity to decrease rather than increase offers, for the same degree of inequality.

9.5.4 STATUS QUO BIAS

Owing to the nature of the choices in this experiment, in which an initial allocation could either be accepted or altered, I had expected to observe a *status quo* bias, reflecting an overall reluctance to change the current state of affairs (Baron & Ritov, 2004; Kahneman et al., 1991; Spranca, Minsk, & Baron, 1991). A *status quo* bias predicts higher offers in the *Take* relative to the *Give* frame, since default offers are higher in the *Take* frame. As displayed in Figure 9.8, mean offers in the *Take* frame were indeed significantly higher than those in the *Give* frame, for both pain ($p < 0.001$, two-tailed permutation test), money gains ($p < 0.001$) and money losses ($p < 0.001$).




Figure 9.8 A *status quo* bias in dictator behaviour. **A** Mean offers of shocks across all dictators (vertical axis; $N = 47$), for each action frame (horizontal axis category). Solid black horizontal bars indicate the mean *default* offer under each frame. Upper dashed horizontal bars show the mean offers that would result from maximally selfish behaviour, lower dashed horizontal bars the mean offers that would result from maximally selfless behaviour. Mean offers were higher in the *Take* frame than in the *Give* frame, consistent with a *status quo* bias. **B** Mean net offers of money across all dictators (vertical axis; $N = 25$) under *Give* (G), *Take* (T) and *Give or Take* (G/T) action frames (horizontal axis category, individual bars), under both Gain and Loss (debt) frames (horizontal axis groupings). Here, taking money is equivalent to giving debt, and *vice versa*. A *status quo* effect is evident for money gains and losses (debt), whereby net offers were higher when the net default offer was high (*Take* frame for money gains and *Give* frame for money losses), compared with when the net default offer was low (*Give* frame for money gains and *Take* frame for money losses). Error bars indicate 95% confidence intervals generated by resampling the data 100,000 times with replacement. Asterisks indicate p values associated with two-tailed permutation tests; *** $p < 0.001$.

9.5.5 MODELLING DICTATOR BEHAVIOUR

Predictions of the Social Discounting model in a Dictator Game are displayed in Figure 9.9, for the special case where u_i and u_j are concave and identical (the same form of utility function for payoffs to self and other). Where $\kappa < 1$, player i will assign more weight to their own payoff than those of the other player, and therefore is motivated to allocate more than half of the benefit to themselves. For $\kappa = 1$, players can maximise their utility by achieving equal payoffs ($s_i = s_j$). For the opposite case of $\kappa > 1$, corresponding to hyperaltruism, player i is motivated to allocate more than half of the benefit to the other player.




Figure 9.9 Social utility functions based on a Social Discounting model. The total utility predicted by a social discounting model of the form shown in Equation 1 (vertical axis, arbitrary units), is plotted as a function of the offer to the receiver (s_j) in a dictator game with an endowment of 24 units ($s_i = 24 - s_j$). Two possible scenarios are shown, either a decreasing offer of painful shocks (horizontal axis, upper scale) or an increasing offer of money (horizontal axis, lower scale). The four lines represent different settings of the social discount parameter, κ , as labeled. In each case, $u_i(s) \equiv u_j(s)$, both given by the quadratic function $u(s) = s - bs^2$, where $b = 0.01$. For $\kappa = 1$ the function has a maximum when allocations to both participants are equal ($s_i = s_j = 12$), corresponding to symmetric inequality aversion. $\kappa < 1$ generates a preference for assigning more than half the benefit to oneself, while $\kappa > 1$ (hyperaltruism) generates a preference for assigning more than half the benefit to the receiver. For each parameter setting, moving away from one's preferred allocation carries a non-linear (increasing marginal) cost.

Of the models tested, only the Social Discounting model can theoretically incorporate both hyperaltruism and increasing marginal inequality aversion, two key patterns observed in the data for pain. With the standard parameter bounds, Fehr-Schmidt utility accounts for neither finding, ERC accounts for non-linear inequality aversion, but not hyperaltruism and Charness-Rabin preferences are capable of incorporating hyperaltruism, but do not incorporate non-linear inequality aversion. I therefore hypothesised that the Social Discounting model would provide the best fit to the observed data. Since the behavioural data are consistent with a *status quo* bias, I modified each model to incorporate this effect, by assigning a cost to altering the *status quo*. (See Methods for detailed model specifications).

Consistent with predictions, the Social Discount model, with the addition of a *status quo* bias, was unequivocally the best fitting model at the group-level (highest maximum likelihood and lowest BIC) for both pain (Figure 9.10A; BIC difference of 124 over the next best model, in this case Charness-Rabin social preferences with a *status quo* bias) and money (Figure 9.11A; BIC difference of 152 over the next best model, in this case ERC with a *status quo* bias). The choice probabilities predicted by the model, formed by taking the mean of individual choice probabilities across participants, are overlaid with the observed data for pain and money in Figures 9.10B and 9.11B respectively. For pain, mean $\kappa=1.17$ (95% CI by resampling = 0.95-1.44; $N=47$), indicating a non-significant trend at the group level towards assigning greater weight to others' pain, i.e. hyperaltruism, in keeping with the behavioural findings above. For money, mean $\kappa=0.81$ (95% CI by resampling = 0.64-0.96, one-sample $p=0.0162$; $N=25$), indicating significantly self-oriented preferences. Also in keeping with the preceding model-free analyses, mean κ for money was significantly lower than that for pain ($p=0.0158$, two-tailed permutation test), supporting the conclusion of more altruistic choices for pain than for money.




Figure 9.10 Model Fits to pain dictator game choices in Experiment 4. **A** BIC values for the alternative models, either with *status quo* bias (light blue bars, right bar of each pair) or without (dark blue bars, left bar of each pair). The social discounting model provided the most parsimonious fit, indicated by the grey dot. **B** Maximum likelihood choice probabilities of the social discounting model, formed by taking the mean of the predicted choice probabilities across participants, are indicated by the symbol $\bullet$, and are overlaid with the observed choice probabilities as displayed in Figure 4, indicated by $\blacktriangle$ for choices to take on pain and $\bullet$ for choices to give pain. A single data point represents a mean choice probability across 47 participants and error bars are 95% confidence intervals around the mean, generated by re-sampling the data 100,000 times with replacement.




Figure 9.11 Model fits to money dictator game choices in Experiment 4. **A** BIC values for the alternative models, either with *status quo* bias (light green bars, right bar of each pair) or without (dark green bars, left bar of each pair). The social discounting model provided the most parsimonious fit, indicated by the grey dot. **B** Maximum likelihood choice probabilities of the social discounting model, formed by taking the mean of the predicted choice probabilities across participants, are indicated by the symbol ●, and are overlaid with the observed choice probabilities, indicated by ◆ for choices to take money, ▲ for choices to give debt, ■ for choices to give money and ● for choices to take on debt. A single data point represents a mean choice probability across 25 participants and error bars are 95% confidence intervals around the mean, generated by re-sampling the data 100,000 times with replacement.

9.5.6 EVALUATION OF EXPERIENCED PAIN ALLOCATIONS.

By generating equivalent outcomes framed as either *Non-Social* (no information given about the other player's allocation) or *Social* (players were informed of the other's allocation), where the latter can be framed as resulting either from the choices made by the dictator (*Social-Intentioned* condition) or from the play of chance (*Social-Chance* condition), I examined the effects of inequality and perceived intentionality on the perception of pain. A prior prediction was that advantageous inequality would reduce and disadvantageous inequality increase pain perception. I also predicted that perceived intentionality, in the *Social-Intentioned* condition, would amplify this effect, resulting in an effect of intentioned advantageous inequality to reduce pain perception, termed a 'kindness effect' and an effect of intentioned disadvantageous inequality to increase it, termed a 'meanness effect'.

Full ratings data was available from 50 participants, 25 dictators and 25 responders. It was expected that receiving more shocks at a given current level would increase the rated intensity of the pain, owing to sensitisation from one shock to the next. Indeed, a repeated measures ANOVA on intensity ratings across *Non-Social* with *Social-Chance* conditions across all participants revealed a significant main effect of number-of-shocks ($F(2,98) = 4.238$, $p = 0.017$; partial $\eta^2 = 0.08$). I predicted that priming inequality in the *Social-Chance* condition would increase the slope of this relationship, manifest as an interaction between condition and number of shocks. In fact there was found no evidence for the predicted number-of-shocks by condition interaction ($F(2, 98) = 0.045$, $p = 0.956$; partial $\eta^2 = 0.001$), and no significant main effect of condition ($F(1,49) = 0.297$, $p = 0.588$; partial $\eta^2 = 0.006$), suggesting that inequality *per se* did not influence pain ratings (Figure 9.12). However, comparison of intensity ratings from responders ($N = 25$) in *Social-Chance* and *Social-Intentioned* conditions did show the predicted number-of-shocks by condition interaction ($F(2, 48) = 11.657$, $p < 0.001$, Greenhouse-Geisser corrected; partial $\eta^2 = 0.327$), as well as a main effect of number-of-shocks ($F(2,48) = 8.026$, $p = 0.004$, Greenhouse-Geisser corrected; partial $\eta^2 = 0.251$) and no significant main effect of condition ($F(1,24) = 1.218$, $p = 0.281$; partial $\eta^2 = 0.048$), indicating that perceiving the other participant as responsible for shock allocations significantly influenced pain perception. Pairwise

comparisons revealed this effect was driven by lower intensity ratings in the *Social-Intentioned* condition for 6 shocks (mean rating difference = 0.645 points on VAS, 95% CI of the difference [0.15 1.13]; $t(24) = 2.72$, $p = 0.012$, two-tailed paired test; Cohen's $d = 0.49$) whilst intensity ratings for 18 shocks did not differ significantly between the two conditions (mean rating difference = 0.046 points on VAS, 95% CI of the difference [-0.22 0.32]; $t(24) = 0.351$, $p > 0.250$; Cohen's $d = 0.03$), indicating a significant 'kindness effect', but no significant 'meanness effect' (Figure 9.12).




Figure 9.12. Rated pain intensity as a function of social context in Experiment 4. Mean pain intensity ratings across participants on a 10-point visual analogue scale (VAS) for the Non-Social condition (all participants, $N=50$, $\blacksquare$) the Social-Chance condition (all participants, $N=50$, $\bullet$) and the Social-Intentioned condition (split into dictators, indicated by $\blacklozenge$ and responders, indicated by $\blacktriangle$, $N=25$ in each group). Data are modelled as a concave psychophysical function of number of shocks (dashed and dotted lines, see 9.3 Discussion). Error bars represent one standard error above and below the mean. In the Social-Intentioned condition, when dictators behaved altruistically, allocating only 6 out of a possible 24 shocks to the receiver, this significantly reduced pain perception relative to the other two conditions.

In the first session, after shock delivery each player rated the intensity of the shocks on a visual analogue scale. Immediately afterwards participants' willingness-to-pay (WTP) to avoid three repeats of the shocks was obtained by means of a first-price "auction". A similar set of effects as observed for subjective ratings was predicted for WTP. However, whilst there was a significant main effect of number-of-shocks on WTP ($F(2,48) = 15.89, p < 0.001$, Greenhouse-Geisser corrected) there was neither a significant main effect of condition ($F(2,48) = 0.202, p = 0.756$) nor a condition by number-of-shocks interaction ($F(4,96) = 0.823, p = 0.513$), suggesting that the motivational value of shocks was not influenced by the social context.

9.5.7 RATINGS OF FAIRNESS, RESPONSIBILITY AND INCLINATION TO PUNISH

In a separate experimental block, participants also made ratings in percentage terms of how much they perceived the other participant as responsible for the allocation, how fair they perceived the allocation to be and the extent to which they felt inclined to punish the other participant. For responsibility ratings made by responders ($N=25$) there was a large significant main effect of condition ($F(2,48) = 157.6, p < 0.001$; partial $\eta^2 = 0.868$), indicating a high degree of credibility in the experimental manipulation. As expected participants rated receiving 18 shocks as less fair than receiving either 12 or 6 shocks, whilst the latter two outcomes were rated as equally fair when collapsing across conditions (estimated marginal mean rating for 18 shocks = 30.8%, 95% CI [21.5-40.2], 12 shocks = 63.9% [54.2-73.6], 6 shocks = 61.3% [52.5 – 70.2]). For punishment ratings there were main effects of number-of-shocks ($F(2,48) = 11.03, p < 0.001$; partial $\eta^2 = 0.315$) and condition ($F(2,48) = 14.2, p < 0.001$; partial $\eta^2 = 0.372$), indicating responders felt more inclined to punish the dictator when they received higher numbers of shocks, irrespective of their causation, and were overall more inclined to punish the dictator in the Social-Intentioned condition, where they believed the dictator to be responsible for the allocations. As expected, there was also a significant number-of-shocks x condition interaction ($F(4,96) = 7.65, p < 0.001$; partial $\eta^2 = 0.242$), driven by a steeper effect of number of shocks in the Social-Intentioned condition, such that responders were most inclined to punish intentional high allocations of shocks. These results serve as a credibility check on the experimental manipulation.

Survey measures were compared against observed behaviour with painful outcomes (mean offers made by dictators, and ratings of pain intensity, fairness, responsibility and inclination to punish by the responders) by means of linear regression. There was no significant relationship between any of the survey instruments used and either the mean offers made by dictators or the ratings made by responders.

9.6 DISCUSSION

Marked inequality in health continues to exist in the UK and is often perceived as unjust. The current experiment finds support for the hypothesis that people are motivated to reduce inequalities in illness, here operationalised as pain. The experiment also tested two hypotheses with respect to how inequality might affect experienced pain: firstly that a perception of having been treated unfairly has an intrinsic effect to exacerbate pain, termed a ‘meanness effect’, secondly that a perception of having been treated with beneficence ameliorates suffering, termed a ‘kindness effect’. The results support a kindness effect, but did not reveal a meanness effect.

The hypotheses above unite existing lines of enquiry from psychology and economics. Psychologists have long been interested in the circumstances under which people are willing to cause or relieve others’ pain. The experiments of Stanley Milgram (unacceptable by many ethical standards, including those of the time: see Baumrind, 1964) reported that under specific conditions, designed to encourage obedience to the wishes of the experimenter, people could be induced to cause high levels of pain to others (Milgram, 1963, 1965). However numerous studies that have hitherto attracted less attention illustrate that humans are generally reluctant to actively harm others (Crockett et al., 2010; Cushman et al., 2012), show empathic responses to others’ pain (Jackson et al., 2005; Singer et al., 2004), and even choose to suffer pain, or pay money themselves to avoid causing pain to others (Batson et al., 1988; Batson et al., 1983; Crockett et al., 2014; Feldman-Hall et al., 2012; Hein et al., 2011). These studies combine to suggest that (at least within the narrow social context of psychological experiments) people find others’ pain intrinsically aversive.

Separately, economists have sought to quantify altruistic human behaviour through formal games, finding that many people prefer to share *money* with others, even when there is no financial reward for doing so (Engel, 2011), a behaviour that has been characterised as a preference for equality. Here I integrate these disparate lines of research, by formalizing the transfer of both pain and money between individuals in economic terms using a modified dictator game.

The data indicate that being allocated more pain than others motivated people to transfer pain to others, while being allocated less pain than others motivated people to take on others' pain, consistent with a preference for equal allocations of pain. Furthermore, increases in pain inequality had an increasing effect on the motivation to redress inequality. In psychological terms this can be viewed in terms of drive reduction: the further one is from a desirable goal, the stronger the motivation to act (Campbell & Kraeling, 1953). In economic terms this finding indicates an increasing marginal aversion to inequality (Bolton & Ockenfels, 2000; Loewenstein, Thompson, & Bazerman, 1989).

I have showed through model-fitting that this behaviour can result from a more fundamental principle of diminishing marginal utility (or in psychological terms, satiety) if it is assumed that social utility is given by a weighted sum of individual utilities. In other words, if people have some degree of concern for others, then diminishing marginal utility creates a motive to donate to those who have least, since they will benefit most from increases in utility. For pain, the model assumes *increasing* marginal *disutility*, such that the motive to share pain arises directly from the escalating cost of pain to each individual. Put simply, those in greatest pain are seen as being in the greatest need of reductions in pain.

For monetary outcomes, people more frequently acted to correct disadvantageous as compared with advantageous inequality, as seen in previous findings where disadvantageous monetary inequality has a larger motivating effect (Loewenstein et al., 1989). For pain, however, there was no significant degree of asymmetry. Consistent with this participants were significantly more altruistic when allocating pain than when allocating money. I also observed hyperaltruism for pain, whereby close to a third of participants (15 out of 47, 32%) retained more than half the shocks across all choices. This pattern was readily modeled by allowing the weighting given to others' utility to exceed that given to one's own utility.

Notably in this experiment the differences in mean offers between money and pain found were subtle, and only significant for the money loss frame (mean offer of shocks = 12.16/24, mean offer of debt = £12.98/£24). This stands in contrast with the more substantial differences in allocations between money and pain seen in conventional dictator game paradigms (A. L. Davis et al., 2011), where mean offers of money were more selfish than seen here (mean offer of cold water immersion time = 28.8s/60s, mean offer of money gain = \$1.80 /\$6). A possible explanation is that the limited choices used for the current study tended to generate a more stereotyped pattern of responding than conventional dictator game choices, and therefore may be less sensitive to differences in altruism between pain and money.

Notably however, the prevalence of hyperaltruistic behaviour seen here is closely comparable to that previously reported for a conventional dictator game with pain, which found that 16 out of 54 (30%) participants allocated less than half of the painful cold immersion time to the other participant (A. L. Davis et al., 2011). Notably both proportions are markedly lower than the rate of hyperaltruism in Experiment 3A (described in the preceding chapter of this thesis), where participants made binary choices between pain for themselves or others, and where 17 out of 27 were hyperaltruistic. Both proportions are also lower than those reported in a recent study which required participants to trade off pain and money, where over half of participants displayed hyperaltruism (Crockett et al., 2014). One explanation for this discrepancy would be that choices in which outcomes are described as allocations of pain to each person tend to prime concerns for equality.

Finally, by implementing selected allocations for real, I was able to examine the effects of inequality and perceived intentionality on the perception of pain itself. Although I found no significant influence of inequality *per se* on pain perception, I did observe a ‘kindness effect’, whereby participants rated a low number of shocks as less intense when they believed that the other participant had chosen to accept a high number of shocks at the same time, thus enduring extra pain for their benefit. This finding may reflect participants’ perception of being in the presence of a supportive other, consistent with existing findings that the presence of a familiar conspecific reduces pain responses in both humans (Krahé, Springer, Weinman, & Fotopoulou, 2013) and animals (F. R. D’Amato & Pavone, 1993; Langford et al., 2010). Lending support to this idea, one previous study found that pain chosen by another participant

was rated as less severe if the receiver believed that the other participant would increase their (the receiver's) monetary payment by giving pain, rather than believing that the other had no apparent motive for giving pain. Reassuringly, this effect suggests that the mere perception of beneficence on the part of caregivers has an intrinsic effect to ameliorate suffering.

Unlike previous studies (K. Gray, 2012; K. Gray & Wegner, 2008) this experiment found no evidence for the converse ('meanness') effect, namely an effect of selfish behaviour on the part of the dictator to increase pain perception for the receiver. An explanation for this null result might be that in the current study intentioned-inequality was in the number of shocks; effects of intentioned-inequality on subjective intensity ratings might therefore be expected to resemble the effect of increasing or decreasing the *number* of shocks. Subjective intensity ratings did increase as a function of the number of shocks (consistent with sensitisation), however this relationship saturated at higher numbers of shocks. Under this account a meanness effect would be equivalent to increasing the number of shocks from an already high level, with little resulting effect on pain perception. To formalise this hypothesis we modeled intensity ratings as a concave psychophysical function of number of shocks, using a Weibull function with two parameters, asymptote, A , and latency, L , constrained to pass through the origin:

$$Rating = A(1 - e^{-\frac{n}{L}})$$

Least squares fits of this function, displayed in Figure 10, illustrate that the saturation hypothesis appears plausible.

A key question for future research is whether relative deprivation or suffering in one modality alters responses to inequality in other modalities. There is already evidence to support this suggestion. For example, Mancini and colleagues report that experiencing pain increases self-oriented behaviour in a one-shot ultimatum game (Mancini, Betti, Panasiti, Pavone, & Aglioti, 2011), and that, somewhat in opposition to the kindness effect shown here, perceived monetary unfairness reduces neural responses to pain (Mancini, Betti, Panasiti, Pavone, & Aglioti, 2014).

GENERAL DISCUSSION

Vitae summae brevis spes nos vetat incohare longam
(Life's short span forbids our embracing far-reaching hopes).
- Horace , Odes (23BC)

This dissertation has examined a hypothesis that unhealthy behaviour, such as tobacco smoking, alcohol abuse or poor dietary choice, can be explained by a tendency to discount future health or illness. The work has also addressed theoretical overlap between temporal and social discounting of pain, with the additional hypothesis that these might be correlated across individuals as part of a 'patient-altruistic phenotype'. Chapters 2 and 3 described systematic reviews of studies measuring temporal discounting for hypothetical health outcomes, and studies comparing measured temporal discounting with health behaviour respectively. Chapter 4 critically appraised a hypothesis that hyperbolic discounting is responsible for a tendency for people to state that they have healthy intentions whilst continuing to perform unhealthy actions. Chapter 5 described general experiment methods used in subsequent chapters, including how Bayesian methods can be used to evaluate alternative behavioural models. Chapters 6 and 7 examined the phenomenon of 'negative temporal discounting' for pain and illness, characterizing this in terms of anticipatory utility. Finally, Chapters 8 and 9 explored choices regarding others' health, and responses to health inequality respectively, both conceptualised in terms of 'social discounting'.

The work supports the conclusion that temporal discounting measured for appetitive commodities, such as food and money, is a viable behavioural correlate of unhealthy choice (Chapter 3). This conclusion can be drawn without specifying the underlying psychological processes linking measured discounting with unhealthy choice. The preference reversals predicted by hyperbolic discounting purport to offer an explanation for the fact that people report healthy intentions which they fail to

carry out. However, the data reviewed here reject too strict an algorithmic interpretation of discounting, and theoretical considerations suggest that self-control failures are attributable to the effects of changes in environmental or motivational states (for example moving from satiety to hunger bringing about a lapse in one's diet plan, or entering a bar causing one to drink alcohol during a period of planned abstinence), rather than to hyperbolic discounting (Chapter 4).

The studies reviewed in Chapters 2 and 3 suggest that discounting for health-related outcomes is a separate construct from that of appetitive rewards, and may require further characterisation before being used as a predictor. The models of dread-discounting for pain and illness tested in Chapters 6 and 7 suggest a promising basis from which to describe individual differences in time preference for aversive health-related outcomes, which might in future studies be capable of accounting for further variance in behaviour, in particular that with an aversive component. However, that these functions explained only a small proportion of the variance in experimental health budgeting choices (Chapter 7) suggests that a still richer set of data than that provided by conventional intertemporal choice tasks might be necessary to predict real-world behaviour.

Finally, Chapters 8 and 9 find that models of social discounting that are formally similar to their temporal counterparts are capable of explained decisions regarding pain in others. However the findings of Chapter 8 find no support for a hypothesis that social and temporal preferences for pain are related, distinct from existing findings with rewarding outcomes. These (and further) conclusions are discussed in more detail below.

10.1 DELAY DISCOUNTING AS A VIABLE BEHAVIOURAL MARKER OF UNHEALTHY CHOICE

The systematic review presented in Chapter 3 revealed evidence for both state and trait-based influences on discounting of future reward. For example, discounting is higher in states of craving for drugs of abuse (e.g. Giordano et al., 2002; S. H. Mitchell, 2004; Yi & Landes, 2012), indicative of state-dependence, however discount rates for appetitive commodities, such as food and money, are stable over time (Odum, 2011), are correlated across different kinds of reward (Odum, 2011) and

show robust relationships with many forms of unhealthy behaviour, supporting the notion of discounting as a trait marker of short-sighted choice. Furthermore, measured discounting for money is able to prospectively predict relevant real-world behaviour. For example in a longitudinal study of 947 adolescents, Audrain-McGovern et al. (2009) found that a single standard deviation increase in baseline discounting increased the odds of subsequent smoking uptake by 11% (OR = 1.11, 95% CI = 1.03, 1.23). These findings suggest that monetary discounting could be measured in young adolescents to determine their risk of engaging in smoking in later adolescence. Higher discounting adolescents might then be offered targeted preventive interventions. Monetary discounting measured amongst current smokers also predicts rates of relapse following smoking cessation programmes (Brown & Adams, 2013; Krishnan-Sarin et al., 2007; MacKillop & Kahler, 2009; Sheffer et al., 2012; Yoon et al., 2007). These findings similarly suggest that discounting could be measured amongst smokers wishing to quit, with more intensive support or longer follow-up being offered to the highest discounters. A similar approach might be taken for other addictive behaviours, such as alcohol abuse.

Discounting for monetary reward has the advantage of being easy to measure. The much-cited monetary choice questionnaire (MCQ) as introduced by Kirby et al (1999) consists of only 27 choice options and uses hypothetical monetary amounts. Also, several studies, have convincingly demonstrated that the results of discounting questionnaires with hypothetical outcomes are not systematically different from those obtained when money is paid out for real (Johnson et al., 2002; Madden et al. 2003; Madden et al., 2004; Lagerio & Madden, 2005). Taken together, discounting forms a promising behavioural marker of unhealthy choice, and warrants further investigation to test whether offering additional interventions to the highest discounters is a cost-effective strategy.

10.2 ENVIRONMENTAL INFLUENCES ON DISCOUNTING

A potential limitation to the usefulness of discounting as a marker of unhealthy behaviour is that high discounting is correlated with socioeconomic variables, which are in turn correlated with health behaviour. Controlling for these variables then inevitably reduces the additional predictive power of discounting. For example Bradford (2010), report associations between discounting and health

behaviour in 978 adults. Associations between discounting and health behaviour are striking, for example, the highest discounting women are 15.3% less likely to have attended for mammography ($p=0.003$). However, when participants' access to health insurance was entered as an additional predictor, the effects were substantially smaller (although many of the associations between discounting and health behaviour did remain significant). Access to health insurance was assumed to influence the availability of preventive health procedures such as cancer screening, and to be associated with higher discounting via an association between discounting and lower socioeconomic status. In support of the latter, participants who had less than a high school level of education had mean discount rates of 4.1% lower than those educated to college level or above ($p=0.011$). Also, participants who reported needing to access their savings within a few months of the experiment had imputed annual discount rates 8.1% higher ($p<0.001$) than those who did not expect to access their savings within the next decade.

A unifying view considers measured discounting as a summary of an individual's beliefs regarding the future. (I use the term 'beliefs' here to refer to probabilistic predictions, updated in the light of evidence, e.g. Mathys, Daunizeau, Friston, & Stephan, 2011. Referring to these predictions as 'beliefs' rather than 'preferences' emphasises that they are both probabilistic and malleable). These beliefs, which putatively lie upstream of measured discounting, would be expected to be influenced by an individual's life history (Del Giudice, 2014; Ellis et al., 2012), including, but not limited to, their socioeconomic position. For example, lower financial resources might act to increase discounting through the greater opportunity cost of delaying money (e.g. Bradford, 2010). Notably, there may also be a bidirectional relationship between low socioeconomic status and discounting, whereby impatience for rewards leads to maladaptive choices such as substance misuse, which in turn are associated with worsening finances and further increases in discounting (e.g. S. Fields et al., 2009; Leitão, Guedes, Yamamoto, & Lopes, 2013). In this sense, discounting can be seen as psychological variable mediating the relationship between low socioeconomic status and unhealthy life choices such as substance misuse.

A further potential link between an individual's life experiences and their level of discounting concerns the level of uncertainty they hold regarding future rewards. The utility of future rewards might be more uncertain than immediate rewards, in the sense of having higher variance (when the variance is known the resulting uncertainty is referred to as risk). If future events tend to evolve with a random component, the uncertainty associated with future events increases with delay (Mathys, Daunizeau, Friston, & Stephan, 2011). To take an example, a decision-maker responding to a discounting questionnaire might have some degree of uncertainty about the subjective utility of a \$20 payout received immediately (if this appears implausible, imagine being paid in a foreign currency, whose worth is uncertain). However, owing to volatility governing future events in their lives (e.g. becoming ill, falling into debt, national economic collapse), uncertainty regarding the utility of the \$20 ought to increase as it is delayed. In combination with risk aversion this motivates delay discounting. In support of this idea, individual discount rates are correlated with risk aversion (Anderhub, Güth, Gneezy, & Sonsino, 2001; Dohmen, Falk, Huffman, & Sunde, 2010; Eckel et al., 2005; Jones & Rachlin, 2009; Leigh, 1986).

The above line of reasoning suggests that people who have experienced an unstable environment, in which future rewards are unreliable, ought to exhibit steeper discounting. In this regard an important direction for future research will be to examine the contribution of the childrearing environment to discounting. The study of self-regulation, often measured by the ability to delay gratification, in developmental psychology has adopted this approach; for example children who experience emotionally close, sensitive, and responsive caregiving have been found to exhibit higher levels of self-regulation (Belsky, Pasco Fearon, & Bell, 2007). Furthermore low childhood self-regulation has been prospectively related to poorer health outcomes later in life (for example Francis, 2009; Moffitt et al., 2011; Seeyave et al., 2009). Future research into delay discounting would benefit from a similar developmental perspective to better understand the origins of trait-level individual differences.

Recent theoretical work in the field of reinforcement learning offers principled approaches by which agents might update their beliefs about the changeability of the environment (e.g. Dayan & Daw, 2008; Mathys, Daunizeau, Friston, & Stephan, 2011), offering a model for how such learning comes to influence intertemporal

choice. Such models propose that when the environment is highly changeable (volatile) agents ought to update their beliefs about their current state more readily in the light of new information (see Behrens, Woolrich, Walton, & Rushworth, 2007; Mathys, Daunizeau, Friston, & Stephan, 2011; R. C. Wilson, Nassar, & Gold, 2010). In everyday terms, changing environments demand flexibility, since past experience rapidly becomes outdated. This can be formalised as an increased learning rate in more unstable settings. These models can be inverted to derive an estimate for an individual's belief about the volatility of the environment based on their learning. Given an estimate of the individual's level of risk aversion, such estimates of environmental uncertainty ought to be reflected in delay discounting choices made within that particular environment.

At present I am setting up a study to investigate the above hypothesis, using a paradigm which requires participants to imagine that they are selling items on eBay (on online marketplace). Participants are first required to track the price of two different items, which they sell on subsequent 'weeks'. For one of the items the price they obtain is stable, for another the price is changeable. More specifically the changeable price evolves according to a Gaussian random walk, such that the participant's uncertainty about the future price ought to increase linearly the further ahead they are required to predict (Mathys, Daunizeau, Friston, & Stephan, 2011). Having experienced the degree of volatility in the price, the participant is asked to state a confidence interval for the price they might obtain at varying times in the future. Finally, they are asked to choose when they would like to sell the item: either immediately or at some future delay (where the future price may be uncertain). By titrating the magnitude of the immediate selling price, we obtain an estimate of participants' discount rates for the different items. The stable-priced item forms a baseline estimate of discounting, while discounting for the volatile-priced item incorporates additional time-dependent risk. We also obtain independent estimates of each participants' degree of risk aversion, and their uncertainty regarding the future price derived from fitting a learning model.

We hope thereby to test the hypothesis that participants' level of discounting is sensitive to their perceived instability of the price for the items, providing a model for how unstable environments might come to be associated with steeper discounting. We also plan to gather data regarding participants' level of uncertainty pertaining to

future events in their own lives, for example their health, wealth and income, with the simple hypothesis that people who experience greater uncertainty over their future resources will discount the future more steeply. Notably this approach, of combining self-reported data with choice-based measures, bridges the social-cognitive methods discussed in the introduction to this work, and the choice-based methods which have formed the focus of this thesis.

10.2 FUTURE DIRECTIONS FOR STUDYING SELF-CONTROL IN HEALTH CHOICE

Many people who report healthy intentions subsequently fail to carry out their intended healthy actions, a fact that has limited the power of previous cognitive models to predict health behaviour (Sheeran, 2002). The observation that temporal discount functions for reward are well-described by a hyperbolic function (Green et al., 1994; Kirby & Herrnstein, 1995; Mazur, 1987), and that this hyperbolic function predicts a tendency to renege on one's best-laid plans (Strotz et al., 1957), has led to the idea that hyperbolic discounting, interpreted algorithmically, can explain the lapses of will which putatively underlie the above intention-action discrepancy (Ainslie, 1992; Ainslie, 1975, 2001).

However, as reviewed in Chapter 4, evidence that hyperbolic discounting is responsible, in an algorithmic sense, for self-control failures is mixed. In Chapter 4 I proposed, following on from previous theoretical accounts, that self-control failures frequently coincide with changes in environmental or motivational states, and that these state changes boost the value of actions leading to immediate reward, producing behaviour incongruent with long-term goals. For example, an environment previously associated with rewarding drug-taking behaviour can lead to drug taking even after prolonged abstinence. This was conceptualised in terms of the interaction between two systems of reinforcement learning, model-free and model-based (for a review see R. J. Dolan & Dayan, 2013). The former corresponds to habit learning and assigns (cached) values to actions in particular cues or contexts without explicitly representing the consequences, while the latter corresponds to goal-directed planning of actions within a cognitive map of the world. I proposed that the values stored by the model-free system might tend to be inaccessible to the model-based system. In

other words the effects of cues on one's behaviour might be difficult to forecast, generating qualitatively similar 'battles of will' as described by Ainslie (Ainslie, 1992; Ainslie, 2001).

10.2.1 'MODEL-BASED' AND 'MODEL-FREE' INTERACTIONS AS NOVEL BEHAVIOURAL PREDICTORS

The account above suggests that measures of the relative contributions to behaviour of model-based and model-free valuations might provide novel behavioural markers to supplement measures of discounting. One approach to measuring the interaction between model-based and model-free decision-making is to directly observe the acquisition of habitual behaviour through repeated training on a given laboratory task. Here, the rate of acquisition of habitual responding may offer a novel measure for predicting field behaviour. For example, outcome-insensitive habits have been demonstrated in humans (Tricomi, Balleine, & O'Doherty, 2009), providing a behavioural counterpart to studies of habit formation, which measure subjective automaticity (Lally, Van Jaarsveld, Potts, & Wardle, 2010). The speed with which responding becomes insensitive to outcome-devaluation might form an individual difference measure capable of predicting real-world habitual behaviour. Also, an important aim for future studies will be to examine habitual or cue-triggered preference reversals in real-time. Along these lines, subjective measures of habitual automaticity in relation to smoking behaviour have been shown to predict goal-incongruent smoking-related responses (Orbell & Verplanken, 2010, 2014).

However, observing habit learning directly is time-consuming. Recent human studies (Daw et al., 2011; Daw et al., 2005; Eppinger, Walter, Heekeren, & Li, 2013; Gläscher et al., 2010; Smittenaar et al., 2013) have used a paradigm with a probabilistic tree structure which separates model-free and model-based control, before habitisation has taken place, depending on whether respondents incorporate the transition structure of the task into their learning (model-based) or learn solely based on the reinforcement obtained in each discrete state (model-free). Humans performing this task generally exhibit some combination of the two modes of control, and the relative contribution of the two strategies may provide a novel behavioural marker. Further studies are required to establish the longitudinal stability of these measures,

and whether they have a trait component, as well as to examine their relationship with habitual behaviour in the field.

10.2.2 RELATIONSHIPS BETWEEN ‘MODEL-BASED’ CONTROL AND DISCOUNTING

Mentally simulating future outcomes decreases measured discount rates (J. Peters & Büchel, 2010) and lesioning neural structures on which this simulation process depends, such as the hippocampus (Hassabis, Kumaran, Vann, & Maguire, 2007; A. Johnson, van der Meer, & Redish, 2007; Schacter, Addis, & Buckner, 2008; Schacter & Schacter, 2008) increases discounting (Mariano et al., 2009). Furthermore, existing studies suggest that the choice of delayed rewards, model-based control and working memory engage overlapping neural substrates: neuroimaging studies have found that the dorsolateral prefrontal cortex (dlPFC) is activated in both model-based learning (Gläscher et al., 2010), and in choosing delayed rewards on intertemporal choice paradigms (McClure et al., 2004, 2007), while disrupting this area (using either transcranial magnetic stimulation or transcranial direct current stimulation) both decreases model-based behaviour (Smittenaar et al., 2013) and increases temporal discounting (Hecht, Walsh, & Lavidor, 2013). Taken together these results suggest that a network subserving mental simulation of the future provides a basis for model-based control and that the engagement of this system tends to generate more patient intertemporal choices.

A key question for future work is how an increased emphasis on model-based control might mediate reduced discounting. A plausible possibility is that mentally simulating the future resolves uncertainty about the utility of larger-later rewards (e.g. Daw et al., 2005). In this regard, an interesting possibility is that environmental conditions might be expected to interact with the deployment of future thinking in a cost-benefit trade-off. For example, if conditions are sufficiently unpredictable, then attempting to explicitly plan out future possibilities is futile, and may even be disadvantageous, given that thinking ahead is computationally (i.e. energetically) costly (see Daw et al., 2005). A corollary of this reasoning is that prolonged exposure to an unstable environment ought to discourage the use of model-based strategies. This possibility would conceptually bind together an unstable childhood environment, diminished cognitive ability and steeper discounting of reward, providing a tentative

theoretical basis for explaining the association between these factors in several psychiatric disorders. For example, people with borderline personality disorder are likely to have experienced childhood abuse (Lewis & Christopher, 1989; Ogata et al., 1990; Zanarini, Williams, Lewis, & Reich, 1997), exhibit below average cognitive function (Swirsky-Sacchetti et al., 1993) and discount the future more steeply than healthy controls (Lawrence et al., 2010).

10.2.3 PRE-COMMITMENT, WILLPOWER AND SOPHISTICATION

The account of self-control presented in Chapter 4 considers a *naïve* decision maker, who cannot forecast the effects of state-changes on their behaviour. However people often demonstrate that they can predict their future tendencies, termed *sophistication*. For example, people sometimes choose paths that remove their opportunity to make myopic choices, an activity referred to as *pre-commitment* (Ariely & Wertenbroch, 2002; Crockett et al., 2013; Kurth-Nelson & Redish, 2010, 2012): a person attempting to abstain from smoking might for instance throw away their cigarette packets. Pre-commitment devices would be expected to obscure real-world relationships between discount rates and myopic behaviour, since at least a subset of sophisticated steep discounters would exhibit far-sighted real-world choices.

In addition to physical pre-commitment, sophisticates can also employ more subtle mental devices to constrain their future choices (Ainslie, 2001; Prelec & Bodner, 2003), by forming a self-fulfilling prediction about their future behaviour, termed an intention (Gollwitzer, 1993) or a resolution (Ainslie, 2001). Gollwitzer (1993) suggests a particular form of commitment device called as ‘implementation intention’, which involves specifying an “if-then” plan as to when, where and how the goal-directed behaviour will be enacted:

The purpose of an implementation intention is to lay down a specific plan that helps to promote the initiation and efficient execution of goal-directed activity (Gollwitzer, 1993).p152.

Forming these ‘if-then’ plans is assumed to cede action control to situational cues, preventing people from being derailed from their goals by contextual factors

(Gollwitzer, 1993). Viewed in terms of the model-free and model-based interactions discussed above, an implementation intention can be envisaged as *pre-assigning goal value* to executing a particular action in response to a cue. There is good evidence that coaching people to form implementation intentions fosters successful behaviour change, across a number of settings not restricted to health (Gollwitzer & Sheeran, 2006; Sheeran et al., 2009).

Interestingly, as Ainslie (Ainslie, 1975, 1992, 2001) and Prelec (Bodner & Prelec, 2003; Prelec & Bodner, 2003) have both discussed, the success of intentions also lies in the belief that they will be consistently carried out (corresponding to the concept of self-efficacy; Bandura, 1977). Should a person fail to execute their intended action this is likely to undermine their belief in their ability to do so in the future. In this manner, the entire value of the ultimate goal (for example losing weight, and living longer) can come to be staked on successfully executing the intended actions on only very few occasions (for example abstaining from biscuits today) (Prelec & Bodner, 2003). For example, eating a single biscuit has a vanishingly small effect on one's future weight gain and/or longevity. If each choice of whether or not to consume an extra biscuit is taken in isolation, the enjoyment of consumption might outweigh the miniscule (and temporally distant) cost, especially if one believes one will abstain from biscuits next time. By contrast, if consuming biscuits on this occasion is taken as signal that one is likely to do so in future, then each individual choice gains greater leverage. In the limit, giving in to temptation this time might be taken to indicate that one will never succeed in resisting, and so by proxy acquires the value of consuming on a whole series of future occasions (e.g. eating a packet of biscuits every day for the next few months), making the health consequences more significant, and well-placed to outweigh the enjoyment of a single instance of consumption. Studying self-signalling and cognitive precommitment experimentally is challenging, but worthy of further study. Furthermore novel interventions might be developed that seek to promote self-signalling processes so as to encourage healthier choice, or adherence to existing behaviour change regimes.

10.3 DREAD AND AVOIDANCE BEHAVIOUR

The work presented in Chapter 6 represents the first attempt to fit alternative dread-discounting functions to choices of delayed pain, and so to quantitatively describe intertemporal preference functions for pain. A goal for future work will be to examine whether dread is capable of predicting avoidance of preventive health behaviours, or indeed other unpleasant tasks. To achieve this, hypothetical questions designed to assess dread, such as the choices between alternative dental appointments offered in Experiment 1B, could be administered to a large sample of participants being offered a preventive procedure, such as a cancer screening appointment. Regression analysis could then be used to identify relationships between uptake of screening and parameters derived from a dread-discounting model. The pronounced individual differences in behaviour seen on such paradigms in this thesis suggest that dread-discounting offers a rich source of individual variability. However any such relationships with real-world behaviour remain to be tested. Further work is also required to establish the test-retest reliability of estimated dread, and whether dread is correlated for different forms of unpleasant outcome, in other words whether dread has features of a trait marker.

10.3.1 AVOIDABLE PAIN

A paradox exists with respect to the relationship between dread and avoidance. On the one hand dread increases the aversiveness of future pain, which ought to motivate avoidance; on the other hand if the pain is in fact *unavoidable* then dread motivates speeding up the onset of pain. It is therefore unclear on *a priori* grounds whether higher dread would predict greater avoidance of unpleasant preventive behaviours, or a tendency to get them over and done with. An important variable would appear to be whether a person believes (implicitly or explicitly) that the dreaded event can be avoided through delaying it.

The notion that pain can be avoided through deferral would be expected to generate discounting of dread itself – since there is a possibility that future dread will not be encountered. In Experiment 1A, a small proportion of participants (2 out of 33) indeed exhibited negative time preference which reverted to positive time preference

at longer delays, consistent with a model in which dread is itself prospectively discounted. The key prediction here is that the negative value function for pain has a maximum at an intermediate time point, as opposed to increasing or decreasing steadily across time (Loewenstein, 1987). Such maxima predict that people will be most motivated to avoid the dreaded event at the point of maximal aversion, when the event is ‘looming’ in the not-too-distant future (See Figure 10.1).

Discounting of dread itself could be studied by eliciting time preferences for pain where delaying pain carries the possibility of avoiding the pain. A particularly interesting approach would be to do so in real-time. This could be implemented in the laboratory, by giving participants the opportunity to pay to avoid a pain which is approaching in time, where there is also a fixed probability per unit time that the pain will simply go away of its own accord. However such a laboratory paradigm could be implemented alongside a study over more naturalistic time scales, of days to weeks. For example a candidate task would ask participants on successive days how much they would be willing to pay to avoid having to perform an unpleasant task (such as drinking an unpleasant drink, performing strenuous physical activity or doing a tedious number-counting task) at the end of the week, where there is also a fixed probability for each day that passes they will be ‘let off the hook’ without having to perform the task.

In such scenarios, where future pain is potentially avoidable, an hypothesis is that greater *discounting of dread* (implying a greater belief in the avoidability of pain) would be expected to be associated with real-world avoidant behaviour. This could be tested by implementing the ‘avoidable pain’ paradigm suggested above in a cohort of participants invited for an unpleasant screening procedure, and examining whether the parameter indexing discounting of dread predicts failure to engage with screening. A further interesting possibility relates to the hedonic impact of future unpleasant events. Discounted dread predicts that future events ought to have maximal negative effects on wellbeing when they are looming in the ‘intermediate’ future. This prediction could be tested by an experience-sampling method, asking participants to rate their hedonic state as a dreaded event approaches in time.




Figure 10.1 Discounted dread and avoidance. Discounting of dread predicts that as pain draws nearer, there is a point of maximal aversion, where dread can be alleviated by either pushing the pain further into the future (thus reducing its probability of actually happening) or advancing the pain to ‘get it out of the way’. Due to negative time preference, at the point of maximal aversion people would be expected to pay more than the immediate cost of the pain itself in order to avoid it. To illustrate this here the present value (vertical axis, note negative scale) of a dreaded event under a Discounted Exponential Dread function is plotted as a function of the decision maker’s position in time (curved line), by comparison to a fixed cost alternative (horizontal line). Say for example that on Monday David agrees to go for a run with a friend on the coming Saturday, which he is dreading due to lack of fitness. David could call the friend to cancel, but this would be embarrassing, indicated by the fixed cost (horizontal line). Early in the week, when the run is still distant, the cost of cancelling appears to outweigh the pain of running. However as Saturday approaches, the dread of running looms larger than the cost of cancelling (shaded area) and David might be tempted to cancel; here the point of maximal dread is on Friday. By contrast, if the run were arranged on the same day, the total dread would be relatively lower, and the embarrassment of cancelling would still tend to outweigh the prospect of running.

10.3.2 WHAT DOES DREAD REPRESENT?

The models of dread introduced in Chapter 5 of this thesis assume that the instantaneous anticipation of pain increases as pain is approach in time, and that when choosing the timing of pain, people take account of a prospective sum of this anticipation. Whilst this model appears to account well for the observed choice data, it is in some respects unrealistic.

Firstly, for real-life events, anticipated over a time scale of days or weeks for example, the build-up of instantaneous anticipation is not likely to be a smooth

function of time. By contrast, one might experience spikes of anticipation when reminded of the upcoming event. For example if, due to a fear of flying, I am dreading taking a plane to travel abroad for my summer holiday, my fear is likely to peak in response to stimuli associated with flying, such as when checking in for the flight online. Nevertheless, it appears sensible to conclude that these ‘reminder’ events become more frequent as the dreaded event approaches, such that an increasing-anticipation-by-time function is not implausible.

Secondly, it appears unlikely that, when choosing the timing of pain, people explicitly think ahead over their future anticipation. It seems more likely that, through experience, people come to learn that deferring unpleasant events is associated with unpleasant anticipation, through which they form a useful heuristic to ‘get the dreaded event over and done with’. Nevertheless, people who find anticipating future events more aversive would be expected to incorporate this heuristic to a greater extent.

Exactly why the anticipation of pain (rather than pain itself) should be something to avoid is unclear. For short-term instances of anticipated physical pain, it seems reasonable to assume that the expectation of pain is intrinsically unpleasant, in other words that organisms are innately programmed with a desire to ameliorate sensations associated with the prediction of pain. In most ecological settings in our evolutionary history this would be achieved by escaping the potentially noxious situation altogether, and would therefore be a protective response. However, as humans we are faced with peculiar scenarios in which the anticipation of inevitable pain can be best minimised by simply facing the pain as soon as possible, rather than attempting to avoid it.

From a normative perspective, the cost of pain, and its anticipation, can be seen to derive from its effect to interfere with cognition and behaviour. The psychiatrist and behavioural economist George Ainslie proposes this idea in his work ‘Breakdown of Will’ (Ainslie, 2001). Ainslie suggests that pain represents a highly salient stimulus, which consumes cognitive resources, exerting its effect through the opportunity cost of being unable to pursue other rewarding thoughts and activities (Ainslie, 2001). Further work on dread might seek to establish whether the anticipation of pain indeed interferes with ongoing reward-seeking behaviour, and whether participants who are more sensitive to such interfering effects are more likely to choose to ‘get pain out of the way’. This could be tested by having participants

perform an incidental reward-seeking task concurrent with the anticipation of pain and subsequently making choices between the same painful outcomes at varying delays.

A further possibility, raised in Chapter 8 of this thesis, is that the desire to ‘have things sooner’ is independent of the nature of the outcome. That is, a tendency to expedite pain might simply be a form of generic impatience, perhaps driven by a cognitive cost associated with keeping track of scheduled future events. The failure to observe significant anticipatory skin conductance responses to delayed shock in Experiment 3A would be consistent with this possibility. This raises an interesting question for future work, and would predict, somewhat paradoxically, that participants who exhibit greater negative discounting for punishment ought also to exhibit greater positive discounting of reward. This could easily be tested and is something I intend to address in future work.

10.4 SOCIAL DISCOUNTING OF PAIN

In Chapter 8 I presented the first description of social discounting for pain, finding that participants were in fact more averse to harming close others than to harming themselves, but that this tendency steeply decays with increasing social distance. Contrary to predictions however, social and temporal preferences for pain were uncorrelated across participants, suggesting that, for pain, deciding for others and deciding for one’s future self engage distinct processes. This finding differs from previous evidence that social and temporal discounting for reward *are* related constructs (Bartels & Rips, 2010; Ersner-Hershfield, Wimmer, & Knutson, 2009; Jones & Rachlin, 2006; Jones & Rachlin, 2009; Pronin, Olivola & Kennedy, 2008) . Further work might focus on establishing the relationships between social and temporal preferences for reward and punishment in process, for example to examine whether factors such as reward magnitude or gain-loss framing influence each form of discounting in a distinct manner. A similar approach has been taken for temporal and probability discounting of rewards (Green, Myerson, & O’Donoghue, 1999; O’Donoghue, Green, & Myerson, 1998).

Healthcare professionals are sometimes required to cause a degree of physical pain or discomfort to a patient within the context of medical procedures that are in the

patient's long-term best interest. Understanding how people evaluate others' pain, and whether such evaluations are in any way biased is therefore of interest. In Chapter 8 I also examined for the first time a hypothesis that people exhibit diminished negative discounting when choosing pain on behalf of another person, and found support for this idea. Whilst people did seek to hasten the occurrence of unavoidable pain for others, when choosing on behalf of close friends or relatives this tendency was diminished relative to when choosing for themselves. An interpretation offered is that people are reluctant to increase the intensity of pain for close others to compensate for a delay. This finding suggests that the actions of caregivers might tend to underestimate the extent to which patients dislike waiting prior to painful experiences, representing a subject for applied study. For example, this could be studied by interviewing caregivers and patients about their preferences for the timing of painful interventions, to test for the hypothesised discrepancy.

10.5 SPREADING AND INEQUALITY AVERSION FOR PAIN

In the experiment described in Chapter 7, involving consumption of a budget of pain relief, many participants chose to distribute their budget evenly across time. I explored through simulation and model-fitting how this behaviour can result from diminishing marginal utility over relief (increasing marginal disutility of pain); this form of utility function means that being left with less relief in any given time period is disproportionately costly, creating a motive to spread relief. In Chapter 9 I applied a theoretically equivalent principle to choices to distribute pain between individuals, using a modified dictator game with painful outcomes. In simple terms this model formalises the idea that those with more pain to begin with are seen as being in greatest need of reductions in pain. An interesting direction for future work will be to examine relationships between the temporal and social spreading phenomena, given their hypothesised common origin in non-linear utility. This could be tested in the laboratory, by eliciting preferred allocations of pain relief (or reward) across time and across individuals in the same participants, as well independently eliciting an instantaneous utility function for increasing pain relief (or reward), for example by measuring risk attitude (e.g. Anderson et al., 2008).

10.6 LIMITATIONS

This thesis has championed the idea simple decision-making tasks can be harnessed as predictors of real-world health behaviour. An outline of the general paradigm is as follows:

1. Identify a ‘problem’ behaviour, e.g. alcohol abuse
2. Identify a decision task thought to engage psychological processes related to the problem behaviour
3. Compare parameters of task performance with measures of the problem behaviour, e.g. correlate discount rates with number of units of alcohol consumed per week
4. Test whether task performance is capable of prospectively predicting health behaviour or response to behaviour change interventions, e.g. whether alcohol dependent individuals with higher discount rates are less likely to sustain abstinence
5. Trial targeting interventions towards high-risk groups, based on task performance, e.g. offering a contingency management programme to individuals with steeper discounting wishing to cut down on drinking

The literature review presented in Chapter 2 examined the extent to which the above steps have been achieved for an existing decision parameter, the temporal discount rate for rewards. The conclusions of this were that steps 1 to 3 have been satisfactorily addressed for discounting, with several authors also having addressed step 4. The results, that discounting is robustly associated with addictive behaviours, and can prospectively predict outcomes in these areas, lend support to the general paradigm of using decision tasks to generate behavioural markers of maladaptive real-world behaviour. Further work is clearly needed to address step 5, and thereby to test whether the predictive power of discounting can be translated into cost effective interventions.

10.6.1 FACE VALIDITY OF DECISION TASKS

The empirical work of this thesis has been concerned with generating novel candidate behavioural markers, relating to temporal and social discounting of pain. A

key issue for such tasks is the extent to which they engage decision processes that are also engaged in real-world health-related decisions. A potential critique is that the laboratory tasks are somewhat rarefied by comparison to the real world scenarios they attempt to emulate. For example, by contrast to real-world pain, a participant receiving shock stimuli in the laboratory need not fear the possibility of actual tissue damage. Less still ought they to fear a threat to their life. This stands in contrast to, for example, the dread of receiving a diagnosis of cancer. As alluded to above, for the case of dread, an important direction for future work will be to examine consistency between responses for hypothetical scenarios and real pain, and also for consistency in response patterns over varying timescales. For example, does a tendency to dread shocks in the laboratory over 30 second delays correlate with a tendency to dread hypothetical dental pain over an imagined delay of 6 weeks or to dread the imagined possibility of receiving a diagnosis of cancer? If consistent responding can be demonstrated across such diverse scenarios this would go a long way to allaying potential concerns that rarefied pain-related decision tasks in the laboratory have poor face validity for everyday choices.

The problem of poor face validity might also be circumvented to some extent by exploiting the remarkable human capacity for imagination. For example, more immersive tasks might be designed in which audiovisual stimuli, or unpleasant pictures are used to evoke more complex sets of emotions pertaining to real-world situations, such as stimuli associated with cancer screening. A tension emerges here between the need to generate controlled tasks designed to probe basic decision-mechanisms and the need for face validity in predicting real-world behaviour. I would advocate that both kinds of task will prove necessary in the endeavour to understand and predict health behaviour. To be of practical use, tasks must also be capable of being transported outside the laboratory and into the clinic. For this purpose a computer-based task involving sounds, pictures and text would be greatly preferable to the use of electric shock stimuli, for example.

10.6.2 PARTICIPANTS AND SAMPLING BIASES

Several of the experiments described here recruit healthy adult participants who have volunteered to take part in psychology experiments. In many respects such

samples might differ from the general population, and indeed from the population of people considering changing their health behaviour. A substantial proportion of these participants tend to be well educated, and often have a personal interest in psychology. A further problem is that some of these individuals take part in a large number of psychology experiments, and as such might be over-trained in responding to similar kinds of choice-based task. A limitation of the current work is that rather scant demographic data about participants were collected.

Another potential issue is selection bias; since the experiments described here all involved painful stimulation, this would be expected to influence the composition of the sample – particularly anxious participants for example might have chosen not to take part. This is potentially problematic for drawing conclusions regarding the distribution of behavioural tendencies in the general population. The necessary requirement for informed consent makes obtaining a truly unbiased sample difficult for experiments involving experienced pain.

Attempts might be made to address these limitations by developing tasks that can be administered online, and therefore made accessible to a much larger sample, covering a wider cross section of society. Researchers at the Wellcome Trust Centre for Neuroimaging have recently taken a similar approach, developing a smart-phone based platform for decision-making studies which has now been accessed by many thousands of users (The Great Brain Experiment, London, UK). Encouragingly, the results, now published, replicate the findings of several laboratory experiments (Brown et al., 2014; McNab et al., 2015).

10.6.3 ON COMPUTATIONAL MODELLING

A potential criticism of computational approaches in psychology, such that taken here, is that they employ an unnecessary degree of mathematical complexity in order to demonstrate a conclusion that might have been more simply stated. For example, it might be argued that the key conclusion of the experiments investigating dread described here is simply that most people dislike delaying inevitable pain, and that the models of dread investigated are little more than an elaborate curve-fitting exercise. My view is that this is a very salutary criticism, and especially ought to serve as a warning to those engaged in computational neuroscience and psychology to ensure that their models can be interpreted based on identifiable features of the data. Nevertheless, I advocate that modelling behavioural and neural data has key

advantages. Firstly building a model forces the experimenter to explicit about the assumptions underlying their theories (Farrell & Lewandowsky, 2010). Secondly, and in a related sense, a well-specified model serves as a rigorous hypothesis regarding how a system will behave in a related experimental setting (Farrell & Lewandowsky, 2010). For example, hyperbolic discounting, predicts that an agent ought to exhibit inconsistent choice over time. Such hypotheses can then be subjected to empirical scrutiny.

10.7 FIN

To sum up, this dissertation has investigated how people make choices regarding the timing of (relief from) pain and illness, both for themselves and for others. The chief contributions of the work have been to appraise discounting as a marker of health choice, to provide empirical validation for existing theoretical models of anticipatory utility for pain and to provide the first direct empirical tests of social discounting for pain. Future work will focus on examining whether parameters of these models are capable of predicting real-world health-behaviour, and on probing the underlying mechanisms of such effects. In this final chapter I have argued however that researchers of intertemporal choice should not limit themselves to measuring discount functions, but should consider the reasons for which people might choose to variously defer or expedite future events, since doing so offers a means to link an individual's behaviour with their life-history.

BIBLIOGRAPHY

- Adams, J. S. (1965). Inequity in social exchange. *Advances in Experimental Social Psychology*, 2, 267-299.
- Adler, N., & Matthews, K. (1994). Health psychology: Why do some people get sick and some stay well? *Annual Review of Psychology*, 45(1), 229-259.
- Ainslie, G. (1992). *Picoeconomics: The strategic interaction of successive motivational states within the person*. Cambridge, UK: Cambridge University Press.
- Ainslie, G., & Haendel, V. (1983). The motives of the will. In E. Gottheil, K. Durley, T. Skodola, and H. Waxman (Eds.), *Etiologic aspects of alcohol and drug abuse* (pp. 119-140) Springfield, IL, USA: Charles C. Thomas
- Ainslie, G. W. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*, 21(3), 485-489.
- Ainslie, G. W. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82(4), 463-496.
- Ainslie, G. W. (2001). *Breakdown of will*. Cambridge, UK: Cambridge University Press.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179-211.
- Allan, J. L., Johnston, M., & Campbell, N. (2010). Unintentional eating. What determines goal-incongruent chocolate consumption? *Appetite*, 54(2), 422-425.
- Allgöwer, A., Wardle, J., & Steptoe, A. (2001). Depressive symptoms, social support, and personal health behaviors in young men and women. *Health Psychology*, 20(3), 223-227.
- Anderhub, V., Güth, W., Gneezy, U., & Sonsino, D. (2001). On the interaction of risk and time preferences: An experimental study. *German Economic Review*, 2(3), 239-253.
- Anderson, K. G., & Woolverton, W. L. (2005). Effects of clomipramine on self-control choice in Lewis and Fischer 344 rats. *Pharmacology Biochemistry and Behavior*, 80(3), 387-393.
- Anderson, L. R., & Mellor, J. M. (2008). Predicting health behaviors with an experimental measure of risk preference. *Journal of Health Economics*, 27(5), 1260-1274.
- Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2008). Eliciting risk and time preferences. *Econometrica*, 76(3), 583-618.

- Andrade, L. F., & Petry, N. M. (2012). Delay and probability discounting in pathological gamblers with and without a history of substance use problems. *Psychopharmacology*, 219(2), 491-499.
- Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., & Weinberg, S. (2001). The Hyperbolic Consumption Model: Calibration, simulation, and empirical evaluation. *The Journal of Economic Perspectives*, 15(3), 47-68.
- Anokhin, A. P., Golosheykin, S., Grant, J. D., & Heath, A. C. (2011). Heritability of delay discounting in adolescence: a longitudinal twin study. *Behavior genetics*, 41(2), 175-183.
- Appelhans, B. M., Woolf, K., Pagoto, S. L., Schneider, K. L., Whited, M. C., & Liebman, R. (2011). Inhibiting food reward: delay discounting, food reward sensitivity, and palatable food intake in overweight and obese women. *Obesity*, 19(11), 2175-2182.
- Ariely, D., & Wertenbroch, K. (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, 13(3), 219-224.
- Armitage, C. J., & Conner, M. (2000). Social cognition models and health behaviour: A structured review. *Psychology and Health*, 15(2), 173-189.
- Asenso-Boadi, F., Peters, T. J., & Coast, J. (2008). Exploring differences in empirical time preference rates for health: an application of meta-regression. *Health Economics*, 17(2), 235-248.
- Audrain-McGovern, J., Rodriguez, D., Epstein, L. H., Cuevas, J., Rodgers, K., & Wileyto, E. P. (2009). Does delay discounting play an etiological role in smoking or is it a consequence of smoking? *Drug and Alcohol Dependence*, 103(3), 99-106.
- Axon, R. N., Bradford, W. D., & Egan, B. M. (2009). The role of individual time preferences in health behaviors among hypertensive adults: a pilot study. *Journal of the American Society of Hypertension*, 3(1), 35-41.
- Bach, D. R., Flandin, G., Friston, K. J., & Dolan, R. J. (2010). Modelling event-related skin conductance responses. *International Journal of Psychophysiology*, 75(3), 349-356.
- Badia, P., McBane, B., & Suter, S. (1966). Preference behavior in an immediate versus variably delayed shock situation with and without a warning signal. *Journal of Experimental Psychology*, 72(6), 847-852.
- Baker, F., Johnson, M. W., & Bickel, W. K. (2003). Delay discounting in current and never-before cigarette smokers: similarities and differences across commodity, sign, and magnitude. *Journal of Abnormal Psychology*, 112(3), 382-392.
- Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37(4), 407-419.
- Bandura, A. (1977). Self-efficacy: toward a unifying theory of behavioral change. *Psychological Review*, 84(2), 191.
- Bargh, J. A. (1992). The ecology of automaticity: Toward establishing the conditions needed to produce automatic processing effects. *The American Journal of Psychology*, 105(2), 181-199.

- Baron, J. (2000). *Thinking and deciding*: Cambridge UK: Cambridge University Press.
- Baron, J., & Ritov, I. (2004). Omission bias, individual differences, and normality. *Organizational Behavior and Human Decision Processes*, 94(2), 74-85.
- Baron-Cohen, S., & Wheelwright, S. (2004). The empathy quotient: an investigation of adults with Asperger syndrome or high functioning autism, and normal sex differences. *Journal of autism and developmental disorders*, 34(2), 163-175.
- Bartels, D. M., & Rips, L. J. (2010). Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology: General*, 139(1), 49.
- Batson, C. D., Dyck, J. L., Brandt, J. R., Batson, J. G., Powell, A. L., McMaster, M. R., & Griffitt, C. (1988). Five studies testing two new egoistic alternatives to the empathy-altruism hypothesis. *Journal of Personality and Social Psychology*, 55(1), 52.
- Batson, C. D., O'Quin, K., Fultz, J., Vanderplas, M., & Isen, A. M. (1983). Influence of self-reported distress and empathy on egoistic versus altruistic motivation to help. *Journal of Personality and Social Psychology*, 45(3), 706.
- Baum, A., Garofalo, J., & Yali, A. (1999). Socioeconomic status and chronic stress: does stress account for SES effects on health? *Annals of the New York Academy of Sciences*, 896(1), 131-144.
- Baum, A., & Posluszny, D. M. (1999). Health psychology: mapping biobehavioral contributions to health and illness. *Annual Review of Psychology*, 50(1), 137-163.
- Baumrind, D. (1964). Some thoughts on ethics of research: After reading Milgram's "Behavioral study of obedience." *American Psychologist*, 19(6), 421.
- Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3), 226-232.
- Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *The Journal of Political Economy*, 96(4), 675-700.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214-1221.
- Belin, D., Mar, A. C., Dalley, J. W., Robbins, T. W., & Everitt, B. J. (2008). High impulsivity predicts the switch to compulsive cocaine-taking. *Science*, 320(5881), 1352-1355.
- Bellman, R. (1956). Dynamic programming and Lagrange multipliers. *Proceedings of the National Academy of Sciences, USA*, 42(10), 767.
- Belsky, J., Pasco Fearon, R., & Bell, B. (2007). Parenting, attention and externalizing problems: Testing mediation longitudinally, repeatedly and reciprocally. *Journal of Child Psychology and Psychiatry*, 48(12), 1233-1242.
- Benartzi, S., & Thaler, R. H. (2007). Heuristics and biases in retirement savings behavior. *The Journal of Economic Perspectives*, 81-104.
- Benzion, U., Rapoport, A., & Yagil, J. (1989). discount rates inferred from decisions - an experimental-study. *Management Science*, 35(3), 270-284.

- Berns, G. S., Chappelow, J., Cekic, M., Zink, C. F., Pagnoni, G., & Martin-Skurski, M. E. (2006). Neurobiological substrates of dread. *Science*, 312(5774), 754-758.
- Bertsekas, D. P. (1976). *Dynamic programming and stochastic control*. Orlando, FL: Academic Press.
- Bickel, W. K., Jarmolowicz, D. P., Mueller, E. T., Koffarnus, M. N., & Gatchalian, K. M. (2012). Excessive discounting of delayed reinforcers as a trans-disease process contributing to addiction and other disease-related vulnerabilities: emerging evidence. *Pharmacology & Therapeutics*, 134(3), 287-297.
- Bickel, W. K., Koffarnus, M. N., Moody, L., & Wilson, A. G. (2014). The behavioral-and neuro-economic process of temporal discounting: A candidate behavioral marker of addiction. *Neuropharmacology*, 76, 518-527.
- Bickel, W. K., & Marsch, L. A. (2001). Toward a behavioral economic understanding of drug dependence: delay discounting processes. *Addiction*, 96(1), 73-86.
- Bickel, W. K., Miller, M. L., Yi, R., Kowal, B. P., Lindquist, D. M., & Pitcock, J. A. (2007). Behavioral and neuroeconomics of drug addiction: competing neural systems and temporal discounting processes. *Drug and Alcohol Dependence*, 90 Suppl 1, S85-91.
- Bickel, W. K., Odum, A. L., & Madden, G. J. (1999). Impulsivity and cigarette smoking: delay discounting in current, never, and ex-smokers. *Psychopharmacology*, 146(4), 447-454.
- Bickel, W. K., Yi, R., Landes, R. D., Hill, P. F., & Baxter, C. (2011). Remember the future: working memory training decreases delay discounting among stimulant addicts. *Biological Psychiatry*, 69(3), 260-265.
- Bjork, J. M., Hommer, D. W., Grant, S. J., & Danube, C. (2004). Impulsivity in abstinent alcohol-dependent patients: relation to control subjects and type 1-/type 2-like traits. *Alcohol*, 34(2-3), 133-150.
- Bleichrodt, H., & Johannesson, M. (2001). Time Preference for Health: A Test of Stationarity versus Decreasing Timing Aversion. *Journal of Mathematical Psychology*, 45(2), 265-282.
- Bobova, L., Finn, P. R., Rickert, M. E., & Lucas, J. (2009). Disinhibitory psychopathology and delay discounting in alcohol dependence: personality and cognitive correlates. *Experimental and Clinical Psychopharmacology*, 17(1), 51.
- Bodner, R., & Prelec, D. (2003). Self-signaling and diagnostic utility in everyday decision making. In I. Brocas, J. D. Carrillo (Eds.) *The Psychology of Economic Decisions. 1: Rationality and Well-being*. New York: Oxford University Press.
- Boettiger, C. A., Mitchell, J. M., Tavares, V. C., Robertson, M., Joslyn, G., D'Esposito, M., & Fields, H. L. (2007). Immediate reward bias in humans: fronto-parietal networks and a role for the catechol-O-methyltransferase 158(Val/Val) genotype. *Journal of Neuroscience*, 27(52), 14383-14391.
- Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1), 166-193.

- Bonato, D. P., & Boland, F. J. (1983). Delay of gratification in obese children. *Addictive Behavior*, 8(1), 71-74.
- Borghans, L., & Golsteyn, B. H. H. (2006). Time discounting and the body mass index: Evidence from the Netherlands. *Economics & Human Biology*, 4(1), 39-61.
- Bornovalova, M. A., Daughters, S. B., Hernandez, G. D., Richards, J. B., & Lejuez, C. (2005). Differences in impulsivity and risk-taking propensity between primary users of crack cocaine and primary users of heroin in a residential substance-use program. *Experimental and Clinical Psychopharmacology*, 13(4), 311.
- Boucsein, W., & Wendt-Suhl, G. (1976). The effect of chlordiazepoxide on the anticipation of electric shocks. *Psychopharmacology*, 48(3), 303-306.
- Bradford, W. D. (2010). The association between individual time preferences and health maintenance habits. *Medical Decision Making*, 30(1), 99-112.
- Bretteville-Jensen, A. L. (1999). Addiction and discounting. *Journal of Health Economics*, 18(4), 393-407.
- Breznitz, S. (1968). 'Incubation of threat' in a situation of conflicting expectations. *Psychological Reports*, 22(3), 755-756.
- Broekmans, S., Dobbels, F., Milisen, K., Morlion, B., & Vanderschueren, S. (2009). Medication adherence in patients with chronic non-malignant pain: Is there a problem? *European Journal of Pain*, 13(2), 115-123.
- Brouwer, W. B., Niessen, L. W., Postma, M. J., & Rutten, F. F. (2005). Need for differential discounting of costs and health effects in cost effectiveness analyses. *BMJ*, 331(7514), 446-448.
- Brown, H., & Adams, J. (2013). The role of time preference in smoking cessation: a longitudinal analysis of data from the Household Income and Labour Dynamics of Australia survey, 2001-08. *Addiction*, 108(1), 186-192.
- Brown, H. R., Zeidman, P., Smittenaar, P., Adams, R. A., McNab, F., Rutledge, R. B., & Dolan, R. J. (2014). Crowdsourcing for cognitive science – the utility of smartphones. *PLoS ONE*, 9(7), e100662.
- Burks, S., Carpenter, J., Götte, L., & Rustichini, A. (2012). Which measures of time preference best predict outcomes: Evidence from a large-scale field experiment. *Journal of Economic Behavior & Organization*, 84(1), 308-320.
- Businelle, M. S., McVay, M. A., Kendzor, D., & Copeland, A. (2010). A comparison of delay discounting among smokers, substance abusers, and non-dependent controls. *Drug and Alcohol Dependence*, 112(3), 247-250.
- Cairns, J. (1992). Discounting and health benefits: another perspective. *Health Economics*, 1(1), 76-79.
- Camchong, J., MacDonald, A. W., Nelson, B., Bell, C., Mueller, B. A., Specker, S., & Lim, K. O. (2011). Frontal hyperconnectivity related to discounting and reversal learning in cocaine subjects. *Biological Psychiatry*, 69(11), 1117-1123.
- Camerer, C., & Loewenstein, G. (2004). *Behavioral economics: Past, present, future*: Princeton, NJ, USA: Princeton University Press.

- Camerer, C., Loewenstein, G., & Prelec, D. (2005). Neuroeconomics: How neuroscience can inform economics. *Journal of Economic Literature*, 43(1), 9-64.
- Camerer, C. F., Loewenstein, G., & Rabin, M. (Eds.). (2004). *Advances in behavioral economics*. Princeton, NJ, USA: Princeton University Press.
- Campbell, B. A., & Kraeling, D. (1953). Response strength as a function of drive level and amount of drive reduction. *Journal of Experimental Psychology*, 45(2), 97.
- Carver, C. S., & White, T. L. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: the BIS/BAS scales. *Journal of Personality and Social Psychology*, 67(2), 319.
- Chang, Y.-P., Wray, L. O., Sessanna, L., & Peng, H.-L. (2011). Use of prescription opioid medication among community-dwelling older adults with noncancer chronic pain. *Journal of Addictions Nursing*, 22(1-2), 19-24.
- Chao, L.-W., Szrek, H., Pereira, N. S., & Pauly, M. V. (2009). Time preference and its relationship with age, health, and survival probability. *Judgment and Decision Making*, 4(1), 1-19.
- Chapman, G. B. (1996a). Expectations and preferences for sequences of health and money. *Organizational Behavior and Human Decision Processes*, 67(1), 59-75.
- Chapman, G. B. (1996b). Temporal discounting and utility for health and money. *Journal of Experimental Psychology: Learning Memory and Cognition*, 22(3), 771-791.
- Chapman, G. B. (2000). Preferences for improving and declining sequences of health outcomes. *Journal of Behavioral Decision Making*, 13(2), 203-218.
- Chapman, G. B. (2002). Your money or your health: time preferences and trading money for health. *Medical Decision Making*, 22(5), 410-416.
- Chapman, G. B. (2003). Time discounting of health outcomes. In G. Loewenstein, D. Read, & R. F. Baumeister (Eds.), *Time and Decision*. New York: Russell Sage Foundation.
- Chapman, G. B. (2005). Short-term cost for long-term benefit: time preference and cancer control. *Health Psychology*, 24(4 Suppl), S41-48.
- Chapman, G. B., Brewer, N. T., Coups, E. J., Brownlee, S., Leventhal, H., & Leventhal, E. A. (2001). Value for the future and preventive health behavior. *Journal of Experimental Psychology: Applied*, 7(3), 235-250.
- Chapman, G. B., & Coups, E. J. (1999). Time preferences and preventive health behavior: acceptance of the influenza vaccine. *Medical Decision Making*, 19(3), 307-314.
- Chapman, G. B., & Elstein, A. S. (1995). Valuing the future: temporal discounting of health and money. *Medical Decision Making*, 15(4), 373-386.
- Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117(3), 817-869.

- Chesson, H. W., Leichliter, J. S., Zimet, G. D., Rosenthal, S. L., Bernstein, D. I., & Fife, K. H. (2006). Discount rates and risky sexual behaviors among teenagers and young adults. *Journal of Risk and Uncertainty*, 32(3), 217-230.
- Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2004). For better or for worse: Default effects and 401 (k) savings behavior. In D. A. Wise (Ed.) *Perspectives on the Economics of Aging* (pp. 81-126): Chicago. IL, USA: University of Chicago Press (for NBER).
- Christensen-Szalanski, J. J. (1984). Discount functions and the measurement of patients' values. Women's decisions during childbirth. *Medical Decision Making*, 4(1), 47-58.
- Chung, S. H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 10(1), 67-74.
- Coffey, S. F., Gudleski, G. D., Saladin, M. E., & Brady, K. T. (2003). Impulsivity and rapid discounting of delayed hypothetical rewards in cocaine-dependent individuals. *Experimental and Clinical Psychopharmacology*, 11(1), 18-25.
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112, 155-159.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933-942.
- Colwill, R. M., & Rescorla, R. A. (1986). Associative structures in instrumental learning. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Volume 20, pp. 55-104): New York, NY, USA: Academic Press.
- Conner, M., & Norman, P. (2005a). *Predicting Health Behaviour* (2nd Edition). Maidenhead, UK: McGraw-Hill International.
- Conner, M., & Norman, P. (2005b). Predicting health behaviour: A social cognition approach. In M. Conner & P. Norman (Eds.), *Predicting Health Behaviour* (pp. 1-27). Maidenhead, UK: Open University Press.
- Cook, J. O., & Barnes, L. W. (1964). choice of delay of inevitable shock. *Journal of Abnormal Psychology*, 68, 669-672.
- Costa, P. T., & McCrae, R. R. (1990). Personality disorders and the five-factor model of personality. *Journal of Personality Disorders*, 4(4), 362-371.
- Crockett, M. J., Braams, B. R., Clark, L., Tobler, P. N., Robbins, T. W., & Kalenscher, T. (2013). Restricting temptations: neural mechanisms of precommitment. *Neuron*, 79(2), 391-401.
- Crockett, M. J., Clark, L., Hauser, M. D., & Robbins, T. W. (2010). Serotonin selectively influences moral judgment and behavior through effects on harm aversion. *Proceedings of the National Academy of Sciences: USA*, 107(40), 17433-17438.
- Crockett, M. J., Kurth-Nelson, Z., Siegel, J. Z., Dayan, P., & Dolan, R. J. (2014). Harm to others outweighs harm to self in moral decision making. *Proceedings of the National Academy of Sciences: USA*, 111(48), 17320-17325.

- Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS One*, 8(3), e57410.
- Cummings, M. K., Becker, M. H., & Maile, M.C. (1980). Bringing models together: an empirical approach to combining variables used to explain health actions. *Journal of Behavioral Medicine*, 3, 123-145.
- Curry, O., Price, M., & Price, J. (2008). Patience is a virtue: Cooperative people have lower discount rates. *Personality and Individual Differences*, 44, 780-785.
- Cushman, F., Gray, K., Gaffey, A., & Mendes, W. B. (2012). Simulating murder: the aversion to harmful action. *Emotion*, 12(1), 2-7.
- D'Amato, F. R., & Pavone, F. (1993). Endogenous opioids: a proximate reward mechanism for kin selection? *Behavioral and Neural Biology*, 60(1), 79-83.
- D'Amato, M., & Gumenik, W. (1960). Some effects of immediate versus randomly delayed shock on an instrumental response and cognitive processes. *The Journal of Abnormal and Social Psychology*, 60(1), 64-67.
- Dallery, J., & Raiff, B. R. (2007). Delay discounting predicts cigarette smoking in a laboratory model of abstinence reinforcement. *Psychopharmacology*, 190(4), 485-496.
- Danaei, G., Ding, E. L., Mozaffarian, D., Taylor, B., Rehm, J., Murray, C. J., & Ezzati, M. (2009). The preventable causes of death in the United States: comparative risk assessment of dietary, lifestyle, and metabolic risk factors. *PLoS Med*, 6(4), e1000058.
- Darzi, A. (2008). *High Quality Care for All: NHS Next Stage Review Final Report*. London, UK: Department of Health.
- Daugherty, J. R., & Brase, G. L. (2010). Taking time to be healthy: Predicting health behaviors with delay discounting and time perspective. *Personality and Individual Differences*, 48(2), 202-207.
- Davis, A. L., Miller, J. H., & Weber, R. A. (2011). Generosity across contexts: Working Paper Series, Department of Economics, University of Zurich.
- Davis, C., Patte, K., Curtis, C., & Reid, C. (2010). Immediate pleasures and future consequences. A neuropsychological study of binge eating and obesity. *Appetite*, 54(1), 208-213.
- Davis, M., Schlesinger, L. S., & Sorenson, C. (1989). Temporal specificity of fear conditioning: Effects of different conditioned stimulus--unconditioned stimulus intervals on the fear-potentiated startle effect. *Journal of Experimental Psychology: Animal Behavior Processes*, 15(4), 295-310.
- Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology*, 44(1), 113-126.
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models. *Decision making, affect, and learning: Attention and performance XXIII*, 23, 3-38.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204-1215.

- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704-1711
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876-879.
- Dayan, P. (2009). Prospective and retrospective temporal difference learning. *Network: Computation in Neural Systems*, 20(1), 32-46.
- Dayan, P. (2012). Models of value and choice. In R. Dolan & T. Sharot (Eds.), *Neuroscience of preference and choice*. (pp. 33-59.). London, UK: Elsevier.
- Dayan, P., & Balleine, B. W. (2002). Reward, motivation, and reinforcement learning. *Neuron*, 36(2), 285-298.
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4), 429-453.
- Dayan, P., Niv, Y., Seymour, B., & Daw, N. (2006). The misbehavior of value and the discipline of the will. *Neural Networks*, 19(8), 1153-1160.
- de Wit, H. (2009). Impulsivity as a determinant and consequence of drug use: a review of underlying processes. *Addiction Biology*, 14(1), 22-31.
- Debreu, G. (1991). The mathematization of economic theory. *The American Economic Review*, 81(1), 1-7.
- Del Giudice, M. (2014). An evolutionary life history framework for psychopathology. *Psychological Inquiry*, 25(3-4), 261-300.
- Dennhardt, A. A., & Murphy, J. G. (2011). Associations between depression, distress tolerance, delay discounting, and alcohol-related problems in European American and African American college students. *Psychology of addictive behaviors*, 25(4), 595-604.
- Department of Commerce and Labor, Bureau of the Census (1906). Mortality Statistics, 1900 to 1904. US Department of Commerce and Labor: Washington, DC, USA.
- DeWall, C. N., & Baumeister, R. F. (2006). Alone but feeling no pain: Effects of social exclusion on physical pain tolerance and pain threshold, affective forecasting, and interpersonal empathy. *Journal of Personality and Social Psychology*, 91(1), 1-15.
- DeWall, C. N., MacDonald, G., Webster, G. D., Masten, C. L., Baumeister, R. F., Powell, C., . . . Tice, D. M. (2010). Acetaminophen reduces social pain behavioral and neural evidence. *Psychological Science*, 21(7), 931-937.
- Dickinson, A., & Balleine, B. (1994). Motivational control of goal-directed action. *Learning & Behavior*, 22(1), 1-18.
- Dickinson, A., Balleine, B., Watt, A., Gonzalez, F., & Boakes, R. (1995). Motivational control after extended instrumental training. *Learning & Behavior*, 23(2), 197-206.
- Dierst-Davies, R., Reback, C. J., Peck, J. A., Nuño, M., Kamien, J. B., & Amass, L. (2011). Delay-discounting among homeless, out-of-treatment, substance-dependent men who have sex with men. *The American Journal of Drug and Alcohol Abuse*, 37(2), 93-97.

- Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review*, 100(3), 1238-1260.
- Dolan, P., & Gudex, C. (1995). Time preference, duration and health state valuations. *Health Economics*, 4(4), 289-299.
- Dolan, P., Hallsworth, M., Halpern, D., King, D., Metcalfe, R., & Vlaev, I. (2012). Influencing behaviour: The mindspace way. *Journal of Economic Psychology*, 33(1), 264-277.
- Dolan, P., Hallsworth, M., Halpern, D., King, D., & Vlaev, I. (2009). *Mindspace: Influencing behaviour through public policy*. London, UK: Cabinet Office, UK.
- Dolan, R. J., & Dayan, P. (2013). Goals and Habits in the Brain. *Neuron*, 80(2), 312-325.
- Dom, G., De Wilde, B., Hulstijn, W., Van Den Brink, W., & Sabbe, B. (2006). Decision-Making Deficits in Alcohol-Dependent Patients With and Without Comorbid Personality Disorder. *Alcoholism: Clinical and Experimental Research*, 30(10), 1670-1677.
- Domjan, M. (2003). *The Principles of Learning and Behaviour*. (5th Edition). Belmont, USA: Wadsworth.
- Duncan, G. H., Bushnell, M. C., & Lavigne, G. J. (1989). Comparison of verbal and visual analogue scales for measuring the intensity and unpleasantness of experimental pain. *Pain*, 37(3), 295-303.
- Eckel, C., Johnson, C., & Montmarquette, C. (2005). Savings Decisions of the Working Poor: Short- and Long-Term Horizons. In J. Carpenter, G. W. Harrison, & J. A. List (Eds.), *Field Experiments in Economics, Research in Experimental Economics* (Vol. 10, pp. 219-260). Greenwich, CT: JAI Press.
- Eisenberg, D. T., MacKillop, J., Modi, M., Beauchemin, J., Dang, D., Lisman, S. A., . . . Wilson, D. S. (2007). Examining impulsivity as an endophenotype using a behavioral approach: a DRD2 TaqI A and DRD4 48-bp VNTR association study. *Behavior and Brain Functions*, 3(2): doi:10.1186/1744-9081-3-2
- Eisenberger, N. I. (2012). Broken hearts and broken bones: a neural perspective on the similarities between social and physical pain. *Current Directions in Psychological Science*, 21(1), 42-47.
- Eisenberger, N. I., Lieberman, M. D., & Williams, K. D. (2003). Does rejection hurt? An fMRI study of social exclusion. *Science*, 302(5643), 290-292.
- Elliott, R., Agnew, Z., & Deakin, J. (2008). Medial orbitofrontal cortex codes relative rather than absolute value of financial rewards in humans. *European Journal of Neuroscience*, 27(9), 2213-2218.
- Ellis, B. J., Del Giudice, M., Dishion, T. J., Figueredo, A. J., Gray, P., Griskevicius, V., . . . Volk, A. A. (2012). The evolutionary basis of risky adolescent behavior: implications for science, policy, and practice. *Developmental Psychology*, 48(3), 598-623.

- Elster, J., & Loewenstein, G. (1992). Utility from memory and anticipation. In E. J. & G. Loewenstein (Eds.), *Choice over Time* (pp. 213-234). NY, USA: Russell Sage Foundation.
- Engel, C. (2011). Dictator games: a meta study. *Experimental Economics*, 14(4), 583-610.
- Eppinger, B., Walter, M., Heekeren, H. R., & Li, S.-C. (2013). Of goals and habits: age-related and individual differences in goal-directed decision-making. *Frontiers in Neuroscience*, 7: 253; doi: 10.3389/fnins.2013.00253
- Epstein, L. H., Richards, J. B., Saad, F. G., Paluch, R. A., Roemmich, J. N., & Lerman, C. (2003). Comparison between two measures of delay discounting in smokers. *Experimental and Clinical Psychopharmacology*, 11(2), 131-138.
- Epstein, S., Epstein, A., & Roupelian. (1970). Heart rate and skin conductance during experimentally induced anxiety: The effect of uncertainty about receiving a noxious stimulus. *Journal of Personality and Social Psychology*, 16(1), 20-28.
- Ersner-Hershfield, H., Wimmer, G. E., & Knutson, B. (2009). Saving for the future self: Neural measures of future self-continuity predict temporal discounting. *Social Cognitive and Affective Neuroscience*, 4(1), 85-92.
- Estes, W. K. (1948). Discriminative conditioning; effects of a Pavlovian conditioned stimulus upon a subsequently established operant response. *Journal of Experimental Psychology*, 38(2), 173-177.
- Estes, W. K., & Skinner, B. F. (1941). Some quantitative properties of anxiety. *Journal of Experimental Psychology*, 29(5), 390-400.
- Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition advancing the debate. *Perspectives on Psychological Science*, 8(3), 223-241.
- Everitt, B. J., Belin, D., Economidou, D., Pelloux, Y., Dalley, J. W., & Robbins, T. W. (2008). Review. Neural mechanisms underlying the vulnerability to develop compulsive drug-seeking habits and addiction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, , 363(1507), 3125-3135.
- Everitt, B. J., & Robbins, T. W. (2005). Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nature Neuroscience*, 8(11), 1481-1489.
- Farrell, S., & Lewandowsky, S. (2010). Computational models as aids to better reasoning in psychology. *Current Directions in Psychological Science*, 19(5), 329-335.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3), 817-868.
- Fehr, E., & Schmidt, K. M. (2006). The economics of fairness, reciprocity and altruism – experimental evidence and new theories. In K. Serge-Christophe & Y. Jean Mercier (Eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity* (Vol. 1, pp. 615-691): Amsterdam, Netherlands: Elsevier.
- FeldmanHall, O., Mobbs, D., Evans, D., Hiscox, L., Navrady, L., & Dalgleish, T. (2012). What we say and what we do: the relationship between real and hypothetical moral choices. *Cognition*, 123(3), 434-441.

- Fenigstein, A., & Vanable, P. A. (1992). Paranoia and self-consciousness. *Journal of Personality and Social Psychology*, 62(1), 129-138.
- Fernie, G., Cole, J. C., Goudie, A. J., & Field, M. (2010). Risk-taking but not response inhibition or delay discounting predict alcohol consumption in social drinkers. *Drug and Alcohol Dependence*, 112(1-2), 54-61.
- Field, M., Christiansen, P., Cole, J., & Goudie, A. (2007). Delay discounting and the alcohol Stroop in heavy drinking adolescents. *Addiction*, 102(4), 579-586.
- Field, M., Santarcangelo, M., Sumnall, H., Goudie, A., & Cole, J. (2006). Delay discounting and the behavioural economics of cigarette purchases in smokers: the effects of nicotine deprivation. *Psychopharmacology*, 186(2), 255-263.
- Fields, S., Leraas, K., Collins, C., & Reynolds, B. (2009). Delay discounting as a mediator of the relationship between perceived stress and cigarette smoking status in adolescents. *Behavioral Pharmacology*, 20(5-6), 455-460.
- Fields, S., Sabet, M., Peal, A., & Reynolds, B. (2011). Relationship between weight status and delay discounting in a sample of adolescent cigarette smokers. *Behavioural Pharmacology*, 22(3), 266-268.
- Fiske, S. T., & Taylor, S. E. (2013). *Social Cognition* (2nd Edition). London, UK: SAGE.
- Fliessbach, K., Weber, B., Trautner, P., Dohmen, T., Sunde, U., Elger, C. E., & Falk, A. (2007). Social comparison affects reward-related brain activity in the human ventral striatum. *Science*, 318(5854), 1305-1308.
- Folkins, C. H. (1970). Temporal factors and the cognitive mediators of stress reaction. *Journal of Personality and Social Psychology*, 14(2), 173-184.
- Ford, E. S., Merritt, R. K., Heath, G. W., Powell, K. E., Washburn, R. A., Kriska, A., & Haile, G. (1991). Physical activity behaviors in lower and higher socioeconomic status populations. *American Journal of Epidemiology*, 133(12), 1246-1256.
- Forder, J., Jnes, K., Glendinning, J. C., Welch, E., Baxter, K., Davidson, J., . . . Dolan, P. (2012). *Evaluation of the personal health budget pilot programme*. London, UK: Department of Health.
- Forsythe, R., Horowitz, J. L., Savin, N. E., & Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic behavior*, 6(3), 347-369.
- Francis, L. A. (2009). Self-regulation and rapid weight gain in children from age 3 to 12 years. *Archives of Pediatrics & Adolescent Medicine*, 163(4), 297-302.
- Franzini, L. R. (1970). Magnitude estimations of the averseness of the interval preceding shock. *Journal of Experimental Psychology*, 84(3), 526-528.
- Frederick, S., & Loewenstein, G. (2008). Conflicting motives in evaluations of sequences. *Journal of Risk and Uncertainty*, 37(2-3), 221-235.
- Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2), 351-401.
- Friston, K. J., Stephan, K. E., Lund, T. E., Morcom, A., & Kiebel, S. (2005). Mixed-effects and fMRI studies. *Neuroimage*, 24(1), 244-252.
- Fuchs, V. R. (1982). *Time preference and health: an exploratory study*. Cambridge, MA, USA.

- Ganiats, T. G., Carson, R. T., Hamm, R. M., Cantor, S. B., Sumner, W., Spann, S. J., . . . Miller, C. (2000). Population-based time preferences for future health outcomes. *Medical Decision Making*, 20(3), 263-270.
- Gigerenzer, G. (2008). Why heuristics work. *Perspectives on Psychological Science*, 3(1), 20-29.
- Gigerenzer, G., Hertwig, R., & Pachur, T. (2011). *Heuristics: The foundations of adaptive behavior*. Oxford, UK: Oxford University Press.
- Giordano, L. A., Bickel, W. K., Loewenstein, G., Jacobs, E. A., Marsch, L., & Badger, G. J. (2002). Mild opioid deprivation increases the degree that opioid-dependent outpatients discount delayed heroin and money. *Psychopharmacology*, 163(2), 174-182.
- Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, 66(4), 585-595.
- Glimcher, P. W. (2004). *Decisions, uncertainty, and the brain: The science of neuroeconomics*. Cambridge, MA, USA: MIT press.
- Glimcher, P. W., Camerer, C., Fehr, E., & Poldrack, R. A. (2009). Introduction: A Brief History of Neuroeconomics. In P. W. Glimcher, C. Camerer, E. Fehr, & R. A. Poldrack (Eds.), *Neuroeconomics: Decision making and the brain* (1st Edition ed., pp. 1-12). London, UK: Academic Press.
- Glimcher, P. W., Dorris, M. C., & Bayer, H. M. (2005). Physiological utility theory and the neuroeconomics of choice. *Games and Economic behavior*, 52(2), 213-256.
- Glimcher, P. W., & Rustichini, A. (2004). Neuroeconomics: the consilience of brain and decision. *Science*, 306(5695), 447-452.
- Gollwitzer, P. M. (1993). Goal achievement: The role of intentions. *European Review of Social Psychology*, 4(1), 141-185.
- Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, 38, 69-119.
- Gray, A. (1982). Inequalities in health. The Black Report: a summary and comment. *International Journal of Health Services: Planning, Administration, Evaluation*, 12(3), 349-380.
- Gray, K. (2012). The power of good intentions perceived benevolence soothes pain, increases pleasure, and improves taste. *Social Psychological and Personality Science*, 3(5), 639-645.
- Gray, K., & Wegner, D. M. (2008). The sting of intentional pain. *Psychological Science*, 19(12), 1260-1262.
- Green, L., Fristoe, N., Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review*, 1(3), 383-389.
- Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, 130(5), 769-792.

- Green, L., Myerson, J., & Ostaszewski, P. (1999). Amount of reward has opposite effects on the discounting of delayed and probabilistic outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(2), 418-427.
- Grillon, C., Ameli, R., Merikangas, K., Woods, S. W., & Davis, M. (1993). Measuring the time course of anticipatory anxiety using the fear-potentiated startle reflex. *Psychophysiology*, 30(4), 340-346.
- Grossman, M. (1972). On the Concept of Health Capital and the Demand for Health. *Journal of Political Economy*, 80(2), 223-255.
- Gunnarsdottir, S., Serlin, R. C., & Ward, S. (2005). Patient-related barriers to pain management: the Icelandic Barriers Questionnaire II. *Journal of Pain and Symptom Management*, 29(3), 273-285.
- Haefner, D. P. (1974). The health belief model and preventive dental behavior. *Health Education & Behavior*, 2(4), 420-432.
- Han, R., & Takahashi, T. (2012). Psychophysics of time perception and valuation in temporal discounting of gain and loss. *Physica A: Statistical Mechanics and its Applications*, 391(24), 6568-6576.
- Hare, R. D. (1966a). Preference for delay of shock as a function of its intensity and probability. *Psychonomic Science*, 5(10), 393-394.
- Hare, R. D. (1966c). Psychopathy and choice of immediate versus delayed punishment. *Journal of Abnormal Psychology*, 71(1), 25.
- Hare, R. D., Krebs, D. L., Creighton, T. D., & Petrusic, W. M. (1966). Latency of self-administered shock as a function of its intensity and probability. *Psychonomic Science*, 6(2), 79-80.
- Harris, A. C., & Madden, G. J. (2002). Delay discounting and performance on the prisoner's dilemma game. *Psychological Record*, 52(4), 429-440.
- Harrison, J. A., Mullen, P. D., & Green, L. W. (1992). A meta-analysis of studies of the health belief model with adults. *Health Education Research*, 7(1), 107-116.
- Hassabis, D., Kumaran, D., Vann, S. D., & Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences USA*, 104(5), 1726-1731.
- Hecht, D., Walsh, V., & Lavidor, M. (2013). Bi-frontal direct current stimulation affects delay discounting choices. *Cognitive Neuroscience*, 4(1), 7-11.
- Heil, S. H., Johnson, M. W., Higgins, S. T., & Bickel, W. K. (2006). Delay discounting in currently using and currently abstinent cocaine-dependent outpatients and non-drug-using matched controls. *Addictive Behavior*, 31(7), 1290-1294.
- Hein, G., Lamm, C., Brodbeck, C., & Singer, T. (2011). Skin conductance response to the pain of others predicts later costly helping. *PLoS One*, 6(8), e22759.
- Heinz, A. J., Peters, E. N., Boden, M. T., & Bonn-Miller, M. O. (2013). A comprehensive examination of delay discounting in a clinical sample of Cannabis-dependent military veterans making a self-guided quit attempt. *Experimental and Clinical Psychopharmacology*, 21(1), 55-65.

- Heyman, G. M., & Gibb, S. P. (2006). Delay discounting in college cigarette chippers. *Behavioural Pharmacology*, 17(8), 669-679.
- Hochbaum, G. M. (1958). *Public participation in medical screening programs: A socio-psychological study*. US Department of Health, Education, and Welfare, Public Health Service, Bureau of State Services, Division of Special Health Services, Tuberculosis Program.
- Hoffman, W. F., Moore, M., Templin, R., McFarland, B., Hitzemann, R. J., & Mitchell, S. H. (2006). Neuropsychological function and delay discounting in methamphetamine-dependent individuals. *Psychopharmacology*, 188(2), 162-170.
- Hofmann, W., Friese, M., & Strack, F. (2009). Impulse and self-control from a dual-systems perspective. *Perspectives on Psychological Science*, 4(2), 162-176.
- Holmes, A., & Friston, K. (1998). Generalisability, random effects & population inference. *Neuroimage*, 7, S754.
- Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica*, 159-174.
- Hoyert, D. L., Kochanek, K. D., & Murphy, S. L. (1999). *Deaths: final data for 1997*. Hyattsville, Maryland: US Department of Health and Human Services, Public Health Service, CDC, National Center for Health Statistics.
- Huckans, M., Seelye, A., Woodhouse, J., Parcel, T., Mull, L., Schwartz, D., . . . Loftis, J. (2011). Discounting of delayed rewards and executive dysfunction in individuals infected with hepatitis C. *Journal of Clinical and Experimental Neuropsychology*, 33(2), 176-186.
- Huys, Q. J., Cools, R., Gölzer, M., Friedel, E., Heinz, A., Dolan, R. J., & Dayan, P. (2011). Disentangling the roles of approach, activation and valence in instrumental and pavlovian responding. *PLoS Computational Biology*, 7(4), e1002028.
- Huys, Q. J., Eshel, N., O'Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, 8(3), e1002410.
- Idler, E. L., & Benyamini, Y. (1997). Self-rated health and mortality: a review of twenty-seven community studies. *Journal of Health and Social Behavior*, 38(1), 21-37.
- Ikeda, S., Kang, M.-I., & Ohtake, F. (2010). Hyperbolic discounting, the sign effect, and the body mass index. *Journal of Health Economics*, 29(2), 268-284.
- Jackson, P. L., Meltzoff, A. N., & Decety, J. (2005). How do we perceive the pain of others? A window into the neural processes involved in empathy. *Neuroimage*, 24(3), 771-779.
- Jacobsen, R., Møldrup, C., Christrup, L., & Sjøgren, P. (2009). Patient-related barriers to cancer pain management: a systematic exploratory review. *Scandinavian Journal of Caring Sciences*, 23(1), 190-208.

- Johnson, A., van der Meer, M. A., & Redish, A. D. (2007). Integrating hippocampus and striatum in decision-making. *Current opinion in neurobiology*, 17(6), 692-697.
- Johnson, M. W. (2012). An efficient operant choice procedure for assessing delay discounting in humans: initial validation in cocaine-dependent and control individuals. *Experimental and Clinical Psychopharmacology*, 20(3), 191-204.
- Johnson, M. W., & Bickel, W. K. (2002). Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the Experimental Analysis of Behavior*, 77(2), 129.
- Johnson, M. W., Bickel, W. K., & Baker, F. (2007). Moderate drug use and delay discounting: a comparison of heavy, light, and never smokers. *Experimental and Clinical Psychopharmacology*, 15(2), 187-194.
- Johnson, M. W., Bickel, W. K., Baker, F., Moore, B. A., Badger, G. J., & Budney, A. J. (2010). Delay discounting in current and former marijuana-dependent individuals. *Experimental and Clinical Psychopharmacology*, 18(1), 99-107.
- Johnson, S., Kotlikoff, L. J., & Samuelson, W. (1987). *Can people compute? An experimental test of the life cycle consumption model*. Cambridge, MA, USA: National Bureau of Economic Research.
- Johnson, W. G., Parry, W., & Drabman, R. S. (1978). The performance of obese and normal size children on a delay of gratification task. *Addictive Behavior*, 3(3-4), 205-208.
- Jones, B., & Rachlin, H. (2006). Social discounting. *Psychological Science*, 17(4), 283-286.
- Jones, B. A., & Rachlin, H. (2009). Delay, probability, and social discounting in a public goods game. *Journal of the Experimental Analysis of Behavior*, 91(1), 61-73.
- Kable, J. W., & Glimcher, P. W. (2010). An “as soon as possible” effect in human intertemporal decision making: behavioral evidence and neural mechanisms. *J Neurophysiol*, 103(5), 2513-2531.
- Kahneman, D. (2011). *Thinking, fast and slow*: Macmillan.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1986). Fairness and the assumptions of economics. *Journal of business*, S285-S300.
- Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. *The Journal of Economic Perspectives*, 5(1), 193-206. doi: 10.2307/1942711
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 263-291.
- Kahneman, D., & Tversky, A. (2000). *Choices, values, and frames*: Cambridge University Press.
- Kalenscher, T., & Pennartz, C. M. (2008). Is a bird in the hand worth two in the future? The neuroeconomics of intertemporal decision-making. *Prog Neurobiol*, 84(3), 284-315. doi: S0301-0082(07)00215-8 [pii] 10.1016/j.pneurobio.2007.11.004

- Kang, M. I., & Ikeda, S. (2013). Time Discounting, Present Biases, and Health-Related Behavior. ISER Discussion Paper No. 885. Available at SSRN: <http://ssrn.com/abstract=2324659>.
- Kass, R. E., & Wasserman, L. (1995). A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *Journal of the American Statistical Association*, 90, 928-934.
- Keramati, M., Dezfouli, A., & Piray, P. (2012). Understanding Addiction as a Pathological State of Multiple Decision Making Processes: A Neurocomputational Perspective. 205-233. doi: 10.1007/978-1-4614-0751-5_8
- Khwaja, A., Silverman, D., & Sloan, F. (2007). Time preference, time discounting, and smoking decisions. *J Health Econ*, 26(5), 927-949. doi: 10.1016/j.jhealeco.2007.02.004
- King, D., Greaves, F., Vlaev, I., & Darzi, A. (2013). Approaches Based On Behavioral Economics Could Help Nudge Patients And Providers Toward Lower Health Spending Growth. *Health Affairs*, 32(4), 661-668. doi: 10.1377/hlthaff.2012.1348
- Kirby, K. N. (2009). One-year temporal stability of delay-discount rates. *Psychonomic Bulletin & Review*, 16(3), 457-462.
- Kirby, K. N., & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science*, 6(2), 83-89.
- Kirby, K. N., & Maraković, N. N. (1995). Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human Decision Processes*, 64(1), 22-30.
- Kirby, K. N., & Petry, N. M. (2004). Heroin and cocaine abusers have higher discount rates for delayed rewards than alcoholics or non-drug-using controls. *Addiction*, 99(4), 461-471. doi: 10.1111/j.1360-0443.2003.00669.x
- Kirby, K. N., Petry, N. M., & Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *J Exp Psychol Gen*, 128(1), 78-87.
- Kobiella, A., Ripke, S., Kroemer, N. B., Vollmert, C., Vollstädt-Klein, S., Ulshöfer, D. E., & Smolka, M. N. (2014). Acute and chronic nicotine effects on behaviour and brain activation during intertemporal decision making. *Addict Biol*, 19(5), 918-930.
- Koffarnus, M. N., Jarmolowicz, D. P., Mueller, E. T., & Bickel, W. K. (2013). Changing delay discounting in the light of the competing neurobehavioral decision systems theory: A review. *Journal of the Experimental Analysis of Behavior*, 99(1), 32-57.
- Kollins, S. H. (2003). Delay discounting is associated with substance use in college students. *Addictive Behavior*, 28(6), 1167-1173.
- Koyama, T., Tanaka, Y. Z., & Mikami, A. (1998). Nociceptive neurons in the macaque anterior cingulate activate during anticipation of pain. *Neuroreport*, 9(11), 2663-2667.

- Kraemer, H. C., Gullion, C. M., Rush, A. J., Frank, E., & Kupfer, D. J. (1994). Can state and trait variables be disentangled? A methodological framework for psychiatric disorders. *Psychiatry Research*, 52(1), 55-69.
- Krahé, C., Springer, A., Weinman, J. A., & Fotopoulou, A. (2013). The social modulation of pain: others as predictive signals of salience—a systematic review. *Frontiers in Human Neuroscience*, 7: 386: doi: 10.3389/fnhum.2013.00386
- Krishnan-Sarin, S., Reynolds, B., Duhig, A. M., Smith, A., Liss, T., McFetridge, A., . . . Potenza, M. N. (2007). Behavioral impulsivity predicts treatment outcome in a smoking cessation program for adolescent smokers. *Drug and Alcohol Dependence*, 88(1), 79-82.
- Kulendran, M., Vlaev, I., Sugden, C., King, D., Ashrafian, H., Gately, P., & Darzi, A. (2013). Neuropsychological assessment as a predictor of weight loss in obese adolescents. *International Journal of Obesity*, 38, 502-514.
- Kulendran, M., Vlaev, I., Sugden, C., King, D., Hussein, M., Le Vay, R., . . . Darzi, A. (2013). Inhibitory control and perception of the future health in obese adolescents and normal weight adolescents. *The Lancet*, 382, Supplement 3(0), S55.
- Kurth-Nelson, Z., & Redish, A. D. (2010). A reinforcement learning model of precommitment in decision making. *Frontiers in Behavioral Neuroscience*, 4: 184 doi: 10.3389/fnbeh.2010.00184.
- Kurth-Nelson, Z., & Redish, A. D. (2012). Don't Let Me Do That!—Models of Precommitment. *Frontiers in neuroscience*, 6: 138 doi: 10.3389/fnins.2012.00138
- Lagarias, J. C., Reeds, J. A., Wright, M. H., & Wright, P. E. (1998). Convergence properties of the Nelder--Mead simplex method in low dimensions. *SIAM Journal on Optimization*, 9(1), 112-147.
- Lagorio, C. H., & Madden, G. J. (2005). Delay discounting of real and hypothetical rewards III: steady-state assessments, forced-choice trials, and all real rewards. *Behavioural Processes*, 69(2), 173-187.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2), 443-478.
- Lally, P., Chipperfield, A., & Wardle, J. (2008). Healthy habits: efficacy of simple advice on weight control based on a habit-formation model. *International Journal of Obesity*, 32(4), 700-707.
- Lally, P., Van Jaarsveld, C. H., Potts, H. W., & Wardle, J. (2010). How are habits formed: Modelling habit formation in the real world. *European Journal of Social Psychology*, 40(6), 998-1009.
- Landes, R. D., Christensen, D. R., & Bickel, W. K. (2012). Delay discounting decreases in those completing treatment for opioid dependence. *Experimental and Clinical Psychopharmacology*, 20(4), 302-309.
- Langford, D. J., Tuttle, A. H., Brown, K., Deschenes, S., Fischer, D. B., Mutso, A., . . . Sternberg, W. F. (2010). Social approach to pain in laboratory mice. *Social Neuroscience*, 5(2), 163-170.

- Lazaro, A., Barberan, R., & Rubio, E. (2001). Private and social time preferences for health and money: an empirical estimation. *Health Economics*, 10(4), 351-356.
- Lawrence, K. A., Allen, J. S., & Chanen, A. M. (2010). Impulsivity in borderline personality disorder: reward-based decision-making and its relationship to emotional distress. *Journal of Personality Disorders*, 24(6), 785-799.
- Ledgerwood, D. M., Alessi, S. M., Phoenix, N., & Petry, N. M. (2009). Behavioral assessment of impulsivity in pathological gamblers with and without substance use disorder histories versus healthy controls. *Drug and Alcohol Dependence*, 105(1), 89-96.
- Leigh, J. P. (1986). Accounting for tastes: correlates of risk and time preferences. *Journal of Post Keynesian Economics*, 9(1), 17-31.
- Leitão, M., Guedes, Á., Yamamoto, M. E., & Lopes, F. D. A. (2013). Do people adjust career choices according to socioeconomic conditions? An evolutionary analysis of future discounting. *Psychology & Neuroscience*, 6(3), 383-390.
- Lettau, M., & Uhlig, H. (1999). Rules of thumb versus dynamic programming. *American Economic Review*, 89, 148-174.
- Levin, D. N., Cleeland, C. S., & Dar, R. (1985). Public attitudes toward cancer pain. *Cancer*, 56(9), 2337-2339.
- Lewandowsky, S., & Farrell, S. (2010). *Computational modeling in cognition: Principles and Practice*: Thousand Oaks, CA, USA: Sage.
- Lewandowsky, S., & Farrell, S. (2011). Maximum likelihood estimation *Computational Modeling in Cognition: Principles and Practice* (pp. 109-145). Thousand Oaks, CA, USA: Sage
- Lewis, J., & Christopher, J. (1989). Childhood trauma in borderline personality disorder. *American Journal of Psychiatry*, 146 (4), 490-495.
- Lewis, E. T., Combs, A., & Trafton, J. A. (2010). Reasons for under-use of prescribed opioid medications by patients in pain. *Pain Medicine*, 11(6), 861-871.
- Lipscomb, J. (1989). Time preference for health in cost-effectiveness analysis. *Medical Care*, 27(3 Suppl), S233-253.
- List, J. A. (2007). On the interpretation of giving in dictator games. *Journal of Political Economy*, 115(3), 482-493.
- Lloyd, D., Di Pellegrino, G., & Roberts, N. (2004). Vicarious responses to pain in anterior cingulate cortex: is empathy a multisensory issue? *Cognitive, Affective, & Behavioral Neuroscience*, 4(2), 270-278.
- Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal*, 97(387), 666-684.
- Loewenstein, G. (1988). Frames of mind in intertemporal choice. *Management Science*, 34(2), 200-214.
- Loewenstein, G. (1996). Out of control: visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65(3), 272-292.

- Loewenstein, G., & O'Donoghue, T. (2004). *Animal Spirits: Affective and Deliberative Processes in Economic Behavior*. Cornell University, Center for Analytic Economics Working Paper 2004-14.
- Loewenstein, G., & Prelec, D. (1991). Negative time preference. *The American Economic Review*, 81(2), 347-352.
- Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, 107(2), 573-597.
- Loewenstein, G., & Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review*, 100(1), 91.
- Loewenstein, G., & Sicherman, N. (1991). Do workers prefer increasing wage profiles? *Journal of Labor Economics*, 9(1), 67-84.
- Loewenstein, G., Thompson, L., & Bazerman, M. H. (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology*, 57(3), 426-441.
- Lucantonio, F., Caprioli, D., & Schoenbaum, G. (2014). Transition from 'model-based' to 'model-free' behavioral control in addiction: involvement of the orbitofrontal cortex and dorsolateral striatum. *Neuropharmacology*, 76, 407-415.
- Luhmann, C. C. (2013). Discounting of delayed rewards is not hyperbolic. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(4), 1274-1279.
- Lynch, J. W., Kaplan, G. A., & Salonen, J. T. (1997). Why do poor people behave poorly? Variation in adult health behaviours and psychosocial characteristics by stages of the socioeconomic lifecourse. *Social Science & Medicine*, 44(6), 809-819.
- MacDonald, G., & Leary, M. R. (2005). Why does social exclusion hurt? The relationship between social and physical pain. *Psychological bulletin*, 131(2), 202-223.
- MacKeigan, L. D., Larson, L. N., Draugalis, J. R., Bootman, J. L., & Burns, L. R. (1993). Time preference for health gains versus health losses. *Pharmacoeconomics*, 3(5), 374-386.
- MacKillop, J. (2013). Integrating behavioral economics and behavioral genetics: delayed reward discounting as an endophenotype for addictive disorders. *Journal of the Experimental Analysis of Behavior*, 99(1), 14-31.
- MacKillop, J., Amlung, M. T., Few, L. R., Ray, L. A., Sweet, L. H., & Munafò, M. R. (2011). Delayed reward discounting and addictive behavior: a meta-analysis. *Psychopharmacology*, 216(3), 305-321.
- MacKillop, J., & Kahler, C. W. (2009). Delayed reward discounting predicts treatment response for heavy drinkers receiving smoking cessation treatment. *Drug and Alcohol Dependence*, 104(3), 197-203.
- MacKillop, J., Mattson, R. E., Anderson Mackillop, E. J., Castelda, B. A., & Donovan, P. J. (2007). Multidimensional assessment of impulsivity in undergraduate hazardous drinkers and controls. *Journal of Studies on Alcohol and Drugs*, 68(6), 785-788.

- MacKillop, J., Miranda, R., Jr., Monti, P. M., Ray, L. A., Murphy, J. G., Rohsenow, D. J., . . . Gwaltney, C. J. (2010). Alcohol demand, delayed reward discounting, and craving in relation to drinking and alcohol use disorders. *Journal of Abnormal Psychology, 119*(1), 106-114.
- Madden, G. J., Begotka, A. M., Raiff, B. R., & Kastern, L. L. (2003). Delay discounting of real and hypothetical rewards. *Experimental and Clinical Psychopharmacology, 11*(2), 139.
- Madden, G. J., Petry, N. M., Badger, G. J., & Bickel, W. K. (1997). Impulsive and self-control choices in opioid-dependent patients and non-drug-using control patients: Drug and monetary rewards. *Experimental and Clinical Psychopharmacology, 5*(3), 256-262.
- Madden, G. J., Raiff, B. R., Lagorio, C. H., Begotka, A. M., Mueller, A. M., Hehli, D. J., & Wegener, A. A. (2004). Delay discounting of potentially real and hypothetical rewards: II. Between-and within-subject comparisons. *Experimental and Clinical Psychopharmacology, 12*(4), 251.
- Mancini, A., Betti, V., Panasiti, M. S., Pavone, E. F., & Aglioti, S. M. (2011). Suffering makes you egoist: acute pain increases acceptance rates and reduces fairness during a bilateral ultimatum game. *PLoS One, 6*(10), e26008.
- Mancini, A., Betti, V., Panasiti, M. S., Pavone, E. F., & Aglioti, S. M. (2014). Perceiving monetary loss as due to inequity reduces behavioral and cortical responses to pain. *European Journal of Neuroscience, 40*(2), 2378-2388.
- Mariano, T., Bannerman, D., McHugh, S., Preston, T., Rudebeck, P., Rudebeck, S., . . . Baxter, M. (2009). Impulsive choice in hippocampal but not orbitofrontal cortex-lesioned rats on a nonspatial decision-making maze task. *European Journal of Neuroscience, 30*(3), 472-484.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman.
- Mason, W., & Suri, S. (2012). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior research methods, 44*(1), 1-23.
- Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. *Frontiers in Human Neuroscience, 5*, 39.
- Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. Commons, Nevin & Rachlin (Ed.), *Quantitative Analysis of Behavior* (pp. 55-73). Hillsdale, NJ: Erlbaum.
- McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience, 27*(21), 5796-5804.
- McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science, 306*(5695), 503-507.

- McDannald, M. A., Lucantonio, F., Burke, K. A., Niv, Y., & Schoenbaum, G. (2011). Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. *Journal of Neuroscience*, *31*(7), 2700-2705.
- McNab, F., Zeidman, P., Rutledge, R. B., Smittenaar, P., Brown, H. R., Adams, R. A., & Dolan, R. J. (2015). Age-related changes in working memory and the ability to ignore distraction. *Proceedings of the National Academy of Sciences USA*, *112*(20), 6515-6518
- McParland, J. L., & Eccleston, C. (2013). "It's Not Fair" Social Justice Appraisals in the Context of Chronic Pain. *Current Directions in Psychological Science*, *22*(6), 484-489.
- Melanko, S., & Larkin, K. T. (2012). Preference for immediate reinforcement over delayed reinforcement: relation between delay discounting and health behavior. *Journal of Behavioral Medicine* *36*(1): 34:43: Retracted Article.
- Metcalf, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: dynamics of willpower. *Psychological Review*, *106*(1), 3-19.
- Miaskowski, C., Dodd, M. J., West, C., Paul, S. M., Tripathy, D., Koo, P., & Schumacher, K. (2001). Lack of adherence with the analgesic regimen: a significant barrier to effective cancer pain management. *Journal of Clinical Oncology*, *19*(23), 4275-4279.
- Milgram, S. (1963). Behavioral study of obedience. *Journal of Abnormal and Social Psychology*, *67*(4), 371-378.
- Milgram, S. (1965). Some conditions of obedience and disobedience to authority. *Human Relations*, *18*(1), 57-76.
- Milne, S., Sheeran, P., & Orbell, S. (2000). Prediction and intervention in health - related behavior: a meta - analytic review of Protection Motivation Theory. *Journal of Applied Social Psychology*, *30*(1), 106-143.
- Miron, D., Duncan, G. H., & Bushnell, M. C. (1989). Effects of attention on the intensity and unpleasantness of thermal pain. *Pain*, *39*(3), 345-352.
- Mitchell, J. M., Fields, H. L., D'Esposito, M., & Boettiger, C. A. (2005). Impulsive responding in alcoholics. *Alcoholism: Clinical and Experimental Research*, *29*(12), 2158-2169.
- Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology*, *146*(4), 455-464.
- Mitchell, S. H. (2004). Effects of short-term nicotine deprivation on decision-making: delay, uncertainty and effort discounting. *Nicotine & Tobacco Research*, *6*(5), 819-828.
- Moallem, N. R., & Ray, L. A. (2012). Dimensions of impulsivity among heavy drinkers, smokers, and heavy drinking smokers: singular and combined effects. *Addictive Behavior*, *37*(7), 871-874.
- Moeller, F. G., & Dougherty, D. M. (2002). Impulsivity and substance abuse: What is the connection? *Addictive Disorders & Their Treatment*, *1*(1), 3-10.

- Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., . . . Caspi, A. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences USA*, 108(7), 2693-2698.
- Montague, P. R., Dolan, R. J., Friston, K. J., & Dayan, P. (2012). Computational Psychiatry. *Trends in Cognitive Science*, 16(1), 72-80.
- Monterosso, J. R., Ainslie, G., Xu, J., Cordova, X., Domier, C. P., & London, E. D. (2007). Frontoparietal cortical activity of methamphetamine-dependent and comparison subjects performing a delay discounting task. *Human Brain Mapping*, 28(5), 383-393.
- Moore, S. C., & Cusens, B. (2010). Delay discounting predicts increase in blood alcohol level in social drinkers. *Psychiatry Research*, 179(3), 324-327.
- Mueller, E. T., Landes, R. D., Kowal, B. P., Yi, R., Stitzer, M. L., Burnett, C. A., & Bickel, W. K. (2009). Delay of smoking gratification as a laboratory model of relapse: effects of incentives for not smoking, and relationship with measures of executive function. *Behavioral Pharmacology*, 20(5-6), 461-473.
- Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, 64(3), 263-276.
- Myerson, J., Green, L., & Warusawitharana, M. (2001). Area under the curve as a measure of discounting. *Journal of the Experimental Analysis of Behavior*, 76(2), 235-243.
- Neal, D. T. (2006). Habits—A repeat performance. *Current Directions in Psychological Science*, 15(4), 198-202.
- Nederkoorn, C., Smulders, F. T., Havermans, R. C., Roefs, A., & Jansen, A. (2006). Impulsivity in obese women. *Appetite*, 47(2), 253-256.
- Neumann Von J, M. O. (1944). Theory of games and economic behavior. Princeton, NJ, USA: Princeton University Press.
- Norton, M. I., & Ariely, D. (2011). Building a better America—One wealth quintile at a time. *Perspectives on Psychological Science*, 6(1), 9-12.
- Nosek, B. A., Hawkins, C. B., & Frazier, R. S. (2012). Implicit social cognition. In S. Fiske, C. N. McCrae (Eds.) *The SAGE Handbook of Social Cognition* (pp. 31-53), London, UK: Sage.
- O'Brien, C. P., Childress, A. R., Ehrman, R., & Robbins, S. J. (1998). Conditioning factors in drug abuse: can they explain compulsion? *Journal of Psychopharmacology*, 12(1), 15-22.
- O'Donoghue, T., & Rabin, M. (2003). Self awareness and self control. In G. Loewenstein, D. Read, & R. F. Baumeister (Eds.), *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*. (pp. 217-243). New York: Russell Sage Foundation.
- Odum, A. L. (2011). Delay discounting: trait variable? *Behavioral Processes*, 87(1), 1-9.

- Odum, A. L., Madden, G. J., Badger, G. J., & Bickel, W. K. (2000). Needle sharing in opioid-dependent outpatients: psychological processes underlying risk. *Drug and Alcohol Dependence*, 60(3), 259-266.
- Odum, A. L., Madden, G. J., & Bickel, W. K. (2002). Discounting of delayed health gains and losses by current, never- and ex-smokers of cigarettes. *Nicotine and Tobacco Research*, 4(3), 295-303.
- Ogata, S. N., Silk, K. R., Goodrich, S., Lohr, N. E., Westen, D., & Hill, E. M. (1990). Childhood sexual and physical abuse in adult patients with borderline personality disorder. *American Journal of Psychiatry*, 147(8), 1008-1013.
- Ogden, J. (2012). *Health psychology: A Textbook*. (5th Edition) Maidenhead, UK, Open University Press.
- Ohmura, Y., Takahashi, T., & Kitamura, N. (2005). Discounting delayed and probabilistic monetary gains and losses by smokers of cigarettes. *Psychopharmacology*, 182(4), 508-515.
- Olsen, J. A. (1993). Time preferences for health gains: an empirical investigation. *Health Economics*, 2(3), 257-265.
- Orbell, S., & Verplanken, B. (2010). The automatic component of habit in health behavior: habit as cue-contingent automaticity. *Health Psychology*, 29(4), 374-383.
- Orbell, S., & Verplanken, B. (2014). The strength of habit. *Health Psychology Review* (Online publication ahead of print): doi:10.1080/17437199.2014.992031.
- Ostaszewski, P., Green, L., & Myerson, J. (1998). Effects of inflation on the subjective value of delayed and probabilistic rewards. *Psychonomic Bulletin & Review*, 5(2), 324-333.
- Ouellette, J. A. (1998). Habit and intention in everyday life: the multiple processes by which past behavior predicts future behavior. *Psychological Bulletin*, 124(1), 54-74.
- Parfit, D. (1971). Personal identity. *The Philosophical Review*, 3-27.
- Pavlov, I. P. (2003). *Conditioned reflexes*. Mineola, NY, USA: Courier Dover Publications.
- Peters, E. N., Petry, N. M., LaPaglia, D. M., Reynolds, B., & Carroll, K. M. (2013). Delay discounting in adults receiving treatment for marijuana dependence. *Experimental and Clinical Psychopharmacology*, 21(1), 46-54.
- Peters, J., & Büchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron*, 66(1), 138-148.
- Petry, N. M. (2001). Delay discounting of money and alcohol in actively using alcoholics, currently abstinent alcoholics, and controls. *Psychopharmacology*, 154(3), 243-250.
- Petry, N. M. (2003). Discounting of money, health, and freedom in substance abusers and controls. *Drug and Alcohol Dependence*, 71(2), 133-141.

- Petry, N. M., & Casarella, T. (1999). Excessive discounting of delayed rewards in substance abusers with gambling problems. *Drug and Alcohol Dependence*, 56(1), 25-32.
- Phelps, E. S., & Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*, 35(2), 185-199.
- Ploghaus, A., Tracey, I., Gati, J. S., Clare, S., Menon, R. S., Matthews, P. M., & Rawlins, J. N. P. (1999). Dissociating pain from its anticipation in the human brain. *Science*, 284(5422), 1979-1981.
- Prelec, D., & Bodner, R. (2003). Self-signaling and self-control. In G. Loewenstein, D. Read, & R. F. Baumeister (Eds.), *Time and Decision* (pp. 277-298). NY: Russell Sage Foundation.
- Price, D. D., McGrath, P. A., Rafii, A., & Buckingham, B. (1983). The validation of visual analogue scales as ratio scale measures for chronic and experimental pain. *Pain*, 17(1), 45-56.
- Pronin, E., Olivola, C. Y., & Kennedy, K. A. (2008). Doing unto future selves as you would do unto others: Psychological distance and decision making. *Personality and Social Psychology Bulletin*, 34(2), 224-236.
- Rachlin, H., & Jones, B. A. (2008). Social discounting and delay discounting. *Journal of Behavioral Decision Making*, 21(1), 29-43.
- Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology*, 25, 111-164.
- Raineri, A., & Rachlin, H. (1993). The effect of temporal constraints on the value of money and other commodities. *Journal of Behavioral Decision Making*, 6(2), 77-94.
- Rand, D. G. (2012). The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, 299, 172-179.
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7), 545-556.
- Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty*, 23(1), 5-32.
- Read, D., Frederick, S., & Airoldi, M. (2012). Four days later in Cincinnati: Longitudinal tests of hyperbolic discounting. *Acta Psychologica*, 140(2), 177-185.
- Read, D., Loewenstein, G., & Kalyanaraman, S. (1999). Mixing virtue and vice: Combining the immediacy effect and the diversification heuristic. *Journal of Behavioral Decision Making*, 12(4), 257-273.
- Read, D., & Roelofsma, P. H. (2003). Subadditive versus hyperbolic discounting: A comparison of choice and matching. *Organizational Behavior and Human Decision Processes*, 91(2), 140-153.
- Read, D., & Van Leeuwen, B. (1998). Predicting hunger: The effects of appetite and delay on choice. *Organizational Behavior and Human Decision Processes*, 76(2), 189-205.

- Redelmeier, D. A., & Heller, D. N. (1993). Time preference in medical decision making and cost-effectiveness analysis. *Medical Decision Making*, 13(3), 212-217.
- Reimers, S., Maylor, E. A., Stewart, N., & Chater, N. (2009). Associations between a one-shot delay discounting measure and age, income, education and real-world impulsive behavior. *Personality and Individual Differences*, 47(8), 973-978.
- Reynolds, B. (2004). Do high rates of cigarette consumption increase delay discounting?: A cross-sectional comparison of adolescent smokers and young-adult smokers and nonsmokers. *Behavioural Processes*, 67(3), 545-549.
- Reynolds, B. (2006a). The Experiential Discounting Task is sensitive to cigarette-smoking status and correlates with a measure of delay discounting. *Behavioral Pharmacology*, 17(2), 133-142.
- Reynolds, B. (2006b). A review of delay-discounting research with humans: relations to drug use and gambling. *Behavioral Pharmacology*, 17(8), 651-667.
- Reynolds, B., & Fields, S. (2012). Delay discounting by adolescents experimenting with cigarette smoking. *Addiction*, 107(2), 417-424.
- Reynolds, B., Karraker, K., Horn, K., & Richards, J. B. (2003). Delay and probability discounting as related to different stages of adolescent smoking and non-smoking. *Behavioural Processes*, 64(3), 333-344.
- Reynolds, B., Leraas, K., Collins, C., & Melanko, S. (2009). Delay discounting by the children of smokers and nonsmokers. *Drug and Alcohol Dependence*, 99(1), 350-353.
- Reynolds, B., Patak, M., & Shroff, P. (2007). Adolescent smokers rate delayed rewards as less certain than adolescent nonsmokers. *Drug and Alcohol Dependence*, 90(2-3), 301-303.
- Reynolds, B., Penfold, R. B., & Patak, M. (2008). Dimensions of impulsive behavior in adolescents: laboratory behavioral assessments. *Experimental and Clinical Psychopharmacology*, 16(2), 124-131.
- Rezvanfard, M., Ekhtiari, H., Mokri, A., Djavid, G. E. , & Kaviani, H. (2010). Psychological and behavioral traits in smokers and their relationship with nicotine dependence level. *Archives of Iranian Medicine*, 13(5), 395-405.
- Roberts, B. W. (2009). Back to the future: personality and assessment and personality development. *Journal of Research in Personality*, 43(2), 137-145.
- Rogers, R. D., & Robbins, T. W. (2001). Investigating the neurocognitive deficits associated with chronic drug misuse. *Current Opinion in Neurobiology*, 11(2), 250-257.
- Rogers, R. W. (1975). A protection motivation theory of fear appeals and attitude change1. *Journal of Psychology*, 91(1), 93-114.
- Rollins, B. Y., Dearing, K. K., & Epstein, L. H. (2010). Delay discounting moderates the effect of food reinforcement on energy intake among non-obese women. *Appetite*, 55(3), 420-425.
- Rosenstock, I. M. (1974). Historical origins of the health belief model. *Health Education & Behavior*, 2(4), 328-335.

- Ross, C. E., & Wu, C.-I. (1995). The links between education and health. *American Sociological Review*, 60(5), 719-745.
- Ross, L. E. (1961). Conditioned fear as a function of CS-UCS and probe stimulus intervals. *Journal of Experimental Psychology*, 61(4), 265.
- Rossow, I. (2008). Alcohol consumption and discounting. *Addiction Research and Theory*, 16(6), 572-584.
- Rushton, J. P., Chrisjohn, R. D., & Fekken, G. C. (1981). The altruistic personality and the self-report altruism scale. *Personality and Individual Differences*, 2(4), 293-302.
- Rust, J. P. (1989). A dynamic programming model of retirement behavior. In D. A. Wise In D. A. Wise (Ed.) *Perspectives on the Economics of Aging* (pp. 359-404), Chicago, IL, USA: University of Chicago Press.
- Samuelson, & P. (1937). A note on the measurement of utility. *Review of Economic Studies*, 4, 155-161.
- Sanfey, A. G., Loewenstein, G., McClure, S. M., & Cohen, J. D. (2006). Neuroeconomics: cross-currents in research on decision-making. *Trends in Cognitive Science*, 10(3), 108-116.
- Sayman, S., & Öncüler, A. (2009). An investigation of time inconsistency. *Management Science*, 55(3), 470-482.
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2008). Episodic simulation of future events: concepts, data, and applications. *Annals of the New York Academy of Sciences*, 1124, 39-60.
- Schacter, D. L., & Schacter. (2008). Episodic simulation of future events. *Annals of the New York Academy of Sciences*, 1124(1), 39-60.
- Schmitt, M., Gollwitzer, M., Maes, J., & Arbach, D. (2005). Justice sensitivity: assessment and location in the personality space. *European Journal of Psychological Assessment*, 21(3), 202-211.
- Schneider, W., & Chein, J. M. (2003). Controlled & automatic processing: behavior, theory, and biological mechanisms. *Cognitive Science*, 27(3), 525-559.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84(1), 1-66.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461-464.
- Seeyave, D. M., Coleman, S., Appugliese, D., Corwyn, R. F., Bradley, R. H., Davidson, N. S., . . . Lumeng, J. C. (2009). Ability to delay gratification at age 4 years and risk of overweight at age 11 years. *Archives of Pediatric and Adolescent Medicine*, 163(4), 303-308.
- Sen, A. K. (1971). Choice functions and revealed preference. *Review of Economic Studies*, 38(3), 307-317.
- Severens, J. L., & Milne, R. J. (2004). Discounting health outcomes in economic evaluation: the ongoing debate. *Value in Health*, 7(4), 397-401.
- Seymour, B., O'Doherty, J. P., Koltzenburg, M., Wiech, K., Frackowiak, R., Friston, K., & Dolan, R. (2005). Opponent appetitive-aversive neural processes

- underlie predictive learning of pain relief. *Nature Neuroscience*, 8(9), 1234-1240.
- Sheeran, P. (2002). Intention—behavior relations: A conceptual and empirical review. *European Review of Social Psychology*, 12(1), 1-36.
- Sheeran, P., Milne, S., Webb, T. L., & Gollwitzer, P. M. (2009). Implementation Intentions and Health Behaviour. In M. Conner & P. Norman (Eds.), *Predicting Health Behaviour* (2nd Edition ed., pp. 277-323). New York, NY: McGraw-Hill.
- Sheeran, P., & Orbell, S. (1996). How confidently can we infer health beliefs from questionnaire responses? *Psychology and Health*, 11(2), 273-290.
- Sheffer, C., MacKillop, J., McGeary, J., Landes, R., Carter, L., Yi, R., . . . Bickel, W. (2012). Delay discounting, locus of control, and cognitive impulsiveness independently predict tobacco dependence treatment outcomes in a highly dependent, lower socioeconomic group of smokers. *The American Journal on Addictions*, 21(3), 221-232.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84(2), 127-190.
- Simon, D. A., & Daw, N. D. (2011). Neural correlates of forward planning in a spatial decision task in humans. *Journal of Neuroscience*, 31(14), 5526-5539.
- Singer, T., Seymour, B., O'Doherty, J., Kaube, H., Dolan, R. J., & Frith, C. D. (2004). Empathy for pain involves the affective but not sensory components of pain. *Science*, 303(5661), 1157-1162.
- Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4(2), 108-131.
- Smith, R., Corrigan, P., & Exeter, C. (2012). *Countering non-communicable disease through innovation: Report of the non-communicable disease working group*. Working Paper presented at the Global Health Policy Summit, London, UK, 2012.
- Smittenaar, P., FitzGerald, T. H., Romei, V., Wright, N. D., & Dolan, R. J. (2013). Disruption of dorsolateral prefrontal cortex decreases model-based in favor of model-free control in humans. *Neuron*, 80(4), 914-919.
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27(1), 76-105.
- Stanger, C., Ryan, S. R., Fu, H., Landes, R. D., Jones, B. A., Bickel, W. K., & Budney, A. J. (2012). Delay discounting predicts adolescent substance abuse treatment outcome. *Experimental and Clinical Psychopharmacology*, 20(3), 205-212.
- Stea, J. N., Hodgins, D. C., & Lambert, M. J. (2011). Relations between delay discounting and low to moderate gambling, cannabis, and alcohol problems among university students. *Behavioural Processes*, 88(3), 202-205.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, 46(4), 1004-1017.

- Stephens, D. W., & Anderson, D. (2001). The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behavioral Ecology*, 12(3), 330-339.
- Stephoe, A., Wardle, J., Pollard, T. M., Canaan, L., & Davies, G. J. (1996). Stress, social support and health-related behavior: a study of smoking, alcohol consumption and physical exercise. *Journal of Psychosomatic Research*, 41(2), 171-180.
- Stillwell, D. J., & Tunney, R. J. (2012). Effects of measurement methods on the relationship between smoking and delay reward discounting. *Addiction*, 107(5), 1003-1012.
- Strotz, R., & H. (1957). Myopia and inconsistency in dynamic utility maximisation. *Review of Economic Studies*, 23, 165-180.
- Stuckless, N., & Goranson, R. (1992). The Vengeance Scale: Development of a measure of attitudes toward revenge. *Journal of Social Behavior & Personality*, 7(1), 25-42.
- Sullivan, M. J., Adams, H., Horan, S., Maher, D., Boland, D., & Gross, R. (2008). The role of perceived injustice in the experience of chronic pain and disability: scale development and validation. *Journal of Occupational Rehabilitation*, 18(3), 249-261.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to Reinforcement Learning*. London, UK: MIT Press.
- Sweitzer, M. M., Donny, E. C., Dierker, L. C., Flory, J. D., & Manuck, S. B. (2008). Delay discounting and smoking: association with the Fagerstrom Test for Nicotine Dependence but not cigarettes smoked per day. *Nicotine and Tobacco Research*, 10(10), 1571-1575.
- Swirsky-Sacchetti, T., Gorton, G., Samuel, S., Sobel, R., Genetta-Wadley, A., & Burleigh, B. (1993). Neuropsychological function in borderline personality disorder. *Journal of Clinical Psychology*, 49(3), 385-396.
- Takahashi, H., Kato, M., Matsuura, M., Mobbs, D., Suhara, T., & Okubo, Y. (2009). When your gain is my pain and your pain is my gain: neural correlates of envy and schadenfreude. *Science*, 323(5916), 937-939.
- Takahashi, T., Ohmura, Y., Oono, H., & Radford, M. (2009). Alcohol use and discounting of delayed and probabilistic gain and loss. *Neurology and Endocrinology Letters*, 30(6), 749-752.
- Takahashi, T., Oono, H., & Radford, M. H. (2008). Psychophysics of time perception and intertemporal choice models. *Physica A: Statistical Mechanics and Its Applications*, 387(8), 2066-2074.
- Tanaka, S. C., Yamada, K., Yoneda, H., & Ohtake, F. (2014). Neural Mechanisms of Gain-Loss Asymmetry in Temporal Discounting. *Journal of Neuroscience*, 34(16), 5595-5602.
- Thaler, R. H. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3), 201-207.
- Thaler, R. H., & Benartzi, S. (2004). Save More Tomorrow™: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112(S1), S164-S187.

- Thaler, R. H., & Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy*, 392-406.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, Connecticut: Yale University Press.
- Tiffany, S. T. (1990). A cognitive model of drug urges and drug-use behavior: role of automatic and nonautomatic processes. *Psychological Review*, 97(2), 147-168.
- Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811), 515-518.
- Tricomi, E., Balleine, B. W., & O'Doherty, J. P. (2009). A specific role for posterior dorsolateral striatum in human habit learning. *European Journal of Neuroscience*, 29(11), 2225-2232.
- Tricomi, E., Rangel, A., Camerer, C. F., & O'Doherty, J. P. (2010). Neural evidence for inequality-averse social preferences. *Nature*, 463(7284), 1089-1091.
- Tuk, M. A., Zhang, K., & Sweldens, S. (2015). The propagation of self-control: Self-control in one domain simultaneously improves self-control in other domains. *Journal of Experimental Psychology: General*, 144(3), 639-654.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453-458.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, 111(3), 757-769.
- Van Der Pol, M. (2011). Health, education and time preference. *Health Economics*, 20(8), 917-929.
- Van Der Pol, M., & Cairns, J. (2002). A comparison of the discounted utility model and hyperbolic discounting models in the case of social and private intertemporal preferences for health. *Journal of Economic Behavior and Organization*, 49(1), 79-96.
- Van Der Pol, M., & Cairns, J. (2008). Comparison of two methods of eliciting time preference for future health states. *Social Science and Medicine*, 67(5), 883-889.
- Van Der Pol, M., & Cairns, J. (2011). Descriptive validity of alternative intertemporal models for health outcomes: an axiomatic test. *Health Economics*, 20(7), 770-782.
- Van Der Pol, M. M., & Cairns, J. A. (2000). Negative and zero time preference for health. *Health Economics*, 9(2), 171-175.
- Van Oers, J., Bongers, I., Van de Goor, L., & Garretsen, H. (1999). Alcohol consumption, alcohol-related problems, problem drinking, and socioeconomic status. *Alcohol and Alcoholism*, 34(1), 78-88.
- Vlaev, I., & Darzi, A. (2012). Preferences and their implication for policy, health and wellbeing. In R. Dolan, & T. Sharot (Eds.), *Neuroscience of Preference and Choice* (pp. 305-330). London, UK: Elsevier.
- Vlaev, I., Seymour, B., Chater, N., Winston, J. S., Yoshida, W., Wright, N., . . . Dolan, R. (2014). Prices need no preferences: social trends determine decisions in experimental markets for pain relief. *Health Psychology*, 33(1), 66-76.

- Vlaev, I., Seymour, B., Dolan, R. J., & Chater, N. (2009). The price of pain and the value of suffering. *Psychological Science*, 20(3), 309-317.
- Vuchinich, R. E., & Simpson, C. A. (1998). Hyperbolic temporal discounting in social drinkers and problem drinkers. *Experimental and Clinical Psychopharmacology*, 6(3), 292-305.
- Ward, S. E., Goldberg, N., Miller-McCauley, V., Mueller, C., Nolan, A., Pawlik-Plank, D., . . . Weissman, D. E. (1993). Patient-related barriers to management of cancer pain. *Pain*, 52(3), 319-324.
- Waschulewski-Floruss, H., Miltner, W., Brody, S., & Braun, C. (1994). Classical conditioning of pain responses. *International Journal of Neuroscience*, 78(1-2), 21-32.
- Washio, Y., Higgins, S. T., Heil, S. H., McKerchar, T. L., Badger, G. J., Skelly, J. M., & Dantona, R. L. (2011). Delay discounting is associated with treatment response among cocaine-dependent outpatients. *Experimental and Clinical Psychopharmacology*, 19(3), 243-248.
- Webb, T. L., & Sheeran, P. (2006). Does changing behavioral intentions engender behavior change? A meta-analysis of the experimental evidence. *Psychological Bulletin*, 132(2), 249-268.
- Weller, R. E., Cook I., E. W., Avsar, K. B., & Cox, J. E. (2008). Obese women show greater delay discounting than healthy-weight women. *Appetite*, 51(3), 563-569.
- White, C., Edgar, G., & Siegler, V. (2007). Social inequalities in male mortality for selected causes of death by the National Statistics Socio-economic Classification, England and Wales, 2001-03. *Health Statistics Quarterly: Office for National Statistics* 38, 19-32.
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Computation*, 22(9), 2452-2476.
- Wing, V. C., Moss, T. G., Rabin, R. A., & George, T. P. (2012). Effects of cigarette smoking status on delay discounting in schizophrenia and healthy controls. *Addictive Behavior*, 37(1), 67-72.
- Winston, J. S., Vlaev, I., Seymour, B., Chater, N., & Dolan, R. J. (2014). Relative valuation of pain in human orbitofrontal cortex. *Journal of Neuroscience*, 34(44), 14526-14535.
- Wunderlich, K., Smittenaar, P., & Dolan, R. J. (2012). Dopamine enhances model-based over model-free choice behavior. *Neuron*, 75(3), 418-424.
- Yankelevitz, R. L., Mitchell, S. H., & Zhang, Y. (2012). Gender differences in factors associated with alcohol drinking: delay discounting and perception of others' drinking. *Drug and Alcohol Dependence*, 123(1-3), 273-276.
- Yates, P. M., Edwards, H. E., Nash, R. E., Walsh, A. M., Fentiman, B. J., Skerman, H. M., & Najman, J. M. (2002). Barriers to effective cancer pain management: a survey of hospitalized cancer patients in Australia. *Journal of Pain and Symptom Management*, 23(5), 393-405.
- Yi, R., & Landes, R. D. (2012). Temporal and probability discounting by cigarette smokers following acute smoking abstinence. *Nicotine and Tobacco Research*, 14(5), 547-558.

- Yoda, T., & Kumakura, T. (2007). Effect of unfairness on customer satisfaction: New insights into customer retention. *Innovative Marketing*, 3(1), 44-55.
- Yoon, J. H., Higgins, S. T., Heil, S. H., Sugarbaker, R. J., Thomas, C. S., & Badger, G. J. (2007). Delay discounting predicts postpartum relapse to cigarette smoking among pregnant women. *Experimental and Clinical Psychopharmacology*, 15(2), 176-186.
- Zanarini, M. C., Williams, A. A., Lewis, R. E., & Reich, R. B. (1997). Reported pathological childhood experiences associated with the development of borderline personality disorder. *American Journal of Psychiatry*, 154(8), 1101-1106.

APPENDIX: BINARY CHOICE OPTIONS

Table A1. Binary choice options for Experiment 1A

Delay Option 1 (Trials)	Delay Option 2 (Trials)	Expected Number of Shocks Option 1	Expected Number of Shocks Option 2
0	4	4	3
0	7	4	3
0	10	4	3
0	13	4	3
0	16	4	3
0	7	5	3
0	13	5	3
0	19	5	3
0	25	5	3
0	31	5	3
0	10	6	3
0	19	6	3
0	28	6	3
0	37	6	3
0	46	6	3
2	6	7	6
2	9	7	6
2	12	7	6
2	15	7	6
2	18	7	6
2	9	8	6
2	15	8	6
2	21	8	6
2	27	8	6
2	33	8	6
2	12	9	6
2	21	9	6
2	30	9	6
2	39	9	6
2	48	9	6
5	9	10	9
5	12	10	9
5	15	10	9
5	18	10	9

5	21	10	9
5	12	11	9
5	18	11	9
5	24	11	9
5	30	11	9
5	36	11	9
5	15	12	9
5	24	12	9
5	33	12	9
5	42	12	9
5	51	12	9
0	4	3	4
0	7	3	4
0	10	3	4
0	13	3	4
0	16	3	4
0	7	3	5
0	13	3	5
0	19	3	5
0	25	3	5
0	31	3	5
0	10	3	6
0	19	3	6
0	28	3	6
0	37	3	6
0	46	3	6
2	6	6	7
2	9	6	7
2	12	6	7
2	15	6	7
2	18	6	7
2	9	6	8
2	15	6	8
2	21	6	8
2	27	6	8
2	33	6	8
2	12	6	9
2	21	6	9
2	30	6	9
2	39	6	9
2	48	6	9
5	9	9	10
5	12	9	10
5	15	9	10
5	18	9	10

5	21	9	10
5	12	9	11
5	18	9	11
5	24	9	11
5	30	9	11
5	36	9	11
5	15	9	12
5	24	9	12
5	33	9	12
5	42	9	12
5	51	9	12
0	4	4	4
0	7	5	5
0	10	6	6
0	13	7	7
0	16	8	8

Table A2. Binary choice options for Experiment 1B

Delay to Earlier Dental Appointment	Delay to Later Dental Appointment	Severity of Dental Pain for Earlier Appointment (%)	Severity of Dental Pain for Later Appointment (%)
0	0	60	41
0	0	60	19
0	1	60	60
0	1	60	55
0	1	60	51
0	1	60	46
0	1	60	37
0	1	60	16
0	5	60	60
0	5	60	55
0	5	60	51
0	5	60	46
0	5	60	37
0	5	60	16
0	13	60	60
0	13	60	55
0	13	60	51
0	13	60	46
0	13	60	37
0	13	60	16
0	32	60	60
0	32	60	55
0	32	60	51
0	32	60	46
0	32	60	37
0	32	60	16
0	89	60	60
0	89	60	55
0	89	60	51
0	89	60	46
0	89	60	37
0	89	60	16
0	237	60	60
0	237	60	55
0	237	60	51
0	237	60	46

0	237	60	37
0	237	60	16
0	0	41	60
0	0	19	60
0	1	55	60
0	1	51	60
0	1	46	60
0	1	37	60
0	1	16	60
0	5	55	60
0	5	51	60
0	5	46	60
0	5	37	60
0	5	16	60
0	13	55	60
0	13	51	60
0	13	46	60
0	13	37	60
0	13	16	60
0	32	55	60
0	32	51	60
0	32	46	60
0	32	37	60
0	32	16	60
0	89	55	60
0	89	51	60
0	89	46	60
0	89	37	60
0	89	16	60
0	237	55	60
0	237	51	60
0	237	46	60
0	237	37	60
0	237	16	60

Table A3. Binary choice options for *self-now-self-later* and *other-now-other-later* conditions in Experiment 3A

Shock intensity (VAS/10) Option 1	Shock intensity (VAS/10) Option 2	Delay (s) Option 1	Delay (s) Option 2
1	5	0	2
3	5	0	2
5	5	0	2
7	5	0	2
9	5	0	2
5	1	0	2
5	3	0	2
5	5	0	2
5	7	0	2
5	9	0	2
1	5	0	2
3	5	0	2
7	5	0	2
9	5	0	2
5	1	0	2
5	3	0	2
5	7	0	2
5	9	0	2
1	5	0	4
3	5	0	4
5	5	0	4
7	5	0	4
9	5	0	4
5	1	0	4
5	3	0	4
5	5	0	4
5	7	0	4
5	9	0	4
1	5	0	4
3	5	0	4
7	5	0	4
9	5	0	4
5	1	0	4
5	3	0	4
5	7	0	4
5	9	0	4

1	5	0	8
3	5	0	8
5	5	0	8
7	5	0	8
9	5	0	8
5	1	0	8
5	3	0	8
5	5	0	8
5	7	0	8
5	9	0	8
1	5	0	8
3	5	0	8
7	5	0	8
9	5	0	8
5	1	0	8
5	3	0	8
5	7	0	8
5	9	0	8
1	5	0	16
3	5	0	16
5	5	0	16
7	5	0	16
9	5	0	16
5	1	0	16
5	3	0	16
5	5	0	16
5	7	0	16
5	9	0	16
1	5	0	16
3	5	0	16
7	5	0	16
9	5	0	16
5	1	0	16
5	3	0	16
5	7	0	16
5	9	0	16
1	5	0	29
3	5	0	29
5	5	0	29
7	5	0	29
9	5	0	29
5	1	0	29
5	3	0	29
5	5	0	29
5	7	0	29

5	9	0	29
1	5	0	29
3	5	0	29
7	5	0	29
9	5	0	29
5	1	0	29
5	3	0	29
5	7	0	29
5	9	0	29

Table A4 Binary choice options for *self-now-other-now* choices in Experiment 3A. Self-now-other-later and other-now-self-later conditions both used the 90 choices shown in Table A3, with the addition of the following 18 choices in which both shock options occurred immediately (delay =0 seconds).

Option 1 Pain Intensity (VAS/10)	Option 2 Pain Intensity (VAS/10)
1	5
3	5
5	5
7	5
9	5
5	1
5	3
5	5
5	7
5	9
1	5
3	5
7	5
9	5
5	1
5	3
5	7
5	9

Table A5 Binary choice options for *self-now-other-now* (social discounting) choices in Experiment 3B

Pain Intensity Self (VAS/10)	Pain Intensity Other (VAS/10)	Social Distance of Other (/100)
1	5	1
3	5	1
5	5	1
7	5	1
9	5	1
5	1	1
5	3	1
5	5	1
5	7	1
5	9	1
1	5	1
3	5	1
7	5	1
9	5	1
5	1	1
5	3	1
5	7	1
5	9	1
1	5	5
3	5	5
5	5	5
7	5	5
9	5	5
5	1	5
5	3	5
5	5	5
5	7	5
5	9	5
1	5	5
3	5	5
7	5	5
9	5	5
5	1	5
5	3	5
5	7	5
5	9	5
1	5	14
3	5	14

5	5	14
7	5	14
9	5	14
5	1	14
5	3	14
5	5	14
5	7	14
5	9	14
1	5	14
3	5	14
7	5	14
9	5	14
5	1	14
5	3	14
5	7	14
5	9	14
1	5	45
3	5	45
5	5	45
7	5	45
9	5	45
5	1	45
5	3	45
5	5	45
5	7	45
5	9	45
1	5	45
3	5	45
7	5	45
9	5	45
5	1	45
5	3	45
5	7	45
5	9	45
1	5	97
3	5	97
5	5	97
7	5	97
9	5	97
5	1	97
5	3	97
5	5	97
5	7	97
5	9	97
1	5	97

3	5	97
7	5	97
9	5	97
5	1	97
5	3	97
5	7	97
5	9	97