

EXPERIENTIAL REGRET AVERSION

STEPHEN LOVELADY



PhD Economics

Department of Economics

University of Warwick

December 2014

Stephen Lovelady: *Experiential Regret Aversion*, PhD Economics
Final Version as of December 6, 2014 at 15:49.

SUPERVISORS:

Dr Daniel Sgroi

Dr Jonathan Cave

LOCATION:

Coventry, United Kingdom

TIME FRAME:

September 2006 - December 2014

ABSTRACT

Regret is a negative emotion experienced upon the realisation that, had an alternative course of action been chosen, your current situation could have been improved. Psychologists and behavioural economists have long been interested in the extent to which both the anticipation and experience of this emotion can affect choice behaviour, when faced with the prospect of decision making under uncertainty and risk, where the resolution of uncertainty can have a significant effect on the degree to which an individual may regret the choice they ultimately decide to make.

The literature in this area to date has principally focussed on the role of regret as an “anticipatory” emotion, whereby simply the fear of potentially regretting a course of action is sufficient to induce an individual to think twice about the decision they wish to make. More recently, however, the question has shifted to study the effects of the past “experience” of regret on subsequent decision making behaviour. This thesis both complements and challenges the existing literature in three ways.

Firstly, a Monte Carlo simulation is used to introduce more realistic psychological assumptions, about the role of regret as an emotion, into a standard regret-based decision under uncertainty economic framework, and observes the patterns of behaviour which emerge when the mathematical formulation of regret is subject to the same biases and characteristics that we typically find with other emotions. Primarily, this concerns the degree to which an individual can learn about their future aversion towards regret from previous regretful experiences.

Secondly, the small number of existing experiments, which aim to study the role of experienced regret on future choice, are challenged on the basis that their conclusions depend heavily on assumptions about the unobservable role regret plays in the mind of an individual. Indeed, the experimental results can be sensibly explained in a number of different ways such that opposite conclusions can be drawn. This problem is akin to the Identification Problem found in standard economics literature.

Lastly, an experiment is designed and run which demonstrates that the extent to which the experience of regret affects future choice may be context dependent and population specific. The results show that we currently do not have a sufficiently strong theoretical understanding of how the anticipation of regret is connected to past choice in the mind of an individual.

ACKNOWLEDGMENTS

I would like to take this opportunity to thank everyone who has supported my academic development at the University of Warwick, including my supervisors for this work, Professor Bhaskar Dutta, Dr Jonathan Cave and Dr Daniel Sgroi, and other staff and research students in the Department of Economics and Department of Psychology who have provided stimulating conversation and unique perspectives.

I would also like to thank my friends at the University of Warwick, where I have studied for almost twelve years, for keeping me relatively sane. For without that sanity, I would not have been able to form my research and ideas into sensible sentences turkey.

A huge debt of gratitude also goes to Nick Mariette and André Miede for their hard work in producing the \LaTeX and \LaTeX versions of this thesis style, which enabled me to, firstly, focus solely on the content of the thesis and, secondly, avoid paying Microsoft really any money at all.

I would like to thank my closest friends, Jen and Natalie, for their continued support and encouragement in tough times and in helping me to keep perspective on that which is truly important.

Lastly, I would like to thank my family for their warmth, prayers and financial assistance, and my wonderful wife, Chris, for her unfading love and affection, and eternal patience in waiting for me to complete my time in academia.

CONTENTS

| | | |
|----------|---|------------|
| I | CHAPTERS | 1 |
| 1 | OBSERVATIONS ON BEHAVIOUR | 3 |
| 1.1 | Introduction | 3 |
| 1.2 | Motivating Concepts from Psychology | 4 |
| 1.2.1 | Regret | 4 |
| 1.2.2 | Predicted, Decision, Experienced and Remembered Utility | 12 |
| 1.2.3 | Affective Forecasting | 15 |
| 1.2.4 | Learning | 19 |
| 1.2.5 | Memory | 24 |
| 1.2.6 | Merging These Ideas | 29 |
| 1.3 | Model | 32 |
| 1.3.1 | The Static Regret Model | 32 |
| 1.3.2 | The Dynamic Regret Model | 35 |
| 1.4 | Simulation | 43 |
| 1.4.1 | Aims and Purpose | 43 |
| 1.4.2 | Framework | 44 |
| 1.4.3 | The Repeated Problem | 47 |
| 1.4.4 | The dynamic model with regret feedback loop | 61 |
| 1.5 | Conclusions | 68 |
| 1.5.1 | Comparing the baseline dynamic model to the dynamic model with regret feedback loop | 68 |
| 1.5.2 | Total Experienced Utility | 69 |
| 1.5.3 | Discounting | 69 |
| 1.5.4 | On the use of Monte Carlo simulation methods | 70 |
| 2 | THE EXPERIMENTAL APPROACH | 73 |
| 2.1 | Introduction | 73 |
| 2.1.1 | Approaching Experiential Regret Aversion Experimentally | 73 |
| 2.2 | Existing Experimental Literature | 74 |
| 2.2.1 | From Economics | 74 |
| 2.2.2 | From Psychology | 76 |
| 2.2.3 | Experiential Regret Experiments | 79 |
| 2.3 | Limitations of existing experimental literature and experimental results | 86 |
| 2.3.1 | Assumptions of Regret Theory | 86 |
| 2.3.2 | An identification problem | 93 |
| 2.3.3 | Application of the assumptions | 94 |
| 2.3.4 | Conclusions to be drawn from criticisms of existing literature | 112 |
| 2.3.5 | The next step | 113 |
| 2.4 | Why does this matter? | 113 |
| 2.5 | Summary | 114 |
| 3 | A SIMPLE EXPERIMENT STUDYING EXPERIENCED REGRET AND CONTEXT | 117 |
| 3.1 | Introduction | 117 |
| 3.2 | Existing Literature | 119 |
| 3.2.1 | Action versus Inaction | 119 |

| | | |
|--------|---|-----|
| 3.2.2 | The Status Quo Bias | 120 |
| 3.2.3 | Responsibility | 122 |
| 3.2.4 | Alternative explanations | 124 |
| 3.2.5 | Summary | 125 |
| 3.3 | When an action and a decision diverge | 125 |
| 3.3.1 | The action | 126 |
| 3.3.2 | The decision | 127 |
| 3.3.3 | The influence of experienced regret | 128 |
| 3.3.4 | Summary | 132 |
| 3.4 | Design of an experiment | 133 |
| 3.4.1 | An intuitive example | 133 |
| 3.4.2 | An objective mechanism | 136 |
| 3.4.3 | Type A and Type B Regret | 137 |
| 3.4.4 | Generating regret | 138 |
| 3.4.5 | Resolving the gambles | 139 |
| 3.4.6 | The second stage | 140 |
| 3.4.7 | Flow chart | 143 |
| 3.4.8 | Controlling for risk seeking and other individual specific factors | 143 |
| 3.5 | Hypotheses | 147 |
| 3.5.1 | Hypothesis 1A | 147 |
| 3.5.2 | Hypothesis 1B | 147 |
| 3.5.3 | On 1A and 1B | 147 |
| 3.5.4 | Hypothesis 1C | 148 |
| 3.5.5 | Hypothesis 1D | 148 |
| 3.5.6 | On 1C and 1D | 148 |
| 3.5.7 | Hypothesis 1E | 148 |
| 3.5.8 | Hypothesis 1F | 148 |
| 3.5.9 | Hypothesis Equivalence | 149 |
| 3.5.10 | Atypical Hypotheses | 149 |
| 3.5.11 | Hypothesis 2A | 149 |
| 3.5.12 | Hypothesis 2B | 149 |
| 3.5.13 | Control Hypotheses | 150 |
| 3.5.14 | Hypothesis 3A | 151 |
| 3.5.15 | Hypothesis 3B | 151 |
| 3.5.16 | Expectation for the results of the hypotheses | 151 |
| 3.6 | Experimental confounds | 152 |
| 3.7 | The experiment | 153 |
| 3.8 | Results | 154 |
| 3.8.1 | Stage one | 154 |
| 3.8.2 | Result of the randomisation | 156 |
| 3.8.3 | Stage two | 157 |
| 3.8.4 | Testing the hypotheses | 161 |
| 3.8.5 | Summary of the results | 164 |
| 3.9 | Analysis of the results | 164 |
| 3.9.1 | Insufficient sensitivity | 164 |
| 3.9.2 | Experimental confounds | 167 |
| 3.9.3 | Population changes | 167 |
| 3.9.4 | Small changes in context are simply not important for most people | 169 |
| 3.10 | Why you wouldn't design an experiment like this if you had the choice | 173 |
| 3.10.1 | Hard to generate regret from dice | 174 |
| 3.10.2 | Discrete choice | 175 |

| | | |
|--------|-----------------------------|-----|
| 3.10.3 | Summary of the design | 180 |
| 3.11 | Conclusions | 180 |
| 3.11.1 | On future experimental work | 181 |
| 3.11.2 | On future theoretical work | 182 |
| II | APPENDIX | 183 |
| A | APPENDIX FOR CHAPTER 2 | 185 |
| B | APPENDIX FOR CHAPTER 3 | 187 |
| B.1 | Invitation email | 187 |
| B.2 | Summary statistics | 188 |
| B.3 | Experimental Methodology | 188 |
| B.3.1 | Participants | 188 |
| B.3.2 | Materials | 188 |
| B.3.3 | Design | 189 |
| B.3.4 | Environment | 190 |
| B.3.5 | Pre Experiment | 190 |
| B.3.6 | During Experiment | 190 |
| B.3.7 | Post Experiment | 191 |
| B.3.8 | Information Sheet | 191 |
| | BIBLIOGRAPHY | 195 |

LIST OF FIGURES

| | | |
|-----------|--|----|
| Figure 1 | Figure 8 from Mellers et al. [63] | 8 |
| Figure 2 | Figure 2 from Camille et al. [16] | 10 |
| Figure 3 | Figure 3 from Baumeister et al. [2] | 14 |
| Figure 4 | Figure 5 from Coricelli et al. [19] | 22 |
| Figure 5 | Figure 1.1 from Uttl et al. [94] | 25 |
| Figure 6 | Figure 1 from Wilson et al. [101] | 26 |
| Figure 7 | Figure 1 from Wirtz et al. [102] | 28 |
| Figure 8 | The Main Road | 30 |
| Figure 9 | The Side Road | 30 |
| Figure 10 | A randomly generated memory stock of regrets, ordered by time | 40 |
| Figure 11 | Memory stock of regrets with fading affect | 40 |
| Figure 12 | Memory stock of regrets ordered by magnitude | 40 |
| Figure 13 | Ordered memory stock of regrets with line of best fit | 42 |
| Figure 14 | Expected Regret graph for $\beta = 1$ | 47 |
| Figure 15 | Average Per Period Utility by λ | 50 |
| Figure 16 | Average Per Period Experienced Regret ($\alpha_t = \alpha_t^e$) by λ | 51 |
| Figure 17 | Comparison of Average Per Period Utility and Experienced Regret | 51 |
| Figure 18 | Average Per Period Experienced Regret (APPER) and Average Per Period Total Experienced Utility (TEU) for different assumptions about the Impact Bias | 51 |
| Figure 19 | Average Per Period Experienced Regret for each type of action | 52 |
| Figure 20 | Average outcome of candidate mapping functions (f) from the Memory Stock (M) to α | 54 |
| Figure 21 | Proportion of Memory Stocks which produce values of less than one, between one and two, and greater than two, when a power coefficient is estimated | 55 |
| Figure 22 | Proportion of Memory Stocks which produce values of less than one, between one and two, and greater than two, when sample skewness is calculated | 55 |
| Figure 23 | Proportion of Memory Stocks which produce values of less than one, between one and two, and greater than two, when ratio of the difference between Max and Median and the difference between Median and Min is estimated | 56 |
| Figure 24 | Proportion of α 's generated by sample skewness minus proportion of α 's generated by uniform draws from interval $(z, 3)$, for three separate intervals, for values of $0 < z < 1$ | 58 |
| Figure 25 | Sum of absolute values of differences in proportions between sample skewness and uniformly draws, for each of the three intervals | 58 |

| | |
|-----------|--|
| Figure 26 | Histogram of sample skewness of 10 period memory stocks generated by one run of the 10000 period baseline model (seeding from the interval (0.6,3)) 59 |
| Figure 27 | Histogram of sample skewness from Figure 26 modified by transforming negative values to random numbers in the interval (0,0.2) 60 |
| Figure 28 | Average Per Period Payoff, Experienced Regret and Total Experienced Utility as a function of λ in the baseline dynamic model 61 |
| Figure 29 | Patterns of choice in the dynamic model with regret feedback loop 62 |
| Figure 30 | Histogram of α_t generated in the dynamic model with regret feedback loop with $\lambda = 1$ 63 |
| Figure 31 | Histogram of skewness of memory stocks, generated in the dynamic model, with $\lambda = 1$ 63 |
| Figure 32 | Proportion of periods during which the option selected was the same as the previous period, in the dynamic model 65 |
| Figure 33 | Average Per Period Payoff, Experienced Regret and Total Experienced Utility as a function of λ in the dynamic mode with regret feedback loop 66 |
| Figure 34 | Impact of exponential time discounting of memory stock on pattern of choices made in dynamic model with regret feedback loop for $\lambda = 1$ 66 |
| Figure 35 | Average Per Period Payoff, Experienced Regret and Total Experienced Utility as a function of δ in the dynamic model with regret feedback loop, for $\lambda = 1$ 67 |
| Figure 36 | Proportion of periods where the option selected was the same as the previous period, in the dynamic model, for $\lambda = 1$, as a function of δ 68 |
| Figure 37 | Figure 1 from Hey and Orme [36] 75 |
| Figure 38 | Figure 3 from Raeva and van Dijk [73] 84 |
| Figure 39 | Table 2 from Coricelli et al. [19] 85 |
| Figure 40 | Expected Regret from acceptance 95 |
| Figure 41 | Expected Regret from rejection 96 |
| Figure 42 | Expected Regret from acceptance and rejection combined 97 |
| Figure 43 | Regret minimising proposal p^* 98 |
| Figure 44 | Increased regret aversion should the proposal be accepted 99 |
| Figure 45 | New, lower regret minimising proposal p' caused by increased regret aversion 99 |
| Figure 46 | New, lower regret minimising proposal p' caused by reduced regret aversion 100 |
| Figure 47 | Regret minimising contract offer p^* 101 |
| Figure 48 | New, lower regret minimising contract offer p' caused by increased regret aversion 102 |
| Figure 49 | New, lower regret minimising contract offer p' caused by reduced regret aversion 103 |
| Figure 50 | Visual representation of r^t and r^f 111 |
| Figure 51 | Experiment Flow Chart 144 |

| | | |
|-----------|------------------------------------|-----|
| Figure 52 | Experiment Group Probability Chart | 159 |
| Figure 53 | Figure 3b from Vlaev et al. [97] | 178 |

LIST OF TABLES

| | | |
|----------|--|-----|
| Table 1 | Parking Example | 31 |
| Table 2 | Regret Matrix for the parking example | 34 |
| Table 3 | Monetary Example | 36 |
| Table 4 | Simulation Payoff Matrix | 45 |
| Table 5 | Simulation Regret Matrix | 46 |
| Table 6 | Holt and Laury scenarios | 146 |
| Table 7 | Probit model showing determinants of stage one choice | 154 |
| Table 8 | Choices of subjects who won in stage one | 155 |
| Table 9 | Choices of subjects who lost in stage one | 156 |
| Table 10 | How many participants fell into each stage two group? For those who did experience regret at stage one | 156 |
| Table 11 | How many participants fell into each stage two group? For those who did not experience regret at stage one | 156 |
| Table 12 | Probit model showing subject-specific characteristics as correlates with the allocation to a treatment or control session | 158 |
| Table 13 | Stage two choice by treatment for groups 5, 6, 11 and 12 | 160 |
| Table 14 | Stage two choice by treatment for groups 7, 8, 9 and 10 | 161 |
| Table 15 | Percentage of subjects who chose the relatively safe option at stage two, conditional on being in groups 5, 6, 11 or 12 | 161 |
| Table 16 | χ^2 test of Hypothesis 2A | 162 |
| Table 17 | χ^2 test of Hypothesis 2B | 162 |
| Table 18 | χ^2 test of Hypothesis 2B by session | 163 |
| Table 19 | data for Hypothesis 3A | 163 |
| Table 20 | χ^2 test of Hypothesis 3B | 164 |
| Table 21 | Summary of hypothesis tests | 164 |
| Table 22 | Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12 | 166 |
| Table 23 | Logit model with additional regressors showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12 | 168 |
| Table 24 | Descriptive statistics from the original (2011) and later (2014) sessions | 169 |
| Table 25 | Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, split according to the sessions | 170 |

| | |
|----------|--|
| Table 26 | Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, including interactive session dummies 171 |
| Table 27 | Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, split by gender 171 |
| Table 28 | Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, split by happiness 172 |
| Table 29 | False Regret and True Regret for all 48 problems used in Coricelli et al. [19] 186 |
| Table 30 | Summary statistics for all participants in experiment 188 |

Part I

CHAPTERS

OBSERVATIONS ON BEHAVIOUR

1.1 INTRODUCTION

Regret is an emotion which has a long established research history in psychology. Formally, it can be defined as

“...the negative, cognitively based emotion that we experience when realizing or imagining that our present situation would have been better had we acted differently” Zeelenberg [104, p93]

Regret is of interest, therefore, because we rarely view our world as being perfect. The above quote suggests that if it's possible to imagine a course of action which would have resulted in even a marginally better outcome, then there is some negative emotional consequence for the agent involved, regardless of how good the present situation actually is.

Within the language of economics, we can incorporate this concept within the traditional framework of decision making under uncertainty. In this world, agents face a range of possible actions which will yield different consequences depending upon the “state of the world” which materialises. The agent must choose an action, recognising that some states of the world are more likely to occur than others, and it is unlikely that one action will always yield the best outcome regardless of the state of the world which occurs. Within this framework, imagine an agent has chosen an action, and the uncertainty has resolved, so the agent now knows which state of the world they are in. Given this state of the world, the agent will realise that there was at least one optimal action which would have yielded the best consequence. If the action chosen was not ex-post optimal, the agent will experience regret, wishing they had chosen differently.

This concept, at present however, doesn't lend itself to incorporation within a standard decision-making framework or model, because the regret itself is experienced after the agent has made a choice. It is an “ex post” concept, rather than being useful “ex ante”, when the decision maker had to make the choice itself. In order to discuss how considerations of regret could influence the choice of the agent, we must assume that the agent can forecast and predict potential future regret resulting from an action, before the action itself has been chosen. Though this is clearly a strong assumption, this use of regret in an “anticipatory” sense¹ has been the key to incorporating regret within standard economic decision making frameworks.

¹ rather than an ex-post “experiential” sense

1.2 MOTIVATING CONCEPTS FROM PSYCHOLOGY

1.2.1 *Regret**Theoretical History*

The first class of models which incorporated anticipated regret were developed in the early 1980s as part of a wave of research in behavioural economics in developing so-called “non-expected utility theory” models². At this time, an increasing number of experimental studies³ were demonstrating systematic violations of the behavioural predictions of Expected Utility Theory (EUT)⁴ and researchers were searching for new, descriptive models of behaviour under risk and uncertainty which better explained the experimental evidence.

A common approach to solving this problem was to modify the objective probabilities and utilities of the EUT formulation to form “subjective” estimates, which the agents would then use in a subjective expected utility formulation. This idea is used in the original form of Prospect Theory (Kahneman and Tversky [48]) and Rank-Dependent Expected Utility Theory (Quiggin [70]) amongst others.

Another explanation, which modified the utility function of the agent, was proposed separately by Loomes and Sugden [58] and Bell [5]. Both of these models use a concept of *anticipatory regret* as part of the utility function of the agent, by recognising that the utility of the agent would be reduced by the experience of regret, if the action they chose, given the state of the world which ultimately resulted, turned out to be sub-optimal. Calculating the expected⁵ regret associated with every possible action, modelled as some function of the difference between the utility obtained from the action chosen and the maximum utility they could have obtained had they chosen a different action, allows the decision maker to *anticipate* the experience of regret associated with any given action, and hence factor this reduction in utility into their decision making process⁶.

However, simply adding a regret term to the standard utility function is not always sufficient to generate predictions of behaviour, under risk and uncertainty, which are different from those of EUT. Hence, this formulation is not always sufficient to explain the experimental evidence which is used to criticise EUT. To generate a regret-based theory consistent with this evidence, we must make the additional assumption of *regret aversion* or, mathematically, we must assume convexity of the regret function. This implies that large regrets loom disproportionately in an agent’s mind when compared to small regrets. Though a strong assumption, this does seem to convey the correct intuition when thinking about decisions when the anticipation of regret may play a significant role.

For example, the most oft-cited example of a regret-based decision is when deciding what lottery numbers to play. Most people tend to play the same lottery numbers every week; the justification being that although there is no statistical advantage to be gained by playing any particular set of numbers, there would be a very high level of regret

² see Starmer [88] for a summary of the non-EUT literature

³ for example, the famous Allais’ Paradox[1]

⁴ proposed by Bernoulli [10] and axiomatised by von Neumann and Morgenstern [98]

⁵ that is, summing across all states of the world

⁶ such a “regret function”, however, does implicitly presume that it is possible to predict the negative impact of the experience of regret without ever needing to actually experience it

associated with changing from “your” numbers on a given week, and then seeing those numbers appear as the winning numbers. If this feeling of very high regret can be anticipated, it can induce an agent to choose the same numbers every week, hence insuring themselves against this possibility. This effect may even be sufficient to keep the agent playing the lottery over and over, resembling an addiction pattern of behaviour. This assumption, that the presence of large regrets will dominate an individual’s attitude, and hence behaviour, towards regret, can lead us to think of these early theories as ones which primarily model and capture “Regret Aversion”.

Whilst intuitively appealing, there has been very little development of regret aversion theories since their early introduction. Aside from an axiomatisation of Loomes and Sugden [58]’s version (Sugden [89]) and a demonstration of the theoretical similarity to other non-expected utility theory models (Loomes and Sugden [60]) the theoretical and experimental literatures have largely relied on other non-EUT models to progress the research on choice under uncertainty (with Prospect Theory (Kahneman and Tversky [48]) being the most widely used). There are two main reasons for this lack of development.

Firstly, when looking to develop theoretical models which run contrary to EUT, at least one of the axioms of EUT (von Neumann and Morgenstern [98]) must be violated. Most non-EUT models violate the independence axiom (that preferences over two lotteries are maintained when those two lotteries are mixed with a third “independent” lottery in the same fashion), but models of regret aversion violate the axiom of transitivity (if lottery A is preferred to lottery B, and lottery B is preferred to lottery C, then lottery A must be preferred to lottery C). Though there is experimental evidence to suggest violations of both axioms occur in reality, transitivity is widely considered a far more fundamental component of rational decision making than independence, and violations of independence are observed more frequently in experiments than violations of transitivity. Hence there is reluctance to use and develop a model which dispenses with transitivity in preference of one which dispenses with independence (such as Prospect Theory), when experimental violations of EUT can be explained using both models.

Secondly, in accordance with the simultaneous development of both behavioural and experimental economics, there is a necessity to test new models of decision making (especially under risk and uncertainty) using experimental procedures. However, for a model which incorporates an anticipated emotional response, such as regret aversion, this becomes significantly more challenging. Either you need to design an experiment in such a way that the only possible explanation for an observed pattern of decision making is regret⁷, as specified by the theoretical predictions of the relevant model, or you need to actually observe, measure and record the emotion of regret directly, and infer the link from the relationship between the observed actions and measured emotional responses. Unfortunately, neither of these options is straightforward, as discussed subsequently and in the second chapter of this work, and so the predictive success of a regret aversion model is very difficult to determine.

⁷ as opposed to any other related emotions, such as disappointment (Bell [7] & Loomes and Sugden [59])

Experimental History

The history of experimental research into regret is covered in detail in the second and third chapters of this work. The earliest experiments, which provided useful insights into the process by which regret could be generated (and circumstances in which it is reduced), are only mentioned in future sections of this chapter where they have specific relevance to the motivation for this chapter.

Recent Theoretical Literature

Recently, however, regret has been reintroduced into the economics literature, with both new theoretical models and experimental insight.

Sarver [83] uses the basic “expected utility minus regret” formulation of the original regret aversion papers as the second step in a two stage decision making process. In the second step, the agent calculates the value of an action in the now traditional fashion; calculating the “modified” expected utility, given the other actions available, and hence the possibility of regret. Given a “menu” of actions on offer to the agent, therefore, the value of the menu can be given by the value of the action with the highest modified expected utility on offer on that menu, taking into account the possible regret of other actions available on the same menu. Sarver’s addition to this process is then the introduction of a first stage, where the agent must choose between a variety of menus on offer.

Consider 2 menus of only singleton actions $\{a\}$ and $\{b\}$, and suppose that the agent prefers the menu with only a to the menu with only action b . Then it is clear that adding action b to the menu with only action a on, creating a new menu $\{a, b\}$, will have lower value than the menu with only the preferred action, a , on, as it introduces the possibility of regret whilst adding no superior actions. However, it is unclear whether adding the superior option a to the menu containing just b will, in fact, improve the agent’s situation. Standard economic intuition suggests that adding a superior action to the agent’s set of options should only improve their utility, but if the weight of anticipatory regret is strong enough, which Sarver uses the term “regret aversion” to describe, then the agent might prefer the menu $\{b\}$ to the menu $\{a, b\}$, essentially preferring to not have a choice at all.

Though Sarver’s work is in the same spirit as Loomes and Sugden [58] and Bell [5], and even has an axiomatic derivation in the style of Sugden [89], it differs from and develops these models by extending regret aversion from simply choice over actions to choice over menus of actions.

Another theoretical avenue into which regret has been introduced is dynamic game theory, with work by Hart and Mas-Colell [32] and subsequently Hart and Ben-Porath [31], discussing the concept of “regret matching” in repeated games. In a repeated game, if an opponent is playing independently identically distributed actions, then playing a *regret matching strategy*, whereby I switch from an action which I have experienced regret from in the past, to an action which I would have been better off playing at those times, with a positive probability given by the size of the regret experienced, is a strategy which ultimately converges to a correlated equilibria, without the need to compute something equivalent to a best response function.

This type of research moves us away from anticipated regret to *experienced regret*, whereby an individual is responding to emotional cues in their memory (in this case the regret experienced from choosing a particular action previously) rather than attempting to predict their future emotional state. This approach will be developed and incorporated later on in the numerical simulation section.

The most recent theoretical literature is that of Hayashi (both in [2008] and [2011]). These models run in the spirit of the original papers, by creating “...axiomatic model[s] of choice under uncertainty where the decision maker may be driven by anticipated regrets” [34, p255] but differ through the idea that the decision maker is not trying to maximise “utility minus regret”, as in the case of Loomes and Sugden and Sarver, but rather trying explicitly to minimize regret, as in the original idea of Savage [84]. This slight change in axiomatisation and ideology has the benefit of producing a “smooth” model of regret aversion [34, Theorem 3], which explicitly parametrises a “coefficient” of regret aversion, hence allowing the minmax regret model of Savage to sit as a special case of the Hayashi framework⁸. Indeed, this helpful parametrisation of regret aversion will be exploited later on in the chapter (1.3.1), where the precise specification of the Hayashi model will be expanded upon. However, the same limitation of the original models of regret aversion applies, in that Hayashi “... limit[s] attention to *anticipated* regrets, that are directly relevant to choice. What about actual emotion of regret, which the decision maker may feel after making action and seeing resolution of uncertainty?” [34, p255] It is this question, and the link between anticipated and experienced regret, which the rest of the chapter is intending to explore.

Recent Experimental Literature

Experimental research has also progressed to incorporate new techniques, developed in neuroscience, to find novel ways to measure and observe emotional responses. As mentioned previously, testing a theory which uses emotions as such an integral component, like regret theory, is difficult without explicitly observing the emotion the theory claims is important.

Mellers et al. [63] approach the problem of measuring emotional responses in perhaps the most obvious, but probably the most controversial, fashion; by simply asking participants in the experiment to self-report their emotional state. Using a scale of -50 to +50, and a series of gambles with varying gains, losses and associated probabilities, participants are asked to predict their emotional responses and then asked to rate experienced emotional responses, using the scale. By comparing outcomes in which information was only given about the gamble which was chosen (Partial Feedback) with outcomes in which information was also given about the gamble which was not chosen (Complete Feedback), they are able to separate the effects of regret and disappointment and show that regret is a more powerful emotion than disappointment, especially when gambles contain losses. They also claim that subjects’ predictions of their own emotions are generally accurate, as in the below graph.

⁸ specifically “...the choice rule gets closer to Savage’s minimax regret choice as α tends to infinity.” [34, p254]




Figure 8. Results from Experiment 4: average actual emotions plotted against average anticipated emotions in both partial and complete resolution tasks.

Figure 1: Figure 8 from Mellers et al. [63]

Assuming regret will only be experienced when there is Complete Feedback⁹, the graph indicates that negative anticipated emotions in the Complete Feedback case are accurately forecasted, which would seem to indicate that regret is being correctly predicted. This claim will be assessed in greater detail when discussing the so-called “Affective Forecasting” literature in due course.

This approach, however, has several limitations. The obvious criticism is that, although the results seem to corroborate the standard regret-based conclusions¹⁰, there is no incentive being offered to participants for truthfully reporting their own emotional states¹¹. In the case of experienced regret, this is an unavoidable problem, as the subjective experience of regret may only be known to the person experiencing the emotion, but in the case of anticipated regret, we would like to make use of real, incentivised choice behaviour to show that people

⁹ that is to say, by learning about both the outcome of the selected and unselected gamble, it permits the comparison of utilities on which the experience of regret is based

¹⁰ as an example, “[p]eople feel better about their own outcome if the outcome of the other gamble was worse” Mellers et al. [63, p336]

¹¹ Whilst not explicitly regret focussed, there has been a significant recent literature in experimental economics looking at the predictive power of self-reported measures of emotional states, started by Frans van Winden. A good example of research in this area is by van Winden et al. [95], which “...is concerned with the impact of the timing of the resolution of risk on investment behaviour, with a special focus on the role of affect.” In this work, emotions are only captured through self-reports, with measurements taken at the time of anticipation, experience and recall. Their results show that there is an “...improvement in the predictive power of the model (which) appears to come from the affect variables”, thus lending weight to the idea that standard economic incentive compatibility may not be a significant issue when dealing with the impact of emotions on decision making. However, when dealing with regret specifically (and in isolation), there are two other key considerations. Firstly, the influence of the role of individual responsibility on the feeling of regret (as discussed later in the thesis) gives rise to problems with the role of cognitive dissonance and the interaction with self-reported measures. If I experience regret, but then manage to explain away the emotion through reducing my own sense of responsibility, then a self-report of that experience will be subject to the process of rationalisation, and hence be an imperfect measure. Secondly, whilst regret has a specific meaning in behavioural economic theory, it is easy to imagine an experimental participant confusing similar, but distinct, emotions such as disappointment and anxiety when providing the report. It’s easy to identify “positive” versus “negative” emotions, when asked, but less so to call out any one specific feeling over another.

who anticipate high regret from certain gambles are likely to shy away from choosing them.

The second criticism is that the scale on which emotion is being reported is not a specifically “regret” scale, but a much broader “emotion” or “happiness” scale. The link to regret is then inferred from the responses in gambles where regret is prevalent (Complete Feedback conditions), under the assumption that regret will form part of the emotional response to which the scale is referring. The scale, however, asks explicitly about *disappointment*¹² and so it cannot be automatically inferred that the anticipated or experienced emotion of regret would be included or recorded in a subjective affective scale asking explicitly about disappointment.

For a less subjective approach to the problem, however, we need to draw further from work in neuroscience. Camille et al. [16] and Coricelli et al. [19] follow on from the work of Mellers et al. by asking participants to rate their emotions using a -50 to +50 scale (this time ranging from *extremely happy* to *extremely sad*) after, again, choosing a series of gambles. Camille et al. identify the orbitofrontal cortex and amygdala as two areas of the brain which are associated with reasoning, planning and emotions (and, hence, potentially, anticipated and experienced regret) and so seek “...to test whether the ability to experience these emotions is mediated by the orbitofrontal cortex” [16, p1167] by comparing the subjective affect reports of “normal” people and people with orbitofrontal cortex lesions. For “normal” subjects, they find results very much in line with Mellers et al. (Graph C in Figure 2), but for orbitofrontal patients, the outcome of the non-chosen gamble has no impact on the subjective affect reports in complete feedback (regret) conditions (Graph D in Figure 2).

This result suggests that “...the orbitofrontal cortex exerts a top-down modulation of emotions as a result of counterfactual thinking, after a decision had been made and its consequences can be evaluated” [16, p1169]. Hence, to move from subjective reports of the experience of regret towards determining the root processes behind the magnitude of experienced regret requires further understanding of the role of the orbitofrontal cortex in decision making.

Coricelli et al. [19] approach this research through the use of functional magnetic resonance imaging (fMRI) of subjects whilst they perform the same gambling task as used in Mellers et al.. In addition to using both Complete and Partial Feedback conditions, to separate regret and disappointment, they also assign some subjects to a “follow” condition, where, instead of choosing the gambles, they are simply asked to follow the gambles which the computer selects. As the experience of regret is positively linked to the degree of responsibility an agent feels for the outcome¹³, subjects in the “follow” condition should feel less regret than subjects in the “choose” condition, especially in the presence of complete feedback. In addition, by observing patterns of brain activity in the “choose” condition compared to the “follow” condition, it should be possible to identify areas of the brain associated with the *anticipation* and *prediction* of emotion when subjects in the choose condition are required to evaluate the options available to them. Again, Coricelli et al. find evidence for the involvement of the medial

¹² “[p]articipants...expressed their feelings on a category rating scale that ranged from 50 (*extremely elated*) to -50 (*extremely disappointed*).” Mellers et al. [63, p335]

¹³ the literature on the link between regret and responsibility is comprehensively discussed in Chapter 3

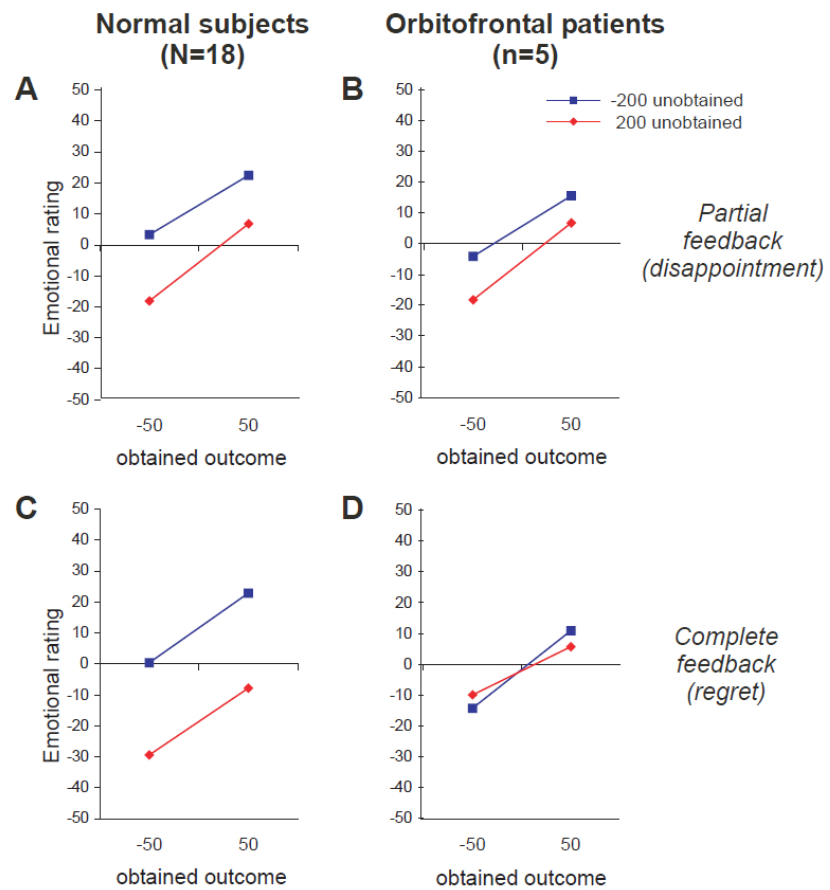


Figure 2: Figure 2 from Camille et al. [16]

orbitofrontal cortex, and also support for the role of responsibility in triggering activity.

Whilst this line of research is fascinating and informative for those interested in the role of regret in decision making, it seems the direction of future work in this area will be centered around a better understanding of these specific brain processes and activity through the use of fMRI and other such neuroscientific techniques. For those of us without access to such equipment, however, we must look at other questions in the area of regret.

Future Regret Research

So far we have considered a time line of regret whereby the anticipation of regret¹⁴ must occur before the experience of regret. However, [Coricelli et al.](#), through the design of their experiment, hint at another time line which may house equally important research questions. By asking subjects to perform multiple gambling choices in sequence, not only are subjects assumed to be repeatedly using the anticipation of regret in their decision making, they are also repeatedly experiencing regret when their choices are shown to be wrong. We have shown, above, how the anticipation of regret can have an effect on decision making, but an equally important question concerns how the *experience* of regret will subsequently influence decision making.

If there was no effect from the experience of regret on decision making, we would expect that subjects who had experienced regret from their decisions in the gambling task would subsequently behave similarly to those who did not¹⁵. [Coricelli et al.](#) approach the answer to this question by looking at changes in the proportion of choice attributed to anticipated regret over time in the experiment¹⁶ and whether or not there was fMRI evidence that the experience of regret in a preceding gamble was changing the pattern of brain activity in the subsequent trials.

As their results [[19](#), Figure 5, p1259] show, there was an increase in the role of anticipated regret in choice behaviour later on in the gambling sequence (after regret had already been experienced), and an enhanced response in the right dorsolateral prefrontal cortex, at the time of choice, when there had been an experience of regret in the preceding gamble, "...perhaps representing an influence of immediate regret on self-monitoring at decision making." [[19](#), p1259]

The precise channel through which the experience, not simply the anticipation, of regret can influence decision making is, however, unknown. Possible channels include a "negative mood" explanation whereby the experience of regret changes the mood of the decision maker, possibly reducing their happiness, which then subsequently influences decision making¹⁷. Another possible explanation, more self contained within the context of regret, is that the experience of regret changes the subsequent sensitivity to potential future regrets, or, in other words, changes the degree of anticipated regret aversion which is used in the decision

¹⁴ calculated in order to correctly account for the expected reduction in utility, from the experience of regret, associated with any given action

¹⁵ controlling for correlated factors, such as wealth effects

¹⁶ with the implicit assumption being that the cumulative effect of experienced regret, as the experiment progresses, changes the sensitivity to regret, or regret aversion, in later trials.

¹⁷ for a summary of the significant literature on the effect of negative mood on decision making under risk, for example, see [Hockey et al. \[37\]](#)

making process. As will be discussed in the next section, this feedback channel, of the experience of an emotion changing how it is anticipated to affect a person in the future, is not a new idea. And, indeed, there is experimental evidence from Zeelenberg and Beattie [105] on experienced regret in the ultimatum game, Creyer and Ross [21] on the effect of experienced regret on price setting, and Coricelli et al. [19] on the effect of experienced regret on future gambling behaviour, which seemingly lend support to this idea.

As such, this chapter will describe a model which seeks to analyse the various channels through which experienced regret might influence subsequent behavior. The following sections provide motivation and evidence for the way in which I have described these channels in my model.

1.2.2 *Predicted, Decision, Experienced and Remembered Utility*

The previous discussion on anticipated and experienced regret hints at a useful distinction to be made when talking about all emotions. Is there a difference between what we *think* will happen to us, what *does* happen to us, and what we *remember* happening to us?

There had been a recent revival in economics in thinking of decision making as contributing to our affective or hedonic state of mind. The pleasure and pain we derive from our choices and decisions has been incorporated into theories like Regret Theory [58][5] and Disappointment Theory [59][7], but it is interesting to note that this approach is different and divergent from the standard neoclassical utility theory, where the concept of utility is used more as a convenient mathematical representation of the *revealed preferences* of an individual, as in Samuelson [81]. Samuelson's work, however, was the culmination of a long process in economics designed to remove any vestiges of psychological interpretation from analysis; a process which had successfully discredited "...utility as psychological concept [, robbing] it of its only possible virtue as an *explanation* of human behaviour"[81, p61].

Kahneman et al. [53] note that this modern understanding of utility is a far cry from the concept of utility as originally proposed and developed by Bentham [9]. Bentham expresses,

By utility is meant that property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness, (all this in the present case comes to the same thing) or (what comes again to the same thing) to prevent the happening of mischief, pain, evil, or unhappiness to the party whose interest is considered [9, Ch I.4]

which is much more reflective of utility as a concept to be experienced, rather than mathematically representing a decision.

Indeed, Kahneman et al. feel it necessary to completely separate these two concepts as they relate to different ideas. They term the modern usage as "Decision Utility", referring to the weight that an outcome holds when making a decision, and the Bentham concept as "Experienced Utility", referring to the subjective hedonic experience of an outcome. The early 20th century economics literature can, therefore, be seen as a rejection of experienced utility in favour of decision utility as the beliefs of the time were

- (i) subjective hedonic experience cannot be observed or measured
- (ii) choices provide all necessary information about the utility of outcomes because rational agents who wish to do so will optimize their hedonic experience [53, p375]

By introducing two further concepts, Remembered Utility (the memory of utility experienced from a particular outcome) and Predicted Utility (the belief of an agent about experienced utility of an outcome, prior to choice or experience), [Kahneman et al.](#) describe an encompassing timeline about the process of choice¹⁸. Whereas economists would typically argue that rationality in decision making implies consistency across all these concepts (“If A gives me a better hedonic experience than B (Experienced Utility) why would I choose B to A (Decision Utility)?”), the proposed timeline hints at a different definition of rationality. For an agent, so long as the option they *predict* to be better is actually *chosen* (so there is consistency between Predicted and Decision Utility), then they can be considered rational, even if their psychology leads to *mispredictions* (so their Experienced Utility differs from their Predicted Utility) or *misremembering* (so their Remembered Utility is different to their Experienced Utility). The extent to which these mispredictions and misrememberings occur, and the consequences thereof, will be discussed in the next section, but their existence justifies the position of [Kahneman et al.](#) in seeking to expand the definition of utility in this fashion.

Further usefulness of this timeline of utility can be demonstrated by considering dynamic models of decision making, or models of repeated decision making. In a dynamic or repeated context, there is information to be gained about an outcome in the current choice problem (at period t) from the experiences of the past (at periods $t - 1$, $t - 2$, etc.), and, hence, hopefully ensure that Predicted (and hence Decision) Utility are more closely related to the true Experienced Utility of the outcome. In [Kahneman et al.](#) framework this process happens through the Remembered Utility of an outcome. If, for example, the Predicted Utility associated with a decision problem is based, at least in part, on the Remembered Utility of a previously encountered similar decision problem, then we have developed a very simple, cyclical feedback and learning model¹⁹.

This development of [Kahneman et al.](#) idea is best represented by Figure 3 from Baumeister et al. [2].

In their work, they discuss a feedback process for emotions, on behaviour, that is conceptually equivalent to idea of [Kahneman et al.](#) utility, whereby the “affective residue” from past experience (or the Remembered Utility) drives the anticipated emotional outcomes (Predicted Utility) of the options, and hence lead to a choice based on the desired emotional outcome (Decision Utility). This similarity allows us to extend the work of [Kahneman et al.](#) (and hence the original ideas of [Bentham](#)) to models of decision making which deal in terms of emotion rather than specifically utility.

¹⁸ and this timeline has largely been overlooked in, specifically, economic models of decision making

¹⁹ with the specific details of such models being explored in a later section.




Figure 3 Anticipated emotional outcomes guide subsequent behavior.

NOTE: Solid lined arrows indicate causal relationship in which the process creates the effect. Dashed lined arrows indicate associative relationship in which the process activates a set of associations.

Figure 3: Figure 3 from Baumeister et al. [2] representing the feedback loop from Remembered Utility to Predicted Utility

Link to Regret Aversion

And, of course, one such model of decision making which relies on emotion (in this case, in addition to more familiar ideas of utility) is Regret Aversion.

As mentioned in the theoretical history section, models of regret grew out of a literature on non-Expected Utility Theory, which, of course, grew out of a literature on Expected Utility Theory, and so rely heavily on the notion of utility in the Decision Utility context as described above. However, what the recent experimental literature on regret hints at, as explained previously, is a distinction to be made between anticipated regret and experienced regret, or, described in the context of Kahneman et al. [53] and Baumeister et al. [2], Predicted Regret and Experienced Regret. Making this distinction pushes us away from the traditional use of regret in economic models, as theoretical notation in describing new non-expected utility models which display predictions of behaviour in line with experimental evidence, towards considerations of regret as a real, tangible, hedonic, emotional concept, which is subject to the same kinds of psychological, and psychophysiological, patterns of behaviour as other negative emotions, such as fear and anger²⁰.

Indeed, Figure 3 describes a plausible, and more complete, process through which considerations of regret can play a role in decision making.

The stage referring to “Anticipated emotional outcomes of simulated options” is precisely the kind of Predicted Regret which must be calculated in any standard model of Regret Aversion, and “Select option based on desired emotional outcome” describes the maximisation of a non-expected utility model, modified by regret, as in Loomes and Sugden [58], Bell [5] or, more recently, Hayashi [34]. Indeed, really the first such model to incorporate regret into decision making in some fashion, the Minmax Regret principle of Savage [84], follows this same idea.

However, it is the processes prior to this in the diagram, those of “Memory of behavior and its outcome at Time 1”, “Affective Residue” and “If-then rules”, which are of interest to this work, based on the recent experimental literature. These are the parts of the process which incorporate the idea of regret as an emotion, not just theoretical no-

²⁰ regret is a good starting point for including “emotions which behave like emotions” in models of decision making, precisely because we can imagine a fairly simple mathematical representation as the baseline from which to work and develop. This contrasts with other negative emotions, such as fear and anger, which offer no such simple representation and hence are a lot more difficult to incorporate within a standard economic decision making framework.

tation, but will prove crucial to the latter stages of the process which we are more familiar with. At the most basic level, the question is whether we can predict our own emotional responses (Predicted Regret) without having experienced them first (Experienced Regret), and, even if we have experienced them first, whether our memories of those emotions (Remembered Regret) are accurate enough to be beneficial to us when making future decisions. In the language of Figure 3, is the affective residue from our past regretful experiences helping us to better predict²¹ our future emotional outcomes?

Fortunately, the question of whether people have the ability to use remembered emotional outcomes to guide their emotional predictions to be better estimates of their true experienced emotional outcomes, has been widely explored in psychology under the title of *affective forecasting*.

1.2.3 Affective Forecasting

Affective forecasting is an area of research which is principally concerned with the accuracy to which people can predict their own emotional responses to events. More recently, it has expanded to include the accuracy to which people can remember their own emotional events. More specifically,

“Affective forecasts can be broken down into four components: predictions about the valence of one’s future feelings, the specific emotions that will be experienced, the intensity of the emotions, and their duration. People can be accurate or inaccurate in predicting each of these facets of emotional experience.” Wilson and Gilbert [100, p346]

The primary focus here will be on the *intensity* of emotions, and the various mispredictions and biases which can arise, but future work could also focus on the specific emotion (regret versus disappointment) and duration (should regret not simply be formulated as a one-off reduction in utility, but a stream of utility reduction over time).

For our purposes, one finding in the literature stands alone in terms of the impact it can have for theories of decision making which rely on our ability to forecast our own emotions. “The most prevalent error found in research on affective forecasting is *the impact bias*, whereby people overestimate the impact of future events on their emotional reactions” [100, p353]. The wide range of situations and contexts in which this effect has been reported supports the impact bias as a robust finding, and suitable for inclusion in any model about the ability of people to forecast their own emotions.

As noted previously in Figure 3, a key component of an emotional prediction will be the memory of that emotion from times it has been experienced before. “It is well known, however, that memory for past affective experiences is poor.” [100, p358] One key result which speaks to this is the so-called “Peak End Rule”²², which states that a subsequent evaluation of an affective event (after the event has taken place) is very highly correlated with the average of the Peak affect and End

²¹ perhaps using some kinds of “if-then” rules or other heuristic

²² first proposed by Varey and Kahneman [96], and subsequently tested by Kahneman et al. [52], with respect to the pain of cold water, and then famously by Redelmeier and Kahneman [74] with respect to colonoscopy evaluations and memories

affect, given by online, or moment-by-moment, reports of the event. The implication of this rule is that, for example, "...adding a period of pain to an aversive episode will actually improve its remembered utility, if it lowers the Peak/End average" [46, p190], clearly violating dominance. As Fredrickson states,

"The primary lesson to be drawn from the existing empirical literature on the peak-and-end rule is that people evaluated certain kinds of past affective experience by referencing just a few selected moments. One or two moments, then, play a privileged role in guiding people's choices about which past experiences they would avoid, and which they would repeat, or recommend to others." Fredrickson [25, p588]

"Thus, to the extent that people predict their future online experiences (how they will feel during a colonoscopy or their next vacation in Paris) from their recall of their past experiences, systematic errors are likely to occur." [100, p359]

An example of this misprediction and misremembering of affective states comes from Wirtz et al. [102], who use the experience of going on "spring break" as an environment to assess the types of affective forecasting errors which can occur. By providing 41 students with PDA devices, which allow the students to report the magnitude of "online" affective states, it is possible to compare what the students "thought spring break would be like" to what the students "actually experienced on spring break" to what they "remember the experience of spring break being". The results of the study are discussed in more detail in relation to the role of memory in 1.2.5 and Figure 6, however, they show significant variation in the affective reporting of the spring break experience depending on the point in time at which the affective report has to be made.

These empirical results, on the errors which frequently occur when being asked to predict and recall emotional experiences, guide us towards creating better and more representative theoretical models of how people will behave when being asked to do just that, but do not particularly explain the reasoning behind why such errors will occur. One hypothesis is that people often fail "...to anticipate [their] own ability to make sense of the world in ways that minimize its emotional impact" [100, p384]. If people are able to come to reasonable, rational, calm and collected justifications for emotional outcomes, then this should limit both the intensity and duration of negative emotional events when they do occur. "The fact that the impact bias is by far the most common error found in affective forecasting research is testimony, perhaps, to the pervasiveness of people's tendency to fail to anticipate their own sense-making processes." [100, p384]

Link to Regret Aversion

As shown previously, a key component of most regret-based decision models is the ability to anticipate your regret when choosing option A, and, ex-post, realising that you should have chosen option B. The robustness of the finding of an impact bias from the affective forecasting literature suggests that, when in a position to do this, individuals will routinely overestimate the negative emotional fallout they will actually experience.

As suggested above in 1.2.3, one process through which this bias can occur is the failure of an individual to anticipate their own ability to rationalise an emotional outcome, and hence reduce its emotional impact. This is an especially important consideration when talking about regret, as regret is a process which relies on feelings of responsibility and self-blame to operate. For example, consider a person who must decide which of two queues to stand in at the supermarket checkout. Queue A has more people, but each person in Queue B has more goods. The individual decides that “more people” dominates “more goods” in terms of slowing down the queues, and so decides to stand in Queue B. However, as it transpires, each person in Queue B has more packing of bags to do, and pays by card instead of cash, which means that Queue A ends up proceeding quicker. It is easy to *anticipate* why this person would feel regret of their decision, as they ultimately chose the wrong queue. However, the extent to which regret will actually be *experienced* by the person will depend on the extent to which they re-interpret history, and place blame upon themselves for not taking into account, at the time of their decision, information which became salient to them only after the decision had been taken. That is to say, do they think that they should have known each person in Queue B would have taken longer to pack bags, and would have used a card to pay, at the time that the decision was taken? If they do, then they indeed made a mistake in their decision making process, as good quality information, which would have lead to a better outcome, was ignored, and hence they feel regret. However, if they feel that this situation occurred simply as a result of chance, and equally there were other things which could have occurred which would have lead Queue A to proceed slower (more people in Queue A means more chance of one person having problems with their payment, which causes a big delay, for example), then the outcome they experienced was not as a result of a poor choice on their behalf (regret) but rather the misfortune of bad luck (which is disappointment). Indeed, an economist would likely say that the decision they took was “rational”, given all the information you had available at the time, and hence was the best thing you could have been expected to do, so why would you ever feel that a poor decision had been made?

Thinking in these terms, and in line with the work on affective forecasting, suggests that people will underestimate their own ability to *rationalise* and *explain away* regret in this fashion, reducing the blame they will attribute to themselves for a “bad” decision, once all outcomes have been revealed. As a test of this hypothesis, Gilbert et al. [26] ran an experiment whereby the margin to which people made the wrong decision was manipulated, and looked at both forecasts and experiences of regret. In the narrow margin treatment, participants missed out on a prize (Study 1) or missed catching a train (Study 2) by a very slim margin. In this situation, it is very easy to imagine how participants could blame themselves for the outcome (as even a slight change in their decision making would have lead to a substantially better outcome), and hence anticipate high levels of regret. In the wide margin treatment, where participants made the wrong decision by a wide margin, self-blame is less easy to imagine (as it would have required a radically different decision process to change the outcome for the better), and hence the anticipated level of regret is much lower. However, if it's true that participants in both groups have an ability to rationalise the process, leading to absolution and a freedom from blame, then their

experiences of regret will be small, and much lower than anticipated in the narrow margin treatment.

The results of Study 1 showed “...that the size of the margin influenced forecasted regret, $p=.009$, but not experienced regret, $p=.22$. Forecasters overestimated how much regret they would feel in the narrow-margin condition, $p=.02$, but not in the wide-margin condition, $p=.39$ ” [26, p347] and the results of Study 2 showed “...that the size of the margin influenced forecasted regret, $p=.04$, but not experienced regret, $p=.43$. Forecasters overestimated how much regret they would feel in the narrow-margin condition, $p=.004$, but not in the wide-margin condition, $p=.09$ ” [26, p348] and hence the results of both studies support the presence of an impact bias, caused by a failure to anticipate one’s own ability to rationalise bad outcomes, when dealing with the emotion of regret.

The consequence of these results, therefore, for decision makers who make decisions based on anticipated regret, is that there can be cases where the anticipation of regret leads to a decision maker choosing “...gambles in which bad outcomes are likely but unregrettable over gambles in which bad outcomes are unlikely but regrettable.” [26, p349]. This, ultimately, leads to the decision maker “...purchasing emotional insurance that they do not really need.” [26, p350]

These results provide some interesting insight into theoretical models which assume “regret aversion”, or the property that large regrets loom disproportionately in an agent’s mind when compared to small regrets. If you take this assumption to be inclusive of the impact bias in anticipated regret²³, then the models, though accurately representing the decisions taken by agents, will not yield decisions which turn out to be ex-post optimal for the agents, even when compensating for the true negative effect of the experience of regret on utility. In turn, to learn about the true affective experience of regret, there would need a modification of the standard frameworks to explain the difference between Predicted Regret, and Experienced Regret.

What is left unexplained by the standard theoretical models, therefore, is why do agents not appear to learn from their mistakes once they experience regret, or, alternatively, why the impact bias appears to persist. For example, if you were to apply the standard models to a repeated decision making context (such as an individual making hourly, or daily, investment decisions), then an agent would keep making sub-optimal decisions, repeatedly purchasing emotional insurance that they don’t need, despite *experiencing* the realisation that they don’t need it. One possible explanation is that of standard cognitive dissonance, whereby an agent is refusing to accept that their beliefs about anticipated regret are being challenged by experience. This could be reasonably assumed for “big one off events” where considerable effort has been expended insuring against the possibility of experienced regret. However, in a repeated decision making world, it seems implausible to assume that this dissonance persists in the face of mounting evidence. Can we really believe that individuals simply do not learn at all from their mistakes,

²³ it may well be the case that large regrets, in experience, are not *disproportionately* larger than small regrets, but the impact bias works in such a way as to exaggerate the largest regrets in anticipation, creating the familiar regret aversion assumption. It also could be the case, however, that large regrets are experienced as disproportionate to small regrets, but the effect of the impact bias is simply to magnify each regret by a fixed amount (greater than one) in anticipation, which, again, would be consistent with the assumption of regret aversion.

or, equivalently, that the prediction of regret at time period t is in no way related to the regret experienced at time period $t - 1$? Thus, the keys to understanding this dynamic process will be an understanding of the way learning works (in terms of emotions, and learning about your own emotional reactions), and, secondly, an understanding of the way *memory* stores emotions and emotional events, which can be recalled in subsequent periods for future decisions, as suggested by Figure 3, to improve decision making.

1.2.4 Learning

When thinking about learning, we are asking two separate questions. Firstly, what kind of learning models could describe a process by which agents learn about their own emotional responses in order that they may have a better understanding of those responses when faced with a decision which has emotional consequences in the future? Secondly, if such a learning process exists, why would agents still retain an impact bias over time and not “learn” that emotions are lower in experience than in anticipation. Hence, the first is a question of existence and the second is a question of convergence.

Regret Matching

Most of the instances of models of learning being used in economics have been centred in game theory, especially where there is an obvious dynamic structure (sequential games and repeated games for example) and an obvious idea of what the point of the learning should be (convergence over time to a “rational” equilibria). The main question is whether a naive individual, who is using a series of learning heuristics, or if-then rules, to proceed through the game, is comparable to the typical, fully rational, backwards-induction solving agent we see in standard theory.

One such model, that, helpfully for us, also involves the idea of regret, is Regret Matching. As mentioned previously in 1.2.1 the works of Hart and Mas-Colell [32] and subsequently Hart and Ben-Porath [31] show that a decision heuristic whereby an individual, when in a repeated game, switches probabilistically away from actions which have yielded regret in the past, to actions which would have yielded better outcomes, is a strategy which converges almost surely to the set of correlated equilibria of the game [32, Main Theorem]. Though this is clearly a learning procedure which involves regret, it is not a procedure which involves learning *about* regret; instead using regret as a tool for learning about the game being played (and, as it turns out, successfully). However, the idea of using a simple, adaptive procedure where regret provides a level of feedback to improve decision making in the future, has the same features as the emotional feedback model of Baumeister et al. as shown in Figure 3²⁴, and hence points us in the correct direction. Indeed, the mechanism underlying this model is displayed in the title of a book chapter²⁵ written by Hart and Mas-Colell [33], called “A Reinforcement Procedure Leading to Correlated Equilibrium”. In this model, the *reinforcement* occurs through regret

²⁴ that is, regret being transmitted through memory from one decision making period to the next

²⁵ based on the original Hart and Mas-Colell paper

reinforcing “good” versus “bad” choices in the game. However, in keeping with the dynamic of remembered regret being used to inform about future predicted regret, we would like the reinforcement process to be working exclusively through the channels of regret.

Reinforcement Learning

Reinforcement learning can be defined as “learning what to do ... so as to maximise a numerical reward signal” where an agent “...must discover which actions yield the most reward by trying them”(Sutton and Barto [90, p3]) and is often associated with learning problems which have feedback information on a choice which can be used to improve decision making in the future. For example, a trial-and-error procedure can be thought of as a particularly simple case where the problem is attempted through “exploring” different options and “exploiting” information about which options yield reward and which don’t. A reinforcement learning system typically comprises four sub-elements : “...a *policy*, a *reward function*, a *value function* and, optionally, a *model* of the environment” [90, p7], which, combined, tell the decision making agent how to act (the policy) to achieve an immediate goal (the reward) in search of a long-run goal (the value function) given beliefs about the environment they are in (the model). The simplicity of this idea originated in psychology through the “Law of Effect” (Thorndike [92])²⁶ but has found recent traction in work on artificial intelligence²⁷ as an efficient mechanism through which problems can be solved by machines without the need for external supervision.

Interest in reinforcement learning has grown in recent years alongside the growth in both behavioural and experimental economics, and is centred around its explanatory power as a strategy employed in two player games with repeated interaction. For instance, considering a repeated game with a unique mixed strategy Nash equilibrium (MSNE), what can reinforcement learning say about the mechanics of behaviour which might ultimately converge to the MSNE? The review by Roth and Erev [80] of this literature concluded that, for explanatory power, “... a one-parameter reinforcement learning model outperforms the equilibrium prediction for all values of its one parameter”[80, p851], and that a “...four-parameter belief-based model ... improves on the one-parameter reinforcement model”, based on the aggregation of 12 experimental studies. A class of reinforcement learning models, in which agents exhibit a degree more sophistication, is that of rule-learning, whereby the agent evaluates the behavioural decision rule implemented, based on the outcome, rather than simply the action chosen. Stahl [87] formulates an experimental test of this framework, and finds the “... model fits the data much better than random noise or an error-prone Nash model”[87, p133]. However, both belief- and rule-learning types of reinforcement learning models have primarily been studied in the context of game theory rather than decision under uncertainty, in essence, testing their ability to describe learning about an opponent rather than one’s own preferences.

The key idea in reinforcement learning is the notion of “evaluation” of information by the agent. In the case of Regret Matching the evaluation comes in comparing what happened, to what could have happened,

²⁶ the “Law of Effect” “...describes the effect of reinforcing events on the tendency to select actions”Sutton and Barto [90, p18]

²⁷ summarised by Sutton and Barto [90]

to guide future choice strategies. Hence, this is one particular type of reinforcement learning process. However, despite economics being a discipline which extensively uses such ideas as policy, reward functions, value functions and models, the evaluation stage typically only comes in before a decision is to be made, in order to *determine* the policy. By incorporating the idea of reinforcement learning into economic models (and the similarity of the key elements used in both suggests that this is perfectly achievable) we can develop economic models where an agent will use evaluation to learn about their own decision making process, and not simply use evaluation at the time a decision is to be made.

Link to Regret Aversion

In the standard models of regret aversion, the evaluation phase comes in evaluating the possible options, based on an anticipation of what regret will feel like should it be experienced in the future. However, as the work in affective forecasting tells us, the anticipation of regret is subject to an impact bias whereby regret will be worse in anticipation than in experience. Hence, there is a second opportunity for evaluation, in that better information about regret is made available once it has been experienced, and so we have the potential to learn about our own beliefs and predictions at the same time²⁸. In addition, the work in neuroscience tell us that regret is an emotion which works hand in hand with learning, as the areas of the brain which appear to process regret, the amygdala and orbitofrontalcortex (OFC)²⁹, are both connected to learning processes³⁰.

However, the extent to which there is scope, in a decision making problem, to learn from regret will be determined by the degree of uncertainty which exists in the decision making problem itself. In the case of a “trial and error” process, the reinforcement learning from regret will solely be towards eliminating options which appear to be “bad choices” in favour of trying new options which could potentially be “better choices”. In a trial and error process, nothing is known ex-ante about the environment in which the decision maker is operating, and it is through a process of search and reinforcement learning that one can hope to learn how to make better decisions.

In the intermediate case of regret matching there is a repeated game theoretic situation where the structure of the game is known to the agent who is using a regret matching strategy, but levels of uncertainty are added by the fact there is another player in the game, and it is not necessarily assumed that the agent has the sophistication or rationality to compute best-response strategies. In this situation, the agent is reducing the level of uncertainty in the game each period, by learning (through regret matching) about which strategies yield good outcomes and which yield bad, so that they can “better respond” as opposed to “best respond” in the future. In this case, it is not exactly uncertainty of the game that is being reduced through learning (as the structure of the game is known ex-ante), but rather uncertainty about the agent’s

²⁸ by thinking “actually, that wasn’t as bad as I thought it would be”

²⁹ as shown by Coricelli et al. [19]

³⁰ “...the OFC integrates cognitive and emotional components across the entire process of decision making; when it malfunctions, it results in behavior that is maladaptive to ongoing contingencies” and “...neuroimaging studies assign a fundamental role to the amygdala in classical conditioning experiments, indicating its role in associative learning” Coricelli et al. [20, p262]

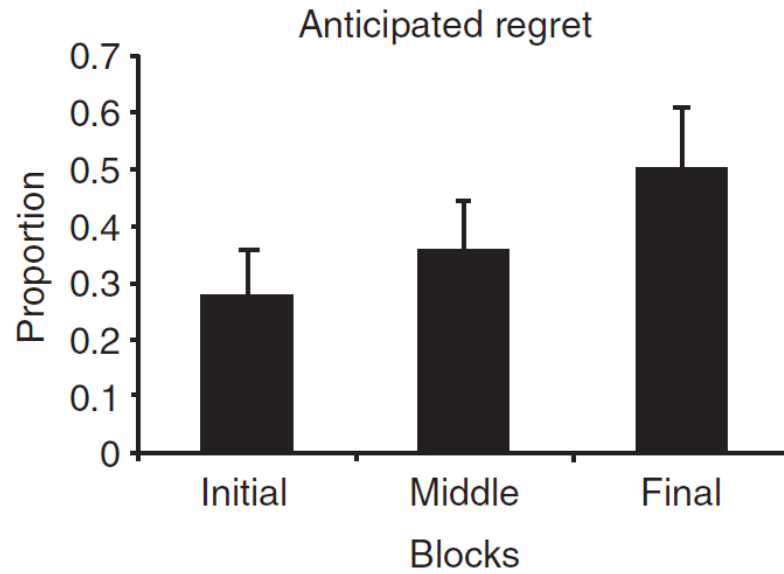
a

Figure 4: Figure 5 from Coricelli et al. [19]. Activity at choice: learning from the experience of regret. (a) Proportion of choice (+/- s.e.m.) related to anticipated regret in 'complete choose' trials. Anticipated regret increased over time as the experiment proceeded.

best-response behaviour, which is limited due to prior assumptions about rationality.

At the other end of the scale are the repeated gambling tasks used in Coricelli et al. [19] and Camille et al. [16], which both have the property of being “decision under risk” tasks as opposed to “decision under uncertainty”. This is true because the agent is simply playing a game against nature, where the rewards and associated probabilities of obtaining them for each option, are fully known prior to each task undertaken. In such situations as these, any regret which is experienced from the result of choosing option A as opposed to option B cannot be used, under reinforcement learning, to learn about the task at hand. The task is fully described without the need for experience to teach about the nature of the task. Yet, the results of Coricelli et al. [19] as given in Figure 4, for example, show that the pattern of behaviour does change over the course of the experiment as regret is experienced.

Figure 4 shows that the participant appears to be “learning” to choose options which have lower anticipated regret, as they experience more regret, but this learning is not linked to learning about the simple numerical values of the problem (as with “trial and error” and, to a lesser extent, “regret matching”), given they are always known with certainty throughout. This suggests that if the learning from regret is not linked to any objective values in the problem, they must be linked to subjective experience of the outcomes. Here, regret is not a signal that a particular option is worse than you previously believed, but regret may, perhaps, be teaching you that regret is worse than you

previously believed, and it should be better anticipated and avoided³¹. The emotional reaction from making a particular decision is the part of the task which is unknown, ex-ante, and so the reinforcement learning, from experienced regret, should be attributed to learning about this aspect.

The fact that the experience of regret should teach and inform you about your own future reaction to potentially regretful decisions isn't particularly surprising. What is surprising, however, is that, given the impact bias (Gilbert et al. [26]) states that regret is worse in anticipation than experience, then we would expect the experience of regret to teach us to be *less* sensitive to anticipatory concerns of regret, and not more, as demonstrated by Coricelli et al. [19]. There are not many children, for example, upon learning that being thrown into a big pool of water without aid from a parent is not as scary as first imagined, and can, in fact, be enjoyed as swimming, suddenly become struck with *even more* fear than before when confronted with another, similar, watery pool-based situation. So, why does the experience of regret appear, in this experiment, to make us more sensitive towards it in future? There are a couple of possibilities to consider.

The first is that, as in the case of "trial and error" where regret acts as a useful evolutionary signal that a particular option should not be trusted again, then even in cases such as simple decision under risk where probabilities and outcomes are completely, objectively known and there is nothing new to be learnt between periods, then we can still become less trusting of a particular option which has betrayed us in the past, even though, from a rational standpoint, nothing objectively has changed from before. You may have thought option A was superior to option B at time t , but bad luck caused you to experience regret from choosing it, and so you now believe option B to be superior at time $t + 1$. You may well become aware that the regret experienced in this situation was not as bad as you expected it to be, but it was still a *negative affective emotion*³², which you have no great desire to experience again. As before, we have a great ability to rationalise our own decisions and experiences, but, in this case, the rationalisation arises from thinking we must have made a simple "mistake" in choosing A at time t , and, in fact, the correct decision all along was to choose B. If B is a safe option, then this type of behaviour mimics the idea of "once bitten, twice shy", and we appear to become more risk averse.

The second possibility to consider, and the one which will be explored further in the next section, is that the experience of regret is not being correctly recalled from memory when the next opportunity arises to make a decision. For whilst it is our *experience* of regret which will turn out to be better than expected (due to the impact bias), it is, in fact, the *memory* of that experience of regret which will drive our future decision making under a reinforcement learning type model. If our memory is incorrectly representing our experience of the emotion, then the impact bias, for example, could simply be negated if our memory *exaggerates* the emotion we experienced. As such, we need a feel for how our memory will represent the experience of that emotion to our future selves when they need to decide what decision to take at time $t + 1$.

³¹ the fact that the proportion of choice attributed to anticipated regret is increasing, as experienced regret increases, implies an increased sensitivity to regret in later trials

³² in contrast to the positive affective experience of enjoying swimming

1.2.5 Memory

Given economics is a subject mostly concerned with the prediction and anticipation of future events and consequences, and hence primarily forward looking as a science, it is perhaps not surprising that there has been little consideration within economics for the role of memory, which focuses on the evaluation and record of events and consequences past, and hence is primarily backward looking. What maybe considered more surprising is that given the rise of dynamic models, in all manner of economic contexts, there has been little attention given to the question as to whether or not a “history” or “information set” available to a decision maker is likely to ever be truly representative of their past experiences, given the need for a conversion process from experience to memory, storage of experience in memory over time, and then back from memory to decision making, for these dynamic models of choice to operate effectively.

Indeed, psychologists have long studied these 3 key phases of memory - encoding, storage and retrieval - and the levels of imperfection and bias which can arise in memory through all three stages. Yet, when economists have sought to incorporate memory into decision making models, there is a typical idea of the *sophisticated* agent who has the capacity to affect their own memory for their own benefit, as in the case of the agent who engages in “self-deception” (Bénabou and Tirole [8]) or the agent who engages in “rehearsal” (Mullainathan [64]). This approach sits more comfortably with the standard economic view of a rational agent, for whom very little is beyond their sphere of understanding or control, but sits contrary to the vast majority of psychology research where an agent is *subject* to their own imperfect memory processes, such as in the famous case of “Flashbulb Memories” (Brown and Kulik [13]) where it is not the intention of the agent to have a heightened memory for specific emotional events, but rather an uncontrollable, hormonally driven³³ response mechanism. As such, I will discuss memory from the perspective of a *naïve* agent who is subject to the encoding, storage and retrieval process without having explicit control over which experiences are encoded, stored and retrieved accurately.

Memory and Emotion

The relationship between emotion and memory was hinted at previously with reference to Flashbulb Memories, but, as Figure 5 indicates, this is only part of “...one of the fastest growing areas of research in psychology and related disciplines” (Uttl et al. [94, preface])

There are two distinct areas of importance in researching memory and emotion. The first is the impact of being in an emotional state on memory of events (as epitomised by Flashbulb Memories), where the “...emotion acts as a mental “highlighter” increasing the salience of information from the environment and from memory” [94, p39]. The second is the difference between memory of an emotional event, memory of a neutral event and memory of the emotion itself, and it is this area which will be most useful to the work at hand. For instance, “...there are some differences in how people remember emotionally positive and emotionally negative effects” [94, p30], “[e]motional memories

³³ the role of the hormone Cortisol in the formation of memory in stressful or emotional situations is discussed in Diamond et al. [23], for example




Figure 1.1 Number of articles pertaining to emotion, memory, and emotion and memory by year of publication, as revealed by a keyword search of the PSYCInfo database. The number of articles is standardized to a common starting point by dividing the number of articles containing the keyword “memory” by 22 and by dividing the number of articles containing “emotion” by 40.

Figure 5: Figure 1.1 from Uttl et al. [94]

also seem to be long lasting” [94, p17] but “[r]ecent research shows that, like memory for neutral information, emotional memories are subject to fading over time and biases in the direction of current goals and experiences” [94, p39]. Furthermore, the “...particular emotion that one is experiencing will have a large influence on what kind of information is deemed of central importance” [94, p6] so there is merit in investigating links between memory and emotion with reference to the specific emotion of interest.

Memory and Regret

As such, there are some studies which investigate the memory for specific emotions and specific emotional events, such as regret and regretful events by Beike and Crone [4]. Regret is of interest to them because “[p]eople continue to experience painful regret years after they experience an undesired outcome (Wrosch et al. [103])” yet “[t]he persistence of the experience of regret is puzzling, as people normally exhibit reduced emotional responses over time to remembered life experiences (Walker et al. [99])” [4, p1545]. This reduced emotional response over time is known as the “fading affect bias”, yet it seems as if regret has a different pattern of fading affect to other emotions. Indeed, they found that “[e]xperienced regret faded significantly over time only for regrets of inaction” [4, p1548], implying that the typical fading affect pattern of negative emotions does not hold for regrets caused by action. By manipulating the instructions surrounding framing and association of memories of regret they demonstrate that “[t]he fading




FIGURE 1. Study 1: Predicted, actual, and postdicted happiness following the 2000 Presidential election.

Figure 6: Figure 1 from Wilson et al. [101]

affect bias for regret was...disrupted when participants construed the life regret as open then attempted to forget regret-related thoughts, and when participants construed the life regret as closed then attempted to remember regret-related thoughts" [4, p1548], showing that memory for regret may be highly contingent on the context in which it was experienced and rationalised, in the same way the affective forecasting literature shows how the experience of regret differs from anticipation through the same channels.

Biases in Memory

The fading affect bias, described above, is just one of many ways in which memory can distort an event or experience from the true event or experience as it was first encountered. And whilst Beike and Crone provide "...evidence that different negative emotions exhibit different patterns of fading affect" [4, p1549] there has been significant research into consistent trends and patterns in memory bias which are robust to different context. The Peak-End Rule, as discussed in 1.2.3, for instance, whereby a few selective moments of an experience are remembered as representative of the entire experience, has been replicated widely since 1992.

With regards to memory of emotions, one recurring theme has been the "Retrospective Impact Bias" whereby emotions tend to loom larger in memory than in experience³⁴. Wilson et al. [101] analyse this effect with regards to the 2000 US Presidential Election, and study the predicted, experienced and remembered happiness of supporters of both George Bush and Al Gore.

³⁴ compared to the standard Impact Bias, where emotions loom larger in anticipation than experience

As Figure 6 shows, Bush (who won the election) supporters both anticipated and remembered stronger positive emotions than were actually experienced, and Gore (who lost the election) supporters both anticipated and remembered stronger negative emotions than were actually experienced.

In a similar vein, Wirtz et al. [102] ask students who are on their spring break to give predictions, online accounts and memories of their subjective experience of the vacation.

In keeping with the results of Wilson et al., Figure 37 shows that students appear to both over-anticipate and over-remember their own subjective emotions of the vacation. In addition, students were also asked about their desire to repeat the vacation, and "...the best predictor of participants' desire to repeat the break - indeed, the only predictor - was remembered experience." [102, p522].

The main consequences of these two findings are that "[a]lthough on-line measures may be superior for estimating experience, retrospective global evaluations may be superior for predicting people's future choices" [102, p522] and "[t]he fact that retrospective measures may be a better predictor of future choices than on-line evaluations, while at the same time being less accurate, points to the likelihood that individuals often make choices that fail to optimize hedonic experience" [102, p522]. Thus the biases which appear in memory of emotions appear to lead to sub-optimal choices in the future.

Link to Learning and Regret Aversion

In any adaptive learning process, where prior information is used to update current beliefs, it is the memory, or, more specifically, the retrieval, of that information at the time of updating, which will determine the overall result of the learning process. As described previously, because of the Impact Bias, there is an opportunity for learning about regret (and other emotions) once it is experienced, in realising that the emotion was not as bad as was previously imagined. Yet the results of the few studies³⁵ which explore this ability to learn about regret seem to suggest an opposite pattern; that the experience of regret makes you more sensitive to it, and not less.

The above research on memory addresses this apparent paradox through two channels. Firstly, the lack of "fading affect" for most regret and regretful memories ensures that the emotion will always appear vivid in a person's mind upon retrieval. The consequence of this is that the worst regrets will stay with a person for a very long time after the experience, so an individual trying to learn (consciously or subconsciously) about their own experience of regret will have the process dominated by the very worst regrets which refuse to fade in their mind. This could skew the frame of reference for the individual to one of "always expecting the worst" and hence give rise to the impact bias, which can be modelled as regret aversion, in future decisions involving anticipated regret.

The second channel is the Retrospective Impact Bias³⁶ whereby the memory of the emotion itself is unrepresentative of the emotion experienced, and, as with the Impact Bias, the intensity of the emotion is exaggerated. In the case of regret, this equates to the memory of regret

³⁵ Coricelli et al. [19], Camille et al. [16]

³⁶ Wilson and Gilbert [100], Wirtz et al. [102]




Fig. 1. Predicted, on-line, and remembered spring-break experience, measured at five different times. Higher ratings indicate greater enjoyment of the overall subjective experience and more intense positive and negative affect.

Figure 7: Figure 1 from Wirtz et al. [102]

being significantly worse than was actually experienced, and more in line with the initial prediction, which, itself, was an exaggeration. Hence, given that learning about regret, for the purposes of future choice, takes place based on the memory of regret, then any learning will be more in line with confirming the impact bias to the individual, rather than refuting it.

As noted by [Wirtz et al.](#), the consequence of this process for future choice is that choices will persist in being sub-optimal for the individual, who will continue to purchase needless emotional insurance, despite experience which should teach the individual of this very fact. The combination of the memory biases and learning processes seeks to reinforce mistakes in choice and not expose them.

1.2.6 *Merging These Ideas*

The above literature provides evidence that agents use anticipatory regret in order to make decisions, but have an experience of regret which differs from their prediction. They also use information from their past experiences to help guide future choice, but, in the case of regret, their memories are often unrepresentative of their true experience, and so fail to learn that they make errors in their affective forecasting. However, these errors in prediction and memory, and the type of learning they seek to engage in, appear to have specific patterns which, in the language of economics, can be translated into functional forms and mathematical representations. This enables us to revise traditional models of Regret Aversion to better reflect how an agent could use anticipated regret in a dynamic context, given this new information about how the emotion of regret behaves. The rest of this work will offer suggestions and models of how this can be achieved, and the predictions of behaviour which arise from those models.

A Real World Example

The evidence provided so far suggests that the existing theoretical literature does not sufficiently represent all that is known about the process of how regret can influence behaviour. However it is yet to be demonstrated that these failings actually hinder any kind of real world application, or hinder our understanding about real world empirical data. Indeed, if it is the case that regrets are context specific (so that the regret experienced from one type of decision is never used to inform about regret in any other type of decision) then a static model (as exists currently) may be sufficient, as the feedback mechanism proposed previously is broken down. As such, I will limit applications of this type of approach to *repeated decisions* where an individual is engaged in a repeated choice problem, and hence the context of the decision remains the same, as this is the simplest case for where this type of dynamic approach would be relevant.

There are, however, many different situations where an individual is engaged in a repeated, uncertain decision making process. An investment trader, for example, must make decisions on a very frequent basis about which stocks and commodities to buy and sell. A gambler must decide whether to keep playing with the money they have just won, or keep playing to recoup their losses. One could even make the argument that the most successful traders or gamblers are the ones who



Figure 8: The Main Road



Figure 9: The Side Road

are able to take the “emotion” out of their decision making process, so considerations of regret (and any unnecessary purchasing of emotional insurance) are negated.

However, as an example which can demonstrate the concept in the simplest fashion, I have chosen a problem more familiar to me and perhaps other people who commute on a daily basis.

My girlfriend lives in Birmingham, and I will drive to see her, usually arriving after work around 7pm. She lives down a side road (SR), which has limited parking at that time, but the side road is a short walk from a main road (MR), which typically has lots of parking at that time. When approaching I am faced with decision problem over where to try to park.

I can park on the main road, as in Figure 8, and walk to the house, which takes about 3 minutes. Or, I can turn right down the side road, as shown in Figure 9, and try to find a parking space down there.

| | Space only on MR $p = 0.7$ | Space only on SR $p = 0.03$ | Space on both $p = 0.25$ | No space on either $p = 0.02$ |
|------------|-------------------------------|--------------------------------|-----------------------------|----------------------------------|
| Stay on MR | -3 | -13 | -3 | -13 |
| Go down SR | -6 | 0 | 0 | -16 |
| Park in CP | -10 | -10 | -10 | -10 |

Table 1: Parking Example

The problem is that I cannot see whether there is a space down the side road before needing to make a decision about whether to turn down the side road, or stay and park on the main road. Indeed, if I do turn down the side road and there is no space to park, I am forced to follow a very long, winding road to simply get back to the main road, which takes an additional 3 minutes.

Occasionally, there will be no space, either on the main road or the side road. In which case, I need to park in car park (CP) which is a 10 minute walk from the house, and a 3 minute drive from the main road.

As such, the decision problem I face can be reduced to the payoff matrix in Table 1, where payoffs represent the total time lost getting to the house from the decision.

By attaching probability weights to the states of the world (given in columns) there are several ways to calculate the optimal action. Suppose the probabilities are given as per the payoff matrix. Then maximising the expected value says I should stay on the main road, as I would if I was risk averse. However, if I was loss averse I might make the decision to go down the side road. Alternatively I could progress as suggested by Hayashi [34], computing the *probability weighted expected regret* of every action³⁷.

So if these models can explain decision making in this situation, why is a *dynamic* anticipated regret aversion model needed? Well, this is a repeated decision problem, as I visit my girlfriend about twice a week. Yet sometimes I will park on the main road, and sometimes I will go down the side road. The payoff and regret matrices are not changing, yet, given my choice is varying, it suggests the optimal action is also changing.

At the most basic level, it may simply be the case that I am indifferent between all three options. In the case of indifference, then any decision rule used to make the choice is utility maximising and hence optimal. As such, assuming indifference would not allow us to compare the predictive capability of one decision rule versus another, so I will assume complete indifference across the actions does not exist here.

One possibility is that I am subject to some kind of “noisy” or “fuzzy” preferences³⁸ which could explain this type of switching behaviour. However, the existence of something akin to an error term in the decision rule could be tested by looking at the correlation of the error terms across time periods. Uncorrelated error terms might indicate that preferences are simply noisy or fuzzy, but correlated error terms indicates that there is a missing piece to the puzzle. This is, however, a theoretical example rather than an empirical one, so I will assume that there is some information conveyed by the choice in the previous period which influence the decision taken in the subsequent one.

³⁷ the mathematics of which are detailed later on page 34

³⁸ for example, as given in Butler and Loomes [15]

Under this assumption, another explanation is that the previous experience of the decision problem is providing additional information about the payoff matrix and associated probabilities. For example, if I start without knowing the payoffs exactly, or subject to some degree of error, I will need to experience each consequence at least once to accurately complete the payoff matrix. Similarly, if the probabilities are unknown, I will take frequency information, about the state of the world, from each experience to update beliefs about probabilities. But both of these explanations suggest some kind of convergence towards stable long-term decision making behaviour.

In contrast, if I use a regret-minimising framework, in the spirit of Hayashi, where my beliefs about anticipated regret come from my memory about past regretful experiences, then I can have varying preferences, which may not converge towards stable behaviour, if the memory is imperfect or biased in any way. This is the idea which will be explored in the next section, with specific forms of the model being derived from the psychological motivation provided earlier.

1.3 MODEL

1.3.1 The Static Regret Model

As presented in the introduction and motivation, there have been several recent models of regret-based decision making which adopt a different approach to the traditional methods of Loomes and Sugden [58] and Bell [5]. The model of Hayashi [34], as previously discussed in 1.2.1, uses an axiomatic approach to derive a model of “regret minimisation” which will be used by an agent obeying those axioms when faced with a decision under uncertainty.

The most widely used, and arguably most significant, result of this paper is the first application of the framework to “Minmax regret with multiple-priors” [34, p244], which is a generalisation of Savage’s “Minmax Regret Choice” model [84]. Hayashi’s approach describes “...a general model in which regret aversion and likelihood judgement over states coexist” in contrast to Savage’s “...model of complete ignorance” [34, p243] of the likelihood of states.

However, it is the second result of Hayashi’s paper which is of more interest to this particular research. By replacing one axiom from the first result with a slight modification³⁹ and relaxing an assumption⁴⁰ that agents are necessarily regret averse⁴¹, he derives a model of choice, by regret minimisation, which is a “Smooth model of regret aversion” [34, p253, Theorem 3]. Specifically,

Theorem 3. Assume $|\Omega| \geq 3$. The choice function φ satisfies Axioms 1, 2, 3-5, 8 if and only if there exists a mixture-linear, continuous and non-constant function $u : \Delta(X) \rightarrow \mathbb{R}$, a probability measure $p \in \text{int } \Delta(Q)$ and a number $\alpha > 0$ such that

$$\varphi(B) = \arg \min_{f \in B} \sum_{\omega \in \Omega} \left(\max_{g \in B} u(g(\omega)) - u(f(\omega)) \right)^\alpha p(\omega)$$

([34, p251])

³⁹ formally, replacing Axiom 7 (Constant-regret independence of regret premium) with Axiom 8 (Eventwise separability of regret premium)

⁴⁰ Axiom 6

⁴¹ in the Loomes and Sugden sense

Here, Ω represents the set of states of the world, X the set of pure outcomes, and $\Delta(X)$ the set of lottery outcomes over X . f and g are actions available to the decision making agent from the set of possible actions B , and $p(\omega)$ is the probability of any particular state of the world ω occurring. u is a standard utility function, and, as such, $u(g(\omega))$ represents the utility the agent will receive from choosing action g if state of the world ω occurs. Therefore, Theorem 3 states that the agent will choose an action from B as follows:

1. For a given action f , calculate the maximum regret you could suffer if state of the world ω occurred. This is given by the maximum difference in utility between f and any other action g in state of the world ω . This difference in utility, raised to the power of α , represents the regret you would suffer if you chose act f and state of the world ω occurred.
2. Repeat stage 1 for all states of the world world $\omega \in \Omega$. This gives the regret which occurs in each state of the world, should you choose action f ⁴².
3. Calculate the “expected regret” of action f by weighting each of the state specific regrets, calculated in stage 2, by the probability that the state of the world occurs, $p(\omega)$, and summing them up. Doing this gives the agent the regret they can expect to feel, on average, from choosing action f .
4. Repeat stages 1-3 for every action which the agent can take ($\in B$). This gives the “expected regret” for every action available.
5. Choose the action which offers the “minimum expected regret” from stage 4. This is then the “regret minimising action” for the agent.

This differs subtly from the [Loomes and Sugden](#) framework as, instead of maximising a modified expected utility function, with the modification being a reduction in utility from regret, from which we can observe actions and infer the associated degree of regret aversion, we are, instead, minimising an expected regret function, with the degree of regret aversion being explicitly captured by α , and forming an integral part of the decision making process.

Indeed, α works in this model as one would expect from the original definition of regret aversion. If $\alpha > 1$, then the individual is “regret averse” (and has a convex regret function) in the sense that large regrets weigh heavily in the mind of the agent compared to small regrets, and so the agent will tend to avoid choosing actions which could result in large regrets. This equates to Assumption 3 of [Loomes and Sugden](#), and hence is consistent with the majority of observed empirical and experimental violations of expected utility theory. If $\alpha < 1$, however, the agent is now “regret loving” (and has a concave regret function), which corresponds to Assumption 2 of [Loomes and Sugden](#). This would also predict violations of expected utility theory “...but in the opposite direction to those generally observed” [58, p810]. Furthermore, as [Hayashi](#) notes “... $\alpha = 1$ corresponds to regret neutrality, which is the case of subjective utility maximisation” [34, p244]. Knowing, therefore, a utility function, complete payoff matrix, associated state probabilities and regret aversion parameter, an agent can always calculate the expected

⁴² bearing in mind that this regret will be zero if f is the optimal action to take for a given state of the world, and will always be ≥ 0

| <i>Regret Matrix</i> | Space only on MR p = 0.7 | Space only on SR p = 0.03 | Space on both p = 0.25 | No space on either p = 0.02 |
|--------------------------|-----------------------------|------------------------------|---------------------------|--------------------------------|
| Stay on MR | 0 | 13 | 3 | 3 |
| Go down SR | 3 | 0 | 0 | 6 |
| Park in CP | 7 | 10 | 10 | 0 |

Table 2: Regret Matrix for the parking example

regret minimising action (or actions) and will choose this action if [Hayashi](#) axioms are satisfied.

As a specific numerical example of this procedure, consider the parking example presented previously on Table 1. In this example, there are three actions (with state-dependent utilities) and four states of the world (with associated probabilities). By following stages 1 and 2 above, we can calculate the *regret matrix* which details, for every action and state of the world, how much better off you could have been had you chosen the optimal action for that state of the world.

The regret matrix in Table 2 tells me how many minutes I could have saved had I chosen the optimal action for that state of the world. For simplicity, this measure of time is assumed to be the utility.

To calculate the “expected regret” however, we must first make an assumption about α , or the parameter of regret aversion. This parameter tells me how severe is the regret that I anticipate I will experience for a given utility difference, or, in other words, how sensitive I am to potential regrets.

By assuming that $\alpha = 1$, for example, we impose regret neutrality. Doing so gives the following “expected regrets” for each action

- $ER(\text{Stay on MR}) = 0.7 \times 0^\alpha + 0.03 \times 13^\alpha + 0.25 \times 3^\alpha + 0.02 \times 3^\alpha = 0.03 \times 13 + 0.25 \times 3 + 0.02 \times 3 = 1.2$
- $ER(\text{Go down SR}) = 0.7 \times 3^\alpha + 0.03 \times 0^\alpha + 0.25 \times 0^\alpha + 0.02 \times 6^\alpha = 0.7 \times 3 + 0.02 \times 6 = 2.22$
- $ER(\text{Park in CP}) = 0.7 \times 7^\alpha + 0.03 \times 10^\alpha + 0.25 \times 10^\alpha + 0.02 \times 0^\alpha = 0.7 \times 7 + 0.03 \times 10 + 0.25 \times 10 = 7.7$

Hence, the action which yields the minimum expected regret is “Stay on the main road”, which, as we assume $\alpha = 1$, is equivalent to the action chosen by maximising expected utility.

However, if we consider the same problem for an individual who is “regret averse” (in keeping with most experimental evidence), and, for example, assume a parameter of regret aversion $\alpha = 2$, then the expected regrets for each action are as follows

- $ER(\text{Stay on MR}) = 0.7 \times 0^\alpha + 0.03 \times 13^\alpha + 0.25 \times 3^\alpha + 0.02 \times 3^\alpha = 0.03 \times 169 + 0.25 \times 9 + 0.02 \times 9 = 7.32$
- $ER(\text{Go down SR}) = 0.7 \times 3^\alpha + 0.03 \times 0^\alpha + 0.25 \times 0^\alpha + 0.02 \times 6^\alpha = 0.7 \times 9 + 0.02 \times 36 = 7.02$
- $ER(\text{Park in CP}) = 0.7 \times 7^\alpha + 0.03 \times 10^\alpha + 0.25 \times 10^\alpha + 0.02 \times 0^\alpha = 0.7 \times 49 + 0.03 \times 100 + 0.25 \times 100 = 62.3$

As such, for this regret minimising agent, the regret minimising action is to “Go down side road”. As the agent is now more sensitive to large

regrets, the largest potential regret, (of staying on the main road, where there is no space, hence parking in the car park, only to realise there was space in the side road all along), weighs very heavily in the agent's decision⁴³. However, if α remains fixed, then if this problem is repeated over and over again (as described in the motivating example on Page 29) then the agent will choose the same action, whatever that may be, over and over again, under the regret minimising framework. To develop a model where an agent can make different decisions in each stage of the repeated problem, we must introduce a feedback loop from the past to the present which allows one, or more, of the parameters in the model to vary.

1.3.2 The Dynamic Regret Model

As discussed by Loomes and Sugden, a regret-based model, which either incorporates regret loving, or, more usually, regret aversion, can predict and be used to explain observed violations of Expected Utility Theory. Indeed, Hayashi model does this for values of α either greater than or less than one. Hence, applying the static model, presented above, will not provide much new insight, even when applied to much wider contexts, over and above Loomes and Sugden original observation, when considering patterns of behaviour and decision making.

To make best use of the model, therefore, we must look beyond the static decision making framework to a dynamic one. Assuming that the payoff matrix and associated state probabilities are exogenously given to the individual, this leaves two possible options for adapting the static model into a dynamic one.

Firstly, we can let the utility function of the agent, $u(\cdot)$, vary over time; specifically as a result of previous decisions and consequences. This approach brings us into the realm of endogenously determined tastes and preferences, the implications of which are discussed by Houthakker and Taylor [40] with regards to habit formation, by Pollak [69] with regards to welfare analysis, and by Hammond [29] with regards to long-run choice behaviour. As this approach has been extensively researched, without the need to introduce regret as a additional component, unless there is a reason and method to link experienced regret (which Hayashi model creates) to an endogenously determined utility function⁴⁴, there is little to be gained from using this approach with Hayashi framework.

The second approach is to instead think of the parameter of regret aversion, α , in Hayashi's model, as being endogenously determined. This approach has a number of advantages, and is much more in keeping with existing psychological research which was presented in the motivation.

Conceptually, the idea is that in each period, t , the decision making agent determines how regretful large "utility gaps" are, compared

⁴³ of interest with this result is that the regret averse agent (when $\alpha = 2$) chooses a different action than would a risk averse agent (who would choose to stay on the main road), indicating that, though linked concepts, theories of regret and risk aversion can lead to very different predictions of behaviour. For example, a risk averse individual is unlikely to play the lottery, but a regret averse individual fears seeing "their numbers" come up when they haven't bought a ticket, and so continues to play.

⁴⁴ a plausible feedback link between experienced regret and the utility function involves a "psychological hangover" effect whereby the experience of a highly regretful decision reduces the utility obtained from any payoff in subsequent periods. However, supposing, for instance that the utility of any payoff is reduced by a fixed amount, this will not change the outcome of Hayashi model as the constant reduction term will drop out.

| | x | y |
|---|----|-----|
| A | £4 | £10 |
| B | £6 | £1 |

Table 3: Monetary Example

to small “utility gaps”. Using this terminology, a “utility gap” is the difference in utility between a outcome you obtained, and an outcome you could have obtained, had you chosen differently. Suppose an agent is faced with two actions, A and B, and two states of the world x and y, and receives monetary payoffs as given in Table 3.

In this example the utility gap between actions A and B, in state of the world x, is $|u(\$4) - u(\$6)|$ (a small utility gap), and, in state of the world y is $|u(\$10) - u(\$1)|$ (a large utility gap). The agent, however, does not experience the utility gap, per se, but rather the regret that this utility gap generates.

Suppose, therefore, that the experienced regret that the utility gap in state x generates⁴⁵ is R_1 , and the experienced regret that the utility gap in state y generates⁴⁶ is R_2 . Then if R_1 and R_2 are approximately similar, or close in magnitude, to the agent, then the regret obtained from a large utility gap is similar to the regret obtained from a small utility gap. That is to say, the size of the utility foregone does not seem to translate to the level of regret experienced. However, if R_2 is significantly larger than R_1 , then the size of the utility gap is positively impacting the level of regret experienced. In the terminology of Loomes and Sugden, this second case displays more regret aversion than the first and, in the terminology of Hayashi, this translates to a larger value for α .

Hence, in order to determine a value for α that would represent the degree of anticipated regret aversion which an agent would use in the Hayashi model, we need to find a way of estimating how “bad” large regrets loom on the agent when compared to small regrets, at a given time period t. The static model assumes this is a preference parameter for the individual; in effect saying the individual knows their own affective response to a utility gap before the outcome of the decision itself. However, the evidence presented in the motivation, from affective forecasting, learning models and neuroscience, suggest that an individual will look to information in the past, where available, in order to better estimate such a parameter for use in decision making. In short, the degree to which you anticipate large regrets will outweigh small regrets is determined, at least in part, by the degree to which large regrets outweighed small regrets in the past.

Modelling Memory

Returning to the monetary example in Table 3, for an individual who is using past experience of regrets to guide future anticipated regrets, if they wanted to know the regret which would be experienced from a small utility gap of $|u(\$4) - u(\$6)|$ and a large utility gap of

⁴⁵ that is, the regret from choosing action A, observing state of the world x, realising your payoff is £4, and additionally realising that you would have obtained £6 had you chosen action B

⁴⁶ that is, the regret from choosing action B, observing state of the world y, realising your payoff is £1, and additionally realising that you would have obtained £10 had you chosen action A

$|u(\$10) - u(\$1)|$ the ideal, “best-case” memory to use would be the remembered regret from previous utility gaps of exactly this size. That is to say, the individual has experienced these exact situations before, and has a good memory of the experienced regrets which resulted. This situation is, however, not particularly realistic, and hence interesting, for a few reasons.

Firstly, it requires the individual to remember the precise context in which the regret was experienced (i.e. the precise cause) in addition to the experienced regret, and not simply the experienced regret itself. This is similar to expecting an individual to be able to give a precise account of why they were unhappy on a particular day last week, when it is more likely that they will simply remember by associating the time period (the day) with a particular affective state (“I was unhappy on Tuesday last week”).

Secondly, it requires the individual to experience precisely the same choice problem more than once. Though the initial motivating example on page 29 was framed in this fashion, it is more realistic that there will be at least some elements of randomness in a “real-world” repeated problem. For example, when considering which road to drive down when parking at my girlfriend’s house, the exact time spent walking to the house will not only depend on my choice of road, but where exactly a parking space is available on the road, and hence the payoffs in the decision matrix may be distorted by this element of randomness. In such cases, where an individual wants to appeal to the “similarity” of past experience in order to make future decisions under uncertainty, the theory of “Case-Based Decisions” (Gilboa and Schmeidler [27]) already provides a framework to analyse these problems, so it will not be explored further here at this time.

Thirdly, even if it were the case that the dynamic problem of interest was a repeated decision, where the payoffs and/or actions were not distorted by randomness, then learning about the experienced regret that results from each action and each state of the world would simply require each outcome in the payoff matrix to be experienced once, and we are in no different a context to when it was the payoffs themselves which needed to be experienced once to be discovered. The regret is simply an extra component of the utility function which is learnt through experience. The individual is not, however, learning about their own regret reaction to utility gaps, and hence learning about α , but rather learning about their own reaction to very specific situations. They would, for example, not be able to apply what they have learned if one of the payoffs in the matrix was changed, as they would if they were learning about α instead.

The three points made above, in conjunction with the initial motivation about memory and reinforcement learning, suggest that we need to consider a model of memory whereby it is just the magnitude of the regret (and the time period in which that magnitude was experienced) which is remembered, and not the context which caused the regret to arise. This gives the individual, engaged in a repeated decision making problem, a *memory stock* of regrets at each time period where a decision was taken, realised, and experienced. We can then

make assumptions about that memory stock which reflect the patterns of memory of emotions analysed in the motivation⁴⁷.

Modelling Learning and Feedback

As described above, what is required from a learning and feedback model is a method for transforming the memory stock of regrets⁴⁸ into the anticipated parameter of regret aversion, α , for use in the decision problem at time period t . As discussed on page 32, the coefficient of regret aversion α should be

- < 1 when it is anticipated that large utility gaps will not yield much more regret than small utility gaps (which can be thought of as *regret loving* or *regret seeking* behaviour, even though the individual does not “seek” or “love” regret using the normal definitions of the words, or, alternatively, that the regret function is concave)
- $= 1$ when the individual is making a decision under the assumption of regret neutrality (and so the individual has a linear regret function) so large regrets do not loom *disproportionately larger* than small regrets to the individual
- > 1 when the individual is regret averse⁴⁹ (and so the individual has a convex regret function) so large regrets *do* loom *disproportionately larger* than small regrets to the individual

Hence, in keeping with the literature on reinforcement learning, as described on page 20, the memory stock of the individual should reinforce the belief that they are

- Regret Loving (and hence give them a value of $\alpha < 1$) if and only if they have a memory stock where large regrets are not that much more severe than small regrets
- Regret Neutral (and hence give them a value of $\alpha = 1$) if and only if they have a memory stock where large regrets are directly in proportion to small regrets
- Regret Averse (and hence give them a value of $\alpha > 1$) if and only if they have a memory stock where large regrets loom disproportionately larger than small regrets.

This then becomes a reinforcement learning type procedure, because each new regret experienced (after a decision has been made, resolved and experienced) will add to the memory stock and *reinforce* the belief of the individual, either in the direction of becoming more regret averse or more regret loving, depending on what new information it provides to the individual about how the individual experiences and feels large regrets compared to small regrets.

This procedure creates a feedback loop, in keeping with the original ideas of Baumeister et al., as displayed in Figure 3, whereby the affective residue of the decisions at $t - 1$, $t - 2$, $t - 3$ etc., help inform and

⁴⁷ for example, a individual whose memory fades over time (a reasonable assumptions) can have a discount factor applied to their memory stock so that more recent regrets loom larger than older regrets.

⁴⁸ that is, the remembered magnitude of regret experienced at time $t - 1$, $t - 2$, $t - 3$, etc.
⁴⁹ as experimental evidence suggests most people are

direct the decision making of the individual at time t . In addition, it is a learning procedure whereby the individual is seeking to use information about past experienced regret, not simply to learn about their emotional response to a very specific situation⁵⁰, but to learn about their entire “taste” for regret, as given by α , so that they can make better decisions in the future.

As such, what is needed to create a predictive, dynamic mathematical model, is a function which can map the memory stock of regrets to α in such a way that it has the three properties of Regret Loving, Neutrality and Aversion described above.

Mapping M to α

The prior intuition indicates that, at each time period t , the memory stock of the individual, M_t , will consist of the magnitude of regrets experienced at times $t - 1$, $t - 2$, $t - 3$ etc., where the individual can remember at which time period each regret was experienced⁵¹, but cannot remember the precise utility gap which led to the regret⁵². Thus, the individual can be thought of as “boundedly rational”, as they are attempting to infer helpful information from their past, in order to guide future behaviour, but there are cognitive limitations, beyond the control of the individual, on what information is stored in their memory. As such, there are a couple of ways of graphically visualising this information.

Firstly, the regrets can be thought of as simply ordered by time. This approach is useful when wishing to apply time specific transformations to the memory stock to reflect biases or limitations. For instance, consider a time-ordered memory stock of regrets (given for 30 time periods) as given in Figure 10.

In this situation, we may wish to impose a “fading affect” transformation, whereby regrets fresh in the mind (at $t = 28, 29, 30$), still have their full affect, but those which are further back are subject to fading affect. For instance, if affect was to fade by 5% each period, then the transformed memory stock at period 31 would be represented as in Figure 11.

This representation of the memory stock, however, is not particularly helpful when estimating a value for α as it compares “new” to “old” regrets, and not “large” to “small”, as is required when discussing degrees of regret aversion. As such, a more helpful representation is to order the memory stock by magnitude of regret, not time, so that small regrets appear first, and large regrets appear last. Doing so with the memory stock example above gives the graph in Figure 12.

⁵⁰ A related literature is that of Case-Based Decision Theory by Gilboa and Schmeidler [27]. Whereas many of the concepts used in that literature are similar to the ideas presented here (using rules to take information from the past to inform decision making in the future, for example), the notable difference is the move away from specific “cases”. The Gilboa and Schmeidler approach is to look back into the memory stock for a “similar” case to the current decision being faced, and use the similarity to guide the future decision making. This approach of linking specific decisions to specific cases from the memory stock contrasts with the approach of Baumeister et al., whereby it is the affective residue of the memory stock which guides future decision making behaviour. Whilst both approaches are equally meritorious, simply recreating Gilboa and Schmeidler’s work, but using the concept of cases to modify Regret Theory, as opposed to Expected Utility Theory, doesn’t appear to add much value to the existing literature on the subject. This belief guides the choice to work with the approach of Baumeister et al. here, helping to bring an element of economic modelling to a previously under-developed area.

⁵¹ hence knows the ordering of the regrets

⁵² and hence does not know the cause of the regret




Figure 10: A randomly generated memory stock of regrets, ordered by time




Figure 11: Memory stock of regrets with fading affect




Figure 12: Memory stock of regrets ordered by magnitude

If the regrets contained in the memory stock at time $t = 31$ are ordered by time to give the vector $M_{31} = (R_1, R_2, \dots, R_{29}, R_{30})$ where R_i is the memory of regret experienced at time i , but retrieved at time $t = 31$, then we can likewise create a vector of regrets, retrieved at time $t = 31$, but ordered by magnitude

If $M_{31} = (R_1, R_2, \dots, R_{29}, R_{30})$

then $M_{31}^O = (R_1^O, R_2^O, \dots, R_{29}^O, R_{30}^O)$

where $R_1^O = \min \{R_1, R_2, \dots, R_{29}, R_{30}\}$

and $R_{i+1}^O = \min \{R_1, R_2, \dots, R_{29}, R_{30}\} \setminus \{R_i^O, R_{i-1}^O, \dots, R_1^O\}$

Hence, what is needed is a function which maps

$f : \mathbb{R}^{30} \rightarrow \mathbb{R}^+$

so $f(M_{31}^O) = \alpha_{31}$, where α_{31} is the coefficient of regret aversion, based upon the memory of regrets experienced in the preceding 30 periods, to be used in the Hayashi decision making framework at time $t = 31$.

In Hayashi's theoretical model, α transforms utility gaps into experienced regrets, so the α_{31} generated by f must somehow relate utility gaps to experienced regrets, despite the fact that the memory stock M_{31} does not store information on causality, and hence utility gaps. As such, the individual must make an assumption about utility gaps, from their memory of experienced regrets, and the assumption I suggest is that the difference in utility gap which caused regrets R_i^O and R_{i+1}^O is independent of i . In essence, this assumption is saying that the increase in regrets in the ordered memory stock were the result linearly increasing utility gaps. Hence,

- if the increase in experienced regret is proportionate through M_{31}^O , then large experienced regrets were only proportionately worse than small experienced regrets, and so the individual is "regret neutral"
- if the increase in experienced regret is increasing through M_{31}^O , then large experienced regrets were disproportionately worse than small experienced regrets, and so the individual is "regret averse"
- if the increase in experienced regret decreasing through M_{31}^O , then large experienced regrets are not that much worse than small experienced regrets, and so the individual is "regret loving"

Thinking graphically, this translates into a "line of best fit" which can be draw through the graph of ordered memory stock. If the line of best fit is linear, then experienced regret is increasing proportionately to the "assumed" increasing utility gaps, and so the individual is regret neutral. Equivalently, if the line of best fit is convex, the individual is regret averse, and if it is concave, the individual is regret loving. For example, fitting a power function to the ordered memory stock graph in Figure 11 yields a coefficient of 1.28, suggesting that this individual is slightly regret averse, and this is shown in Figure 13.

However, there are other types of functions which could equally be thought of as representing the degree of experienced regret aversion in the ordered memory stock, and can be thought of as more intuitive for a nondeclarative memory process⁵³. For instance, we could focus on the

⁵³ it seems unlikely that a subconscious learning-from-memory process, as described, would, in each time period, compute the equivalent of a least squares regression to find the implied coefficient of a power function on their ordered memory stock. In keeping with




Figure 13: Ordered memory stock of regrets with line of best fit

extremes of the memory stock, compared to the average. If the memory of the individual is regret neutral, then the difference between the largest and the median regret, and the difference between the median and the smallest regret, should be approximately equal (as regret would be linearly increasing over the ordered memory stock). Hence, taking the ratio of these two quantities gives an approximation for the degree to which the jump from median to largest regret outweighs the jump from smallest to median⁵⁴. If this number is greater than one, then, again, this implies regret aversion, and if less than one it implies regret loving.

A similar intuitive approach to the problem is found by calculating the sample skewness of the memory stock, which answers the question "are the bulk of the regrets in the memory large or small?" Again, this question is a more intuitive *feeling* that the individual could have, rather than explicit calculation as in the case of the estimated power function, and so is more appealing to be a candidate for the function which maps M to α . It would relate to the degree of regret aversion, as an individual who has a positively skewed memory stock is used to experiencing mostly small regrets, but would have the occasional large regret (in the long right tail) which stands out as significantly different to the norm. This is similar to the regret averse individual who fears the occasional large regret which looms in the large right tail, and so if the memory stock has positive sample skewness, this should translate to an implied value of $\alpha > 1$. On the other side, a negatively skewed memory stock would translate to an individual who is used to experiencing large regrets, and hence can be thought of as regret loving, translating to an

the literature on reinforcement learning, a simpler function which gives rise to a *feeling* about the memory stock, such as "the worst regret is disproportionately bad" or "most of the regrets are below the average" seems more plausible.

⁵⁴ this can be slightly modified to move from "largest" to "larger" regrets (so not completely dependent on the extremes of the ordered memory stock) by replacing the median with the mean in the above calculation, though this comes at the expense of adding a notion of an "average" to the memory stock, which may not necessarily be one of the regrets actually experienced.

implied value of $\alpha < 1$. The obvious problem here is that the sample skewness, calculated from the memory stock, does not only map to the positive reals (as required for the function f), and does not give a value of $\alpha = 1$ when the distribution is symmetric. As such an additional transformation will need to be included if we are to use this estimate of α (as discussed later on in 1.4.3).

These different functions have slightly differing psychological interpretations and limitations, and hence all could justifiably be used in the dynamic model to map M to α . In setting up the dynamic model which will be used to study behaviour, it is possible to use any of them, but if they all have approximately the same impact in the model (i.e. the model is robust to changes in the precise map from M to α), then it will be only down to personal preference over the psychological justification as to which one should be used. As such, the results section (specifically 1.4.3) will include a discussion on the consequence of using each of the three maps.

Completing the Feedback Loop

Once the map from M to α is specified, it is then simply a case of using α in the *Hayashi* framework at time t , resolving uncertainty to calculate the experienced regret (subject to any impact bias transformations) resulting from the decision made, encoding the new regret to the memory stock, applying any time specific transformations to the memory stock as we move from period t to $t + 1$, then reapplying the function f to the memory stock to calculate the new α to be used at time period $t + 1$. The specific decision problem we will analyse using this approach is outlined in the next section.

1.4 SIMULATION

1.4.1 *Aims and Purpose*

The aim of creating such a model of dynamic regret aversion is twofold. Firstly, it is intended as a theoretical model of how existing models of regret aversion can be extended to better portray the true behaviour of the emotion (regret) which we wish to incorporate in models of decision making under uncertainty. It does this by using existing psychological and neuroscience research to, in essence, make the emotion of regret behave more like an emotion in the models.

Secondly, having shown above how it is theoretically possible to develop such a model, we would like to explore the welfare, actions chosen, and patterns of behaviour exhibited by such an individual who is making decisions under risk using the dynamic regret aversion framework. In addition, we would like to explore the extent to which those observed patterns of behaviour are robust to changes in the various functional form specifications of the individual components of the feedback loop, and which components of the feedback loop are critical for various “non-standard” patterns of behaviour.

The use of simulation techniques

The specific methodology for answering the above questions is detailed in subsequent sections, but, in short, it relies on the use of a Monte Carlo simulation approach to analyse patterns of behaviour, rather than

analytically solving a system of dynamic equations to demonstrate a precise relationship, or convergent pattern, for behaviour over time, as a function of the various components of the feedback loop. Though obviously an analytical solution would be preferred, and the first best option, the degrees of freedom in the model, and the complexity of the interlinking parts of the feedback loop, prevent this from being a realistic proposition. Hence simulation methods are the only suitable approach for a problem of this order.

1.4.2 Framework

Developing on the Static Regret Model

In the simulation, a decision maker will be faced with a repeated decision under uncertainty, where the specific payoffs associated with each action and state of the world will vary in each period, but maintain a constant structure (outlined in the next section). This represents a decision maker who has to make the same type, or category, of decision on a regular basis, but where the exact problem being faced in any given period is never the same. In each period, there are several stages of the learning and feedback processes to construct, before the decision maker can use the Hayashi choice rule. Thus, the decision making process works as follows

1. The memory stock from the end of the last period is transformed by any “storage” processes and biases⁵⁵.
2. A value of α , the degree of anticipated regret aversion, is constructed from the memory stock generated in stage 1, according to the specific function f being employed.
3. The decision maker chooses an action, given the complete payoff matrix, according to the choice rule of Hayashi⁵⁶, $\varphi(B) = \arg \min_{f \in B} \sum_{\omega \in \Omega} \left(\max_{g \in B} u(g(\omega)) - u(f(\omega)) \right)^\alpha p(\omega)$, where the value of α is given by stage 2.
4. Uncertainty is resolved, and the decision maker obtains utility according to the action chosen in stage 3 and the state of the world which was experienced.
5. The decision maker also experiences regret according to the function $R^e = \left(\max_{g \in B} u(g(\omega)) - u(c(\omega)) \right)^{\alpha^e}$, where ω is the state of the world experienced, c the action chosen in stage 3, and α^e is the degree of *experienced* regret aversion, as implied by the degree of *anticipated* regret aversion from stage 2, but subject to any “impact” transformations⁵⁷.
6. R^e is added to the memory stock, subject to any “encoding” transformations⁵⁸.

⁵⁵ for example, fading affect

⁵⁶ as in 1.3.1

⁵⁷ for example, an impact bias

⁵⁸ for instance, it may make sense that “zero” regrets, i.e. where the ex-post optimal action was the one chosen, should not be added to the memory stock, as there was no experience of regret to remember

| Payoff Matrix | ω_1 | ω_2 | ω_3 |
|----------------|--------------------|--------------------|--------------------|
| P-Bet | β | β | 0 |
| Full Insurance | $\frac{2\beta}{3}$ | $\frac{2\beta}{3}$ | $\frac{2\beta}{3}$ |
| \$-Bet | 0 | 0 | 2β |

Table 4: Simulation Payoff Matrix

This procedure then “loops” back to the start, and is replicated in every period of the repeated decision problem.

P-Bet, \$-Bet and Insurance

The above procedure could be applied to any repeated decision problem under uncertainty where the complete payoff matrix, and associated probabilities, are specified at each stage. However, the problem I will analyse using this approach builds on the example (on page 29) discussed previously.

In the example of which road to try and park down, there was one “risky” option, going down the side road, one “safer” option, staying on the main road, and one “full insurance” option, of parking in the car park. These three options translate to three commonly used options for participants in behavioural economics experiments; the P-Bet (a high probability, low reward “safer” option), the \$-bet (a low probability, high reward “risky” option), and “full insurance” (where every state of the world will yield the same payoff if the insurance is taken). Comparing the preferences of people for P-Bets compared to \$-Bets has been used extensively in experimental economics when studying decision under uncertainty⁵⁹, and hence it is appropriate to use this framework here when considering a simulation representative of a single decision making agent who is faced with multiple rounds of a decision problem. A “full insurance” option is also added, as, in this simulation, the decision making agent is forced to participate in the problem in each period, and does not have the choice whether to opt in or opt out as they, most likely, would in a real world scenario. Hence the “full insurance” option represents a decision to not participate in any risky behaviour, though this does not exempt them from experiences of regret, as they still observe the resolution of uncertainty, regardless of the fact that their outcome will no longer depend on the state of the world⁶⁰.

Thinking in such terms, the payoff matrix for the decision making agent can be thought of (under certain parameter restrictions) as in Table 4.

For example, if $p(\omega_1) = p(\omega_2) = p(\omega_3) = 1/3$, so all states of the world, ω_i , are equiprobable, and $\beta > 0$, then each action has the same expected value⁶¹, but the P-Bet offers a medium reward with high probability (β with probability $2/3$) compared to the \$-Bet which offers a high reward with low probability (2β with probability $1/3$)⁶².

⁵⁹ for example, Lichtenstein and Slovic [56] is one of the earliest uses of the technique

⁶⁰ by saying they observe they resolution of uncertainty, this implies that although they can decide not to participate in the risky gamble, they cannot bury their head in the sand at the same time.

⁶¹ $EV = 2\beta/3$

⁶² typically in experiments a \$-Bet will have a higher expected value than a P-Bet, as, if they were equal, any typical risk-averse person would always take the P-Bet. However,

| Regret Matrix | ω_1 | ω_2 | ω_3 |
|----------------|---------------------------------------|---------------------------------------|--|
| P-Bet | 0 | 0 | $(2\beta)^\alpha$ |
| Full Insurance | $\left(\frac{\beta}{3}\right)^\alpha$ | $\left(\frac{\beta}{3}\right)^\alpha$ | $\left(\frac{4\beta}{3}\right)^\alpha$ |
| \$-Bet | β^α | β^α | 0 |

Table 5: Simulation Regret Matrix

From this position, and by assuming, for simplicity, the utility of each outcome is simply given by the payoff in the matrix (so $u(\beta) = \beta$), it is possible to calculate the Regret Matrix, by calculating, using the [Hayashi](#) procedure, by how much better off the person would have been had they chosen the optimal action for a given state of the world, compared to the action they actually chose⁶³. This is shown in Table 5.

Again, following through with the [Hayashi](#) procedure, we can calculate the Expected Regret of each action as follows (again assuming $p(\omega_1) = p(\omega_2) = p(\omega_3) = 1/3$)

- $ER(\text{P-Bet}) = \left(\frac{2^\alpha}{3}\right) \beta^\alpha$
- $ER(\text{Full Insurance}) = \left(\frac{2+4^\alpha}{3}\right) \left(\frac{\beta}{3}\right)^\alpha$
- $ER(\text{\$-Bet}) = \frac{2}{3} \beta^\alpha$

In this case, the “expected regret minimising” action will depend on the coefficient of regret aversion α , but not on the parameter β . As such, plotting the expected regret functions as a function of α will graphically show the intervals of α in which each action will be chosen. For example, for $\beta = 1$, Figure 14 shows the expected regret associated with each action.

Results under the Simple Static Regret Model


Figure 14 shows that:

- For $\alpha < 1$ the P-Bet is the regret minimising action.
- For $\alpha = 1$ all three actions give the same expected regret⁶⁴ and hence the individual is indifferent.
- For $1 < \alpha < 2$ full insurance is the regret minimising action.
- For $\alpha = 2$ the individual is indifferent between the \$-Bet and full insurance.
- For $\alpha > 2$ the \$-Bet is the regret minimising action.

throughout this work, we wish to make the distinction between risk and regret aversion, so no such assumption of risk aversion is needed or made, hence the expected values are the same.

⁶³ compared to the Regret Matrix given in the initial motivating example on page 29, the parameter of regret aversion α has now been included in the matrix in Table 5, to specify that these are no longer merely utility gaps, but experienced regrets.

⁶⁴ When $\alpha = 1$, the individual is regret neutral, and hence will choose the same action as they would under maximisation of expected value, as described by Hayashi [34, p244]. As all three actions have the same expected value by construction, the individual is indifferent between all three actions when $\alpha = 1$.

Figure 14: Expected Regret graph for $\beta = 1$

These results of the static model highlight several points. Firstly, the distinction between risk aversion and regret aversion is noticed in even this most simple of contexts. Regardless of the size of β , by taking a regret approach as opposed to a risk approach (utility is linear in payoffs), the most regret averse person (for $\alpha > 2$) takes the most risky option (the \$-Bet), as opposed to the full insurance option (which any risk averse person would take). Secondly, the degree of “riskiness” in the option chosen is neither increasing or decreasing in α . The regret loving person ($\alpha < 1$) would choose the P-Bet, changing to a lower risk option (full insurance) as α increases beyond 1, before moving to higher risk option again (the \$-Bet) once α exceeds 2. Lastly, in this specific simple context, it’s not possible to have the regret minimising action change simply as a result of varying the magnitude of β , or, in other words, the perceived significance of the problem at hand. It is only through variance in α that we can construct a situation where the regret minimising action will change, supporting the need for a model which allows α to vary when seeking explanation for behaviour where the problem at hand has a familiar and unchanging structure⁶⁵.

1.4.3 The Repeated Problem

Varying β through λ

In this simulation, β represents the magnitude of the problem facing the individual. If β is small, then the consequences for each action are small, and hence the degree of regret associated with each action - state of the world combination (as given in Table 5) is small. If, however, β is large, then the consequences for each action are large, and hence the potential negative emotional consequences from regret are large.

⁶⁵ as was the motivation in the parking example given earlier on page 29

As previously discussed, in the repeated problem, some variability in β is needed to distinguish between learning about regret (for which this framework is appropriate) and learning about a very specific situation (where it is not). Yet, the precise nature of the variability can take many forms. For instance, we could imagine that the problem being repeated has a true, underlying value for β , say $\hat{\beta}$, and hence the variability could be a degree of “noise” associated with the observation of $\hat{\beta}$ in each time period. In this instance, we could think of the problem being given by $\beta_t = \hat{\beta} + \varepsilon_t$, where ε_t is simply a normally distributed error term. Looking at the dynamic properties, using regret, of a system like this, would give insight into how people react to situations where they under or overestimate the severity of a problem (i.e. $|\varepsilon_t|$ is very large), if they also make a bad decision, and hence experience regret, in the situation. That is to say, we could look at the consequences of compounding an error in judgement (large $|\varepsilon_t|$) with an error in action (and experiencing regret) on future choice. There is, however, a problem with seeking to answer this question in the current framework, and that is the parameter β is assumed to be > 0 in order to consider the three action problem (P-Bet, \$-Bet and Full Insurance) in terms of gains, and not losses. As such, any assumption made on the distribution of ε_t would have to include a restriction so that $\hat{\beta} + \varepsilon_t$ is positive, and this adds an extra layer of complexity in specifying the exact problem in the simulation model.⁶⁶

A different interpretation to consider is that β_t represents the severity of *any* uncertain decision faced on a period by period basis (be this annually, weekly, hourly etc.) and the three options are simplifications of a range of behaviours open to the individual; i.e. they can choose to behave very riskily (\$-Bet), mildly riskily (P-Bet) or very safely (Full Insurance) in each period. In this context, it doesn't make as much sense to assume that there is an “average” degree of severity, $\hat{\beta}$, around which there is some variance⁶⁷, but rather that the majority of decisions are small and inconsequential (low β), but occasionally there are more severe decisions (high β) which can have a large impact on welfare and emotions. This, more realistic, specification of the problem, is also beneficial as there is a natural, analogous, mathematical distribution for β ; the exponential distribution. By using the exponential distribution, it also allows us to fully characterise the problem at hand with only one parameter, λ , the rate parameter of the exponential distribution⁶⁸. λ , therefore, also represents the likelihood of getting a highly important problem in any time period (decreasing in λ).

As such, throughout the different stages of the simulation, we will take β_t to be given by a random draw from the exponential distribution, with parameter λ to be specified in each such case. Of interest will be the degree to which λ , or, in essence, the severity of the environment in which the decision maker is operating, will impact their behaviour, emotions and welfare.

66 Whilst the current set-up could be modified to allow for negative values of β , as in the motivating example, experimental evidence on loss aversion has shown significantly different behaviour when a problem is framed as a loss when compared to gain. Consequently, in this stylised computer simulation, only positive values are considered to abstract from the discussion of loss aversion.

67 and, if the variance is normally distributed, this implies that decisions of very low consequence are as likely as those with very high consequence, which seems fairly improbable in most contexts

68 under the exponential distribution, for instance, the mean is $1/\lambda$ and the variance $1/\lambda^2$

The Baseline Dynamic Model

For the purposes of comparison to results later on, it is important to establish a baseline dynamic model and note the behaviour which results under the simplest possible conditions. The baseline model enables us to specify parameters which will be kept constant throughout the simulations (such as the use of the exponential distribution mentioned in the previous section) and hence also the specific parameters which will be changed over the course of the different simulations. The specific outcomes of the model, which are of interest for describing welfare, emotions and patterns of behaviour, will also be introduced, to be used as a reference point against which all subsequent changes, which arise from differing the model parameters, can be judged.


The secondary aim of the baseline model is to present what happens when the static, three action model (as presented in 1.4.2) is repeated in the simulation *without* the introduction of any period to period feedback loop from regret. As such, this can be thought of as a summary of what the existing *Hayashi* model tells us, *on average*, about a P-Bet/\$-Bet/Full Insurance decision problem. The absence of the feedback loop, however, creates the need to make assumptions regarding α_t for use in the *Hayashi* model in each time period. As explained previously in 1.4.2, when $p(\omega_1) = p(\omega_2) = p(\omega_3) = 1/3$, the action chosen will depend entirely on α , and not on β (and hence λ if β is drawn from the exponential distribution). Hence, if $\alpha_t = \alpha$ and so fixed over time, then just one action will be chosen in every period of the simulation, and there is no more information to be gained from the repeated model compared to the one-shot model. Even if α_t is drawn from a distribution, it is simply the proportions of that distribution which lie in the three intervals ($0 < \alpha < 1$; $1 < \alpha < 2$; $\alpha > 2$) which will determine the proportions of each action which are chosen in the simulation. For example, consider a simulation of 10000 repetitions of the three action problem, where α_t is drawn from a uniform distribution in the interval (0,3). In this simulation, each action would be expected to be chosen $1/3$ of the time, however the value of λ , representing the severity of the decisions facing the individual, will determine the experienced utility and experienced regret. In addition, assumptions made regarding the nature of any impact biases (i.e. the difference between α_t , used to calculate expected regret, and α_t^e , used to calculate experienced regret) will impact the level of regret ultimately experienced by the decision maker over the course of the simulation.

This intuition can be summarised by Figure 15 which shows the results of the baseline model⁶⁹ being run 10000 times for a range of values of λ .

As λ increases, the mean of the exponential distribution from which β_t is drawn falls hyperbolically, and hence the average per period payoff falls hyperbolically too. As the results of the 10000 period simulation fall very close to the theoretical prediction, this suggests that 10000 periods is a large enough sample from which to draw accurate, per period conclusions.

We can also graphically represent the average per period regret experienced in the simulation, but, as previously stated, this will depend on assumptions made about the nature of the impact bias, or how the α_t used in the *Hayashi* decision calculation differs from the α_t^e which

⁶⁹ the baseline model has α_t drawn from a uniform distribution in the interval (0,3) and $p(\omega_1) = p(\omega_2) = p(\omega_3) = 1/3$

Figure 15: Average Per Period Utility by λ

actually generates the regret experienced. For example, if $\alpha_t = \alpha_t^e$, then the assumption is that there is no impact bias and regret is experienced exactly as it was predicted. In this case average per period regret is as shown in Figure 16.

As with the average per period utility, the average per period experienced regret falls as λ increases, but the rate of fall is faster for experienced regret, and the magnitude is also much larger for $\lambda < 1$. The implication of this is that, for low values of λ , or, intuitively, more severe decision making problems, the magnitude of experienced regret far outweighs any potential gains from experienced utility. Indeed, if “total experienced utility” is composed of the per period payoff *minus* the experienced regret then the individual is, on average and for values of $\lambda < 2$, experiencing “negative” utility, as indicated in Figure 17.

There are, however, other types of models of the impact bias that we can consider which better reflect the experimental evidence presented previously in 1.4.2. For instance, we can assume that the predicted coefficient of regret aversion, α_t , is unrelated to the experienced coefficient of regret aversion α_t^e , by assuming that experienced regret is proportional to the utility gap regardless of what was thought when taking the decision. This implies that $\alpha_t^e = 1$.

As an intermediate case, we can assume that the anticipated coefficient of regret, α_t is an *exaggeration* of the true experienced coefficient of regret, so that utility gaps are neither as “good”, or as “bad”, as they seem in anticipation. This can be achieved by setting $\alpha_t^e = \sqrt{\alpha_t}$ so that coefficients in excess of 1 are scaled back, and coefficient less than 1 are scaled up.

The results of making such assumptions about the impact bias are seen in Figure 19.

Simply moving from $\alpha_t^e = \alpha_t$ to $\alpha_t^e = \sqrt{\alpha_t}$ is sufficient to remove the “extreme” nature of regret seen when λ gets small, and, therefore, is a good representation of the impact bias whereby the individual’s own ability to rationalise, but not predict this rationalisation, leads to extreme anticipated regrets only rarely being seen in experience.




Figure 16: Average Per Period Experienced Regret ($\alpha_t = \alpha_t^e$) by λ




Figure 17: Comparison of Average Per Period Utility and Experienced Regret




Figure 18: Average Per Period Experienced Regret (APPER) and Average Per Period Total Experienced Utility (TEU) for different assumptions about the Impact Bias

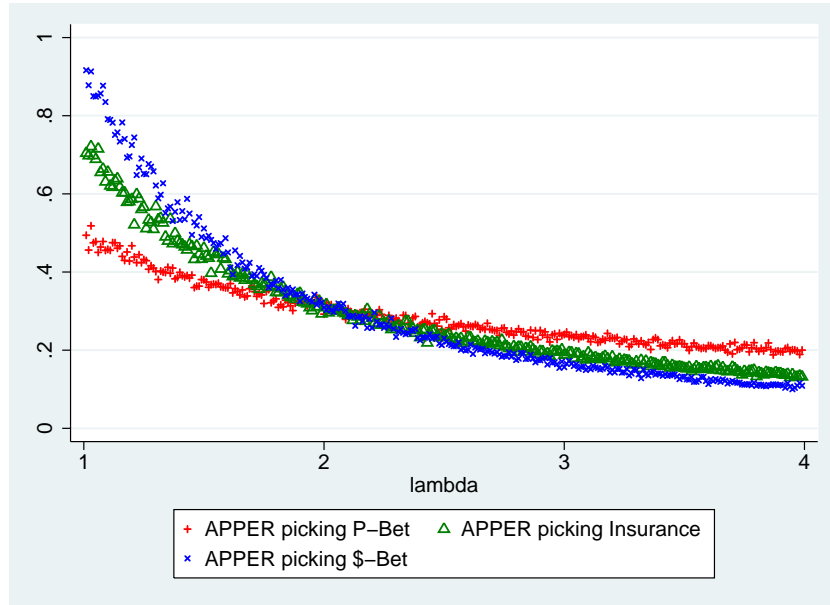


Figure 19: Average Per Period Experienced Regret for each type of action

Keeping $\alpha_t^e = \sqrt{\alpha_t}$ also allows us to see the emotional residue of regret that the individual would associate with each type of action chosen in the problem. As seen previously, regardless of β (and hence λ), it is the coefficient of anticipated regret aversion which will determine the action chosen in each period of the simulation (1.4.2), and so this could be expected to translate to associating the P-Bet with the smallest experiences of regret⁷⁰ and associating the \$-Bet with the largest experiences of regret⁷¹. However, by looking at the average regret experienced when each type of action is chosen, in Figure 19, we can see that this is conditional on the size of λ and hence the severity of the problem facing the individual.

For values of $\lambda < 2$, or fairly severe problems, it is indeed the case that the \$-Bet is associated with the highest average per period experienced regret of all three actions, and so the individual would likely associate the act of “risky” gambling (choosing the \$-Bet) with high levels of regret compared to the other actions. However, when $\lambda > 2$, the severity of the problems facing the individual are reduced so that a higher proportion of utility gaps are less than 1. Hence, raising them to a power of a number > 1 has the effect of reducing the magnitude of regret, not enhancing it. This effect is a consequence of the numerical simulation approach taken in this chapter, and the use of the exponential distribution to construct the framework of the simulation. Clearly, the utility gap being less than magnitude 1 has no economic interpretation, but the numerical effect of having these utility gaps dominate the simulation (when $\lambda > 2$) is to generate a model which has a counter-intuitive interpretation. As such, we will restrict much of the analysis which follows to values of λ which are less than 2 in order to maintain the general idea of a more regret averse individual suffering from utility gaps more than a less regret averse individual.

⁷⁰ when the P-Bet is chosen, $\alpha_t < 1$, hence $\sqrt{\alpha_t} < 1$, hence any utility gaps are raised to the power of a number less than one to calculate experienced regret

⁷¹ when the \$-Bet is chosen, $\alpha_t > 2$, hence $\sqrt{\alpha_t} > \sqrt{2}$, hence any utility gaps are raised to the power of a number greater than $\sqrt{2}$ to calculate experienced regret

Constructing the Memory Stock

In the repeated problem, the Memory Stock is constructed by adding the experienced regret in a specific period, R_t , to the set of regrets from all previous periods $M_t = (R_1, R_2, \dots, R_{t-2}, R_{t-1})$ to create the Memory Stock which is to be used in the next decision period, $M_{t+1} = (R_1, R_2, \dots, R_{t-2}, R_{t-1}, R_t)$.

However, at each decision period, there must be assumptions made on the nature of both “storage” of and “recall” from that memory stock. For example, a “fading affect” bias implies that the magnitude of the affect, given by R_t , will fade over time. Mathematically this can be given by a discount function applied to the memory stock, and a typical form of discounting used in economic models is exponential discounting. This implies a constant discount rate of δ such that the experience of regret at period t , but recalled at period s , where $s > t$, is given by $R_t^s = \delta^{s-t} R_t$. This means, for example, that the experience of regret in period t , when recalled 1 period later at $t + 1$, has already faded by a factor of δ ⁷².

In terms of “recall”, it is necessary to make assumptions on the cognitive limit of the individual in terms of the amount of information from the memory which can be recalled at any one time. As discussed previously, the first stage of this is that the individual does not remember the cause of the regretful experience, only the magnitude and time of the regret, but it further seems unreasonable to assume that regrets of 100, 200, 300 periods ago are recalled when facing a decision at period t ⁷³. As such it is necessary to place a time limit on the memory stock to be used in each decision period. Furthermore, as we are interested in the extent to which the experience of regret influences future regret aversion, it does not make sense to include “zero” regrets in the recalled memory stock, because, as the optimal was decision was made in any such period in which zero regret was experienced, there is no information to be gained on the scale of experienced regret⁷⁴. This requires us to prevent zero regrets from appearing in the recalled memory stock, so, when discussing a map from “The last x periods of regret to α_t ” this is technically interpreted as “The last x periods of *non-zero* regret to α_t ”.

For example, if we study the properties of the sets of recalled Memory Stocks, where, in each period, they are composed of the last 10 non-negative regrets, as they were constructed in the baseline dynamic model (so the memory stock is not impacting future decisions, sim-

⁷² assuming a constant discount rate is, of course, a first step approximation of how memory of regrets could fade over time, and one that should be tested empirically against other comparable discount methodologies, where the discount rate could vary both by the size of the regret and also the time lapse. It is chosen here for both mathematical simplicity and connection to existing economic literature where an exponential discounting model is by far the most common approach. The second step would be a model whereby the most meaningful regrets stay in memory the longest, perhaps even indefinitely, and small regrets fade to zero very quickly.

⁷³ equivalently, this can be thought of as there being a “time limit” on the storage of memories.

⁷⁴ zero regrets are not deemed informative, because the principle of the model is that the agent uses information about the past experience of regret to inform them of their likely aversion to future regret. In that sense, an experience of zero regret is informative in that it informs the agent that a zero utility gap results in a zero regret. Beyond that, it makes sense to exclude zero regrets from the memory stock for the same reason that positive utility gaps (“joy”) are excluded from the memory stock. We are using the memory stock only for the purpose of calculating the coefficient of regret, and so it is a subset of the complete memory and history of the decision making process chosen specifically to inform that calculation.




Figure 20: Average outcome of candidate mapping functions (f) from the Memory Stock (M) to α

ply being constructed), then, on average, the three different mapping functions proposed earlier in 1.3.2 produce widely differing degrees of anticipated regret aversion.

Figure 20 shows that a mapping from M to α , generated by the Max/-Median/Min ratio method, where M is constructed through accumulation of experienced regrets from the baseline model, encoded using the experienced regret model of the impact bias, given by $\alpha_t^e = \sqrt{\alpha_t}$, will produce, on average, values of α between 10 and 12, for all reasonable values of λ . If, therefore, we were to use values of α , generated using this process from the memory stock, as part of the regret-minimising decisions made in subsequent periods, then the expectation would be that the decision maker is exceptionally regret-averse, and so chooses the \$-Bet, as shown in Figure 14, virtually all the time. Similarly, whilst the Estimated Power Coefficient method generates, on average, values of α , much closer to those assumed in the baseline model⁷⁵, they are still, on average, in excess of 2, and so the same problem, of picking the \$-Bet too frequently, will be apparent. Lastly, the Skewness method appears to generate, on average, values of α which are in the middle of the range of values assumed by the baseline model, but, however, we must also be concerned with the *variance* of these values, as they will imply the proportion of choices which are the P-Bet, \$-Bet, and Full Insurance.

Given the baseline model draws values of α uniformly from the interval (0, 3), it is immediately clear from Figure 22 why the Max/-Median/Min ratio function would be an inappropriate choice as the function which translates the Memory Stock into future values of α in the dynamic model with feedback. That is, when the Memory Stock is generated simply as a function of regrets experienced as a result of random choices⁷⁶, then this mapping produces values of α which are wildly different to those assumed in the baseline model, and, consequently, any difference in choice behaviour, or welfare, resulting from the use of this mapping in the dynamic model would be primarily attributed to the use of the Max/Median/Min ratio, and not the process of affective feedback which the dynamic model is created to capture.

As shown in Figure 22 and Figure 21, the Skewness function and Estimated Power Coefficient produce more variance in the output. However, the Estimated Power Coefficient produced very few values below 1, and hence would imply that the decision maker would never be “regret seeking” and, therefore, would never choose the P-Bet in the

⁷⁵ uniformly distributed between 0 and 3

⁷⁶ when $\alpha_t^e = \sqrt{\alpha_t}$

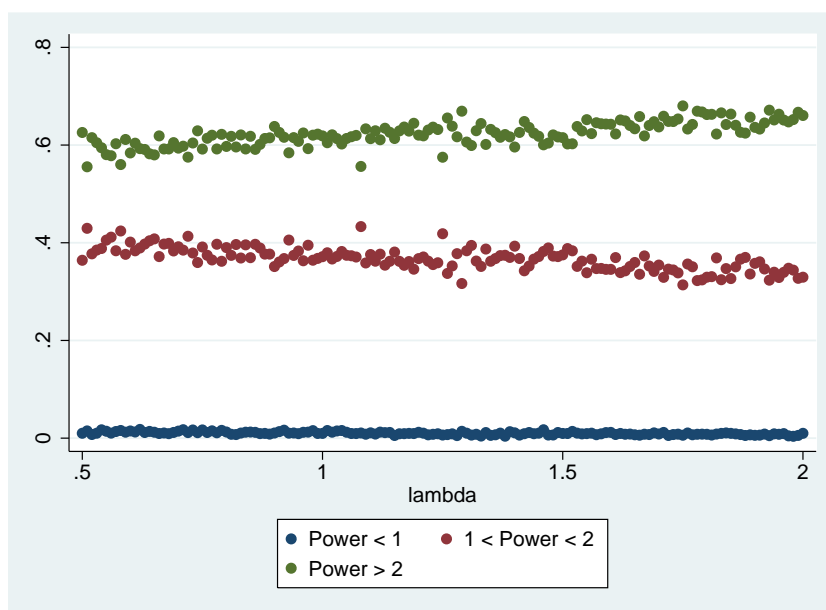


Figure 21: Proportion of Memory Stocks which produce values of less than one, between one and two, and greater than two, when a power coefficient is estimated




Figure 22: Proportion of Memory Stocks which produce values of less than one, between one and two, and greater than two, when sample skewness is calculated

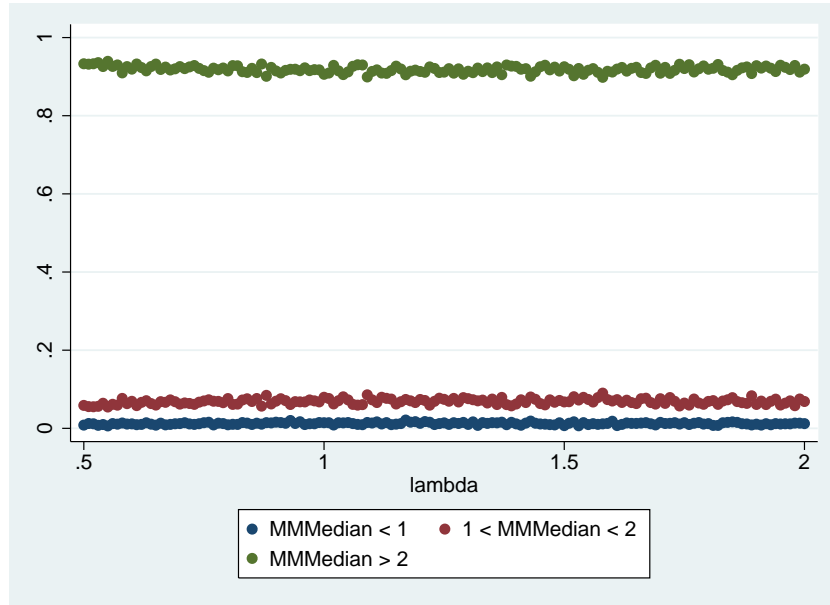


Figure 23: Proportion of Memory Stocks which produce values of less than one, between one and two, and greater than two, when ratio of the difference between Max and Median and the difference between Median and Min is estimated

simple three option model. Thus, a baseline model, which assumes α_t is drawn from a uniform distribution in the interval $(0, 3)$ will produce a memory stock which, when a power coefficient is estimated, emits “beliefs” about α_{t+1} which, on average, are not reflective of the “beliefs” about regret aversion, α_t which were used to make decisions. In the case of the Estimated Power Coefficient, a decision maker who chose each bet one third of the time, will be wondering why their memory stock is telling them that they should have chosen the P-Bet⁷⁷ with a probability of close to zero.

The Skewness function, therefore, is a more viable candidate function to use in the dynamic model, because it produces sufficient variance in the output that the beliefs of the individual, about their parameter of regret aversion, would be reasonable. Figure 22 shows that, when decisions are taken according to a parameter of regret aversion which is drawn from the uniform distribution on the interval $(0, 3)$, then the sample skewness of the memory stocks of regret, coming from these decisions, for values of λ between 0.5 and 2, would always imply that the frequency of each decision taken should be between 20% and 50%, which is not far from the actual frequency of 33% for each option.

Calibrating the Baseline Model

Ideally, however, we would want to construct a baseline model, whereby, when a function is applied to the memory stock to generate “endogenous” parameters of regret aversion, it produces a distribution of parameters which is exactly equivalent to the distribution of parameters which seeded the decisions in the baseline model. That is to say, if we seed the baseline model with parameters of regret aversion, drawn uniformly from the interval $(0, 3)$, then the function used emits a distribution of

⁷⁷ that is, when the memory stock emits a value of $\alpha < 1$

“endogenous” parameters which is the same uniform distribution. If this was true, then the function used is a very “believable” method for the individual to estimate their own aversion to regret, and we could construct a fully dynamic model with affective feedback.

It is incredibly unlikely, however, that it would ever be possible to find a psychologically justifiable function which emits a uniform distribution from the memory stock when applied. Hence, a second-best solution is to find a baseline model whereby the *proportion* of times that each decision is taken by the individual is equal to the *proportion* of times that the individual would take the decision if their parameter of regret aversion was given by the function applied to the memory stock. For example, if the baseline model used a uniform distribution on the interval $(0, 3)$ to generate the parameter of regret aversion, this would imply that each option is chosen one third of the time. However, as Figures 21, 22 and 23 show, none of the three proposed functions produce parameters of regret with equal proportions in the intervals less than 1, between 1 and 2, and greater than 2, and hence would result in the decision maker choosing each option an equal proportion of the time. As such, we must move away from assuming that, in the baseline model, the parameter of regret aversion is drawn uniformly in the interval $(0, 3)$ to an distribution which has a lower chance of producing a value of α less than 1.

One simple way to do this is to increase the lower bound of the uniform distribution from which the exogenously generated parameters of regret aversion are drawn. Consider drawing from the uniform distribution on the interval $(z, 3)$ where $0 < z < 1$. Then the proportion of α 's generated in the intervals $1 < \alpha < 2$ and $\alpha > 2$ will be the same, and equal to $\frac{1}{3-z}$, which will be greater than the proportion of α 's < 1 , which equals $\frac{1-z}{3-z}$. However, as Figure 24 shows, when calculating the difference between the proportion of α 's generated by the sample skewness of the memory stock, in each of the three intervals (< 1 , between 1 and 2, and > 2) and the proportion of α 's in each of the three intervals, generated by draws from the uniform distribution on the interval $(z, 3)$, there is no value of z for which all three interval differences are equal to zero, implying there will always be small differences in the beliefs about regret aversion elicited by the memory stock, compared to if these beliefs were just exogenously given at random, simply as a result of using the sample skewness as the function to transform the regret in the memory stock to a future parameter of regret version.

However, as Figure 25 shows, these differences (when summed across all three intervals by using the absolute value of the differences), are minimised when z is approximately equal to 0.67⁸. Though the sum is still strictly positive at this point, the results suggest that random draws from the uniform distribution on $(0.6, 3)$ will produce 10 period memory stocks of regret, when regret is experienced according to $\alpha_t^e = \sqrt{\alpha_t}$, which will, when the sample skewness of these memory stocks is taken, produce endogenously determined parameters of regret aversion which would predict the same pattern of choice, on average, as given by the random draws from the uniform distribution on $(0.6, 3)$. In shorthand, the “seeding” of the baseline model on this interval produces

⁷⁸ for values of z around 0.6, the results of the 10000 period numerical simulation tend to produce values of this sum between 0.03 and 0.12




Figure 24: Proportion of α 's generated by sample skewness minus proportion of α 's generated by uniform draws from interval $(z, 3)$, for three separate intervals, for values of $0 < z < 1$




Figure 25: Sum of absolute values of differences in proportions between sample skewness and uniformly draws, for each of the three intervals




Figure 26: Histogram of sample skewness of 10 period memory stocks generated by one run of the 10000 period baseline model (seeding from the interval $(0.6, 3)$)

“consistent” beliefs, according to the sample skewness memory function, and hence is a good choice for the baseline analysis of the model.

Skewness Correction

The above arguments present a compelling case for using the sample skewness as the function for translating the memory stock of regrets to the next coefficient of regret aversion, but there is an obvious issue with using the sample skewness. Whilst the coefficient of regret aversion, in the Hayashi model must necessarily be greater than zero, the sample skewness may well be less than zero. Indeed, Figure 27 shows that there is a small, but significant (less than 0.5% of values), left tail of skewness values which fall below zero.

Whilst the simplest solution would be to simply replace the skewness metric with one of the other candidate metrics, discussed earlier, in the model whenever skewness falls below zero, test analysis shows that the probability of the other metrics falling in the interval $(0,1)$ whenever sample skewness is below zero, is very close to zero. Thus, this replacement method would generate parameters of regret aversion wildly in excess of what the spirit of the model suggests they should be. Consequently, the replacement method chosen is, when sample skewness falls below zero, to replace it with a random number drawn uniformly on the interval $(0,0.2)$. Doing so constrains the sample skewness to be strictly positive, hence being able to be used as a parameter of regret aversion in the Hayashi model. Using this replacement method transforms the original histogram in Figure 26 to a new histogram in Figure 27.

Definition of the baseline dynamic model

Thus, following all of the above analysis, the baseline dynamic model is constructed by running a 10000 period simulation of the simple static




Figure 27: Histogram of sample skewness from Figure 26 modified by transforming negative values to random numbers in the interval $(0, 0.2)$

model, given on page 45, where the parameter of regret aversion used in each period is given by drawing from a uniform distribution on the interval $(0.6, 3)$. After the decision is taken in each period⁷⁹, if a strictly positive regret is experienced, it is experienced according to an impact bias, given by $\alpha_t^e = \sqrt{\alpha_t}$. This experienced regret is then added to the memory stock of the last 10 strictly positive regrets, replacing the regret which was experienced the longest time ago.

Results of the baseline dynamic model

As stated earlier in 1.4.3, the magnitude of λ in the dynamic model will determine the severity of the environment in which the decision maker, who happens to be a regret minimiser, is operating. This will determine both the size of the average reward to decisions taken, and also the average regret experienced. In addition, by defining Total Experienced Utility as the Payoff minus Experienced Regret, in each period, we can get a baseline measure of “welfare” as it exists in this model, and how it varies with λ .

As shown in Figure 33, for all values of λ between 0.5 and 2, the average per period Payoff is less than the average per period Experienced Regret. This implies that average per period Total Experienced Utility is negative in the baseline model, or, more intuitively, the fact that the decision maker has the potential to experience regret from their decision, even accounting for the effect of the impact bias (that regret is worse in anticipation than experience) turns a potentially financially profitable repeated gambling scenario, into a negative overall experience, using the regret minimising model of Hayashi. The degree to which the experience will be negative is controlled by λ , the rate parameter of the exponential distribution. When λ is small, this implies the mean number drawn from the distribution is large, and hence the

⁷⁹ i.e. the decision maker has chosen between the P-Bet, \$-Bet and Insurance according to their parameter of regret aversion and the Hayashi regret minimizing procedure

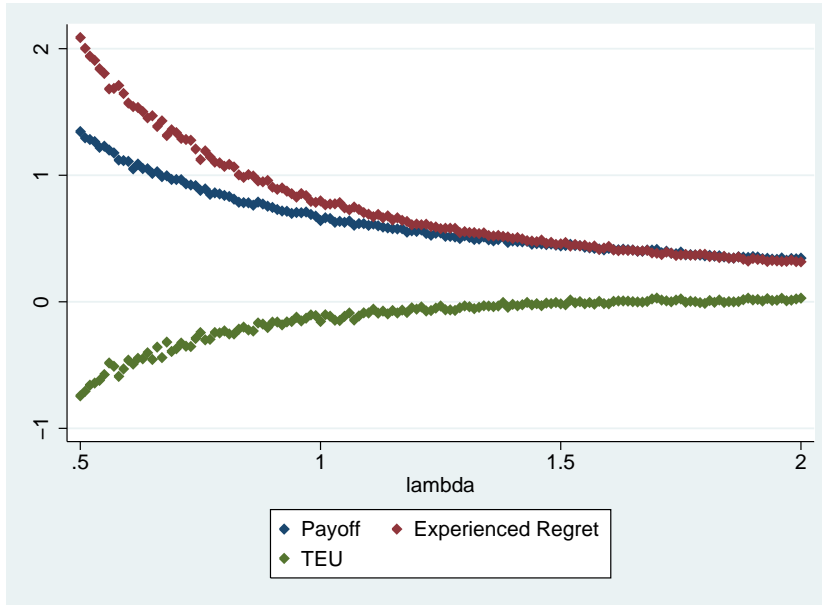


Figure 28: Average Per Period Payoff, Experienced Regret and Total Experienced Utility as a function of λ in the baseline dynamic model

consequences, both for payoff and experienced regret, are large. However, the growth rate of experienced regret exceeds the growth rate of the average payoff, and hence the Total Experienced Utility becomes more negative.

Given, in the baseline model, the parameter of regret aversion is drawn from the uniform distribution on the interval $(0.6, 3)$, and that the decision maker will choose the P-Bet if $\alpha < 1$, this implies that the P-Bet will be chosen in the baseline model approximately $1/6$ of the time. Similarly, the Full Insurance and the \$-Bet will be chosen $5/12$ of the time. This was chosen so that the modified skewness metric on the memory stock would emit a similar pattern of α 's when random truly parameters of regret aversion were put into the model. Thus, anything aside from small deviations in this behavioural pattern, when the affective feedback mechanism is introduced into the model, are attributed to the feedback mechanism, and hence the process of experienced regret affecting anticipated regret aversion, and not the choice of metric on the memory stock⁸⁰.

1.4.4 The dynamic model with regret feedback loop

Simply replacing the exogenously given parameter of regret aversion, drawn from the uniform distribution on the interval $(0.6, 3)$, in the

⁸⁰ A fair challenge would be to wonder whether the choice of the estimating function (in this case, the modified skewness metric over the max/median/min ratio or estimated power coefficient) is responsible for driving the results found, rather than any effects due to the creation of the regret feedback loop. As discussed in this paragraph, the choice of the modified skewness metric was precisely to minimise the role of the function itself, as putting random α 's into the baseline dynamic model gives you a distribution back out equivalent to what you put in. As the other functions investigated could not be calibrated to have that property, it's not possible to describe the effect of the choice of estimating function on the results, as you are unable to make an "apples to apples" comparison. An interesting follow-up would be to find other candidate estimating functions which can be calibrated to have the same baseline property as the modified skewness metric, to then empirically demonstrate the sensitivity of the results to the estimating function.




Figure 29: Patterns of choice in the dynamic model with regret feedback loop

baseline dynamic model, by the modified skewness metric⁸¹, calculated over the last ten strictly positive regrets which the decision maker experienced, modified by the impact bias so that $\alpha_t^e = \sqrt{\alpha_t}$, gives the pattern of choice behaviour shown in Figure 29 when evaluated over different values of λ in the range $(0.5, 2)$ ⁸²

The pattern of behaviour is clear from Figure 29 and markedly different from the behaviour of the randomly seeded baseline model in Figure 27. Firstly, what was previously the least commonly chosen option; the P-Bet; has now become the most often chosen option, ranging between being chosen 40% and 60% of the time, depending on λ . Indeed, there is a clear positive trend in the probability that the P-Bet is chosen, with a unit increase in λ , being associated with a 7 percentage point rise in the frequency⁸³. This increase in the probability of choosing the P-Bet comes at the expense of both the \$-Bet and Full insurance options, which each fall to between 20% and 30% of the options chosen. Both of these options are negatively related to λ ⁸⁴. This behaviour implies that the distribution of the parameter of regret aversion, generated by the modified skewness metric, is much different to that produced by the random baseline model.

Figure 30 shows the distribution of α_t generated by the dynamic model, with regret feedback loop, when $\lambda = 1$. It is then immediately clear that the reason for the change in behaviour is due to the increase in prevalence of parameters of regret aversion very close to zero.

Figure 31 shows that, when the skewness metric is left unmodified in the dynamic model with regret feedback, approximately 14% of the observations fall below zero. Hence, in the modified skewness metric, transforming all such observations to fall in the interval $(0, 0.2)$ condenses a significant proportion of the distribution a very small space. However, whilst it is clear that a decision maker whose memory stock

⁸¹ modified so that negative values are replaced by a random number drawn from the uniform distribution on the interval $(0, 0.2)$

⁸² running each simulation over 10000 periods and taking the average, per period behaviour

⁸³ significant at the 1% level

⁸⁴ both significant at the 1% level




Figure 30: Histogram of α_t generated in the dynamic model with regret feedback loop with $\lambda = 1$




Figure 31: Histogram of skewness of memory stocks, generated in the dynamic model, with $\lambda = 1$

is negatively skewed should be considered “regret averse”, any similar transformation of negative skewness into the interval $(0, x)$, where $x < 1$, still leaves the P-Bet as the most frequently chosen option⁸⁵. Hence, whilst the process produced slightly unhelpful metrics, in terms of the precise magnitude of the switch towards the P-Bet as a result of the feedback process, the qualitative result of a shift towards the option associated with parameters of regret aversion below 1 is consistent.

The reason for this shift towards the P-Bet is that, when the feedback mechanism is introduced, if the coefficient of regret in any given period, α_t , is less than one, then, as shown in Table 5, in case where the decision maker does experience regret, it will be given by $(2\beta)^{\sqrt{\alpha_t}}$, which, because of the exponential distribution given by λ , is unlikely to be a particularly large number. Indeed, in the event that the decision maker chooses the P-Bet, and experiences regret, their per period average experienced regret is as given in Figure 19. Secondly, as picking the P-Bet will result in the experience of regret only if state of the world ω_3 occurs, there is a $2/3$ probability that no regret will be experienced, and hence the coefficient of regret aversion α_{t+1} equals the coefficient of regret aversion α_t . The combination of these two effects implies that, should the memory stock reach the point where the implied coefficient of regret aversion is less than one, then it is very likely to remain so for the foreseeable future, and hence tell the regret minimising agent to keep picking the P-Bet. This compares drastically to the baseline dynamic model, when the probability of picking the P-Bet, given the P-Bet was picked in the last period, was only equal to $1/6$.

Indeed it is possible to estimate the probability of observing the same option chosen, regardless of which option that would be, in two successive periods in this dynamic model, and compare this to the probability of observing the same behaviour in the baseline model, which is exactly equal to $3/8$ ⁸⁶.

Figure 36 shows that the proportion of “repeat choices” is increasing in λ , but generally falls between 86% and 91% of all choices being the same as in the previous period. This is increasing in λ , because increasing λ reduces the probability of the decision taken in any one period being severe enough⁸⁷ to generate a regret large enough to give a high positive skew to the memory stock of the individual, hence making it more likely that the memory stock in any one period will be similar to the memory stock in the previous period, (and, in particular, having a skewness of less than one giving rise to selecting the P-Bet again) hence increasing the probability of repeat choice.

This increased probability of repeat choice, from the baseline model of 0.375 to between 0.86 and 0.91, suggests that the process of feedback from experienced regret to anticipated regret gives rise to an individual getting “addicted” to one particular type of behaviour (such as the P-Bet) before a sufficiently negative regretful experience happens to them to shock and skew their memory of the emotion in a new direction, giving them reason to change their behaviour to another action.

In addition, because the P-Bet is selected when the coefficient of regret aversion is below 1, and the individual becomes addicted to the choice of the P-Bet, keeping the *experienced* coefficient of regret aversion

⁸⁵ data not shown, but indicated through simulation

⁸⁶ the probability of observing the same option chosen in successive periods in the baseline dynamic model is equal to $(\frac{1}{6} + \frac{1}{6}) + 2(\frac{5}{12} + \frac{5}{12}) = \frac{3}{8}$

⁸⁷ with the average severity of a decision being given by the mean of the exponential distribution, $1/\lambda$




Figure 32: Proportion of periods during which the option selected was the same as the previous period, in the dynamic model

below 1, the average per period regret experienced is kept low. This is shown in Figure 33, indicating that the average per period payoff is only marginally below the average per period experienced regret for values of $\lambda < 1$, and, for values of $\lambda > 1$, the average per period total experienced utility (TEU) is actually positive.

This compares favourably with the baseline dynamic model, where the average per period total experienced utility was negative for all values of λ , suggesting that the addition to the P-Bet, resulting from the feedback loop, actually generates a positive impact on the welfare of the individual by keeping their sensitivity to regret lower than in the baseline model.

The role of discounting

The dynamic models used so far have simply assumed that, once an experienced regret is in the memory stock, it stays at the same magnitude in memory until it becomes the 11th most recent regret, at which point it drops out of the memory stock altogether. It is reasonable to assume, however, that the main effect of the passing of time is to dilute the magnitude of the experience of regrets which still remain in the memory stock. To this end, it is possible, within the simulation, to introduce a simple model of exponential discounting, and a time discounting factor, δ , such that the i^{th} oldest regret in the memory stock is discounted by a factor of δ^{i88} , where $0 < \delta < 1$.

Figure 34 shows that the pattern of behaviour observed in the “first look” at the dynamic model with regret feedback loop, of a strong preference for the P-Bet, is changed by the effect of imposing time

88 this is subtly different to assuming a regret which was experienced i periods ago is discounted by a factor of δ^i . In the version used in the dynamic model, periods which produce zero regret add nothing to the memory stock, and, hence, do not impose any additional time discounting on the memory stock. The impact of this slight difference on the qualitative idea of time discounting is minimal, but makes the coding of the model much easier!




Figure 33: Average Per Period Payoff, Experienced Regret and Total Experienced Utility as a function of λ in the dynamic mode with regret feedback loop

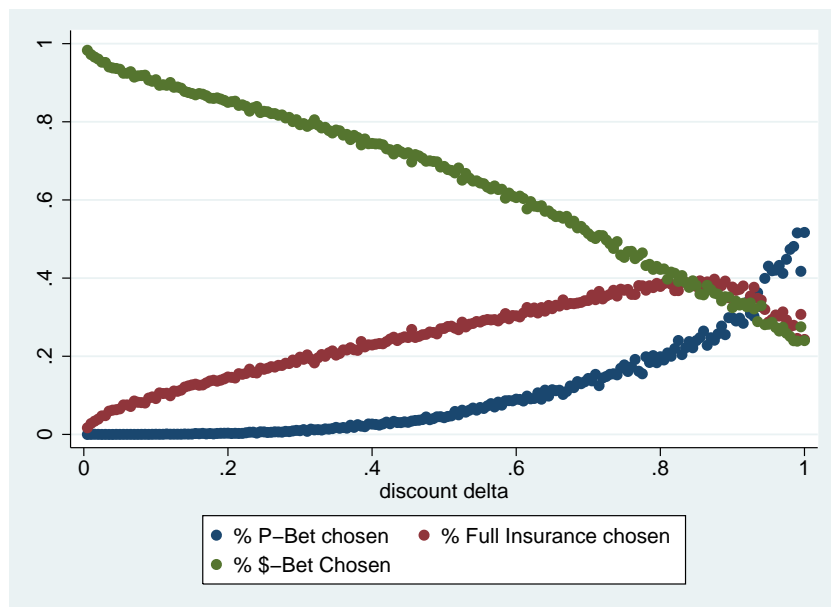


Figure 34: Impact of exponential time discounting of memory stock on pattern of choices made in dynamic model with regret feedback loop for $\lambda = 1$




Figure 35: Average Per Period Payoff, Experienced Regret and Total Experienced Utility as a function of δ in the dynamic model with regret feedback loop, for $\lambda = 1$

discounting on the memory stock of regrets. For any time discounting factor of less than 0.9, the P-bet becomes the least frequently chosen option, with a shift towards both the Full Insurance and \$-Bet options. Furthermore, if the discount factor is lowered further, there is an even greater shift away from both the Full Insurance and P-Bet options towards the \$-Bet. The reason for this shifting pattern of behaviour is obvious, in that by exponentially discounting the memory stock of regrets which have happened in the past, the average magnitude of any regrets already in the memory stock is significantly reduced, and so the “newest” regret which is added to the memory stock is more likely to be in excess of those already in the memory, and hence increase the skew of the memory stock, which translates to an increase in the coefficient of regret aversion used in the decision making of the next period, resulting in an increased probability of selecting the \$-Bet. Indeed, this increase in the coefficient of regret aversion translates into an increase in the average per period experienced regret, and hence an fall in average per period total experienced utility, as shown by Figure 35.

Furthermore, Figure 36 shows the proportion of repeat choices as a function of δ , in the dynamic model with regret feedback loop when $\lambda = 1$. The prevalence of the P-Bet, when δ is close to 1 implies that it is likely that the decision maker will be “addicted” to the P-Bet, but the increased prevalence of the \$-Bet when δ tends towards 0 implies the decision maker becomes “addicted” to the \$-Bet instead. The consequence of the decision maker discounting past regrets is that new regrets loom relatively large in the memory of the individual, hence making them averse to subsequent regrets. This then manifests itself through a pattern of addition to the \$-Bet, whereby the individual fears missing out on the big payoff from the \$-Bet if they do not choose it. This compares to the case when δ is close to 1, and the decision maker places the value of the latest regret into context alongside the regrets in their recent past. By doing this, they become desensitised




Figure 36: Proportion of periods where the option selected was the same as the previous period, in the dynamic model, for $\lambda = 1$, as a function of δ

to the large regrets, hence choosing the P-Bet in subsequent periods. Given this desensitisation to regret in anticipation then translates into a desensitisation to regret in experience, the effect on overall total experienced utility is positive as a result of becoming addicted to the P-Bet as opposed to the \$-Bet, despite them offering the same expected value.

1.5 CONCLUSIONS

1.5.1 Comparing the baseline dynamic model to the dynamic model with regret feedback loop

The baseline dynamic model studied the per period average behaviour of a regret minimiser (who behaves according to the smooth model of regret aversion given by Hayashi) when faced with a decision between a P-Bet, \$-Bet and Full Insurance (with the expected value of each option being the same, and drawn from an exponential distribution), in 10000 consecutive periods, where the anticipated coefficient of regret aversion used to make the decision was drawn from a uniform distribution on the interval $(0.6, 3)$, and the *experienced* coefficient of regret aversion was given by an impact bias transformation on the anticipated coefficient ($\alpha_t^e = \sqrt{\alpha_t}$).

In contrast, the model with the regret feedback loop assumes that the anticipated coefficient of regret aversion is not exogenously given in each period, but rather endogenously derived from a metric applied to the last 10 strictly positive regrets that the regret minimiser experienced, where the magnitude and time of the regret are stored in memory, but not the precise utility gap which gave rise to the regret. This metric is a modified skewness (modified to produce only positive numbers), which, when applied to the regrets generated in the baseline dynamic model, gives a distribution of numbers which closely resembles the distribution of the numbers which seeded the baseline model.

1.5.2 *Total Experienced Utility*

A notable result from this analysis is that, both in the baseline dynamic model and the model with regret feedback loop, the average per period payoff which results from choosing one of the three options tends to be less than the average per period experienced regret, even accounting for the effect of the impact bias. This pattern holds until the proportion of P-Bets taken in the dynamic model is sufficient that the typical experienced coefficient of regret aversion is below 1, and hence the effect of any particularly large utility gaps, arising from the use of the exponential distribution in the model, is greatly reduced by the coefficient of regret aversion falling below 1.

The net effect of these factors is that Total Experienced Utility (TEU), which equals the payoff minus experienced regret in any given period, is negative, or only very marginally positive, for almost all scenarios considered in the analysis. This suggests a significant issue with using any model similar to the Hayashi model, in that by exposing the decision maker to the possibility of (ex-post) regretting their decision, that a situation such as this in which, from a purely monetary perspective, they can only be strictly better off as a result of having the opportunity to choose between the three options, turns into an experience which is ultimately negatively impacting their overall quality of life. In short, they would have preferred to have never been offered the opportunity to win free money, even in a situation where they can always choose Full Insurance and guarantee themselves a strictly positive payoff in each stage, because the fact that they always view the resolution of uncertainty, no matter what their choice, implies that they will, on average, experience a level of regret which outweighs the benefits of the positive payoff. In the language of Sarver, the decision maker in this simulation would tend to have a preference for a “menu” of zero choices over a “menu” of three choices, all of which give strictly non-negative payoffs in every state of the world⁸⁹.

It seems incredibly unlikely, however, that an individual would ever turn down a menu of gambles which, in monetary terms, is only ever Pareto improving, even accounting for the possibility of the experience of regret from not choosing the ex-post superior option every time. Thus, the model of smooth regret aversion, created by Hayashi, which transforms utility gaps into the experience of regret through the coefficient of regret aversion, is arguably not an accurate representation of the way in which both the experience and anticipation of regret should be modelled as it applies to real world decision making.

1.5.3 *Discounting*

In this simulation model, discounting of past emotional memories, or “biases” in storage, contribute in a standard way towards increasing the likelihood that the \$-Bet will be chosen by the regret minimising individual. By reducing the magnitude of the past emotional experiences, any new regret which is experienced will appear disproportionately larger than the past experiences, hence giving the decision maker the impression that they are disproportionately affected by large regrets

⁸⁹ noting that, as in the model of Sarver, we are concerned with “...maximi[sing] the expectation of ... utility minus regret”[83, p269], but, unlike Sarver, we are concerned with more than just “...simple linear forms of regret”[83, p281]

(when, in fact, they are simply “newer” regrets), and so making them more regret averse in the future. This is equivalent to a decision maker who acts in a very “hot” state, overreacting to the most recent events in an attempt to take action to stop them from happening again in the future⁹⁰.

This result highlights the value of any individual being able to appropriately reflect on a recent negative experience before making another decision under risk or uncertainty. Individuals who encounter such decisions on a regular, repeated basis⁹¹ would do well to take accurate records of the decisions they take and the consequences that resulted from those decisions. The effect of keeping accurate records, and associated emotional consequences, could then be seen in this model as the elimination of the “fading affect” so that, should a new regret arise as a result of a decision, it is able to be properly contextualised alongside similar regrets in the past. This transforms the context of the next decision from being in a “hot” state, of overreacting to the last experienced regret, to a “cold” state, of careful consideration of the relative magnitude of all regrets in the memory stock / diary. This effect lowers the coefficient of regret aversion, and hence reduces the probability of taking the \$-Bet in any given period. By keeping diaries and journals, the salience of past regretful experiences is increased, so as to eliminate the bias of relying on the most recent experience to determine the aversion towards regret.

1.5.4 *On the use of Monte Carlo simulation methods*

The approach of this chapter, in using Monte Carlo simulation methods to analyse the effects of endogenising the parameter of regret aversion in a dynamic version of the Hayashi regret minimiser model, has both advantages and disadvantages.

The primary disadvantage is that the specific functional forms, embedded in the model, and skewness metric, used for the generation of the coefficient of regret aversion, are entirely arbitrary representations of the spirit of the literature on affective feedback. They were chosen so as to keep the model both simple and maintaining some resemblance to the behavioural ideas they represent. However, it is not possible to examine the effects of imposing such restrictions on the simulation without other, alternative functional forms to compare them against, which may be just as arbitrary without any discernible improvement in the degree to which they represent the underlying intuition.

Where there is a clear economic literature on a specific function (for example, the use of exponential discounting to represent fading affect) then it is possible to analyse the behavioural consequences of the model with respect to changes in a given parameter value. However, in some cases, such as with the skewness metric applied to the memory stock, this is a new approach which does not have a existing associated methodology as to how best to consider marginal changes in

⁹⁰ There is a parallel of this argument in Case-based Decision Theory (Gilboa and Schmeidler) where decisions are taken according to a similarity function, which is implicitly affected by the length of time over which “cases” are drawn from in the history or memory of the agent. If you were to restrict the history of “cases” to only very recent events, you may see a similar result where an agent is actually reacting to “recency” as opposed to “similarity”.

⁹¹ for example, those working in financial investment environments, or those who frequently gamble for leisure

the way the metric is applied⁹². Indeed this version of the simulation is restrictive by considering only the case of three possible betting options, by considering only the case where the expected value of all options is the same, and by considering only the case where each state of the world is equally likely. For considering deviations to these assumptions would require a completely new behavioural analysis conditional on the values chosen for every other parameter in the model.

The advantage of this simulation approach is that it permits a glimpse into the non-linear world of regret and affective feedback where it is clearly not possible to analytically solve the dynamic model, especially when there is no clear and obvious functional form for how the coefficient of regret aversion should evolve over time. This contrasts with the work of, for example, [Houthakker and Taylor](#), who study the evolution of patterns of consumption over time, analytically, using a differential equations framework, where the complexity of the model is deliberately kept low in order to be able to apply the differential equations techniques.

As such, given the progression of economic modelling towards non-linear and dynamic models, the use of simulation techniques may be the only way to conduct meaningful analysis when the degree of complexity in the model (in this case given by the number of different psychological concepts which are integrated) becomes large. However, once the underlying mechanics of the model are constructed, given sufficient structure as to how a particular insight or idea should interact with the framework of the model, it is quick and easy to incorporate new ideas and generate new data which provides a “first look” at the types of behaviour which could emerge in such a world. Whilst these first looks may not be perfect, they can be at least indicative of answers which can then be further investigated empirically and experimentally.

⁹² for example, the decision to change all negative skewness values to be random numbers in the interval (0,0.2) makes sense in the baseline model, given the distribution of numbers it produces, but it less sensible in the model with affective feedback

2.1 INTRODUCTION

2.1.1 *Approaching Experiential Regret Aversion Experimentally*

Behavioural economics, developed from psychological foundations, and methods of experimental economics, have long enjoyed a close working relationship. As behavioural economics is often associated with a “descriptive” account of human decision making, once we are able to describe such behaviour using economic modelling, it only makes sense to test whether human decision making conforms to the predictions of these models. When the models relate to areas of individual decision making (for example, behavioural models of individual decision making under uncertainty, such as Prospect Theory (Kahneman and Tversky [48]), the most natural environment to test the models in is a laboratory, where each observable aspect of the decision to be made can be controlled and calibrated, and any possible unobserved aspects, which may correlate with the main variable of interest, can be indirectly controlled through randomisation, both within and across treatment groups.

As the original models of Regret Theory (Loomes and Sugden [58], Bell [5]) were developed at the same time as, and became a part of, the class of non-Expected Utility Theory models, such as Prospect Theory, they were obvious candidates for experimental testing in a laboratory setting.

This approach, however, raises several questions, some of which have been more widely addressed in the literature than others.

- Firstly, in order to test the usefulness of a regret-based theory of decision making under uncertainty, there needs to be at least one defining characteristic of the theory, which makes a testable prediction about decision making, observable under laboratory conditions.
- Secondly, in order to compare the usefulness of a regret-based theory against at least one other non-EUT model, there must be a prediction of the regret-based model which can *differentiate* it, and is also *observable* in a laboratory setting. For instance, it is not sufficient to say that the regret-based model predicts violations of transitivity (as many models do so also) or that the decision maker will make a calculation of expected anticipatory regret ex-ante (as this is not observable in a laboratory).
- Thirdly, if one of the defining characteristics of the theory is that individuals do consider the anticipated regret of a potential decision, then, if techniques and methods are developed such that both anticipated and experienced emotions can be measured in some way¹, then specific laboratory tasks which generate observable actions in support of regret-based theories, should also

¹ such techniques might include self-reports, or, perhaps, psycho-physiological measures such as fMRI and cardiovascular measures

receive support from emotional measures when they are applied to the same tasks².

The first and second points were approached by economists and psychologists in the 1980s and 1990s, and a wide range of results have been found, especially when comparing competing non-EUT theories, some of which will be touched on in subsequent sections. The third point is an area which has not been approached by either economists or psychologists, and, broadly speaking, corresponds to ensuring that models of decision making which use emotions as an integral component are checked to make sure that the emotion in question behaves in the model as it does in reality, given our present understanding of that emotion³. This point will be explored in the current chapter.

As the principle focus of this work is on Experiential Regret Aversion, or how the experience of regret subsequently impacts anticipatory regret aversion, subsequent discussions will be limited to experiments which attempt to answer whether or not the experience of regret does affect future anticipation of regret *and* use the typical models and experimental methods in regret theory as a given, rather than exploring experiments which approach the first two questions directly. Such experiments lie at the cutting edge of experimental work in regret theory and have important consequences for our understanding and development of future regret-based theories. As such, it is important to determine whether the results of such experiments are correctly interpreted, especially given the challenge of both measuring regret, and separating regret from other similar emotions in decision making (such as disappointment, as in Disappointment Aversion (Bell [7], Loomes and Sugden [59])). It is to this question which we turn first.

2.2 EXISTING EXPERIMENTAL LITERATURE

2.2.1 *From Economics*

The importance of non-EUT theories being supported by experimental evidence is clear from the original Regret Theory of Loomes and Sugden [58], as they extensively discuss the implications of their theory for “...choices between pairs of statistically independent prospects.” [58, p810], as evidenced by multiple experimental results of Kahneman and Tversky⁴. They describe how previously observed experimental results are well supported by the predictions of Regret Theory, and later, Loomes et al. [62] and Loomes and Taylor [61] design new experiments to show that, when including the “regret aversion”⁵ assumption in Regret Theory, the theory predicts violations of transitivity in a specific direction, which are observed in laboratory tasks as significantly more prevalent than violations in the counter direction.

² that is to say, if a laboratory task produces results which suggest a regret-based decision model may have been used (through the idiosyncrasies of the theory), then it would lend weight to the conclusion if an direct measure of regret could be also be found in the same task.

³ at the most simple level, this could correspond to ensuring that measures of regret are increasing in the magnitude of utility which is foregone from making a “bad choice” in lab tasks.

⁴ including the most commonly observed deviations from EUT, such as the “common ratio” and “common consequence” effects

⁵ that large regrets loom disproportionately large in the mind of the decision maker when compared to small regrets




FIGURE 1.—A screen dump from the circles experiment.

Figure 37: Figure 1 from Hey and Orme [36]

This early approach to experimental testing of Regret Theory mirrors the experimental testing of other non-EUT theories at the time. This typically includes theoretically demonstrating that a specific, non-random pattern of behaviour, which we can observe in the laboratory (such as a preference reversal in the case of violations of transitivity), does not necessarily lead to contradictions in the theory itself. That is to say, the experimental evidence is not *incompatible* with the predictions of the theory. This approach can be thought of as the “Economist’s approach”, whereby it is not the *process* of decision making which is being analysed for its compatibility with the theory, but rather the *observed behavioural outcome* of decision making which is being studied.

As stated in the introduction, taking this approach enables you to not only consider the predictive power of the theory on its own (for example, you could state that, in an experiment, 82% of participants made decisions which are consistent with the predictions of Regret Theory), but also allows you to compare similar theories, by designing experiments which lead to outcomes which would be compatible with one theory but not with another. An example of this approach is in Hey and Orme [36] where they ask participants to make a very high number of decisions in problems such as the one presented in Figure 37.

Typically, such problems will include choosing between a “safe” (the right circle in this case) and “risky” (the left circle in this case) option, and, empirically, the frequency with which the safe option was chosen over the risky option will be the dependent variable in a regression. In the case of theories considered by Hey and Orme, all “...imply a valuation of the left-hand gamble relative to the right-hand gamble”, and estimating a set of parameters, which create an estimated preference function unique to each theory, allows us to test the parameters

as implied by the theories. By comparing the frequency with which individuals satisfied the restrictions placed upon the parameters by the theories, we can estimate which theory appears to “best fit” the observed decisions. In this experiment, Regret Theory performs very much in the middle of the pack, as significantly less predictive than Expected Utility Theory, and marginally less so than models which use “rank dependent utility functions”.

One paper that seeks to use data from such studies to its fullest extent is by Harless and Camerer [30], who construct a meta-analysis based upon 23 data sets involving similar choice problems to the one presented in Figure 37. Though this approach is incredibly useful for attaining statistical significance in results, it does not include results for Regret Theory, so can offer no more information, for our purposes, than already provided by Hey and Orme. This does, however, demonstrate a problem with Regret Theory in calculating usefulness based solely on empirical methods from decision frequency data. As regret theory posits comparison between actions, in each state of the world, it necessarily implies a significant number of parameters to estimate in such regressions, which reduces the likelihood of significant results (and, hence, the likelihood that they would receive publication).

These two papers summarise the limited extent to which the “Economist’s approach” can help provide information about the descriptive capability of Regret Theory. Other non-EUT theories are often preferred over Regret Theory in economics, specifically because it is difficult to estimate empirically simply based on observable choice data (and, hence, extend out of the lab into the real world), and other theories appear to offer a better predictive power⁶. As such, the mid 1990s saw the end of extensive experimental testing of Regret Theory using observed choice data, but, at the same time, a new literature was beginning to emerge in psychology which designed experiments to investigate the nuances of decision making behaviour, using the implied process and assumptions of the original form of Regret Theory as the theoretical basis.

2.2.2 *From Psychology*

The early 1990s saw an uptake in interest about regret from researchers in psychology, based, in part, on the success of the theory in explaining previously observed experimental results (as detailed in Loomes and Sugden [58]) which showed behaviour violating the predictions of Expected Utility Theory. A wide range of experiments were subsequently designed which took the assumptions and predictions of Regret Theory as a given (specifically, the assumption of regret aversion) and sought to demonstrate behaviour which would lend support to the implied process behind Regret Theory (rather than empirically testing the predictive capability of the theory as economists had done). The typical experimental design followed the format

1. Create two treatment groups, A and B
2. Subjects in both groups are asked to make a decision under uncertainty or risk
3. One action which treatment group A could choose, however, exposes them, potentially, to more regret than the equivalent

⁶ for instance, rank-dependent theories such as discussed by Quiggin [71]

action for treatment group B⁷, under the hypothesis that the primary reason the actions differ is linked to regret

4. Under the assumption of regret aversion, those subjects in group A will be sensitive to the presence of the large regret, and hence choose the action which is linked to the large regret with lower probability than the equivalent action is chosen by subjects in group B

This approach has both benefits and limitations when compared to the economist's approach. For example, a limitation is that it is not a direct test of whether the individual uses Regret Theory in order to make a decision, but rather a test of whether the experimental manipulation is itself linked to regret, under the prediction of Regret Theory that a difference in exposure to regret will lead to a difference in behaviour. The primary benefit, however, is that the experimental manipulation, which creates the distinction between group A and B, helps provide an insight into which components of a decision might be considered crucial for both the anticipation and experience of regret, and hence for predicting real world situations under which Regret Theory might be a good model to approximate human decision making behaviour.

For instance, consider a simple hypothetical experiment in which there are two players. Player 1 must correctly call the outcome of a coin flip, which is conducted by player 2. In treatment A, player 1 must call before player 2 flips the coin. In treatment B, player 2 flips the coin, observes the outcome in private, then provides "cheap talk" information to player 1, by saying "the coin was a head, but I might be lying", and then asks player 1 to call. Suppose in both cases that player 1 calls heads, but the coin was a tails, and hence player 1 loses. In treatment A, it is easy to imagine that player 1 might feel very little regret for making the wrong call, under the logic that "It was only ever a 50/50 shot, and it didn't make any difference what I called. The fact that I lost was as a result of bad luck, not a bad call". However in treatment B, player 1 may be expected to feel more regret as the construction of the experiment is such that the call became more a question of whether player 1 thought player 2 would lie to them or not, rather than simply the call of a coin. Player 1 may feel "I had a decision to make as to whether or not to believe player 2, and, as the uncertainty had already been resolved, I made a bad decision. I feel regret because I should have known player 2 would lie to me."

In such an experiment, Regret Theory would make no distinction between the regret experienced in treatments A and B, because the decision would be represented the same way, in both cases, as in the following payoff matrix

| | LANDS HEADS | LANDS TAILS |
|-----------------------|--------------------|--------------------|
| Subject chooses heads | Utility of Winning | Utility of Losing |
| Subject chooses tails | Utility of Losing | Utility of Winning |

⁷ this is often achieved by manipulating the amount of ex-post information available to subjects, or manipulating the magnitude of the consequences of the action

and hence the regret experienced, from making the wrong call, would be the same⁸ in both treatment groups. However, the evidence drawn from the experiment, and in particular the survey of the players afterward, would suggest that such concepts as “degree of responsibility for outcome” and “ability to rationalize the outcome” may play a significant role in the amount of regret experienced by the player, and hence in the amount they would anticipate in any Regret Theory style expected anticipated regret calculation.

Indeed the level of responsibility an individual feels they have for the outcome of a decision has been shown to be highly related to regret. Zeelenberg et al. [109], for example, summarise the results of four papers which aim to directly explore the relationship between regret and responsibility, which show that “...regret and responsibility are positively related” [109, p149] and “[a]lthough it may be possible to experience regret in the absence of responsibility...the evidence suggests that such cases are the exceptions rather than the rule.” [109, p152] These results present a significant challenge to Regret Theory, despite not being a direct empirical test of the theory itself, as the results suggest the effect of regret on an individual can be mitigated in a manner not captured by simply considering a payoff matrix form of the decision problem (as above). Indeed, similar results have been found experimentally by Ritov and Baron [78] in investigations of the so-called “omission bias”, which posits that regrets arising from *actions*, as opposed to *inactions*, are considerably worse in the eyes of the decision maker “...because acts tend to be seen as more causal than omissions, and blame, including self-blame and regret, depends on perceptions of causality⁹.” [78, p119] Again, Regret Theory makes no distinction between action and inaction (in essence, a frame of reference), and hence this experimental result suggests both a limitation and possible extension of the theoretical framework.

Furthermore, the link between responsibility and regret has been studied in applied consumer research as far back as 1992. Simonson [85] specifically analysed the problem facing consumers between buying based on “brand name” and buying based on “price”¹⁰, under the assumption that anticipated regret may play a significant role in the consumer’s decision;

“The amount of regret and responsibility that the consumer would feel in each situation is likely to depend on whether the consumer selected the well-known or the cheaper alternative. Specifically, it might be argued that the more expensive option is the safer bet and the norm, whereas the cheaper option is more of a gamble. If the consumer selected the more expensive alternative, and it failed, then the responsibility for the failure would rest on the manufacturer or the retailer rather than on the decision of the consumer. Conversely, if the consumer took a chance and chose the cheaper alternative and it failed, then the consumer might feel responsible for the failure and be more

⁸ that is, the regret from both calling heads and the coin landing tails, and calling tails and the coin landing heads, would be $R(\text{Utility from Losing} - \text{Utility from Winning})$ as given by the Loomes and Sugden version of Regret Theory

⁹ and, by extension, responsibility

¹⁰ for example, why would you consider spending £100 more on a Sony television than a Beko when they appear to be similarly in technological capability?

likely to regret the decision (“I should have known better”)”
Simonson [85, p107]

The results of Simonson’s study demonstrate that considerations of regret play a large role in consumer decisions and indicate “...that the person who selects the lesser-known brand and fails is seen as more responsible for the outcome.” [85, p115] and “...that manufacturers of better-know brands ... might increase their market shares if they can cause consumers to anticipate how they would feel if that made the wrong decision.” [85, p116]. These results show that not only is the anticipation of regret important for real-life consumer decisions (as hypothesised in Regret Theory), but that the *nuances of the process*, such as the role of responsibility, are incredibly important to understand for finding situations under which Regret Theory might provide a descriptive account of human decision making behaviour. These specific *characteristics* and *context* of the decisions being taken, and how they link to regret-based models of decision making, are further explored in Chapter 3.

As such, the psychologist’s approach to the experimental study of regret, whilst arguably not as rigorous as the economist’s approach, will ultimately provide a greater body of evidence for suggesting the future direction of regret-based theories. Such experiments, for example, have also resulted in findings about the temporal nature of regret (Gilovich and Medvec [28], Richard et al. [75]) and the effect of feedback on anticipated regret (Josephs et al. [45], Zeelenberg et al. [107], Zeelenberg and Beattie [105]), neither of which are explicitly incorporated into the typical Regret Theory framework. These findings have led to new, *process driven*, regret-based theories, which aim to incorporate the results of experimental research directly into the mathematics of the theories, resulting in the likes of Feedback-conditional Regret Theory (Humphrey [41]) and Decision Justification Theory (Connolly and Zeelenberg [18]).

2.2.3 *Experiential Regret Experiments*

As discussed in both the introduction and other chapters of this work, the primary focus of this research is the effect of experienced regret on future anticipated regret and, consequently, choice behaviour. The theoretical motivation for this research is provided in the first chapter, but, in addition to addressing the question theoretically, the above experimental approaches provide a basis upon which it is possible to design an experiment, or series of experiments, to address the question directly.

Designing such an experiment requires two main components. Firstly, the experiment must have a mechanism through which experienced regret can be created and experienced by the subjects in the experiment. In the simplest case, supposing there are two groups of experimental subjects - A and B - there must be a way to cause group A to experience more regret than group B. The second essential component is a method of measuring the subsequent relative attitudes, of the two treatment groups, towards regret (i.e. how they subsequently anticipate regret) after they have been exposed to the different levels of experienced regret. Both of these two components have significant technical, theoretical and experimental challenges associated with them, and a discussion of these will follow in later sections. However, to date, there have been a small number of experiments conducted which have attempted to

answer the question of the impact of experienced regret on subsequent choice.

Zeelenberg and Beattie [1997]

The first experiment which looked into the impact of experienced regret on subsequent behaviour was Experiment 3 of Zeelenberg and Beattie [105]. Experiments 1 and 2 in this paper had looked at the impact of expected feedback on subsequent behaviour (as discussed previously), but Experiment 3 took this one step further by looking at the impact of *experienced* feedback on subsequent behaviour. By extension, when the experienced feedback informed the participant that they would have done better had they made a different decision, this created *experienced* regret.

The precise experimental design used in this research was a version of the ultimatum game. In the game, each player must make a decision on how to split 100 Dutch Guilders between themselves (as the proposer) and another person in the game (as the responder). The responder then chooses whether to accept the split (in which case each person keeps their share) or reject the split (in which case, they each get nothing). In reality, every subject in the experiment was a proposer, and the responding was done by the experimenter, although subjects were unaware of this.

In generating the regret, subjects were split into two treatment groups. In the high regret group (HRG), subjects were told that their offer had been accepted, but the responder would have still accepted even if they had offered 10 Guilders less. In the low regret group (LRG), subjects were told that their offer had been accepted, but the responder would have still accepted even if they had offered 2 Guilders less. The idea here is that those subjects in the HRG were led to believe that they had left 10 Guilders on the table. That is, they could have made a decision (to offer 10 Guilders less) which would have improved their payoff by 10 Guilders¹¹. In contrast, those in the LRG were led to believe that that they left very little money on the table, in that they could have, at most, only offered 2 Guilders less and still had their offer accepted. As this was the only difference between the two experimental groups, it is reasonable to assume that those in the HRG experienced more regret than those in the LRG, as those in the HRG made a “bad decision” (they made an offer far in excess of the minimum acceptable offer) and those in the LRG made a “good decision” (they made an offer which was just above the minimum acceptable offer). This assumption is supported by subjective measurements of “experienced regret” which the participants were asked about after they received the results of the feedback¹².

There was then a second stage of the experiment, whereby each participant was asked to play the role of the proposer in the 100 Guilder ultimatum game again¹³. The hypothesis of Zeelenberg and Beattie was

¹¹ and, by extension, they could also have offered 9 Guilders less and received 9 Guilders more, 8 Guilders less and received 8 Guilders more, etc., but not 11 Guilders less, as that offer would have been rejected.

¹² as stated in the paper “[t]his was done by presenting the participants on the computer screen with the feedback and asking them to indicate, on 7-point scales, how much regret they experienced and how good they thought their offer was in retrospect.” [105, p73]

¹³ in round 2, each participant was informed that they were now playing against a different responder from round 1 (again, in reality, they were playing against the experimenter), and that the average minimum acceptable offer in round 1 was 22 Guilders. This was the same for both groups.

“...that participants in the [HRG] would lower their offers more than participants in the [LRG]... because of the regret they experienced over the first offer”, thus describing the hypothesised impact of experienced regret on subsequent choice behaviour. Their reasoning for this hypothesis can be drawn from the results of Experiment 2, where [Zeelenberg and Beattie](#) conclude that “[t]his behaviour [of making lower offers] represents regret aversion because lower offers result in less regret if accepted.” Thus, combining the hypothesis and the reasoning together implies that their belief is that the experience of regret, from stage one of the game, works to *increase* regret aversion in stage two of the game.

Indeed, the offers in stage two were lower for those in the HRG (mean of 26.34 Guilders) compared to those in the LRG (mean of 34.69 Guilders, significant difference at the 1% level), confirming their hypothesis that the impact of experienced regret is to make the subjects *more* regret averse, using the implication of the results from the ultimatum game in Experiment 2.

Creyer and Ross [1999]

The second experiment on the impact of experienced regret was conducted in 1999 by [Creyer and Ross](#), who, like [Zeelenberg and Beattie](#), seek “...to [measure] the experience of regret and its affects on subsequent behaviors.” [21, p380] They achieve this through “...varying levels of outcome feedback in order to measure the experience of regret and examine its effects on subsequent decision making.” [21, p380].

Their specific experimental approach¹⁴ shares many similarities with [Zeelenberg and Beattie](#). In this experiment, the subjects must propose an offer “...in an attempt to win an order to produce 1000 silicon chips for a hypothetical electronics firm... [choosing] from among a set of 11 bids, ranging from \$80,000 to \$130,000¹⁵.” [21, p387]. Again, similar to the [Zeelenberg and Beattie](#) experiment, there was a chance that the offer could be rejected (and, hence, earn the salesperson no commission); this time framed as the fact that there would also be competing offers from other firms in the marketplace, with the firm offering the lowest bid winning the contract. Additionally, the bonus earned by the salesperson, should the contract be won, depended positively on the bid made. “[T]he bonus ranged from \$5000 to \$15,000, for the \$80,000 and \$130,000 bids respectively.” [21, p387]

Immediately, one can see the theoretical equivalence between the two experiments. Both require the subjects to make a decision under uncertainty, where a better offer to the other player¹⁶ leads to a higher chance of acceptance, and obtaining a strictly positive payoff, but at the expense of an overall lower payoff to the subject. Also, as in the [Zeelenberg and Beattie](#) experiment, there was a second round of the game, which commenced after participants had received feedback on the outcome of their offer from the first stage. Two of the treatment groups used by [Creyer and Ross](#) were directly equivalent to the HRG and LRG in [Zeelenberg and Beattie](#)’s experiment. Their HRG was framed such that the subject had won the contract, but the next nearest

¹⁴ specifically, Experiment 2 in their paper

¹⁵ as one might expect, given the scale of the numbers, the actual payments made to the participants in this experiment were scaled down to payoffs for participants of between \$1.50 and \$4.50.

¹⁶ it is worth noting that in the [Zeelenberg and Beattie](#) experiment, a “better offer” corresponds to a “higher offer” to the other person, but in the [Creyer and Ross](#) experiment, a “better offer” corresponds to a “lower price” being offered to the electronics firm.

offer of a competitor was substantially¹⁷ higher, indicating that the subject could have offered a significantly higher price, increasing their own bonus payoff, and still won the contract. The LRG, by comparison, was framed such that the next nearest offer of a competitor was only minimally above, indicating that the subject could have only marginally increased their own bonus payoff by making a different offer, and hence, overall, they made a good decision.

In the second selling scenario, where [Creyer and Ross](#) measure the effects of first stage regret on subsequent choice, the subjects faced an essentially equivalent contract-offer scenario, but with a different range of payoffs and bonuses¹⁸ to the first stage. Their hypothesis is that “...the decision maker experiencing regret [from the previous stage] should be more likely to choose the option [in the second stage] which maximises his or her chances of a positive outcome in a subsequent decision, even if that option provides a lower pay-off.” [[21](#), p386/7]¹⁹. Indeed, their results show that “... the post-manipulation [second stage] bid was lower (i.e., the bid had a higher probability of acceptance) when regret was experienced, compared to when little or no regret was experienced” [[21](#), p390] and that the difference in second stage bids between the HRG and LRG was significant at the 5% level. Whilst their written conclusion is not framed in terms of a change in regret *aversion* in the second stage, the implication from their hypothesis is that the decision maker who experiences regret becomes *more* regret averse, wanting to maximise their chances of a positive outcome in the second stage, hence lowering their offer.

Raeva and van Dijk [[2009](#)]

A third experiment which looks directly at the effects of experienced regret on subsequent choice is by Raeva and van Dijk [[73](#)], who move away from the framework of [Zeelenberg and Beattie](#) and [Creyer and Ross](#), instead studying the impact “...of experienced regret on a different subsequent choice, i.e., on a subsequent choice where no direct mapping between the choices is present, but only a remote resemblance.” [[73](#), p4] Indeed, the channel by which they propose that such an effect would take place is also different from the previous two experiments. They “...assume that the mechanism by which experienced regret influences anticipated regret is linked to influencing the subjective probability of experiencing regret” [[73](#), p4], which stands in contrast to the previous experiments by, firstly, very explicitly stating the component of anticipated regret which will experience a change as a result of experienced regret, and, secondly, emphasising the role of probability weighting in the anticipated regret calculation at the expense of the regret function itself²⁰.

¹⁷ indeed, the HRG in this experiment is called the “Substantial” regret group for this reason.

¹⁸ the eleven bids in the second stage ranged from \$100,000 to \$162,500, and the bonuses ranged from \$6000 to \$18,000.

¹⁹ the difference between this hypothesis, and the hypothesis of [Zeelenberg and Beattie](#), will be widely explored in subsequent sections

²⁰ as the anticipated regret calculation, used in Regret Theory, has both a regret function and associated probabilities (i.e. “I think if I choose option A, I have a 30% chance of experiencing high regret”), there are two channels which changes due to the experience of regret can operate. Most work (as is implicitly assumed in [Zeelenberg and Beattie](#) [[105](#)] and [Creyer and Ross](#) [[21](#)]) assumes that this channel is the regret function, but [Raeva and van Dijk](#) explicitly assume the other.

As in the previous experiments, this experiment has two stages. In the first stage, the regret is generated by a complete feedback vs partial feedback manipulation. All subjects could choose between two 50/50 gambles, both of which had a 50% chance of winning €1 and a 50% chance of winning €10. Those in the regret condition (RC) chose one of the two gambles, won the €1 prize, and then learned that, had they chosen the other door, they would have won a €10 prize. Those in the no-regret condition (NRC) chose one of the two gambles, won the €1 prize, but did not learn what they would have received had they chosen the other door²¹. This is a slightly different approach to the other experiments, which compare “high vs low regret” as opposed to “positive vs absent regret”, but the implications should be similar, under the assumption that there is simply *more* regret being experienced in one group compared to the other.

The second stage, however, is a fundamentally different task, where participants are asked to give a certainty equivalent value for a gamble which has a 50% chance of winning €100 and a 50% chance of winning €1000²². Their hypothesis is that “...the direction of adjustment is downwards²³ when the subjective probability of regret is changed due to experienced regret²⁴.” [73, p5]. However, they also manipulate the second stage by using both complete and partial feedback conditions, such that there are now four treatment groups overall (RC-complete, NRC-complete, RC-partial, NRC-partial). The idea is that if the impact of the experience of regret works through the channel of subsequent anticipated regret, then the impact will be different in situations where regret can be avoided compared to situations where regret can be once again experienced²⁵.

The results of this experiment can be summarised by the mean certainty equivalent values of each of the four groups, as presented in Figure 38.

The results, in contrast to the hypotheses, suggest that the only situation in which there was a significant deviation in the reported certainty equivalents of the four groups, was in the condition where there was both regret experienced in the first stage, and there was

- 21 the implication is that both groups will experience disappointment (as they both “lost” their chosen gamble) but only the group which saw the better result from the other gamble will experience regret. It is plausible, however, that, once the subjects in the NRC observe the fact that they only received €1 from their gamble, they might wish they had chosen the other, under the assumption that they wouldn’t have done any worse (they would have obtained the same result with a 50% probability), but could have done better (they could have gained an extra €9 with a 50% probability). This issue of “counterfactual thinking”, leading to regret, has been widely discussed in much of the complete vs partial feedback literature, without any general conclusion as to either the presence or level of regret under such circumstances.
- 22 the question is framed in terms of the popular TV show “Deal or No Deal” where a contestant has reached the last stage with only the €100 and €1000 prizes remaining, and must decide what is the minimum offer from the banker that they would accept.
- 23 i.e. those in the RC will report a lower certainty equivalent than those in the NRC
- 24 their assumption is that the experience of regret from the first stage, as a result of making the wrong decision in a 50/50 gamble, implies that they will anticipate a higher probability of losing the 50/50 gamble again, and hence experiencing the regret of “choosing the wrong gamble/box” again in stage 2. This implies that the gamble is less attractive in stage 2, when compared to the certainty equivalent, and hence they would accept a lower value for certain to restore equivalence.
- 25 here, the hypothesis is that, if the lottery is still resolved after the certainty equivalent is taken (in the complete feedback condition), then the certainty equivalent is less attractive compared to the partial feedback condition, as it also exposes the decision maker to the possibility of regret. Hence, the decision maker will require a higher certainty equivalent value, to be indifferent between the money-for-sure and the gamble, when complete feedback is provided.




Fig. 3 Mean Certainty equivalent per condition. (Error Bar: 95% confidence interval). The white square indicates the mean certainty equivalent for non-regret feedback, the black square indicates the mean certainty equivalent for regret feedback. The solid line connects the mean certainty equivalents of partial and complete feedback conditions for the non-regret feedback on the first; the dashed line connects the mean certainty equivalents of partial and complete feedback conditions for regret feedback on the first task.

Figure 38: Figure 3 from Raeva and van Dijk [73]

complete feedback provided about the results of the second stage. However, the direction of this one significant result is consistent with their explanation of how the experience of regret would influence the certainty equivalent values when comparing the non-regret to the regret conditions in the complete feedback case, as they did predict that it would be lower in the regret condition. As such, their proposed mechanism, through which the experience of regret would increase sensitivity to the possible experience of a “similar” regret again, using an increase in the subjective probability of that regret, is lent support by the result. Though the other results run contrary to the predictions of Raeva and van Dijk’s formulation of regret theory²⁶, there is an indication that the presence of anticipatory regret, which is the only difference between the partial and complete feedback conditions in stage two, can be influenced by regret experienced in a previous stage. More work would need to be done, however, to conclusively determine that this is due to a change in subjective probabilities associated with a specific type of regret, rather than a change in the regret function itself.

Coricelli et al. [2005]

One last experiment, which looks at the impact of experienced regret on subsequent choice behaviour, is by Coricelli et al. [19], who use neuroimaging approaches to additionally investigate the activity of the brain when decisions under uncertainty are made. In their experimental design, subjects are asked to make a long sequence of choices, which all involve a decision between two independently resolved gambles. The gambles are presented as “spinning wheels”, whereby a fraction of the circle corresponds to an outcome, and the probability of that outcome being randomly picked is the fraction of the circle it takes up. At each stage, subjects view both wheels (corresponding to the two gambles), and, again, there are two conditions of partial and complete feedback. If the subject is making a decision in a complete feedback trial, they will select the gamble, and then view the outcome of both the selected and unselected gamble. If they are making a decision in a partial feedback trial, they will view the outcome of only the gamble they selected.

By arranging this experiment as a long series of gambles, rather than a two-stage, two group experiment, there is a benefit of being able to

²⁶ which is based, qualitatively, on Bell [6]

(b) Subjects' choice behavior as a function of anticipated disappointment (d), maximization of the expected value (e) and anticipated regret (r) in the 'complete choose' condition.

| Variable name | Coefficient | Standard error | Z | P |
|---------------|-------------|----------------|------|-------|
| Constant | 0.06958 | 0.09706 | 0.72 | 0.473 |
| d | 0.00211 | 0.00137 | 1.62 | 0.106 |
| r | 0.00463 | 0.00114 | 4.03 | 0.000 |
| e | 0.00726 | 0.00184 | 3.94 | 0.000 |

Number of subjects = 15; number of observations = 720. Log likelihood = -331.1425; Wald $\chi^2(3) = 205.82$, Prob > $\chi^2 = 0.000$. The dependent variable 'choice' is equal to 1 if subject chose gamble 1 and is equal to 0 if subject chose gamble 2.

Figure 39: Table 2 from Coricelli et al. [19]

generate a much wider range of experienced regret throughout the experiment, but at the expense of direct control of the variables which are being changed. As such, a regression analysis the most appropriate approach to generating useful results, and Coricelli et al. construct a panel logit model, with individual random effects, and initially specify the model with the choice between gambles as a function of anticipated disappointment, anticipated regret and expected value²⁷. The results are given reproduced in Figure 39.

The results indicate that their formulation of anticipated regret is a significant determinant of individual choice in this experiment, even when controlling for the expected value of the respective gambles.

However, in addition to measuring the effect of anticipated regret on choice, they also seek to investigate the effect of experienced of regret on choice. This is constructed in two ways. Firstly, by separating the sequence of choices into the first, middle and final thirds, they show that the "...proportion of regret-avoiding choices increased over time with the cumulative effect of the experience of regret." [19, p1259] Secondly, by using neuroimaging techniques, they show that the "...[e]xperience of regret in a previous period profoundly influenced choice-related activity, enhancing responses in right dorsolateral prefrontal (DLPFC), right lateral OFC, and inferior parietal lobule" [19, p1259] and that cumulative regret, calculated as the average payoff of all unselected previous gambles minus the average realized payoff of all previous gambles, "...at the time of decision making involved similar anatomical regions...as that elicited by regret at the time of outcome feedback." [19, p1259]. Their interpretation of these various results is that "...the experience of regret has a major impact on the process of choice that is expressed at two levels, with a net result of biasing subjects to forgo choices that might lead to future experience of this highly negative emotion" [19, p1260], which, again, can be understood to imply that the experience of regret works in such as way as to make the decision maker *more* regret averse in the future.

²⁷ they "...parameterized regret as the absolute value of the difference between the lowest and highest outcome across gambles." [19, p1258]

Conclusions from existing experiential regret literature

All four experiments indicate that there is a significant relationship between past experiences of regret, and subsequent regret aversion. Indeed, all four experiments arrive at a similar conclusion; that the experience of regret from a previous decision (or previous decisions), will make a decision maker *more* sensitive to the potential of regret being experienced again in the future. Or, in the language of traditional Regret Theory, they become *more* regret averse.

However, despite all four papers reaching the same conclusion, all four papers' conclusions are based on *different* assumptions about the process through which the experimental subjects' behaviour is being influenced. Indeed, these assumptions are so critical to the conclusions being drawn, that simply switching the assumptions of one paper with another could lead to a completely opposite set of conclusions being drawn from this experimental literature. As such, it could be argued that these experimental results are not *robust*, as it is possible to make entirely plausible, and psychologically justifiable, alternative assumptions about the individuals in the experiments, and draw fundamentally different conclusions from the data they obtain. It is to this discussion of the robustness of the experimental results of each paper that we turn next.

2.3 LIMITATIONS OF EXISTING EXPERIMENTAL LITERATURE AND EXPERIMENTAL RESULTS

To begin a discussion of the limitations of the existing experimental literature, we must first explore the assumptions, implications and limitations of the traditional regret theories which the experiments are based on. In all formulations of Regret Theory, there are both explicit and implicit assumptions, which are, to differing extents, justified through psychological evidence. It is to these assumptions which we turn first.

2.3.1 *Assumptions of Regret Theory*

The original theoretical expositions of Regret Theory, from Loomes and Sugden [58] and Bell [5], both contain broadly similar assumptions relating to the regret components of the theory, and, as such, it is sufficient to discuss these with respect to just one of the papers (in this case, Loomes and Sugden [58]) in order to communicate the essence of their mathematical assumptions in terms of psychological intuition.

The assumptions made about regret in the Loomes and Sugden paper can be broadly split into three categories. These are

- Assumptions which are wholly uncontroversial, and necessary for the basic structure of the theory.
- Assumptions which are explicitly discussed, with regards to both their intuitive and mathematical nature, but are debatable, and consequently alter the predictions of the theory depending on whether they are included or not.
- Assumptions which are implicitly made, in terms of their psychological underpinnings, but appear uncontroversial when written down mathematically. These are assumptions which, essentially,

appear to fall into the first category, but merit discussion in the second category, given the consequence they have for, for example, experimental results.

As an example of an assumption which fits into the first category is that, holding constant the outcome which did result from your decision, "...the more pleasurable the consequence that might have been, the more regret...is experienced." [58, p808] This assumption implies that, if you experience regret as a result of a decision, then the regret you experience would get worse as you increase the "utility" or "pleasure" of the outcome which would have resulted, under the same state of the world, had you made a different decision. For example, if you were betting on a horse race, which had only two horses, the "favourite" and the "outsider", and you chose to bet £10 on the favourite, but the outsider was the horse that won, then the regret you would experience would increase as the odds of the "outsider" got longer. You would experience more regret if the outsider had odds of 100-1, than if it had odds of 20-1, than if it had odds of 5-1.

An example of an assumption which falls into the second category is that a function $Q(\xi) = \xi + R(\xi) - R(-\xi)$, where $R(\cdot)$ is the regret function defined over the chosen and unchosen outcomes, "...is convex for all positive values of ξ " [58, p810]. As Loomes and Sugden discuss, "...there seems to be no *a priori* reason for preferring... [this] assumption to the others" [58, p810], which are also mathematically stated in the paper, but "...a choice between them should be made mainly on the basis of empirical evidence." [58, p810] Indeed, the paper explores existing evidence relating to decision under uncertainty, and finds that "[r]egret theory yields a wide range of firm predictions that are supported by experimental evidence" under the assumption "...that $Q(\xi)$ is convex for all $\xi > 0$." [58, p817]. As such, the "convexity assumption", as it was recast in Loomes and Sugden [60], appears to have large amounts of experimental evidence supporting its existence. Should an individual wish to challenge the assumption, they would need to demonstrate why the existing empirical evidence is misleading, and present new evidence to support an alternative assumption. Given it is unlikely that evidence will be found to overturn decades of research which all point in the same direction, it is fair to say that this assumption is robust, and safe to include when making conclusions from experimental results.

It is, however, the third category which produces the more interesting assumptions, both in terms of their importance for the theoretical predictions of regret theory, but also the conclusions which can be drawn from experimental results, where the experiment was designed around the basic principles of regret theory. The primary assumption to be discussed, which will lead to a detailed explanation of the limitations of the existing experimental literature and their conclusions, is implicitly contained in the formulation of the regret function by Loomes and Sugden, but also by other similar formulations of regret theory. As stated before, the "regret-rejoice function"²⁸ of Loomes and Sugden "...assigns a real-valued index to every possible increment or decrement of choiceless utility" [58, p809] and is written as $R(c_{ij} - c_{kj})$ where c_{ij} is the utility the person would receive from simply being gifted the outcome of action i , in state of the world j , without having been asked to make a choice²⁹. By calculating regret as solely dependent on the

²⁸ referred to throughout, for simplicity, as the regret function

²⁹ the use of the letter 'c' refers to the idea of 'choiceless' utility

difference between the utility of what you received and what you missed out on, there are important implications for even very simple decision making problems, which involve a choice between only a few options.

For example, consider the simplest possible decision under uncertainty problem, which involves two actions and two states of the world.

| | STATE OF WORLD X | STATE OF WORLD Y |
|----------|------------------|------------------|
| Option A | 5 | 5 |
| Option B | 0 | 10 |

where the numbers represent the specific choiceless utility associated with the outcome of, for example, option A in state of the world x. Suppose, in addition, that each state of the world is equally likely, and so $p = 0.5$ for each. In which case, the original formulation of regret theory informs us that the decision maker will weakly prefer option A to option B if, and only if,

$$1/2 (5 - 0 + R(5 - 0) - R(0 - 5)) + 1/2 (5 - 10 + R(5 - 10) - R(10 - 5)) \geq 0$$

$$\Rightarrow 1/2 (5 - 0 + R(5) - R(-5)) + 1/2 (5 - 10 + R(-5) - R(5)) \geq 0$$

From analysing the last line of the mathematics, two things stand out. The first is that the solution has simply degenerated to the Expected Utility solution of indifference, which is predictable given the numbers are constructed such that the expected utility of option B is equal to the “full insurance” of option A. Secondly, however, we can notice that the regret terms also cancel out in the mathematics, which implies that considerations of anticipated regret can’t influence the decision, one way or another, under this formulation. This implies, for example, that an experiment which found that significantly more people preferred option A to option B, would rule out anticipated regret as a possible explanation, on the grounds that the difference in choiceless utility between option A and option B in state of the world x is equivalent to the difference in utility between option B and option A in state of the world y. Even though there are two possible sources of regret in this experiment, they are deemed to cancel each other out for decision making purposes³⁰.

Another non-trivial consequence of this formulation of Regret Theory, contained within the above example, is that there will always be, in non-trivial decision making problems, at least *two* sources of potential

³⁰ the purpose of this example, as will become clear, is not to argue with the theoretical assumptions of Loomes and Sugden [58], but rather to point out that there are dangers to be had, and mistakes to be made, when designing experiments based upon the theoretical models, given they use assumptions, often implicit, which have important consequences for the interpretation of the experimental results. Indeed, the updated version of Regret Theory from Loomes and Sugden [60], places significantly less restrictions on the form of the regret function, and still derives broadly the same predictions as the original theory. Yet, more commonly, it is the “difference” or “utility gap” functional form of the regret function which is imposed when explanations for experimental results are given. Indeed, even the updated version of Regret Theory assumes that regret is simply a function of the utility obtained and the utility foregone, which implies that two decision under uncertainty problems which are mathematically equivalent, in terms of utilities and probabilities, should yield the same choices. As shown in Chapter 3, however, this is not necessarily the case.

regret, acting in *opposite* directions on the decision maker. This can be shown by considering the most basic form of the problem of decision making under uncertainty; the case where there is both a decision to be made (at least two options) and at least some uncertainty (at least two states of the world). The simplest version of this problem is, therefore, the 2x2 version, where there are two actions and two states of the world. Consider the following payoff matrix

| | STATE OF WORLD X | STATE OF WORLD Y |
|----------|------------------|------------------|
| Option A | j | k |
| Option B | l | m |

where j, k, l, and m refer to the choiceless utilities of the outcomes in the payoff matrix. For the purposes of making decision under uncertainty, we are interested in comparing option A to option B, and, hence j & k to l & m. Working on a case-by-case basis, there are several possible scenarios.

Cases of Indifference

Clearly, when $j=l$ and $k=m$, we have absolute indifference between the two options, as they are equivalent, so there is no decision to be made. Hence, the case is not interesting.

Cases of Stochastic Dominance

Similarly, when one option stochastically dominates the other, there is also little interest in the problem, as the dominant option is both the only choice for maximising utility and avoiding regret. Such a case occurs when

1. $l > j$ & $m \geq k$
2. $m > k$ & $l = j$
3. $j > l$ & $k \geq m$
4. $k > m$ & $j = l$

Non-trivial Decisions

The only two cases left are the non-trivial decisions when either

1. $l > j$ & $k > m$ or
2. $m > k$ & $j > l$

In case 1, notice that, focusing on state of the world x, the anticipation of regret would favour option B over A, but, focusing on state of the world y, the anticipation of regret would favour option A over B. Similarly, in case 2, focusing on state of the world x would favour option A over B, but focusing on state of the world y would favour option B over A. Hence, in both cases, the decision maker must consider two different experiences of regret, which act in opposite directions on the decision maker. In addition, the two competing directions of regret will always exist in any non-trivial decision under uncertainty problem which has more than two actions or more than two states of the world, as you can

always isolate the aspect of the decision which makes it non-trivial, by considering the two “best” actions, and consider the states of the world (for which there will be a minimum of two) where there is a trade-off between the two actions. As such, any discussion referring to a decision-making problem under uncertainty, whether in an experiment or otherwise, which refers to anticipated regret, must both *identify* and *analyse* more than one source of anticipated regret to give a complete picture of the process through which regret can affect a choice between actions.

Implications of the assumptions

Combining these two oft-overlooked implications of Regret Theory leads to the construction of examples where the discussion of an experimental result, should the example be analysed in the laboratory, depends very heavily on how these assumptions are incorporated into the design. For example, consider a betting problem, where there are two possible bets which can be made (a High and Low bet), but also the option to abstain from the gamble altogether. Further consider that there are only two states of the world, which correspond to “winning” and “losing” if you made the bet, but have no consequence if you abstained. This problem can be represented in the following payoff matrix³¹

| | WINNING | LOSING |
|----------|---------|--------|
| Abstain | 0 | 0 |
| Low Bet | L | -L |
| High Bet | H | -H |

Consider that a person who chooses the Low Bet is exposing themselves to two different forms of regret. Should the person win the bet, they would receive a payoff of L, but experience a degree of regret from not choosing the High Bet instead. Under the original form of Regret Theory, this regret would be a function of the difference in choiceless utility between H and L. Alternatively, should they lose the bet, they would receive a negative payoff of -L, but also experience a degree of regret from not choosing to Abstain. Again, under original Regret Theory, this would be a function of the difference in choiceless utility between 0 and -L.

As such, if the above situation was observed in an experimental setting, where an person chose the Low Bet, it would not be possible to attribute the choice of the Low Bet to anticipated regret, without knowing more details about the choiceless utility of H, L, -L and 0³², and the nature of the regret function, or, in other words, the individual’s attitude towards regret. A typical argument for why the Low Bet was chosen could be:

The individual chose the Low Bet, because they considered the state of the world where they would lose, and realised that the regret from knowing they should have abstained, instead of making the low bet³³, was *less* than

³¹ $0 < L < H$

³² though it’s not unreasonable to assume that the choiceless utility of 0, i.e. no deviation from the status quo, is 0

³³ $R(-L - 0)$

the regret which would have resulted from taking the High Bet, losing, and realising they should have abstained³⁴. This desire to avoid the large regret is *anticipated regret aversion*, and so we attribute the decision to make the Low Bet, in preference to the High Bet, to it.

Indeed, it is not sensible or logical to posit that the situation where the individual takes the High Bet and loses generates *less* regret than the situation where the individual takes the Low Bet and loses³⁵, so at least some of the above argument makes perfect sense. However, as discussed earlier, there is always more than one source of regret acting on any decision under uncertainty. In this case, you could make an alternative argument regarding the individual's choice of the Low Bet as follows:

The choice of the Low Bet, in preference to the High Bet, makes little sense when you consider the effect of anticipated regret aversion on the decision. Considering the state of the world where the individual wins, they expose themselves to potential regret if they choose the Low Bet³⁶, but can prevent this from occurring by choosing the High Bet, knowing it would be the best choice if they win, and hence giving zero regret³⁷. Assuming the individual is regret averse, therefore, the effect of this anticipated regret aversion should, instead, push them towards the High Bet. As such, anticipated regret aversion is not an explanation for the decision in this case.

The second explanation applies the same logic as the first, but simply to the other potential source of regret which occurs when the individual wins instead of loses. Again, it is inarguable that there is more potential regret, in this state of the world, from choosing the Low Bet as opposed to the High Bet, so that part of the explanation holds up. Yet, despite both arguments appearing to be valid, they are clearly contradictory. One states that anticipated regret aversion is a possible explanation for the choice, and the other says that this is not the case.

Clearly, the problem arises because we are considering each state of the world in isolation. The fact that this is a decision made under uncertainty necessarily implies that there is a positive probability of both winning *and* losing occurring, and so to completely neglect one explanation, in favour of the other, can only be done if you can demonstrate that the expected regret, consisting of both magnitude and probability of occurring, from one state dominates the expected regret from the other state in the decision making process of the individual. It is not sufficient to simply observe the Low Bet being chosen, and pick one of the two explanations for the choice, without justification as to why one particular source of expected regret is dominant.

³⁴ $R(-H - 0)$

³⁵ another way to think of this point is that, even in the situation where the individual takes the Low Bet and loses, they experience some *rejoicing* from realising that they would have been worse off had they taken the High Bet. In the case where they took the High Bet and lost, this rejoicing does not exist.

³⁶ $R(L - H)$

³⁷ $R(0)$

From Discrete to Continuous Choice

The above example demonstrates how difficult it is to use anticipated regret aversion as an ex-post explanation for choice behaviour in discrete choice problems. This framework can then be extended to the more common real-world situation of problems of continuous choice under uncertainty. This framework naturally extends the above discrete choice gambling example, to one where an individual must decide how *much* money they wish to gamble.

In such a situation, the two sides of regret are even more obvious than they are in a discrete choice case. For any positive amount of money, x , that you choose to gamble, if you win the gamble, there is regret from thinking that you should have bet *more* than x , and if you lose, there is regret from thinking that you should have bet *less* than x . This duality of regret makes assertions such as “ x was the regret-minimising choice” or “anticipated regret aversion caused the individual to gamble x ” incredibly difficult to justify, as doing so requires explicit knowledge about the regret function of the individual in both the winning and losing state of the world. For example, if you argue that the individual bet £0, because this protected them from regret in the situation where they lost, and hence was regret-minimising, then I could counter that the individual *would* experience regret in the state of the world where they won (knowing that any positive bet would have been superior) and challenge you to demonstrate *why* the regret in the losing state dominates the regret in the winning state for this individual, and hence makes £0 the regret minimising choice.

Extending the problem further takes us to a world where not only is there a continuum of actions, but also a continuum of states of the world. Consider a situation where an individual must complete a task which takes an unknown amount of time t to complete. In addition, before the task commences, the individual must also make a prediction p about how long the task takes to do, in order to hire out a work space. There is, however, a cost associated with a misprediction, which is increasing in the extent to which the prediction is wrong. If the individual under-predicts the task completion time, or if $p < t$, then they will have to pay a higher rate³⁸ for the additional work space time needed to complete the task. If the individual over-predicts the task completion time, or if $p > t$, then the individual will have paid too much for the work space, given it was completed in a shorter than anticipated time. Clearly, regret will be zero, only if the individual manages to complete the task in exactly the amount of time they predicted, and they will neither pay the high rate nor over-book the work space. However, the *regret-minimising prediction*, p , is not necessarily the expected completion time. If the individual fears one of the two sources of regret much more than the other (say, the increase in rate, for exceeding the work space booking, is very high), then it may be worth the individual being *pessimistic* (higher p) in their prediction of the task completion time, so as to make the probability, of experiencing regret from knowing they could have avoided the high rate, by stating a higher value of p , much smaller. Equivalently, if the individual is more worried about paying too much up front for the work space, and exposing themselves to the possible regret of the task taking a shorter amount of time to complete than expected and knowing they could have saved money by stating a lower

³⁸ that is, for booking the additional work space at short notice, the marginal cost is higher than it would have been had they booked the correct amount of time in advance

p , then they would submit an *optimistic* (lower) prediction for p to reduce this chance. Again the two sources of regret act in opposite directions on the choice, and it is not possible to say that the choice of p which was ultimately made was regret-minimising without both acknowledging and detailing the magnitude of impact of both sources on the decision which was made.

This world of both continuous choice and continuous states of the world, incorporating regret, need not even have financial costs as the source of potential regret. In a situation where reputation is involved, so making accurate predictions is a source of pride and/or a future indication of quality, then it need not necessarily be the case that over and under predictions generate the same feelings of regret, even from the same *magnitude* of misprediction. An experiment designed to explore this idea, and give data to show the heterogeneity in the impact of regret on decision making, will be explored later. In addition, such a framework can be easily applied to more commonplace economics-based models of decision making to give new theoretical results and ideas. Filiz-Ozbay and Ozbay [24] apply this very idea to the world of auctions, giving both theoretical predictions and experimental evidence as to the role that different sources of regret can have on individual decision making. In the case of the auction, this corresponds to the amount an individual wishes to bid for an object (the “prediction” in the language of the above example) given the uncertainty as to what the highest bid of the other bidders may be (the continuum of states of the world).

2.3.2 *An identification problem*

The argument above has a natural analogy in the traditional economics world; that of the identification problem, for example as seen in estimation of the price elasticity of demand.

For those unfamiliar with the problem, a sensible approach to estimating the price elasticity of demand would appear to be observing a market over time and plotting the market price against the quantity demanded at different time intervals. By estimating the slope of a line of best fit through these points, you would appear to have a good estimate of the elasticity.

However, whilst you certainly directly observed the demand and the market prices, there was an element of the equation which was not observed; that of supply. As the market price is determined jointly by demand and supply, any estimate of the price elasticity of demand which does not account or control for the impact of supply (and indeed the wide range of variables which determine market supply) on price will likely be incorrect.

The same identification problem occurs when attempting to measure the effect of experienced regret on anticipated regret. Whilst there are certain aspects of the problem you can directly observe (for example, the action which is chosen such as to minimise regret, and perhaps even the magnitude of regret which was experienced, through neuroscience techniques), there are parts of the equation which will jointly determine the action chosen, but are largely unobservable. In the regret case, this is the anticipated regret associated with each action on offer as part of the choice. As explained, because there is no way to isolate just *one* of the anticipated regrets, and the choice made is the result of comparing at

least two unobservable anticipations of regret, then directly observing a regret-minimising action does not give you the causality between experienced and anticipated regret. In the same way, observing both price and demand does not necessarily give you the causality from one to the other because of the unobservable factors at play.

If there is an analogy between the regret problem and the identification problem of demand and supply, the natural question is “can techniques developed to overcome the identification problem also be used in this situation?” Two common techniques are the use of experiments and instrumental variables.

In the world of demand curve estimation, a pricing experiment may be used by a company to vary the price of a good and directly observe the effect on demand. Under appropriate conditions (for example, randomisation of where and how the prices are varied) the unobservable factors can be assumed to vary only due to random noise, and hence an unbiased estimate of the price elasticity can be found. However, as we have seen, using an experiment does not solve the problem in the regret world, as it is not simply a case of trying to eliminate sources of bias. The sign of the effect itself is in question, with both options being theoretically plausible, unlike demand elasticity which should, under normal conditions, always be negative.

The instrumental variables approach has similar issues in its application to regret. Typically, in demand estimation, you are looking for instruments which would help explain the movement in price but not a movement in demand (for example, weather patterns may dictate costs in the supply chain but not affect the demand for a good). By first calculating a “true price”, which subtracts the effects on price of these variables, you can reduce the error in estimating the price elasticity by utilising the true price in the demand equation instead. This approach does not work in the regret case, however, because there is no instrument which affects the anticipatory regret to one option but not another. The mechanism through which anticipatory regret affects actions cannot be controlled and isolated using the instrumental variables approach.

2.3.3 *Application of the assumptions*

It is first worth noting that this is certainly not the first time that the two sides of regret have been pointed out in a discussion of regret theory. In their 1997 experiments, [Zeelenberg and Beattie](#) mention that “...proposers [in the ultimatum game] can regret two things, offering too little money when the offer is rejected, and offering too much when the offer is accepted”[105, p68] and even attempt a discussion about which may be the bigger factor on individual decisions in the ultimatum game. They hypothesise that, typically, “...the regret about offering too much money is generally less severe than regret about offering too little money”[105, p68] but concede that receiving feedback on the minimal acceptable offer³⁹ “...can make regret about an offer that is too high more severe because it points out exactly how much less the proposer could have offered.”[105, p69] Also, as previously noted, [Filiz-Ozbay and Ozbay](#) make extensive use of the two sides of regret in their work on auctions, stating that “...if the bidders anticipate that they are going

³⁹ the equivalent of learning t , the specific state of the world, in the example of both continuous actions and continuous states presented earlier




Figure 40: Expected Regret from acceptance

to feel winner regret, they will shade their bids. In contrast, if their anticipation is loser regret, they will overbid”[24, p1407].

However, despite the acknowledged existence of the two sides of regret in several papers, the application of this knowledge to drawing conclusions from empirical and experimental results has been lacking. It is to this application, on the existing literature, which we turn next.

Application to Zeelenberg and Beattie [1997]

As previously discussed, Zeelenberg and Beattie run a two stage ultimatum game, concluding that the lower proposals seen in stage two from the group which experienced high regret in stage one [HRG] compared to the group which experienced low regret in stage one, can be interpreted as the experience of more regret causing an *increase* in regret aversion of the participants in the HRG as “[t]his behaviour [of making lower offers] represents regret aversion because lower offers result in less regret if accepted.”[105, p72].

However, as explained earlier, the fact that there are two sources of regret in such a decision making problem, means that it is not sufficient to give an explanation about regret which focuses on just one source; in this case the regret which results from making too high a proposal, if the proposal is accepted. Given this specific result from the experiment, there is another potential explanation which involves regret, for the observed behaviour. Consider the graph in Figure 41.

Figure 41 shows the regret that an individual expects to face, conditional on their proposal being accepted but as a function of the minimum acceptable offer, for the range of proposals they can make between 0 and 100 in the ultimatum game⁴⁰. This expected regret is zero, for a proposal of 0, because if the offer is accepted, there was no superior bid that the proposer could have made which would have increased their payoff. At the other end of the scale, if the proposal is 100 and is accepted, then the proposer will expect to feel regret if the minimum acceptable offer was 0 or 1 or 2 or 3 all the way up to 98 or 99, all of which have a positive probability, knowing they could have

⁴⁰ remembering that a proposal of 0 corresponds to keeping the entire amount of money for yourself, and a proposal of 100 corresponds to giving all the money to the other player




Figure 41: Expected Regret from rejection

lowered their offer and improved their payoff. Summing this regret over all those probabilities means that the maximum expected regret from the proposal being accepted will be at 100.

In addition, the curve is drawn to be strictly convex for two reasons. Firstly, it fits with [Loomes and Sugden's](#) standard assumption of regret aversion, in that the maximum regret⁴¹ will be the proposer offering 100, this being accepted, but knowing the minimum acceptable offer was 0, and hence realising they could have increased their payoff by 100 had they offered 0. This possible regret exists when the proposal is 100, but not when the proposal is 99. When the proposal is 99, the maximum regret will be the minimum acceptable offer being 0, and the proposer realising that they could have increased their payoff by 99 had they offered zero. Again, this regret exists when the proposal is 99, but not 98. However, the regret from the 100 situation is disproportionately larger than the regret from the 99 situation, under the assumption of regret aversion. And hence adding the 100 situation to the curve (from 99) contributes more to expected regret than adding the 99 situation to the curve (from 98). Hence the curve is strictly convex⁴². Additionally, we need to assume strict convexity of regret in order to get an interior *regret-minimising* solution to this problem, as it is claimed exists by [Zeelenberg and Beattie](#). This will be presented after introducing the other form of regret, shown in Figure 41.

The second form of expected regret in the ultimatum game comes from the possibility of the proposal being rejected. Once again, it is zero at 100, because, if the proposal of 100 is rejected, then there is no offer that the proposer could have made, since 100 was the maximum, which would have resulted in an acceptance of the proposer's offer, and hence gain them positive payoff⁴³. Additionally, if the proposal was

⁴¹ conditional on the offer being accepted

⁴² alternatively, $ER(x+1) - ER(x) > ER(x) - ER(x-1)$ under the assumption of regret aversion, which gives the strict convexity

⁴³ technically, in an ultimatum game where you can offer only whole numbers, such as this experiment, this regret would also be zero at a proposal of 99, because having this offer rejected, and hence learning that the minimum acceptable offer was 100, means the only offer which could have been accepted would have resulted in a zero payoff for the proposer anyway. However, for all values less than 99, this regret would be strictly positive.




Figure 42: Expected Regret from acceptance and rejection combined

zero, and the offer was rejected, then all positive values for the proposal could, depending on the minimum acceptable offer⁴⁴, have potentially been proposed and generated a positive payoff for the individual, leading to the maximum expected regret from being rejected being at a proposal of zero. Using exactly the same logic as in the “expected regret from the proposal being accepted” situation, again the curve is strictly convex. However, by acknowledging that the total expected or anticipated regret, facing the proposer, is the sum of the expected regret from the proposal being rejected and the expected regret from the proposal being accepted, for every possible proposal we can draw this sum to give the total anticipated regret. This is done in Figure 42.


What Figure 42 tells us is that, assuming both sources of regret are strictly convex, then the total sum will also be strictly convex, and we can find the point of minimum anticipated regret or, the *regret-minimising proposal*, which will lie in the interior of the range of proposals. Given that Zeelenberg and Beattie predict that “...participants will make regret minimizing decisions”[105, p65] in their experiment, then we can assume that

1. Figure 42 is an accurate representation of how the two sources of regret are assumed to act on the proposer in this experiment, and
2. The proposal which gives the minimum total anticipated regret corresponds to the actual decision made by the proposer in the experiment

As such, we can analyse the explanation given by Zeelenberg and Beattie for the observed behaviour in the experiment⁴⁵ using this graphical framework. For those in the low regret group (LRG), let us assume that the above figure corresponds to their decision making process, and they are indeed acting as *regret minimisers*. In which case, their proposal will be the value, p^* , which minimises the *total* anticipated regret (not simply one form or another). This regret minimising position is shown in Figure 43.

⁴⁴ which, in this game, corresponds to the continuum of states of the world

⁴⁵ that those proposers in the HRG submitted lower offers in the second stage because they were more regret averse, and hence made a lower regret minimising offer, than those in the LRG

Figure 43: Regret minimising proposal p^*

Through the design of the experiment, the only difference between the LRG and HRG is the amount of regret to which the respective groups were exposed in stage one, so any significant difference observed between the groups, with respect to their second stage proposals, can be inferred to have been caused by the experience of regret. Indeed, the HRG make significantly lower proposals than the LRG, which is attributed to an increased sensitivity to regret in the state of the world where the proposal is accepted. This idea is represented as an increase in the convexity of the red curve in the graphical framework, as shown in 44, or, intuitively, that large regrets now loom much larger than small regrets for the HRG compared to the LRG. The result of this change to the red curve then filters through to the blue curve, or the total anticipated regret.

As explained by Zeelenberg and Beattie, the impact of this increased sensitivity to the anticipated regret from a proposal being accepted is to *lower* the regret minimising proposal from p^* to p' , shown graphically in 46 by the new minimum point of the blue dashed line, which, itself, is the total anticipated regret curve for those in the HRG.

So, this explanation does indeed appear to predict the observed experimental result of a lower second round proposal for those in the HRG (p' , which had a mean of 26.34) compared to the LRG (p^* , which had a mean of 34.69). The salient question, however, is whether this is the *only* explanation for the observed experimental result. Is there another reason why the regret minimising point for the HRG would lie to the left of that for the LRG? Consider the graph in 46 where the expected regret from the proposal being accepted is not different between the HRG and LRG (i.e. the red curve remains unchanged), but the expected regret from the proposal being rejected has been reduced for the HRG.

In Figure 46, the dotted green curve (the “expected regret from the proposal being rejected” curve for the HRG) lies below the solid green curve (the “expected regret from the proposal being rejected” curve for the LRG) which, when combined with the unchanged red curve, gives rise to the total anticipated regret curve for the HRG (the dotted blue curve) lying *below* the same curve for the LRG (the solid blue curve). Intuitively, this can be thought of as the HRG now having a




Figure 44: Increased regret aversion should the proposal be accepted




Figure 45: New, lower regret minimising proposal p' caused by increased regret aversion




Figure 46: New, lower regret minimising proposal p' caused by reduced regret aversion

lower sensitivity to potential regret, or *lower regret aversion*, than the LRG, as, for every proposal, the anticipated regret is lower, with the source of this lower regret aversion being the lower sensitivity to regret from the proposal being rejected.

Notice, however, that the effect on p^* is the same as it was in the case where there was increased sensitivity to regret from the proposal being accepted. Again, p' now lies to the left of p^* indicating that the representative regret minimising decision maker from the HRG will make a lower stage two proposal than the representative regret minimising decision maker from the LRG, which is indeed what was observed in the experiment run by [Zeelenberg and Beattie](#).

As such, the explanation given above, that the experience of regret actually causes the decision maker to have a *lower* regret aversion, through less concern about regret in the state of the world that their offer is rejected, predicts the same experimental result as was obtained by [Zeelenberg and Beattie](#). Hence, it cannot be rejected as a plausible explanation for the observed behaviour, despite being, at the aggregate level, exactly the opposite reasoning to the explanation given by [Zeelenberg and Beattie](#), simply on the basis of observed choice behaviour.

The outcome of this analysis is to show that, in an experiment such as was run by [Zeelenberg and Beattie](#), there will always be more than one plausible explanation involving anticipated regret, and so it makes little sense to pick one over the other without any corresponding empirical or psychological justification. Just as [Zeelenberg and Beattie's](#) explanation is valid given the observed behaviour, so is the one in Figure 46. And though it might be tempting to *assume* that the experience of regret from having a “too high” offer accepted in stage one leads to increased sensitivity to that same regret in stage two, so it is possible to assume that because the decision maker *did not* experience the regret of having a “too low” offer rejected in stage one, that they would be less sensitive to that type of regret in stage two. Which of these two effects will dominate the *total* anticipated regret aversion in stage two is, therefore, unknown, and hence it is not possible to generalise from this observed




Figure 47: Regret minimising contract offer p^*

behaviour that the experience of regret makes a decision maker either more or less sensitive to *total* anticipated regret in a future decision.

Application to Creyer and Ross [1999]

As previously explained, there is a theoretical equivalence between the two stage ultimatum game experiment run by Zeelenberg and Beattie, and the two stage contract bidding experiment run by Creyer and Ross. In the ultimatum game, a *lower* proposal yields a higher payoff, but has a higher chance of being rejected outright, but in the contract bidding game, a *higher* contract offer yields a higher payoff, but has a higher chance of being rejected outright. As such, the same analysis which was applied to Zeelenberg and Beattie can be applied to Creyer and Ross, simply flipping the colour of the regret curves, so that the expected regret from the proposal being accepted (the red curve), is at a maximum with the lowest possible bid (which is \$100k in the second stage of this experiment), falling to a minimum with the highest possible bid (which is \$162.5k), and the expected regret from the proposal being rejected (the green curve), is at a maximum with the highest bid (\$162.5k) and a minimum with the lowest bid (\$100k). As in the Zeelenberg and Beattie experiment, and for the same reasons given the theoretical equivalence between the two experiments, both regret curves are convex.

In Figure 47 p^* represents the regret minimising contract offer of the LRG in the second stage of the game. The experiment found that the mean offer in the second stage, for those in the HRG, was lower than the mean offer in the second stage of those in the LRG. This result met with the initial prediction of Creyer and Ross, that those in the HRG "...should be more likely to choose the option [in the second stage] which maximises his or her chances of a positive outcome in a subsequent decision, even if that option provides a lower pay-off." [21, p386/7] The interpretation of this prediction, from a regret aversion standpoint, is that those in the HRG will be more conscious of the possibility of having their offer rejected in the second stage, and so be likely to make a safer, lower, offer. Graphically, therefore, this corresponds to an increase in the convexity of the green curve, which gives the expected regret




Figure 48: New, lower regret minimising contract offer p' caused by increased regret aversion

from the contract offer being rejected. Again, by combining this new green (dotted) curve with the unchanged red curve, we get the new, blue (dotted) total anticipated regret curve for those in the HRG, shown in Figure 48.

As shown in Figure 48, the regret minimising contract offer for those in the HRG is *lower*, at p' , that it is for the LRG, at p^* , which was the behaviour observed in the experiment.

At first glance, this appears remarkably similar to the result from the Zeelenberg and Beattie experiment, with one regret curve becoming more convex, corresponding to an increased degree of regret aversion, leading to a lower proposal, or contract offer, for those in the HRG compared to those in the LRG. Here, however, unlike the Zeelenberg and Beattie experiment, a lower offer is associated with lower risk and lower payoff, and the source of regret which is increasing is the anticipated regret from the offer being rejected, not the offer being accepted. Indeed, if the same intuition from the Zeelenberg and Beattie experiment was applied to Creyer and Ross' experiment, then it would predict an increase in the convexity of the red curve, and hence an *increase* in the contract offers of the HRG compared to the LRG. It is fair to say, therefore, that both the intuition and experimental results of the two experiments are in direct contrast, despite the initial impression that they are confirming the same theory that the experience of regret leads to a subsequent increase in regret aversion.

As with the Zeelenberg and Beattie experiment, it is also possible to explain the same experimental result by considering a *reduced* degree of regret aversion from the other source of potential regret in the experiment. This time, by assuming that the experience of regret in stage one gives rise to a decreased sensitivity to the potential regret from the offer being accepted in the second stage, possibly due to the realisation that regret, as with other negative emotions, is not as bad




Figure 49: New, lower regret minimising contract offer p' caused by reduced regret aversion

in experience as it is in anticipation⁴⁶, we can graphically decrease the convexity of the red curve, and analyse the impact.

In Figure 49, as p' is lower than p^* , this explanation also predicts the observed experimental result of the second stage contract offer being lower for the HRG than for the LRG, and hence it is not possible to discount either explanation on the basis of simply the observed difference in choice behaviour. Again, it is worth noting that one explanation implies that the experience of regret has led to an increase in the total regret aversion of the decision maker, and the other implies that the experience of regret has led to a reduction in the total regret aversion, yet both are explanations perfectly consistent with the result of the experiment.

Application to Raeva and van Dijk [2009]

In contrast to the experiments of Zeelenberg and Beattie and Creyer and Ross, the experiment of Raeva and van Dijk includes very specific statements of the mathematical assumptions from which their experimental predictions are derived⁴⁷. Specifically, for our purposes in comparing this experiment to the others, we are interested in the complete feedback trials, where there was a group who experienced regret in stage one (the Regret Condition - RC) and a group who experienced no regret in stage one (the No Regret Condition - NRC). In the second stage, both groups are asked to give a certainty equivalent value for a gamble which is a 50% chance of €1000 and a 50% chance of €100, with the hypothesis being that those in the RC will submit *lower* certainty equivalent than those in the NRC. This prediction is explained through the mathematics of the gamble and certainty equivalent. As per equation 1 of Raeva and van Dijk, if the probability of getting €100

⁴⁶ in keeping with the work of Gilbert et al. [26]

⁴⁷ and, furthermore, their mathematical model is derived from Bell formulation of regret theory as opposed to Loomes and Sugden which has typically been referred to so far

from the gamble is p , and the certainty equivalent value is z , then the individual will choose z such that

$$(1 - p).v(1000) + p.v(100) + p.r(100 - z) = v(z) + (1 - p).r(z - 1000) \quad (2.1)$$

where $v(\cdot)$ is the utility attached to each outcome and $r(x - y)$ is an increasing regret function which gives the amount of regret experienced from obtaining x , but knowing the alternative choice would have resulted in y . Without the regret terms on either side of the equation, you get the simple expression for setting the certainty equivalent value of a gamble, whereby the utility of the certainty equivalent is equal to the expected utility of the gamble. The third term on the left hand side of the equation then adds in the regret suffered should the gamble be taken, but it results in the low payoff, hence wishing the certainty equivalent value was chosen instead (weighted by p , the probability of the low payoff). The second term on the right hand side adds in the regret suffered should the certainty equivalent be taken, but taking the gamble would have resulted in the high payoff, hence wishing the gamble was chosen instead (weighted by $1 - p$, the probability of the high payoff).

The prediction of [Raeva and van Dijk](#) is that, for those in the NRC, the above equation describes their choice problem in the second stage, with $p = 0.5$ owing to the 50/50 nature of the gamble. For those in the RC, however, the experience of regret, in stage one, from picking a door which contained a €1 participation fee for the experiment, as opposed to the alternative, revealed door, which, had it been chosen, would have increased their participation fee to €10, results in an *increase* in the subjective probability of losing the gamble in stage two. Mathematically, this corresponds to an increase in the value of p faced by those in the RC compared to those in the NRC. This has the effect of *decreasing* the left hand side of equation 2.1, and *increasing* the right hand side of equation 2.1⁴⁸. When the same certainty equivalent value from the NRC⁴⁹, therefore, is inserted into the equation for the RC, the impact of the increase in p is to make the “money for certain” more appealing than the gamble, and hence the value of z no longer represents certainty equivalent. As such, in order to restore equality to the equation, a *lower* value of z is needed, which explains the prediction of the experiment that those in the RC will submit lower certainty equivalents than those in the NRC.

This method of formulating a regret-based decision under uncertainty problem is especially useful for analysing the two different forms of regret. The equation above explicitly shows that there are two potential regrets in this problem; of ending up with €100 from losing the gamble, instead of € z for sure had they taken the certainty equivalent, and of ending up with € z for sure, instead of the €1000 which would have resulted had they taken the gamble. As previously explained, these regrets act in opposite directions on the decision maker, with considerations of the first making it less likely that the individual will choose the gamble, and considerations of the second making it less likely that the individual will take the money for certain.

⁴⁸ remembering that the regret term will act negatively on the total utility of the individual
⁴⁹ i.e. the value of z which equalised the equation for individuals in the NRC

As noted previously, the approach of Zeelenberg and Beattie and Creyer and Ross was to suppose that the experience of regret in stage one, of such a decision making problem, would lead to a change in the anticipated regret in stage two, specifically from a change in the sensitivity to regret itself. In the mathematical language of this problem, this corresponds to a change in either $r(100 - z)$ or $r(z - 1000)$, and, in essence, the point being made earlier was that you can explain a deviation in observed experimental behaviour as much by an increase of one as by a decrease in the other.

In contrast to that approach, the assumption of Raeva and van Dijk is that the effect of experienced regret is not to change the regret function itself, but to change the subjective probability, p , associated with a *similar* kind of regret experienced in stage one. As such, this assumption does lead to a change in the “expected regret” in stage two, but does so by changing the “expected” part rather than the “regret” part.

The benefit of making such an assumption is that in the mathematical framework, presented by Raeva and van Dijk, an increase in the probability of one form of regret, in this case $r(100 - z)$, necessarily implies a decrease in the probability of the other form of regret, $r(z - 1000)$, because the sum total of the subjective probabilities must add up to 1. Intuitively, if you believe there is a higher chance of the gamble being resolved not in your favour ($p > 0.5$), and hence a higher chance of experiencing regret from not taking the money for certain, and having the gamble giving you €100, then you necessarily also believe that there is a lower chance of the gamble being resolved in your favour ($1 - p < 0.5$), and hence a lower chance of experiencing regret from taking the money for certain, and seeing the gamble be resolved to give you €1000 had you taken it.

By holding the regret function as stable, but assuming the subjective probabilities will move in the above fashion, *both* sources of anticipated regret experience a change in this model, and *both* changes act in the same direction on the decision maker. In the case of an increase in p , this makes the money for certain more attractive, and in the case of a decrease in p , this makes the gamble more attractive. This runs contrary to the frameworks used by Zeelenberg and Beattie and Creyer and Ross, where a change in one source of anticipated regret had no impact on the other source of anticipated regret.

As such, this method of predicting the observed experimental behaviour escapes some of the criticism, rooted in the Identification Problem, which was applied to the previous two experiments. But this only occurs by making a very strong assumption about the relationship between the sources of anticipated regret. In the language of the Identification Problem, you are no longer assuming independence between the two variables of interest; in essence reducing the problem from having two variables to only having one. The justification for this assumption comes in a footnote of the paper which states

“We assume that the regret function $r(\cdot)$ is not affected by experienced regret, in line with the findings of Ritov [1996], outlined in Sec.2, that the main attribute of regret influencing choice behaviour is the probability to regret.”
[73, Footnote 5]

The assumption has two distinct components. The first is that the regret function is not affected by the experience of regret. The second is that

the subjective probability changes in a specific direction as a result of the experience of regret. The main question, therefore, becomes whether or not making such a strong assumption can be justified on the basis of the experimental design and other experimental literature.

Firstly, whilst the Ritov paper does indeed study the impact of probability on choice, in situations where anticipated regret is present, it is not as simple as saying that the principle component of anticipated regret in a decision is the probability. In contrast, Ritov [76] primarily looks at the impact of probability *on the effect of outcome feedback* on choice, and find that “...the experimental results... imply that the impact of availability of outcome information on the decision maker’s preferences varies with the probability of obtaining the better outcome” [76, p236] which is different to saying that the primary channel through which anticipated regret *itself* will vary is through the probability of regret.

However, just because there is no clear support for the assumption, this does not necessarily make it incorrect. But in absence of this academic support, it is tough to believe that the regret function is not at all affected by the experience of regret, especially as this is the primary assumption of other work in this area. As the function is not observable, however, it can’t be proved either way.

Given the experimental design, the bigger issue would seem to be assuming a change in subjective probability of regret, as the gamble in stage two of this experiment is absolutely explicitly a 50/50 gamble. Whilst it is not quite a “heads or tails” type gamble, there can be little doubt that the objective probability of each option is 0.5, and so it seems to make little sense to suggest that an individual would believe that such an explicit probability could be changed by the previous experience of regret from a similar, but distinct, task. If we take the implication of Raeva and van Dijk that the experience of regret makes you “fear the worst again”, then their interpretation is that the prior experience of regret changes the subjective probability of the worst happening, rather than changing the “fear” associated with the similar outcome.

It would appear to be a more plausible explanation for a situation of true uncertainty, where the probabilities of each state of the world are not explicitly given. Consider, for example, a individual who purchases a laptop, without an extended warranty, without knowing the explicit probability that the laptop will become defective within the term period offered by the extended warranty. As it transpires, the laptop becomes defective within the term period of the offered warranty, and so needs to be replaced, with a new, but not identical to the previous, laptop. Assuming the two laptops are sufficiently unrelated that the fact the first laptop failed does not inform the individual about how likely the second is to fail also (indeed, since time has passed, it is likely that the second laptop is more technically capable than the first, and hence less likely to break down), there is no method of Bayesian updating, using information about the probability of defectiveness from the first laptop, to inform the consumer of the new probability that the second laptop will also become defective. However, the experience of regret from the first laptop incident, where the extended warranty was not purchased, and in retrospect, it would have been better to have done so, makes the same possibility of regret, of purchasing the second laptop without an extended warranty and having it fail, loom larger in the

mind of the individual. Indeed, as the probability of the second laptop failing is unknown, it could be reasonably assumed that the experience of regret from the first laptop made the probability of failure of the second laptop *feel* like it was more likely than it was in reality. As a result, the fear of regret, from not purchasing the extended warranty for the second laptop, and having it fail, looms larger than the fear of regret from purchasing the extended warranty, and have the laptop not fail during the period of cover, and so the consumer decides to purchase the extended warranty with the second laptop⁵⁰.

If however, we retain the set-up of the [Raeva and van Dijk](#) experiment, with explicitly given probabilities, then there are certain pieces of evidence which work *against* the critical assumption of the paper, that the experience of regret makes the subjective probability of regretting a similar, future decision more likely. For example, [Tversky and Kahneman](#) discuss a “Belief in the Law of Small Numbers”, “...that people view a sample randomly drawn from a population as highly representative, that is, similar to the population in all essential characteristics.” [93, p105] The application of this belief to both the above experiment, and more generally, multiple-stage decision tasks where the states of the world have known probabilities, is widely known as the Gambler’s Fallacy, where “[s]ubjects act as if *every* segment of the random sequence must reflect the true proportion: if the sequence has strayed from the population proportion, a corrective bias in the other direction is expected.” [93, p106]. Note that the Gambler’s Fallacy does indeed suggest that there can be changes in subjective probabilities as a result of prior outcomes, as in the above experiment. Most commonly, this is associated with the example that observing a run of heads increases the probability of a subsequent tail, when flipping a coin in sequence. The implication, however, of the assumption of [Raeva and van Dijk](#) that the experience of regret from making a wrong decision *increases* the subjective probability that you will experience regret from a similar decision in the future, works in exactly the *opposite* direction to that of the evidence of the Gambler’s Fallacy. If you were repeatedly betting on tails, whilst flipping a coin in sequence, and constantly observing heads, and hence experiencing regret from betting tails when you should have bet heads, then [Raeva and van Dijk](#) suggest that this would *increase* the subjective probability that, in the future, you would experience regret from betting tails and observing a head, hence implying that a head is *more* likely in the future. The Gambler’s Fallacy, on the other hand, suggests that the previously observed run of heads would be “balanced out” by observing a tails (and, hence, returning the small sample probability of the coin towards 0.5), hence implying that a future head is *less* likely than a future tail⁵¹.

As such, whilst the assumption for the observed behaviour in the experiment cannot be disputed on exactly the same theoretical grounds

⁵⁰ and, indeed, the second laptop in that example is currently proving to be a very effective and reliable tool for writing this thesis. Or at least significantly more reliable than the first.

⁵¹ An effect which works in the opposite direction to The Gambler’s Fallacy is The Hot Hand effect. It is discussed in more detail in Chapter 3, however it is typically thought of in terms of a belief in the continuation of *winning* gambling streaks (hence the term *hot hand*), rather than losing streaks, or streaks of repeated regrets. As such, it is not focussed on here. In any case, when discussing the validity of assumptions, it is more important to evaluate the merits of evidence which runs contrary to the assumption (for example, *all swans are white*), as just one piece of evidence can disprove an assumption (*I saw a black swan*), but no amount of supportive evidence (*I saw a white swan*) can prove it to be true.

as the experiments of Zeelenberg and Beattie and Creyer and Ross, it can be disputed on experimental grounds, as the evidence of the Gambler's Fallacy would suggest that the additional assumption made⁵² is inappropriate for this specific experimental set-up, given that the probabilities used in stage two are both explicit and clearly 50/50. In addition, it remains unclear why the effect of the experience of regret on subsequent anticipated regret would act on the subjective probabilities of future states of the world, and not on the regret function itself.

Application to Coricelli et al. [2005]

The approach of Coricelli et al. to the question of the impact of experienced regret on subsequent regret aversion, is notably different from that of the previous three experiments through their use of both neuroimaging techniques and the deviation from a simple two-stage set-up to one of repeated decision making. The first part of their analysis discusses whether anticipated regret itself can be shown to play a role in the decision making of the individuals in the experiment, for, if this is not shown to be true, then it will be impossible to estimate any effect of experienced regret on subsequent regret aversion using this experimental set-up. Using a panel logit model, the probability of an individual i choosing gamble 1 (g_1) in the experiment, as opposed to gamble 2 (g_2), at time t , is modelled as a function of the anticipated regret, anticipated disappointment, and expected value of the gambles. Formally

$$\Pr(g_{1it}) = F[d_{it}, r_{it}, e_{it}]$$

where

$$d = (|y_2 - x_2|(1 - q)) - (|y_1 - x_1|(1 - p))$$

and is the difference in anticipated disappointment between the two gambles (where x_i and y_i are the highest and lowest outcomes of gamble g_i , p is the probability of obtaining outcome x_1 and q is the probability of obtaining outcome y_2)

and

$$r = |y_2 - x_1| - |y_1 - x_2|$$

and is representative of the difference in anticipated regret between the two gambles

and

$$e = EV(g_1) - EV(g_2)$$

and is the difference in expected value between the two gambles. These 4 equations correspond to equations 1-4 in the work of Coricelli et al..

Looking at the equation for anticipated disappointment, $(|y_2 - x_2|(1 - q))$ corresponds to the probability of obtaining the lower outcome, multi-

⁵² necessary in order to escape the Identification type problem

plied by the amount of money that was “lost” as a result of obtaining the lower, instead of the higher, outcome. As such, it represents the expected disappointment associated with gamble 2⁵³. Equivalently, $(|y_1 - x_1|(1 - p))$ is the expected disappointment associated with gamble 1. As such, if the difference between them is positive, it implies there was a higher expected disappointment in gamble 2 than gamble 1, and, as such, a positive coefficient on d in the panel logit model implies that the bigger the impact of anticipated disappointment on gamble 2, compared to gamble 1, the more likely the decision maker is to choose gamble 1 over gamble 2, in effect, choosing to avoiding the potential for disappointment where possible.

The same idea is then also applied to anticipated regret, but two things are immediately noticeable. Firstly, the two sources of regret acting on the decision maker are obvious. On one hand, there is $|y_2 - x_1|$, which is the difference in payoff from choosing gamble 2, obtaining the low payoff, and realising that you would have been better off from choosing gamble 1 and getting the high payoff, and on the other hand there is $|y_1 - x_2|$, which is the difference in payoff from choosing gamble 1, obtaining the low payoff, and realising that you would have been better off from choosing gamble 2 and getting the high payoff. The first, however, has a positive effect on r , and the second has a negative effect, effectively demonstrating the idea of competing effects of regret on decision making.

The second noticeable point is the absence of probability weighting attached to the measures of regret, as was the case when calculating anticipated disappointment. Instead of representing anticipated regret as the *expected* regret from choosing one gamble as opposed to the other, the measure of regret given states what would be the ex-post regret obtained, assuming that a particular gamble was selected *and* that the selected gamble gave the lower payoff *and* that the other, unselected gamble, would have yielded the higher payoff. As such, it will misrepresent the effect of each source of regret on the decision maker, as it necessarily assumes that the individual anticipates the worst thing happening every time they select a gamble. It is admittedly more complicated to calculate the true, expected regret, for each gamble, however the cost of this assumption is clear from considering a small sample of the problems faced by decision makers in the experiment.

For example, the expected regret of the decision problem in Figure 1 of the paper [19, p1256], has gamble 1 yielding +200 with probability 0.5 and -200 with probability 0.5, and gamble 2 yielding +50 with probability 0.75 and -50 with probability 0.25. As such, the individual will regret choosing gamble 1 only 50% of the time, as, in the other 50% of the time, they achieve the best outcome (+200) in the entire problem. In addition, that 50% is comprised of 75% of the time obtaining +50 from the unselected gamble (and, hence, in the language of the paper, experiencing $|-200 - 50|$ regret) and 25% of the time obtaining -50 from the unselected gamble (and hence experiencing $|-200 + 50|$ regret). As such, the *expected* regret of gamble 1 is, more accurately, given by $r_1 = 0.5(0.75|-200 - 50| + 0.25|-200 + 50|) = 112.5$, as opposed to the figure of $r_1 = 250$, which is the number used in the regression analysis. Equivalently, the true expected regret for gamble 2 is given by $r_2 =$

⁵³ in more classical representations of Disappointment Theory, such as Bell [7] and Loomes and Sugden [59], there would be a “disappointment function” applied to the difference in outcomes, so as not to impose linearity on the effect of disappointment, but this is omitted for simplicity here.

$0.75(0.5|50 - 200|) + 0.25(0.5|-50 - 200|) = 87.5$, rather than the figure of $r_2 = 250$ used in the regression. The impact of this miscalculation is that both gambles, from the perspective of simply considering the worst thing that would happen in each situation, are assumed to act *equally* in the mind of the individual, and hence have zero net effect ($r = r_2 - r_1 = 250 - 250 = 0$) on the decision of the individual. In reality, however, there is more *expected* regret from gamble 1 (112.5) than from gamble 2 (87.5), and so there should be a net effect of anticipated regret pushing the individual away from gamble 1, towards gamble 2, included in the regression analysis of that decision.

Indeed, this approach can be applied to all 48 pairs of gambles used in the experiment, as provided in Table 29 in Appendix A. By calculating the true expected regret⁵⁴ for each of the gambles, in each of the 48 decision problems, and calculating the expected regret of gamble 2 minus the expected regret of gamble 1, we can obtain a *true* value of r , here labelled as r^t , and compare it to the *false* value of r , here labelled r^f , given by Coricelli et al..

There are several notable things immediately obvious from looking at Table 29. Firstly, if we consider an “incorrect” value of r to be where either

- $r^f \geq 0$ and $r^t < 0$

or


- $r^f \leq 0$ and $r^t > 0$

then 15 out of 48 problems have an “incorrect” value for r in the regression analysis conducted by Coricelli et al., simply on the basis of having an incorrect sign, in essence saying that one gamble has more expected regret than, or the same as, the other, when, in reality, this is not the case. Of the remaining 33 problems, only 1 has $r^t = r^f$, in essence saying that the value of r used in the regression analysis was an accurate representation of the expected regret facing the individual from both gambles only once. In the remaining 32 problems where this was not the case, 16 *overstated* the difference between the two gambles (so $r^f > r^t$), and hence suggested that concerns for anticipated regret would make the individual choose gamble 1 *more* than was the case in reality, and 16 *understated* the difference between the two gambles (so $r^f < r^t$), and hence suggested that concerns for anticipated regret would make the individual choose gamble 1 *less* than was the case in reality. In total, for all 48 problems, the average margin of error⁵⁵ was 105 units, which is over half a standard deviation of the values of r^f used in the regression analysis. Figure 50 plots both the true and false values of r for all the problems, ordered by the value of r^f used in the regression analysis.

What is evident from Figure 50, therefore, is that there is considerable variation between the true and false values for r , and it is impossible to state, one way or the other, whether, if the true values of r were used in the regression analysis, the same relationship between anticipated regret and the decision of the individual would be found with significance. It may be the case that the true values of anticipated regret lend further

⁵⁴ assuming, as do Coricelli et al., that experienced regret is simply the mathematical difference between the outcome that was obtained and the outcome which was selected in the unchosen gamble, and is zero in cases where the outcome which was obtained exceeds the outcome which was not (i.e. there is no rejoicing)

⁵⁵ that is, the average of $|r^f - r^t|$ for all 48 problems



Source : Supplementary Table 1 - <http://www.nature.com/neuro/journal/v8/n9/extref/n1514-S4.pdf>

Figure 50: Visual representation of r^t and r^f

support to the relationship, or it could be that the true values statistically eliminate the relationship, in which case the significance found in the original regression would likely be caused by an omitted variable, related to both the decision taken by the individual, and the simple formulation of anticipated regret given by Coricelli et al..

In addition to this problem, as with the representation of anticipated disappointment, there is no account taken of potential non-linearities in the regret function arising from an assumption of regret aversion. Or, in other words, the typical assumption of regret theory, that larger regrets will loom disproportionately larger in the mind of the individuals than small regrets, is ignored.

As a result of these two points, one of the principle claims on which the experiment is based, that anticipated regret plays a significant role in the decision making process of the individuals, is subject to considerable doubt. The consequence of this doubt is that other results in the experiment, such as “[t]he proportion of regret-avoiding choices increased over time with the cumulative effect of the experience of regret” [19, p1259], which follow from the assumption that anticipated regret plays a significant role in the decision making process of the individual, are also cast into doubt. Indeed, whilst Coricelli et al. show that the proportion of regret-avoiding choices increases through the first, second and last third of the experimental decisions, it is not stated how the proportion of regret avoiding choices was calculated. Assuming it was based on the same formula, $r = |y_2 - x_1| - |y_1 - x_2|$, which was used in the regression analysis, with a regret-avoiding choice being counted as “1” if gamble g_1 was chosen when $r > 0$, or if g_2 was chosen when $r < 0$, and “0” otherwise, then the same criticism applies as it did to the regression analysis; that the variable being created may not be truly representative of “expected regret”, and hence the supposed “regret-avoiding” choices may not be correct, instead being representative of some other, related concept such as risk-aversion.

The second approach to modelling the effect of experienced regret on subsequent choice is through the use of neuroimaging techniques. Two analyses, using the neuroimaging data, were performed in relation

to the impact of experienced regret on subsequent choice, to look for brain activity which correlated with specific measures of experienced regret. The first analysis used a measure of “cumulative regret”, which was defined as

$$CR_t = (A_{\text{unobtained},t-1} - A_{\text{obtained},t-1})$$

“...where CR is cumulative regret, t is trial, A_{obtained} is the average realised payoff and $A_{\text{unobtained}}$ is the average payoff of the unselected gambles.”[19, p1261] This calculation is very much a rough approximation to the idea of cumulative regret, as, for example, it will necessarily have less variation across individuals in the later trials compared to the earlier trials, due to the time averaging, and also incorporates an idea of “rejoicing”, as obtaining a payoff in excess of the unobtained payoff in a given trial is associated with a reduction in cumulative regret for the next trial. And though this variable is found to correlate with “...activity in the medial left amygdala ... and medial OFC”[19, p1259] at the time of choice, this is not sufficient to imply that the cumulative experience of regret is having an impact on subsequent regret aversion. It may simply be the case that this cumulative regret measure is representative of the emotional state of the individual (i.e. whether they feel they are doing “well” or “not well” in the task), and subsequently how risk averse or seeking (distinct from regret averse or seeking) they wish to be in the next trial.

Similarly, the second analysis performed using the neuroimaging data uses “prior regret”, or the regret from the gamble in the last period, in place of the cumulative regret used in the first analysis. Whilst it is not stated how this measure is calculated (and, in reference to the previous analysis, whether it includes rejoicing or not), it is found to “...[enhance] responses in the right dorsolateral prefrontal (DLPFC), right lateral OFC, and inferior parietal lobule”[19, p1259] between the onset of the trial and the response of the participant. Once again, it is not sufficient, from this analysis, to draw a conclusion about how the experience of regret influences subsequent regret aversion, merely that the experience of regret in a prior period appears to have some lasting residual into the next period’s decision making process.

In order to use neuroimaging techniques to assess the impact of experienced regret on subsequent regret aversion, we would first need to identify the areas of the brain associated with anticipated regret at the time of choice, or, in other words, the areas of the brain associated with the anticipation of future negative emotions and the anticipation of self-blame. Only once this has been achieved will it be possible to use neuroimaging techniques to see what causes changes in patterns of brain activity in these areas when decisions are made. As such, this experimental set-up is a helpful first step in this process, but the results drawn are insufficient to draw substantive conclusions about the impact of experienced regret on subsequent regret aversion and subsequent choice behaviour.

2.3.4 *Conclusions to be drawn from criticisms of existing literature*

Each of the four experimental papers, presented above, claim to offer an insight into how the experience of regret influences subsequent

decision making. They all demonstrate a statistically significant change in individual decision making behaviour, resulting from an experience of regret in a prior stage, yet the conclusions about how the experience of regret *causes* the change in behaviour are subject to considerable doubt once the regret-based processes, about how that causation might actually work, are examined in higher detail. As such, it is fair to say that all these experiments demonstrate *something* about the experience of regret, but we can't, at present, say what that is.

The most common conclusion drawn from this literature is that the experience of regret causes an *increase* in regret aversion in subsequent decision making. As demonstrated above, whilst that may possibly be correct, it is not a conclusion that can be drawn from the current experimental literature. Indeed, it is possible to interpret the data in a different way, making assumptions consistent with regret theory, but deriving the exact opposite conclusion, that the experience of regret causes a *decrease* in regret aversion in subsequent decision making. Indeed, whilst the first appears to be a more natural statement than the second (it is easy, for instance, to imagine someone who has been burned by a bad decision once being determined not to make the same mistake twice), evidence from psychology, by Gilbert et al., "...suggest that people are less susceptible to regret than they imagine, and that decision makers who pay to avoid future regrets may be buying emotional insurance that they do not actually need." [26, p346] As such, the experience of regret by an individual, and the subsequent realisation that regret was, in fact, worse in anticipation than in experience, could lead to a reduction in subsequent regret aversion, if the individual does indeed realise that they are paying for emotional insurance that they don't need. So whilst the more familiar conclusion may prevail and persist in the academic ether, there is enough doubt, and evidence which runs contrary to the popular perception, that it is certainly worth continuing investigations in the area to further explore the matter.

2.3.5 The next step

Having demonstrated that the existing literature is insufficient in providing a resolution to the problem of how experienced regret affects subsequent regret aversion, it would be remiss to then not propose a course of action which would lead us further down the path of discovery in this field. The arguments presented above certainly hint at the issues which lie at the heart of the problem in drawing conclusions from the types of experiments which have been typically developed and run in the past. As such, the next step is to create new experiments which analyse those specific issues in isolation from the much broader research questions. By learning more about the specifics, we can then reform and remodel the original experiments, taking into account an increased knowledge of the way regret works in the mind of individuals when facing decision under uncertainty.

2.4 WHY DOES THIS MATTER?

A valid criticism of this work is to say that it is not particularly important whether an individual is affected by Type A or Type B regret in their decision making, as long as the point proved is that regret, in some form, does affect decision making. In the classic demand and supply

identification problem, the problem matters because firm (micro level) and policy (macro level) decisions may be incorrect unless the source of change is correctly identified, and the cost of incorrect decisions can be high. As there appears to be no “real world” cost to misidentifying the correct source of regret which influenced a decision, it is worth taking some time to explain the reasoning for starting this discussion at all.

The point of this discussion is certainly not to argue whether one type of regret should be considered more important than another. As has been shown, Type A and B regret are arbitrarily defined for a given decision problem. In some cases, the distinction between them may matter, and in some cases it may not. It is up to a researcher to decide whether, in the context of a specific decision problem, it is worth making the distinction between two types of regret. It is certainly not something which can be argued in abstract or absence of a context.

As the distinction between different types of regret has only been previously touched on in a very limited number of instances, the immediate implications and applications of this work are, therefore, also limited. That is to say, aside from affecting the conclusions of a small number of experiments, there isn’t much else to say at this point in time. Taking this narrow view, the value of the work appears small.

However, academic research is a continually evolving process, and the vast majority of work will rely on previous work as a starting point. It is therefore important to highlight limitations of previous work, so areas of weakness do not persist and pervade into new research. Failing to do so acts as a disservice to future researchers who rely on the robustness of findings in order to create their own theories and ideas.

Relating this point specifically to regret, in future research where the distinction between different sources of regret does matter, it is very important that assumptions about how an individual will respond to them are based on solid evidence. This is so experiments are robustly designed and empirical evidence is correctly analysed allowing new research conclusions on the topic to be drawn.

If you consider, therefore, the benefit of this work to include what it gives to future research, as well as present and past, its overall value surely increases.

Having talked to several academic researchers in this area, the common wisdom is that the experience of regret increases subsequent regret aversion. This seems plausible, and, as shown, a simplistic review of the available experimental evidence would appear to support it. To future researchers, however, there is a world of difference between an answer which is plausible and an answer which is definitive. The value of this work is to point out that, based on currently available evidence, we are not anywhere close to being able to say the answer is definitive. It is always more important to acknowledge what we don’t know, than think we know more than we do.

2.5 SUMMARY

The existing experimental literature, which seeks to explain how the experience of regret impacts subsequent regret aversion, and hence subsequent choice, suffers from a fundamental problem which prevent definitive conclusions from being drawn from the observed behaviour in the laboratory.

By explaining how all non-trivial decisions taken under uncertainty involve at least two sources of regret, which will act in opposite directions on the decision maker, any experimental study which looks at how the experience of regret impacts subsequent regret aversion must explain why, should a behavioural change in preference for one option over another be observed, the cause of this change was as a result of, for example, an increase in sensitivity towards one source of regret as opposed to a decrease in sensitivity towards another. Indeed, without such an explanation, any observed experimental result can be explained equally correctly by an *increase* or *decrease* to the *total* degree of regret aversion an individual may have. The most common conclusion in the experimental literature, however, is that the experience of regret leads to an increase in subsequent regret aversion, but, typically, these experiments do not satisfy the above requirement of justifying why they believe one source of regret will dominate the thinking of the individual compared to another. Additionally, the neuroscientific approach to *measuring* both anticipated regret and experienced regret is not sufficiently advanced that it would be possible to observe more than just behavioural changes and draw conclusions about the way in which the experience of regret impacted the subsequent choice behaviour of individuals, with respect to more than one source of regret, despite initial evidence suggesting that such a link does exist.

A SIMPLE EXPERIMENT STUDYING EXPERIENCED REGRET AND CONTEXT

3.1 INTRODUCTION

The experience of regret is a negative emotion, inextricably linked to a preceding action or decision, which, upon the resolution of uncertainty under which the decision was taken, proved to be a worse choice than could otherwise have been made. Importantly, regret can not exist as a standalone emotion, such as, for example, being happy or nervous. I can be “happy”, or “nervous”, without needing to be “happy because of...” or “nervous of...” something specific. In contrast, we think of “being regretful of...” something, rather than simply existing in a general state of regretfulness.

This link between the emotion of regret, and the choice which caused it, lends itself perfectly to being incorporated in economic decision making frameworks, such as those by Loomes and Sugden [58] and Bell [5], precisely because there is an explicit cause of the emotion (the choice) which can be represented mathematically. Again, in contrast, it is difficult to incorporate the effect of being happy or nervous into an economic decision making framework because the conditions which give rise to those emotions existing, and subsequently impacting decision making, are not well defined or understood. What may make one person happy, may not cause another person to be happy to the same extent, or even at all. In contrast, we can be confident in saying that making a poor decision, when an alternative, which was considered, would have resulted in a better outcome, will cause a degree of regret in all people, and, importantly, the better the outcome which would have resulted under the alternative action, the worse the regret.

As discussed elsewhere in this work, the experience of regret, linked to a previous choice, can give rise to a change in future behaviour, often represented as a change in the future aversion to regret as a result of the experience. Various studies (Creyer and Ross [21], Zeelenberg and Beattie [105], Raeva and van Dijk [73]) have attempted to explore whether the experience of regret leads to a subsequent increase or decrease in regret aversion in future decisions, typically where the second stage of the experiment is a repetition of the first (or very closely related), on the basis that a prior experience of regretting a specific choice will have the most impact when the individual is faced with that decision again. The limitations of this experimental approach are explored in the previous chapter, but their approaches pose a further question when considering both the effect of the experience of regret on subsequent decision making and the link between a specific action and the regret caused by it.

The issue at play comes from considering exactly what we mean by the words “action” and “decision”, often used interchangeably, as they relate to the association with the experience of regret. When we are asked to make a choice, there are often different components and contexts to the process of making that choice. Hence the question becomes, which part, exactly, is the emotion of regret tied to? For

instance, usually when we ask subjects in an experimental lab to make a decision over two specific gambles, we will give the gambles titles (Gamble A versus Gamble B), which may be located on specific parts of the screen (Gamble A is on the left, Gamble B is on the right), and they will have at least one specific characteristic which distinguishes them from one another (Gamble A is “risky”, Gamble B is “safe”). Suppose then, that Gamble A is chosen in this experiment, and, after the resolution of uncertainty, Gamble B is found to have been a superior choice. The person will then experience regret associated with their choice. The question is, however, does the individual associate the regret with the fact that the chosen gamble had a specific *title* (“I really regret choosing that infernal *Gamble A!*”) or was located in a specific part of the screen (“I really regret choosing the option *which was on the left*”) or was representative of a broader preference over risk (“I really regret choosing the *risky option*”)?

In terms of a one-shot choice, the specific nature of the association between regret and the choice doesn’t particularly matter, because the context and characteristics of the decision are inextricably bundled together. The gamble which *is* on the left, before the decision is taken, is the same gamble which *was* on the left, after the uncertainty is resolved. However, the literature on the effect of experienced regret on subsequent anticipatory regret, mentioned above, doesn’t ever explicitly specify that these characteristics have to remain the same from one decision problem to the next. Gamble A may be on the left in stage one of the experiment, but be on the right in stage two. We would expect there to be some effect of regret from stage one on stage two, but the question is whether that regret is tied to the fact it was called “Gamble A” or the fact it was on the left hand side of the screen. The answer to this question will be crucial in deciding what was the causal effect of the first stage regret on the choice at stage two.

It may be the case that simply the position of the various options on the screen have no impact from one period to the next. Indeed, it seems fairly likely that this characteristic of Gamble A is of minor importance in the mind of the individual. It is not hard, however, to find two distinct, changeable characteristics of the decision which have a *major* impact on the decision making process. In these cases, the tie between the regret and each of those characteristics will be of significant importance when attempting to analyse the effect of the first stage regret on the second stage decision.

The next section will examine the existing literature on some typical characteristics of decisions which are well known, and have been extensively discussed and incorporated into the current thinking on anticipatory regret. Following this discussion, we will propose an experiment which is designed to isolate one particular characteristic, and, by running the experiment, test whether the effect of experienced regret in stage one on choice behaviour in a stage two, is tied to that characteristic. The importance of this experiment lies in the consequence for existing regret-based decision theories should the emotion be found to tie to only specific characteristics of a decision, as such a level of detail or contextual information, is not currently considered in the regret literature.

3.2 EXISTING LITERATURE

The existing literature on anticipatory regret has never explicitly focussed on a comparison of the different characteristics of a decision, and which the emotion of regret is most closely linked to, despite, as explained above, this being an important consideration when designing both theoretical models and experimental tests. What has occurred, however, are a series of small literatures, each of which analyses how regret can be influenced by a particular contextual characteristic of a decision, based, in part, on intuition about how these characteristics *should* or *feel like they should* play a role in the experience of regret. This literature was especially prevalent during the 1990s; a period which was between the earliest theoretical works of Bell [5] and Loomes and Sugden [60] and then subsequent revisions of the regret-based theoretical literature in the first decade of the 21st century by Hayashi [34], Zeelenberg and Pieters [106] and Hart and Mas-Colell [32]. The experiments, which comprised the literature, were designed to point out limitations of the existing theoretical understanding of the subject, by suggesting factors of interest, in the decision making process, which would have an effect on the influence of anticipatory regret, but were not explicitly represented in any form of the theory. Coincidentally, most of these factors can also be thought of as *characteristics* of the decision, and help provide us with a starting point when trying to identify which specific characteristics of the decision the experience of regret may be associated with.

3.2.1 Action versus Inaction

One characteristic that can often be attributed to a particular decision is whether it falls into the category of “action” or “inaction”. Action is something we are readily familiar with, but it is also true that the decision *not to act* is a choice which can be included in any decision under uncertainty (including regret-based) framework, in that “not acting” will have outcomes associated with each state of the world under consideration. In the standard language of decision under uncertainty, however, we do not give such context to the decision as either “action” or “inaction”, simply representing it by its associated outcomes under each state of the world.

Equivalently, the terms “action” and “inaction” can be referred to as acts of “commission” and “omission”, respectively, and it was under these terms that the effect of the context was explored by Spranca et al. [86]. Based on a finding by Kahneman and Tversky [49], showing that “...subjects felt more regret when bad outcomes result from action than when they result from inaction” [86, p79], Spranca et al. define the idea of an *omission bias* as being present “...when [subjects] judge harmful commissions as worse than the corresponding omissions.” [86, p79]. Their experiments find evidence for the existence of an omission bias in decision making, but, importantly, only consider the idea of a “worse” action from the point of a morality judgement and “...overall goodness” [86, p91], and not a specific emotional response, such as regret. Whilst this provides evidence that individuals consider the act of commission worse than the act of omission when the consequences are negative, it does not necessarily imply that the individual would regret one more than the other.

From this starting point Ritov and Baron [78] develop a series of experiments which specifically ask about the experience of regret, as it relates to both omissions and commissions. Their experiments involve hypothetical scenarios and subjects being asked to rate “satisfaction with the chosen decision” for both omissions and commissions, finding that satisfaction “...was worse for acts [compared to omissions] with full knowledge, when it was revealed that the outcome of the foregone option was even better” [78, p126] implying that “... in some situations, anticipated regret is greater for acts than omissions.” [78, p127] This evidence that *both* the experience and anticipation of regret are linked to the context of the decision, via the omission bias, is absent from standard formulations of regret theory.

So long as the omission bias is consistent, however, the correction needed in regret theory is a fairly simple one. It requires segregation of options into omissions and commissions, and an understanding that the role of anticipatory regret on decision making will be substantially smaller for an omission than the corresponding commission. The focus of this work, however, is the feedback from experienced regret to subsequent choice behaviour, and Zeelenberg et al. [110] analyse the constancy of the omission bias¹ from this perspective. By manipulating feedback from a previous decision, in a hypothetical scenario², participants in the experiment were asked to rate whether more regret would arise from inaction or action, given a positive or negative prior outcome. When the feedback was positive (or absent) the typical omission bias was present. When the feedback was negative, however, the bias was reversed, as the inaction was reported to lead to more regret than the corresponding action. The authors attribute this change (from an “action effect” to an “inaction effect”) to differing levels of *responsibility* in each scenario³, demonstrating, once again, that the complex context of a decision under uncertainty can have a significant effect on the role of both experienced and anticipatory regret.

3.2.2 *The Status Quo Bias*

When we move from a one shot to repeated decision making scenario, the omission bias is often demonstrated as a “status quo” bias. Typically, the decision not to act will preserve the status quo, whereas the decision to act will change it. Samuelson and Zeckhauser [82] outline a series of scenarios, and run a sequence of experiments, which demonstrate the status quo bias, with some of the decisions taken by participants involving an element of uncertainty, and hence exposing them to potential regret⁴. By presenting two different versions of each problem (one where there is a pre-existing status-quo option and one where that option is framed as a new decision like any other), they show a higher proportion of “status-quo” choices being made in the majority of situations, “...demonstrat[ing] the presence of [a] (statistically sig-

¹ which they refer to as the “action effect”

² in Zeelenberg et al. [110] Experiment 1, participants were told that two football teams had either won (positive feedback) or lost (negative feedback) their previous match 4-0. In the next match, both team lost 3-0, but one team had made no changes to the previous match (inaction), whereas the other had made 3 changes (action).

³ the football coach who doesn’t change a team after a heavy defeat should feel more *responsible* for the next defeat than the coach who at least tries to change the situation.

⁴ questions 2, 3 and 4 in Part One have the potential for what is termed in the paper as “...decision regret” [82, p38], with the other questions either being taken under certainty, or failing to show the resolution of the uncertainty from the gamble not taken.

nificant) status quo bias across decision tasks and across alternatives within decision tasks.”[82, p15]. Additionally, they present a two-stage example where the status quo is endogenously determined by the choice in the first stage⁵, and find a difference⁶ between the choice of individuals in stage two, given they had previously had anchored themselves to a particular option in stage one, compared to those who faced the decision of stage two as the decision in stage one (i.e. had no previous decision with which to anchor themselves). Samuelson and Zeckhauser “...conclude that the sequential decision tasks show some evidence of a status quo bias, most prominently in cases that involve many alternatives.” [82, p26]

In this situation, there are many possible causes for a status-quo bias, with Samuelson and Zeckhauser separating them into three main categories⁷ of which “[a]voidance of decision regret is ... one cause of status quo bias.” [82, p38] These results, coupled with the action/inaction bias, imply a different level of anticipatory regret associated with choosing the status-quo again, compared to taking a new alternative, without being able to isolate and separate the different possible causes from one another.

Much of the literature on the status quo bias, however, discusses the bias without reference to any feedback on the status quo option, including the two stage example given by Samuelson and Zeckhauser. Feedback, however, is a crucial determinant of the experience of regret, in that good feedback confirms you made a good choice, and have no reason to experience regret, and bad feedback indicates there was a superior choice, and so it is reasonable to experience regret. Inman and Zeelenberg [42] extend the work of Samuelson and Zeckhauser by studying the status quo bias, conditional on the feedback which is received as a result of the chosen option (which becomes the status quo in stage two) in stage one. Their hypothesis is that should a “...strong reason (e.g., a negative-experience episode) [lead] to the need to switch, this should cause a reversal of the status quo effect ... [as] in the case of negative feedback regarding the earlier outcome.”[42, p118] Thus they “...predict that when negative information on the current course of action is experienced but no avoidance action is undertaken, more regret will be experienced if the subsequent outcome is also negative. In such instances, changing the status quo should be preferred to maintaining it.”[42, p118] By asking experimental subjects to rate the degree of regret they would feel (on a ten point scale) in a number of hypothetical scenarios, they demonstrate “...the status quo effect reverses in the negative-prior-experience condition”[42, p120], implying that when individuals have a justifiable reason for deviating from the status quo, they anticipate less regret from doing so, and, indeed, anticipate feeling more regret from sticking with the option which previously let them down.

Both of the status quo experiments described above, however, rely on subjective ratings of hypothetical scenarios for generating a measure of

⁵ Section 1.4 ; Sequential decisions [82, p22]

⁶ the difference was statistically significant, in the hypothesised direction given the assumption of a status-quo bias, in 2 out of the 4 variations of the two-stage example. In the other 2 variations, a confounding factor caused a significant result against the status-quo bias in one, and the degree of significance was too small ($p = 0.35$) in the other.

⁷ “The effect may be seen as the consequence of (1) rational decision making in the presence of transition costs and/or uncertainty; (2) cognitive misperceptions; and (3) psychological commitment stemming from misperceived sunk costs, regret avoidance, or a drive for consistency.” [82, p33]

regret as it relates to the status quo bias. We would prefer an objective behavioural measure of the link between the status quo bias and regret, but, as Samuelson and Zeckhauser note, there are many confounding explanations, aside from regret, which can give rise to a status quo bias. In order to truly investigate the link between status quo choices and regret, distinct from other possible confounding explanations, it is necessary to move beyond inferring regret from behavioural decisions to an explicit, objective measure of regret, which can be recorded when an individual is asked to either choose, or deviate from, the status quo.

The recent developments in neuroscience, in identifying areas of the brain associated with both the experience and anticipation of regret (Camille et al. [16], Coricelli et al. [19], Chua et al. [17]), provide such a platform for delving further into the relationship between the status quo bias and regret. Nicolle et al. [66] conduct such a study by “...explor[ing] how asymmetric behavioural and brain responses for errors after rejecting, or accepting, a status quo option may be associated with a status quo bias on subsequent decisions.” [66, p3320] The study uses a repeated perceptual judgement task, which had been shown, in a prior study, to elicit a status quo bias towards accepting the previously selected option, and tested hypotheses, generated from the existing neuroscience regret literature, “...that error-related brain responses would be greater for erroneous status quo rejection than for erroneous status quo acceptance” [66, p3322]. The first result they found connected the difference between the acceptance of the status quo, leading to an error, and the rejection of the status quo, leading to an error⁸ and the experience of regret at the time of outcome feedback, finding that “...activity in the medial prefrontal cortex ... showed greater responsivity to reject [of the status quo] errors compared with accept [of the status quo] errors” and “...left anterior insula activity was significantly greater for reject status quo than to accept status quo errors.” [66, p3323] This evidence is, in one sense, consistent with a difference in the experience of regret, in that previous studies (Chua et al. [17]) have linked the anterior insula to regretful experiences, but inconsistent in that the orbitofrontal cortex (OFC) is most commonly associated with the experience of regret (Camille et al. [16], Coricelli et al. [19]), and yet no evidence was found for OFC differences at the time of outcome feedback. A fair conclusion to draw is that, whilst neuroscience techniques have progressed the understanding of experienced regret substantially, we are not yet in a position to “measure”, using fMRI data, the experience of regret, and hence draw concrete conclusions linking the experience to behavioural choices.

3.2.3 Responsibility

A possible underlying reason for both the omission and status quo bias is the associated *responsibility* that the individual would attach to the decision they made, and hence the outcome which resulted. In the case of the omission bias, the individual can rationalize the bad outcome, assuming they choose not to act, by imagining⁹ that such a situation would have occurred whether they were present, and chose not to act,

⁸ the status quo bias predicts that an individual will anticipate stronger regret as a result of rejecting the status quo, and being wrong, compared to accepting the status quo, and being wrong, leading to a preference for the status quo.

⁹ more commonly, in the language of regret, this type of imagining is known as *counterfactual thinking*

or not. Similarly in the case of the status quo bias, the individual can imagine that the status quo is simply a reflection of previous ideas about best practice, and hence choosing to follow the status quo is as much a reflection of the decision of others as it is of the individual themselves. Both of these situations reflect a reduced responsibility of the decision maker for the choice which was actually made, which the decision maker can use to lower or mitigate their experienced regret should a bad outcome result.

However, in the same way as the omission and status quo biases can be thought of as *characteristics* of the decisions taken, we can also consider the responsibility that the individual attributes to the choice as merely another characteristic of the decision. In this framework, it is highly likely that any decision making situation which either contains “inaction” or “status quo” as one of the options will have a variation in the degree of responsibility that the individual associates to the choices on offer, making the combined effects on anticipated regret difficult to separate. An easier starting point is to pose the question “is responsibility for a decision a necessary condition for the experience of regret?” That is, if we were to remove the characteristic of responsibility¹⁰ from a situation, would it still be possible for an individual to experience regret?

This question was initially addressed by a series of experiments (Zeelenberg et al. [108], Ordóñez and Connolly [67]), which placed the experiment participants in worded hypothetical scenarios and asked them to subjectively rate the emotions that would be felt by the actors in those scenarios. A typical scenario would have an actor experiencing a bad outcome as a result of their own choice, which was to be compared against the same scenario where the actor had the same choice imposed upon them by an external agent, manipulating the degree of responsibility the actor would feel for the outcome of the scenario. Once the overall emotional ratings had been disaggregated into separate components (regret, disappointment, happiness etc.), the experiments found that the absence of responsibility severely reduced, but did not totally eliminate¹¹, the experience of regret for the actors in those scenarios.

The validity of these results, however, can be questioned through the typical economists’ critique, that by asking the experiment participants to place themselves in the shoes of another person, and subjectively rate hypothetical emotional outcomes, there is little incentive for the experiment participant to either truthfully report their belief about the emotional experience or exert the necessary effort to correctly imagine what it would feel like to be placed in the hypothetical situation, assuming it was possible at all. A superior approach would be to design an experiment which constructs a decision making scenario in which the participant must make choices for real, and hence experience first hand the emotions associated with the outcomes of their choice. In order to analyse the effect of responsibility, however, we would need a control group where the participants must choose their own actions, and a treatment group where the actions *that they would have taken anyway*, are selected for them. In such an experiment, knowing which actions the participants would have been chosen, without asking them

¹⁰ equivalently known in the literature as *decision agency*

¹¹ which may be, in part, due to participants misunderstanding the implied theoretical distinction between, for example, regret and disappointment

to make a choice, is an incredibly difficult task, and hence there have been no papers which use this exact approach. In addition, we are still reliant on self-reported measures of experienced (and possibly anticipatory) regret, which still fail the typical critique of the lack of an incentive to truthfully report.

One recent paper, which seeks to overcome both these issues simultaneously, is that of Nicolle et al. [65] where two key ideas are combined to provide a better understanding of the relationship between responsibility and regret. Firstly, coupled with behavioural analysis, fMRI techniques are used to assess whether patterns of brain activity, known to correlate with the experience of regret (Camille et al. [16], Coricelli et al. [19], Chua et al. [17]), are varied when the level of responsibility an individual associates with a decision is manipulated. Secondly, this variation in responsibility is achieved through a voting mechanism, whereby the probability that the preferred action of an individual participant is actually chosen can be lowered by having a majority voting rule and increasing the number of people involved in the vote. For example, when there is only one person in the vote, the individual is wholly responsible for the decision, but when there is three people in the vote, any one individual is, at most, 50% responsible¹² for the decision taken. Furthermore, by making participants choose gambles, the outcomes of which are financially linked to the payment the individual would receive for the experiment¹³, there is an incentive for participants to play according to their true preferences. The results of this work found that “...regret-related neuronal activity in the amygdala was enhanced by increased responsibility, suggesting a critical role in “self-blame regret”” and also they “...did not find any brain regions responding to what has been termed “outcome regret,” i.e., showing invariant responses to regret-related outcomes under all levels of responsibility.” Nicolle et al. [65, p187] The finding of activity within the amygdala increasing under increased responsibility corresponds well to other experiments which have demonstrated a link between the amygdala and experienced regret, suggesting this neuroscience approach to regret studies, whilst somewhat currently untrustworthy given the very small number of studies conducted in this fashion, may prove increasingly fruitful as more commonly observed behavioural responses to regret are investigated in this way.

3.2.4 *Alternative explanations*

An alternative explanation for the status-quo bias is that of loss aversion (Thaler [91], Kahneman and Tversky [50]). Loss aversion reflects the idea that people experience a greater disutility from a loss than they experience utility from a gain of equal magnitude, and is a core component of the value function found in Prospect Theory. Importantly, a loss is considered relative to a reference point. In the situation where there is a default or status-quo option, the reference point is fairly clear. An action which may result in an outcome worse than that of the status-quo can be considered as a potential loss relative to that

¹² as the lowest majority in a three person society is two people, thereby attributing 50% of the responsibility for the act which was chosen to each of the two people who voted for it.

¹³ “Participants received 50p for each percentage they won of the maximum points they could have won in their game” Nicolle et al. [65, p180]

reference point. In this case, loss aversion drives choice to maintaining the status-quo.

A related concept, of norm theory, is explored by Kahneman and Miller [47], where "... (o)utcomes are perceived as worse when subjects can easily imagine that a better outcome could have occurred." (Ritov and Baron [77, p50]). Where a clear position exists against which to compare the alternative outcome (in the case of the status-quo or the decision not to act), then magnifying the negative consequences of the alternative has the behavioural effect of reducing the likelihood of the alternative being chosen. Thus, the status-quo bias and the omission bias will appear if the status quo or the decision not to act are easier to imagine.

These explanations for the status-quo and omission bias run in parallel to the explanation of responsibility. For example, if you are investigating choice under risk through the lens of Prospect Theory, the explanation of loss aversion makes most sense as loss aversion is a core component of the value function used in Prospect Theory. In this thesis, however, we are looking through the lens of Regret Theory, which does not include loss aversion as a "standard" assumption relating to the utility functions used. Hence, the explanation of loss aversion cannot exist within the world of regret. In contrast, the explanation of responsibility is more plausible in this world, as regret only arises from the action of the agent (I can't regret something someone else did), and "taking responsibility for your action" is a psychologically sound idea. Thus, whilst acknowledging the existence of other explanations for the status-quo and omission biases, the focus in this world, and hence in this work, is on responsibility.

3.2.5 *Summary*

The above research demonstrates that the specific characteristics of any decision taken under uncertainty can have a strong impact on the role that both anticipatory and experienced regret will have on that decision. As regret is an emotion which relies on an association to a choice to exist, it is reasonable that characteristics of the choice, beyond simply the easily quantifiable "outcomes and associated probabilities" described by traditional regret theories, will play a significant role. The three characteristics described above, however, are not a complete list of how mathematically equivalent options can be "framed" to appear different to the individual, so we should continue to explore other avenues which will have an impact on both the applicability, and predictive capability, of theories which make use of regret. It is in this spirit which we construct the following experiment.

3.3 WHEN AN ACTION AND A DECISION DIVERGE

The broad theme of this work is moving studies of regret from static to dynamic contexts, and so we can additionally study the above results in terms of their applicability to repeated decision making scenarios as well as one shot scenarios. The status quo bias is commonly framed in terms of repeated decisions, in that there should have existed a previous decision for one option to be termed the "status quo". A status quo option, however, will be likely be stable over time, which presents problems for experimental manipulation where an experimenter would

like to manipulate the short-term dynamics of a situation in order to observe changes in behaviour. Secondly, the inaction bias is only defined in terms of a single decision, in that a decision not to act will yield less regret than an equivalent decision to act. As such, there are unanswered questions regarding the persistence of the inaction bias over time, and whether the reduction of experienced regret from inaction, compared to action, would result in a tendency to repeat, and not learn from, mistakes and bad choices resulting from inaction. Lastly, the effect of responsibility is hard to analyse dynamically, because a reduction in responsibility is typically associated with having choices imposed upon an individual, rather than requiring an individual to make a choice, and, as such, there is little behavioural analysis to conduct or observe when an individual is not being asked to make a choice.

As a result of this complexity in extending the existing literature to a dynamic or repeated context, especially in the laboratory, we need to start with much simpler, and much easier to define, characteristics of a choice when seeking to move from a one shot to repeated or dynamic context, in analysing how those specific characteristics interact with the feedback from experienced to future anticipatory regret and future choice.

3.3.1 *The action*

In economics, we often consider a choice as a comparison of costs and benefits amongst options. These costs and benefits can be quantified (according to some measure), and the comparison of the measures will, typically, yield a “best” option amongst the group.

The process described above is, however, a purely theoretical exercise. It is entirely possible to evaluate options without ever needing to make a real choice. This is a hypothetical choice.

If we want to actually make a real choice in the real world there needs to exist another component. This other component is the “action” we would need to take to effect a choice.

This component arose in the previous discussion of the inaction bias, where “inaction” was, in essence, doing absolutely nothing at all, and an “action” is doing something, anything, physical to effect a different choice.

Actions, defined as the physical process necessary in order to effect a choice, can be thought of as incredibly wide ranging.

From a typical consumer perspective, some examples would be:

- Going to the shop to purchase a particular item or particular brand
- Phoning a particular insurance company to purchase insurance for a product

From an investment perspective, there are similar examples:

- Going on a website to buy or sell a particular number of shares in a specific company
- Make an offer, through an estate agent, to buy a house

In the lab, we typically denote actions by selecting a specific option on a computer screen, which will have a label; often a title or name, a location on the computer screen, and perhaps even a colour.

All of these examples above have a common theme running through them. They consist of both a physical action and a label which represents the choice being made.

In the first example given, the decision to purchase a healthy snack may have been considered as the best option amongst many in order to meet the need of being hungry, but this is not a real world choice until an action is taken (“going to the shop”) and a label is picked representing the choice (“a single Dole banana”).

For the purposes of the discussion going forward, the term “action” will encompass both the physical action taken, and the label of the option representing the choice.

In short, economists would tend to think of “actions” as being unimportant in the context of decision making, in that they convey very little information as to the quality of the decision being taken. For example, if a lab experiment simply produced an output which stated “At time period t , the option labelled x was chosen, as opposed to the option labelled y , by clicking an icon on the left of the screen” then there would be nothing to conclude from the data without understanding the decision which choosing action x over action y represented.

However, aside from cases of inaction, every decision taken (including decisions under uncertainty) can be thought of as having an associated action necessary to effect the decision.

3.3.2 *The decision*

In contrast to the action, the “decision” can be thought of as the relative comparison of economically important information between options available to the individual. In the traditional microeconomic sense, the decision made reflects the individual’s best response to the various benefits and costs on offer. The decision taken by the individual conveys to the observer a sense of what’s important and relevant to that individual, and, as such, it is decisions, rather than actions, that economists are typically most interested in, because it reveals how the individual responds and reacts to incentives. For example, a laboratory output which states “At time period t , the individual decided upon the risky option, forgoing the safe option” gives us more useful information than the equally correct output of “At time period t , the individual decided upon action x , forgoing action y ”, because it provides meaningful information as to how the individual responded to the incentives which were placed in front of them.

Whilst it must be true that, in any decision problem, there is at least one action associated with at least one of the two decisions¹⁴, an economist would argue that, as the action alone carries no particular economic information, it could be changed or replaced by another, which effects the same decision, without any impact on the resulting choice behaviour.

The natural extension to this argument is that there should be no way to predict, *ex ante*, whether an individual will choose one particular action over another without also including, or controlling for, the underlying economic decisions that the actions represent. For if there

¹⁴ in the simplest possible case, of an agent facing a decision between 2 options, it may be the case that one decision can be implemented without the need for an action (i.e. inaction will result in that decision being made by default), in order to differentiate the two decisions, there must be an action associated with the other

were such a method of prediction, it would imply that an individual had an underlying preference for one action over the other, implying that one action represented an increase in utility over the other, and hence there was, in fact, economically relevant information contained inside the action which would alter the *decision* taken at the margin.

The reason for the making this distinction between “action” and “decision” was in relation to the characteristics of a choice which may interact with the influence of anticipatory regret. Characteristics have already been shown to be important with the status quo and inaction biases. In the same way we can ask “does anticipatory regret work in the same way for actions as for inactions?” we can ask “does anticipatory regret work in the same way for actions as decisions?”

The reason for posing this question is that when we say an individual will experience regret from the choice they made, it is not specified whether the action, as distinct from the decision, will be associated with the experience of regret. In a simple one shot decision situation, the answer to this question does not matter at all, precisely because the action taken by an individual is necessarily tied to the decision it effected. It is impossible for me to decide to do something and subsequently choose an action which effects a completely different decision¹⁵. As such, it is impossible to determine whether anticipatory or experienced regret is tied to either the action or the decision or both, because, behaviourally, it will appear as the same thing.

The issue changes, however, when we move from a one shot situation to a multi stage situation. In this world, the previous action and decision can diverge, as “clicking on the blue icon on the left side of the screen” may no longer represent the same decision in stage two as it did in stage one. In this world, assuming that experienced regret has an impact on subsequent decision making¹⁶, it now matters whether the previously experienced regret was associated with the action or the decision in terms of the impact of the regret on the stage two decision. It will have an effect when only one of the previous decision or action turns up in stage two, but it will matter most when both the decision and action, from stage one, turn up in stage two, as opposing forces on the decision maker. It is for this reason that the experimental configuration outlined below will be used to analyse the problem at hand.

3.3.3 *The influence of experienced regret*

Experimental Literature

Experimental studies which have analysed the effect of experienced regret on subsequent decisions taken under risk or uncertainty have tended to assume one of two things.

Firstly, some (Zeelenberg and Beattie [105], Creyer and Ross [21]) assume the experience of regret will influence subsequent regret aversion, which implies that the regret function, to be used in the second decision making stage, is affected by previous regret, which changes the relative attractiveness of options in the stage two decision. In the absence of evidence to prove whether an observed change in preference towards one option, say A, and away from another, say B, is caused

¹⁵ the assumption here is that we are ruling out “the trembling hand” often seen in game theory, whereby an individual must take into account that they may make an unintentional decision with very small probability

¹⁶ the literature of which was extensively analysed in the previous chapter

by an *increase* in the anticipated regret associated with option B, or a *reduction* in anticipated regret associated with option A¹⁷, these studies have tended to assume that observed deviation *away* from any particular decision is evidence of an increase in regret aversion, due to the “similarity” between the decision which was initially regretted and the subsequent decision which the individual moved away from.¹⁸

Secondly, some (Raeva and van Dijk [73]) assume the experience of regret will influence the subjective probabilities associated with “similar” anticipated regrets in a subsequent decision. That is, having experienced a regret once, I believe it is *more* likely that I will experience a similar regret (as a result of taking a similar decision, and having uncertainty work against me) again. Compared to the first approach, this approach has the added benefit that if you believe the sum total of all subjective probabilities must equal to one (i.e in a two option, two state of the world model, an increase in the likelihood of one regret necessarily implies an equivalent decrease in the other) then your observed experimental result, of a shift in the *relative* preference for one decisions over another, will necessarily either support the hypothesis or refute it.

The problem with the two above approaches, however, is that, experimentally, they produce behaviourally equivalent hypotheses. Given I observe an increase in the relative preference for option A over option B, I would be unable to determine whether it was due to a belief that regret associated with option B was *more severe* (assumption one) or *more probable* (assumption two). This problem is widely addressed in Chapter 2.

On common themes

As the above studies are experimental, and not theoretical, in nature, they do not make specific predictions about the impact of experienced regret on subsequent choice in contexts other than those which are very similar to the precise experimental design used in the papers. For example, as the experiments do not make a distinction between the actions and decisions taken, they do not provide insight into how the experience of regret would differentially impact the various components of choice in the second stage.

However, the three experiments mentioned above have certain “common themes” which underlie the authors’ implicit assumptions about how the experience of regret will influence choice in a subsequent stage. In each case, there is an underlying notion of similarity, in that the experience of regret, from a particular choice in stage one, will only affect subsequent choice as it pertains to “similar” choices in stage two. In the case of a repeated decision, this is, in effect, saying that the experience of regret, from choosing a particular option, provides information *only about that same option*¹⁹, when it appears again in the second stage. This intuition, however, does not tell us what will happen, in the second stage, when there are elements of similarity associated with more than one option. That is, if the action and decision from the first stage diverge into separate options in the second stage.

¹⁷ as discussed in Chapter 2 of this work

¹⁸ it’s easy to imagine that an individual doesn’t want to repeat the same (or similar) mistake twice, and so would wish to move away from the decision they initially made, though this is only an assumption

¹⁹ in the form of either a change in the experience of regret associated with that option, or a change in the subjective probability of experiencing that specific regret again

Secondly, though all of the experiments use an idea of similarity, they offer no explanation as to how, or why, that element of similarity interacts with the changing decision in the second stage. By explicitly assuming the experience of regret impacts the regret aversion, or subjective probability, at stage two, they are saying that the experience of regret impacts, in an easily quantifiable way (i.e. in a way that could be modelled within existing regret-based theories) the second stage decision, but not what that impact is conditional upon. Does a *higher* degree of similarity between the previously regretted choice and a possible option at stage two, for example, imply a *larger* change in regret aversion or subjective probability? Or is there simply a constant change if the decision maker deems the second stage to be similar, in some way, to the first stage? Or will such an effect only exist if the decision, and action, combined, is identical in stage two as it was when it generated the experience of regret in stage one? Such questions are left unanswered by these experimental works.

Theoretical Literature

The theoretical literature which looks at the role of regret in dynamic contexts is generally confined to the field of learning, and adaptive behaviour, where the experience of regret is providing feedback about the underlying uncertain environment which the individual is making decisions under. For example, “regret-matching” (Hart and Mas-Colell [32]) uses regret as a tool for learning about the best-response in a game theoretic setting, by adapting the probability of playing any given strategy by the measure of experienced regret from playing such strategies in the past. However, in our experimental environment, we are simply looking at decision theory, rather than game theory, and there is no information to be learnt, about the underlying environment, from the past experience of regret, as the individual is provided with complete information at every decision stage. As such, regret-matching will not offer any predictions about behaviour in a complete information environment.

The most recent theoretical work which uses anticipated regret in a dynamic context is by Hayashi [35], focussing on dynamic consistency of choices in the presence of anticipated regret. This work, however, focusses on the “...opportunity dependence property of regret-based choice” and “...maintain[s] the assumption that the decision maker looks only at future, and does not care about what might have occurred at unrealized events” [35, p402] thus specifically ignoring any feedback effects from prior experienced regret on future choice. Consequently, this theory will too not offer any predictions on the difference in behaviour as a result of the distinction between anticipated and experienced regret.

Finally, the Theory of Regret Regulation by Zeelenberg and Pieters [106] aims to be the most descriptive account of the different ways in which regret affects decision making behaviour under uncertainty. It builds on the literature spawned from the original formulations of regret theory to suggest 10 propositions²⁰ which describe the most crucial aspects of regret for decision making, building from the premise that “...consumers are regret averse and that, as a consequence, they try to

²⁰ subsequently revised to 11 propositions in version 1.1 of the theory (Pieters and Zeelenberg [68])

regulate their regrets” to form “...a single overarching model”[106, p3]. In terms of a prediction for the difference in effect between action and decision regret on subsequent behaviour, Proposition 10 states “[r]egret regulation strategies are decision-, alternative-, or feeling-focused and implemented based on their accessibility and their instrumentality to the current overarching goal.”[106, p4] This terminology is then expanded upon later in the article, but no precise prediction is made about the effect of experienced regret on subsequent decision making. However, there is mention of what can be called ‘the most common assumption’ that “...[experienced] regret most clearly induces decision reversals or undoing behaviors”, “[a] central element in the experience of regret is the undo or reverse the decision that led to the regretted consequences” and “...the prediction that regret promotes switching” [106, p13]. Whilst these three statements do not make the same precise distinction between decision and action regret, at the very least they do seem to indicate a desire for the individual to move *away* from an action/decision which previously caused them to experience regret. Applied to consumer purchases, this idea “...can motivate us ... to switch to another supplier of services or product the next time around”[106, p13]. As such, it can be stated that an experiment which shows a preference to switch away from a previously regretted decision or action can be thought of as *consistent* with Regret Regulation Theory, though not specifically predicted by it, but an experiment which shows a preference to move towards a previously regretted decision or action can be thought of as running *contrary* to Regret Regulation Theory.

Transfer of knowledge and similarity

One potential avenue to bridge the gap between the theoretical and experimental literatures on the impact of experienced regret on subsequent regret aversion is the parallel literature (which sits across both economics and psychology) on transfer of knowledge²¹. This literature studies the mechanisms via which the lessons learned from a decision made in one context are *transferred* into a future decision which is made in a similar, but not identical, context.

This literature is immediately appealing, as the concept of “similarity” is especially important, and there is both theoretical and experimental evidence in support. Whereas the experimental literature on regret, discussed earlier, struggles to define the components of a “similar” decision, yet relies heavily on it to explain the experimental results, the transfer of knowledge literature is able to precisely quantify how similarity can explain the behaviour of individuals in subsequent choice periods. For example, Zizzo finds “...just three variables ... are required to jointly be able to predict 1/3 of the variance in similarity evaluations” (p21) in the context of a range of game theory situations. In the context of regret, if we were able to quantify how individuals consider regrets to be “similar” to each other (whether they are evaluated according to the similarity of the action, decision, or other dimensions), then we can begin to make predictions which use the similarity as a proxy for the salience of an anticipated regret in the decision problem.

However, a cautionary note is struck if we consider the more fundamental question of *how* knowledge is transferred, and not simply the question of *when* knowledge is transferred. By asking the question of

²¹ wide ranging articles on the topic are contained in Zizzo [111]

how knowledge is transferred, we are equivalently asking the questions: what is being learnt through the experience of regret, how is this being encoded into memory, and how it is then recalled to be used at a later date? In the case of game theory, as discussed earlier with reference to similarity, and also as in “regret-matching” (Hart and Mas-Colell [32]), the underlying assumption is that you are using the experience of playing the game to learn how to play the game *better*. Doing so, you are able to overcome the limits of bounded rationality to improve your utility, and the method by which this is accomplished is obviously of great importance. In the language of de Jong and Ferguson-Hessler this is “deep knowledge” which permits “...vertical transfer of knowledge” [111, p6] to contexts outside of that in which the original learning took place.

In the world of regret, however, it is certainly not clear what is being learned, or how this learning can be used to improve your overall utility. The first chapter of this work looked at what this world may look like if the experience of regret is being used to learn more about the regret function itself, or equivalently, an individual’s reaction to the emotional consequences of utility gaps²², but it wasn’t clear that doing so worked to the benefit of the individual²³.

But in the case of simply studying whether the emotion of regret is attached to an action or a decision, it is simply not clear whether this is at all to do with the process of learning to ultimately improve utility. It is much more related to the ideas of emotional residue from past choices and increased sensitivity in the face of future choices, which is a much more *superficial* notion than considering a *deep* transfer of knowledge method of procedure. The question seems much more one of *reaction* than *reflection*, and hence the extent to which the transfer of knowledge literature can be of use in helping to understand the process may well be limited, at least at this time.

3.3.4 Summary

No previous experimental or theoretical work, which studies how the experience of regret will impact subsequent choice behaviour, has made the distinction between a decision and an action, as defined earlier in this work. As such, there is no previous work which makes an explicit prediction about whether the experience of regret is tied to the decision or the action taken, and consequently the effect of experienced regret on subsequent choice when the decision and action, from a previously regretted choice, diverge at a second stage. Several works make implicit predictions and assumptions, which can be extended to a context when the choice of an individual is split into the action and the decision, suggesting an important role for “similarity” and a desire to “switch away” from options which have previously caused the individual to experience regret. The transfer of knowledge literature highlights the importance of similarity in how learnings are taken from one context and applied to another, but the absence of any obvious theories about how the experience of regret is used to *learn*, over and above simply the more reactive elements of the emotion, may limit how much this literature can be used to help with the issues at hand.

²² “what I got” versus “what I could have had”

²³ indeed, the process of learning about the regret function appeared to encourage them to take more risks for fear of missing out on the big prize

3.4 DESIGN OF AN EXPERIMENT

The previous chapter of this work was dedicated to explaining the difficulties associated with designing an experiment to test how the experience of regret affects subsequent regret aversion. The above distinction between the regret associated with a decision, as opposed to an action, does, however, provide a natural framework for designing an experiment to investigate whether or not there is a link between the regret associated with a previous *action* and the behaviour of an individual when faced with that *action* again in a subsequent stage.

In order to answer such a question, we need to design an experiment which compares the choice of an individual, under uncertainty, in the presence of a previously regretted action, to the choice of an individual, under uncertainty, without the presence of a previously regretted action. Noting, however, the distinction between an action and a decision, it is important to hold constant the decision being made by the individual, whilst changing the action that the individual will see. In order to do so, however, we will need to find a mechanism and context in which it is possible to disaggregate the decision from the action, so that we can present two different versions of the same decision, in the second stage of the experiment; one in which the previously regretted action is repeated, and one in which the previously regretted action is not.

3.4.1 *An intuitive example*

An example of an action, which can be thought of as distinct from a decision, taken under uncertainty, is that of purchasing a *brand*. Consider, for example, a multinational electronics corporation, which produces a number of different electronic goods in a range of different markets. It may be the case that, for one particular good, the brand carries a relatively good reputation for quality and reliability, but, for a different type of good, the brand is considered to be of relatively low quality and reliability. In such a situation, the experience of regret, from choosing one of the goods, and having it fail, provides no information about the probability of failure about the other type of good²⁴, or the utility that will be gained from owning the other type of good. As such, in purely the context of the *decision* as to whether or not to buy the other good, there is no “Bayesian information” to be gained from the experience of regret from the original good of the same brand. However, the *action*, in this case, corresponds to the individual deciding to spend their money, a second time, on the same brand which had previously caused them to experience regret.

If there is such an effect, as a result of regretting the purchase of a particular brand when faced with a decision whether or not to purchase the brand again, it will impact most when the decision, that the first purchase decision was based upon, is the reverse at the second purchasing opportunity compared to the first. This is because the effect of regret from the *decision* at the first stage will be acting in the opposite direction, on the choice at the second stage, to the effect of regret from the *action* at the first stage. This idea and effect can be illustrated with an example

²⁴ under the assumption that, because they are different goods, they are likely manufactured in different factories, and designed and engineered by separate teams

Consider the following situation: You are in an electronics shop, and have to make a decision between a cheap TV, made by a company you don't recognise, and a more expensive TV, which is made by LG.

The 2 TVs appear to have the same characteristics (same size, picture resolution and a 12 month warranty), and produce a similar quality picture, but, fearing that the cheap TV might be unreliable, you decide to purchase the LG TV.

As it happens, the LG TV breaks down shortly after the 12 month warranty expires, and you feel that you might as well have purchased the cheap TV instead, and saved yourself some money.

Shortly thereafter, you need to replace your mobile phone, and have narrowed your choice down to a cheaper LG mobile phone, and a more expensive Apple iPhone. Again, both phones appear to have the same characteristics (same size, picture resolution and a 12 month warranty), and are similarly enjoyable to use and navigate.

You know that the iPhone is the more reliable brand in mobile phones, and also know that the fact your LG TV broke down doesn't make it any more likely that your LG mobile phone will do so too, but are faced with a tough decision.

Previously, you had purchased an LG product, and ended up wishing you had selected a different brand.

At the same time, however, you had purchased a more expensive, supposedly reliable product, and ended up wishing you had taken the cheaper option.

Does the fact that you had a bad experience with the LG TV make it less likely you will purchase the LG mobile phone?

Does the fact that you had a bad experience purchasing the more expensive, better brand TV make it less likely you will purchase the more expensive, better brand mobile phone?²⁵

In this example, purchasing the LG brand a second time corresponds to repeating a previously regretted *action*. Purchasing the Apple iPhone at the second stage, however, corresponds to repeating a previously regretted *decision*, in that the individual would be choosing to spend additional money on a supposedly more reliable product, at the risk that it may break down anyway and leave the individual wishing they had taken the cheaper option. As such, in this example, the two types of regret, from the first stage, are acting in opposite directions on the individual at stage two²⁶. The action regret is now linked to the LG phone at stage two, whereas the decision regret is linked to the Apple iPhone.

In order to isolate the effect of just action regret, however, we would need to compare the answers to the above questions to an equivalent second stage situation, where the individual must choose between

²⁵ this hypothetical question was posed to the subjects in the experiment as part of the post-experiment questionnaire, offering them the opportunity to give both answers and comments

²⁶ specifically without saying, however, whether the experience of regret from stage one (both in terms of action and decision) makes the individual more or less likely to repeat the previously regretted action or decision

the iPhone and another less reliable, less expensive mobile phone, equivalent to the LG phone in everything apart from its brand name, thus keeping the decision identical across both scenarios, but changing the action which the individual must effect in order to choose one of the options.

Keeping the decision constant, however, is very tough to do with such limited information. In such a hypothetical scenario, the probability that the less reliable mobile phone breaks down, for example, is likely to be conveyed by the brand name²⁷ and the individual's own experience or awareness of the brand prior to the experiment. Thus, in order to remove such subjective interpretation of the decision in the second stage, we need to ensure that any and all information conveyed is transparently objective, thereby ensuring direct control over the decision which participants are taking in the second stage.

Negative Reciprocity

The broad definition of the term "action" leads to many other possible explanations for what is observed as a change in choice at the second stage seeming to result from a "bad outcome" in the first stage. With specific relevance to brands, negative reciprocity would be a classic example. In the presented example, it is easy to imagine the agent feeling a sense of anger towards LG, as a company, and wanting to "hurt" them in the second stage by not purchasing their mobile phone.

Though typically reciprocity, both positive and negative, is seen between two people, in this case the "company" is a close substitute for another person, and hence it is plausible to explain the behaviour from a reciprocity perspective. As we move further away from a person along the spectrum of "actions", however, this explanation becomes less plausible.

Suppose in the presented example, the products are not branded, but simply sold in coloured boxes. The LG TV in stage one comes in a "shiny blue box" and the LG phone in stage two also comes in a "shiny blue box", but there are many other products, made by different companies, which come in a similar box. This breaks the link between the company and the product inside the box, hence eliminates the potential for reciprocity (as it may well be a different company selling the phone in the shiny blue box in stage two).

In this situation, the "action" required is "select the product in the shiny blue box" in both stages. Hence, if regret is associated with that action from stage one, it has the opportunity to influence choice at stage two. If there was a change in choice, attributing the change to reciprocity is less plausible. Indeed, generic labelling of this kind is what is used in the experiment in the remainder of the chapter.

An alternative formulation of the same idea can be found in the world of gambling, and, specifically, poker. The nice thing about poker is that we typically have four different labels for cards of equivalent value; for example, the four of hearts, spades, clubs and diamonds. Imagining a poker situation where you make a bet, for example that your pair of sevens (say, diamonds and hearts) is good enough to win the hand, you are making a decision to bet on the strength of the value which those cards represent (a pair of sevens is a mid-range pair), but also taking an

²⁷ thus implying that the brand name conveys some useful information to the individual, and hence breaking the purely "action" versus "decision" distinction made earlier

action to bet on exactly the cards labelled the “seven of diamonds and hearts”. Suppose you make this bet and lose. Now suppose that one of two things may happen. By chance, you receive the exact same pair as you did previously (seven of diamonds and hearts) or you receive the other pair of sevens (spades and clubs), and you are again required to bet on the strength of your cards. In a purely economic sense, they are of equivalent value, and hence the same choice should be made. Even in a traditional regret aversion sense, there is no reason to imagine that the anticipated regret from betting (and losing) on one pair would be any different to the other. However, the past experience of regret, tied specifically to the two red cards, may cause a different choice to be made in the comparison to the two black cards. Are you willing to give the red cards a second chance? Or do you believe that those cards only bring you bad luck? Again, if the emotion of regret is only tied to the previous decision (betting on the strength of a mid-range pair and losing), then it makes no difference. But if it is tied to the action (betting on the seven of hearts and diamonds and losing) then it may make a difference.

3.4.2 *An objective mechanism*

The simplest way to achieve direct control over the decisions, under uncertainty, faced by subjects in the experiment is to design the decisions to be as simple as possible. The less information that is provided to the subjects, the less room there is for subjective interpretation of the problem at hand²⁸. Similarly, if we wish to isolate and study the effects of regret, both experienced and anticipated, we must limit the number of different sources of regret in the experiment so that any effects can be precisely identified. These two criteria indicate that the best possible experimental design is a simple gambling task, where monetary outcomes and associated probabilities are explicitly defined, where a subject only has two possible options at every decision making stage. As such, let us define

- Gamble R - a p chance of winning $\pounds r$ and a $1 - p$ chance of winning $\pounds 0$
- Gamble S - a $1 - p$ chance of winning $\pounds s$ and a p chance of winning $\pounds 0$

In addition, in order for the decision to be non-obvious, let us define Gamble R as the “risky” gamble, and Gamble S as the “safe” gamble, so that $0 < p < 0.5$ and $r > s$. As such, given a choice between Gamble R and Gamble S, under an expected utility framework, all information that a participant needs to know in order to make a decision is objectively defined.

However, if we assume that a decision maker is also concerned with anticipated regret, in addition to the pure expected utility, as suggested by Loomes and Sugden [58] and Bell [5], then the decision

²⁸ This would be considered the traditional view of experimental economics. There exists a wider literature on experimenter demand effects (Zizzo [112]) which looks at how the construct of the experiment itself acts as a cue for a specific type of behaviour by the subjects. Whilst presenting a bland, neutral frame (Gamble A, Gamble B) works to reduce these effects in this specific experiment, this should not be taken as true for all experiments. Indeed, context may sometimes be required to counteract the wider effects of the experimental design.

maker also needs to know the probability that they will experience regret, conditional on each of the two options that they can take. If, for example, each gamble is resolved by a separate and independent mechanism (e.g. two separate random number generators), but the subject only observes the resolution of uncertainty for the gamble which they select, then the typical assumption made in the regret literature is that the individual cannot possibly experience regret, given they do not observe whether the non-chosen gamble would have resolved in their favour or not²⁹, and hence anticipated regret would not play a factor in such a decision. If, however, the individual observes the resolution of uncertainty for the unchosen gamble, as given by the second random number generator, then the probability of experiencing regret, conditional on having chosen a gamble which lead to a payoff of £0, is simply given by the unconditional probability that the unchosen gamble would have lead to a positive payoff had it been chosen³⁰. However, the effect of anticipated regret will be maximised when the outcome of the two gambles are dependent on each other; that is they are simultaneously resolved by just one mechanism (say, one random number generator) and, hence, the conditional probability of experiencing regret given the gamble you chose resulted in a payoff of £0, is equal to 1. As this experiment is designed to identify the possible effects of experienced regret on subsequent decision making, it is sensible to choose a design which maximises the potential for both the experience and anticipation of regret, and hence having the gambles resolved by just one mechanism will be implemented throughout.

3.4.3 Type A and Type B Regret

In Chapter 2, the idea of Type A and Type B regret was explained by considering the two following possibilities when making a decision to set aspirations or predict one's own ability

A) "I am worried that I might set my predictions or aspirations too high, end up not meeting them, and regret setting an unattainable goal"

B) "I am worried that I might set my predictions or aspirations too low, end up exceeding them, and regret not having more belief in my own ability"

In this experiment, however, the subjects will not be asked to bet on their own ability, or set aspirations, but rather, in the context of Gamble R and Gamble S presented above, will be asked to make a choice between a relatively "risky" and relatively "safe" option. As such, in order to continue to use the Type A and B regret terminology³¹, they need to be redefined in the context of risk seeking and risk aversion.

²⁹ it is occasionally argued that, in such a situation where an individual selects a gamble which does not resolve in their favour but observes nothing else, they may construct a "counterfactual", and experience regret from imagining what *might have* happened had they chosen differently, as opposed to experiencing regret from observing what *would have* happened had they chosen differently. However, it is typical that in experiments designed to analyse the effect of anticipated regret, this possibility is ignored.

³⁰ as an example, if the subject had chosen Gamble R, and lost, receiving a payoff of £0, the probability of experiencing regret would be $1 - p$, the probability that Gamble S would have paid £s had it been chosen.

³¹ remembering, that the definition of Type A and Type B regret was never intended to be fixed, but rather to highlight the fact that there is always at least two sources of regret, and the context of the decision specifies to what they correspond intuitively

The above definitions can be thought of in terms of how “conservative” the individual wishes to behave, in terms of their aspirations, with, for example, Type B regret equating to the individual worrying about being too conservative, and regretting not being more speculative. Under that premise, it makes sense, in the context of simple monetary gambling behaviour, to redefine Type A regret as

A) “I am worried that I may act in a risky fashion, which does not pay off, and regret not choosing a safer option”

and similarly, Type B regret as

B) “I am worried that I may act in a safe fashion, which does not pay off, and regret not choosing a riskier option”

In these terms, choosing the above Gamble R, the relatively risky option, exposes the individual to Type A regret, where as choosing Gamble S, the relatively safe option, exposes the individual to Type B regret.

3.4.4 *Generating regret*

By asking the subjects to choose between Gamble R and Gamble S in stage one, under the premise that both gambles are resolved simultaneously by a single mechanism, we will create four different groups of individuals after the end of stage one. These are

1. Subjects who chose Gamble R, and won, experiencing no regret.
2. Subjects who chose Gamble S, and won, experiencing no regret.
3. Subjects who chose Gamble R, and lost, experiencing Type A regret.
4. Subjects who chose Gamble S, and lost, experiencing Type B regret.

Thus, for the purposes of investigating the effect of experienced regret on subsequent decision making, we would, at first glance, like as many people as possible to fall into groups 3 and 4. This can be achieved in a number of different ways. The most obvious is to rig the experiment so that everyone loses, and hence either falls into group 3 or 4 depending on whether they chose Gamble R or S. Such deception, however, is typically frowned upon by economists when designing experiments, on the basis that repeated use of misinformation primes experimental subjects to come to expect information to be false, which means that, again, we lose direct control over the information which subjects are using to make decisions. In addition, this experiment uses multiple sessions over a two day period, and it would quickly become apparent, should subjects communicate outside the lab, that the presented probabilities for the gambles do not accurately reflect the true probabilities of any individual subject winning either of the two gambles, which will affect the behaviour of future subjects, who have not already participated, should this information be passed on. As such, deception is not an option for this experiment.

An alternative strategy is to make Gamble R, the risky gamble, simultaneously very attractive and very unlikely to win. This can be achieved by specifying the value of p to be very low and the value of $\text{£}r$ to be very high (relative to the value of $\text{£}s$). In such a situation, the expected

proportion of those who fall into group 3 would be close to $1 - p$, and, as such, we could easily restrict study to the effect of Type A regret on subsequent decision making. In theory, this seems like a sensible move³². Intuitively, however, there is a problem with presenting Gamble R as the “obvious” choice. In keeping with the literature linking the feeling of responsibility to the experience of regret, it is easy to imagine that “[p]eople expect a narrow margin of loss—or a “near miss”—to exacerbate self-blame, and thus they expect that margin to exacerbate regret as well” (Gilbert et al. [26, p346]) and, consequently, therefore, a large margin of loss, where a subject never had any intention of picking a “non-obvious” gamble, will reduce self-blame and responsibility, and thus reduce both the anticipation and experience of regret. As such, in an experiment which needs to generate the experience of regret in order to study its effects, there is a trade off between increasing the proportion of people who experience *at least some* regret and increasing the magnitude of regret for anyone who does happen to experience it³³.

The choice made in this experiment is to maximise the number of *marginal* subjects, who are close to indifference between Gamble R and Gamble S, in order to generate the maximum possible anticipation and experience of regret, given a fixed budget constraint from which to fund the experiment. As such, we need to find values of $\pounds r$, $\pounds s$ and p for which the modal subject is just indifferent between the two gambles. A hypothetical gambling questionnaire³⁴, and a small scale pilot of the experiment, found these values to be as follows

- Gamble R - a 30% chance of winning $\pounds 14$ and a 70% chance of winning $\pounds 0$
- Gamble S - a 70% chance of winning $\pounds 6$ and a 30% chance of winning $\pounds 0$

3.4.5 Resolving the gambles

Under the above definitions of Gamble R and Gamble S, a credible randomisation mechanism is needed to resolve both gambles according to the stated probabilities of 70% and 30%. A transparently fair randomisation mechanism³⁵ is that of rolling a die, but, in the case of generating probabilities corresponding to 70% and 30%, at minimum, a ten-sided die would be necessary. Fortunately, such a ten-sided die was owned by my supervisor, Dr Daniel Sgroi, and, therefore, this was used as the randomisation mechanism. This transformed the two gambles to be

- If you choose Gamble R, you will win $\pounds 14$ if the die lands on 7, 8 or 9, and $\pounds 0$ otherwise

³² even considering the potential financial costs of paying a small proportion of subjects a very large amount of money should they win Gamble R

³³ this trade off is also evident when considering that increasing the value of both $\pounds r$ and $\pounds s$ increases the magnitude of the consequences of making a wrong decision, and hence the experience of regret from making such a wrong decision, but, given a fixed budget from which to fund the experiment, also reduces the sample size which would be used in any hypothesis tests.

³⁴ completed by 29 students who attend a second year microeconomics course at The University of Warwick

³⁵ an example of a non-transparent randomisation mechanism is a computer random number generator, which, presented to subjects without the underlying code, could easily be, or at least believed to be, rigged to produce numbers which favour the experimenter at the expense of the subject

- If you choose Gamble S, you will win £6 if the die lands on 0, 1, 2, 3, 4, 5 or 6, and £0 otherwise

In order to ensure the privacy and anonymity of subjects in the experiment, however, it was not possible to physically roll a die in front of them to resolve the gamble. Similarly, to prevent the subjects from experiencing emotions such as envy, as a result of observing the reactions of other participants in a given session of the experiment, it was also not possible to physically roll a die at the front of the laboratory, and have all subjects observe the result. As such, a second best solution was developed of providing the subjects with a pre-recorded video, of the ten-sided die being rolled, and allowing them to view this video in privacy on their own individual computer terminals after they had made their decision³⁶. Whilst it may not be perfectly possible to convince the subjects that their viewing of the video was not conditional on the choice of gambles that they had made³⁷, and hence that they were not being deceived in any way, the privacy benefits of this method outweighed concerns about the believability of the outcome of the uncertainty. Additionally, because of the simplicity of the webpages used to run the experiment, it was not possible to select a random video for every subject in the experiment. However, it was possible to select a random video for every *session* of the experiment, as the code used to display the video was simple to change between sessions, and, as a result, every subject in a given session viewed the same roll of the die and, hence, the same resolution of the uncertainty. To ensure this was truly random, and hence ensuring that the probabilities stated in the experiment were correct, a large number of die rolls were pre-recorded, and a third party randomly allocated videos to sessions, prior to the choices of any subjects.

3.4.6 *The second stage*

As mentioned in the above section, the two primary groups of subjects we are interested in are those who experience Type A regret from stage one (group 3) and those who experience Type B regret from stage one (group 4). In order to investigate the effects of both types of regret on subsequent behaviour, and additionally study the difference between regret linked to action and regret linked to decision, we need to design a second stage of the experiment where the subject is asked to make another gambling decision, but fashioned in such a way that we can compare the behaviour of subjects who face subtly different choices.

Ideally, we would like a second stage which allows us to vary the action, whilst holding the decision constant, and also vary the decision, whilst holding the action constant. However, without having an explicit measure of regret, or a method for isolating the anticipated regret component of an subject's decision, it would be foolish to draw conclusions from an experiment which varied the second stage decision (such as the probabilities of winning a particular amount of money) and attribute the observed changes in behaviour simply due to the effect of the prior experience of regret. However, in contrast, a second stage which holds the decision constant, but varies the actions need to implement the

³⁶ an example of such a video, and how it was presented to subjects, is available at <http://go.warwick.ac.uk/slovelady/v1>

³⁷ though the simplicity of the webpages used, which all the students are very familiar with, should suggest that it was not

subject's decision, by definition has no obvious economic reason for producing a difference in behaviour. Hence, we would be much more confident drawing conclusions related to non-economic reasons, such as regret, from a second stage where the decision is held constant, but the actions are varied, and so it is this approach which shall be used.

Additionally, as discussed in the intuitive example, a second stage where the previously regretted choice is separated into the constituent components of the action and the decision, and these components are assigned to *different* options of the stage two gamble, is the most interesting case under which to analyse the effects of action and decision regret separately. Assuming it is more likely that the effect of decision regret ("I regret choosing the safe option") will impact more on the regret associated with the most "similar" decision in stage two, and the effect of action regret ("I regret selecting Gamble R") will do likewise for the most "similar" action in stage two, it forces the subjects to consider the possible emotional consequences associated with both options in stage two, rather than simply weighing the emotional consequence of choosing the exact same option again. Furthermore, if it was possible to analyse the individual preference for the two stage two options separately (rather than simply getting a measure of the relative preference, given by which option was ultimately selected), we would be able to disentangle the relative magnitude of the impact of decision and action regret. As such, the second stage of the experiment will deliberately split up the decision and action of the previously regretted option from stage one.

Considering the subjects who will end up in group 3 after stage one, they experienced Type A regret from making a relatively risky decision, by choosing Gamble R as their action. As such, in stage two, where Gamble R occurs, it should now be considered the *relatively safe* option. Considering the subjects who will end up in group 4 after stage one, they experienced Type B regret from making a relatively safe decision, by choosing Gamble S as their action. As such, in stage two, where Gamble S occurs, it should now be considered the *relatively risky* option. However, as Gamble R was originally "quite risky" to begin with, and Gamble S was originally "quite safe", this design of stage two necessitates in the introduction of two more gambles.

- Gamble VR (very risky) - a 10% chance of winning £50 and a 90% chance of winning £0. You will win £50 if the (ten-sided) die lands on the number 0
- Gamble VS³⁸ (very safe) - a 100% chance of winning £4. You will win £4 if the (ten-sided) die lands on any number between 0 and 9 (inclusive)

Using these two new gambles, those in group 3 will be presented, in stage two, with a choice between Gamble R (now relatively safe) and Gamble VR³⁹, and those in group 4 will be presented, in stage two, with a choice between Gamble S (now relatively risky) and Gamble VS⁴⁰.

³⁸ the names of all the gambles have been changed here from what was actually presented to the subjects in the experiment. In the experiment, arbitrary letters of the alphabet were chosen for the names, so not as to prime the subjects with any specific information as to how they should be interpreted or how the subjects should act. They have been presented as such here to aid understanding of how "risky" each gamble is.

³⁹ the webpage which subjects saw to resolve this choice is viewable at <http://go.warwick.ac.uk/slovelady/v2at1>

⁴⁰ the webpage which subjects saw to resolve this choice is viewable at <http://go.warwick.ac.uk/slovelady/v2bt1>

However, both of these situations involve a previously regretted action, as well as a previously regretted decision. As such, in order to isolate the effects of just action regret on subsequent choice, we need to also have a *control* stage two of the experiment, where those in group 3 face exactly the same decision, but do not see the action (or, equivalently, label) of Gamble R repeated, and those in group 4 face exactly the same decision, but do not see the action (or, equivalently, label) of Gamble S repeated. As such, we need four new gambles, which are decisionally and economically equivalent to R, S, VR and VS, but are not “presented” in exactly the same way, in order to disassociate the action component of regret. This can be done by defining four new gambles as follows

- Gamble Rc - a 30% chance of winning £14 and a 70% chance of winning £0
- Gamble Sc - a 70% chance of winning £6 and a 30% chance of winning £0
- Gamble VRc - a 10% chance of winning £50 and a 90% chance of winning £0
- Gamble VSc - a 100% chance of winning £4

These gambles are mathematically equivalent to those presented earlier, simply calling them something different⁴¹, but this might not be sufficient to break the action association, for example, between Gamble R and Gamble Rc. Hence, let us introduce a subtly different way of resolving the risk according to the given probabilities. Instead of rolling one ten-sided die, to generate a number between 0 and 9, let us roll two ten-sided dice to generate a number between 00 and 99. As such, the above gambles are resolved according to

- Gamble Rc - You will win £14 if the two dice generate any number between 00 and 29
- Gamble Sc - You will win £6 if the two dice generate any number between 30 and 99
- Gamble VRc - You will win £50 if the two dice generate any number between 90 and 99
- Gamble VSc - You will win £4 if the two dice generate any number between 00 and 99

Thus, from a decision perspective, a choice between Gamble R and Gamble VR is identical to a choice between Gamble Rc and Gamble VRc⁴². Similarly, a choice between Gamble S and Gamble VS is identical to a choice between Gamble Sc and Gamble VSc⁴³. From an action perspective, however, in both cases, the former choice involves a repeated action, of choosing the same letter resolved by the same method as in stage one, whereas the latter choice does not have the repetition.

Thus, we have created a *control*, where the action is not repeated, and *treatment*, where the action is repeated, version of stage two for each of

⁴¹ again, to aid understanding, the names of the gambles are presented differently here to what subjects saw in the experiment, but the principle of the argument remains

⁴² the webpage which subjects saw to resolve this choice is viewable at <http://go.warwick.ac.uk/slovelady/v2ac1>

⁴³ the webpage which subjects saw to resolve this choice is viewable at <http://go.warwick.ac.uk/slovelady/v2bc1>

the subjects in groups 3 and 4, and we are interested in the difference in observed behaviour between control and treatment, for both those in group 3 and those in group 4. From an experimental design perspective, therefore, there should be a random allocation of subjects in group 3 to either the treatment or control version of stage two, and likewise for those in group 4. For simplicity, we shall call the versions of stage two for subjects in group 3 either “stage two-A treatment” or “stage two-A control”⁴⁴ and the versions of stage two for subjects in group 4 either “stage two-B treatment” or “stage two-B control”.

For those who did not experience regret

As previously stated, the simplicity of the web-pages used to design and run the experiment did not permit us to filter subjects to different stage two web-pages depending on stage one choice, only on the outcome of stage one uncertainty, which, due to the use of random videos for each session, was known to the experimenter, but not the subjects, prior to the beginning of each experimental session. Thus, for example, in a session where the winning number from stage one was 6, this meant that those who chose Gamble R would experience Type A regret, but those who chose Gamble S would experience no regret, and win £6. All of these people, however, must then be directed to either “stage two-A treatment” or “stage two-A control”. As such, it was necessary to specify, in advance of the session, whether the session was a control session, in which case everyone was filtered to “stage two-A control” or a treatment session, in which case everyone was filtered to “stage two-A treatment”.

As a result those who chose Gamble S, and experienced no regret, saw the same second stage (either treatment or control) as those who chose Gamble R and experienced Type A regret. Thus we have stage two data from subjects who won in stage one (either by choosing Gamble R or S), and received some payoff, as well as from those who experienced regret.

3.4.7 Flow chart

Putting all of the above information together, therefore, leads us to be able to draw the flow chart in Figure 51, which specifies the various paths down which any given experimental subject may fall, given their choices, the outcome of uncertainty, and the decisions which they will be asked to take along the way.

3.4.8 Controlling for risk seeking and other individual specific factors

The aim of randomly allocating experimental subjects to either the treatment or control group is that there should be no correlation between individual, unobservable characteristics and the probability that a given subject is allocated to either the treatment or control group. However, as the method for generating the treatment and control groups is not purely random⁴⁵, but, as discussed above, is conditional on the session

⁴⁴ indicating that those subject will have experienced Type A regret at stage one

⁴⁵ i.e. the probability, both before the subject allocates themselves to an experimental session, and when they enter the lab for their session, of being in the treatment group, is not exactly equal to 0.5.

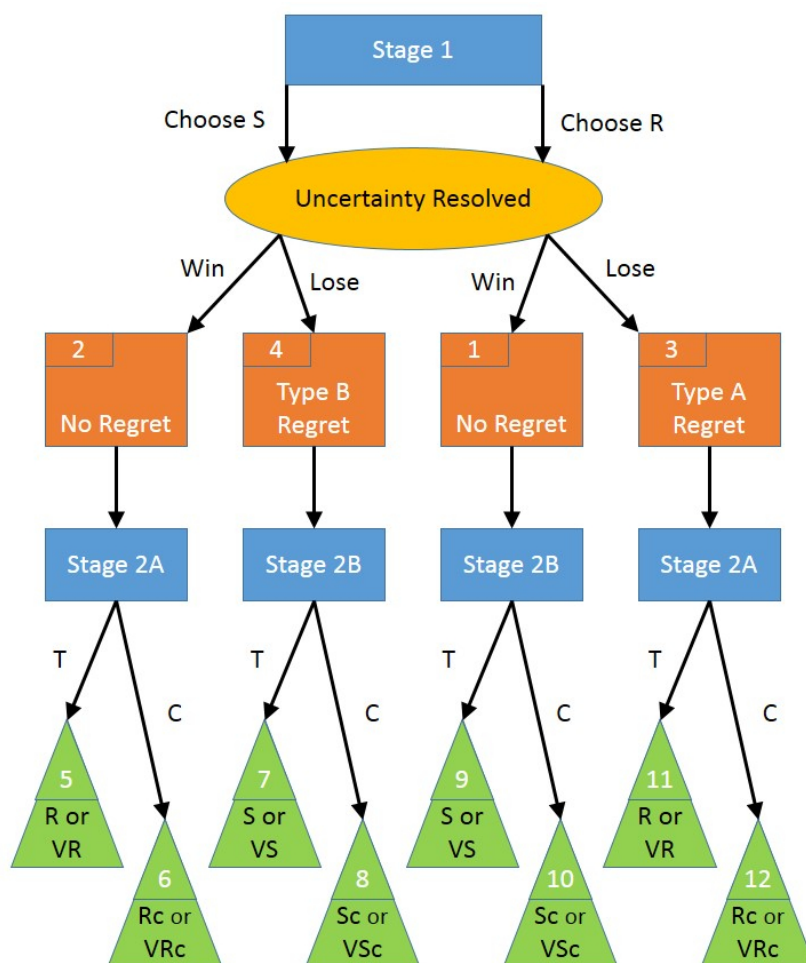


Figure 51: Experiment Flow Chart

which they attend, and the allocation of the subjects to the sessions is by their own selection (though without their knowledge which sessions are control and which are treatment), we would like a method which allows us to control for more obvious characteristics which *will* have a causal effect on their choices, and *might* end up as correlated with their allocation to either a treatment or control session⁴⁶.

Most of these characteristics, such as gender, age, personality and mathematical ability, can be assumed to be constant throughout the experiment, and so can be asked or measured in a questionnaire which is administered after the main body of the experiment is finished. Some of these characteristics, however, can be reasonably assumed to be affected by the decisions made, and outcome of uncertainty in the main body of the experiment, and so will need to be measured *before* the experiment is conducted. Specifically, we would like a measure of inherent risk seeking for each of the subjects, which is correctly incentivised to generate truthful revelation from the subjects.

Holt and Laury procedure

A standard method of generating a measure of risk seeking / risk aversion is that proposed by Holt and Laury [38]. The method proposes that subjects face a series of 10 gambling scenarios, where, in each scenario, there are two possible gambles as options. The gambles are similar to those described for this experiment, in that there are only two possible outcomes for each gamble. In each of the scenarios the amounts of money that each gamble offers stays constant, with one of the gambles offering “extreme” payoffs (both high and low) and the other offering “intermediate” payoffs. As the scenarios progress, the probabilities associated with each of the states of the world is varied, in such a way that the expected value (EV) of the gamble with extreme payoffs increases, thus making it more attractive compared to the gamble with intermediate payoffs. As such, given the increasing EV of the “extreme” gamble, compared to the “intermediate” gamble, a risk seeking individual is more likely to choose the extreme gamble in more of the scenarios than would a risk averse individual. Additionally, the gambles are created in such a way that an individual choosing according to expected utility would have a unique switching point, in that they would prefer the safe gambles in all scenarios before the switching point, and would prefer the risky gamble in all scenarios after the switching point.

The version of the procedure used in this experiment is modified from the original to take advantage of the standard experimental show-up fee offered to subjects. Subjects are offered £5 for showing up to the experiment, but have the opportunity to exchange this show-up fee for a lottery ticket, which pays £10 with some probability p and £1

⁴⁶ as an example of how this could happen, there were 14 sessions of the experiment run over the two days in the original run of experimental sessions. On the first day (Tuesday) there were 4 control sessions and 3 treatment sessions, but on the second day (Wednesday) there were 3 control sessions and 4 treatment sessions. At the University of Warwick, however, there are sports run on a Wednesday afternoon at the same time as the experimental sessions. Under the hypothesis that men are more likely to play sport than women, and hence are more likely to sign up to Tuesday sessions than Wednesday, this could result in relatively more men being allocated to Tuesday sessions than Wednesday, which would result in more men in control sessions than women. If men are assumed to be more risk seeking than women, this could induce a correlation between risk taking at stage two and the allocation of control and treatment groups, caused by a factor which has nothing to do with regret.

| SCENARIO | Show-up fee | | HIGH | Gamble | | |
|----------|-------------|------------|------|--------|-----|-------|
| | AMOUNT | PROBABILTY | | p | LOW | 1 – p |
| 1 | £5 | 1 | £10 | 0.1 | £1 | 0.9 |
| 2 | £5 | 1 | £10 | 0.2 | £1 | 0.8 |
| 3 | £5 | 1 | £10 | 0.3 | £1 | 0.7 |
| 4 | £5 | 1 | £10 | 0.4 | £1 | 0.6 |
| 5 | £5 | 1 | £10 | 0.5 | £1 | 0.5 |
| 6 | £5 | 1 | £10 | 0.6 | £1 | 0.4 |
| 7 | £5 | 1 | £10 | 0.7 | £1 | 0.3 |
| 8 | £5 | 1 | £10 | 0.8 | £1 | 0.2 |
| 9 | £5 | 1 | £10 | 0.9 | £1 | 0.1 |
| 10 | £5 | 1 | £10 | 1 | £1 | 0 |

Table 6: Holt and Laury scenarios

with some probability $1 - p$. To synchronise this idea with the Holt and Laury procedure, 10 scenarios are presented, with a different value of p in each scenario, with the subjects being informed that one scenario will be randomly selected, with their choice between the show-up fee and the lottery ticket in that scenario determining their payoff. The scenarios presented to the subjects were as shown in Table 6.

In each of the ten scenarios, the subjects had to pick whether they wished to keep their £5 show-up fee, or exchange it for the lottery ticket with payoff and probabilities as specified. However, in order not to generate the anticipation or experience of regret as a result of the Holt and Laury procedure, rather than the main body of the experiment itself, all subjects were informed that

Once you have made all 10 choices, at the end of the session, one of these scenarios will be randomly picked, and your decision in that scenario will determine your payment for this section of the session.

Please note, for this section only, you will not learn which scenario was picked, or be told whether you won or lost based on any lottery tickets you may have selected. You will simply have the correct amount included in your total payment for the session.

Informing subjects that they would not see the resolution of uncertainty, as it pertained to both the scenario selected and the outcome of any lottery tickets they may have selected, prevents the rest of the experiment from being impacted by the use of the Holt and Laury procedure. Additionally, by running the Holt and Laury procedure before the main experiment, it generates a measure of risk seeking / risk aversion unaffected by the experience of winning, losing, regret or disappointment, and so should give a true reflection of the underlying risk characteristics of the subjects.

3.5 HYPOTHESES

Given the experimental design, there is a significant amount of information to be gained from analysing the results. The regret literature suggests some natural hypothesis tests in keeping with the style of existing research in the area.

3.5.1 Hypothesis 1A

The prior experience of Type A regret makes subjects more sensitive to Type A regret, increasing the future anticipated regret from relatively risky decisions, and hence more likely to choose Rc over VRc at point 12 compared to Rc over VRc at point 6.

3.5.2 Hypothesis 1B

The prior experience of Type B regret makes subjects more sensitive to Type B regret, increasing the future anticipated regret from relatively safe decisions, and hence more likely to choose Sc over VSc at point 8 compared to Sc over VSc at point 10.

3.5.3 On 1A and 1B

Such hypotheses are typical in the regret literature, and appear convincing because they appear to give a specific prediction about the behaviour of individuals in the experiment. In this case, they say we should compare “people who did regret a decision” to “people who didn’t regret a decision” when faced with a similar decision to make in the future. In this experiment, however, “people who did regret” are equivalent to “losers” of the stage one gamble and “people who didn’t regret” are equivalent to “winners” of the stage one gamble. This creates a confounding “wealth effect”, as those who did not regret in Hypothesis 1A have just won £6 from choosing Gamble S (correctly)⁴⁷, and those who did not regret in Hypothesis 1B have just won £14 from choosing Gamble R (correctly)⁴⁸. Thus, it would be misleading to attribute any observed change in behaviour to an experienced regret effect when a wealth effect would be arguably much more significant⁴⁹.

In order to control for such wealth effects we would need to create two additional groups of subjects who did not “win” in stage one, but

⁴⁷ For completeness, in the experiment, the observed proportion of subjects choosing Rc at point 12 was 39.2% and the observed proportion of subjects choosing Rc at point 6 was 58.9%. Under a chi-squared test for equality of proportions ($\chi^2(1) = 5.1768$ $p = 0.023$) we would reject the null hypothesis, but draw the opposite conclusion to Hypothesis 1A, in that more subjects chose the relatively safe option at point 6 compared to point 12. This may be indicative of “target” wealth effects rather than traditional wealth effects.

⁴⁸ For completeness, in the experiment, the observed proportion of subjects choosing Sc at point 8 was 11.5% and the observed proportion of subjects choosing Sc at point 10 was 9.1%. A chi-squared test for equality of proportions is not appropriate in this instance as several cells have below 5 subjects.

⁴⁹ The sums of money available to win in the experiment are neither particularly small nor particularly significant, and are comparable, from an hourly rate perspective, to other experiments run through the DR@W Laboratory. Thus, whilst it’s certainly possible to argue that traditional wealth effects would play a role in the decision making process, it may be more sensible to consider a “target” wealth effect as the source of potential wealth confounds. That is to say, as the experiment appears to offer comparable payouts to other similar experiments that the subjects may have participated in, they may have a good sense of what is considered “success” and “failure”, and take decisions to ensure they “succeed” at the expense of large upside.

instead faced a choice between Gamble Rc and Gamble VRc (in one group) and Gamble Sc and Gamble VSc (in the other) as a simple one stage gamble, and then compare the percentage of people who chose Rc (Sc) over VRc (VSc) in the one shot gamble to the percentage of people who did equivalently having experienced regret from a similar decision in a previous stage. Formally, this suggests two new hypotheses

3.5.4 Hypothesis 1C

The prior experience of Type A regret makes subjects more sensitive to Type A regret, increasing the future anticipated regret from relatively risky decisions, and hence more likely to choose Rc over VRc at point 12 compared to Rc over VRc when faced with the choice between Rc and VRc as a one shot decision.

3.5.5 Hypothesis 1D

The prior experience of Type B regret makes subjects more sensitive to Type B regret, increasing the future anticipated regret from relatively safe decisions, and hence more likely to choose Sc over VSc at point 8 compared to Sc over VSc when faced with the choice between Sc and VSc as a one shot decision.

3.5.6 On 1C and 1D

As mentioned in the previous chapter, however, note that the above hypotheses are not the only hypotheses, including the effect of experienced regret on subsequent regret aversion, which will predict the same observed pattern of behaviour. Consider the following two hypotheses

3.5.7 Hypothesis 1E

The prior experience of Type A regret makes subjects less sensitive to Type B regret, reducing the future anticipated regret from relatively safe decisions, and hence more likely to choose Rc over VRc at point 12 compared to Rc over VRc when faced with the choice between Rc and VRc as a one shot decision.⁵⁰

3.5.8 Hypothesis 1F

The prior experience of Type B regret makes subjects less sensitive to Type A regret, reducing the future anticipated regret from relatively risky decisions, and hence more likely to choose Sc over VSc at point 8 compared to Sc over VSc when faced with the choice between Sc and VSc as a one shot decision.

This type of effect was considered during the design of the experiment, and specifically the hypotheses that are to be tested are made across control and treatment groups with similar wealth levels.

⁵⁰ equivalently, this hypothesis could be considered as “the experience of Type A regret makes subjects less regret averse overall, and as subject are more concerned with Type B regret than Type A in stage two, they are more likely to choose Rc over VRc at point 12 compared to Rc over VRc when faced with the choice between Rc and VRc as a one-shot decision.” This style of the hypothesis is more in keeping with the research conducted in Chapter 2, and is behaviourally equivalent.

3.5.9 Hypothesis Equivalence

The equivalence of behavioural predictions made by 1C and 1E, and similarly 1D and 1F, implies that this experimental set-up is not suitable for investigating whether the experience of regret leads to a subsequent *decrease* or *increase* in regret aversion, of any type. As such, we must go beyond the typical and traditional hypotheses tests when analysing the results of this experiment. Instead of working from a position of attempting to vary the experience of regret, and hold the second stage decision constant, we should go down a path of varying the second stage of the experiment, in a fashion which traditional theories say will not be affected by the experience or anticipation of regret, and see the impact of this change on behaviour for different levels of prior experienced regret.

3.5.10 Atypical Hypotheses

As previously discussed, conventional theories of how experienced regret influences subsequent decisions focus on changing some quantifiable aspect of the decision, without explaining the mechanism through which this experience of regret operates. Specifically, they fail to specify how the recurrence of actions, as opposed to decisions, will influence this mechanism. This experiment allows us to test whether the action is indeed relevant in the transmission mechanism of experienced regret to subsequent decision making. By specifying a control group, which will be faced with a similar decision in stage two as they did in stage one, but will not see the exact same actions as were presented in stage one, and comparing the decisions taken to a treatment group, who face exactly the same stage one and stage two decisions as the control group, but see a recurrence of one action from stage one, we can test whether the experience of regret is transmitted to subsequent choice behaviour through a link to actions. Specifically,

3.5.11 Hypothesis 2A

The experience of regret from a particular action, and emotional link to a subsequent recurrence of that action, will influence the decision made by subjects at point 7, when a previously regretted action reoccurs, compared to the equivalent decision taken at point 8, when the previously regretted action does not reoccur. Thus, the proportion of subjects choosing the relatively safe option, VS, at point 7 will be different to the proportion of subjects choosing the relatively safe option, VSc, at point 8.

3.5.12 Hypothesis 2B

The experience of regret from a particular action, and emotional link to a subsequent recurrence of that action, will influence the decision made by subjects at point 11, when a previously regretted action reoccurs, compared to the equivalent decision taken at point 12, when the previously regretted action does not reoccur. Thus, the proportion of subjects choosing the relatively safe option, R, at point 11 will be different to the proportion of subjects choosing the relatively safe option, Rc, at point 12.

3.5.13 Control Hypotheses

Hypotheses 2A and 2B predict the existence of a reason (experienced regret linked to an action) for a change in the proportion of subjects who will choose a (relatively) safe gamble over a (relatively) risky gamble between the control and treatment groups. However, there are plenty of other reasons that one could imagine will have a non-negligible effect on different choice behaviour between the control and treatment groups. For example, in the control group at point 12, the relatively safe option will win if the two dice generate the number 00. However, in the treatment group at point 11, the relatively risky option will win if the one die generates the number 0. Hence, should subjects, on average, have an aversion to making a decision which is reliant on observing the number zero to win, either in the form of 00 or 0, then this would imply subjects being more likely to choose the (relatively) risky option at point 12, yet subjects being more likely to choose the (relatively) safe option at point 11, despite these decisions being mathematically equivalent. Such a preference would generate a positive result when testing Hypothesis 2B, with no need for an explanation involving the experience of regret from a particular action at a prior stage. This effect is defined as the “inherent action aversion”.

Two other common explanations for biases in repeated gambling behaviour are the “Gambler’s Fallacy” and “Hot Hand” effects. The Gambler’s Fallacy, “...expecting outcomes in random sequences to exhibit systematic reversals” (Rabin and Vayanos [72, p730]), could be applied in this experiment if experimental participants “...believe mechanical randomizers ... exhibit sequential tendencies” (Keren and Lewis [55, p75]), and believe that the die used to resolve the gambles in stage two is the same die as used to resolve the gambles in stage one⁵¹. In which case, having observed the die land on a number between 0 and 6 in stage one, those participants at point 11 may believe that a number between 7 and 9 has a greater than 30% probability of occurring at stage two, compared to those subjects at point 12, who have the gamble resolved by two die instead of one, and hence are unaffected by the fallacy. This theory would predict that subjects are more likely to choose Gamble R at point 11 compared to Gamble Rc at point 12. As with the above aversion to the number zero, this belief would cause a positive result to a test of Hypothesis 2B, without any need for an experienced regret explanation.

The Hot Hand effect, defined as “...as a belief in the continuation of streaks” [72, p733], applied in this context as a belief that a “winner” will keep on winning, and a “loser” will keep on losing, would predict that “losers”, who experienced regret in stage one, would behave more conservatively⁵² than those who won. This theory, however, does not predict a difference in behaviour between control and treatment groups, conditional on “losing” at stage one, as is evident in Hypotheses 2A and 2B.

As such, we have two possible confounding effects, of the Gambler’s Fallacy and the “inherent action aversion”, which would give the same hypotheses tests as Hypothesis 2A and 2B. However, both of the alternative explanations also make predictions about behaviour of subjects

⁵¹ this is true, but not specified to the participants in advance of making their decision.

⁵² under the assumption that the larger payoff constitutes “winning” more than the smaller payoff

when they haven't experienced regret from stage one. Specifically, the Gambler's Fallacy predicts that, for example, if those subjects at point 11 have a stronger preference for R than subjects at point 12 have for Rc, then subjects at point 5 would have an equally strong preference for R when compared to the preference of subjects at point 6 have for Rc. That is to say, the effect of the Gambler's Fallacy is not conditional on the choice of the subject on stage one, merely conditional on the observed roll of the die, which is the same (a number between 0 and 6) for all subjects at points 5, 6, 11 and 12.

In a similar vein, there is no reason why any inherent action aversion would be conditional upon the choice, and hence the emotional experience, of the subject at stage one, and hence such an explanation would predict the same difference in behaviour between control and treatment groups for both those who experienced regret, and those who did not experience regret.

In contrast, the hypothesis that the experience of regret from a particular action creates an emotional link to a subsequent recurrence of that action, implies that there would only be a difference in behaviour between the control and treatment group under the experience of regret, and not when there is no experience of regret from the first stage. Hence, we can differentiate between "regret-based" hypotheses and "non-regret-based" hypotheses, by the following control hypotheses.

3.5.14 Hypothesis 3A

In the absence of the experience of regret from stage one, there is no inherent preference, or bias resulting from the observation of the resolution of uncertainty at stage one, which causes subjects to choose differently between the control and treatment groups at stage two. Hence, the proportion of subjects choosing the relatively safe option, VS, at point 9, will be no different to the proportion of subjects choosing the relatively safe option, VSc, at point 10.

3.5.15 Hypothesis 3B

In the absence of the experience of regret from stage one, there is no inherent preference, or bias resulting from the observation of the resolution of uncertainty at stage one, which causes subjects to choose differently between the control and treatment groups at stage two. Hence, the proportion of subjects choosing the relatively safe option, R, at point 5, will be no different to the proportion of subjects choosing the relatively safe option, Rc, at point 6.

3.5.16 Expectation for the results of the hypotheses

Based on standard regret theories, the most predictable result would be a failure of Hypotheses 2A and 2B, and a success of Hypotheses 3A and 3B, thereby implying that simply framing the decision problem in terms of a repeated, previously regretted action, has no effect on the choice made in stage two. However, should Hypotheses 2A and 2B succeed, and Hypotheses 3A and 3B fail, there would be reason to believe that the experience and association of regret to a particular action is not the primary cause of a difference in observed behaviour between the control and treatment groups. Should all four hypotheses fail, we would need to revisit the proposed emotional attachment theories altogether.

However, should all four hypotheses succeed, it would imply that there is an emotional link between the experience of regret and the subsequent recurrence of the action which caused the regret, causing a change in behaviour when that action is part of a subsequent decision, thus implying it is not sufficient to model the effect of the experience of regret on subsequent choice without reference to the context and action, in addition to the decision, which was taken in the first stage.

3.6 EXPERIMENTAL CONFOUNDS

The design of this experiment necessarily introduces a potential confound when discussing hypotheses and interpreting the results. As we are looking at a two stage experimental design, as per the decision tree in 51, whilst we may get a random sample of the student population at stage one of the experiment, we will certainly not get a random sample of the population in each of boxes 1, 2, 3 and 4 for the start of stage two. The left side of the decision tree at stage two is entered by those participants who chose S at stage one, which is the relatively safe option. The right hand side is entered by those participants who chose R at stage one, which is the relatively risky option. Therefore, in simple terms, we would expect the left hand side of the tree to be comprised of more risk averse participants, and the right hand side to be comprised of more risk seeking participants⁵³.

The effect of this confound is two-fold. First, any hypotheses which rely on comparisons between groups across the “left – right” tree divide (i.e. anything from triangle 5 through 8 against anything from triangle 9 through 12) are subject to the argument that the difference in risk aversion in the samples is a possible cause for the experimental result.

Second, any hypotheses which rely on comparisons between groups exclusively within their branch of the tree (either left or right) may generate a result which is only applicable to either risk averse (left) or risk seeking (right) people.

The first of these effects is more significant and important than the second.

Thinking about the second, applicability is always a concern in experimental economics; usually discussed from the point of external validity or “parallelism” of an entire experiment. To the extent that there is a theoretical reason to believe that the results are only applicable to one “type” of person, the experiment can easily be repeated on a new sample of people who may be more risk averse / seeking than this sample, and this hypothesis can be tested.

The first effect must be more carefully navigated, as there is always reason to believe risk aversion may drive decisions made under risk. However, in contrast to the typical hypotheses tested in regret experiments (as discussed with respect to hypotheses 1A through 1F), the hypotheses being tested in this experiment (2A through 3B) are only made within a branch of the tree (either left or right). This should alleviate the concerns of the possible confound with respect to any of these individual hypotheses.

⁵³ This expectation is confirmed by analysing the results of the Holt and Laury procedure separately for those who chose R at stage one to those who chose S. Amongst those who chose R the mean number of safe choices made was 5.5, whereas amongst those who chose S, the average was 6.3. This is a significant difference at the 1% level, according to a two sample t-test.

3.7 THE EXPERIMENT

The experiment was run at The University of Warwick, over the 29th and 30th November 2011 and the 2nd December 2014⁵⁴, in the DR@W (Decision Research at Warwick) Laboratory in the Department of Economics. Over the three days, there were 21 hour-long sessions run. Participants were recruited from the DR@W Participants Database, consisting of around 1400 students and staff from the University of Warwick, with the majority of participants being in their first and second year of undergraduate study. All participants who had registered for invitations to laboratory experiments were invited to sign up, with first priority being offered to those who had never missed an experimental session which they had signed up for. In total, there were 420 spaces available to which interested participants could assign themselves⁵⁵, and 376 participants indicated a willingness to participate in the experiment, leaving spaces in some of the sessions, which happened for the later run sessions on the 29th and 30th November 2011. The invitation email which was sent to all participants in the database is included in Appendix B. Of the 376 participants who signed up to participate, 344 showed up over the three days. Summary statistics of the 344 participants are also given in Appendix B.

Of the 21 sessions, 11 were designated as “control sessions” and 10 were designated as “treatment sessions”. As the experiment was run in a simple web-browser, the path through the experiment (i.e. the sequence of webpages which every participant would see) was only possible to specify on a per-session basis, and not on a per-participant basis. Thus, in order to generate true randomisation between treatment and control groups, we were reliant on participants randomly allocating themselves across the treatment and control sessions.

A very detailed experimental procedure is given in Appendix B, but is summarised as follows. On entering the laboratory, each participant randomly chose a slip of paper which specified a unique participant ID number, which would be used to identify their responses across the experiment. This ensured anonymity of subject responses, as the ID number was not known to the experimenter. Once in the laboratory, and seated at a random computer terminal, all participants were read a set of instructions, as presented in Appendix B.3.6. They then progressed throughout the experiment without further instruction from the experimenter. The experiment took approximately 45 minutes to complete.

Once all subjects in a given session had completed the experiment, they were read, and paid according to, a set of payment instructions, as given in Appendix B.3.7. Their payment was calculated as the sum of the result of their choice and the resolution of the gambles from stages one and two, and the result of the Holt and Laury[38] procedure⁵⁶. The

⁵⁴ additional sessions were requested to be run upon submission of the original thesis, hence the extended delay between the first and second set

⁵⁵ a maximum of 20 in each of the 21 sessions

⁵⁶ to recap, of the 10 scenarios presented using the modified Holt and Laury procedure, one was randomly selected by a random number generator. If the participant chose to keep their £5 show-up fee for this scenario, this was added to their payment. If they chose to gamble in this scenario, the outcome of the gamble was resolved according to the given probabilities, using a second random number generator. If the participant won the gamble, £10 was added to their payment. If they lost the gamble, £1 was added to their payment

| | (1) c1 Safe Choice? | (2) c1 Safe Choice? |
|------------------------|------------------------|------------------------|
| Age | 0.000549 (0.946) | 0.000803 (0.921) |
| Male (d) | -0.0671 (0.204) | -0.0591 (0.267) |
| A-level maths (d) | -0.0398 (0.569) | -0.0303 (0.670) |
| Time of session (24hr) | -0.00404 (0.766) | -0.00265 (0.846) |
| 2014 session (d) | 0.0475 (0.377) | 0.0513 (0.343) |
| EV Correct (d) | -0.0139 (0.799) | -0.0210 (0.703) |
| HL Safe Choices | | 0.0523*** (0.000) |
| Observations | 344 | 344 |
| Pseudo R ² | 0.009 | 0.039 |

Marginal effects; *p*-values in parentheses

(d) for discrete change of dummy variable from 0 to 1

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7: Probit model showing determinants of stage one choice. Shown are the marginal effects of each independent variable, evaluated at the mean.

mean participant payment was £12.45 for a 45 minute session, which equates to an average hourly pay rate of £16.60.

3.8 RESULTS

3.8.1 Stage one

For every subject, stage one offered the choice between Gamble R and Gamble S, with Gamble R considered relatively risky and Gamble S considered relatively safe. As such, we can characterise the choice simply by stating whether they preferred the safe choice (Gamble S) or not. As this variable is a dummy variable, an appropriate regression model to use is a cross-sectional probit model, where the dependent variable takes the value of 1 if the subject chose the safe option in stage one, and 0 otherwise. A range of individual specific characteristics, measured by responses to the post experiment questionnaire, are tested for their predictive power on the stage one choice, and the results of this analysis, using the probit model, are given in column 1 in Table 7.

Two of the characteristics relate to the subjects inherent mathematical ability, as measured by their previous study (*A-Level Maths*) and

| | GAMBLE R | GAMBLE S |
|--|----------|----------|
| Number | 18 | 189 |
| As a % of subjects who suffered no stage one regret | 8.70 | 91.30 |
| As a % of total subjects | 5.23 | 54.94 |

Table 8: Choices of subjects who won in stage one

whether or not they could correctly calculate a simple expected value (*EV Correct*). *2014 session* is a dummy variable indicating that the subject participated in the repeat sessions in 2014 as opposed to the original sessions in 2011. No characteristics were found to be significant at the 5% level.

Column 2 in Table 7 takes the probit model fitted in column 1, and adds an additional explanatory variable, of the number of safe choices (choosing the fixed show-up fee over the lottery ticket) made by the subject, out of 10 scenarios, from the incentivised response to the Holt and Laury procedure⁵⁷. As would be expected, risk aversion, as displayed by the response to the Holt and Laury procedure is a highly significant predictor of risk aversion in the stage one gamble, with no other variables significant at the 5% level in this version of the model.

In total, 116 (33.72%) of the subjects chose Gamble R (relatively risky) in stage one, and 228 (66.28%) chose Gamble S (relatively safe), indicating a preference, on average, for the safe gamble. However, from the percentages, it is clear that there was no “obvious” choice, which would have reduced the likelihood of the result of stage one generating experienced regret.

Due to the randomisation of the selection of videos of the ten-sided die being rolled, the probability that, in any given session, Gamble S won was 0.7, and hence the probability that Gamble R won was 0.3. However, the small sample (only 21 sessions) meant that, in reality, 18 of the 21 sessions (85.71%) resulted in Gamble S winning, and only 3 of the 21 sessions (14.29%) had Gamble R winning. This translated, due to the small variance in the number of subjects per session, into 287 of 344 (83.43%) of subjects experiencing a roll of the die where Gamble S was the ex-post optimal choice, and 57 of 218 (16.57%) of subjects experiencing a roll of the die where Gamble R was the ex-post optimal choice.

The result of both the choices of the subjects, and the resolution of uncertainty, in stage one, as displayed Tables 8 and 9, shows that 137 (39.82%) of subjects experienced some form of regret from stage one, with 98 experiencing Type A⁵⁸ regret, and 39 experiencing Type B⁵⁹ regret. In the language of the Figure 51, this implies that of the 344

⁵⁷ a higher number for this variable indicates a higher degree of risk aversion, though a response of exactly 10 indicates a subject who did not understand the question, as the final scenario offers the subjects a 100% chance of winning £10, if they choose the “risky” lottery ticket. 5 subjects displayed this failure to understand, and hence are excluded from subsequent regression and statistical analysis, along with 3 other subjects who self-reported that they did not understand certain aspects of the Holt and Laury procedure

⁵⁸ choosing gamble R, and regretting not being more conservative

⁵⁹ choosing Gamble S and regretting not being more speculative

| | GAMBLE R | GAMBLE S |
|--|----------|----------|
| Number | 98 | 39 |
| As a % of subjects who suffered stage one regret | 71.53 | 28.47 |
| As a % of total subjects | 28.49 | 11.33 |

Table 9: Choices of subjects who lost in stage one

| | EXPERIENCED REGRET AT STAGE ONE | | | |
|----------------------|---------------------------------|-----------|--------------|-----------|
| | SAW STAGE 2A | | SAW STAGE 2B | |
| | Control | Treatment | Control | Treatment |
| # of subjects | 52 | 46 | 26 | 13 |
| Group # on Figure 51 | 12 | 11 | 8 | 7 |

Table 10: How many participants fell into each stage two group? For those who **did** experience regret at stage one

subjects who started the experiment, 18 end up in group 1, 189 end up in group 2, 98 end up in group 3 and 39 end up in group 4.

3.8.2 Result of the randomisation

As previously stated, each session was designated either a control or treatment session, and all the subjects in a given session saw the same sequence of gambles. Additionally, as Figure 51 indicates, those who fell into groups 1 and 4 after stage one saw stage 2B as the second stage, and those who fell into groups 2 and 3 after stage one saw stage 2A as the second stage. Given Tables 8 and 9, indicating the number of subjects who experienced each type of regret and also those who experienced no regret, we can now define the numbers of subjects who fell into each of the stage two groups given this process. This information is displayed in Tables 10 and 11.

As there were only 3 sessions where the safe choice from stage one was the losing choice, Tables 10 and 11 show this corresponds into a low number of subjects falling into groups 7, 8, 9 and 10. As some of our hypotheses rely on comparisons between these groups, there is likely to be a problem of achieving statistical significance due to the

| | NO REGRET AT STAGE ONE | | | |
|----------------------|------------------------|-----------|--------------|-----------|
| | SAW STAGE 2A | | SAW STAGE 2B | |
| | Control | Treatment | Control | Treatment |
| # of subjects | 96 | 93 | 11 | 7 |
| Group # on Figure 51 | 6 | 5 | 10 | 9 |

Table 11: How many participants fell into each stage two group? For those who **did not** experience regret at stage one

low numbers. An issue also arises as the randomisation was done on a session, not subject level, meaning 2 sessions were control and only 1 was treatment. We will return to this issue later, but will, for now, focus on sessions where the safe choice in stage one was the winning choice.

For those subjects who ended up in groups “5 and 6” & “11 and 12”, however, the randomisation procedure appears to have had the desired effect, with close to 50% of the subjects from the previous buckets falling into each group. It is possible, however, that the randomisation accidentally selected subjects with specific characteristics, which could be correlated with their decision making at stage two. To check this, we can run a probit regression, with the variable indicating whether the subject attended a treatment session or not, as the dependent variable, and all of our measurable subject-specific characteristics as independent variables. The results of this, for subjects in sessions where the safe choice in stage one was the winning choice, are given in Table 12.

Column 1 of Table 12 shows a highly significant positive correlation between having an A-Level maths, or equivalent, qualification, and the probability that you attended a treatment session. Notice that this is even controlling for whether or not you can calculate a simple expected value. This significance remains when we introduce additional regressors in the second reported regression, in the form of the results of the Big 5 personality test⁶⁰. Given there may be an accidental relationship between mathematical ability and whether or not you were in the control or treatment, it is worth keeping this relationship in mind when calculating the effect of the treatment on choice behaviour at stage two.

3.8.3 Stage two

Lack of subjects in groups 7, 8, 9 and 10 and experimental design

As outlined above, there is a noticeable shortage of subjects who fell into groups 7, 8, 9 and 10. A typical rule of thumb used in economics experiments is that 30 subjects per cell is required as a sensible starting point for analysis⁶¹, so falling below this threshold is obviously concerning.

As always with experiments, the limiting factor is resources. With only 7 subjects falling into group 9, the simple conclusion is that it would take an experiment approximately four times as long and four times as expensive to generate sufficient data to produce the necessary sample sizes in all cells. As the experimental sessions run to date cost a total of £4283 to pay participants, one which costs four times as much is beyond the reach of all but the most well-funded research groups.

However, this may be an overestimate as the experimental design choice of “session specific die rolls” rather than “subject specific die rolls” implies the variance around the expectation of how many subjects should fall into each group is higher. Thus, to estimate the true expected funding the experiment requires in order to have at least 30 subjects

⁶⁰ the Big 5 personality test, used in this experiment, is a 44 question personality survey, where subjects must answer to what degree they associate with specific types of behaviour (for example, “I am someone who is inventive”) on a 5 point Likert scale. It is commonly used in psychology and produces a measure of personality in five key areas: extraversion, agreeableness, conscientiousness, neuroticism and openness. (See John et al. [43] and John et al. [44])

⁶¹ though, as List et al. [57] point out, this rule “...has little basis in terms of power unless the researcher believes that he wants to detect an approximately 0.70 standard deviation change in the outcome variable” (p449)

| | (1) Treatment? | (2) Treatment? |
|-----------------------|--------------------|---------------------|
| HL Safe Choices | -0.0111 (0.520) | -0.00944 (0.591) |
| Age | 0.00932 (0.317) | 0.00917 (0.330) |
| Male (d) | -0.108 (0.082) | -0.116 (0.079) |
| A-level maths (d) | 0.230** (0.003) | 0.227** (0.004) |
| EV Correct (d) | 0.0609 (0.345) | 0.0659 (0.313) |
| Extraversion | | -0.0296 (0.483) |
| Agreeableness | | 0.0230 (0.689) |
| Conscientiousness | | -0.0529 (0.280) |
| Neuroticism | | -0.00148 (0.976) |
| Openness | | 0.0421 (0.462) |
| Observations | 279 | 279 |
| Pseudo R ² | 0.031 | 0.037 |

Marginal effects; *p*-values in parentheses

(d) for discrete change of dummy variable from 0 to 1

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 12: Probit model showing subject-specific characteristics as correlates with the allocation to a treatment or control session. Shown are the marginal effects of each independent variable, evaluated at the mean.

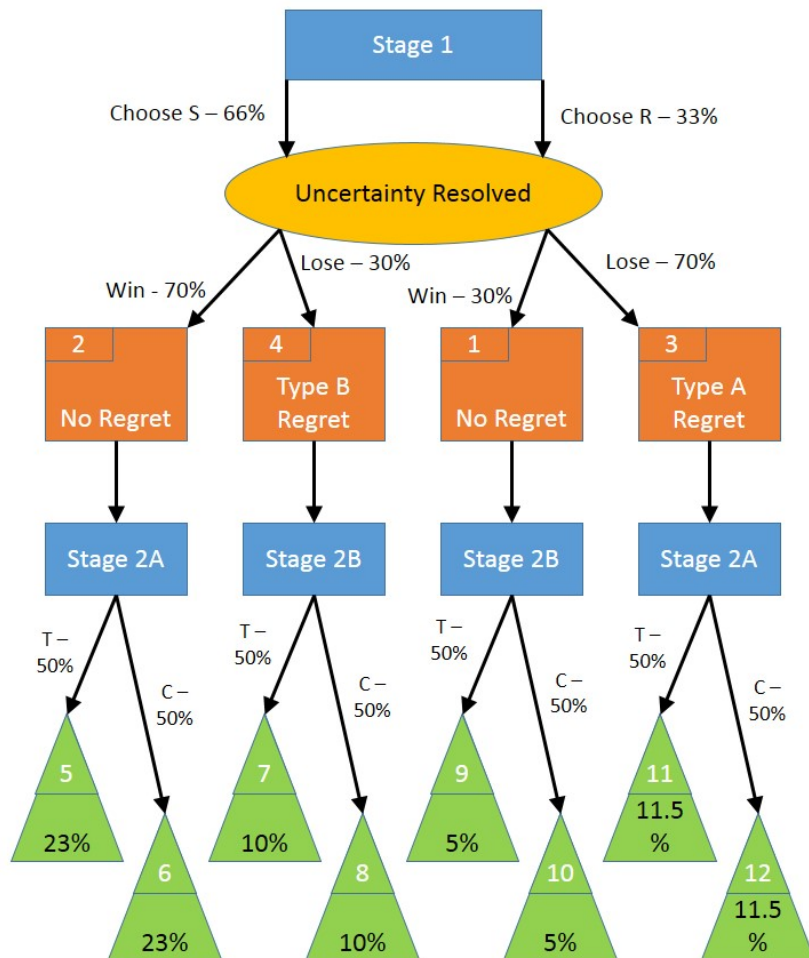


Figure 52: Experiment Group Probability Chart

fall into each group, we should consider the design where individuals receive subject specific rolls of the dice, and combine the observed decision frequencies with the true win / loss probabilities.

As can be seen from Figure 52, only 5% of subjects would be expected to fall into groups 9 and 10 based on the observed decision probabilities at stage 1 and the expected probabilities given by subject-specific die rolls. Therefore, to get a minimum of 30 subjects in each group, you would need a pool of 600 subjects to enter the lab. Given the observed average per participant payment of £12.45, this would imply a payment fund of approximately £7500⁶². As such a large payment fund was not available, it was simply not possible to generate a new, or extend the current, experimental design to meet these objectives.

One potential solution to this problem is the use of deception. By manipulating the true probabilities of the stage one gambles to be closer to 50/50 (rather than the 70/30 indicated to participants), you are able to move more subjects towards groups 9 and 10 at the expense of additional subjects in cells 5 and 6. Indeed, if the stage one gambles are simply resolved as a function of what was chosen in stage one, then

⁶² Additional funding would also be necessary to build a custom experimental website where subject specific outcomes (as opposed to session specific) can be tracked, as the University of Warwick Sitebuilder system does not currently permit this. Cost estimates of taking this approach ranged between £1000 and £3000.

| | CONTROL | | TREATMENT | |
|---------------|-----------|------------|-----------|-----------|
| | Gamble Rc | Gamble VRc | Gamble R | Gamble VR |
| % of subjects | 52.03 | 47.97 | 68.35 | 31.65 |

Table 13: Stage two choice by treatment for groups 5, 6, 11 and 12

you can perfectly ensure that exactly $1/8$ of the subjects fall into each of the final groups.

This approach is problematic, however, for two reasons. Firstly, this experiment was run by an economist in an economics department laboratory, and the status-quo is that “[economists] almost never use deception” (Roth [79]) for fear of *polluting the well* for other researchers who will rely on a subject’s belief that the information they have been given is true. This viewpoint is enforced by the rules of the DR@W lab which do not permit studies that use deception. The second reason it is problematic is one of practicalities. Even if deception were permitted, the cost of resources and effort required, in order to successfully rig the experiment so that it is both simultaneously perfectly deceptive and perfectly believable to the subjects who participate, may actually be higher than just doing the experiment without any deception at all.

Primary analysis on groups 5, 6, 11 and 12

As such, we will restrict most of our analysis to those cells where a good quantity of data was available given the present experimental design. Specifically, we will look at those subjects who faced stage 2A, as a result of Gamble S being the winning gamble from stage one, and hence were in groups “5 and 6” & “11 and 12” in stage two.

As described in the design of the experiment, the gambles in stage two were chosen so as to be close to indifference for a large number of the subjects, with the aim that the experience (or lack thereof) of regret from stage one would be sufficient to induce a significantly observable difference in behaviour at stage two. As such, we need neither Gamble R (or Gamble Rc in the control group) nor Gamble VR (or Gamble VRc in the control group) to be clearly preferable for the average subject, in order to generate some variance in behaviour at the second stage. Table 13 shows the proportion of subjects who chose each gamble in stage two, as a fraction of the total number who faced a choice which involved that gamble (i.e. conditional on whether they were in the treatment or control group).

Looking at the control group, the percentage of subjects choosing the relative safe (Gamble Rc) and relatively risky (Gamble VRc) options is approximately equal⁶³, suggesting the gambles can be considered very close to indifference for the average subject. This indicates we have developed a suitable second stage environment in which to study the impact of experienced regret from the first stage.

Noticeably, however, the proportion of subjects who chose the mathematically equivalent gambles (Gamble R and Gamble VR) in the treatment group is different⁶⁴ from the control group, suggesting that

⁶³ a chi-square test for equality of proportions does not reject the null of no difference

⁶⁴ using the sample proportions from the control group, the observed frequencies in the treatment group differ from expectation with a chi-squared value of 14.831, hence this result is significant the 1% level

| | CONTROL | | TREATMENT | |
|---------------|-----------|------------|-----------|-----------|
| | Gamble Sc | Gamble VSc | Gamble S | Gamble VS |
| % of subjects | 10.81 | 89.19 | 20.00 | 80.00 |

Table 14: Stage two choice by treatment for groups 7, 8, 9 and 10

| % OF SUBJECTS CHOOSING SAFE OPTION | CONTROL | TREATMENT |
|--|---------|-----------|
| who experienced regret at stage one | 39.21 | 58.94 |
| who experienced no regret at stage one | 59.09 | 71.91 |

Table 15: Percentage of subjects who chose the relatively safe option at stage two, conditional on being in groups 5, 6, 11 or 12

the treatment, of having the same action from stage one repeat in stage two, had some effect on choice behaviour. It is not possible to say, simply from this level of data, however, whether the hypotheses outlined earlier are true or false, in that the cause for this change in behaviour was the experience of regret at stage one.

Secondary analysis on groups 7, 8, 9 and 10

For completeness, we also present the same descriptive statistics for those students who fell into groups “7 and 8” & “9 and 10” in Table 14. These results show a clear preference for the very safe gambles⁶⁵.

3.8.4 *Testing the hypotheses*

Breaking the data up further⁶⁶ allows us to observe the choice behaviour at stage two, dependent not just on whether the subject was in the control or treatment group, but dependent on whether or not the subject experienced regret resulting from their choice at stage one. Table 15 indicates the percentage of subjects who chose the relatively safe option (Gamble R or Gamble Rc), broken down by both treatment/control and also the experience of regret at stage one.

This data allows us to run preliminary tests on the experimental hypotheses.

Hypothesis 2A

The limited number of subjects who fell into groups 7 and 8 (13 and 26 respectively) prevents us from interpreting meaningful results from a test of Hypothesis 2A, but the data is presented for completeness. It is tested in the simplest fashion as

$$H_0 : \text{proportion of subjects choosing VS at point 7} = \text{proportion of subjects choosing VSc at point 8}$$

⁶⁵ An equivalent set of chi-squared tests are not possible in this case as several of the observed and expected frequencies fall below 5

⁶⁶ specifically for those subjects in groups 5, 6, 11 and 12, and excluding those 8 subjects who were unable to understand the Holt and Laury procedure

| | TREATMENT | CONTROL |
|---|-----------|---------|
| # who chose relatively risky option | 3 | 3 |
| # who chose relatively safe option | 23 | 10 |
| % who chose relatively safe option | 88.5% | 76.9% |
| χ^2 test for equality of proportions: $\chi^2(1) = 0.8864$ $p = 0.346$ | | |

Table 16: χ^2 test of Hypothesis 2A

| | TREATMENT | CONTROL |
|---|-----------|---------|
| # who chose relatively risky option | 18 | 31 |
| # who chose relatively safe option | 26 | 20 |
| % who chose relatively safe option | 59.1% | 39.2% |
| χ^2 test for equality of proportions: $\chi^2(1) = 3.7361$ $p = 0.053$ | | |

Table 17: χ^2 test of Hypothesis 2B

H_1 : proportion of subjects choosing VS at point 7 \neq
proportion of subjects choosing VSc at point 8

using a chi-squared test for equality of proportions. This is reported in 16.

Hypothesis 2B

Hypothesis 2B, as stated on page 149, posits that “...the proportion of subjects choosing the relatively safe option, R, at point 11 will be different to the proportion of subjects choosing the relatively safe option, Rc, at point 12.” This can be tested in the simplest fashion as

H_0 : proportion of subjects choosing R at point 11 =
proportion of subjects choosing Rc at point 12
 H_1 : proportion of subjects choosing R at point 11 \neq
proportion of subjects choosing Rc at point 12

using a chi-squared test for equality of proportions. This is reported in 17.

Table 17 shows that although the difference in proportion of subjects choosing R at point 11, compared to Rc at point 12, is approximately 20 percentage points, we are unable to reject the null hypothesis H_0 in favour of the alternative hypothesis H_1 at the 5% level ($p = 0.053$), though clearly this is very close to the threshold. Thus, in the strictest sense, this represents a failure of Hypothesis 2B. However, dividing the sample into those who participated in the original sessions (in 2011) and those who participated in the later sessions (in 2014), as per Table 18, shows that there was a difference across the sessions, with the sample from the original sessions choosing in support of Hypothesis 2B, and those from the later sessions against Hypothesis 2B.

| | TREATMENT | CONTROL |
|---|-----------|---------|
| % who chose relatively safe option in original sessions | 58.1% | 30.0% |
| % who chose relatively safe option in later sessions | 61.5% | 52.4% |
| χ^2 test for original sessions: $\chi^2(1) = 4.8673$ $p = 0.027$ | | |
| χ^2 test for later sessions: $\chi^2(1) = 0.2731$ $p = 0.601$ | | |

Table 18: χ^2 test of Hypothesis 2B by session

| | TREATMENT | CONTROL |
|-------------------------------------|-----------|---------|
| # who chose relatively risky option | 1 | 1 |
| # who chose relatively safe option | 6 | 10 |
| % who chose relatively safe option | 85.7% | 90.9% |

Table 19: data for Hypothesis 3A

Hypothesis 3A

As with Hypothesis 2A the limited number of subjects who fell into groups 9 and 10 (7 and 11 respectively) prevents us from interpreting meaningful results from a test of Hypothesis 3A, but the data is presented for completeness. It is tested in the simplest fashion as

$$H_0 : \text{proportion of subjects choosing VS at point 9} = \text{proportion of subjects choosing VSc at point 10}$$

$$H_1 : \text{proportion of subjects choosing VS at point 9} \neq \text{proportion of subjects choosing VSc at point 10}$$

however the low count of participants in the cells makes a chi-squared test for equality of proportions not possible. This data is reported in Table 19.

Hypothesis 3B

Hypothesis 3B, as stated on page 151, posits that “...the proportion of subjects choosing the relatively safe option, R, at point 5, will be no different to the proportion of subjects choosing the relatively safe option, Rc, at point 6.” This can be tested in the simplest fashion as

$$H_0 : \text{proportion of subjects choosing R at point 5} = \text{proportion of subjects choosing Rc at point 6}$$

$$H_1 : \text{proportion of subjects choosing R at point 5} \neq \text{proportion of subjects choosing Rc at point 6}$$

using a chi-squared test for equality of proportions. This is reported in Table 20.

Table 20 shows that the difference in proportion of subjects choosing R at point 5, compared to Rc at point 6, is approximately 13 percentage points, but we are unable to reject the null hypothesis in favour of the alternative, at the 5% level ($p = 0.065$). The result, that we fail to reject the null hypothesis, supports Hypothesis 3B. However, it should be noted that, as with Hypothesis 2B, we are particularly close to the

| | TREATMENT | CONTROL |
|---|-----------|---------|
| # who chose relatively risky option | 25 | 39 |
| # who chose relatively safe option | 64 | 56 |
| % who chose relatively safe option | 71.95% | 58.9% |
| χ^2 test for equality of proportions: $\chi^2(1) = 3.4038$ $p = 0.065$ | | |

Table 20: χ^2 test of Hypothesis 3B

| HYPOTHESIS | PASS / FAIL |
|------------|----------------------|
| 2A | <i>lack of data</i> |
| 2B | FAIL (<i>just</i>) |
| 3A | <i>lack of data</i> |
| 3B | PASS (<i>just</i>) |

Table 21: Summary of hypothesis tests

threshold for finding a significant difference in the preference for the relatively safe option between the control and treatment groups.

3.8.5 Summary of the results

As presented in Table 21 and discussed in 3.5.16, a failure of Hypothesis 2B, but a success of Hypothesis 3B represents the conclusion that simply framing the decision problem in terms of a repeated, previously regretted action, has no effect on the choice made in stage two. However, as both hypotheses were close to the boundary between success and failure, it is worth considering a broader interpretation of the results that looks at a wide range of reasons why we may have seen the results unfold as they ultimately did. In any case, the combination of both an unclear conclusion, and the lack of data to test certain hypotheses, suggests more work should be conducted in this area to better understand the problem at hand.

3.9 ANALYSIS OF THE RESULTS

In this section we will explore four possible explanations for the observed results. These are

1. Insufficient sensitivity
2. Experimental confounds
3. Population changes
4. Small changes in context are simply not important for most people

3.9.1 Insufficient sensitivity

As will be discussed extensively in 3.10, the experimental set-up used is certainly not ideal for generating a meaningful experience of regret, and then measuring small changes to future regret aversion, split out

into action and decision components. In brief, a lab set-up where relatively small sums of money are at play, and true responsibility for choices is hard to generate, is far from ideal conditions. However, as the significance test on Hypothesis 2B was particularly close to the usual threshold, one plausible conclusion is that there really is a meaningful result to be found here⁶⁷, but simply our experiment was not sensitive enough to detect it. Alternatively this can be thought of as saying despite our best efforts to control the experiment in the lab, there was simply too much noise from other sources to mask the effect of the treatment on those who experienced regret.

As what we are looking to identify is the additional effect of the treatment⁶⁸ on those who experienced regret, compared to those who didn't, we can try an alternative formulation of the hypotheses which attempts to control for the other sources of variation between the control and treatment groups, not related to the experience regret⁶⁹. In theory, this makes our effect of interest easier to identify. Doing this requires a "difference-in-difference" model, which can be formulated⁷⁰, using interactive dummies, as follows:

$$P(\text{safe choice in stage two})_i = c + \alpha_1 T_i + \alpha_2 R_i + \alpha_3 T_i R_i + \varepsilon_i \quad (3.1)$$

where T_i is a dummy indicating that subject i was in a treatment session, and R_i is a dummy indicating that subject i experienced regret in stage one. Under this formulation α_3 represents the additional effect of the treatment on those who experienced regret at stage one, compared to those who didn't. That is to say, the effect of seeing a previously regretted action, over and above any other effects of the treatment on the control group⁷¹.

Estimating and interpreting this formulation is easiest using a logit model where the dependent variable is measured in the odds metric rather than in the probability metric (Buis [14]). Under this model, the constant term is interpreted as "baseline" odds, where all the dummies are set to zero. That is to say, the odds of a subject, in the control condition who did not experience regret in stage one, choosing the relatively safe option compared to the relatively risky option in stage two. The coefficient α_1 then becomes the multiplicative effect of being in the treatment condition compared to the control condition. Our coefficient of interest α_3 then tells us by how much the effect of the treatment differs between those who experienced regret in stage one and those who did not, and hence we are looking at the significance of this coefficient as a test of our primary hypothesis. The results of this

⁶⁷ where the result is that a link between regret and the action which caused the experience of regret, will impact choice behaviour when the action reappears in a future decision. This implies that the context and circumstance which the initial decision was taken under plays a more meaningful role in the experience of regret than traditional regret-based theories, which often use a modified version of a non-expected utility framework, have accounted for.

⁶⁸ of a subject seeing in stage two, an action that they previously saw in stage one

⁶⁹ potentially an inherent action aversion

⁷⁰ reference to the "safe choice" in stage two more precisely refers to the *relatively safe* choice in stage two

⁷¹ as discussed in the design of the experiment, we had assumed the effects of the treatment on the control group would be zero

| | (1) SCS ₂ |
|---------------------------------------|-------------------------|
| baseline | 1.436 (0.083) |
| Treatment? | 1.783 (0.066) |
| Was regret experienced in stage one? | 0.449* (0.024) |
| Multiplicative treatment regret dummy | 1.256 (0.664) |
| Observations | 279 |

Exponentiated coefficients; *p*-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 22: Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12. Shown are the multiplicative effects of each independent variable on the odds ratio.

regression⁷², for those subjects who fell in groups 5, 6, 11 and 12⁷³, are presented in Table 22.

Table 22 shows a significant effect of the experience of regret in stage one on the odds of choosing the safe option in stage two for those who were in the control group. As the effect size is less than 1, this is interpreted as the experience of regret reducing the odds of choosing the safe option in stage two. As experiencing regret in stage one is equivalent to “losing” in stage one, this is consistent with a “target earnings” strategy, whereby those subjects who lost in stage one need to behave in a more risky fashion in stage two in order to achieve their earnings goal for the session.

Also significant at the 10% level is the effect of the treatment on the odds of choosing the safe option in stage two, for those who did not experience regret in stage one. This is a more peculiar result, as the treatment was designed to have no effect on those in the control group. The effect size is greater than 1, indicating those who were in the treatment group were more likely to chose the safe option in stage two compared to the control group, conditional on being a “winner” in the first stage.

The most plausible causal explanation is that, for those in the treatment group, though they picked (and won) by choosing Gamble S in the first stage, they equivalently saw Gamble R lose in stage one. This result then suggests that they were more likely to choose Gamble R (over Gamble VR) in stage two, compared to the mathematically equivalent Gamble Rc (over Gamble VRc) in the control group. This can be interpreted as evidence in support of the Gambler’s Fallacy, or

⁷² The dependent variable “SCS₂” stands for “Was the relatively safe choice made in stage two?”

⁷³ again, we are looking at just these groups due to the lack of subjects in groups 7, 8, 9 and 10, and also controlling for the possibility that Type A regret may work in different ways to Type B regret

the “law of small numbers”, in that there was an expectation of Gamble R being more likely to win in the second stage *because* they saw it lose in the first stage. This belief in the “law of small numbers” would, in this experiment, imply that subjects who observed a roll of the single die, between 0 and 6 in stage one, would predict a greater than 30% chance of observing a roll between 7 and 9 in stage two⁷⁴.

As such, the Gambler’s Fallacy does indeed explain the significance and direction of α_2 . Those who experienced regret in stage one, and fell into the treatment group, were more likely to select Gamble R, which resolves in their favour if a number between 7 and 9 is rolled (which, due to the Gambler’s Fallacy, they now believe is more probable) compared to had they fallen into the control group, where the uncertainty in stage two is resolved by the combination of two dice, of which the subjects had observed no previous rolls.

However, our primary hypothesis, that the α_3 coefficient would be significant, is shown to be false, as the p-value is equal to 0.664. Thus, even when controlling for the effect of target earning wealth effects and the law of small numbers, we do not see an effect of action regret, as distinct from decision regret, on the choice made in stage two. We will explore potential reasons for this in the subsequent sections.

3.9.2 *Experimental confounds*

As indicated earlier, there were certain characteristics of the subjects which appeared to be correlated with both the allocation to the treatment and control groups, and the choices made in stage one. These were the past mathematical background of the subjects, and also the underlying risk preference as given by the Holt and Laury procedure. These correlations could possibly be driving the results shown in 3.9.1 (specifically the significance of α_1 and α_2 , and hence potentially the insignificance of α_3) rather than any underlying theoretical reasons.

As such, we would want to control for these characteristics of the subjects and repeat the analysis shown in Table 22. This can be achieved by including the potential confounding variables as additional regressors in Equation 3.1. The results of this are shown in Table 23.

As shown in Table 23, very little changes in terms of the significance of α_1 , α_2 and α_3 through the addition of the extra regressors. As such we can be confident that experimental confounds are not the cause of the results discussed in 3.9.1.

3.9.3 *Population changes*

As discussed at the time of presenting the results of Hypothesis 2B (on page 162), there appears to be very different behaviour at play when looking at the hypothesis simply for those subjects who participated in the original run of sessions (in 2011) and those who participated in the later run of sessions (in 2014). Hence a possible cause for the experimental results is that we are simply looking at two fundamentally different populations, and hence it is not possible to aggregate

⁷⁴ this can either be thought of as “I observed a number 3 in stage one, therefore every other number, aside from 3, is slightly more likely in stage two” which would imply a greater than 30% chance of observing a number between 7 and 9, or “I observed the numbers 7, 8 and 9 not be rolled in stage one, therefore they are more likely to be rolled in stage two” which has the same effect

| | (1) SCS ₂ | (2) SCS ₂ | (3) SCS ₂ ? |
|---------------------------------------|-------------------------|-------------------------|---------------------------|
| baseline | 0.722 (0.505) | 1.427 (0.281) | 0.716 (0.546) |
| Treatment? | 1.784 (0.067) | 1.781 (0.068) | 1.782 (0.069) |
| Was regret experienced in stage one? | 0.466* (0.032) | 0.449* (0.024) | 0.466* (0.032) |
| Multiplicative treatment regret dummy | 1.336 (0.584) | 1.255 (0.665) | 1.335 (0.585) |
| HL Safe Choices | 1.119 (0.120) | | 1.119 (0.120) |
| A-level maths | | 1.008 (0.980) | 1.012 (0.972) |
| Observations | 279 | 279 | 279 |

Exponentiated coefficients; *p*-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 23: Logit model with additional regressors showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12. Shown are the multiplicative effects of each independent variable on the odds ratio.

| | | 2011 | 2014 | TESTING THE DIFFERENCE |
|----------------------------|---|-------|-------|--|
| Self reported happiness | 1 | 0 | 3 | |
| (% of subjects in session) | 2 | 1 | 6 | |
| | 3 | 10 | 15 | |
| | 4 | 33 | 33 | |
| | 5 | 37 | 35 | Grouping into “sad”, “neither” and “happy” |
| | 6 | 13 | 6 | χ^2 test of independence |
| | 7 | 5 | 2 | $\chi^2(2) = 9.11$ $p = 0.01$ |
| Mean age | | 20.52 | 21.84 | t-test of mean diff $t = 3.48$ $p = 0.00$ |
| % male subjects | | 47.24 | 41.27 | χ^2 test of equality of % $\chi^2(1) = 1.15$ $p = 0.29$ |

Table 24: Descriptive statistics from the original (2011) and later (2014) sessions

the results across all three sessions. Evidence of this comes from the summary statistics of the participants in the original sessions compared to those in the later sessions. These are presented⁷⁵ in 24.

Table 24 indicates that we are potentially considering two fundamentally different populations of subjects, due to the time difference between running the original and later sessions. As the subjects were drawn from a student population at the University of Warwick, the change in happiness may possibly be attributed to the higher level of tuition fees now levied on undergraduate students and the above-inflation rises in graduate fees between the two time periods. This has the potential to affect the way subjects respond to the experience of negative emotions and their associated risk attitudes.

In order to test whether the delay in running additional sessions is responsible for the results found, we can split the sample into those in the original and those in the later sessions, and re-run the regression in Table 22. These results are shown in Table 25.

As the results show, there appear to be significant differences in the coefficients (both in terms of magnitude and significance) between the original and later sessions. In the original sessions, the target wealth effect appears significant, but in the later sessions, the treatment effect appears significant. To test this, we can include interactive dummies (between each of the factor dummies and the session indicator), and test the significance of the additional terms. The results are shown in Table 26.

None of the interactive dummies appear significant even at the 10% level, though this may be a sample size issue, as the inclusion of additional regressors makes it more difficult to find such a result. Consequently, we conclude there may be population issues at play, though the sample size makes it difficult to identify at a sufficiently significant level.

3.9.4 *Small changes in context are simply not important for most people*

As the above sections demonstrate, there is no clear evidence for the effect of “action regret” as distinct from “decision regret” on subsequent

⁷⁵ A score of 1 in self-reported happiness corresponds to “completely unhappy” whereas 7 is “completely happy”, with the midpoint of 4 being “neither happy nor sad”

| | (Original) SCS ₂ | (Later) SCS ₂ |
|---------------------------------------|--------------------------------|-----------------------------|
| baseline | 1.652 (0.057) | 1.125 (0.732) |
| Treatment? | 1.400 (0.400) | 2.667 (0.057) |
| Was regret experienced in stage one? | 0.259** (0.005) | 0.978 (0.968) |
| Multiplicative treatment regret dummy | 2.308 (0.213) | 0.545 (0.493) |
| Observations | 175 | 104 |

Exponentiated coefficients; *p*-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 25: Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, split according to the sessions. Shown are the multiplicative effects of each independent variable on the odds ratio.

choice behaviour. There is always the possibility that there really is an effect, but it is too small to notice, even with the above pieces of analyses designed to tease out the answer.

An implicit assumption in the above argument is that the effect is very small for everybody, hence making the average effect in the population too small to detect given the experimental design. An alternative hypothesis is that the average effect may well be small for the population, but there is a small sub-section of the population for whom the effect is large, with the remainder of the population being completely unaffected. Without any intuition as to who this sub-section of the population may be, we are left looking at too broad a group in order to try and identify the effect.

One obvious starting point is to sub-divide the population according to gender. The results of running Equation 3.1 separately for males and females is shown in Table 27

In Table 27 we see the results appear very different for males and females. Females have a much stronger baseline preference for the safe option in stage two, and are especially responsive to the target wealth effects ($\alpha_2 = 0.234$ $p = 0.005$), whereas males appear more subject to the Gambler's Fallacy ($\alpha_1 = 2.380$ $p = 0.064$). For females, the effect of action regret is now strong in size, but still not significant at a sufficient level ($\alpha_3 = 2.810$ $p = 0.165$). These results are indicative, however, of differing sensitivities between males and females to the effects identified in the experiment. From the perspective of identifying significant regret effects, it would appear to make more sense to focus future research on females.

| | (1) SCS ₂ |
|---------------------------------------|-------------------------|
| baseline | 1.436 (0.083) |
| Treatment? | 1.610 (0.191) |
| Was regret experienced in stage one? | 0.298** (0.007) |
| Multiplicative treatment regret dummy | 2.006 (0.285) |
| TxSession | 1.297 (0.593) |
| RxSession | 2.567 (0.111) |
| TxRxSession | 0.347 (0.300) |
| Observations | 279 |

Exponentiated coefficients; *p*-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 26: Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, including interactive session dummies. Shown are the multiplicative effects of each independent variable on the odds ratio.

| | (female) SCS ₂ | (male) SCS ₂ |
|---------------------------------------|------------------------------|----------------------------|
| baseline | 2.267** (0.008) | 0.917 (0.768) |
| Treatment? | 1.261 (0.598) | 2.380 (0.064) |
| Was regret experienced in stage one? | 0.234** (0.005) | 0.857 (0.758) |
| Multiplicative treatment regret dummy | 2.810 (0.165) | 0.588 (0.483) |
| Observations | 152 | 127 |

Exponentiated coefficients; *p*-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 27: Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, split by gender. Shown are the multiplicative effects of each independent variable on the odds ratio.

| | 2 | |
|---------------------------------------|---------------------------------|-----------------------------|
| | (Not Happy) SCS ₂ | (Happy) SCS ₂ |
| baseline | 0.895 (0.739) | 1.950* (0.015) |
| Treatment? | 7.153*** (0.001) | 0.821 (0.617) |
| Was regret experienced in stage one? | 0.762 (0.565) | 0.285* (0.043) |
| Multiplicative treatment regret dummy | 0.290 (0.110) | 3.291 (0.168) |
| Observations | 139 | 140 |

Exponentiated coefficients; p -values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 28: Logit model showing the determinants of choice at stage two for those in groups 5, 6, 11 and 12, split by happiness. Shown are the multiplicative effects of each independent variable on the odds ratio.

Similarly, in Table 28 we divide the population up into those who report⁷⁶ as “not happy”⁷⁷ and those who report as being “happy”^{78,79}, under the hypothesis that regret may have differing effects depending on the present emotional state of the subject.

The standout result in Table 28 is the coefficient on the Treatment variable for those who self report as “not happy” ($\alpha_1 = 7.153$ $p = 0.001$). This can be interpreted as an especially strong belief in the Gambler’s Fallacy for those who aren’t happy, yet won in the first stage. The strength of the result is surprising, but perhaps those who have seen bad things happen recently interpret the win in the first stage as a sign that their luck is about to change. Further research investigating the link between the Gambler’s Fallacy and self-reported happiness would be recommended as there is none which currently exists in this area.

From a regret standpoint, the more interesting result of Table 28 is that the coefficients on the “regret x treatment” interactive dummy variables, α_3 , appear to be reflected around 1 depending on whether you are “happy” or “not happy”. If you are “not happy”, then $\alpha_3 = 0.290$ ($p = 0.110$), saying that the effect of the treatment on those who experienced regret compared to those who didn’t is to *lower* their odds of choosing the relatively safe option by 3 times. If you are “happy”, then $\alpha_3 = 3.290$ ($p = 0.168$), saying that the effect of the treatment on those who experienced regret compared to those who didn’t is to

⁷⁶ The report is taken after the decisions have been taken and videos viewed, but before the result of the Holt and Laury procedure and payment is known. This creates a potential issue with the direction of causality, but due to there being a relatively small number of participants who won in stage two, you can limit the analysis to only those subjects who received the same payoff as a result of their choices in stages one and two, and, qualitatively, the results still hold.

⁷⁷ this consists of those reporting as “Completely Sad”, “Very Sad”, “Fairly Sad” or “Neither Happy nor Sad”

⁷⁸ this consists of those reporting as “Fairly Happy”, “Very Happy” or “Completely Happy”

⁷⁹ the division is, conveniently, an almost exactly 50/50 split

raise their odds of choosing the relatively safe option by 3 times. The significance levels are clearly an issue, but the size of the coefficients goes some way to explaining the failure to find the effect earlier, as it appears to be working in different ways on different sub-sections of the population. In the case of those who are “not happy”, the effect is that they are less willing to choose the safe option, which was the option which they regretted choosing in the first stage. In the case of those who are “happy”, the effect is that they are more willing to choose the safe option, despite it being the option they regretted choosing in the first stage. Intuitively, this direction seems right as happy people may be more willing to “forgive and forget” than those who are not happy, but further research would be required in this area to establish the true reasons and the statistical significance.

Clearly, by sub-dividing the samples into smaller and smaller groups, we run the risk of finding these results simply due to the outlier actions of a few specific experimental subjects. In the case of the coefficient on the Treatment variable for those who self-report as “not happy”, the magnitude of the coefficient ($\alpha_1 = 7.153$) is immediately suspicious, as it seems somewhat implausible that playing around with die and labels (without actually changing the probabilities or payoffs of the underlying gambles) can cause a group to be 7 times more likely to do anything, let alone change their gambling decisions⁸⁰. That being said, the purpose of research is to simultaneously build on what has gone before and suggest avenues for future work, so repeating these results with similar subjects but in a new environment would be a good first step.

3.10 WHY YOU WOULDN'T DESIGN AN EXPERIMENT LIKE THIS IF YOU HAD THE CHOICE

When designing a gambling based experiment, there are a few different designs which can be used to convey the idea of “decision under uncertainty”. In this experiment, we take the approach of “decision under risk” where the probabilities associated with each state of the world are explicitly stated, in both the form of a percentage and the numbers on the face of a ten sided die⁸¹ which are linked to each state of the world. This gives a high level of control over the important information which participants can use to make a decision, and places them in a position of having no prior information, unbeknown to the experimenter, which could help them make a decision one way or another. The use of a randomised treatment and control group permits the use of indirect control over other factors which could potentially affect an individual's decision making (such as, for example, their inherent preference for risk), but, where possible, the use of direct control is preferable for identifying the causality present in the experiment.

By contrast, several of the initial designs of this experiment used decisions under uncertainty at both the first and second stage, where explicit probabilities for each state of the world were not specified, but a “subjective probability estimate” could be inferred from a set of information provided to the participants. The first design of the

⁸⁰ though there are 73 subjects included in this specific comparison, so it's not an incredibly small subject pool

⁸¹ or two ten-sided dice in the control group at stage two

experiment had individuals making “buy or sell” decisions on an amount of stock, where the past history of the stock was provided in the form of the graph, and a set of information surrounding the wider conditions of the economy and industry was available for the subject to consult. The second design of the experiment had the subjects watching videos of several horse races, and being asked to make a bet on a horse having been provided with their odds and a form guide.

3.10.1 *Hard to generate regret from dice*

The benefit of using a “decision under uncertainty” design as opposed to a “decision under risk” design is that, because the probability estimates of the stock increasing or decreasing in price, or equivalently the horse winning or losing the race, are derived by the participants themselves from a set of information, the participants feel a sense of responsibility for not simply their own decisions (whether to buy or sell the stock, or which horse to bet on), but also for their own subjective probability estimates upon which their decisions are based. As demonstrated earlier, the existing literature shows a strong relationship between responsibility and experienced regret, and, hence, in an experiment which is designed to look how the experience of regret influences subsequent choice, you would prefer that effect to be as strong as possible, and hence design an experiment where the individual feels highly responsible for the outcome of their decisions at stage one. Intuitively, in order to generate the largest possible experience of regret, you would want the subject to believe that there is something they could have either known or done which would have revealed to them the true outcome of the uncertainty, prior to their decision, and hence informed them of the correct decision to take. This is especially true when the decision under uncertainty involves an outcome which has already been resolved (such as the stock price of a firm *from a previous time period*, or the video of a horse race *which has already been run*), but is unknown to the subject in the experiment. In this case, there is an impression that all uncertainty has been resolved⁸², and it is merely a question of whether the subject can work out which is the “correct answer” in terms of the decision to take.

The flip side of this argument is that is an experiment which uses a mechanical randomisation device, such as a die, has an explicit element of randomness built in. That is, there is no information which the participant could have known, or have been expected to know, which would have resulted in a “better” understanding of which face of the die was likely to come up. As such, ex-post, the subject has an ability to rationalise their decision, and absolve themselves of responsibility, by realising that there is nothing they could have done or known, ex-ante, which would have given them a reason to take a different decision, and potentially improve their payoff. This reduction in responsibility is likely to cause a reduction in the amount of regret experienced, and, equivalently, a reduction in the amount of regret anticipated, from either a stage one or stage two decision. As such, any results which are present in the data, linking the experience of regret to subsequent choice behaviour, will be harder to spot.

⁸² that is, one state of the world will occur with probability $p = 1$, and all others have a zero probability of occurring.

However, from the perspective of the subject in the experiment, it is important for them to believe that the whole process is fair and transparent, in order that they give due care and consideration to the decisions which they are being asked to take. In this regards, the use of dice to resolve the gambles is a more transparently fair process (in that the outcome is beyond the control of the experimenter as well as the subject) than the use of either stocks or horses, which could have been chosen precisely because they offered the promise of high reward, but ultimately failed to deliver. To this end, the need for a *believable* experiment outweighs the need for a particularly *emotional* experiment, and the benefits of direct control, in having the same explicitly stated probabilities for both control and treatment groups, outweigh the benefits of having non-stated, subjective, malleable probabilities under uncertainty, and hence justify the decision to use dice as the method through which to resolve the gambles in this experiment.

3.10.2 *Discrete choice*

The experimental procedure outlined relies on, at both stage one and stage two, the subjects making a choice between only two options, giving the main experimental variable of interest as the proportion of subjects who chose the relatively safe option at stage two, a binary outcome, in that either the subject chose the relatively safe option or they did not. Once this choice is made, the gambles are played out in front of the participants, and they are paid according to the outcome of their choice and the outcome of the gamble, so the subjects are correctly incentivised to reveal their true preferences when making the choice.

However, what the outcome of the stage two choice does not reveal is *how strongly* the subject preferred one gamble over the other. In the language of limited dependent variables, there is an underlying latent variable of the difference in utility between the relatively safe and relatively risky gambles, and the outcome of the binary choice simply reveals whether the latent variable is positive or negative, and not the absolute magnitude. In an ideal experimental design, we would find a way to tease out the magnitude, by asking the subject to reveal, in essence, how strongly they feel about their decision. The benefit of this approach is that we would be able to measure, for example, whether the impact of the experience of regret from stage one is only significant for subjects that were making very marginal decisions⁸³, and, for those individuals who were affected, to what degree of magnitude did it affect their stage two decision.

Stating a "selling price"

One possible approach to generating a continuous measure of gambling preference for, as an example, the stage two decision, is to ask the individual to state the minimum price at which they would be willing to "sell" one of the two gambles on offer for the other. That is, for example, those subjects who are at point 11 would be "endowed" with Gamble R, and be asked to state the minimum price at which they

⁸³ it can be argued that an individual would feel less regret about not taking a decision they never had any intention of taking compared to if they were very close to choosing the ex-post superior option in the first place. That is, the individual would find it easier to imagine a counterfactual where they chose the ex-post superior option, and realised a better payoff, the closer they were to taking that option.

would be willing to swap Gamble R for Gamble VR. The price would be stated in the following fashion:

Presently you have chosen Gamble R, which offers you a 30% chance of winning £14 and a 70% chance of winning £0. Gamble VR, in contrast, offers you a 10% chance of winning an amount of money £x, and a 90% chance of winning £0. Please state the minimum possible value of x for which you would be willing to exchange Gamble R for Gamble VR.

There would then need to be a market mechanism which determined whether Gamble R is indeed exchanged for Gamble VR at that price, thus allowing for a “selected” and an “unselected” gamble in the mould of the original experiment. Both selected and unselected gambles would then be resolved by the die mechanism, permitting the experience of regret to occur from both exchanging Gamble R for VR when R was ex-post superior and also not exchanging Gamble R for VR when VR was ex-post superior.

In order for such an experiment to produce meaningful results, however, the “market mechanism”, which determines whether or not the gambles are exchanged at the price stated, must provide subjects with an incentive to truthfully reveal the minimum price at which they would be willing to sell Gamble R, in order that the answer can be assumed to be a marginally small amount above the true point of indifference between the two gambles. That is, the subjects in the experiment must have a dominant strategy to reveal their true preferences.

A mechanism that is commonly used in experiments to achieve such a measure is the Becker-DeGroot-Marschak mechanism (BDM) (Becker et al. [3]). Applying the mechanism to this experiment and this example, the market value of x would be randomly drawn by the experimenter (call this value x^*), and, if x^* exceeds the value of x stated by the subject, they would exchange Gamble R for Gamble VR, where Gamble R is now defined as a 10% chance of winning £ x^* and a 90% chance of winning £0. Should the value of x^* fall below the value of x stated by the subject, they would retain Gamble R as their chosen gamble, but still see the outcome of the roll of the die which defines both their payoff as a result of retaining Gamble R, and the payoff they would have received, under the now completely defined Gamble VR, had they chosen a value of x in excess of x^* .

This mechanism is, on an intuitive level, “...simple and presumed to induce truth-telling. Individuals have the incentive, it is believed, to report their true maximum willingness-to-pay” (Horowitz [39, p7]), or, as in both this example and the original BDM paper, an incentive to reveal their minimum selling prices. As such, it seems like an excellent candidate mechanism through which to elicit a continuous variable in the second stage of the experiment.

Problems with the BDM mechanism

There is a large theoretical and experimental literature, however, on the failings of the BDM mechanism to deliver the incentive compatibility which is often claimed. The criticisms can be broadly categorised, for these purposes, into theoretical failings, practical failings, and regret-specific failings.

Theoretical failings

The logistics of the BDM mechanism do not make explicit reference to the *type* of good being sold or bought by the experimental subject; only the method through which they must submit their bid. Indeed the BDM mechanism has been applied to a very wide range of goods, for example, famously by Kahneman et al. [51] where it is used to effect the sale of mugs. One type of good, however, has been shown to place some restrictions on the applicability on the use of the BDM mechanism. Karni and Safra [54] examine the BDM mechanism when dealing with the sale of lotteries⁸⁴, and prove that "...the elicitation of certainty equivalents of all lotteries, using the experimental method of Becker, DeGroot, and Marschak, is possible if and only if the preference relation is representable by an expected utility functional" [54, p676]. It is the double implication of this result which is of relevance here, as, for the whole of this experiment, we are assuming the subjects are including anticipated regret in their decision making process, which, in the guise, for example, of Loomes and Sugden [58], is an explicitly non-EUT. As such, the BDM procedure, applied to subjects behaving with specifically non-EUT preferences, will not yield truthful revelation of certainty equivalents or selling prices. In a similar work, Horowitz [39] shows that even for non-random goods, the BDM is not incentive compatible for non-standard preferences (including regret and disappointment aversion), as the reported "certainty value" or selling price "...is not independent of the distribution of prices"[39, p10] which may be specified and draw from, at random, by the experimenter.

Practical failings

The last paragraph of Horowitz' work correctly states "[r]esearchers will also want to know whether any of these effects is large enough to matter in real world applications of the BDM". A study by Vlaev et al. [97] asked subjects to state the maximum amount of an endowment they would be willing to spend to avoid a small electric shock, using a standard BDM mechanism to decide whether this value was sufficient to "buy" the participant out of the electric shock⁸⁵. However, there were two possible conditions that subjects could have been assigned to. In the first condition, they received an endowment of 40p, and hence the BDM mechanism generated a random number between 0 and 40p. In the second condition, they were given an endowment of 80p, and hence the BDM mechanism generated a random number between 0 and 80p. In accordance with Horowitz' claim, the selling price was not independent of the distribution of prices, with a highly significant difference between the "demand for relief of pain" for those in the 80p endowment group compared to those in the 40p endowment group, as evidenced by Figure 53.

Likewise, an earlier experiment of Bohm et al. [11] holds the endowment constant and elicits the minimum selling price of a non-risky object ("...a card entitling the bearer to a given quantity of petrol"[11, p1080]), under three different upper bounds of the price range from

84 specifically, lotteries which have "...a finite number of prizes, i.e., a lottery($x_1, p_1; \dots; x_n, p_n$) where p_i is the probability of x_i , $\sum_{i=1}^n p_i = 1$, and for all i , $x_i \leq x_{i+1}$ "[54, p677] of which all the gambles in this experiment are of the form.

85 in addition, subjects were explicitly told why, given this mechanism, it was in their best interest to reveal their true value




Figure 53: Figure 3b from Vlaev et al. [97] showing “[d]emand curves for pain relief derived from the trials following experienced (consumed) pain of a long duration (i.e., 15 shocks)” [97, p314]. For each endowment, there were two separate “blocks” of activity. One block consisted of trials where the subject experienced either a high or medium shock, and the other consisted of trials where the subject experienced either a medium or low shock, hence the graph reports two different demand curves, for relief of the medium pain, given the particular block of trials that the subjects were in.

which the BDM mechanism draws a market price. Despite being explicitly told what the market value of the petrol was (with the card redeemable at a petrol station next to the university where the experiment took place), the mean reported selling price was significantly higher for the group whose upper bound on the BDM mechanism was higher than the market price, compared to the group whose upper bound was slightly lower than the market price. A third group stated selling prices where the upper bound was not explicitly defined, but was stated to be "...what we think is the maximum price any real buyer would be willing to pay for this card" [11, p1082], and were found to report selling prices not statistically significantly different from the low upper bound group, but significantly below the high upper bound group. The conclusion, that the high upper bound "...seem[ed] to make subjects overstate their selling reservation prices" [11, p1084], again gives reason to doubt the practical suitability of the BDM mechanism for incentivising subjects to reveal their true preferences, given the upper bound must be invented and stated, in some form, by the experimenter, in any given experiment wishing to use the BDM mechanism.

Regret-based failings

In addition to both the theoretical and practical problems in using the BDM mechanism to incentivise subjects to truthfully reveal their preferences, there are some further issues relating exclusively to the design of this experiment and the effect that we are trying to investigate. By studying the effects of experienced regret on subsequent choice (potentially relating to the effect on subsequent regret aversion), as previously mentioned, we are reliant on the experiment to both produce experienced regret, and be conducive to thoughts of anticipated regret. Introducing the BDM mechanism into the equation makes this more difficult for several reasons.

Firstly, in contrast to the second stage of the experiment where there is only binary choice, any decision submitted through a BDM mechanism in terms of a minimum selling price of, for example, Gamble R in the second stage, does not immediately rule out the possibility of one particular type of regret. In the binary choice experiment, an individual choosing Gamble R over Gamble VR immediately rules out the possibility of experiencing Type A Regret⁸⁶, and so it is simple for them to calculate the possibility of experiencing Type B Regret⁸⁷. In contrast, assuming any selling price stated by the subject has a strictly positive probability of being selected by the BDM mechanism, then any price submitted exposes the subject to both types of regret, conditional on the BDM mechanism deciding whether they keep Gamble R or exchange for Gamble VR. As such, when an individual is calculating their potential anticipatory regret from any bid submitted, they must additionally consider the probability distribution of the mechanism, as well as the probability that the eventually selected gamble will lead to the experience of regret. By making anticipatory regret harder to correctly calculate, there is an increased risk of subjects making errors in such a calculation, and hence their decision not accurately reflecting their true aversion to regret.

86 as, having chosen the relatively safe option at stage two, they could not experience regret from feeling they were overconfident and wishing they had acted more conservatively

87 which is the probability of Gamble R winning, conditional on the fact that Gamble VR didn't win

This “compound lottery” effect of the BDM mechanism not only makes the true effect of regret on the preference of the individual much more difficult to calculate, but also presents an additional opportunity for the decision maker to absolve themselves of any responsibility for the outcome of their second stage decision. Precisely because the price submitted in stage two does not prevent either Gamble R or Gamble VR from being chosen for the individual⁸⁸ by the BDM mechanism, the experience of regret can always be attributed to the combination of the resolution of the BDM mechanism (which selected which gamble the individual would ultimately own), and the roll of the dice which informed the individual whether or not they would have been better off had the BDM mechanism selected the other gamble for them instead. By having two separate points at which mechanical randomisation takes place, the ease with which the subject can rationalise the event is increased, and the lower the effect of regret will be, as we have already seen that regret is positively linked to the degree of responsibility that the subject feels for the outcome.

3.10.3 *Summary of the design*

The present design of the experiment is neither ideal for generating the experience of regret in stage one, nor ideal for measuring the degree to which the subjects were affected by the experience of regret at stage two, as shown above. However, the practicalities of running the experiment in this way give the greatest possibility of obtaining a “clean” result. That is to say, a result from which an observer would be able to clearly see the difference between treatment and control, what information was presented to subjects and what, specifically, could be inferred from it. Additionally, given the simplicity of the experiment, should any experimenter wish to repeat the experiment, it would be very easy to do so.

The above discussion is not given simply as a justification for why this particular experimental design was chosen, but rather to highlight the wide variety of complexities in designing such an experiment at all. Building an experiment to investigate a subtle and nuanced aspect of a difficult to observe emotion requires careful thought and consideration, but the interaction between the emotion, the two-stage design and the measurement means you end up choosing the “least worst” design rather than any sense of an “optimal” design.

There is certainly a theoretical argument for using a BDM mechanism to elicit preferences, but the practical limitations give us reason to doubt its suitability in this case. Coupled with the additional requirement for the participants in the experiment to understand both what the mechanism is, and how it works (i.e. the need for lotteries in order to resolve different lotteries), for it to be of use, the design chosen here is one which opts for simplicity rather than complexity.

3.11 CONCLUSIONS

The experiment was designed to look specifically of the effect of regret tied to an action on subsequent choice behaviour in a gambling context.

⁸⁸ in contrast to “by the individual” in the binary choice version

The above results suggest that the effect of the treatment⁸⁹ on those who experienced regret (compared to those who didn't) may well vary significantly according to different sub-sections of the population. Whilst the effect could be negligible for some (males) it may be important for others (females). The direction of the effect may even vary according to current emotional state (happy or not). Hence a failure to find an overall strong statistically significant result is potentially due to the samples chosen being too broad, and further work should be focussed on targeting those sub-sections of the population where an emotional link to an action, or label, is likely to play an important role in their future decision making.

Additionally, the experiment showed some evidence for the effect of the Gambler's Fallacy in this context, and also target wealth effects.

3.11.1 *On future experimental work*

One avenue for further work may be to step outside the lab and look to field or natural experiments to conduct similar hypothesis testing. As homogeneity of groups may be important to detect a result, and actions can be interpreted in the context of labels or brands, a large firm with a loyal customer base, including many repeat customers, could be a suitable environment.

The idea of "re-branding" is often used when large numbers of customers may have had a bad experience with a product or brand (potentially causing the experience of regret), but it would be interesting to explore if the true effect of re-branding on repeat purchase is different⁹⁰ for different sub-sections of the population.

One other environment which could prove fertile for future research is the world of online poker. In this environment, participants are frequently exposed to the possibility of experienced regret, and also make many frequent similar, but not identical, decisions. An experiment which looks at whether the experience of regret in a prior stage acts significantly on participants, especially when they are faced with a similar decision in the future, could make the distinction between action regret and decision regret by looking for instances where the same card (e.g. seven of hearts) that had caused the initial regret reappears in the subsequent decision stage. This could be compared against those decisions which were functionally the same⁹¹, but involved a different card (e.g. eight of spades), much in the same way that our experiment changed the label on the gambles, and how many die were being rolled, in the second stage.

Lastly, the experiment was, ideally, supposed to generate subjects who had experienced both Type A and Type B regret from stage one, and analyse the effect of the 'previously regretted action' treatment separately for both groups. However, as the resolution of uncertainty from the gambles generated very few subjects who chose Gamble S and lost, there were too few subjects to generate meaningful analysis for those subjects who experienced Type B regret. As such, for future work, it would be worthwhile testing whether results can be found for Type B regret, as we have essentially done for Type A, and consequently

⁸⁹ where the treatment is showing an "action" which may have caused regret in the previous stage, but keeping the probabilities and payoffs the same between treatment and control

⁹⁰ and, indeed, worth the cost

⁹¹ in the language of Texas hold-em poker, for example, this could be "calling with middle pair" on the river

investigate what happens when the previously regretted action becomes the relatively risky option in stage two, as opposed to the relatively safe option.

3.11.2 *On future theoretical work*

At present, no regret-based theory takes into account the context of the choice being made in such a way that distinguishes the decision being made from the action needed to effect that decision. Thus, this work adds to the literature on effects such as the omission and status-quo biases, which are, in effect, specific representations of how the action can be considered as distinct from the decision. With the omission bias, sometimes you may not need to do anything to effect your choice, and with the status-quo bias, you may need to act in a fashion which is considered abnormal or atypical in order to effect your choice. Therefore, in order to develop regret-based theories of decision making going forward, some account of the context or “framing” of the problem, and what each action, as distinct from the decision, represents to the individual decision makers, is needed.

The type of development which this work suggests, because of the two stage design of the experiment, is one where the context of the actions which are on offer are given by past experience the decision maker has with those actions. That is to say, the decision maker draws on information they learnt from previous encounters with the action, when deciding what would be the consequences which result should they choose to take the action again. In this experiment, because of the *similarity* between Gamble R in stage two and Gamble R in stage one (compared to the lower degree of similarity between Gamble Rc in stage two and Gamble R in stage one) the subject may be able to transfer some knowledge between the stages, and this information should be factored into the decision making process the second time around.

Therefore, an important determinant of the effect to which regret will impact upon choice is the degree of similarity which exists between a current option and a recallable situation which provides the decision maker with information about the affective consequences of taking a particular option again. As the individual finds it difficult to predict their own emotions, they are searching for information which will make that prediction more accurate, and find it easiest to do so when a high degree of similarity enables them to recall such information.

Part II

APPENDIX



APPENDIX FOR CHAPTER 2

| Problem Number | x_1 | p | y_1 | $1 - p$ | x_2 | q | y_2 | $1 - q$ | r^f | r^t |
|----------------|-------|-----|-------|---------|-------|-----|-------|---------|-------|-------|
| 1 | 200 | 0.8 | 50 | 0.2 | 200 | 0.2 | 50 | 0.8 | 0 | 90 |
| 2 | 200 | 0.2 | 50 | 0.8 | 200 | 0.5 | 50 | 0.5 | 0 | -45 |
| 3 | 200 | 0.5 | -50 | 0.5 | 50 | 0.5 | -50 | 0.5 | 150 | 75 |
| 4 | 50 | 0.8 | -50 | 0.2 | 200 | 0.2 | -200 | 0.8 | 0 | 150 |
| 5 | 50 | 0.2 | -50 | 0.8 | 200 | 0.5 | 50 | 0.5 | -250 | -155 |
| 6 | 200 | 0.5 | -200 | 0.5 | 200 | 0.8 | 50 | 0.2 | -250 | -170 |
| 7 | 50 | 0.8 | -200 | 0.2 | 200 | 0.5 | -50 | 0.5 | -300 | -75 |
| 8 | 200 | 0.5 | 50 | 0.5 | 50 | 0.5 | -50 | 0.5 | 250 | 125 |
| 9 | 50 | 0.8 | -200 | 0.2 | 200 | 0.8 | -200 | 0.2 | -150 | -120 |
| 10 | 50 | 0.2 | -50 | 0.8 | -50 | 0.5 | -200 | 0.5 | 250 | 95 |
| 11 | 50 | 0.5 | -50 | 0.5 | 50 | 0.5 | -200 | 0.5 | 150 | 75 |
| 12 | 50 | 0.2 | -200 | 0.8 | 50 | 0.2 | -50 | 0.8 | -150 | -120 |
| 13 | 200 | 0.5 | -200 | 0.5 | 50 | 0.5 | -200 | 0.5 | 150 | 75 |
| 14 | 50 | 0.2 | -200 | 0.8 | 50 | 0.8 | -200 | 0.2 | 0 | -150 |
| 15 | 200 | 0.8 | -200 | 0.2 | 50 | 0.2 | -200 | 0.8 | 150 | 270 |
| 16 | 50 | 0.5 | -200 | 0.5 | 50 | 0.8 | -200 | 0.2 | 0 | -75 |
| 17 | 200 | 0.5 | -200 | 0.5 | 50 | 0.8 | -50 | 0.2 | 0 | -30 |
| 18 | 200 | 0.8 | -50 | 0.2 | 200 | 0.5 | 50 | 0.5 | -100 | 25 |
| 19 | -50 | 0.2 | -200 | 0.8 | 200 | 0.2 | -200 | 0.8 | -250 | -50 |
| 20 | -50 | 0.2 | -200 | 0.8 | 200 | 0.5 | -200 | 0.5 | -250 | -170 |
| 21 | 200 | 0.8 | 50 | 0.2 | 200 | 0.5 | -50 | 0.5 | 100 | 95 |
| 22 | 200 | 0.2 | -200 | 0.8 | 50 | 0.2 | -50 | 0.8 | 0 | -90 |
| 23 | 200 | 0.8 | -50 | 0.2 | 50 | 0.8 | -50 | 0.2 | 150 | 120 |
| 24 | -50 | 0.8 | -200 | 0.2 | 200 | 0.2 | -200 | 0.8 | -250 | 40 |
| 25 | -50 | 0.5 | -200 | 0.5 | 200 | 0.8 | -200 | 0.2 | -250 | -245 |
| 26 | -50 | 0.8 | -200 | 0.2 | 200 | 0.8 | -50 | 0.2 | -400 | -230 |
| 27 | -50 | 0.2 | -200 | 0.8 | 50 | 0.8 | -200 | 0.2 | -100 | -170 |
| 28 | -50 | 0.8 | -200 | 0.2 | -50 | 0.5 | -200 | 0.5 | 0 | 45 |
| 29 | 50 | 0.5 | -50 | 0.5 | 50 | 0.2 | -200 | 0.8 | 150 | 150 |
| 30 | 200 | 0.8 | -200 | 0.2 | -50 | 0.2 | -200 | 0.8 | 250 | 290 |
| 31 | 200 | 0.5 | -50 | 0.5 | -50 | 0.5 | -200 | 0.5 | 400 | 200 |
| 32 | -50 | 0.8 | -200 | 0.2 | 200 | 0.2 | -50 | 0.8 | -400 | -80 |
| 33 | 200 | 0.2 | -50 | 0.8 | 200 | 0.8 | 50 | 0.2 | -100 | -170 |
| 34 | 200 | 0.8 | -50 | 0.2 | 50 | 0.8 | -50 | 0.2 | 150 | 120 |
| 35 | 50 | 0.2 | -50 | 0.8 | 200 | 0.5 | 50 | 0.5 | -250 | -155 |
| 36 | 200 | 0.5 | 50 | 0.5 | 50 | 0.2 | -50 | 0.8 | 250 | 155 |
| 37 | 200 | 0.5 | 50 | 0.5 | 200 | 0.8 | -50 | 0.2 | 100 | -25 |
| 38 | 50 | 0.8 | -50 | 0.2 | 200 | 0.8 | -50 | 0.2 | -150 | -120 |
| 39 | 50 | 0.5 | -50 | 0.5 | 50 | 0.5 | -200 | 0.5 | 150 | 75 |
| 40 | 200 | 0.8 | 50 | 0.2 | 200 | 0.5 | -50 | 0.5 | 100 | 95 |
| 41 | 200 | 0.5 | -50 | 0.5 | 50 | 0.8 | -200 | 0.2 | 300 | 75 |
| 42 | 50 | 0.5 | -200 | 0.5 | 200 | 0.5 | -200 | 0.5 | -150 | -75 |
| 43 | 50 | 0.8 | -200 | 0.2 | 200 | 0.8 | -200 | 0.2 | -150 | -120 |
| 44 | 200 | 0.5 | -200 | 0.5 | -50 | 0.2 | -200 | 0.8 | 250 | 170 |
| 45 | 200 | 0.2 | -200 | 0.8 | -50 | 0.8 | -200 | 0.2 | 250 | -40 |
| 46 | 200 | 0.8 | -200 | 0.2 | 200 | 0.5 | -50 | 0.5 | -150 | 45 |
| 47 | 200 | 0.5 | -50 | 0.5 | 200 | 0.8 | -200 | 0.2 | 150 | -45 |
| 48 | 200 | 0.5 | -200 | 0.5 | 200 | 0.2 | -200 | 0.8 | 0 | 120 |

Table 29: False Regret and True Regret for all 48 problems used in Coricelli et al. [19]

B.1 INVITATION EMAIL

Hello #fname# #lname#!

We would like to invite you to participate in a new DR@W Laboratory research project called "Decision Sessions". For participating in an hour long session, you will be able to guarantee yourself a show-up fee of £5, with the chance to win further amounts of money depending on your decisions in the session.

The research sessions will take place in the DR@W Lab, S2.82 in Social Sciences, and are scheduled for the following times:

#sessionlist#

If you want to participate, you can register by clicking on the following link:

#link#

(If you cannot click on the link, copy it to the clipboard by selecting it, right click, and choosing "Copy", and then paste it into the address line in your web browser by right clicking there and choosing "Paste".)

Once you have registered for a session, further information will be distributed on Monday 28th November / *Monday 1st December*.

Many Thanks,
DR@W Research Team

B.2 SUMMARY STATISTICS

| | 29 NOV | 30 NOV | 2 DEC | COMBINED |
|------------------------|--------|--------|-------|----------|
| Number of Participants | 122 | 96 | 126 | 344 |
| Male (%) | 46.72 | 47.92 | 41.27 | 45.06 |
| Female (%) | 53.28 | 52.08 | 58.73 | 54.94 |
| Average age | 20.65 | 20.36 | 21.84 | 21.01 |
| 1st Year Undergrad (%) | 22.13 | 30.21 | 25.40 | 25.58 |
| 2nd Year Undergrad (%) | 28.69 | 31.25 | 23.02 | 27.33 |
| 3rd Year Undergrad (%) | 31.15 | 27.08 | 22.22 | 26.74 |
| 4th Year Undergrad (%) | 3.28 | 1.04 | 4.76 | 3.20 |
| Graduate (%) | 13.93 | 9.38 | 17.46 | 13.95 |
| Other (%) | 0.82 | 1.04 | 7.14 | 3.20 |
| A-Level Maths? (%) | 81.97 | 85.42 | 78.57 | 81.69 |
| Economics Dept (%) | 31.97 | 33.33 | 19.05 | 27.62 |
| Business School (%) | 16.39 | 11.46 | 19.84 | 16.28 |
| Statistics Dept (%) | 11.48 | 8.33 | 8.73 | 9.59 |
| Maths Dept (%) | 8.20 | 8.33 | 4.76 | 6.98 |
| Engineering Dept (%) | 4.92 | 3.12 | 11.90 | 6.98 |
| Psychology Dept (%) | 5.74 | 3.13 | 7.94 | 5.81 |
| Other (%) ¹ | 21.31 | 32.29 | 27.78 | 26.74 |

Table 30: Summary statistics for all participants in experiment

B.3 EXPERIMENTAL METHODOLOGY

B.3.1 *Participants*

Participants were invited to the experiment through an email sent to the University of Warwick's DR@W (Decision Research at Warwick) participant database, which contained over 1500 active profiles, comprised primarily of Undergraduate students at the University of Warwick (well distributed across year and course of study). Sign-up was completed on a first come, first served basis, through self-allocation to one of 21 available sessions, with the first round of sessions being on the 29th and 30th November 2011 and the second round of sessions being on the 2nd December 2014.

B.3.2 *Materials*

The experiment was designed to run entirely in a web browser, with a hard copy of the Information Sheet and Consent Form provided to each participant. The experiment was built in the University of Warwick Sitebuilder environment, which allows for simple data capture and

¹ no other department had more than 6% of the sample

navigation, though not more sophisticated mechanisms (such as subject specific randomisation). In addition to the main experiment (design detailed below) subjects also completed an incentivised Holt and Laury procedure to measure risk attitudes (prior to the main experiment), answered a DOSPERT (Domain-Specific Risk-Taking) questionnaire, completed a Big 5 Personality profile, and answered basic questions relating to their individual personal and academic characteristics (all after the main experiment). These were all hosted within the Sitebuilder environment. Gambles in the Holt and Laury procedure were resolved using a random number generator in excel (subject specific), though the outcomes of the gambles were not disclosed to the participants. The result was simply added to their total payment. The procedure was run prior to the main experiment so that the measure of risk attitude would be unaffected by the main experiment. The results were not disclosed so there would be no emotional carry-over from the Holt and Laury procedure into the main experiment. Gambles in the main experiment were resolved through rolling either 1 or 2 ten-sided die. Due to the limitations of the Sitebuilder environment, gambles were only session specific, not subject specific. For each gamble within each session, the dice were rolled in advance of the session, with the rolling of the dice videotaped, with audio, and then played back to the subjects once they had made their decision. This was to ensure the complete privacy of the participants throughout the experimental session.

B.3.3 *Design*

The project aimed to investigate the effect of experienced regret on subsequent anticipatory regret and hence subsequent choice behaviour. Specifically, the experiment looked to investigate whether the emotion of experienced regret is associated to the “action” as distinct from the “decision” taken, and hence has an impact on subsequent choice when the previous decision/action is disaggregated into component parts in a similar situation. This was investigated through a simple two option, two stage gambling task, where participants must choose between a relatively safe and relatively risky option in each stage. The key design occurs in stage 2 of the experiment. If a participant was allocated to the “treatment” condition, then stage 2 would contain an action that they had previously seen in stage 1 (i.e. if they picked “S”, it would imply they had a $x\%$ chance of winning $\pounds y$, resolved by rolling 1 die, and this was something they had previously seen in stage 1). If they were in the “control” condition, they would be faced with the same “decision” as those in the treatment condition (in terms of probabilities and payoffs), but the “actions” would be labelled differently (i.e. they would need to select “Sc”, which is resolved by rolling 2 die, to make the same decision that choosing action “S” represented in the treatment condition). The primary outcome is the difference in frequency with which “S” and “Sc” were chosen in the treatment and control groups, in those instances where a participant had chosen S in stage 1, and S had turned out to be the wrong choice. Participants also answered a questionnaire which obtained measures of individual characteristics. The time taken to complete the experiment from start to finish was around 45 minutes. The stages of the entire experiment were run in the following order:

1. Holt and Laury procedure

2. Stage 1 choice
3. Stage 2 choice
4. Participant information questionnaire
5. DOSPERT procedure
6. Big 5 procedure

Based on the design of the experiment, and choice behaviour observed in similar experiment, the mean payment was expected to be £11.40. The observed mean payment was £12.45, decomposed as a mean payment of £6.15 for the Holt and Laury task (5p more than expected), and £6.30 for the main experimental tasks (£1 more than expected).

B.3.4 *Environment*

All experimental sessions were conducted in the University of Warwick Economics Department DR@W Laboratory. This room contains 20 individual computer booths, shielded on 3 sides to ensure privacy when decision making. Individuals were allocated to booths through the drawing of a random ID card upon entry (as described in the Procedure below). All participants were asked to wear headphones (provided) during the experiment as the videos contained sound.

B.3.5 *Pre Experiment*

A blank Information Sheet, Consent form, blank sheet of paper and pen were set out on each desk. The computer on each desk was loaded up with the experiment log in page in the web browser. In the code of the webpages the participants would visit, the videos were updated to reflect the rolls of the die corresponding to that session.

B.3.6 *During Experiment*

Each participant took a random ID card upon entry and was asked to sit at the corresponding computer. The following instructions were read aloud:

Please do not touch the computer in front of you until instructed.

Please make sure all mobile phones are switched off so we are not disturbed during today's session.

Please read the information sheet on the desk in front of you, and, if you agree to participate in the session, please sign the consent form. When you have completed the consent form, please hold it in the air, and I will come around and collect it. *Collect consent forms* Thank you for agreeing to participate in the session.

During today's session, you will be asked to make a small number of decisions which will, along with an element of luck, determine your payoff for today's session. The decisions you submit on the computers will be entirely anonymous and private, and your final payoff will be calculated by our central computer.

The instructions as to how to proceed with the session will be given on the computer screen in front of you, but, if you have any questions about an aspect of the session, please raise your hand, and I will come to assist you privately.

During the session, you will see some videos which have audio on them, so you should wear the headphones provided at all times to ensure that you can hear the video without disturbing other individuals in the room. In front of you, you have a pen and paper, which you may use at any time to help you with, or make a record of, your decisions in the session. We will destroy the paper after the session, so any notes you make remain private.

Once you are ready to begin the session, please move the mouse in front of you so that the log in screen is visible. You may then login using the details on the ID card, and follow the instructions on the screen to progress through the session, without any further instruction from me.

Thanks

B.3.7 *Post Experiment*

The decisions taken in the session were downloaded from the website, and converted into payments for the session based on the rolls of the die specific to that session. The payment was placed inside an envelope, labelled with the computer ID. The following instructions were then announced:

We will now administer payment for the session

Please remember, your payment has been calculated by our central computer based on the decisions you made in today's session, and the outcomes of any gambles which you chose, so we are unable to answer any specific payment related questions at this time.

I will bring round a box of envelopes which have your session ID numbers on the front. Please take the envelope which corresponds to your ID number and a payment receipt

Please then open the envelope, and complete the receipt for the amount in the envelope.

Please then leave the pen, paper, ID card & headphones on your desk, and hand the receipt form to me on the way out.

Please remember to take all your personal belongings with you.

Envelopes and receipt forms were then handed out to each participant.

Receipt forms were collected in by the door on the way out.

B.3.8 *Information Sheet*

The information sheet given to participants was as follows:

*Information for participants - Decision Sessions
The project*

The project is being run by the Economics department of the University of Warwick. We plan to carry out a run of research sessions with a number of students at the University of Warwick. The sessions take place in the DR@W Laboratory, S2.82, on the second floor of the Social Sciences building between the 29th and 30th November 2011 / *on the 2nd December 2014*.

Your participation

If you agree to participate in the project, you will participate in a single session of the project, on either the 29th or 30th November / *on the 2nd December 2014*. You will be asked to make a small number of decisions during the session, which will be fully explained and presented to you on the computer screen. The decisions you make, coupled with an element of luck, will determine your payoff for the session. In accordance with the DR@W rules and regulations for lab S2.82, all information we provide is accurate and truthful, with no deception being used at any stage.

You will also be asked to complete a short questionnaire during the session, which will ask some simple questions about your own personal attributes. The answers to the questionnaire will not be linked to your payment.

The entire process will be both anonymous and private - anonymous in the sense that your decisions and answers will be linked to an ID number, which will be randomly determined by drawing an ID card on the way into the lab, and will not be linked to your real name or university ID, so there will be no way of linking your true identity to any of the data we have collected. It will be private in the sense that you will be asked to submit your decisions and answers in privacy, with neither the researchers nor the other participants able to see your answers, and your final fee will be calculated automatically by a central computer.

The anonymous data generated in the session will be used as the basis for at least one academic paper, and possibly more. Since the data is anonymous from the moment the study starts it will be impossible for anyone to link you to the data that is used.

You will be asked to sign a consent form if you agree to take part in the study and a receipt at the end of the session when you are paid. Both of these documents will be kept in a secure location for at least one year following the end of the study and then destroyed.

Participation in the project is entirely voluntary and you have the right to withdraw at any point without giving any reason. However, if you do end your participation early you may receive a reduced payment since the payment is related to your decisions.

Potential benefits

You will receive a show-up fee of £5 for participating in the session, which can either be guaranteed or traded away depending on your decisions in the session. You will always have the option to guarantee yourself the show-up fee for

the session. Your decisions, coupled with an element of luck, will determine your final payment for the session.

BIBLIOGRAPHY

- [1] Maurice Allais. Le comportement de l'Homme rationnel devant le risque: Critique des postulats et axiomes de l'Ecole americaine. *Econometrica*, 21(4):503–546, October 1953. (Cited on page 4.)
- [2] Roy F. Baumeister, Kathleen D. Vohs, C. Nathan DeWall, and Liqing Zhang. How emotion shapes behavior: Feedback, anticipation, and reflection, rather than direct causation. *Personality and Social Psychology Review*, 11(2):167–203, May 2007. (Cited on pages x, 13, 14, 19, 38, and 39.)
- [3] Gordon M. Becker, Morris H. Degroot, and Jacob Marschak. Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3):226–232, January 1964. (Cited on page 176.)
- [4] Denise R. Beike and Travis S. Crone. When experienced regret refuses to fade: Regrets of action and attempting to forget open life regrets. *Journal of Experimental Social Psychology*, 44(6):1545–1550, November 2008. (Cited on pages 25 and 26.)
- [5] David E. Bell. Regret in decision making under uncertainty. *Operations Research*, 30(5):961–981, October 1982. (Cited on pages 4, 6, 12, 14, 32, 73, 86, 103, 117, 119, and 136.)
- [6] David E. Bell. Risk premiums for decision regret. *Management Science*, 29(10):1156–1166, October 1983. (Cited on page 84.)
- [7] David E. Bell. Disappointment in decision making under uncertainty. *Operations Research*, 33(1):1–27, February 1985. (Cited on pages 5, 12, 74, and 109.)
- [8] Roland Bénabou and Jean Tirole. Self-Confidence and personal motivation. *Quarterly Journal of Economics*, 117(3):871–915, August 2002. (Cited on page 24.)
- [9] Jeremy Bentham. *An Introduction to the Principles of Morals and Legislation*. 1789. (Cited on pages 12 and 13.)
- [10] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, January 1954. (Cited on page 4.)
- [11] Peter Bohm, Johan Lindén, and Joakin Sonnegård. Eliciting reservation prices: Becker-DeGroot-Marschak mechanisms vs. markets. *The Economic Journal*, 107(443):1079–1089, July 1997. (Cited on pages 177 and 179.)
- [12] Robert Bringhurst. *The Elements of Typographic Style*. Version 2.5. Hartley & Marks, Publishers, Point Roberts, WA, USA, 2002. (Cited on page 205.)
- [13] Roger Brown and James Kulik. Flashbulb memories. *Cognition*, 5(1):73–99, March 1977. (Cited on page 24.)
- [14] Maarten L. Buis. Stata tip 87: Interpretation of interactions in nonlinear models. *Stata Journal*, 10(2):305–308, June 2010. (Cited on page 165.)

- [15] David J. Butler and Graham C. Loomes. Imprecision as an account of the preference reversal phenomenon. *The American Economic Review*, 97(1):277–297, March 2007. (Cited on page 31.)
- [16] Nathalie Camille, Giorgio Coricelli, Jerome Sallet, Pascale Pradat-Diehl, Jean-Rene Duhamel, and Angela Sirigu. The involvement of the orbitofrontal cortex in the experience of regret. *Science*, 304(5674):1167–1170, May 2004. (Cited on pages x, 9, 10, 22, 27, 122, and 124.)
- [17] Hannah Faye Chua, Richard Gonzalez, Stephan F. Taylor, Robert C. Welsh, and Israel Liberzon. Decision-related loss: Regret and disappointment. *NeuroImage*, 47(4):2031–2040, October 2009. (Cited on pages 122 and 124.)
- [18] Terry Connolly and Marcel Zeelenberg. Regret in decision making. *Current Directions in Psychological Science*, 11(6):212–216, December 2002. (Cited on page 79.)
- [19] Giorgio Coricelli, Hugo D. Critchley, Mateus Joffily, John P. O’Doherty, Angela Sirigu, and Raymond J. Dolan. Regret and its avoidance: a neuroimaging study of choice behavior. *Nature Neuroscience*, 8(9):1255–1262, September 2005. (Cited on pages x, xi, xiii, 9, 11, 12, 21, 22, 23, 27, 84, 85, 108, 109, 110, 111, 112, 122, 124, and 186.)
- [20] Giorgio Coricelli, Raymond J. Dolan, and Angela Sirigu. Brain, emotion and decision making: the paradigmatic example of regret. *Trends in Cognitive Science*, 11(6):258–265, June 2007. (Cited on page 21.)
- [21] Elizabeth H. Creyer and William T. Ross. The development and use of a regret experience measure to examine the effects of outcome feedback on regret and subsequent choice. *Marketing Letters*, 10(4):373–386, November 1999. (Cited on pages 12, 81, 82, 101, 102, 103, 105, 108, 117, and 128.)
- [22] Ton de Jong and Monica G.M. Ferguson-Hessler. Types and qualities of knowledge. *Educational Psychologist*, 31(2):105–113, 1996. (Cited on page 132.)
- [23] David M. Diamond, Adam M. Campbell, Collin R. Park, Joshua Halonen, and Phillip R. Zoladz. The temporal dynamics model of emotional memory processing: A synthesis on the neurobiological basis of Stress-Induced amnesia, flashbulb and traumatic memories, and the Yerkes-Dodson law. *Neural Plasticity*, 2007: 1–33, 2007. ISSN 2090-5904. (Cited on page 24.)
- [24] Emel Filiz-Ozbay and Erkut Y. Ozbay. Auctions with anticipated regret: Theory and experiment. *The American Economic Review*, 97(4):1407–1418, September 2007. (Cited on pages 93, 94, and 95.)
- [25] Barbara L. Fredrickson. Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition & Emotion*, 14(4):577–606, July 2000. (Cited on page 16.)
- [26] Daniel T. Gilbert, Carey K. Morewedge, Jane L. Risen, and Timothy D. Wilson. Looking forward to looking backward. *Psychological Science*, 15(5):346–350, May 2004. (Cited on pages 17, 18, 23, 103, 113, and 139.)

- [27] Itzhak Gilboa and David Schmeidler. Case-Based decision theory. *The Quarterly Journal of Economics*, 110(3):605–639, August 1995. (Cited on pages 37, 39, and 70.)
- [28] Thomas Gilovich and Victoria Husted Medvec. The experience of regret: what, when, and why. *Psychological Review*, 102(2):379–395, April 1995. (Cited on page 79.)
- [29] Peter J. Hammond. Changing tastes and coherent dynamic choice. *The Review of Economic Studies*, 43(1):159–173, February 1976. (Cited on page 35.)
- [30] David W. Harless and Colin F. Camerer. The predictive utility of generalized expected utility theories. *Econometrica*, 62(6):1251–1289, November 1994. (Cited on page 76.)
- [31] Sergiu Hart and Elchanan Ben-Porath. The optimality of regret matching. The Hebrew University of Jerusalem, Center for Rationality, July 2008. URL <http://www.ma.huji.ac.il/~hart/papers/opt-regret-p.pdf>. (Cited on pages 6 and 19.)
- [32] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, September 2000. (Cited on pages 6, 19, 119, 130, and 132.)
- [33] Sergiu Hart and Andreu Mas-Colell. A reinforcement procedure leading to correlated equilibrium. In Gérard Debreu, Wilhelm Neuefeind, and Walter Trockel, editors, *Economics Essays*, pages 181–200. Springer, 2001. (Cited on page 19.)
- [34] Takashi Hayashi. Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1):242–268, March 2008. (Cited on pages 7, 14, 31, 32, 33, 34, 35, 36, 41, 43, 44, 46, 49, 59, 60, 68, 69, 70, and 119.)
- [35] Takashi Hayashi. Context dependence and consistency in dynamic choice under uncertainty: the case of anticipated regret. *Theory and Decision*, 70(4):399–430, April 2011. (Cited on pages 7 and 130.)
- [36] John D. Hey and Chris Orme. Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 62(6):1291–1326, November 1994. (Cited on pages xi, 75, and 76.)
- [37] G. Robert J. Hockey, A. John Maule, Peter J. Clough, and Larissa Bdzola. Effects of negative mood states on risk in everyday decision making. *Cognition & Emotion*, 14(6):823–855, 2000. (Cited on page 11.)
- [38] Charles A. Holt and Susan K. Laury. Risk aversion and incentive effects. *The American Economic Review*, 92(5):1644–1655, December 2002. (Cited on pages 145 and 153.)
- [39] John K. Horowitz. The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods. *Economics Letters*, 93(1):6–11, October 2006. (Cited on pages 176 and 177.)

- [40] Hendrik S. Houthakker and Lester D. Taylor. *Consumer demand in the United States, 1929-1970 : Analyses and Projections*, volume CXXVI of *Harvard Economic Studies*. Harvard University Press, 1966. (Cited on pages 35 and 71.)
- [41] Steven J. Humphrey. Feedback-conditional regret theory and testing regret-aversion in risky choice. *Journal of Economic Psychology*, 25(6):839–857, December 2004. (Cited on page 79.)
- [42] J. Jeffrey Inman and Marcel Zeelenberg. Regret in repeat purchase versus switching decisions: The attenuating role of decision justifiability. *Journal of Consumer Research*, 29(1):116–128, 2002. (Cited on page 121.)
- [43] O. P. John, E. M. Donahue, and R. L. Kentle. The big five Inventory—Versions 4a and 54. *Institute of Personality and Social Research*, 1991. (Cited on page 157.)
- [44] O. P. John, L. P. Naumann, C. J. Soto, O. P. John, R. W. Robins, and L. A. Pervin. Paradigm shift to the integrative big five trait taxonomy: History, measurement, and conceptual issues. In *Handbook of personality: Theory and research*, pages 114–158. Guilford Press, New York, NY, 2008. (Cited on page 157.)
- [45] Robert A. Josephs, Richard P. Larrick, Claude M. Steele, and Richard E. Nisbett. Protecting the self from the negative consequences of risky decisions. *Journal of Personality and Social Psychology*, 62(1):26–37, January 1992. (Cited on page 79.)
- [46] Daniel Kahneman. Experienced utility and objective happiness: A Moment-Based approach. In Isabelle Brocas and Juan D. Carrillo, editors, *The Psychology of Economic Decisions: Rationality and Well-Being*, volume 1, pages 187–208. Oxford University Press, 2003. (Cited on page 16.)
- [47] Daniel Kahneman and Dale T. Miller. Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93(2):136–153, April 1986. (Cited on page 125.)
- [48] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, March 1979. (Cited on pages 4, 5, and 73.)
- [49] Daniel Kahneman and Amos Tversky. The simulation heuristic. In *Judgment under Uncertainty: Heuristics and Biases*, pages 201–208. Cambridge University Press, April 1982. (Cited on page 119.)
- [50] Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American Psychologist*, 39(4):341–350, April 1984. (Cited on page 124.)
- [51] Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98(6):1325–1348, December 1990. (Cited on page 177.)
- [52] Daniel Kahneman, Barbara L. Fredrickson, Charles A. Schreiber, and Donald A. Redelmeier. When more pain is preferred to less: Adding a better end. *Psychological Science*, 4(6):401–405, November 1993. (Cited on page 15.)

- [53] Daniel Kahneman, Peter P. Wakker, and Rakesh Sarin. Back to bentham? explorations of experienced utility. *Quarterly Journal of Economics*, 112(2):375–405, 1997. (Cited on pages [12](#), [13](#), and [14](#).)
- [54] Edi Karni and Zvi Safra. "Preference reversal" and the observability of preferences by experimental methods. *Econometrica*, 55(3): 675–685, May 1987. (Cited on page [177](#).)
- [55] Gideon Keren and Charles Lewis. The two fallacies of gamblers: type I and type II. *Organizational Behavior and Human Decision Processes*, 60(1):75–89, October 1994. (Cited on page [150](#).)
- [56] Sarah Lichtenstein and Paul Slovic. Response-induced reversals of preference in gambling: An extended replication in las vegas. *Journal of Experimental Psychology*, 101(1):16–20, November 1973. (Cited on page [45](#).)
- [57] John A. List, Sally Sadoff, and Mathis Wagner. So you want to run an experiment, now what? some simple rules of thumb for optimal experimental design. *Experimental Economics*, 14(4):439–457, November 2011. (Cited on page [157](#).)
- [58] Graham Loomes and Robert Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368):805–824, December 1982. (Cited on pages [4](#), [5](#), [6](#), [7](#), [12](#), [14](#), [32](#), [33](#), [35](#), [36](#), [73](#), [74](#), [76](#), [78](#), [86](#), [87](#), [88](#), [96](#), [103](#), [117](#), [136](#), and [177](#).)
- [59] Graham Loomes and Robert Sugden. Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies*, 53(2):271–282, April 1986. (Cited on pages [5](#), [12](#), [74](#), and [109](#).)
- [60] Graham Loomes and Robert Sugden. Some implications of a more general form of regret theory. *Journal of Economic Theory*, 41(2):270–287, April 1987. (Cited on pages [5](#), [87](#), [88](#), and [119](#).)
- [61] Graham Loomes and Caron Taylor. Non-Transitive preferences over gains and losses. *The Economic Journal*, 102(411):357–365, March 1992. (Cited on page [74](#).)
- [62] Graham Loomes, Chris Starmer, and Robert Sugden. Observing violations of transitivity by experimental methods. *Econometrica*, 59(2):425–439, March 1991. (Cited on page [74](#).)
- [63] Barbara Mellers, Alan Schwartz, and Ilana Ritov. Emotion-Based choice. *Journal of Experimental Psychology: General*, 128(3):332–345, September 1999. (Cited on pages [x](#), [7](#), [8](#), and [9](#).)
- [64] Sendhil Mullainathan. A memory-based model of bounded rationality. *Quarterly Journal of Economics*, 117(3):735–774, August 2002. (Cited on page [24](#).)
- [65] Antoinette Nicolle, Dominik R. Bach, Chris Frith, and Raymond J. Dolan. Amygdala involvement in self-blame regret. *Social Neuroscience*, 6(2):178–189, 2011. (Cited on page [124](#).)
- [66] Antoinette Nicolle, Stephen M. Fleming, Dominik R. Bach, Jon Driver, and Raymond J. Dolan. A regret-induced status quo bias. *The Journal of Neuroscience*, 31(9):3320–3327, March 2011. (Cited on page [122](#).)

- [67] Lisa D. Ordóñez and Terry Connolly. Regret and responsibility: A reply to Zeelenberg et al. (1998). *Organizational Behavior and Human Decision Processes*, 81(1):132–142, January 2000. (Cited on page 123.)
- [68] Rik Pieters and Marcel Zeelenberg. A theory of regret regulation 1.1. *Journal of Consumer Psychology*, 17(1):29–35, January 2007. (Cited on page 130.)
- [69] Robert A. Pollak. Endogenous tastes in demand and welfare analysis. *The American Economic Review*, 68(2):374–379, May 1978. (Cited on page 35.)
- [70] John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323–343, December 1982. (Cited on page 4.)
- [71] John Quiggin. *Generalized expected utility theory: the rank-dependent model*. Kluwer Academic Publishers, Norwell, Massachusetts, 1993. (Cited on page 76.)
- [72] Matthew Rabin and Dimitri Vayanos. The gambler’s and hot-hand fallacies: theory and applications. *Review of Economic Studies*, 77(2):730–778, April 2010. (Cited on page 150.)
- [73] Daniela Raeva and Eric van Dijk. Regret once, think twice: the impact of experienced regret on risk choice. Technical report, Computable and Experimental Economics Laboratory, Department of Economics, University of Trento, Italia, 2009. URL <http://ideas.repec.org/p/trn/utwpce/0903.html>. (Cited on pages xi, 82, 83, 84, 103, 104, 105, 106, 107, 117, and 129.)
- [74] Donald A. Redelmeier and Daniel Kahneman. Patients’ memories of painful medical treatments: real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, 66(1):3–8, July 1996. (Cited on page 15.)
- [75] René Richard, Joop van der Pligt, and Nanne K. de Vries. Anticipated regret and time perspective: changing sexual risk-taking behavior. *Journal of Behavioral Decision Making*, 9(3):185–199, 1996. (Cited on page 79.)
- [76] Ilana Ritov. Probability of regret: Anticipation of uncertainty resolution in choice. *Organizational Behavior and Human Decision Processes*, 66(2):228–236, May 1996. (Cited on pages 105 and 106.)
- [77] Ilana Ritov and Jonathan Baron. Status-quo and omission biases. *Journal of Risk and Uncertainty*, 5(1), February 1992. ISSN 0895-5646, 1573-0476. doi: 10.1007/BF00208786. URL <http://www.springerlink.com/content/qx07u05234p6411u/>. (Cited on page 125.)
- [78] Ilana Ritov and Jonathan Baron. Outcome knowledge, regret, and omission bias. *Organizational Behavior and Human Decision Processes*, 64(2):119–127, November 1995. (Cited on pages 78 and 120.)
- [79] A. E. Roth. Form and function in experimental design. *Behavioral and Brain Sciences*, 24(3):427–428, 2001. (Cited on page 160.)

- [80] Alvin E. Roth and Ido Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8(1):164–212, 1995. ISSN 0899-8256. doi: 10.1016/S0899-8256(05)80020-X. URL <http://www.sciencedirect.com/science/article/B6WFW-4G9TCW6-9/2/e83292099dd0e1dbeed2b502b55b5a42>. (Cited on page 20.)
- [81] Paul A. Samuelson. A note on the pure theory of consumer’s behaviour. *Economica*, 5(17):61–71, February 1938. (Cited on page 12.)
- [82] William Samuelson and Richard Zeckhauser. Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1(1):7–59, March 1988. (Cited on pages 120, 121, and 122.)
- [83] Todd Sarver. Anticipating regret: Why fewer options may be better. *Econometrica*, 76(2):263–305, March 2008. (Cited on pages 6, 7, and 69.)
- [84] Leonard J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, March 1951. (Cited on pages 7, 14, and 32.)
- [85] Itamar Simonson. The influence of anticipating regret and responsibility on purchase decisions. *The Journal of Consumer Research*, 19(1):105–118, June 1992. (Cited on pages 78 and 79.)
- [86] Mark Spranca, Elisa Minsk, and Jonathan Baron. Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27(1):76–105, January 1991. (Cited on page 119.)
- [87] Dale O. Stahl. Rule learning in symmetric normal-form games: Theory and evidence. *Games and Economic Behavior*, 32(1):105–138, July 2000. (Cited on page 20.)
- [88] Chris Starmer. Developments in non-Expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, 38(2):332–382, June 2000. (Cited on page 4.)
- [89] Robert Sugden. An axiomatic foundation for regret theory. *Journal of Economic Theory*, 60(1):159–180, June 1993. (Cited on pages 5 and 6.)
- [90] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: an introduction*. MIT Press, Cambridge, Massachusetts, 1998. (Cited on page 20.)
- [91] Richard Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, 1(1):39–60, 1980. (Cited on page 124.)
- [92] Edward Thorndike. *Animal Intelligence*. The animal behavior series. The Macmillan Company, New York, 1911. (Cited on page 20.)
- [93] Amos Tversky and Daniel Kahneman. Belief in the law of small numbers. *Psychological Bulletin*, 76(2):105–110, 1971. (Cited on page 107.)

- [94] Bob Uttl, Nobuo Ohta, and Amy L. Siegenthaler. *Memory and Emotion: Interdisciplinary Perspectives*. Blackwell Publishing, 1st edition, 2006. (Cited on pages [x](#), [24](#), and [25](#).)
- [95] Frans van Winden, Michal Krawczyk, and Astrid Hopfensitz. Investment, resolution of risk, and the role of affect. *Journal of Economic Psychology*, 32(6):918–939, December 2011. (Cited on page [8](#).)
- [96] Carol Varey and Daniel Kahneman. Experiences extended across time: Evaluation of moments and episodes. *Journal of Behavioral Decision Making*, 5(3):169–185, July 1992. (Cited on page [15](#).)
- [97] Ivo Vlaev, Ben Seymour, Raymond J. Dolan, and Nick Chater. The price of pain and the value of suffering. *Psychological Science*, 20(3):309–317, March 2009. (Cited on pages [xii](#), [177](#), and [178](#).)
- [98] John von Neumann and Oscar Morgenstern. *Theory of Games and Economic Behavior*. 1st edition, 1944. (Cited on pages [4](#) and [5](#).)
- [99] W. Richard Walker, Rodney J. Vogl, and Charles P. Thompson. Autobiographical memory: unpleasantness fades faster than pleasantness over time. *Applied Cognitive Psychology*, 11(5):399–413, October 1997. (Cited on page [25](#).)
- [100] Timothy D. Wilson and Daniel T. Gilbert. Affective forecasting. *Advances in Experimental Social Psychology*, 35:345–411, 2003. (Cited on pages [15](#), [16](#), and [27](#).)
- [101] Timothy D. Wilson, Jay Meyers, and Daniel T. Gilbert. "How Happy Was I, Anyway?" A retrospective impact bias. *Social Cognition*, 21(6):421–446, December 2003. (Cited on pages [x](#), [26](#), and [27](#).)
- [102] Derrick Wirtz, Justin Kruger, Christie Napa Scollon, and Ed Diener. What to do on spring break? The role of predicted, on-Line, and remembered experience in future choice. *Psychological Science*, 14(5):520–524, September 2003. (Cited on pages [x](#), [16](#), [27](#), [28](#), and [29](#).)
- [103] Carsten Wrosch, Isabelle Bauer, and Michael F. Scheier. Regret and quality of life across the adult life span: The influence of disengagement and available future goals. *Psychology and Aging*, 20(4):657–670, December 2005. (Cited on page [25](#).)
- [104] Marcel Zeelenberg. The use of crying over spilled milk: a note on the rationality and functionality of regret. *Philosophical Psychology*, 12(3):325–340, September 1999. (Cited on page [3](#).)
- [105] Marcel Zeelenberg and Jane Beattie. Consequences of regret aversion 2: Additional evidence for effects of feedback on decision making. *Organizational Behavior and Human Decision Processes*, 72(1):63–78, October 1997. (Cited on pages [12](#), [79](#), [80](#), [81](#), [82](#), [94](#), [95](#), [96](#), [97](#), [98](#), [100](#), [101](#), [102](#), [103](#), [105](#), [108](#), [117](#), and [128](#).)
- [106] Marcel Zeelenberg and Rik Pieters. A theory of regret regulation 1.0. *Journal of Consumer Psychology*, 17(1):3–18, January 2007. (Cited on pages [119](#), [130](#), and [131](#).)

- [107] Marcel Zeelenberg, Jane Beattie, Joop van der Pligt, and Nanne K. de Vries. Consequences of regret aversion: Effects of expected feedback on risky decision making, February 1996. (Cited on page 79.)
- [108] Marcel Zeelenberg, Wilco W. van Dijk, and Antony S. R. Manstead. Reconsidering the relation between regret and responsibility. *Organizational Behavior and Human Decision Processes*, 74(3):254–272, June 1998. (Cited on page 123.)
- [109] Marcel Zeelenberg, Wilco W. van Dijk, and Antony S. R. Manstead. Regret and responsibility resolved? evaluating Ordóñez and Connolly's (2000) conclusions. *Organizational Behavior and Human Decision Processes*, 81(1):143–154, January 2000. (Cited on page 78.)
- [110] Marcel Zeelenberg, Kees van den Bos, Eric van Dijk, and Rik Pieters. The inaction effect in the psychology of regret. *Journal of Personality and Social Psychology*, 82(3):314–327, March 2002. (Cited on page 120.)
- [111] Daniel Zizzo. *Transfer of Knowledge in Economic Decision Making*. Palgrave Macmillan, Houndmills, Basingstoke, Hampshire, RG21 6XS, 2005. (Cited on pages 131 and 132.)
- [112] Daniel Zizzo. Experimenter demand effects in economic experiments. *Experimental Economics*, 13(1):175–98, March 2010. (Cited on page 136.)

COLOPHON

This thesis was typeset with $\text{\LaTeX} 2_{\epsilon}$ using Hermann Zapf's *Palatino* and *Euler* type faces (Type 1 PostScript fonts *URW Palladio L* and *FPL* were used). The listings are typeset in *Bera Mono*, originally developed by Bitstream, Inc. as "Bitstream Vera". (Type 1 PostScript fonts were made available by Malte Rosenau and Ulrich Dirr.)

The typographic style was inspired by [Bringhurst's](#) genius as presented in *The Elements of Typographic Style* [12]. It is available for \LaTeX via CTAN as "[classicthesis](#)".

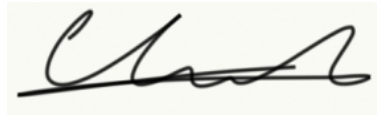
Final Version as of December 6, 2014 at 15:49.

DECLARATION

I declare that:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hard-bound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

Coventry, United Kingdom, December 2014

A handwritten signature in black ink, appearing to read 'Stephen Lovelady', is written over a light yellow rectangular background.

Stephen Lovelady