LANCASTER UNIVERSITY

DOCTORAL THESIS

Navigating Risk: Structural Modelling and Experimental Insights into Economic Decision-Making Under Uncertainty.

Author:
Nathan Nabil

Supervisors:
Konstantinos Georgalos
David Peel

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy

June 2024

Declaration of Authorship

I, Nathan Nabil, declare that this thesis titled, "Navigating Risk: Structural Modelling and Experimental Insights into Economic Decision-Making Under Uncertainty." and the work presented in it are my own. I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this or any other university.

Acknowledgements

I would like to express my deepest gratitude to my Supervisors, Konstantinos Georgalos, and David Peel, for their ongoing support, their confidence in me, and their constant willingness to share their expertise and experience with me.

I am grateful to everyone in the Lancaster University Economics department, for creating a welcoming, open environment for idea sharing and joint problem solving. A special mention to Giorgio Motta, Olivier Cardi, and Dakshina De Silva for always being there to answer my questions, provide additional support when required, and giving me the confidence to continuously progress throughout the PhD.

A big thank you to my Mother, Sister, and Partner, for putting up with my constant rambling on about economic decision-making under risk, and for always being there for me when things got tough - I do not think I could have done this without your support and your confidence in me. And thank you to my friends for sticking by me throughout the process.

Finally, an additional, and incredibly grateful, mention to my supervisor Konstantinos Georgalos. Thank you for putting up with my never ending questions, for giving me confidence when I doubted my abilities, and for everything you have taught me over the years. Konstantinos has been the most influential person in shaping the economist that I have become, from listening to me screaming at my R console when I was adamant I would never learn how to program, to providing structural, informative, and intuitive explanations about concepts I once thought to be inconceivable. I am very proud of my skill-set now, in my ability to program on various software's, in my ability to articulate the intuition behind complex econometric models, and in my ability to interpret results and generate innovative insights; I can safely say that this is all because of Konstantinos. Thank you for everything, and I hope we can work together on a project again in the future.

Contents

List of Figures

List of Tables

List of Abbreviations

AIC Akaike Information Criterion

BHM Bayesian Hierarchical Modelling

C Commitment (Chapter 2)

C Control (Chapter 1)

CE Certainty Equivalent

CL Cognitive Load

CPT Cumulative Prospect Theory

DI Dynamic Inconsistency

EUT Expected Utility Theory

HG Habitual Gambler

LL Log-Likelihood

LOOCV Leave One Out Cross Validation

LRT Likelihood Ratio Test

MLE Maximum Likelihood Estimation

MSE Mean Squared Error

NC No Commitment

NCL Not Cognitive Load

PG Problem Gambler

SOGS South Oaks Gambling Screen

Introduction

This thesis comprises four chapters, each focusing on different aspects of economic decision-making under risk. It explores various dimensions of risky decision-making, each contributing uniquely to the various choices we make and the outcomes we experience in everyday life. We investigate structural patterns in the financial decisions of individuals with impulse control disorders. We analyse how discrepancies between ex-ante risky financial strategies and ex-post actions can be modeled. Additionally, we examine how subjects alter their decision-making strategies under time pressure or when faced with increasingly complex tasks, and identify the structural models that best explain these changes.

In the first chapter we investigate various functional specifications of well established decision-theory models, namely Cumulative Prospect Theory (CPT henceforth). Using an existing data set, we seek to determine which specification is able to best explain the decision-making strategies of individuals with gambling disorders, and determine whether a parametric specification is able to identify structural patterns that are otherwise not as obvious with non-parametric methods. We find that a CPT model with a Power utility function and a Prelec two-parameter weighting function was the data generating process that had the best descriptive and predictive capacity. We highlight how parametric methods can help in disentangling probability distortion from probability elevation, which generated interesting insights into what characterises the choices of problem gamblers. We identify structural and statistically significant differences between the risk preferences of problem gamblers from non-gambling controls.[1]

The second chapter aims to test the empirical robustness of CPT in a dynamic setting. Using the theoretical predictions of the Casino Gambling Model, we introduce a three-stage experimental design that estimates risk preferences and classifies subjects into behavioural types that exist in the literature. We find that existing classifications are only able to explain the decisions made by a small proportion of the sample, and so we generate theoretical predictions for a new classification based on

[1] Chapter 1 is a solo-authored paper.

our results as well as those found in existing studies. We find that dynamic inconsistencies exist, whereby subjects deviate from their ex-ante strategies when making decisions in real-time. We explore the financial welfare implications of this behavior and discuss how different types and structures of commitment devices can help overcome these discrepancies.[2]

The third chapter tests CPT's explanatory power in tasks where there is increased time pressure or increased complexity. We propose that when subjects engage in tasks that overwhelm their cognitive loads, they switch from using sophisticated compensatory models (like CPT) to using simplification strategies such as heuristics. We define the theoretical assumptions of a cognitive toolbox model of heuristics and compare its explanatory power to that of CPT using Bayesian inference techniques. Our analysis is based on a meta-analysis of six existing data sets that differ in the domain in which they operate, as well as in their complexity. We find that in more traditional settings, CPT has more explanatory and predictive power; however, as the complexity of a task increases, subjects switch to using an adaptive toolbox of heuristics. We identify the size of toolboxes and the types of heuristics that make up each toolbox. We create a complexity index based on multiple characteristics of complexity, and find a positive correlation between complexity and toolbox usage.[3]

The final chapter is a short chapter which serves as a byproduct of chapter 3. This chapter expands on the complexity analysis by introducing various decision-theory models that have since been manipulated to capture complexity preferences. Using the same Bayesian inference techniques, we test four model variants and again find that an adaptive toolbox of heuristics is able to best explain choices when complexity increases.[4]

[2]Chapter 2 is a solo-authored paper, where Dr Konstantinos Georgalos provided the experimental interface on Python.

[3]Chapter 3 is a co-authored paper with Dr Konstantinos Georgalos.

[4]Chapter 4 is a co-authored paper with Dr Konstantinos Georgalos.

Chapter 1

Decision-Theory Modelling of Risk Preferences in Problem Gambling

Abstract

This chapter investigates the explanatory power of risk preferences in elucidating problem gambling behaviour. We utilise the Ring et al. (2018) choice data, which examines gambling decisions of three distinct treatment groups: Problem Gamblers (PG), Habitual Gamblers (HG), and a non-gambler control group (C). Employing advanced econometric techniques, we assess the efficacy of parametric models in capturing the behavioural choices of different groups. Specifically, we explore four model variants of Cumulative Prospect Theory (CPT) alongside a variant of the Markowitz utility model. Through rigorous model comparisons, utilising Maximum Likelihood Estimation (MLE) techniques, the Akaike Information Criterion (AIC), and the Leave-One-Out Cross validation routine (LOOCV), we identify the optimal data generating process - a CPT framework featuring a flexible power value function coupled with a two-parameter Prelec (1998) risky weighting function. Our findings highlight the advantages of parametric approaches while cautioning against the limitations imposed by linear utility and representative agent assumptions. Notably, we discern distinct patterns among PG, HG, and C groups, revealing that the PGs exhibit pronounced probability distortion in the gain domain, heightened sensitivity to loss payoffs, and a propensity to underweight probabilities in the loss domain.

1.1 Introduction

The evolution of decision-theory has allowed for advancements in the modelling of consumer behaviour. Attempts to mirror the choices of economic agents parametrically has provided deeper insights into how individuals make decisions in various economic domains. In the context of decision-making under risk, there has been a focus on parameterising models that assess the utility individuals receive from varying monetary amounts, how we perceive objective probabilities, and how individual preferences change when dealing with wins and losses. Tversky and Kahneman (1992)'s Cumulative Prospect Theory (CPT) combines these components and is recognised as the most prestigious model of human decision-making to date. The model has provided explanations for various economic paradox regarding investor and saver behaviour (Henderson, 2012), as well as explaining health decisions (Schwartz, Goldberg, and Hazen, 2008), and has provided insights into the relationship between risk preferences and perceptions of climate change (Osberghaus, 2017).

Multiple functional forms of CPT have been put forward to explain the behaviour of economic agents (Kahneman and Tversky 1979; Prelec 1998; Gonzalez and Wu 1999; Bouchouicha and Vieider 2017). Identifying the explanatory optimal functional form can provide further insights into the behaviours under investigation. Conducting empirical investigations that adopt a random functional form without informed insights may limit the robustness of results if one cannot confirm that the adopted functional form is the most probable data generating process (Stott 2006). Whereas applying informed, context-specific functional specifications increases the likelihood of accurately identifying preferences. We can robustly evaluate the effectiveness of different specifications by assessing which is able to give the best explanatory and predictive account of the empirical data. Not only does this allow one to identify the model that provides the best representation of the data, but will also foster insights into why competing models were not as effective, what implications this has for the individual decisions in that context, and how one can optimise the model further to fit consumer behaviour more accurately.

Behavioural economics has been built on the foundation of bounded rationality and individual heterogeneity, so when implementing a functional form, accounting for the individual and situational context is key. Whilst examinations on the effectiveness of functional forms of CPT and other decision-making theories have taken place (Blondel 2002; Stott 2006), these assessments are based on the more traditionally "standard" economic agents. The objective of this chapter is to build upon previous research and determine which functional forms of CPT offer the best explanation of the decision-making processes of individuals typically characterised as more impulsive and risk-seeking. A number of individuals in our economy deviate from classical economic theories excessively, and understanding the risk-preferences of those on the more extreme end of the "rationality spectrum" may provide interesting insights into their decision-making processes.

The subgroup this chapter seeks to examine are those with gambling addictions. Given the rise in gambling following the coronavirus pandemic, being able to better characterise the decisions of individuals in this domain may compliment the psychology and neuroscience literature in better understanding the disorder. Gambling disorders can be as distressing as life-threatening addictions (alcoholism and drug dependency), yet the adverse implications associated with the impulse-control disorder have been overlooked. From an economic perspective, problem gambling has resulted in the displacement of local residents, productivity losses, increased crime and higher unemployment rates. Similarly, from a biological welfare perspective, it has been shown to increase emotional pain, depression and anxiety (National Research Council, 1999). It is also associated with various negative externalities, where 7% of adults in Great Britain reported they were negatively affected by someone else's gambling problem (Davies, 2020). The majority of research attempting to better understand the problem lies in psychology, neuroscience and medicine, where these fields have sought to understand why individuals gamble and the pathology of the addiction. We wish to provide a new complementary approach, using advanced econometric modelling, to identify if there is a structural association between individual risk-preferences and being classified as a PG.

This chapter utilises the experimental data of Ring et al. (2018) to identify structural differences between the risk-preferences of PGs from non-gambling control groups. The data includes certainty equivalents from lottery choice tasks across three treatments groups: Problem Gamblers (individuals with a gambling addiction), Habitual Gamblers (those who gamble but not excessively), and non-gamblers (those who have never gambled). Using maximum-likelihood estimation approaches, as well as the cross-validation prediction routine, we fit four CPT specifications and a Markowitz (1952) model of utility to the experimental data to identify the most probable data generating process. The chapter highlights the benefits of parametric approaches, and identifies limitations associated with restrictive assumptions such as strict linear utility and representative-agent modelling.

We find that subjects in all three groups are best characterised with a power utility function and a Prelec (1998) two-parameter probability weighting function, where around 52% of subjects exhibited linear utility. Our results suggest that PGs differ from non-gambling controls in that they (1) distort probabilities more in the gain domain, (2) have an increased sensitivity to increases in loss amounts, and (3) underweight the whole probability spectrum in the loss domain.

The rest of the chapter will proceed as follows: Section 1.2 provides a comprehensive review of the literature on the psychology of gambling, and the success of decision-theory functional forms. Section 1.3 discusses the experimental data. Section 1.4 explains the preference functional under consideration. Section 1.5 provides our methodology. Section 1.6 presents the results, and Section 1.7 provides a concluding discussion.

1.2 Existing Literature

This section will provide a holistic review of the available literature on gambling addiction and decision-making under risk. Firstly, we will discuss evidence in the psychology and neuroscience field, providing a focus on why individuals gamble, and what may be driving the addiction. We will then discuss how the economics literature, namely in the domain of decision-making under risk, provides complimentary inputs to the research. Specifically, we will address the identification and characterisation of preferences, and how structural economic frameworks can help us better understand the decisions of PGs.

1.2.1 Psychology and Neuroscience

Much of the literature on problem gambling assumes that its impulsivity stems from neuropsychological impairments (Brand et al., 2005), with a focus on impairments in the prefrontal cortex domains. In the prefrontal cortex, there are 2 underlying mechanisms that could influence risk-taking behaviour: The orbitofrontal (ventromedial) prefrontal cortex (Brand et al. 2005; Cavedini et al. 2002; Labudda et al. 2007), and the dorsolateral prefrontal cortex (Bechara et al. 1994; Brand et al. 2005; Gelskov et al. 2016).

Deficits in the orbitofrontal prefrontal cortex functions are shown to leave individuals susceptible to uncontrolled behaviour, with hindered feedback processing mechanisms. This suggests that when presented with new information, PGs are inclined to neglect or manipulate this information. The information could be the outcome of a previous bet, or the objective probabilities of winning. Additionally, this hinders their ability to consider the future consequences associated with their actions. Cavedini et al. (2002) suggest that there is a link between PG and other disorders (e.g. OCD, drug addiction) in which they all have a diminished ability to evaluate future consequences, which can be explained by abnormalities in the orbitofrontal cortex function. The concept of feedback processing, and the inability to evaluate future consequences, are widely recognised in the psychology literature, and there are arguments that support these as being underlying causes of gambling. Clark and Dagher (2014) explain that impulsiveness, in a gambling context, can be characterized by an over-weighting of potential rewards relative to losses. As potential benefits are immediate, but many of the indirect costs (long-term gambling-related harm) are delayed, individuals with orbitofrontal prefrontal cortex impairments are likely to have disturbances in their feedback processes and thus are unable to process the potential costs that are associated with their decisions. In turn they prioritise the immediate potential rewards, which results in increased risk-taking (Shead, Callan, and Hodgins, 2008). Bechara et al. (1994) tested this idea empirically and found that individuals with prefrontal deficits had unstable representations of future outcomes, as they were unable to retain this aspect of information in their working memory long enough for reasoning strategy to be applied. There is an abundance of literature highlighting that PG is in part described by a high sensitivity for reward and a neglection of future consequences. (see Goudriaan et al. (2004) for a comprehensive review of the findings.)

Impairments in the dorsolateral prefrontal cortex are also able to explain, in part, the pathology of problem gambling. Changes in brain activity within this domain can lead to impairments of planning capacity, and the generation and use of unstable cognitive strategies that can also result in disadvantageous decision-making (Brand et al., 2005). PGs can be characterised by hypersensitivity to the most appetitive and risky bets, which comes from an executive cortico-stratal network including the dorsolateral prefrontal cortex (Gelskov et al., 2016). Again, the psychology literature has provided explanations as to how this works. With regards to deficits in one's ability to use cognitive strategies or to plan, the term idiosyncratic beliefs has been widely used to explain impulsive behaviour. Idiosyncratic beliefs are illogical beliefs that your subjective understanding of something is more accurate than proven and objective facts. Delfabbro, Lahn, and Grabosky (2006) tested this empirically on adolescent problem gamblers and found that idiosyncratic beliefs came to over-ride more objective considerations in the PG group. The PG group rated themselves as more skilful than others when gambling on activities where no such skill was possible. They believed that certain numbers on a six-sided die were easier to obtain, thus attaining more optimistic views on their chances of winning. Their results coincide with that of Limbrick-Oldfield et al. (2020) who found that PGs used prior feedback from gambles to inform their next choice, even though the tasks were completely independent. Similarly, Lopez-Gonzalez, Griffiths, and Estévez (2020) investigated PGs and how they perceived the role of knowledge in sports betting markets. Subjects first undertook cognitive behavioural therapy and then were tested for cognitive distortions. Their results showed severe distortions in their ability to distinguish between luck and skill. They were convinced that the betting products were not designed to giver bookmakers an advantage, and that the success of bookmakers came from their access to better quality sports information. Some participants were adamant that their addiction obstructed their rational way of thinking and their ability to use their skillset to predict the game. They felt that when they first saw the gamble they were in a rational state, but the addiction took over and they lost control over the betting. In essence, they believed that if they did not have an addiction, they could have used their knowledge and skills to increase their winnings. Clark (2010) identified two beliefs that were prominent in PGs: The near miss effect, and the effect of personal control. The near-miss effect allows an individual to believe that if they were "close to winning" on one gamble, then in the next they are sure to win. This leads to inevitable "loss-chasing" behaviour. Breen and Zuckerman (1999) show that loss-chasing is identified as a central characteristic of behaviour in PGs and leads to more frequent involvement, increased persistence and elevated monetary risk in an effort to salvage their losses. On the other hand, the effect of personal control allows an individual to believe that their strategic actions and their knowledge can influence the decision of a gamble, even though these beliefs have no objective influence on the likelihood of winning. Clark (2010) emphasises that these cognitive distortions lead to the irrational creation of skill-oriented behaviours which promote excessive gambling.

The psychology and neuroscience literature have provided evidence into the "why", whereas the economics literature, particularly in the field of decision-making under risk, may offer insights into the "how". More specifically, it looks to identify if there are preference-based, structural explanations for excessive gambling.

1.2.2 Economics

This subsection will provide a review of the literature on decision-theory, decision-making under risk, existing attempts to model consumer behaviour parametrically, and the limited literature on decision-theoretical approaches to explaining PGs risk preferences.

Expected Utility Theory (EUT) (Bernoulli, 1738) laid the foundation for understanding decision-making under risk and uncertainty. EUT posited that an individual accepts risk not only by the evaluation of gains and losses, but by the utility gained from the risky action itself. A key assumption being that marginal utility diminishes as outcomes increase. Von Neumann and Morgenstern (1947) provided a modification of EUT which suggested that under 4 axioms of rational behaviour, an individual who is faced with a risky decision with some probabilistic outcome will act as an expected utility maximiser. Despite EUT's prominence, its descriptive validity has been questioned, with economic paradoxes arising that could not be explained by its assumptions (Allais 1953). Kahneman and Tversky (1979) developed Prospect Theory, later adapted to Cumulative Prospect Theory (Tversky and Kahneman, 1992) which was able to provide explanations for paradoxes that EUT could not solve, and is recognised as the most complete model for descriptive decisions under risk and uncertainty to date. CPT was formed under three primary assumptions. The first, initially proposed in the Markowitz (1952) model of utility, suggests that all outcomes must be evaluated relative to a reference point in wealth, thereby analysing relative gains and losses rather than absolute gains and losses. This allowed for the subjective transformation of absolute changes in wealth to relative changes in utility. The second assumption implies that one evaluates outcomes using a value function characterised by both loss aversion and diminishing sensitivity, proposing that individuals tend to be risk seeking (risk averse) in the loss (gain) domain, and that we experience greater disutility for a loss than the utility gained for an equivalent gain. The third assumption, originally proposed by Preston and Baratta (1948), suggests that there is a discrepancy between the objective probability of an outcome, and the subjective probability we assign to it. Probability distortion, refers to our subjective overweighting of smaller probabilities and underweighting of large ones. CPT models this conjecture using a probability weighting function, which transforms objective probabilities into decision-weights to capture these preferences.

Tversky and Kahneman (1992) suggested that the value function be of a power form, and provided an original probability weighting function. However, various functional specifications are able to satisfy the assumptions of CPT, and the literature has sought to identify the utility and probability weighting functions that are most empirically plausible and robust (Conte, Hey, and Moffatt, 2011), where the effectiveness of a functional form has been largely dependent on the decision-making context.

Stott (2006) provided the first comprehensive examination of which specific forms of CPT give the best explanatory account of experimental data, placing an emphasis on accounting for stochastic choices. In other words, capturing the fact that humans make mistakes. The results of Stott (2006) identified that a model incorporating a power value function with a single parameter risky weighting function by Prelec (1998) provided the most accurate explanatory account of the data. The study emphasised that neglecting the parametric approach leaves one at an explanatory disadvantage when analysing risk preferences of any kind.

The results obtained from the existing literature concerning which parametric forms of CPT provide the best explanatory account of data generally favour more flexible specifications, however it is largely dependent on the structure of choice tasks at hand. Camerer and Ho (1994) and Birnbaum and Chavez (1997) found that the flexible power function provides a reliable and accurate representation of the data, with Balcombe and Fraser (2015) highlighting that the power form was far superior to the other forms they consider. However, Blondel (2002) found that the exponential function overruled the power function, whilst Bouchouicha and Vieider (2017) found the logarithmic utility function dominated the power and exponential function in a setting where larger payoffs were involved.

Regarding the explanatory performance of various risky weighting functions, Gonzalez and Wu (1999) found that the two-parameter function adopted by Prelec (1998), and the 'linear in log odds' weighting function (Goldstein and Einhorn, 1987), performed better than the single parameter functions of Tversky and Kahneman (1992) and Prelec (1998). Although they highlight that their data structure could not discriminate between the "linear in log odds" model and the two-parameter Prelec functional form. The explanatory dominance of two-parameter weighting functions is further supported by Bleichrodt and Pinto (2000) and Sneddon and Luce (2001) who argue that two-parameter specifications outperform the single-parameter ones as a result of their flexibility and ability to disentangle probability distortion from elevation. Balcombe and Fraser (2015) also find that the two-parameter form of the Prelec (1998) weighting function was always superior to its nested single-parameter counterpart. However, Wu and Gonzalez (1996) who pool participant data and then proceed to fit functional forms find that, whilst the two parameter functions outperform their nested counterparts, when adjusting for degrees of freedom in each model, it was in fact the single parameter Tversky and Kahneman (1992) risky weighted function that outperformed the rest. This exemplifies the importance of penalising overly complex models.

Probability distortion is a primary driver of the CPT model, but what if subjective probabilities are not driving the preferences of PGs. Regarding the argument on probabilistic education, Smith (2001) find generally (not in relation to problem gambling) that subject error rates in mathematical questions regarding gambling odds are greater than in standard probability questions. However, Delfabbro, Lahn, and Grabosky (2006) find little evidence that young PGs have a poorer understanding of objective odds on gambles, and found that on some tasks concerning binary odds, the PGs were found to be more accurate than non-gamblers in their sample. Shead, Callan, and Hodgins (2008) support these claims when struggling to find evidence that PGs differ from non-gamblers in their understanding of objective probabilities. Interestingly, Lambos and Delfabbro (2007) find that PGs score higher on a cognitive bias examination, but were unable to attribute this difference to a poorer knowledge of objective probabilities or numerical aptitude. Pelletier and Ladouceur (2007) highlight that the importance of knowledge of mathematics as a protective factor against excessive gambling is questionable.This does not mean we can rule out probability distortion as a driver, but is a meaningful insight into the underlying drivers of probability distortion.

However, given how little we know about the structural preferences of PGs, it is important to account for non-CPT models of utility. Specifically, we need to account for models that do not include probability distortion as a primary driver of behaviour. Markowitz (1952) model of utility suggests that individuals use a fourfold pattern of risk when making decisions over probabilistic outcomes. The framework suggests that individuals are locally risk seeking in the gain domain over small monetary outcomes, but as the stake size increases, individuals switch to risk aversion. Whilst in the loss domain, individuals are initially risk averse over small monetary losses, and then risk seeking as the loss increases. Scholten and Read (2014) highlight that Markowitz' conjecture has been overlooked in decision-theory analysis, with probability distortion being too often the primary focus. However, empirical investigations have been able to provide support for parametric specifications that satisfy the assumptions of the fourfold pattern of risk (Abdellaoui, Barrios, and Wakker 2007; Peel and Zhang 2012; Georgalos, Paya, and Peel 2021; Bruhin, Fehr-Duda, and

Epper 2010; Vieider 2012). The current chapter will assess the explanatory power of four CPT specifications and a Markowitz specification to gain further insights into the decision-making processes of PGs.

To the best of our knowledge, there are only two existing attempts to estimate parameters of decision-theory models in the domain of PGs (Ligneul et al. 2012; Ring et al. 2018). Both studies involve PGs and non-gambling controls making repeated decisions between binary monetary lotteries and different sure outcomes. Ligneul et al. (2012) find that PGs have globally shifted preferences towards risky options, and that probability distortion differences were not significant. Whilst Ring et al. (2018) find PGs were only more risk taking in the gain domain, with excessive probability distortion prevailing predominantly in the gain domain. We further contribute to this literature by applying more rigorous econometric estimation and prediction routines, and by relaxing restrictive assumptions. Firstly, the two aforementioned studies impose a representative agent assumption, whereby they assume all subjects can be represented as a single agent, without accounting for within-group heterogeneity – we relax this assumption and carry out the analysis at the subject-level. Secondly, both studies impose linear utility assumptions, which restricts the ability to assess preferences over increasing or decreasing monetary amounts. We allow for flexible parameters in the value function.

1.3 Experimental Data

We use the experimental data of Ring et al. (2018) which comprises of certainty equivalents obtained from a choice menu task. The design was initially proposed by Bruhin, Fehr-Duda, and Epper (2010), and was later adopted by Vieider et al. (2015) and Bouchouicha et al. (2019). This section will outline the participant selection process and the experimental design.

The participant selection process included the recruitment of 74 participants with a mean age of 38.9. This consisted of 25 problem gamblers, 23 habitual gamblers and

TABLE 1.1: **Ring et al. (2018) choice tasks.**

Gains	Losses
(5, 1/2; 0)	(-5, 1/2; 0)
(10, 1/2; 0)	(-10, 1/2; 0)
(20, 1/2; 0)	(-20, 1/2; 0)
(30, 1/2; 0)	(-30, 1/2; 0)
(30, 1/2; 10)	(-20, 1/2; -10)
(30, 1/2; 20)	
(20, 1/8; 0)	(-20, 1/8; 0)
(20, 1/8; 5)	(-20, 1/8; -5)
(20, 2/8; 0)	(-20, 2/8; 0)
(20, 3/8; 0)	(-20, 3/8; 0)
(20, 5/8; 0)	(-20, 5/8; 0)
(20, 6/8; 0)	(-20, 6/8; 0)
(20, 7/8; 0)	(-20, 7/8; 0)
(20, 7/8; 5)	(-20, 7/8; -5)

Notes: There are 14 tasks in the gain domain, and 13 in the loss domain.

a control group of 26 participants classified as non-gamblers. The participant selection procedure was based on the Diagnostic and Statistical Manual of Mental Disorders (DSM-IV-TR). From the sample of applicants, 25 fulfilled at least three of the criteria which classified them as problem gamblers, 23 fulfilled less than three but were gambling at least once a week, being categorised as the habitual gambling control group, and 26 were gambling less than once a month and were therefore recruited as the non-gambling control group. The South Oaks Gambling Screen (SOGS) was implemented to obtain a continuous variable for an individual's gambling behaviour, thus confirming that the three groups were significantly different with regards to their SOGS scores. Finally, all three groups were matched based on characteristics (independent of gambling behaviour) that potentially could have affected task performance: e.g., income, age, education. This chapter uses the data points from 27 of their choice tasks, of which 14 were for pure gains and 13 for pure losses.[1] An overview of the decision task is illustrated in Table 1.1. The binary gambles are modelled as $(x, p; y)$ where a subject can receive a payoff x with probability p, or payoff y with probability $1 - p$.

[1] Ring et al. (2018) implement an additional two mixed tasks to estimate loss aversion, however as our focus is parametric, and we do not wish to over-parameterise the model, we drop these tasks and focus exclusively on risk preferences; namely utility curvature and probability weighting. Additionally, Ring et al. (2018) find that the loss aversion parameter does not add much to the analysis.

The choice task set up requires less tasks and less data points compared to traditional binary choice tasks. Instead of selecting a preference between two binary lotteries, subjects were asked to elicit certainty equivalents for each gamble. This was done by making repeated decisions between binary monetary lotteries and varying sure monetary amounts. The certainty equivalent represents the point of indifference between the lottery and the sure amount. For example, in the first task in the gain domain of Table 1.1, subjects were presented with an urn with eight balls inside, numbered one to eight, and told any ball could be chosen at random with equal probabilities. They were presented with a lottery whereby if balls one to four are selected, the subject would receive €5, whereas if balls five to eight were selected, they would receive €0. Subjects were then presented with a list of sure amounts, ranging from the smallest outcome in the lottery to the largest outcome, so in this case, from €0 to €5. Subjects were told to select, for each sure amount, whether they prefer the lottery or the sure amount.[2] Subjects could only switch at one point, and the certainty equivalent was calculated as being the mean of the two values between which they switched. For example, if at €2.50 and all lower amounts, the subject preferred the lottery, but at €3, switched to preferring the sure amount, then the certainty equivalent would be 2.75. As illustrated in Table 1.1, over the 14 tasks in the gain domain, the payoffs and probabilities varied such as to capture subject level preferences for changes in outcomes and probabilities. In the loss domain, the tasks followed the same structure but only for negative monetary amounts. In the gain domain, payoffs varied between €0 and €30, whilst in the loss domain, payoffs varied between -€20 and €0. In both conditions, probabilities vary by intervals of 0.125, as the number of balls in the urn was always 8, but the number of balls with which they could win or lose varied.

Whilst the payoffs were generally low, the relative increments in payoffs may still have captured sensitivity to outcome changes, as changes from €10 to €30 could have been viewed as triple the monetary payoff. Similarly, many casino gambles involve relatively low payoffs. For these reasons, we see it important to relax the linear utility assumption of Ring et al. (2018), and allow the parameter in the value

[2] It is likely that subjects would start by selecting the lottery, and as the sure amounts increase towards €5, switch to preferring the sure amount.

function to vary freely.

1.4 Preference Functionals

This section will present the 5 model variants under consideration. Parametrically, Cumulative Prospect Theory's predictions can be captured by combining two key functions. The first of which, the utility/value function, evaluates the preference for a monetary amount that will be given with some probability. The literature suggests that the function generally exhibits an S-shape (Dacey 2003; Fishburn and Kochenberger 1979; Kahneman and Tversky 1979) which represents risk aversion for gains and risk seeking for losses. Estimating the parameters of the value function allows us to determine the convexity/concavity of the utility function. The CPT analysis will focus on four value functions: Power, Exponential, Logarithmic and Linear. The latter is a replication of the Ring et al. (2018) analysis. The preference functions will be fit to the data at the subject level to determine the theoretical specification with the most explanatory power.

The most widely used value function is the power function, primarily due to its flexibility and facilitation of constant relative risk aversion (CRRA) preferences. Kahneman and Tversky (1979) propose a power function in their development of CPT. It forms the basis of a Cobb-Douglas utility function and has been praised on its explanatory performance in the vast literature (Luce 1991; Tversky and Kahneman 1992; Wakker and Tversky 1993; Stott 2006; Balcombe and Fraser 2015). Equation 1.1 represents its form in the gain domain.

$$V(x)^{+} = x^{r_G} \tag{1.1}$$

Where x is the monetary incentive and r is our parameter of interest; representing an individual's constant relative risk aversion and is independent of x.[3] r must be greater than 0 to satisfy monotonicity (non-decreasing function) and we can elicit

[3]The subscript G represents the parameter value in the gain domain, and the subscript L, represents the parameter value in the loss domain. We assume subjects will have different preference parameters across the two domains. When we use r, we are referring to the generic risk coefficient independent of the domain

FIGURE 1.1: **Power function plots**

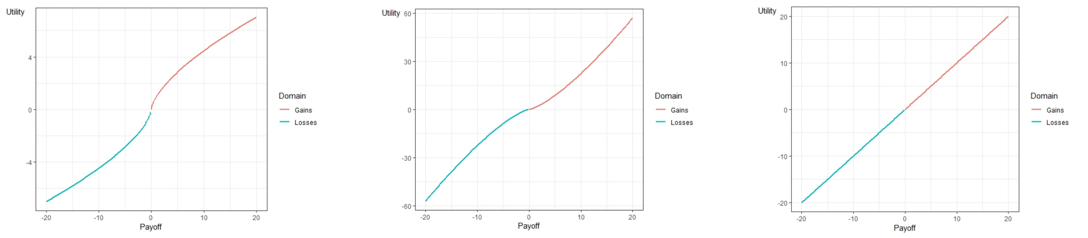

(A) $r < 1$ (B) $r > 1$ (C) $r = 1$

Notes: Plot for when $r = 0.65$, $r = 1.35$, and $r = 1$ respectively.

individual preferences by assessing how r deviates from 1. $r_G > 1$ represents convex preferences in the gain domain, suggesting that marginal utility increases with payoff increases, and is associated with risk-seeking behaviour. $r_G < 1$ represents concave preferences in the gain domain, accommodating the notion of diminishing marginal utility, and is associated with risk-averse behaviour. If $r = 1$, the subject is characterised as an expected value maximiser, associated with risk neutrality and linear utility. Equation 1.2 represents its form in the loss domain, where the same parameter restrictions apply.

$$V(x)^- = -(-x)^{r_L}, \quad \text{if } x < 0 \tag{1.2}$$

However, preferences are reversed, such that $r_L > 1$ represents concavity and risk-averse behaviour, whilst $r_L < 1$ represents convexity and risk-seeking behaviour. Figures 1.1a, 1.1b and 1.1c illustrate the power function for when $r < 1$, $r > 1$, and $r = 1$ respectively.

Ring et al. (2018) impose a linear utility assumption, which is illustrated in Figure 1.1c. We aim to test the robustness of this assumption, and so in addition to the flexible power function, we also test its nested counterpart, linear utility, as a separate specification. The linearity assumption suggests that the value function takes the form $V(x) = x$.[4]

[4]This assumption neglects changes in preferences over outcomes, and resides in assuming any changing risk attitudes are a result of probabilistic sensitivity.

Certainty equivalent (CE) functions will be fitted to each parametric specification. A CE refers to the amount of money a decision-maker is happy to receive so as to be indifferent between a sure monetary outcome and the CPT utility of a gamble. CE's are advantageous in that they permit a stochastic structure to be developed and allow for the quantitative measure of preferences, rather than relying solely on binary choices. The flexible adjustment to subject-level preferences allows for the estimation and prediction of the point at which a subject will switch from a lottery to a sure amount. The CE's are used directly to estimate CPT parameters. Mathematically, the CE of a gamble can be interpreted as the inverse of the expected utility of a gamble, as in equation 1.3 :

$$CE = V^{-1}(EU) \tag{1.3}$$

where $EU = p \cdot V(x) + (1-p) \cdot V(y)$. $V(x)$ is the utility gained from winning, $V(y)$ is the disutility received from losing. Whilst p and $1-p$ represent the corresponding probabilities of each outcome respectively. The certainty equivalents in the gain and loss domain are shown in equations 1.4 and 1.5.

$$ceg(eu) = eu^{\frac{1}{r_G}} \tag{1.4}$$

$$cel(eu) = -(-eu)^{\frac{1}{r_L}} \tag{1.5}$$

ceg is the CE for gains, and *cel* is the CE for losses. The second value function under examination is the exponential function. This function is able to accommodate increasing relative risk aversion ($IRRA$) and constant absolute risk aversion $CARA$. Wakker and Tversky (1993) suggest that the value function should always take an exponential form if preferences are invariant under the addition of a positive constant to outcomes. It has, however, been known to have certain drawbacks in the form of boundedness of utility. The function takes the form of equations 1.6 and 1.8 in the gain and loss domain, with equations 1.7 and 1.9 as their corresponding certainty equivalents.

$$V(x)^{+} = 1 - e^{-r_G x} \tag{1.6}$$

$$ceg(eu) = -\frac{\log(1 - eu)}{r_G} \tag{1.7}$$

$$V(x)^{-} = -(1 - e^{r_L x}), \quad \text{if } x < 0 \tag{1.8}$$

$$cel(eu) = \frac{\log(1 + eu)}{r_L} \tag{1.9}$$

Where $r > 0$ to satisfy monotonicity. However, under the assumption that $r > 0$, we will only be able to elicit risk averse preferences in gains and risk seeking preferences in losses. To ensure we can recover the reverse preferences, whilst still ensuring the function is non-decreasing, we follow Wakker (2010) and fit an alternative form for when $r < 0$ such that we are able to elicit convex (concave) preferences in the gain (loss) domain.[5] Equations 1.10 and 1.12 are the functional form in the gain and loss domain, with equations 1.11 and 1.13 as their respective certainty equivalents.

$$V(x)^{+} = e^{-r_G x} - 1 \tag{1.10}$$

$$ceg(eu) = -\frac{\log(1 + eu)}{r_G} \tag{1.11}$$

$$V(x)^{-} = -(e^{-r_L x} - 1), \quad \text{if } x < 0 \tag{1.12}$$

$$cel(eu) = \frac{\log(1 - eu)}{r_L} \tag{1.13}$$

This specification permits risk seeking preferences in gains, and risk-averse preferences in losses, to be recovered. When $r = 0$, the model converges to linear utility. Figure 1.2 illustrates the convexity and concavity of this function when $r > 0$ and $r < 0$ respectively.

The third value function, well known as the first utility function, was proposed by Bernoulli in the 18*th* century and takes a logarithmic form. The logarithmic function accommodates when incremental utility is proportional to incremental wealth when measured as a proportion of existing wealth. Simply put, the wealthier an individual is, the more likely they are to take on risk as long as the rewards are

[5] Note that Wakker (2010) show there are two ways of doing this. One way is to propose a new parametric form as we show here, and another is to normalise the function over r, as we do with the logarithmic function, at which point when $r < 0$, monotonicity is still satisfied. Both methods are interchangeable and yield the same results.

FIGURE 1.2: **Exponential function plots.**

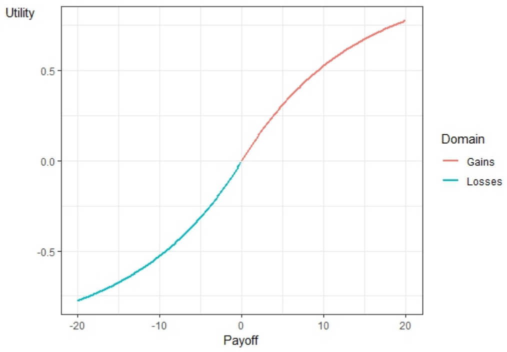

(A) $r > 0$

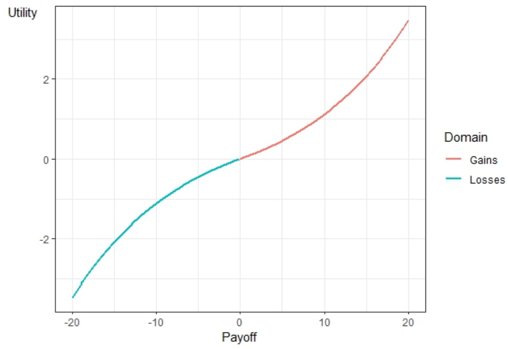

(B) $r < 0$

Notes: Plots for when $r = 0.075$ and $r = -0.075$ respectively.

substantial enough (and vice versa). This function is also able to capture Increasing Relative Risk Aversion ($IRRA$) and Decreasing Absolute Risk Aversion ($DARA$) (Bouchouicha and Vieider, 2017). Although this function has not had as much success as the power and exponential functions in the decision-theory literature when stakes are generally lower, dependent on relative perceptions of payoffs, it still may play a part in explaining the risk-attitudes of our subjects. Equations 1.14 and 1.16 illustrate the functional form in the gain and loss domain, and the $CE's$ are equations 1.15 and 1.17.

$$V(x)^{+} = \frac{\log(1 + r_G x)}{r_G} \tag{1.14}$$

$$ceg(eu) = \frac{e^{eu} - 1}{r_G} \tag{1.15}$$

$$V(x)^{-} = -\frac{\log(1 - r_L x)}{r_L}, \quad \text{if } x < 0 \tag{1.16}$$

$$cel(eu) = -\frac{e^{eu} - 1}{r_L} \tag{1.17}$$

Where when $r > 0$ ($r < 0$), the function is concave (convex) in gains and convex (concave) in losses.[6] Figure 1.3a and 1.3b illustrate the logarithmic function when

[6]Without normalising over r, whenever $r < 0$, monotonicity would not be satisfied, so normalisation is essential.

FIGURE 1.3: **Logarithmic function plots.**

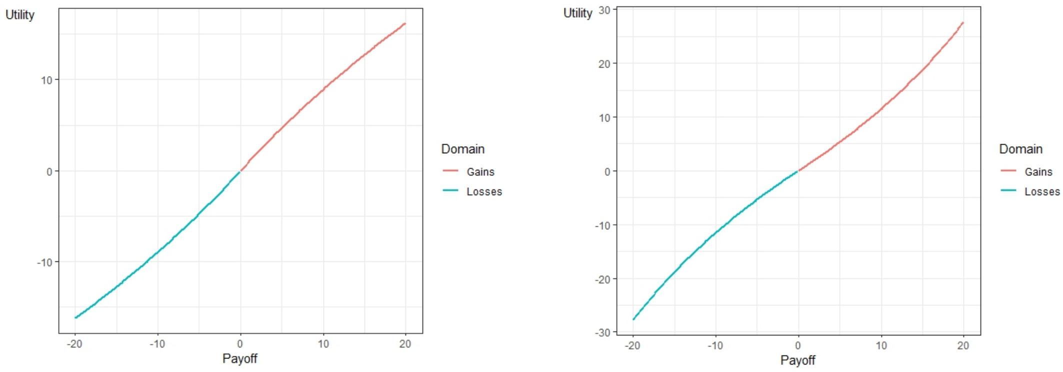

(A) $r > 0$ (B) $r < 0$

Notes: Plots for when $r = 0.025$ and $r = -0.025$ respectively.

$r > 0$ and $r < 0$ respectively.

Cumulative Prospect Theory assumes that a utility function is estimated in conjunction with a probability weighted function to determine an individual's subjective perception of probabilities. The weighted probability function transforms the probability of obtaining the monetary amount into a decision weight that lies between 0 and 1. We combine the four aforementioned value functions with equation 1.18: the two-parameter weighting function by Prelec (1998).

$$w(p) = e^{-\delta(-\log(p))^{\gamma}} \tag{1.18}$$

Where $\gamma > 0$ and $\delta > 0$. The lower bounds have been set to ensure the function is monotonically increasing. γ represents the curvature of the weighting function, whilst δ measures the elevation of the weighting function. We assume unique weighting functions, and therefore unique parameters, in the gain and loss domain. Gonzalez and Wu (1999) highlight that elevation is logically independent from curvature, and that this should be reflected in two independent variables within the weighting function. γ, the more commonly used parameter within the weighting function, represents diminishing sensitivity, thus predicting how the function is concave and then convex, and how subjects become less sensitive to probabilities as they move away from the reference point probabilities (0 and 1). However, as Gonzalez

FIGURE 1.4: **Prelec (1998) two-parameter weighting function plot**

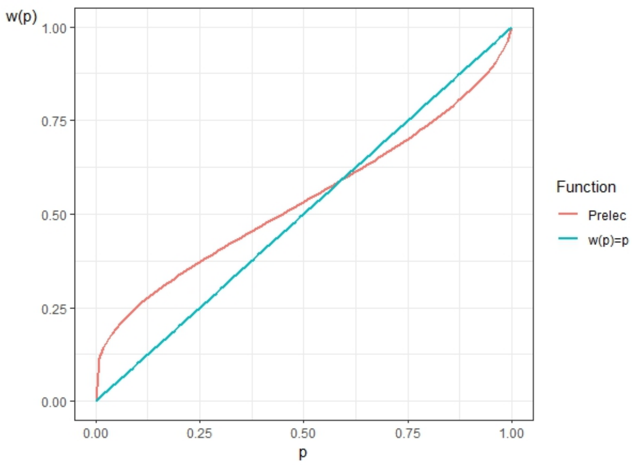

Notes: Plot for when $\gamma = 0.65$ and $\delta = 0.8$. γ represents probability distortion, and δ represents probability elevation.

and Wu (1999) illustrate, this alone provides an incomplete account of the weighting function, as it does not account for the level of overweighting or underweighting relative to the objective probability. A function could be concave and then convex, represented with an inverse S-shaped graph, but we do not know whether this is always below the 45 degree line, above the 45 degree line, or cutting it at a certain point. It is δ that captures the absolute value of the weighting function, which represents global overweighting or underweighting of the probability scale. This represents the attractiveness of taking risk in the chance domain, where, in the gain domain, a higher (lower) value of δ_G represents a shifted downward (upward) weighting function, a more underweighted (overweighted) probability scale and thus a less (more) attractive risk. However, in the loss domain, this higher δ_L represents a more attractive risk, as the underweighted probability scale makes losses feel less likely. Figure 1.4 illustrates a weighting function where $\gamma = 0.65$ and $\delta = 0.8$.

Only the two-parameter function by Prelec (1998) will be fit to the experimental data. We prioritise this functional form for various reasons: Firstly, due to the extreme, but seemingly ambiguous decisions of problem gamblers, it is important to model the individual preferences in the most flexible way possible. Therefore, we wish to provide a specification that can disentangle probability distortion from elevation (Gonzalez and Wu 1999; Abdellaoui, Vossmann, and Weber 2005; l'Haridon et al. 2010). The literature has illustrated how the two-parameter specifications in the probability domain are advantageous when engaging in subject-level analysis. Secondly, if CPT ends up having an explanatory advantage over models that neglect probability distortion, pinpointing where in the weighting function this explanatory power is coming from will contribute to a more holistic analysis. Thirdly, whilst there are alternative two-parameter risky weighting functions, adopting the two-parameter Prelec function will permit for more direct comparison of results with the Ring et al. (2018) analysis. Bouchouicha and Vieider (2017) also confirmed that when adopting alternative two-parameter functions, their results were unchanged, however Prelec's function provided the best-fit to their data.[7] This concludes the CPT specifications included in the analysis.

As aforementioned, the literature suggests that PGs have an adequate understanding of objective probabilities in a gambling context, so it is important to assess the descriptive validity of non-EU models that do not incorporate probability distortion. Markowitz (1952) proposes a four-fold pattern of risk attitudes. For the current analysis, we employ the parametric form and restrictions of an expo-power specification used in the Georgalos, Paya, and Peel (2021) analysis. The expo-power function was originally proposed by Saha (1993) and has since been proven to capture the assumptions of the Markowitz model of utility. Peel and Zhang (2012) demonstrate that assuming a power value function, when the true data-generating process is of an expo-power form, can lead to misleading implications regarding the properties of the value function and the degree of probability distortion. Equations 1.19 and 1.20 represent the functional form in the gain and loss domain respectively.

$$V(x)^{+} = 1 - e^{-\alpha_G x^{\eta_G}} \tag{1.19}$$

$$V(x)^{-} = -(1 - e^{-\alpha_L (-x)^{\eta_L}}), \quad \text{if } x < 0 \tag{1.20}$$

Where, as usual, the first equation is for the gain frame $(x > 0)$ and the second for the loss frame $(x < 0)$. The power parameter, η, can be interpreted as a measure of absolute risk aversion ($\eta < 1 \rightarrow$ DARA, $\eta = 1 \rightarrow$ CARA, $\eta > 1 \rightarrow$ IARA), and the α parameter as a measure of relative risk aversion ($\alpha < 0 \rightarrow$ DRRA, $\alpha > 0 \rightarrow$ IRRA).

[7] We ran a separate horse race between other one and two parameter weighting specifications, but the two parameter Prelec (1998) function was the most descriptive on all occasions, so we ignore the other weighting functions in the proceeding analysis.

FIGURE 1.5: **Markowitz (1952) utility function**

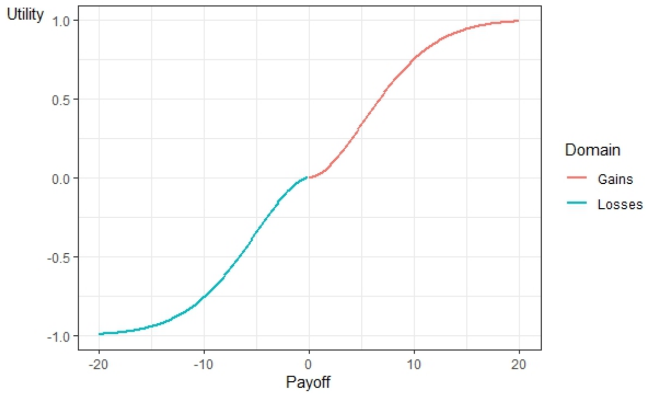

Notes: This plot represents the fourfold pattern of risk, where η = 1.75 and $\alpha = 0.05$.

Georgalos et al., (2021) suggest that when η is above unity, or when the agent is exhibiting $IARA$ preferences, the expo-power function can be categorised as consistent with Markowitz's conjectures. Under their specification, with $\eta > 1$, the agent is risk seeking over gains when $\eta_G - 1 - \alpha_G \eta_G x^{\eta_G} > 0$, risk averse when $\eta_G - 1 - \alpha_G \eta_G x^{\eta_G} < 0$, and risk-neutral when $\eta_G - 1 - \alpha_G \eta_G x^{\eta_G} = 0$. Under these assumptions, if a participant has $\eta_G > 1$, they can be categorised as exhibiting preferences that coincide with the Markowitz model of utility. That being, locally risk seeking then risk averse as stakes increase in the gain domain, and vice versa in loss domain. Figure 1.5 illustrates the Markowitz model of utility plot for $\eta = 1.75$ and $\alpha = 0.05$.

Equations 1.21 and 1.22 represent the corresponding certainty equivalents in the gain and loss domain.

$$CEG(EU) = \left(-\frac{\log(1 - eu)}{\alpha_G}\right)^{\frac{1}{\eta_G}} \tag{1.21}$$

$$CEL(EU) = -\left(-\frac{\log(1 + eu)}{\alpha_L}\right)^{\frac{1}{\eta_L}} \tag{1.22}$$

The five aforementioned specifications will be fitted to the Ring et al. (2018) data set to determine which model provides the best fit for each participant in terms of descriptive and predictive capacity. Additionally, individual parameter estimates will be determined based on the data-generating process that is most compatible with each individual's decision-making process. The following section will provide the econometric specification.

1.5 Econometric Specification

Our econometric specification comprises of two main components. The first uses Maximum Likelihood Estimation (MLE) procedures to determine each models' explanatory power. The second formulates a Cross Validation routine to determine each models' predictive power using Leave-One-Out-Cross-Validation (LOOCV) techniques.

As an additional extension to the Ring et al. (2018) analysis, we proceed to estimate our models using MLE, rather than Non-linear Least Squares (NLS) due to the advantages it has in terms of using distributional information. It has been shown that NLS may be less efficient than MLE as it does not use information about the distribution of the dependent variable (see Greene (2003)). Similarly, MLE has been shown to produce more reliable estimates in external investigations (Patmanidis et al., 2019). However, we do not expect differences in results between the two approaches to be statistically significant.

Similarly, we opt for a parametric approach over a non-parametric approach for three key reasons. Firstly, given the small sample size, it will be easier to capture variations in preferences using a parametric approach. Secondly, a key contribution of our analysis is regarding the relaxation of homogeneity assumptions, so a parametric approach is required to highlight variations in our CRRA parameter. Finally, as the literature on PG risk attitudes is limited, we wish to identify where underlying characteristics stem from. A parametric approach can disentangle utility curvature, probabilistic sensitivity, and probability elevation, such as to increase interpretability of results.

The analysis will be based on the certainty equivalents calculated at the subject level over each prospect. This allows us to recover quantitative measures regarding when an individual will switch from a sure outcome to the risky gamble. Each model will attempt to estimate parameters given the subject-level certainty equivalents, such as to determine a parameter set that maximises the likelihood of observing the experimental data with some noise. We first develop a stochastic structure to account for the individual mistakes in decision-making (Bouchouicha and Vieider, 2017). Equation 1.23 are the certainty equivalents predicted by each of our models.

$$\hat{ce}_i = V^{-1}(w_s(p_i) \cdot V(x_i) + (1 - w_s(p_i)) \cdot V(y_i)) \tag{1.23}$$

Where the CE is a function of the win and loss utilities and their respective weighting of probabilities. Note that for the linear utility specification, $V(x) = x$, and under the Markowitz model of utility, $w(p) = p$. Individual choices are stochastic, and subject-level decisions are subject to some degree of estimation error, so our actual certainty equivalent, as shown in equation 1.24 will be equal to the CE's predicted by our model, plus some error term.

$$ce_i = \hat{ce}_i + \epsilon_i \tag{1.24}$$

This error term provides a degree of theoretical consistency that each participant has over all choices made. Consistent choices (a smaller error) will bring the actual certainty equivalent closer to that predicted by our models. We assume this error term is normally distributed with mean zero $\epsilon_i \sim N(0, \sigma_i^2)$.

To determine each model's degree of accuracy, we need to estimate σ, the standard deviation, to deduce how much noise there is in the data for each participant. When a model's purpose is to simulate a process that generated the data, some of the information will inevitably be lost in the process. A standard deviation that is closer to zero, for a given set of parameters, corresponds to less noise, suggesting the actual choices are closer to the theoretically optimal ones. To estimate the standard deviation, we first express the probability density function (PDF), as in equation 1.25 for a given subject and prospect. The PDF is advantageous for this analysis, in that its value at any point in the whole sample can provide a relative likelihood that the value of the predicted random variable from the model will equal the variable generated by the data.

$$\psi_{in}(\theta, \sigma_i) = \frac{\phi}{\sigma_i}\left(\frac{\hat{ce}_i - ce_i}{\sigma_i}\right) \tag{1.25}$$

The PDF $\psi(.)$, is expressed for a given subject n and prospect i, with ϕ being the standard normal density function. The parameter θ represents the vector of parameters to be estimated $(r, \gamma, \delta, \alpha, \eta)$ and σ_i represents the standard deviation. The subscript i allows the error to depend on the specific prospect. We do this by allowing it to depend on the difference between the high and low outcome in each prospect: $\sigma_i = \sigma|x_i - y_i|$. By doing this, the error term can account for when there is a larger gap between the high and low outcomes.

The various specifications are fit at the subject-level to avoid representative agent violations (Navarro et al., 2006). The parameters are estimated by standard maximum likelihood procedures. By taking the product of all the density functions across prospects for each subject, we obtain equation 1.26: the likelihood function:

$$L_n(\theta) = \prod_{i \in i,n} \psi_{in}(\theta, \sigma_{i,s}) \tag{1.26}$$

The subscript n captures the subject specific likelihood, whilst the subscript s on the error term allows us to control for heteroskedasticity across the two domains, as the precision parameters will be different for gains and losses. The log-likelihood (LL) function, equation 1.27, is obtained by taking log of equation 1.26.

$$LL(\theta) = \sum_{n=1}^{N} \log(L_n(\theta)) \tag{1.27}$$

The estimation of this function is referred to as the Maximum Likelihood Estimator (MLE), where the maximized parameter values therefore have the highest likelihood of generating the observed data within the parametric constraints of the given theory (Stott, 2006).[8]

We report a measure of goodness of fit, namely the Akaike Information Criterion (AIC) (Akaike, 1973), to penalise models with more parameters. The explanatory performance of models is affected by the number of free parameters it obtains, in

[8] All statistical analyses were performed using R Statistical Software (v4.3.1; R Core Team 2023).

that under MLE, a more complex model will obtain a stronger, but potentially biased, measure of fit. AIC is an adjustment technique that accounts for the degrees of freedom in the model. The AIC identifies which model is descriptively superior rather than the likelihood of a model being true. The AIC calculates the model's prediction error by estimating the relative amount of information lost in any data generating process, and thereby predicts the relative quality of a statistical model for a given data set. It uses the maximised log-likelihood whilst accounting for the number of free parameters in the model, therefore penalizing models that are overfitted. The AIC's will be calculated manually using equation 1.28.

$$\mathrm{AIC} = -2LL(\theta) + 2k \tag{1.28}$$

Where $LL(\theta)$ is the maximised log likelihood and k is the number of free parameters. After compensating for degrees of freedom, a lower value of the AIC indicates a better fit to the data.

We estimate parameters in two settings. One in which we assume our parameter estimates are identical in the gain and loss domain, and one in which we relax this assumption to identify differences across domains. Whilst not estimating loss aversion directly, we can indirectly identify an individual's aversion to losses by estimating each parameter in the gain and loss domain separately. Additionally, we wish to assess if a model's performance is hindered by a homogeneous parameter assumption. Finally, we look to identify the predictive performance of our models. When model complexity is high relative to the number of tasks, estimation procedures are susceptible to overfitting. This occurs when the model captures not only the underlying patterns (signal) but also the random noise in the data, leading to potentially misleading parameter estimates and results. The risk of overfitting is particularly pronounced when there are many parameters or when non-linear relationships are permitted in the model. To address this issue, we utilise Leave-One-Out Cross-Validation (LOOCV) to assess how well each model predicts out-of-sample data.

LOOCV, the most robust and computationally expensive of the cross validation methods (Stone, 1974), works by first dividing the data set into folds (subsets). These comprise of the training sets and the test set. We use MLE to retrieve parameter estimates from the data in the training sets. We then fit these parameters to our various models to identify their accuracy in predicting the results from our test set. In our specific case, there are 27 data points, and as LOOCV is a k-fold cross validation classification in which $k = n$, we have 27 folds. Therefore, each iteration consists of estimating on 26 data-points, and using the results to predict the final data point. 27 iterations take place, to ensure each data point acts as the test set. For each subject, n, we calculate the mean squared error (MSE) of the prediction, for each iteration, j, using equation 1.29.

$$\text{MSE}_{nj} = (CE_{nj} - \hat{CE}_{nj})^2 \tag{1.29}$$

Where CE represents the actual certainty equivalent for the k_{th} element of the data set, and $\hat{CE}$ represents the certainty equivalent predicted by the model using the estimates from the training set. MSE assesses how close estimates or forecasts are to actual values. The squared element of the function ensures our trained models have no outlier predictions with huge errors by placing a larger weight on said errors. The mean of each subjects 27 MSE's is our statistical metric to assess a models predictive capacity, with a lower MSE representing a more accurate prediction and a more reliable data generating process.

1.6 Results

This section will provide the results from the three primary focuses of this chapter. Firstly, we will provide the AIC results from all models at the subject level in the analysis where we assumed homogeneity over gain and loss domain parameters, and compare these to the AIC results from all models at the subject level where we assume heterogeneity in gain/loss parameters. This will highlight the importance of allowing for heterogeneity across domains. We do not report the standard log-likelihood coefficients due to our models having different numbers of parameters, so to reduce the chances of selecting the wrong optimal model as a result of over-fitting, we focus primarily on the AIC results which were derived from our log-likelihood coefficients. Secondly, we will report the MSE results from our cross-validation prediction routine to identify differences in-sample and out-of-sample. As a key element of our analysis is to assess the benefits and hindrances of assuming linearity in utility, we will also report two equivalence tests for our Power model and our Linear specification. These will comprise of a likelihood ratio test and a data simulation routine to identify whether individuals adopt linear utility and whether these assumptions are replicable. Following the model comparison analysis, we will recover parameters to identify statistically significant differences across groups in terms of individual risk preferences.

1.6.1 Homogeneity Vs Heterogeneity in Gain and Loss Domain

Starting with our AIC analysis in which we impose homogeneity in risk parameters across the gain and loss domain, we find that the Power specification's explanatory power dominates substantially across subjects and groups. For the Problem Gambler group, 68% of participants were best categorised with a Power specification, whilst the Exponential Model and Markowitz model best explained 12% of the Problem gamblers. The linear utility model was only able to explain the behaviour of 8% of the problem gamblers, and the Logarithmic model failed to explain the behaviour of any of the problem gamblers.

For the Habitual gambler group, again the Power specification came forward as the leading data generating process, accounting for 69.57% of this subgroup, followed by the Exponential model, explaining 13.04% of the behaviour in this group, whilst the Linear Utility Model and the Markowitz model accounted for 8.7% each of the groups behaviours. The results of the Control group follow the same pattern, with the Power model explaining the behaviour of 62% of the subjects, followed by the Exponential model, Linear, and Markowitz model with percentages 15.38%, 15.38% and 7.69% respectively.

We find that these results alter substantially when we allow for heterogeneity in risk preferences across domains. For the PG group, now 40% are best categorised by the Exponential specification, 36% with the Power, 16% with a Markowitz model,

FIGURE 1.6: **AIC Results**

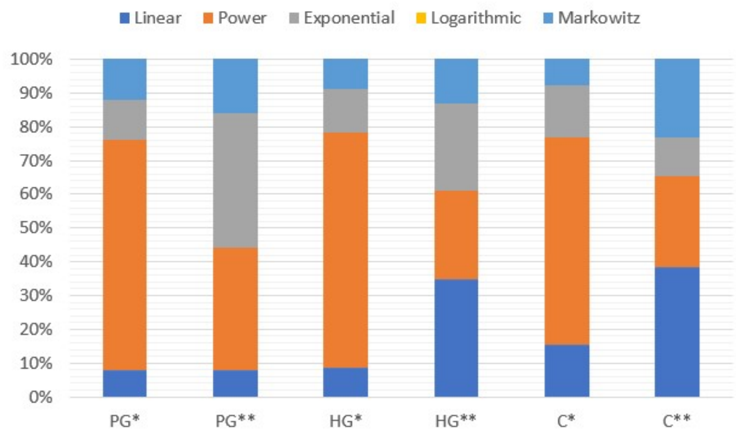

Notes: Percentage of individuals in each group that are best described by each model based on AIC results. * represents the same gain/loss parameter. ** represents different gain/loss parameters.

and 8% with linear utility. Whereas in the HG group, the Linear utility model dominates in terms of descriptive power, accounting for 34.78% of the subjects, with the power and exponential models accounting for 26.09% of the subjects each, and the Markowitz model explaining the behaviour of 13.04% of the subjects. Again, the linear specification dominates for the Control group, accounting for 38.46% of the subjects' behaviours, whilst the power model can only explain that of 26.92% of the subjects, Markowitz explaining 23.08% and the Exponential accounting for 11.54% of individual behaviour. Figure 1.6 illustrates the results of both the homogeneity and heterogeneity analysis across each of the 3 groups.

With regards to overall explanatory power and whether it is more effective to assume heterogeneity or homogeneity in gain/loss parameters, we utilise the likelihood ratio test (LRT) between each model that assumes heterogeneity in risk preferences across domains, and its nested counterpart which imposes a homogeneity assumption. For all 5 of our models, 100% of individuals were better characterised by a heterogeneous model, with p-values of 0.00 showing for all participants over all 5 models. We therefore highlight the importance of accounting for heterogeneity in risk preference parameters across domain. For this reason, our analysis from here on will relax the homogeneity assumption, and we will only report results from our heterogeneous specifications.

Throughout the analysis, we identified very similar AIC results for the Power,

FIGURE 1.7: **LOOCV Results**

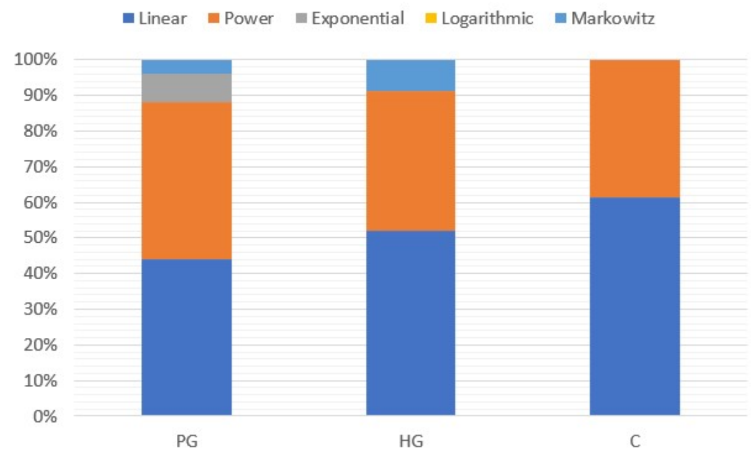

Notes: Percentage of individuals that are best described by each model based on our cross validation prediction results.

Linear, and Exponential models. We therefore assess the predictive capacity of each model in case it can illuminate additional insights about the data-generating processes.

1.6.2 Out-of-Sample Prediction

Our MSE results, from our LOOCV analysis, show that the Exponential Model, the Logarithmic Model and the Markowitz specification all struggled to predict the CEs of subjects, and it was the Linear utility specification and the Power Specification that prevailed. We find that for the PGs, the linear model and the Power model were joint top in their predictive power, each predicting the behaviour of 44% of the subjects with the highest accuracy. However, when analysing the HGs and the Control Group, the linear specification marginally dominates, predicting the behaviour of 52.17% and 61.54% of these groups respectively with the smallest error. The Power specification accounted for 39.13% of subjects in the HG group, and 38.46% of those in the Control group. Figure 1.7 illustrates the results from our MSE predictions.

Interestingly, although the exponential model seemed to work well in an explanatory setting, it struggled in prediction; perhaps due to overfitting, which is likely given the small sample size and limited data points. The majority of subjects are best explained by either a flexible power utility model, or its nested counterpart, the linear utility model.

1.6.3 Equivalence

The similarity in log-likelihoods, AIC statistics, and MSE statistics across the power model and linear utility model suggests there is potential equivalence between the two specifications. That being, subjects characterised by the power utility function may obtain a CRRA parameter that is equal to 1. If this is the case, then we can confirm that the the linear utility assumption is robust. In this section we test for equivalence to identify if this holds for all subjects. If it does not hold for all, it is safer to implement a flexible power function, as if any subjects are characterised with a non-linear value function, imposing linear utility will inflate their probability weighting estimates and restrict the reliability of results.

We provide two equivalence measurement techniques. Firstly, we provide the results of a likelihood ratio test to identify if there are statistically significant differences between the Linear utility model and the Power model in terms of their maximum likelihood coefficients. We then follow Pachur, Suter, and Hertwig (2017) and use simulations to determine equivalence. This involves estimating parameters using the power model, simulating synthetic data using these parameter coefficients, and then using the synthetic choice data to estimate parameters from the linear utility model. If simulated estimates are identical to that of the actual estimates, it confirms the hypothesis of equivalence for these subjects (Pachur, Suter, and Hertwig, 2017), providing support for linear utility. Additionally, it would confirm that adopting a flexible power function would not bias the weighting function parameters for those characterised as linear utility maximisers.

The likelihood ratio test (Vuong, 1989) is a hypothesis test that allows you to select the "best" model between two nested models. The test uses the log-likelihood values, at the subject level, to determine if the power model is the statistically more probable data-generating process. Let us denote the log-likelihood for subject n for the power model as $LL_{nPower}(\theta)$, and the log-likelihood for subject n for the linear model as $LL_{nLinear}(\theta)$. Under the null hypothesis that the true utility function is linear, we compute the log-likelihood test statistic, Λ, at the subject level as:

$$\Lambda_n = -2\left(\log LL_{nLinear}(\theta) - \log LL_{nPower}(\theta)\right) \tag{1.30}$$

where Λ follows a chi-squared distribution with degrees of freedom equal to the difference in the number of parameters between the two models. In our case, this is 1, $\Lambda \sim \chi^2(1)$. Our likelihood ratio tests show that for 45/74 of the subjects, the power model was not significantly better than its nested linear utility model ($p > 0.05$). For these subjects, we want to test whether their risk aversion parameter (r) is significantly different from 1 to confirm that their utility is in fact linear. To do this, we calculate confidence intervals (CI) in the gain and loss domain:

$$CI = r \pm z_{\alpha/2} \cdot \mathrm{SE}(r) \tag{1.31}$$

where $z_{\alpha/2} \approx 1.96$ is the critical value corresponding to the 95% confidence interval ($\alpha = 0.05$), and SE corresponds to the error term on r, the risk aversion parameter. We find that 41 of the 45 participants in the gain and loss domain had confidence intervals surrounding one. This suggests there are no statistically significant differences between the power model and the linear utility model for 41/75 of the subjects.

Finally, for our simulation routine, we estimate parameter sets (θ_P) using the power model (P), and simulate synthetic data, D_{PS}, using θ_P. We then estimate parameters from data set D_{LS}, using the linear utility model (L). This it to identify if adopting a power utility function, when an individual has linear utility, biases the parameter estimates of the probability weighting function. Comparing the simulated estimates to the original parameter estimates of the linear utility model, we find that $p > 0.7$ for all parameters of interest (γ_G, γ_L, δ_G, δ_L). This confirms equivalence, and that implementing a flexible power function will not bias the weighting estimates. [9]

1.6.4 Parameter Estimates

This section recovers and analyses parameter estimates using the Power model. We decide to utilise the power utility function for all participants for 3 main reasons:

[9] Additionally, we run a parameter interval test over the full sample to see if the absolute difference in the parameter values is less than 0.05, we get that γ_G is equivalent for 61 participants, γ_L is equivalent for 66 participants, δ_G for 62 participants, and δ_L for 69. Therefore for the majority of, but not all, subjects, a linear utility assumption would not bias the weighting function parameter estimates. However, it does restrict the identification of outcome sensitivity, where these preferences are likely to have been picked up by the precision parameters.

TABLE 1.2: **Summary Statistics:** μ

Parameters	Group μ			Welch two sample t-test p-value		
	PG	HG	C	PG vs HG	PG vs C	HG vs C
γ_G	0.548	0.823	0.765	0.042**	0.073*	0.654
γ_L	0.7079	0.721	0.785	0.459	0.27	0.319
δ_G	1.03	1.079	1.368	0.408	0.039*	0.0614*
δ_L	1.2721	0.8574	0.8966	0.015**	0.013**	0.6037
r_G	1.982	1.215	1.408	0.07*	0.126	0.776
r_L	1.2582	0.872	0.998	0.021**	0.071*	0.828

Notes: Mean parameter values, μ, for each subgroup and subsequent t-test p-values for group comparisons. * indicates statistical significance at the 5% level, and ** indicates statistical significance at the 10% level. Again, the subscript G represents the gain domain, and L, represents the loss domain.

(1) The linear utility model is a nested counterpart of the power function, so even if individuals adopt linear utility, this will still be captured in the power model. (2) Using both the linear model and the power model for estimation based on each individual's most accurate data-generating process could lead to inconsistent results. This is because slight deviations from 1 in our r parameter for subjects with linear utility may be transferred to the weighting parameters, leading to inconsistent and incomparable results across subjects and groups. (3) If we use the linear utility model for all subjects, those better characterised by a power value function may have probability weighting estimates that are biased and overfitted.

The analysis will focus on identifying statistically significant differences between the three groups in terms of their levels of risk aversion and their subjective perceptions of probabilities (distortion and elevation). In doing so, we can attempt to uncover the cognitive patterns associated with gambling addiction.

Table 1.2 provides the mean parameter values for each group, as well as the p values from our Welch two-sample t-test in order to identify if any group differences are statistically significant.

Given that we find large variations in parameter estimates within-groups, in Table 1.3 we report the median values for the preference parameters, as well as the p values from a Wilcoxon signed rank test (Wilcoxon, 1992).

TABLE 1.3: **Summary Statistics:** *Median*

Parameters	Group median			Wilcoxon signed rank test p-value		
	PG	HG	C	PG vs HG	PG vs C	HG vs C
γ_G	0.455	0.852	0.75	0.057*	0.0674*	0.7868
γ_L	0.89	0.764	0.721	0.718	0.947	0.756
δ_G	0.837	1.045	1.193	0.724	0.048**	0.173
δ_L	1.289	0.843	0.91	0.044**	0.074*	0.521
r_G	1.019	1.059	1.25	0.592	0.99	0.2412
r_L	1.288	0.95	1.014	0.052*	0.089*	0.383

Notes: Median parameter values for each subgroup and subsequent Wilcoxon signed rank test p-values for group comparisons. * indicates statistical significance at the 5% level, and ** indicates statistical significance at the 10% level. Again, the subscript G represents the gain domain, and L, represents the loss domain.

Gain Domain CRRA

Contrary to the assumptions of Ring et al. (2018), who first run their parametric analysis with a linear value function, and then with $r = 0.5$ and $r = 0.75$ to confirm the robustness of their assumption, we find that, with a flexible power value function, most participants exhibited convex functions in the gain domain ($r_G > 1$), suggesting that marginal utility does not diminish as payoffs increase, and that all groups were generally more risk seeking than risk averse. Implementing a linear utility assumption, or only testing where $r < 1$ will therefore restrict the analysis in picking up any risk seeking preferences, and consequently transfer these effects to the probability weighting parameters.

From Table 1.2 we can see that convexity was much more pronounced for the problem gamblers group, who had a mean r_G value of 1.98, as opposed to slightly more "rational" values of 1.21 and 1.41 for our habitual gambler group and control group respectively. However, to account for extreme parameter values that some subjects exhibited, we report the median values, and find that the PG, HG and C group have median r_G values of 1.019, 1.059, and 1.25 respectively. Clearly there was substantial heterogeneity within each of the groups regarding the CRRA parameter in the gain domain, especially in the PG group, where some of the participants exhibited extreme risk seeking behaviour, thereby inflating the group mean. This provides further support for our subject-level analysis, and is likely why statistically significant differences between the groups did not prevail. Nonetheless, Figure 1.8a highlights the convexity in the CRRA function, and illustrates the increased sensitivity some of the PGs exhibited in the gain domain. As aforementioned, this was not the case for the whole group, therefore the subject level analysis, as well as relaxing the linear utility assumption, was appropriate to pick up on this behaviour.

Loss Domain CRRA

When focusing on the loss domain, however, statistically significant difference did prevail and we find that the PGs have more convex CRRA functions, suggesting an increased sensitivity to losses, which is widely associated with risk aversion. We can see that the mean r_L value for the problem gamblers was 1.26, as opposed to the neutral/slightly concave values of 0.872 and 0.998 for the habitual gamblers and control group respectively. The median values coincide with the means, whereby the PGs, HGs and Cs have median parameter estimates of 1.28, 0.95, and 1.01 respectively.

This result is illustrated in Figure 1.8b. These results combined, emphasise that some, but not all, PGs are characterised with a higher sensitivity to outcome changes across all domains. Figure 1.8a and 1.8b illustrate the utility functions for each group in the gain and loss domain.

Weighting function gains

With regards to the probability distortion parameter, γ_G, in the gain domain, on average the PG group had a lower mean and median values (0.548, 0.455) than the HG and C group (0.832, 0.852; 0.765, 0.75), suggesting that PGs distort subjective probabilities to a larger degree than the HG and C group, thus exhibiting a more pronounce S-shape in their weighting function. A difference that is statistically significant in both cases. With regards to the probability elevation parameter in the gain domain, we do not find many differences. The PGs have a slightly lower mean and median values (1.03, 0.837), than the HG (1.079, 1.045) and C (1.386, 1.193) groups, but this difference is only statistically significant between the PG and the C group.

FIGURE 1.8: **Mean CRRA plots.**

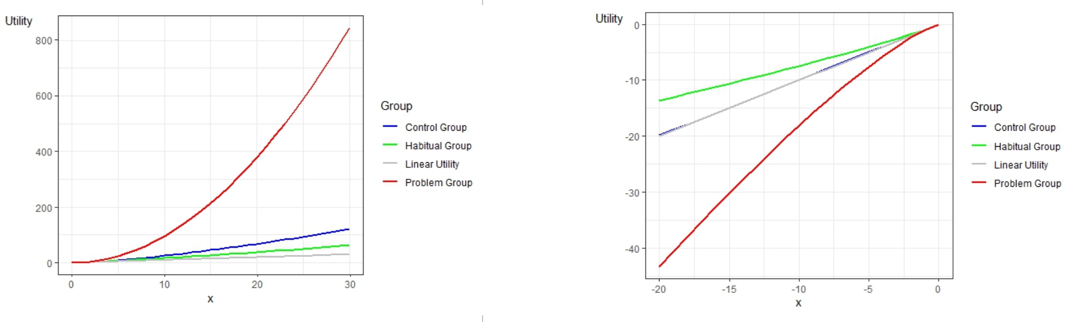

(A) Gain Domain

(B) Loss Domain

Notes: Utility function plots of the mean CRRA parameters for each group in the gain and loss domains. Note the result in the gain domain is not statistically significant, but we provide the plot to illustrate the extreme risk seeking behaviour of some of the PG group.

Weighting function losses

In the loss domain, when looking at γ_L, we find no systematic differences between the groups, which reinforces the results of Ring et al. (2018). However, we do find systematic differences between the groups when looking at δ_L. The problem gambler group tend to underweight the probability of a loss at all probability ranges, represented by a higher mean and median δ_L value (1.27, 1.29), than the HG and C groups (0.86, 0.84; 0.9, 0.91). This finding suggests that problem gamblers find the risks in the loss domain more attractive than non-gambling control groups. Figure 1.9a and 1.9b illustrates this result by plotting the mean parameter values for our weighting functions.

Comparing results to the Ring et al. (2018) analysis

Comparing our results to those found by Ring et al. (2018), we support their results regarding probability distortion in the gain and loss domain, however we find some interesting differences in other areas. Firstly, in the gain domain, whilst Ring et al. (2018) find that PGs overweight the whole probability scale, we find no systematic difference in probability elevation between the groups. On the other hand, in the loss

FIGURE 1.9: **Mean weighting function plots.**

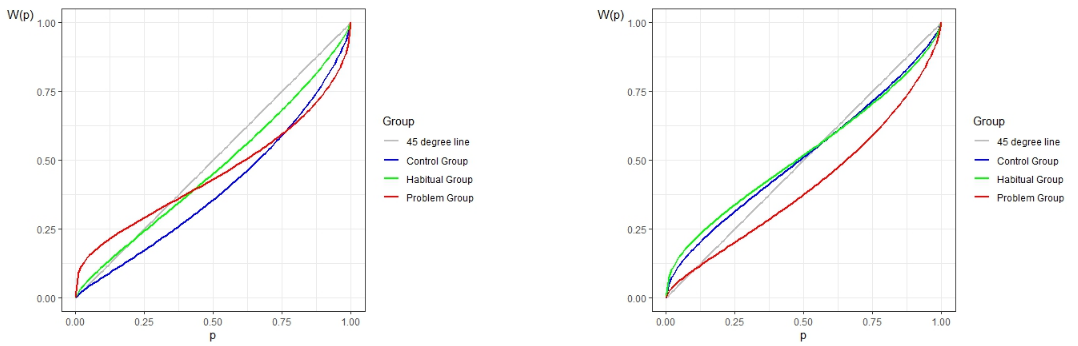

(A) Gain Domain

(B) Loss Domain

Notes: Weighting function plots for group means of the distortion and elevation parameters in the gain and loss domain.

domain, they find no difference in elevation, however, we find that PGs significantly underweight the whole probability scale.

These contrasting results are likely due to one of three reasons. (1) The Ring et al. (2018) parametric analysis was done using NLS, whilst we adopt maximum likelihood estimation routines. (2) They assume a representative agent for each group, whilst our analysis accounts for heterogeneity at the subject level. (3) They impose a linear utility assumption, meaning any subject level deviations from linear utility are likely to be picked up by a parameter in the weighting function. We relax this assumption.

1.7 Discussion

It is evident that problem gambling, as an impulse-control disorder, entails a great degree of ambiguity and irrationality in decision-making, therefore if we wish to understand what characterises their preferences, our modelling and estimation procedures, as well as our experimental designs, need to be exceedingly robust.

With regards to our functional form analysis, following our Maximum-likelihood estimation routine, accounting for overfit models using the AIC criterion, a cross-validation routine to understand predictive power, and two equivalence tests, we find that the majority of subjects, across all groups, are best characterised with either a variable power value function or its linear nested counterpart. Our parameter analysis and group-comparisons find that PGs are more sensitive to changes in monetary payoffs in the loss domain than non-gambling controls. Controversially, this suggests that PGs receive a greater disutility from increases in losses than non-gamblers. Intuitively, this makes sense, as the focus of gambling is to acquire a desirable gain, thus when dealing with pure losses, those that are more familiar with the feeling of losses, may be more inclined to try and avoid them. In the gain domain, we find that on average, the three groups do not differ much in their subjective utilities of varying monetary amounts, however we do find that a proportion of the PGs exhibit particularly high sensitivity to gain amounts. This heterogeneity within the PG group is not a surprising finding, as the drivers, and therefore characteristics, of the disorder are likely to differ between different PGs. In this case, it is likely that for some in the group, their decisions are highly influenced by their risk preferences, whilst for others, their disorder may be influenced by external factors. We also find that PGs, when engaging in risky decisions with pure gain amounts, distort probabilities more than non-gamblers. That is, they distort probabilities closer to 0 and 1 more. Additionally, in the loss domain, PGs significantly underweight the whole probability scale more than non-gamblers. We can think of this as a downward shift in the weighting function, and represents more risk-seeking behaviour. Abdellaoui, Vossmann, and Weber (2005) find there was more elevation for loss probabilities than for gains. A result that coincides with those of our HG and C groups, but opposes our findings for the PG group, which suggests that this variable may be an influential characteristic in problem gambling.

We are particularly interested in the results from the loss domain, as we find that PGs are more sensitive to outcome changes, but also underweight the probability scale. Again, whilst this may seem contradictory, when we disentangle the assumptions that the parameters impose, it provides interesting implications. Combined, these results suggest that PGs have a great disutility for monetary amounts, however, their subjective perceptions of the likelihood of these losses occurring is distorted, leading them to take more risks in an attempt to avoid the high loss, even though the disutility they would receive if they were to lose would be of a greater magnitude. A combination of loss aversion and risk-seeking behaviour may be driving their risky choices.

In addition to contributions in the realm of PGs risk attitudes, we provide a number of methodological contributions to the literature. Firstly, we emphasise the risks of imposing restrictive assumptions on structural models; namely linear utility. Although the majority of subjects exhibited linear utility, there were still a large number of subjects who's CRRA parameters deviated from one. Imposing this assumption therefore disregards the heterogeneity of preferences that different individuals have, and may result in this behaviour being picked up by the probability weighting parameters. Even though the experimental payoffs were particularly small, there were still relative increments which can pick up relative attitudes towards monetary increases. A flexible parameter approach allows the model to vary freely, thus not restricting a holistic analysis of behaviour. Especially given the volatility of gambling disorders, we see a flexible parameter approach as imperative, which is supported by the increased sensitivity to loss amounts that the PG group exhibited in our analysis.

Secondly, we recommend the use of subject-level estimation over imposing a representative agent assumption. Whilst this is still an ongoing debate in the literature (Kirman 1992; Nilsson, Rieskamp, and Wagenmakers 2011, pg. 87), due to the lack of information we have on the risk attitudes of PGs, assuming they are all characterised by the same decision-making process could be restrictive. In our analysis, especially in the gain domain, we find that there is significant heterogeneity within PG group. Pooling individuals will likely restrict the identification of within-group differences. Thirdly, we highlight the benefits of a parametric analysis in this context. Parametric approaches are able to account for the stochastic choices that individuals make, and can disentangle the risk-preference parameters that provide unique implications to subject's decisions. Finally, we highlight the importance of assuming different preference parameters in the gain and loss domain, as subject level preferences varied substantially across the domains.

On the discussion of experimental designs, Ring et al. (2018)'s introduction of a third control group (HG's), allowed for a more precise investigation into how psychological disorders influence choice. The identification and separation of regular gamblers from those with a disorder permitted the direct assessment of PG behaviour. To the best of our knowledge, it is the only available data set to have done so. It would be interesting to explore how the same treatment group set up would operate with higher payoffs and multi-outcome choice tasks, such as to gain a more comprehensive understanding of the underlying differences between the groups. Fehr-Duda and Epper (2012) suggest that the primitive nature of certain choice tasks may hinder a model's ability to disentangle utility curvature from probability weighting. Our results further highlight that we require designs that are also able to disentangle probability distortion from probability elevation. It may be that multi-outcome decision tasks are required in order to robustly disentangle the three variables.

As we wanted to avoid over-parameterisation, our analysis did not involve mixed tasks, therefore we were unable to directly estimate loss aversion. However, given our findings, it would be interesting to estimate a loss aversion parameter in a more holistic choice task, with multiple mixed gambles, to determine where significant differences prevail when gains and losses are at stake simultaneously. An analysis of gains and losses separately may only capture a fraction of the broader PG behaviour, as most real-world gambles are associated with potential gains and losses simultaneously. Nonetheless, we are still able to provide valuable insights into their decision-making processes.

The purpose of this research was to tackle both the methodological and policy implication issues involved in the study of pathological gambling. We therefore conclude by highlighting the implications that our results provide with regards to the psychological/ medical issue at hand. Takahashi et al. (2010) have delivered essential findings regarding individual risk preferences and the neuro-biological nature of the disease. Using positron emission tomography, they directly investigated whether dopamine D1 and D2 receptors in the brain are associated with the transformation of probabilities into decision weights. They found that individuals with lower striatal D1 receptor density show a more pronounced overestimation of low probabilities and underestimation of high probabilities. This could explain in part as to why pathological gamblers are more sensitive to changes in probabilities in the gain domain, thus providing an explanation for their irrational, overly optimistic nature in an actuarially unfair gambling setting. Moreover, D1 receptors, a natural regulator of neuronal growth and development responsible for the transmission of signals between neurons, are believed to be released when certain drugs are taken, and, when dis-functioning, can cause different diseases and disorders, including addiction (Mishra, Singh, and Shukla, 2018). The fact that our store of dopamine decreases with age also beckons for age as a variable to be included in future studies.

On the other hand, with regards to monetary rewards, Adinoff (2004) looks into the brains reward system, particularly the mesolimbic pathway, which plays a crucial role in processing rewards and punishment. The study suggests that addiction is associated with heightened sensitivity or dysregulation in this pathway. This heightened sensitivity might lead to increased responsiveness to gambling-related rewards/losses, contributing to the addictive behaviour, and providing an explanation as to why we find that PGs have an increased sensitivity to outcome changes loss domain, and why some PGs have extreme sensitivity to outcome changes in the gain domain. Interestingly, several neuro-economics studies have proposed that serotonin and dopamine affect the curvature of the CPT value function (see Berns, Capra, and Noussair 2007; Takahashi 2008; Zhong et al. 2009).

Finally, it is interesting to reflect upon our finding that PGs had an increased sensitivity to both outcomes and probabilities in the loss domain, suggesting indirect loss aversion with risk seeking tendencies. Due the loss domain not entailing any gains, there is no thrill or reward-seeking potential, and therefore they may be more loss averse than non-gamblers. Note, we are not implying gambling addiction is directly associated with a higher sensitivity to losses, as addiction is not a static problem, it comes from a dynamic sequence of events that leads to vulnerability. In a static environment, there can be no "loss chasing", which is well recognised as a driver of gambling addiction (Bibby, 2016), however what we can infer is that it is not loss chasing alone that characterises PG behaviour, as in a static environment, increased disutility prevailed for these individuals. Future research calls for the need to assess these preferences in a dynamic setting that includes mixed gambles such that we can estimate loss aversion and see how risk seeking behaviour manifests over time.

Our findings on the sensitivity problem gamblers have to monetary rewards allows us to provide meaningful policy implications. Increased regulation regarding reward mechanisms to limit the intensity or frequency of reward-based stimuli may reduce these effects. Similarly, marketing regulations on tighter restrictions on gambling advertising, especially concerning the use of reward-based messaging, could help early-stage addicts by mitigating this impact on potentially vulnerable individuals. Whilst with regards to our finding that PGs distort probabilities more in the gain domain, we advocate for policies that require gambling establishments to provide clearer and more transparent information about the probabilities of winning or losing, as well as for marketing to avoid the use of misleading or exaggerated claims that could reinforce distorted probability perceptions. Due to the subconscious nature of probability distortion, reinforcing the correct objective probabilities may mitigate some of the effect.

The combination of behavioural economics, psychology, and neuroscience has potential for ground-breaking discoveries regarding how and why PGs gamble. This analysis has filled a gap in the existing literature whilst creating new pathways for future research in terms of measurement procedures, elicitation tasks and a deeper understanding of the disorder. Risk preference analysis is able to compliment existing psychological findings in identifying the patterns of choices that PGs make, and how this may be influencing the addiction.

Chapter 2

Dynamic Inconsistencies in Risky Choice: Testing The Casino Gambling Model

Abstract

This chapter investigates dynamically inconsistent choices in decision-making under risk, based on the Barberis (2012) Casino Gambling Model (CGM). The CGM posits that individuals differ in their awareness of the discrepancy between their intended strategies and actual choices, and their ability to commit to these strategies. The model classifies subjects into two broader types: Naive, and Sophisticate. Our three-stage experimental design tests the model's theoretical predictions for each type, examining individual dynamic inconsistencies, awareness of self-control issues, demand for pre-commitment, and the impact of semi-binding pre-commitment strategies on financial welfare. While some subjects align with the model's theoretical assumptions, the majority exhibit behaviour that represents a merger of the original type classifications. We propose a new type classification and generate theoretical predictions for this group. We find higher levels of dynamically inconsistent behaviour than recent studies, both when commitment is, and is not, available. We show that financial welfare increases with commitment, but only for agents who would engage in the risky action regardless of commitment availability. Overall, our findings, along with those from existing studies, suggest that commitment devices fostering intrinsic motivation may be more effective than those imposed extrinsically.

2.1 Introduction

In the intricate landscape of human decision-making lies the everlasting discrepancy between professed intentions and subsequent behaviours. This inherently complex phenomenon is regularly referred to as the intention-action gap in the behavioural science and psychology literature. As humans, we have our own best intentions at heart, but for some reason, we have an innate inability to stick to ex-ante strategies or plans that we put in place to optimise our lives. We put in place a cognitive plan to eat healthily and exercise more, yet we find ourselves eating chocolate and binge watching our favourite programme. We express our intentions to want to save more, invest more wisely, or stick to a budget to achieve our financial goals, yet impulsive spending fosters an immediate deviation from this plan. The psychology literature has prioritised understanding the cognitive drivers of this behaviour, traditionally arguing that this "intention-action gap" is a result of self-control, temptation, and underlying emotional states (Gul and Pesendorfer, 2001). The economics literature, on the other hand, has sought to characterise these choices based on the expected utility of competing outcomes, as well as the subjective probability of achieving the desired goal, to provide context-based explanations of our inconsistent behaviour.

Traditionally, economic theory has been based on the normative assumption of rationality, suggesting the representative agent will act in exact accordance with their self-interest. This assumption of "Homo Economicus" has been largely critiqued (Thaler, 2000), with many empirical and experimental studies providing evidence to oppose these assumptions. Amidst the assumptions of classical economic theory lies the principle of dynamic consistency, suggesting that a decision-maker will follow through on a predetermined plan. Multiple experimental studies have provided evidence of violations of this principle (Cubitt, Starmer, and Sugden 1998; Busemeyer et al. 2000; Nebout and Dubois 2014), and empirical observations have shown that these violations play a large role in gambling and investment markets (Barberis and Xiong, 2009). This behaviour is defined as dynamic inconsistency (DI), reflecting a changing nature of economic agent's preferences over time (Moloi and Marwala, 2020), leading to deviations from initial strategies, and making one susceptible to intertemporally sub-optimal decisions.

This chapter studies dynamically inconsistent behaviour in decision-making under risk. When deciding whether to take on a risk, our decision is based on whether the expected outcomes are more desirable than doing nothing at all. In order to assess the expected outcomes, we seek to generate a decision-making strategy that generates a positive expected-outcome distribution. With this strategy in mind, we have cognitively minimised our potential losses, making the risky action more attractive. In a casino, an individual may enter with the intention of leaving if they lose 20% of their money, and continue playing if they are winning. With this strategy, the potential gains may outweigh the potential losses, motivating one to enter a casino. Similarly, an investor may purchase a stock with a strategy whereby they will liquidate the asset if it depreciates, and hold onto it if it appreciates. Again, this strategy generates a positive outcome distribution, leading to the purchase of the stock. In our investor market example, deviations from these trading strategies, also referred to as the disposition effect, represent a largely documented example of economic agent discrepancies between planned and actual choices (Shefrin and Statman 1985; Odean 1998; Weber and Camerer 1998).

This chapter tests the experimental predictions of a theoretical model that aims to explain this type of behaviour: The Barberis (2012) Casino Gambling Model (CGM). The model extends Cumulative Prospect Theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992) from its traditionally static environment, by applying its assumptions to a sequential choice problem to provide an explanation for why individuals gamble in casinos. It proposes that dynamic inconsistencies in risky choice arise as a result of probability distortion and loss aversion. Given CPT's explanatory power in a static context, it makes sense that its assumptions should hold in a dynamic environment (Hotaling and Busemeyer, 2012). The Barberis (2012) model categorises individuals into behavioural types who differ in their knowledge of their dynamic inconsistency problem, and their ability to find a means of committing to their initial strategy. Given an individual's CPT parameter set (utility curvature, probability distortion, loss aversion), the model predicts how different behavioural types will devise an ex-ante strategy, and how they will behave ex-post.

There are studies that have used features of the casino gambling model to elicit individual preferences for different types of strategies (Ebert and Voigt, 2023), whilst others have sought to identify if discrepancies between planned and actual choices exist using different dynamic choice tasks (Andrade and Iyer 2009; Barkan and Busemeyer 1999; Heimer et al. 2023). With regards to the commitment side of the model, studies have explored how individuals devise optimal stopping rules (Antler and Arad, 2023), and whether individuals are more likely to accept risk if they are able to commit to a strategy (Heimer et al. 2023; Hey and Paradiso 2006, Bettega et al. 2023). However, to the best of our knowledge, the theoretical predictions of the model have never been tested.

This chapter provides an experimental contribution which tests the predictions of the Barberis (2012) model. With 71 students from Lancaster University, our 3-stage experimental design aims to test the key assumptions of the model. Following a parameter estimation routine, we fit subject-level parameters to the various data-generating processes proposed by the model in order to identify behavioural types. We find, on average, that the original classifications struggle to identify structural differences between the types in various settings, and so we proceed by introducing a new classification that represent the experimental results from the study in question, as well as results from existing examinations. We find that subjects exhibit dynamic inconsistency, whereby they deviate from their initial strategies ex-post. Additionally, we find that there is a strong demand for pre-commitment. When this commitment is only semi-binding, subjects deviate from their ex-ante pre-committed strategies within the constraint of their commitment plan, exhibiting further dynamically inconsistent behaviour.

The rest of the chapter proceeds as follows: Section 2.2 discusses the theoretical framework, highlighting how CPT can be extended to a dynamic setting, and what this implies for the CGM. Section 2.3 provides our simulation methodology and results, and Section 2.4 details the experimental design. Section 2.5 explains our methodology for identifying behavioural types and estimating risk preference parameters in the Casino Gambling model. Section 2.6 presents our results, Section 2.7 introduces a new behavioural type, and Section 2.8 concludes.

For a comprehensive review of the literature, which disentangles theoretical contributions to dynamically inconsistent choices in risk-taking, empirical and experimental evidence of these inconsistencies, and intervention options (the use of commitment devices) in experimental and empirical settings that aim to mitigate axiom violations, please see Appendix A.1.

2.2 Theoretical Framework: Model of Casino Gambling

Our theoretical framework is based on the Barberis (2012) model of casino gambling which uses a dynamic version of CPT to identify behavioural types and predict how different types will make risky choices dependent on their CPT parameters. The model provides an explanation as to why certain individuals may enter a casino, and why they may continue to gamble in the face of losses. We will start off by explaining how CPT works in a static environment, such as to familiarise ourselves with how its core components provide an explanation for dynamic choices.

2.2.1 Static CPT

We now extend the CPT framework introduced in Section 1.4, by incorporating a loss aversion parameter. For our utility function, we impose a power utility function and find our first two parameters of interest: α and λ.

$$u(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda(-x)^{\alpha} & \text{if } x \leq 0 \end{cases} \tag{2.1}$$

where x represents a monetary amount, and α is our CRRA parameter, which determines how individuals value various monetary amounts, and is the parameter which forms our S-shaped utility function. λ is our loss aversion parameter, which incorporates the assumption that individuals are more sensitive to losses than equivalent gains. The lambda value magnifies the dis-utility of receiving a loss, where $\lambda \geq 1$. We assume the CRRA parameter α, is equivalent in the gain and loss domain, where $0 < \alpha < 1$.[1] Our final parameter of interest, δ, is found in the probability

[1] This is imposed to replicate the Barberis (2012) analysis.

weighting function which takes the functional form of that proposed in Tversky and Kahneman (1992).

$$w(p) = \frac{p^{\delta}}{(p^{\delta} + (1-p)^{\delta})^{\frac{1}{\delta}}} \tag{2.2}$$

where p represents the probability of receiving an outcome (e.g. x), and δ represents our probability distortion parameter, where $\delta > 0.278$ to satisfy monotonicity, such that our objective probability p is transformed into the subjective probability $w(p)$.

Combining the aforementioned elements, we get, for a single monetary amount, the CPT value function, $CPT = w(p) \cdot u(x)$. A CPT agent will use this function to assist their decision-making process, in that when making a choice on certain risky decisions, she will evaluate the CPT values of all the available options and choose the option that maximises her CPT utility. This is most commonly used in the experimental literature with binary choice tasks, where individuals have to make a choice between two specific lotteries.

For example, imagine an agent is asked to make a decision between 2 choices, A or B. In option A they have a 50% chance of receiving £10, but also a 50% chance of losing £8, and in option B they can choose to remain at £0 with 100% probability. These options, or lotteries, are written in the following notation $(10, 0.5; -8, 0.5)$. For option A the CPT value function would look like:

$$CPT = w(0.5)u(10) - \lambda(w(0.5))u(8) \tag{2.3}$$

If we assume that individuals have parameter values $(\alpha, \delta, \lambda)$ = (0.7,0.8,2), then the CPT utility of option A is CPT= -2.226, and the CPT utility of option B is CPT = 0. Therefore the individual in this case would have a preference for option B, even though the expected value of option A is greater. However, if we ask the same question, but to an individual who has parameter values $(\alpha, \delta, \lambda)$ = (0.9,0.6,1.1), then option A would now yield CPT = 0.663, and option B is still CPT = 0, thus preferring option A. Due to individual differences in how subjects value monetary amounts, their subjective perceptions of probabilities, and their distaste to losses, some individuals will engage in riskier activities than others.

Finally, with regards to the weighting of probabilities, it is important to note how probabilities are transformed by the weighting function when there are more than two options. CPT's rank-dependant probability weighting posits that to obtain a probability weight π_i, in the gain (loss) domain, we take the sum of the probabilities for all outcomes equal to or greater than (smaller than) x_i and apply the weighting function to each. We then we take the sum of the probabilities of all outcomes strictly greater than x_i, apply the weighting function to each, and then compute the difference.[2] Such that for a lottery $(10, 0.3; 5, 0.6; 0, 0.1)$, we would weight the functions accordingly:

$$\pi(x) = \begin{cases} w(0.3) & \text{for } x = 10 \\ w(0.3 + 0.6) - w(0.3) & \text{for } x = 5 \\ 1 - w(0.3 + 0.6) & \text{for } x = 0 \end{cases} \tag{2.4}$$

2.2.2 Dynamic CPT: Casino Gambling

Barberis (2012) extends CPT to a dynamic setting, and utilises its core assumptions to provide an explanation as to why certain individuals enter casinos and engage in gambling activity, even given the negative expected return associated with such games. The probability distortion component of CPT provides an explanation as to why individuals engage in buying lottery tickets, as the the overweighting of extremely low probabilities creates a positively skewed "game". However, Barberis (2012) highlights that it is not as straightforward to apply this argument to casino gambles due to their negatively skewed nature. Nonetheless, CPT's assumptions still hold.

The framework consists of an experimental casino with T + 1 dates, in which an individual must first make a decision as to whether they wish to enter the casino, and then make sequential decisions as to whether to accept or reject a gamble. We proceed with 5 decision periods (T=5), where in each period, the individual has to make a choice as to whether to accept or reject a 50:50 bet to win or lose a fixed amount £h. See Figure 2.1 below for a visual representation of the casino using a binomial tree. Figure 2.1b gives the visual representation of the tree, and figure 2.1a

[2] Note that we propose different weighting functions for gains and losses.

illustrates how we index the nodes of the tree. At time t=0, the individual is at the far left node of the binomial tree (0,1), and is offered the first 50:50 bet. If it is turned down, we say they have "declined to enter the casino" and the game is over. However, if they accept, then nature will decide whether they move up the tree or down the tree. If nature moves up and to the right (1,1), then they have won £h, but if nature moves down and to the right (1,2), then they have lost £h. We can think of this as a coin flip to determine whether you win or lose. At your new position in the tree, you are then offered the same 50:50 bet again; if you decline, then you will leave with either £h or -£h dependant on the previous move by nature, but if you accept the second gamble, then nature will decide whether to move up (down) and to the right, representing a further win/loss of £h. You then repeat the game until you either reach time T (the end of the tree), or you decide to leave the game at some date prior to T. For example, if you win three times in a row and then decide to leave (W,W,W), then your winnings will be £3h (at node 3,1). If you lose 5 times in a row (L,L,L,L,L), then your take-home will be -£5h (at node 5,6), if your outcomes are (W,L,W,L,W), then your take home would be £h (at node 5,3). Getting from t=0 to t=5 is a step by step process, moving from the far left of the tree to the far right of the tree, making a decision at each stage as to whether to accept or reject the 50:50 gamble.

Barberis (2012) models the tree with 5 decision periods to represent the finite time-horizon associated with trips to the casino - individuals may run out of money, or they may need to leave for other commitments; ultimately they will not remain in the casino forever, and this is captured by our period T. Each node is noted as a pair of numbers, "(t,j)" following the orginal notation. t ranges from 0 to T and corresponds to the time period. At the first node at the beginning of the binomial tree (0,1), t=0. j, on the other hand, ranges from 1 to t+1, and represents the vertical position of the node. The highest position on any vertical axis is always j=1. Then at t=1, once the individual has entered the casino and nature has played once, then j=1 if nature moved upwards, and j=2, if nature moved downwards. The lowest node in any column is therefore j = t+1.

According to the theory, a CPT agent will decide whether to enter or not based on the CPT value of their specified strategy. A crucial component of our analysis

FIGURE 2.1: **Binomial tree**

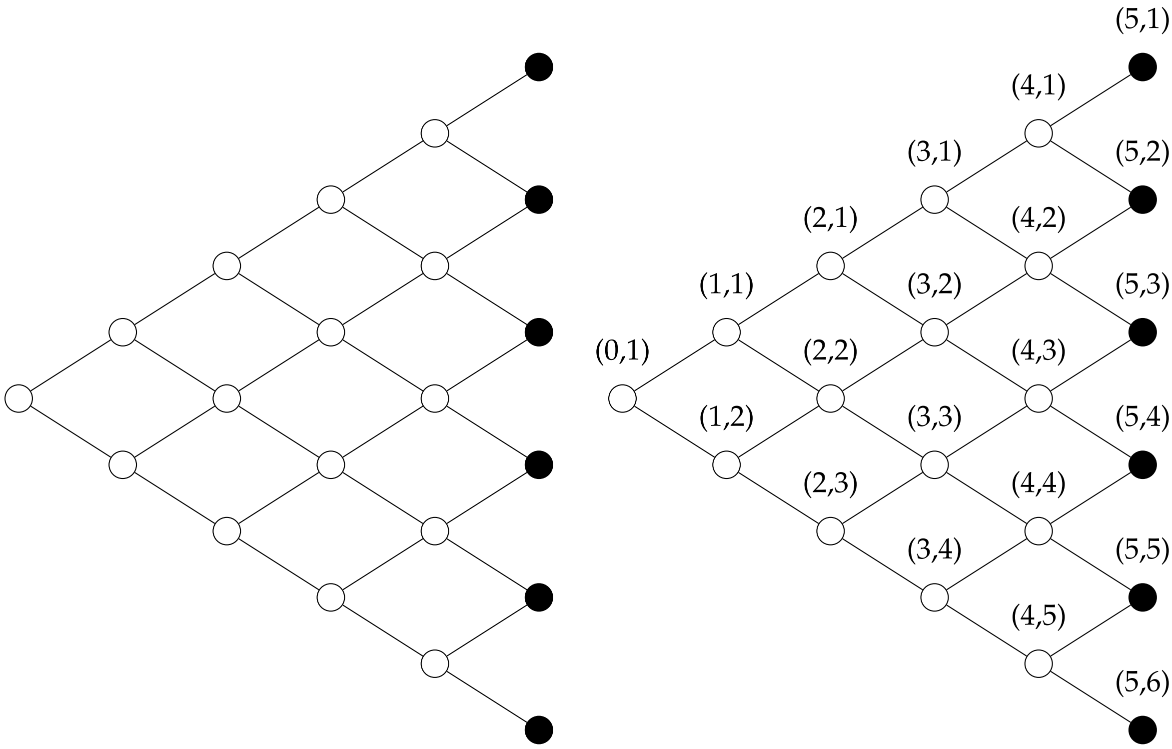

(A) Visual representation (B) Indexing

Notes: Panel A illustrates the binomial tree that represents 5 stages of a sequential decision-making game. Panel B shows how we index the tree.

is that individuals will devise a strategy that maximises the CPT value of their accumulating winnings or losses at the moment they decide to leave the casino. This plan (plan and strategy are interchangeable terms) is initially generated at t=0, and if they are able to find a plan that yields a CPT value that is greater than 0, then they will enter. Of course this is dependent on their parameter values. We will now break each component down, as each plays a role in explaining dynamically inconsistent choices. Let us start by visualising the payoffs from the binomial tree as a matrix, for simplicity. See Table 2.1.

There are in total 11 rows to represent payoffs [-10,10] including zero, and 5 columns, representing the time period in which a payoffs can be earned. We use $h = 2$ to demonstrate as this was the amount used in our experiment. Since the payoff can go either up or down, we set the zero point at the middle of the matrix, and we allow moves up and down the tree. The highest payoff is therefore in row 1, and

TABLE 2.1: **Payoff matrix**

Row	Payoff	t=0	t=1	t=2	t=3	t=4	t=5
1	10	0	0	0	0	0	10
2	8	0	0	0	0	8	0
3	6	0	0	0	6	0	6
4	4	0	0	4	0	4	0
5	2	0	2	0	2	0	2
6	0	0	0	0	0	0	0
7	-2	0	-2	0	-2	0	-2
8	-4	0	0	-4	0	-4	0
9	-6	0	0	0	-6	0	-6
10	-8	0	0	0	0	-8	0
11	-10	0	0	0	0	0	-10

Notes: The payoff matrix illustrates the payoff an agent would receive if they were to leave at the respective node.

the lowest payoff in row 11. Following the same structure as the binomial tree, negative payoffs are in the lower part of the matrix, and the positive payoffs are in the higher part of the matrix. For any given path that nature can make e.g. L,L,W,L,W (Lose,Lose,Win,Lose,Win) the row number increases by 1 if L and decreases by 1 if W, with t increasing by 1 in each period. The higher the row number, the lower the payoff. We generate paths of length 5 such that there are $2^5 = 32$ potential paths that nature can select from. Now, for the example path (L,L,W,L,W), we can formulate Table 2.2.

In this example, and given that the agent's strategy is to play at every node, then we can see that they will end at [Row, t] = [7,5] in Table 2.1, generating a payoff of -2, which is equivalent to node (5,4) in our binomial tree in Figure 2.1b.

However, this is assuming that the individuals plan is to play at every node, which is unlikely. Remember, individuals will elicit a plan that maximises the CPT

TABLE 2.2: **Path and payoff example**

Move	Row No.	t	Cum. payoff
L	7	1	h[7,1]=-2
L	8	2	h[8,2]=-4
W	7	3	h[7,3]=-2
L	8	4	h[8,4]=-4
W	7	5	h[7,5]=-2

Notes: The example illustrates the path-dependent cumulative payoff received throughout a game given the respective moves by nature.

TABLE 2.3: **Decision-nodes**

Row	t=1	t=2	t=3	t=4	t=5
1	0	0	0	0	
2	0	0	0	4	0
3	0	0	3	0	
4	0	2	0	8	0
5	1	0	7	0	
6	0	6	0	11	0
7	5	0	10	0	
8	0	9	0	13	0
9	0	0	12	0	
10	0	0	0	14	0
11	0	0	0	0	

Notes: Nodes where individuals are required to make a decision

value of their accumulated winnings or losses at the moment they leave the casino. Let us consider the matrix corresponding to the decision nodes in Table 2.3.

If the individual decides to enter the casino, then there is a possibility that they will make a decision at any of the 14 decision-nodes in Table 2.3. They will therefore devise a strategy in which they will either play, or leave, at each of the 14 nodes. Table 2.4 represents a strategy matrix. Let us denote this matrix as s. This is an 11x4 matrix where the decision-nodes take values 1 or 0, representing "play" and "leave" respectively. We now ignore the 5th column as no decision is to be made here as they have to leave the casino at time T, and we ignore the entry decision as the entry decision is based on the CPT value that will be generated from the optimal strategy, that is, they will enter if they find a strategy, based on decisions to be made at each of the 14 nodes, that yields a positive CPT value.

From the strategy in Table 2.4, in any node where there is a 1, then they would continue playing if they ever arrive at this node, and in the nodes where there is a 0, they would leave the casino with their accumulated winnings up until that node. However, Table 2.4 only shows one specific strategy, and given that there are 14 potential decision-nodes, excluding the entry decision, and at each decision node there are two actions they can make (play, leave), we have a total of 2^{14} potential strategies. This equates to 16,384 potential strategies. However, many of these strategies are not feasible. For example, if in position [8,2] of Table 2.4, their strategy is to leave (marked with a 0), then any strategies that have 1s in positions [9,3] and [10,4]

TABLE 2.4: **Strategy Matrix "s"**

Row	t=1	t=2	t=3	t=4	
1	0	0	0	0	
2	0	0	0	1	0
3	0	0	1	0	
4	0	1	0	1	0
5	1	0	1	0	
6	0	1	0	1	0
7	1	0	1	0	
8	0	0	0	0	0
9	0	0	0	0	
10	0	0	0	0	0
11	0	0	0	0	

Notes: Matrix form example of 1 out of the 801 feasible strategies a subject can elicit.

are not feasible, as the agent has already decided to leave the casino prior to these nodes. Therefore, after we eliminate all non-feasible strategies, we are left with 801 unique and feasible strategies. According to Barberis (2012), an agent calculates the CPT value of all potential strategies and chooses the one that maximises their utility. If there is a strategy that provides an outcome distribution with a CPT value that is greater than 0, then they will enter the casino and plan to play according to that strategy.

However whether or not they stick to this plan is another matter, where not sticking to their initial plan represents a violation of the dynamic consistency axiom; a departure that can be rationalised by Cumulative Prospect Theory. For example, let us suppose that h=2. An individual at t=0 makes a plan as to whether they will play or leave at each node. Now lets focus on the node (4,5) of the binomial tree as highlighted in red in Figure 2.2 as an example (in Table 2.1 this is position [10,4]).

From the perspective of $t = 4$ at node (4,5), an individual deciding to leave the casino will leave with -£8. If they choose to play, there is a 50% chance of ending with -£10 at node (5,6) and a 50% chance of ending with -£6 at node (5,5).

From the perspective of $t = 0$ (0,1), the individual devises a plan considering the probabilities of reaching different nodes. The probability of reaching the least desirable node (5,6) is 1/32, and the probability of reaching node (5,5,) is 5/32. Additionally, the probability of arriving at node (4,5) is 1/16. When planning at $t = 0$, the individual evaluates the choice to play or leave at node (4,5) using equations 2.1

FIGURE 2.2: **Binomial tree decision at node (4,5).**

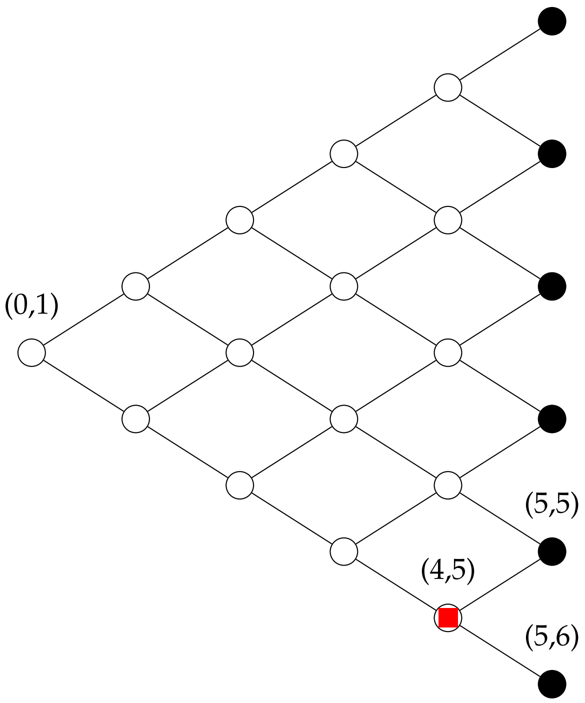

and 2.2 to maximise their CPT utility. CPT suggests that individuals overweight low probabilities in a rank-dependant manner, making the low probability of -£10 subjectively higher and less appealing. Consequently, many individuals with high probability distortion parameters would plan to leave at node (4,5) because CPT(Exit) > CPT(Play).

However, upon actually reaching node (4,5) the probabilities change to 50% for each outcome. This higher probability is no longer overweighted, leading the agent to re-evaluate the decision. At this point, for many parameter values, CPT(Play) > CPT(Exit), making gambling more appealing than initially planned. The discrepancy highlights dynamic inconsistency: the individual's ex-ante strategy at t=0 was to exit at node (4,5), but their ex-post choice at t=4 is to continue gambling due to the changed perception of probabilities.

The analysis of dynamic inconsistencies required two pieces of information: first, the ex-ante plan at t=0 of an agent, and second, their ex-post behaviour. From this, we can start to make assumptions about the heterogeneity of individual behaviour. A primary contribution of the CGM is that it utilises the identification of behavioural types to explain how and why some individuals enter casinos, why some individuals spend too much time in there, and why some do not enter at all. Following the

Barberis (2012) specification, we proceed with two behavioural types, who differ in their awareness of their time inconsistencies, and their ability to commit to an initial plan. The first behavioural type is the "naive" individual, who are not aware of their time inconsistencies, and believe that their actions will follow their initial t=0 plan, but when proceeding with the game, they deviate from this plan ex-post. The second is the "sophisticate" who is split into two based on whether pre-commitment is avaiable. The "sophisticate without commitment", unlike the naive, is aware of their time-inconsistency, however, they are unable to find a way to commit to their initial strategy, leading to the use of backward induction to develop a strategy at t=0 that reduces the potential side-effects of their inconsistency. They follow their backwards induction t=0 plan. The "Sophisticate with commitment" is aware of their time inconsistencies, but is able to find an exogenous means of committing to their initial strategy. We assume this is with an external commitment device. Therefore their t=0 strategy is forward looking and will follow the same pattern as a naive individual with the same risk preference parameters. The main difference is that, because of their exogenous commitment device, they are able to follow their t=0 strategy in all periods t>0.

Whether an individual enters the casino, and how they decide to play, is dependant on their behavioural type and their risk preferences, therefore within-group heterogeneity will still prevail for individuals with different risk preferences. Prior to being able to identify types and estimate CPT parameters, we need to know for which parameters will each of the three types enter the casino in the first place. As we are directly testing the assumptions of the CGM, we initially follow the same parameter structure. The preference parameters are therefore those aforementioned in equations 1 and 2, and are $\alpha \in [0,1]$, $\delta \in [0.3,1]$ and $\lambda \in [1,4]$. For our simulation to determine the range of preference parameter triples $(\alpha, \delta, \lambda)$, we separate the parameter intervals into 20 equally spaced points, such that there are $20^3 = 8000$ potential parameter triples. For each of the parameter triples we evaluate whether the agent is able to find a plan for which their CPT value is greater than 0. We will now expand by showing how each of the behavioural types, dependant on their preference parameter triples, devises an ex-ante plan, decides whether they or not to enter, and how they play ex-post.

FIGURE 2.3: **Loss exit strategy example.**

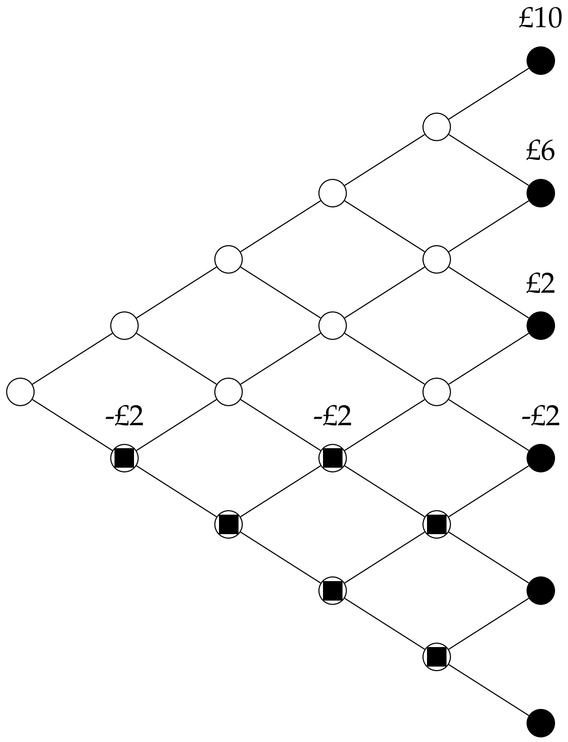

Naive

At t=0, a naive agent goes through all potential plans of action, and selects the plan that maximises their CPT utility. We use the notation $S_{(0,1)}$, to denote the set of all possible plans at node (0,1) on the binomial tree. For each plan $s \in S_{(0,1)}$, we use the random variable G_s to represent the accumulated winnings or losses that the individual faces by following the specific plan. For example, if we have the exit strategy (plan) shown in Figure 2.3 where the black nodes represent the individuals exit nodes, then we can solve G_s by filtering through the possible payoffs and respective probabilities, and sub these into our CPT equations 2.1 and 2.2, to determine the CPT value of this strategy. Figure 2.3 represents a single example of a loss-exit plan, suggesting that if the individual starts losing, their plan is to exit, and if they start winning, they plan to continue. We expect this to be a commonly used plan by our subjects. Again we assume that T=5 and that h = 2. Therefore, from this plan, the payoffs they could leave the game with are (£10, £6, £2, -£2). However, all outcomes do not have the same likelihood of occurrence. Whilst some outcomes, e.g. £10, can only occur with 1 path by nature (W,W,W,W,W), other payoffs, e.g. £6, can occur with more than one path (W,W,W,W,L), (W,W,L,W,W),(W,W,W,L,W), and

(W,L,W,W,W). Additionally, some payoffs, e.g. -£2, can occur by the game ending at more than one of the exit nodes (1,2), (3,3) and (5,4). We therefore simulate the probability of each outcome occurring, given the strategy[3], and yield the accumulated winnings or losses as: G_s $(10, 1/32; 6, 4/32; 2, 5/32; -2, 22/32)$, which corresponds to the following CPT function:

$$\begin{aligned} CPT(G_s) = \ & u(10)w(1/32) \\ & + u(6)(w(5/32) - w(1/32)) \\ & + u(2)(w(10/32) - w(5/32)) \\ & + u(-2)w(22/32) \end{aligned}$$

Where the negative outcome has it's own weighting function, and is not ranked with the positive outcomes. The individual therefore solves equation 2.5 for each of the 801 potential plans, and will enter the casino if they are able to find a plan that creates a gambling experience where the CPT value is greater than 0, given their preference parameter triple.

$$\max_{s \in S(0,1)} \mathrm{CPT}(G_s) \tag{2.5}$$

To determine the number of preference parameter triples for which a naive decision-maker will enter the casino, we use simulations.[4] Running through each of the 801 plans, we calculate the CPT value of each plan for each of the 8000 parameter triplets. Remember, an important concept is that, because of the weighting function, individuals are able to generate positively skewed gambling experiences from a specific plan.[5] Following the simulation, we find that a naive agent would be willing to

[3]To do this, we generate paths of length 5 and at each step we check whether the motion went up or down, calculate the cumulative payoff and check whether according to their plan if they want to continue or exit the casino. Dependent on this plan, we record the cumulative payoff at the stopping node, for that particular path. We repeat the process 10,000 times and are able to generate an empirical distribution which approximates the probability or reaching a particular decision node. Using the empirical distribution, we can generate the gamble that corresponds to a particular plan. This is easy to do manually for one gamble, however given there are 32 potential paths, and 801 potential plans, a simulation is required to yield this for each plan

[4]Our simulation generates a matrix, with 16 columns, and 801 rows, where each row represents the gamble that has been generated from a specific plan.

[5]The type of optimal plan they generate is dependant on their parameter triples, and our simulation determines the number of parameter triples for which a naive agent enters, and what plan they enter with (e.g. gain-exit or loss-exit).

enter the casino for 1792 out of the possible 8000 parameter triples considered. [6]

Interestingly, even under the assumption that the naive individual is risk averse, loss averse and the casino only generates 50:50 bets with zero-expected value, there are still a wide range of parameter values for which the agent is willing to enter the casino. We can now classify plans into gain-exit and loss-exit. As aforementioned, a loss-exit (gain-exit) plan is any in which the expected length of time in a casino, conditional on exiting with a loss (gain), is less than the expected length of time in the casino conditional on exiting with a gain (loss). The choice of plan is dependant on an individuals preference parameter values. For example, an individual with $(\alpha, \delta, \lambda)$ = (0.2,0.9,1) receives the highest CPT value by generating a gain-exit plan, whereas an individual with $(\alpha, \delta, \lambda)$ = (0.9,0.6,1.2) maximises their CPT value with a loss-exit plan. Of the 1792 parameter triples for which a naive agent will enter, 1110 generate a loss-exit plan, and 682 generate a gain-exit plan.

Once a naive enters the casino, they follow the same procedure for all of the available nodes. As the tree shortens, the number of permissible plans decrease, and the probability of reaching specific nodes change, therefore naive agents will likely deviate from their t=0 plan and devise a new strategy that is optimal from their current time-period.[7]

For example, let us take the parameter triple (0.9, 0.6, 1.2) again. With these preference parameters, a naive agent maximises their CPT utility when choosing the generated loss-exit plan in Figure 2.4a. We can see that as a result of a high

[6]Note that Barberis (2012) finds that the naive agent enters for 1813 triples. Our results differ for two reasons. Firstly the Barberis (2012) theoretical analysis is based on h=10. As we run an incentivised experiment, having a maximum payoff per subject of £50+ is not financially feasible, so we proceed with h=2. We therefore run our simulation under these figures. However,when we replicate our simulation with h=10, we get that the naive agent enters for 1763 parameter triples. Secondly, we account for a special case of when $\alpha = 0$, as in this case, any increment in payoff would simply yield a utility of 1, which results in a non-increasing utility function. For when $\alpha = 0$, we take the log of the payoff to ensure utility is increasing. Nonetheless, when we replicate the simulation with h=10, without taking the log of utility when $\alpha = 0$, we find the naive agent enters the casino for 1813 parameter triples and are able to replicate the original analysis.

[7]At time t=1, the individual has entered the casino, and are now at node (1,1) or (1,2). From the perspective of either of these nodes, there are 9 possible future nodes in which they may need to make a decision, therefore at node (1,1) for example, there are $2^9 = 512$ potential strategies/plans the agent will consider when deciding upon whether to keep playing or stop. After eliminating non-feasible nodes, at either of the nodes at t=1 there are 96 potential strategies to consider. If they are able to generate a new plan, from the perspective of their new position in time t, that yields a CPT value that is greater than the CPT value they would yield from leaving at that node, then they will play. At t=2, in node (2,1) for example, there are 5 possible decision-nodes following, so they consider $2^5 = 32$ possible plans, which becomes 17 potential strategies after eliminating non-feasible nodes. It is easy to see there are are 4 possible plans at each of the t=3 nodes, and once we get to t=4, we are left with a one-shot 50:50 gamble.

FIGURE 2.4: **Loss exit strategy vs actual play.**

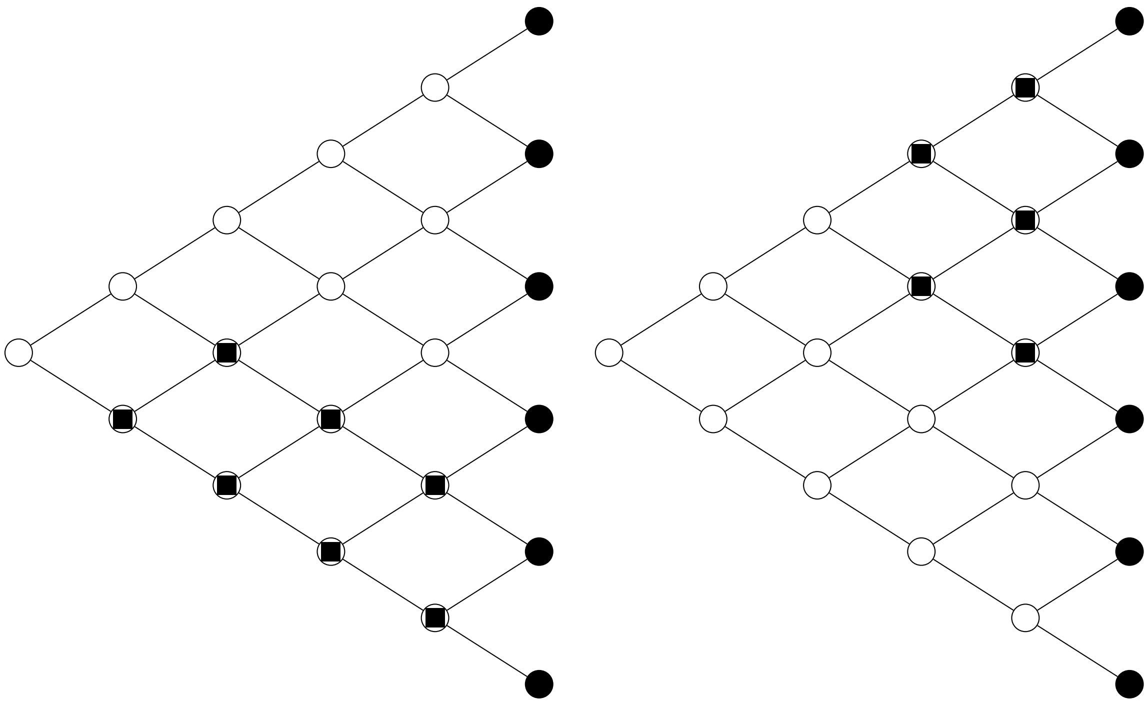

(A) Loss Exit Initial Plan

(B) Actual Play

Notes: Panel A represents the Loss exit strategy generated by, $(\alpha, \delta, \lambda)$ = (0.9,0.6,1.2). Panel B represents the models predictions of ex-post actions given the same parameter triple. Remember a white node indicates play, and a black node indicates exit.

α, marginal utility does not diminish excessively, so the utility received from the higher outcomes makes the prospect desirable. There is also a low δ, such that this subject overweights the tails of the probability distribution to a large degree, suggesting their subjective likelihood of receiving the desirable higher amount is increasing, again making the prospect even more attractive. Finally, λ is relatively low, so the disutility of experiencing a small loss will not suffice as to deter the individual from seeking out the subjectively likely desirable gain. Therefore, with the positively skewed outcome distribution generated from their plan in Figure 2.4a, they decide to enter. However, as they are unaware of their dynamic inconsistency problem, what they actually do ex-post is illustrated in Figure 2.4b, which represents a gain-exit strategy, even though they planned to follow a loss-exit strategy.

On the other hand, an individual with preference parameter triples (0.2,0.9,1) maximises their CPT utility with a gain exit plan, as illustrated if Figure 2.5a. As you can see, they are deciding to enter with a plan that has a negatively skewed outcome distribution. This may sound counter-intuitive, but their parameter values makes this appealing. As the α is relatively low, it implies the marginal utility they receive from higher outcomes diminishes rapidly, implying a large loss does not decrease utility significantly, and a large gain is not much more desirable than a small gain. Similarly, the low λ means they do not fear the losses excessively. A high δ implies the low probability of a large loss is barely overweighted, and the low probability of a larger gain is not that desirable. They therefore find this attractive as the highly likely small potential gain is desirable, and they are not put off by a small loss. We can interpret this as a subject wanting to take any gain they can and then leave, but also that they do not fear losses, so they will keep gambling until they reach that marginal gain. Again, however, when they enter, they may deviate from this plan if their updated preferences do not align with their initial strategy. In our example, their actual play will look like Figure 2.5b, where they only deviate at one point. Other parameter triples may lead to more extreme deviations, and some may lead to none at all.

Sophisticate Without Commitment

As defined previously, a sophisticate without commitment is an individual who is aware that they violate the axiom of dynamic consistency, and whilst it is in their best intention to try and commit to their t=0 plan, they are unable to find a means of doing so. As a result, this individual uses backward induction to devise a strategy, as it is the only way they will be able to generate a plan that they know they can stick to. This is because, due to the backward induction procedure, the individuals plan will represent what they "know" they will do in future rounds.

To identify the sets of parameter triples for which a sophisticate will enter the casino, and to determine the type of plan they follow, we follow a similar exhaustive search approach as that for the naive individual, except with two main differences: firstly, the process now begins at the final nodes in t=4 and rolls back to the beginning. This represents the backward induction approach, rather than the forward-looking approach of the naive individual. Secondly, the set of available gambles is now parameter specific. That is, instead of generating all possible plans and identifying which parameter triples give these plans a positive CPT value, we build the

FIGURE 2.5: **Gain exit strategy vs actual play.**

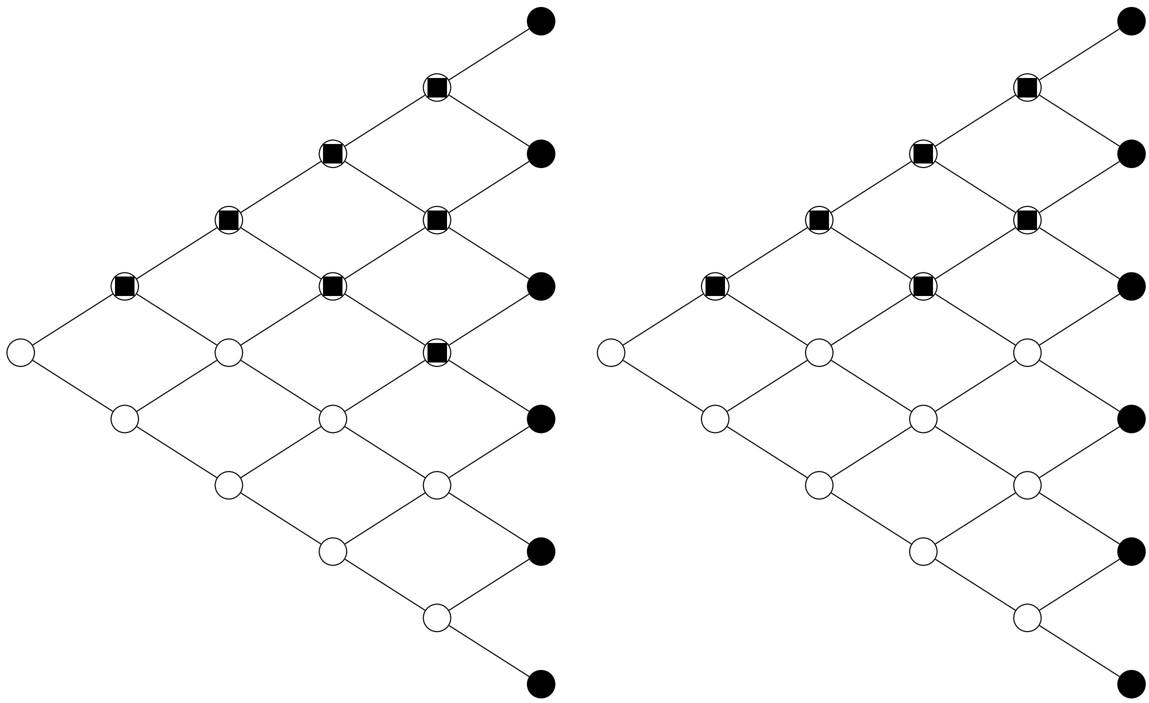

(A) Gain-exit Initial Plan

(B) Actual Play

Notes: Panel A shows a Gain exit strategy generated by, $(\alpha, \delta, \lambda) = (0.2, 0.9, 1)$. Panel B represents a subjects ex-post actions with the same parameter triple. Remember a white node indicates play, and a black node indicates exit.

strategy profile from the end of the tree to the start of the tree for a given parameter set. In other words, we identify for each set of parameters, whether the individual will gamble or play at each node, starting from t=4, and work backwards to build the optimal plan.

Using the notation from 2.3, we can think of the sophisticate without commitment's decision-making process to devise a plan as a 5 step process. In step 1, the decision maker has to solve the t=4 gambles at nodes 4,8,11,13,14. The gamble at node 4 is a 50:50 gamble between £10 and £6, against a certain £8. At node 8, it is a 50:50 gamble between £6 and £2, against a sure £4, and so on. The decision-maker, given her parameter triples, identifies whether they will continue or leave, makes a note of this decision, and then moves back one period.

In step 2, in t=3, specifically nodes 3,7,10 and 12, she considers both decisions, conditional on her choice in step 1. For example, at node 3, the individual now makes a decision as to whether to continue or gamble, but their CPT value of continuing or leaving is conditional on the choices she knows she will make at nodes 4 and 8. This individual therefore "prunes the tree" and makes a decision. By pruning, we mean removing branches that she knows she will never reach due to her knowledge of her future actions. If for example, this individual knows they will leave at node 4, but continue at node 8, then their gamble becomes $G = (8, 1/2; 6, 1/4, 2, 1/4)$. In step 3, they move back one period again, and now at t=2, they follow the same procedure, pruning the trees conditional on their decisions in t=3 and t=4, and identifying their optimal strategy at nodes 2,6 and 9. In step 4 they do the same thing again at nodes 1 and 5, and in step 5, they identify their optimal choice at node 0. For each parameter triple, there will be a single optimal plan. As the individual uses backward induction, their decisions in actual-play will be time-consistent, that being, they will follow their t=0 strategy.[8]

Once the individual has completed the backward induction process, they will have created a mapping of available options given their awareness of their dynamic inconsistencies. Therefore, for each of the 8000 parameter triples, there is a single plan. For 646 of these parameter triples, the plan generates a positive CPT value, and will therefore incentivise the sophisticate to enter the casino under their optimal plan.[9] For all parameter sets for which the sophisticate without commitment enters, they follow a gain-exit strategy.

Sophisticate With Commitment

The final behavioural group which we examine is that of the Sophisticate with commitment. This individual takes properties from the Naive group and the Sophisticated without commitment group. On the side of the latter, this individual is aware of their individual time inconsistencies, that is, they know, without assistance, their ex-ante strategies will not coincide with their ex-post actions. The key difference is

[8] The main difference between the naive individual and the sophisticate, is that the naive individual considers every available plan and chooses the plan that maximises their CPT value, whereas the sophisticate works backwards and prunes the trees based on their triples, and is left with a single plan. Our simulation generates the corresponding gambles at each node, conditional on individual choices in future periods.

[9] As with the naive, we receive a different number of entry triples than Barberis (2012), who receives 753 triples. The reasons for this are explained in the footnote under the naive section, attributed a different payoff scale, and a log transformation for when $\alpha = 0$.

that the Sophisticate with commitment is able to find an exogenous means of committing to their initial strategy. The literature refers to this means as a commitment device. At t=0, this individual therefore solves the same problem as the Naive individual, using a forward looking approach to identify the set of strategies, for their preference parameter sets, that generate a positive CPT value. That is, they solve equation 2.5. They will enter the casino if they identify a strategy that generates a positive CPT value over the game, and due to the exogenous commitment device, they are time consistent and follow this strategy ex-post.

As their t=0 problem is identical to that of the Naive, the sophisticate with commitment enters the casino for the same number of parameter triples as the Naive, and for the same strategy type, that is, they enter the casino for 1792 of the 8000 parameter triples, and for 1110 of these they follow a loss-exit strategy, and for 682 of these they follow a gain-exit strategy. However, when the naive agent proceeds by deviating from this strategy in t>0, the commitment device assists the sophisticate with commitment in following this strategy.

The sophisticate with commitment (SC) differs from the sophisticate without commitment (SNC) in their initial strategy elicitation. Without commitment, a sophisticate uses backward induction, mapping their future decisions such that they can create a strategy, conditional on what they know they will do in the future. With commitment, the sophisticate is able to adopt a forward looking approach, maximising their expected outcome distribution as they can rely on the commitment device to ensure the strategy is carried out. With commitment, the sophisticate will therefore enter the casino for more parameter triples than without commitment. Note that although we imply that Naive, Sophisticate with commitment, and Sophisticate without commitment, are 3 types of individuals, this is for theoretical purposes. In reality, the sophisticate is a single classification, and whether they commit or not is dependant on the availability of a devise. When a devise is available and free to use, we assume the sophisticate will always opt for the commitment device.

For clarity on commitment devices, this can be any exogenous means of ensuring execution of a plan. In the context of gambling in a casino, this could be leaving your debit/credit card at home and taking a set amount of cash, essentially setting an upper bound on how much can be lost. It could be that you arrange to be picked up at a set time, imposing a social commitment, and reducing the time spent in the casino. In our case, we give individuals the choice as to whether or not they would like to enforce their initial strategy, that being, they elicit an ex-ante strategy by determining, at each of the nodes in Table 2.3, whether they would like to continue or exit at this node, and if they choose to enforce this strategy, they will automatically be removed from the game if an "exit" node is reached.

2.3 Simulations

Prior to estimating parameters and identifying types given the predictions of the theoretical model, we first generate synthetic data to see if we can robustly recover simulated parameters. Simply put, if we adopt the design imposed by Barberis (2012), and generate synthetic data for a parameter set b_i and model m_i, are we able to recover parameter set b_i and classify them as type m_i with the design proposed in the model?

2.3.1 Simulation Methodology: Trembling Binary Choice Model

We now proceed to fit a stochastic structure to the casino task in order to identify types and estimate CPT parameters. Utilising the assumptions specified in Section 2.2.2 for each types structural behaviour, we use simulations to identify an agent's actions dependent on their type and parameter triple. That being, how each type will generate an ex-ante strategy, whether or not they will enter, how they will play ex-post, how this differs when given the opportunity to commit, and whether they opt for commitment.

In order to assess the feasibility of model identification, we need to develop a stochastic structure to the deterministic decision-making theories, to account for stochastic behaviour (Stott, 2006). We consider two potential stochastic specifications. The first, which assumes a constant error term, also known as a tremble, was introduced by Selten (1975) when discussing the trembling-hand equilibrium, and has been extended to the "constant error" model of Harless and Camerer (1994) and Wakker, Erev, and Weber (1994). This assumes a decision-maker will make the correct decision with probability $1 - \omega$, and with probability ω, they make a mistake.

We can interpret this error as a subject temporarily losing concentration at the time of solving the decision problem, and therefore choosing randomly (Loomes, Moffat, and Sudgen 2002; Moffatt and Peters 2001).

The second stochastic specification assumes agents make mistakes via some probabilistic rule, where the mistake is a result of a numerical error made in the computation of the differences between the two or more competing outcomes. This can be modelled with a Probit transformation (Hey and Orme 1994; Carbone and Hey 2000).

Both stochastic specifications provide important behavioural implications regarding potential error stories in decision-theory modelling. Although Carbone (1997) finds evidence to reject the constant error approach, Moffatt and Peters (2001) introduce the idea of a combination of the probabilistic rule with a tremble, and find that following a probabilistic rule, whilst ignoring the tremble, resulted in an upward bias estimate of the computational error parameter. They suggest that any subject-level mistakes arising as a result of a "tremble" is picked up by the error term in the probabilistic rule (we call σ). Loomes, Moffat, and Sudgen (2002) tested this hybrid specification on existing data sets and found it improved the explanatory power of the decision-theory models.

A probit transformation is feasible for the naïve's entry decision and ex-post choices, as well as the sophisticates ex-ante plan, entry decision, and ex-post choices. However, given the assumptions of the model, it is not possible to determine a probabilistic rule for the naïve t=0 ex-ante plan, or the sophisticate with commitments t=0 ex-ante plan and ex-post choices. This is because the theory predicts that these subjects do not make node-level decisions when selecting their strategies, instead, they filter through all 801 potential strategies, and select the strategy that maximises their CPT value. For this reason, we are unable to determine the probability of a subject selecting a specific exit-node, but rather we determine the probability of entry based on the CPT value generated from the optimal strategy. As such, for these node-level decisions, we impose a constant-error rule. For the decisions that permit a probit transformation, the constant error is combined with a probabilistic rule. To implement the stochastic specification, we follow the Moffatt and Peters (2001) stochastic "trembling binary choice model".

The probabilistic rule follows the probit transformation function, whereby we use the cumulative distribution function (CDF) of the standard normal distribution to determine the probability of a subject entering or playing. This transforms the deterministic choices (enter/don't enter, play/exit) into a probability between 0 and 1, depending on the differences in utility between the competing binary options. The cumulative normal distribution maps the differences in CPT utility to to a probability of choosing one option over another. Our binary choice error, ϵ, is normally distributed with mean 0 and standard deviation, σ, $\epsilon \sim \mathcal{N}(0, \sigma^2)$.

If we assume the variable y is a binary dependent variable representing whether the subject played or exited, where $y = 1$ indicates the individual plays, and $y = 0$ indicates they exit, then equation 2.6 is the probit transformation model.

$$
\begin{aligned}
&P(y_i = 1|\beta_i) = \Phi(x_i, \beta_i) \\
&P(y_i = 0|\beta_i) = 1 - \Phi(x_i, \beta_i) \quad \text{where } x_i = \frac{CPT_{\text{enter}} - CPT_{\text{exit}}}{\sigma} \quad \text{and } \epsilon \sim N(0, \sigma)
\end{aligned}
\tag{2.6}
$$

for $i = 1, \ldots, n$

Where i is a subscript representing subject-level choices, for sample size n. β represents the vector of parameters to be estimated, and Φ is the standard normal cumulative distribution function. x_i is a vector of explanatory variables representing the differences in CPT utilities of the binary outcomes with some error. For example, let us assume that a subject is deciding whether to enter the casino or not. They have generated a strategy that yields a CPT value of 2.65. Whereas the CPT value of not playing at all would be 0. We assume the individual will decide to enter if the difference between the CPT values of the two options, plus a random term, ϵ, is greater than 0. Therefore, the probability of the subject entering is:

$$
P(y = 1_i|\beta_i) = P(CPT_{\text{enter}} - CPT_{\text{dont enter}} + \epsilon > 0|\beta_i) \quad \text{where } \epsilon \sim N(0, \sigma^2) \tag{2.7}
$$

In terms of the the standard normal CDF, this becomes:

$$P(y_i = 1|\beta_i) = \Phi\left(\frac{CPT_{\text{enter}} - CPT_{\text{no enter}}}{\sigma}\right) \tag{2.8}$$

With the denominator being the standard deviation of the binary choice error, ϵ. We can now follow on from equation 2.6 and add the constant error tremble. The constant error, ω, comes from a subject making a random, non-computational mistake, whether that be by losing concentration or an external distraction. Independently, the constant error can be modelled as:

$$\begin{aligned} P(y_i = 1|\beta_i) &= (1-\omega)\cdot\delta + \omega\cdot(1-\delta) \\ P(y_i = 0|\beta_i) &= \omega\cdot\delta + (1-\omega)\cdot(1-\delta) \end{aligned} \tag{2.9}$$

Where δ represents a vector of deterministic predictions of the model and can take values 0 or 1. 0 represents a non-entry prediction, and 1 represents an entry prediction. For example, if ω = 0.1, and the deterministic prediction is to enter, then the probability of entering becomes $P(y = 1|\beta)$ = 0.9, and the probability of not entering becomes $P(y = 0|\beta)$ = 0.1. When $\omega = 0.5$, the subject's choice is completely random, whilst $\omega = 0$ suggests the individual is perfectly fit to the proposed data generating process.

Combining the constant error model with the probit transformation model, we form our stochastic specification in equation 2.10.

$$\begin{aligned} P(y_i = 1|\beta_i) &= (1-\omega)\cdot\Phi(x|\beta_i) + \frac{\omega}{2} \\ P(y_i = 0|\beta_i) &= (1-\omega)\cdot(1-\Phi(x|\beta_i)) + \frac{\omega}{2} \end{aligned} \tag{2.10}$$

Again β is the vector of parameter triples, x_i is the vector of explanatory variables representing the differences in CPT utilities of the binary outcomes. Φ is the cumulative distribution function, and ω represents is the constant error. If $omega = 0$, then the stochastic specification becomes the nested standard probit transformation model, as there is no tremble. However, as $\omega \to 1$, the probabilities approach 0.5, indicating complete random choice.

To recover parameters and assess each models fit to the data, we use the log-likelihood function to determine the likelihood that a subjects' choices can be explained by the corresponding data generating process. Given the stochastic specification, our log-likelihood, for each subject, becomes

$$\mathcal{LL}(D_i|\beta_i) = \sum_{j=1}^{m} \log(\hat{p}_j)$$
$$\text{where } \hat{p}_j = \begin{cases} P & \text{if } \delta = D \\ 1 - P & \text{if } \delta \neq D \end{cases} \tag{2.11}$$

Where $D_{i,j}$ represents the choice of subject i, for a decision j, and P represents the probabilistic predictions of the theoretical model, as calculated in equation 2.10

We aim to identify the parameters that maximised the likelihood for each model, and classify types based on the model that recovers the parameter set with the lowest absolute log-likelihood value.

2.3.2 Simulation Results

Our simulation determined that effective parameter recover was not feasible directly from the theoretical design. We find that for all 8000 parameter triples for the naive plan, only 33 of the 801 potential strategy sets emerged as optimal strategies. Meaning that on average, 242 parameter triples predict the same strategy. This result holds for the naive actual play where across the 8000 parameter triples, there are 28 different strategies, such that on average, 285 parameter triples predict the same strategy. When combining the naive plan and actual play strategies, again we observe substantial homogeneity in joint strategies across the parameter set.

Again this holds for the sophisticates, in that across the 8000 parameter triples, there are 7 unique strategies that are adopted. When combining the SC plan and action (identical to naive plan) with the SNC plan and action, we still find multiple pairings of the same triples predicting the joint equivalent strategies.

Our results from the simulation confirm this result. We first attempt parameter recovery using the constant error specification. Depending on the parameter triple simulated, we find that there are sometimes over 200 parameter triples yielding the lowest absolute likelihood value, with parameter intervals spanning the full parameter space. Whilst the correct parameter always fell within this range, it was unidentifiable. Secondly we implement the probabilistic rule to create more variability in choices, and find, over 1000 simulations, on average we are able to identify simulated parameter triple less than 50% of the time. When combining the constant error model with the probabilistic rule, we are able to identify the simulated parameter triple slightly more often. Nonetheless, as a result of multiple triples predicting the same strategies in various conditions, we confirm that direct parameter recovery was not feasible.

Note, we are not implying this has implications for the empirical validity of the Barberis (2012) CGM. It does however, have implications for experimental designs in dynamic decision-making. The objective of this exercise was to inform the experimental design of this chapter, where it is now evident that the Barberis (2012) task alone will not provide robust estimates. We therefore proceed by estimating parameter triples from an independent certainty equivalent elicitation task using a binary lottery choice task design. Using this routine, our simulations provide support for the robustness of parameter recovery. We then proceed to fit estimated triples to the decision-tree design in order to identify types.

2.4 Experimental Design

The experiment took place with 71 students from Lancaster University, where individuals completed the experiment in the Lancaster University Experimental Economics Lab. The experiment comprised of 3 main stages. The first stage was a real effort task where individuals could earn an endowment to play the game. The second stage (the main game), involved playing the 5-period decision-tree, 16 times over two different settings (8 times in each setting). The third stage was a static risk preference elicitation task. Each stage will now be explained, with key insights justifying the design.

2.4.1 Part 1: Real Effort Task

The exploration of loss aversion in experimental settings poses several intricate challenges. Presenting hypothetical payoffs tends to result in arbitrary choices (Slovic, 1969), while requiring participants to risk their own funds raises profound ethical concerns. Moreover, endowing individuals with money results in a lack of personal attachment, thereby restricting the elicitation of loss aversion and increasing risk-taking behaviour (Thaler and Johnson, 1990). The most effective method identified thus far involves employing real-effort tasks (Mccabe et al. 1994; Carpenter and Gong 2015). This approach allows participants to earn money through straightforward, low-effort tasks, fostering a psychological connection to the earnings and replicating a more genuine loss experience. The endowment effect suggest that individuals attach a higher subjective value to possessions that they believe "belongs" to them (Kahneman, Knetsch, and Thaler, 1991). Therefore, the first stage of our experiment will consist of a real-effort task whereby subjects will generate the funds they need to enter the casino, and to compensate them for any incurred losses in the later stages. As our intended effect is to make subjects feel more entitled to their endowment, it is crucial to ensure that the task entails a positive cost of effort for the subjects. Charness, Gneezy, and Henderson (2018) provide a comprehensive overview of the real-effort experiments used in the literature. Our task follows a similar design to that of Abeler et al. (2011) which involves counting the number of zeros in tables that consisted of 25 randomly ordered 0's and 1's. The characteristics of this task make it pragmatically effective in our setting, as no prerequisite knowledge is required, it is boring and holds no significance, and offers limited learning opportunities from experience. It is also important to ensure that we do not fatigue the participants at this stage, thus we restrict the task to being 1 minute long, as was done in Lezzi, Fleming, and Zizzo (2015).

The participants therefore had 1 minute to complete as many tables as possible, where for each table they had to choose a number, between 0 and 25, that corresponded to the number of 0's they identified in the table. Once they had chosen, they were presented with a new, randomly distributed table, where they had to count again. Figure 2.6 shows how the table was presented to the subjects. Participants "passed" the real effort task, and received an endowment of £10, if they answered at least 4 tables correctly. As it is inevitable that some participants will perform better than others due to unobserved abilities, assigning endowments based on the specific

FIGURE 2.6: **Real Effort Task**

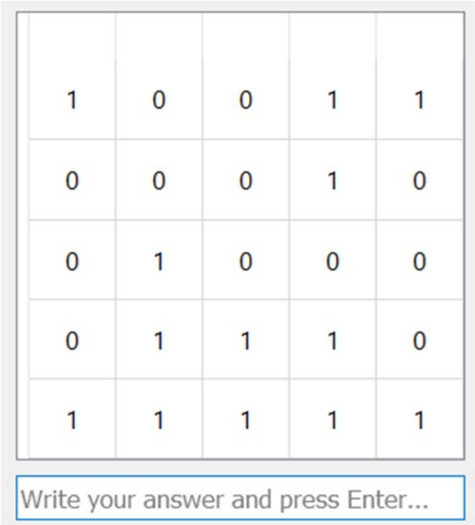

Notes: Subjects were required to identify the number of 0's in a table. They had 1 minute to complete as many tables as possible.

number of correctly answered tables would lead to different participants having different endowments for the next stages. Carpenter and Huet-Vaughn (2019) highlight the hindrances associated with this approach and the likelihood of omitted variable bias occurring. Therefore, to restrict any unintended effects, we imposed that any "pass" would result in the earning of a set endowment, such that all participants would receive either 0 or the same endowment. In this case, if subjects pass, they will receive an endowment of £10 to play the Barberis casino gamble task, as well as the risk preference elicitation task. This also served as a competency test, such that we would only proceed with subjects who had basic numeric reasoning abilities.

2.4.2 Part 2: Casino Play

The second stage of our experiment involved the participants playing in two versions of the binomial tree game, where in each version, they have to play 8 rounds. The first version neglects any form of commitment to initial strategies, and the second version gives individuals the option to commit. We adopt a within-subject design, such that all individuals engage in both versions of the game. The two versions, alongside the within-subject design, are crucial in order to effectively identify the various behavioural types under consideration, given the theoretical assumptions of each model. This will become clear shortly.

FIGURE 2.7: **Strategy Selection.**

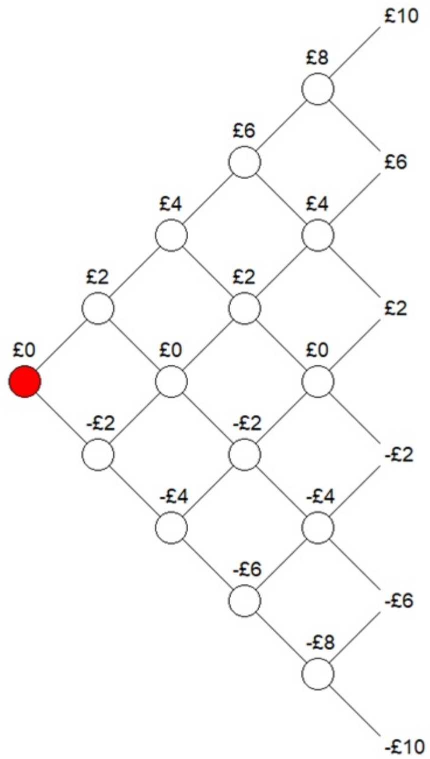

Notes: Subjects were required to select the nodes where they wished to exit.

Casino "No Commitment"

In the first version of the game, the participants first played three practice rounds in order to familiarise themselves with the process. After playing the practice rounds, the subjects were then asked to elicit their initial strategies. More specifically, they were asked to think as if they had to instruct the computer on how to play on their behalf, and to select, on an interactive binomial tree, the nodes at which they would want the computer to leave or continue playing, if they were to land on those nodes. This was to capture subjects' initial strategies at t=0. Following Johnson et al. (2021), subjects made their decisions by giving instructions to the computer. Ebert and Voigt (2023) highlight that it is crucial that subjects choose the order in which they devise their strategies by themselves. Given that we wish to differentiate between forward-looking strategies and backward induction strategies, we allowed for flexible strategy selection. That being, subjects were able to start selecting exit nodes from any position on the tree. Figure 2.7 illustrates the interactive screen in which they selected their exit nodes.

FIGURE 2.8: **Example Decision Tree Task.**

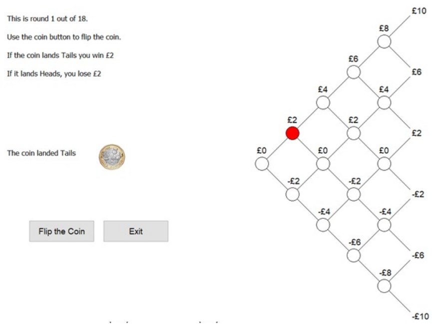

(A) t=1 after a win at t=0

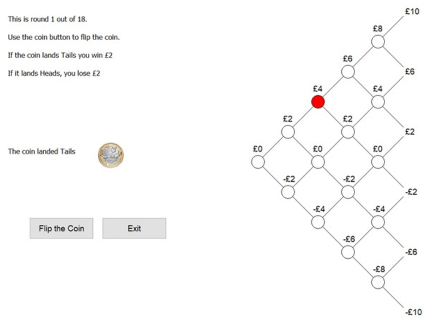

(B) t=2 after a win at t=1

Subjects were informed that there is a 1/n chance that this round could be their payment round, which would determine their final payoff. This incentive compatible approach ensures they reveal their true strategies.

Following the submission of their initial strategies, they were required to play 8 rounds of the game, where they were not shown their initial strategy, and had to make a decision, in each round, as to whether they wished to enter the game, and if so, make sequential decisions at each node dependent on nature's move. Subjects were informed that the computer would not play on their behalf, and that one of these 9 rounds (strategy + 8 rounds) could be selected for real at the end with a probability of 1/n, and would determine how much they walk away with. Figure 2.8 provides an example from the experimental interface.

Due to there being no form of commitment mentioned or available, this stage identified whether subjects were merely Naïve, or Sophisticated, based on their ex-ante strategies, as well as their ex-post choices. Once subjects had completed the 9 rounds, they could move onto the second version.

Casino "Commitment"

In the second version of the game, subjects were required, again, to elicit initial strategies in the same manner as before, except this time they were told they would be given the opportunity to enforce this strategy in the game, whereby if they selected "commit" then the software would automatically end the game if they reached one of the nodes specified as a “leave” node in their initial strategy. The intuition behind this is explained in the next section. After eliciting their initial strategies, subjects then made the decision as to whether they wanted to enter the game or not. If they entered, they then had to make a decision as to whether they wanted to enforce commitment to their elicited strategy or not. They then played the game out for another 8 rounds. Again, any of these 9 rounds (the strategy + 8 games), could have been selected at random at the end to be played for real.

Rationale and Justification for Design

The theory defines a sophisticate with commitment as an individual who is aware of their dynamic inconsistencies, but is able to find a form of commitment that assists them in abiding by their initial strategies. A sophisticate without commitment is an individual who is also aware of their dynamic inconsistencies, but is unable to find a means of committing to their strategy. We therefore assume that the "commitment device" is exogenous, and that the sophisticate with and without commitment are the same individual, but in different conditions, and will elicit different ex-ante strategies based on the condition in which they find themselves. Using this design, we can elicit different initial strategies for each of the types, and identify structural differences in how sophisticates play based on whether they have the option to commit and their preference parameter triples. In the "no-commitment" version of the game, we expect the naïve individuals to elicit a forward-looking strategy, and for those who enter, to deviate from this strategy ex-post. Whilst a sophisticate will devise an initial strategy using backwards induction, to account for their dynamic inconsistencies, and tailor their strategy to how they know they will act in future periods. They will enter if they are able to find a strategy that yields a CPT value greater than 0, and will play according to their strategy without deviations. It is likely a large number of sophisticates will not enter here. In the commitment version of the game, the naïve individuals will again elicit a forward-looking strategy, they will enter, decline the commitment device, and deviate from their strategies ex-post. However, the sophisticates will now devise a forward-looking strategy, which would be identical to a naïve strategy with the same preference parameter triples, and will enter if their strategy yields a CPT value greater than zero. As they are now able to commit to their ex-ante strategies, they no longer require backward induction, and there are a larger number of potential strategies that generate a positive

TABLE 2.5: **Risk preference elicitation task.**

Gains	Losses	Mixed
(5, 1/2; 0)	(-5, 1/2; 0)	(10, 1/2; -x)
(10, 1/2; 0)	(-10, 1/2; 0)	(5, 1/2; -x)
(10, 1/2; 5)	(-10, 1/2; -5)	
(10, 1/8; 0)	(-10, 1/8; 0)	
(10, 7/8; 0)	(-10, 7/8; 0)	

Notes: The 12 certainty equivalent elicitation tasks. 5 were in gains, 5 were in losses, and 2 were in the mixed domain.

payoff distribution. After entering, they will choose to commit, and the computer will end the game when they reach their bounds. We expect some sophisticates who did not enter in the non-commitment version, to now enter given they are granted a commitment opportunity.

2.4.3 Part 3: Risk Preference Elicitation Task

The final stage of the experiment involved subjects completing a series of certainty equivalent tasks that allow us to measure their risk and loss preferences in a static environment. This was required as parameter estimation was not possible from the sequential casino design.[10] To reduce negative effects arising from individual boredom, as well as cognitive fatigue, we avoid multiple binary lotteries and follow the approach of Bruhin, Fehr-Duda, and Epper (2010), which was later adopted by Vieider et al. (2015), Bouchouicha et al. (2019) and Ring et al. (2018), to elicit individual certainty equivalents. This approach required subjects to complete fewer tasks, as it involved individuals selecting a value of indifference rather than selecting a preferred lottery, thereby generating cardinal data instead of ordinal data.

We elicit certainty equivalents for a total of 12 binary prospects, of which 5 are pure gains, 5 are pure losses, are 2 are mixed. The tasks included in the analysis are presented in Table 2.5.

Participants were presented with each of the 12 tasks in Table 2.5, where the first column represents the tasks in the gain domain, column 2 represents tasks in the loss domain, and column 3 in the mixed domain. The three domains allow for the elicitation of three CPT behavioural parameters, namely utility curvature, probability

[10]See Section 2.3 for more details.

weighting, and loss aversion. By varying outcomes and probabilities in a structured way, our simulation routine identified that our 12 task design was able to accurately estimate risk preferences.

Figure 2.9a shows an example of the instructions for the first task the gain domain. In this task, subjects were presented with a lottery, represented with balls on a screen. They were asked whether they prefer the lottery (15,1/2;0,1/2), or a list of sure amounts, ranging between the win and loss amount of the lottery. For each element of the sure amount list, they had to decide whether they preferred the lottery or that specific sure amount. The point where the individual switched from preferring the sure amount to preferring the lottery determined the certainty equivalent, whereby we take the mean of the two values between which they switched. Subjects completed 5 tasks in the gain domain, where payoffs and probabilities varied across tasks. In the loss domain, the task followed the same format, except with pure negative outcomes, as shown in figure 2.9b.

In the mixed domain, instead, subjects were presented with a lottery (x,p;y,1-p), where x and p are given. They had to select the y amount, where $y < 0$, that made them indifferent between the lottery, and a certain £0. An example is shown in Figure 2.9c.

Subjects were informed that one of the tasks would be played out for real, with probability 1/n, and if selected, would determine their earnings (Show-up fee + Endowment from Stage 1 +/- Selected task for real).

Every task had the same probability of being extracted for real play. This provided an incentive to respond according to one's true preferences, and is the standard procedure in the literature (Baltussen et al. 2012; Bruhin, Fehr-Duda, and Epper

FIGURE 2.9: **Risk preference elicitation task example.**

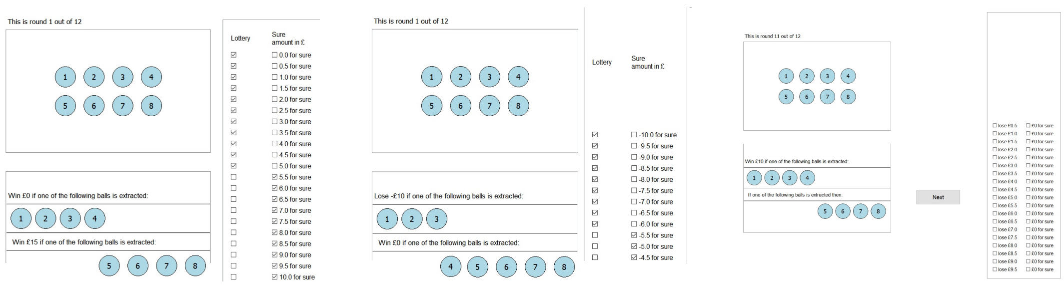

(A) Gains example (B) Losses example (C) Mixed example

2010; Cubitt, Starmer, and Sugden 1998).

2.5 Methodology

The following section outlines our methodology for recovering parameters from the risk preference elicitation task, and identifying types using a model prediction exercise.

The risk preference elicitation task recovers subject-level certainty equivalents for each decision-task. The tasks included gains, losses, and mixed prospects, with varying monetary amounts and probabilities. Therefore quantitative CEs allow for the recovery of CPT parameters. We use MLE to estimate parameters. Please refer back to Section 1.5 for an in depth explanation of the MLE econometric specification. For each subject, we recover the unique preference parameter triple that corresponds to their maximised log-likelihood value.

2.5.1 Prediction routine

Given that the parameter triples are estimated with some error in an independent task, utilising the trembling binary choice model for type identification would add additional error stories (ω, σ) to the estimation. This would inevitably create excessively noisy results. Instead, we simplify the exercise and adopt a basic prediction routine.

We aim to evaluate the predictive accuracy of the two models in question: the naive model, N, and the sophisticate model, S, using the vector of parameters β, which were estimated using the maximum likelihood estimation routine. For each parameter set β, we derive theoretical predictions from both models and compare these predictions against the experimental data.

Let $\boldsymbol{\beta} = (\beta_1, \beta_2, \ldots, \beta_k)$ be the vector of estimated parameters. $T_N(\beta_i)$ and $T_S(\beta_i)$ are the theoretical predictions from models N and S respectively, given a parameter set β_i. Finally, $\mathbf{D} = (D_1, D_2, \ldots, D_n)$ is the vector of experimental data points, where each D_i corresponds to a vector of subject $i's$ observed data.

For each $\beta_i \in \boldsymbol{\beta}$, we compute the theoretical predictions from both models:

$$\mathbf{T}_N = (T_N(\beta_1), T_N(\beta_2), \ldots, T_N(\beta_k))$$

$$\mathbf{T}_S = (T_S(\beta_1), T_S(\beta_2), \ldots, T_S(\beta_k))$$

For each data point $D_j \in \mathbf{D_i}$, for each subject, we calculate the mean squared error between the experimental data and the theoretical predictions from each model:

$$\begin{aligned} \text{MSE}_{Ni} &= \sum_{j=1}^{n} (D_j - T_N(\beta_i))^2 \\ \text{MSE}_{Si} &= \sum_{j=1}^{n} (D_j - T_S(\beta_i))^2 \end{aligned} \tag{2.12}$$

Given that our data comprises primarily of 0s and 1s, we classify types based on the model that elicits the lowest mean squared error.

2.6 Results

All subjects correctly identified at least 4 tables in the real-effort task, and were therefore endowed with £10 in order to complete the experiment.

2.6.1 Parameter Estimates

Our risk preference elicitation task identified early on that most subjects were in fact risk seeking, and were hitting the bounds of the α parameter, resulting in higher noise. Given that gambling is associated with risky behaviour, this is intuitive, and we therefore proceed to estimate types with the upper bound of 1, such as to replicate the Barberis (2012) model, as well as with an upper bound of 2, such that we are able to capture risk-seeking preferences. We find that, when relaxing the bound on α, the subject level standard errors on the parameter estimate are statistically significantly better (Welch t-test, $p<0.05$; Mann-Whitney U test: $p<0.05$). Additionally, when identifying types by minimising the prediction error between the subject-level data and the theoretical predictions of the Barberis (2012) model, we find the prediction error is significantly lower ($p<0.05$) when we allow for risk-seeking behaviour.

FIGURE 2.10: **Kernel Density Estimate Plots.**

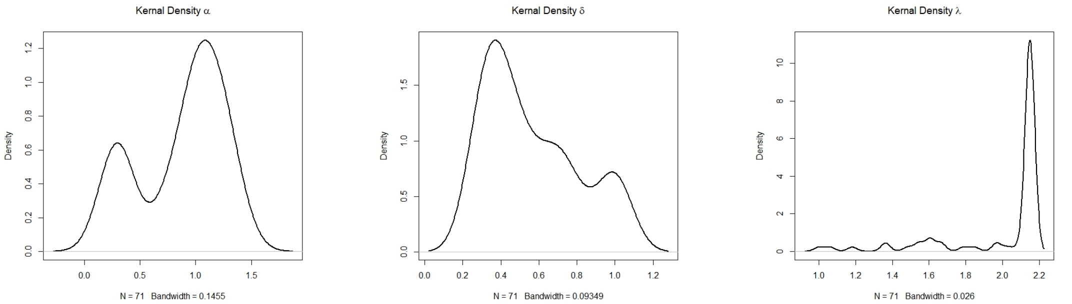

(A) Kernel Density Estimate Plot for α

(B) Kernel Density Estimate Plot for δ

(C) Kernel Density Estimate Plot for λ

Notes: Panel A is the density estimate for the CRRA parameter, Panel B is for the probability distortion parameter, and Panel C is for the loss aversion parameter.

Similarly, we identify that for any λ >2.15, as long as α is greater than 0.1, then there is no difference in decision-strategies or choices in the theoretical model. For this reason, we impose an upper bound limit of 2.15 in order to allow for more variability of lambda within this range.

Prior to classifying types, we analyse the mean and standard deviation of the parameters over all participants, where we find that (μ_α, SD), = (0.873,0.379), (μ_δ, SD) = (0.564,0.244), and (μ_λ, SD) = (2, 0.200). Figure 2.10 illustrates the density plot for each parameters kernel density estimate (Weglarczyk, 2018).

Given that we have now increased the upper bound of α from 1 to 2 in order to capture risk-seeking preferences, the theoretical predictions of the model will have changed. More specifically, the number of parameter triples for which a naïve or sophisticate will enter the casino will have changed. The original specification imposes that 0<α<1, 0.3<δ<1, 1<λ<4, however our updated specification now imposes the following bounds: 0<α<2, 0.3<δ<1, 1<λ<2.15. Additionally, we modify the parameter intervals for α such there are 40 equally spaced points instead of 20, to allow for more variability within the bounds. Therefore, we now have 16,000 possible parameter triples, instead of the original 8000. Under the new parameter intervals, we find that the naive agent will enter for 11,171 triples, and the sophisticate will enter for 4823 triples.[11]

[11] Note: the increase in the number of parameter triples is not directly comparable to the original model, as we change both the limits and the number of spaced points. We provide the statistics purely for reference.

2.6.2 Type Identification

Based on the estimates from our risk preference elicitation task, we identify that 56 out of the 71 participants were best classified with the naïve model, and 15 out of the 71 participants are best classified with the sophisticate model. For each subject's optimal model, on average, the model predicts 70% of the subject's node-level decisions correctly. We find that each subjects' prediction error for their classified model is statistically significantly better than the predictions from the opposing model (t-test, p<0.001; Mann Whitney U Test, p<0.001).

2.6.3 Commitment

Regarding commitment decisions, we find that 59/71 participants opted for the commitment device, and 12/71 declined the devise, where 100% of the sophisticates committed, and 79% of the naïve subjects committed. If we remember the theoretical predictions of the model, it states that sophisticates will always choose the commitment devise when available, but the naïve subjects will decline it. However, we find that most naives in fact choose to commit. This has interesting implications, which we will discuss in more detail in Section 2.7.

2.6.4 Parameter/Type Relationship

We now look to see if there is a relationship between risk preferences and subject-types. Table 2.6 provides the mean (μ), standard deviations (σ) and median (M) values for the preference parameters at the type level.

TABLE 2.6: **Mean and Median parameter values per group.**

	Naive	Sophisticate	Significant Differences
μ_α / σ_α	0.978 / 0.326	0.479 / 0.3	$p < 0.001^{**}$
M_α	1.077	0.359	$p < 0.001^{**}$
μ_δ / σ_δ	0.497 / 0.177	0.812 / 0.29	$p < 0.002^{**}$
M_δ	0.447	1	$p < 0.001^{**}$
μ_λ / σ_λ	2.03 / 0.255	1.903 / 0.3	$p = 0.262$
M_λ	2.15	2.15	$p = 0.381$

Notes: Our statistical significance tests are a two-sided Welch Sample t-test for the means, and Mann Whitney U Test for the medians. ** Represents statistically significant differences at the 5% level.

FIGURE 2.11: **Group level kernel density estimates.**

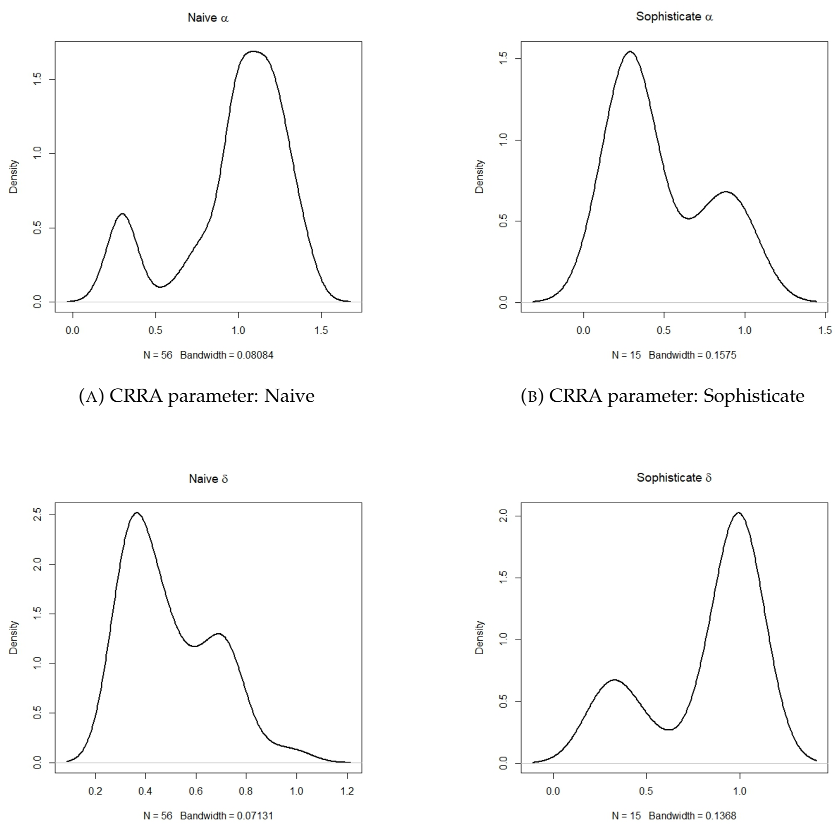

(A) CRRA parameter: Naive

(B) CRRA parameter: Sophisticate

(C) Distortion parameter: Naive

(D) Distortion parameter: Sophisticate

Notes: Plots for the Kernel Density Estimates of the CRRA parameter and the probability weighting parameter for each type.

We find that the naïve group are associated with less risk averse behaviour than the sophisticates, and in many cases, the naïve subjects exhibited risk seeking behaviour (α>1). Additionally, the naïve group tend to distort probabilities more than the sophisticates. We find no differences in the loss aversion parameters, where all subjects tend to, as suggested by the theory of loss aversion, receive a disutility for losses that is double the utility they receive for respective gains. Figure 2.11 provides the density plots for the CRRA parameter and the probability distortion parameter for each type.

2.6.5 Entry Parameters

In the experimental setting, subjects were allowed to decide as to whether they wished to enter the casino at the beginning of each of the 16 rounds. Therefore, for each subject, we have 16 entry decisions, of which 8 are in the no-commitment

FIGURE 2.12: **Parameter estimates 3D scatter plot.**

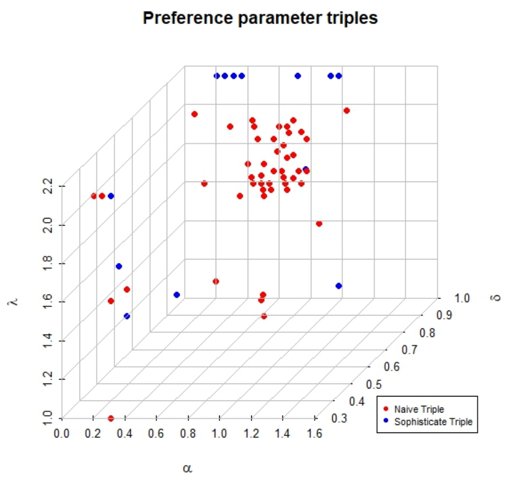

Notes: Parameter triples elicited by our subjects. A red node indicates the model classified them as naive, and a blue node, sophisticate.

rounds, and 8 are in the commitment rounds. We find that on average, the naïve subjects enter the casino 98% of the time in the no-commitment round, and 98% of the time in the commitment rounds. The sophisticates entered almost 100% of the time in both the no-commitment and commitment conditions. However, the theoretical predictions of the models suggest that only 35/56 of the Naïve's should have entered in both the commitment and no commitment conditions, whilst only 2/15 of the sophisticates should have entered.[12]

Figure 2.12 illustrates the parameter triples for our types, where red nodes illustrate the parameter triples for subjects classified by the model as naive, and blue nodes denote the sophisticate triples.

2.6.6 Strategy Type

Subjects' were asked to elicit initial strategies twice: once where commitment was not offered or ever mentioned (NC Condition), and one in which they were told they would have the option to make this strategy binding during real play (C Condition). In the NC setting, we find that 80% of the naives elicit a loss-exit strategy; one in

[12]There are three potential explanations for this. Firstly, it could be that the contextual independence of the risk preference elicitation task from the casino task has resulted in biased parameter estimates. Secondly, it could be that the CGM's assumption regarding strategy elicitation are hindered. Finally, in could be in part due to the experimental setting in which they are being asked to gamble, rather than making the decision in a real-world scenario, which fosters risk-seeking behaviour.

FIGURE 2.13: **Strategy type elicitation.**

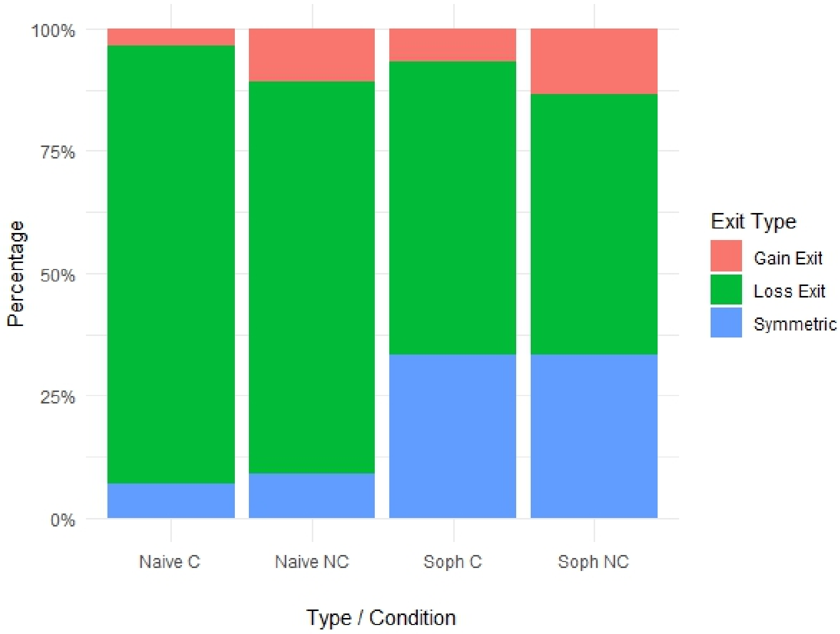

Notes: The strategy types that subjects elicit in both conditions. Naive C is a subject classified as naive in the commitment condition. NC represents no commitment. Soph represents the sophisticate type.

which subjects intend to play for longer in the gain domain than they do in the loss domain. 11% of the naïve's elicit a gain exit strategy, and 9% elicit a neutral strategy. For the sophisticates in the no-commitment setting, 53% elicit a loss-exit strategy, 13% elicit a gain exit-strategy, and 34% elicit a neutral strategy.

In the commitment setting, 90% of naives elicit loss-exit strategies, 4% elicit gain-exit strategies, and 6% elicit neutral strategies. Whereas for the sophisticates, in the commitment setting, 60% elicit loss-exit strategies, 7% elicit gain-exit strategies, and 34% elicit neutral strategies. Figure 2.13 provides a graphical representation of the strategy type frequencies for each subject-type in both of conditions. We do not observe any tangible differences between the groups, but we do find, aggregating over all participants, that the availability of a commitment device decreased the likelihood of observing a gain-exit strategy.

2.6.7 Strategy Changes

We compute the hamming distance between each subjects' no-commitment strategies and commitment strategies to identify if the availability of a commitment device led to statistically significant changes in strategy selection (Bookstein, Kulyukin, and Raita, 2002). Given there are only 14 potential nodes at which participants were required to make strategy-level decisions, we compute hamming distances at the subject level. Initially looking at subjects with hamming distances that are greater than 1, we find that 48% of naives and 40% of sophisticates adopt at least minor strategy changes. However, if we look at subjects' with hamming distances greater than 4, or in other words, subjects who changed their strategies on at least 35% of the nodes, we find only 12.5% of naives adopted major strategy changes, and over 30% of sophisticates adopted major strategy changes.

Figure 2.14 provides the density of strategy changes at the subject-level for hamming distances greater than 1. Figure 2.14 gives light to the differences in median strategy changes illustrated in Figures 2.16 and 2.17. It is also interesting to note that of those sophisticates who do not change their strategy at all, 50% of them have strategies in which they selected nearly every node (13/14 nodes selected), suggesting their ideal plan would be to almost not enter at all. Yet in the no-commitment round, they proceed to play in all rounds. This, again, provides insights to the original classification of types, and whether the strict "naive" and "sophisticate" classifications are restrictive. We expand in Section 2.7 on the more flexible generation of a new type.

Whilst we follow the original Barberis (2012) classification of strategy types (e.g. gain-exit, loss-exit, symmetrical), we acknowledge that there are deeper classifications within these strategies. Ebert and Voigt (2023) assign strategy types into clusters and identify four primary strategy types: Buy-and-hold, Never-start, Take-profit, Stop-loss. We replicate this analysis with our experimental data to identify unbiased assignments of strategies across all participants, in both conditions. To do so, we perform a K-means clustering unsupervised algorithm to identify clusters of subject-level strategies in both conditions. The algorithm partitions data into

FIGURE 2.14: **Strategy changes across conditions.**

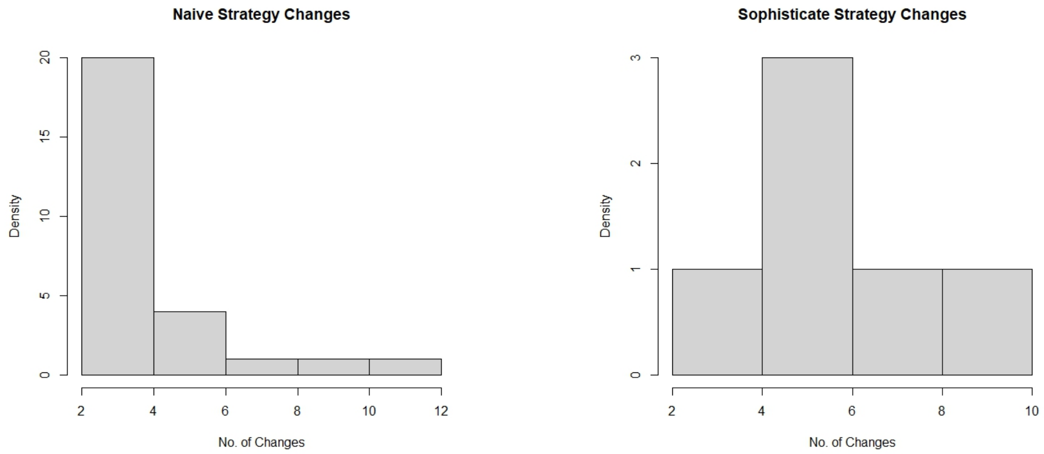

(A) Naive

(B) Sophisticate

Notes: Density of strategy changes from No commitment condition to Commitment condition. Panel A represents the density of strategy changes for those classified as Naive, and Panel B, for those classified as Sophisticate.

$K \in \mathbb{N}^+$ clusters with high within-cluster similarity and high-between cluster differences, where we set K=4[13]. Figure 2.15 illustrates the median clusters found over all participants in the No Commitment condition.

We find four distinct types of exit strategies. Figure 2.15a (Cluster 1) suggests subjects would exit anytime they fall into the loss domain, given they had experienced a gain at some point. We define this behaviour as "Exit Losses Post-Gain". Whereas Figure 2.15b (Cluster 2) represents "never enter" strategies. Subjects in cluster 3 (Figure 2.15c), look to avoid the extremes, such that any time the monetary payoff is 4+, they should take their winnings, and anytime it reaches -4, they should cut their losses and exit. We classify these subjects as "avoid extremes". Whilst subjects in cluster 4 (Figure 2.15d), are "pure loss-exit" strategists, that being, they continue playing for as long as they are objectively winning, and exit after an objective loss in any period, even if the cumulative earnings represent a gain.

Although we observe strategy changes across the NC and C conditions, the clusters remain almost identical, with the exception of cluster 2 in the C condition, where subjects play in the first node, and play once more if they win in t=0. We also observe

[13]Given that Ebert and Voigt (2023) identified four strategy types, and our sample size is limited, we restrict the number of clusters to 4

FIGURE 2.15: **Median Clusters: NC**

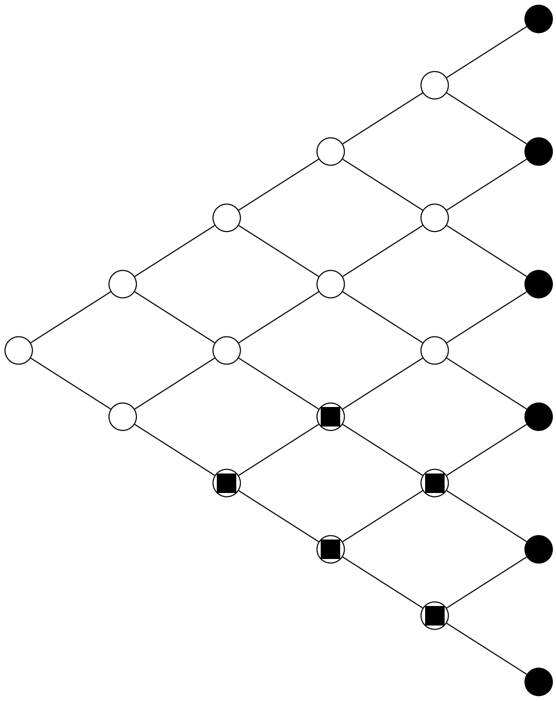

(A) Median Cluster 1: NC

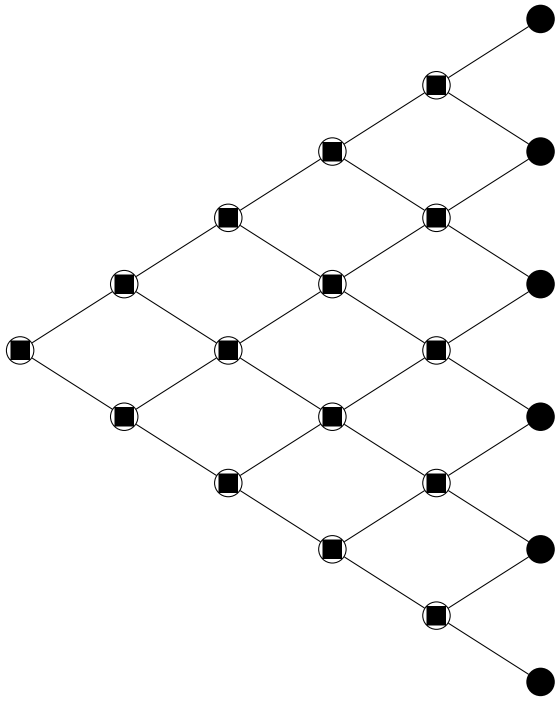

(B) Median cluster 2: NC

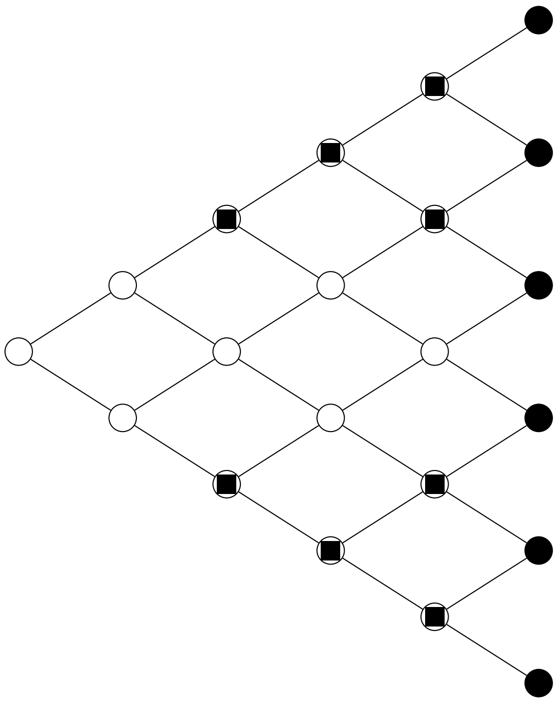

(C) Median Cluster 3: NC

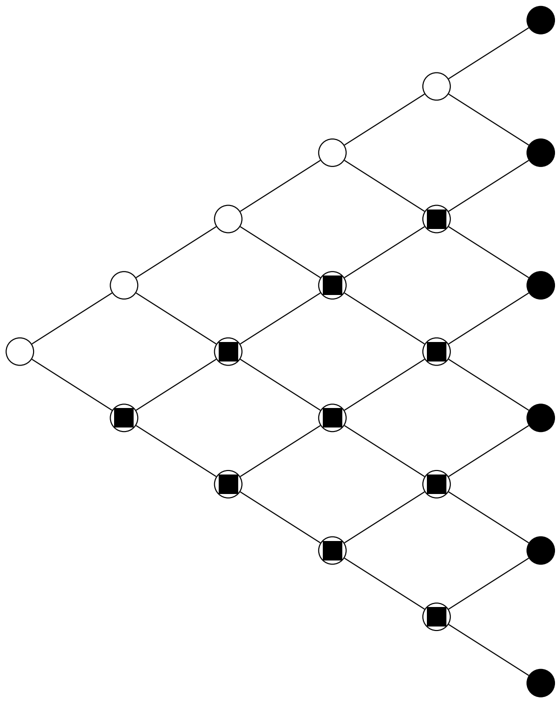

(D) Median cluster 4: NC

TABLE 2.7: **K-means clustering segmentation of strategy types.**

Cluster	N_{nc}	N_c	Categorisation	Naïve$_{nc}$	Soph$_{nc}$	Naïve$_c$	Soph$_c$
1	39	16	Exit in Loss Post-Gain	62.6%	26.7%	23.2%	20%
2	14	13	Never enter	16%	33.3%	14.3%	33.3%
3	9	32	Avoid Extremes	10.7%	20%	48.2%	33.3%
4	9	10	Pure Loss-Exit	10.7%	20%	14.3%	13.3%

Notes: Column 1 represents the cluster number for our k=4 clusters. Columns 2 and 3 represent the sample size for each cluster, with subscript nc indicating it was in the no commitment condition, and subscript c indicating it was in the commitment condition. Column 4 provides the economic interpretation of the each cluster, and columns 5:8 provide the density of each of our types in both the no commitment and commitment condition.

changes in the density of each cluster.

Table 2.7 provides these results, and provides the density of each type by strategy classification. Again, we do not find any patterns or statistically significant differences between the types.

2.6.8 Ex-ante: Methods

A primary assumption of the theoretical model is that the naïve subjects will elicit strategies with a forward-looking approach in both the no-commitment and the commitment setting. Whilst the sophisticates should elicit strategies using backward induction in the no-commitment setting, but switch to a forward-looking approach in the commitment setting. By disentangling our analysis such as to identify model prediction rates purely over strategies, we find in the no-commitment setting that 92% of those classified as naïve generated their strategies using more of a forward-looking approach, whilst for the other 8% it is unclear as to whether they used backward induction or some other method. 40% of those classified as sophisticates generated strategies using some form of backward induction.[14] In the commitment setting, the number of naives generating strategies that resemble a forward-looking

[14] We find that 30% of sophisticates, that are not included in the 40% who use backward induction, indeed use some form of backward induction in their actual play. Of course this will lead to dynamic inconsistencies if their initial strategy elicitation did not involve backward induction.

FIGURE 2.16: **Median ex-ante plan for Naive subjects.**

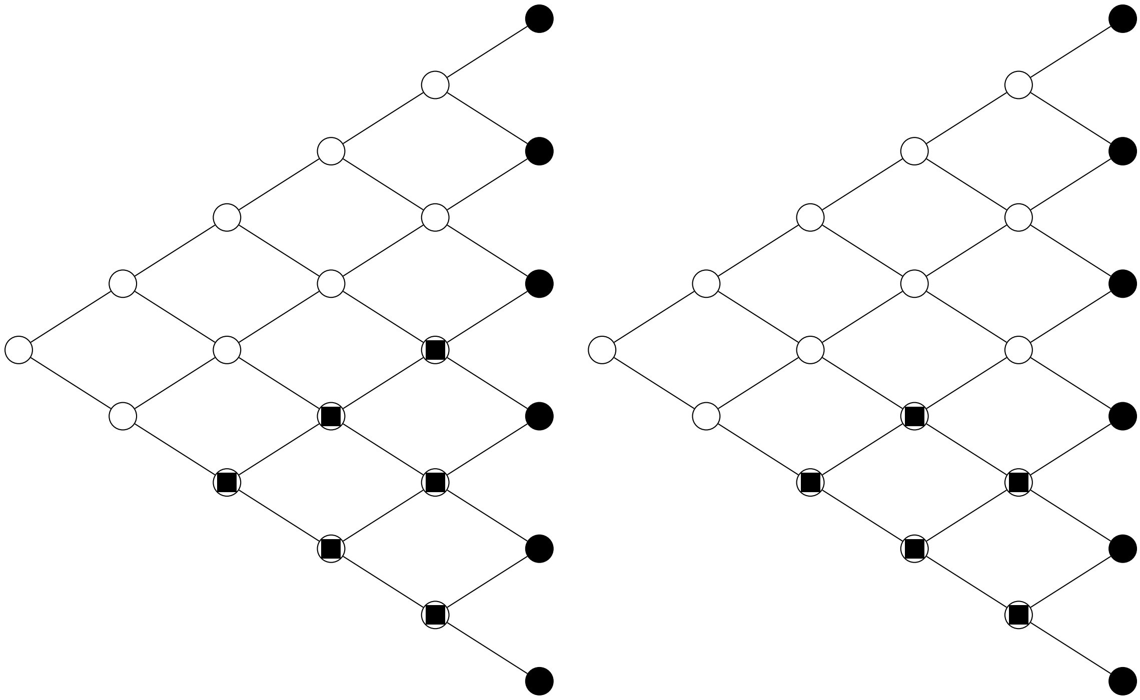

(A) No Commitment (B) Commitment

approach now increases to 100%, and interestingly, the number of sophisticates using backward induction now drops to 11%, providing support for the theoretical assumptions of the sophisticate type. Aggregating over each type, Figure 2.16 illustrates the average strategy used by those classified as naive, in both conditions, and Figure 2.17 plots the same for those classified as sophisticates. Whilst the strategy types (e.g. loss-exit / gain-exit) are similar across types, we can see from Figures 2.16 and 2.17 that the strategies at the node-level varied according the theoretical assumptions of the model. That being, on average, that the availability of a commitment device will not significantly alter the strategies of a naive agent, but it will for those classified as sophisticated.

2.6.9 Ex-post: Methods

We now analyse the model level predictions for the subject-level actual play. We find that in the no-commitment setting, 80% of the naive subjects follow the theoretical assumptions given their type, in that their t>0 ex-post choices resemble a forward looking approach such that deviations from initial strategies prevail, dependent on

FIGURE 2.17: **Median ex-ante plan for Sophisticate subjects.**

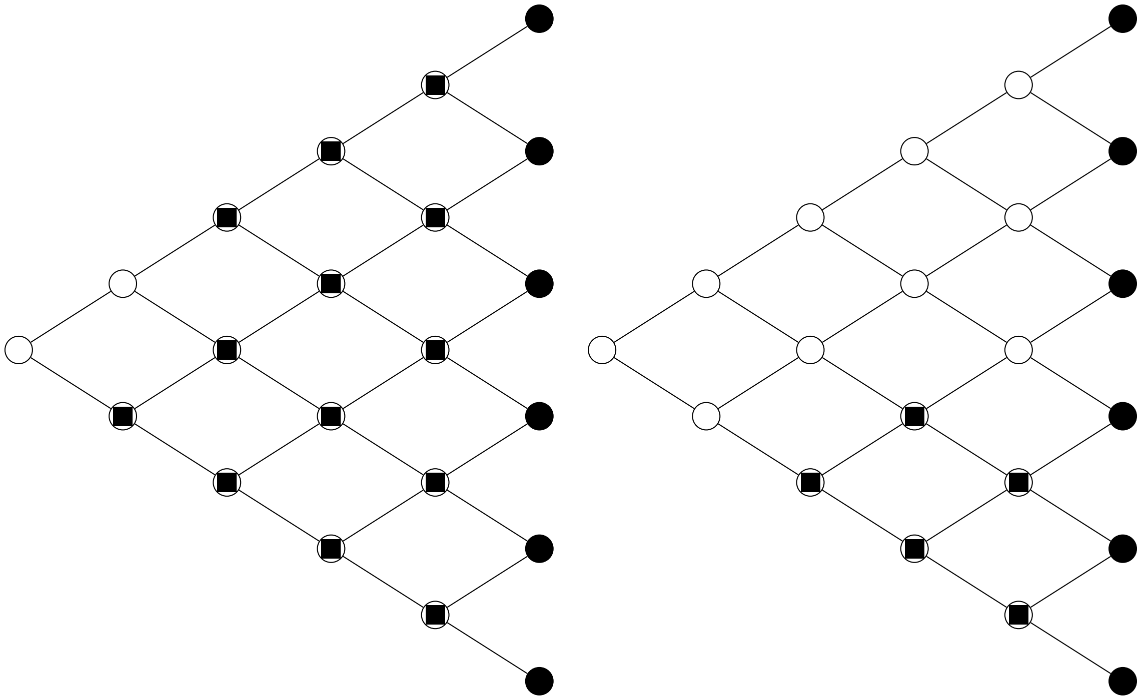

(A) No Commitment

(B) Commitment

the subject-level parameter triples. However, we find that only 6% of the sophisticates follow their t=0 ex-ante strategy, implying deviations from initial strategies ex-post, when t>0[15].

In the commitment setting, however, we find that 100% of sophisticates obtain relatively low prediction error scores for their actual play - that being, the theoretical predictions of the model, given their parameter triples, are relatively accurate. On the other hand, 65% of naive subjects satisfy the theoretical predictions of the model. The rise in prediction error here is likely a result of the commitment device objectively reducing the time-horizon and subsequently altering the expected outcome distribution of the gamble. Therefore, as a result of naive subjects opting for the commitment device, they either exited as a result of the device hitting the bounds, or they deviated within the devices constraints as a result of a change in the probabilities and the outcome distribution of the available gamble. Therefore, naive subjects who committed were likely unable to satisfy the predictions of the original model.

[15]Prediction-level assumptions are based on prediction errors of less than 25%

2.6.10 Strategy type: Ex ante Vs Ex Post

What underpins dynamically inconsistent choice in decision-making under risk is deviations from initial strategies. An agent may decide to engage in a risky action as their initial strategy generates a positively skewed outcome distribution. However, as t increases, distributions evolve, leading to deviations from ex-ante strategies over time. We now analyse whether a subjects' initial strategy type is carried out in actual play. In other words, if an individual elicits a loss exit strategy, do they then proceed to exit more in the loss domain or the gain domain. We find that for both types, in the no-commitment stage, whilst most subjects enter the casino with a loss-exit strategy in mind: 80% of the naïve's and 53% of the sophisticate's, there is a notable discrepancy between this strategy type and their actual play. Of those who are classified as naïve and elicited a loss-exit strategy, 45% exited more in the gain domain, 35% exited more in the loss domain, and the rest represent symmetrical exit-strategies. Of the sophisticates who elicit a loss-exit strategy, 50% exit more in the gain domain, 37.5% exit in losses, and the rest are symmetrical. A chi-square test, and fisher test, confirms in the no-commitment stage, that there is no observed association between eliciting a loss-exit strategy, and playing in a loss-exit manner (Naïve: $p=0.2603$, $p=0.266$; Sophisticate: $p = 0.4638$, $p=0.6357$).

In the commitment round, although commitment devices were technically semi-binding, in that subjects could still deviate within their bounds, most subjects exited in the domain that their initial strategy suggested. 68% of Naïve's stick to their initial strategy type in actual play ($p<0.001$, $p<0.001$), and 73% of sophisticates do also ($p=0.03$, $p=0.02$). Figure 2.18 illustrates the discrepancy between initial strategies and actual play for the naïve subject in both conditions, and Figure 2.19 illustrates the same for the Sophisticates.

2.6.11 Dynamic Inconsistencies: NC

As of yet we have discussed deviations from strategy type, but not deviations from objective strategies at the node-level. In the no-commitment round, we find that 79% of naive subjects and 67% of sophisticate subjects deviate from their initial plan at least once. The mean and median number of deviations across the two groups is not

FIGURE 2.18: **Naive: Strategy type vs actual play.**

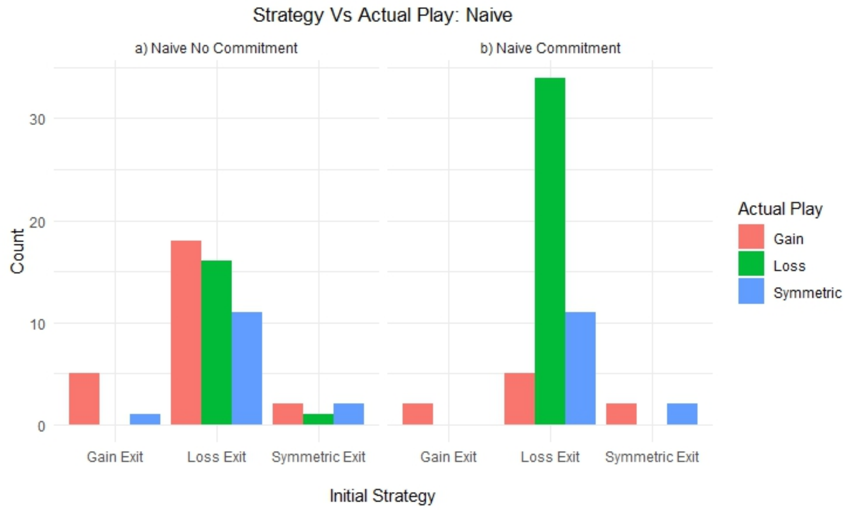

Notes: Panel A on the left represents the No Commitment condition, and Panel B on the right represents the Commitment condition.

statistically different (Naive: μ=2.85 M=2; Soph: μ=2.13 M=2) (p=0.246, p=0.297). However, when we aggregate types and separate subjects into those that opted for commitment, and those who did not in C condition, we find statistically significant differences in the NC condition. Remember, at this point, pre-commitment had not been mentioned, so this is based purely on subject behaviour. For those who have a demand for commitment (any who opt for commitment in the C condition), the mean and median number of deviations are 2.54 and 2 respectively, where only 42% of subjects deviated at more than 2 nodes. However, for those without demand for commitment, the mean and median number of deviations is 3.5 and 4 respectively.[16] The differences in the NC condition are statistically significant at the 10% confidence level (p=0.1, p=0.1). Figure 2.20 illustrates the kernel density estimates for the number of deviations per subject, in the NC condition, where we classify subjects into those who have demand for commitment, and those who do not.

Ebert and Voigt (2023) develop a measure of a subject's dynamic consistency, which takes insights from the Fischbacher, Hoffmann, and Schudy (2017) measure of dynamic consistency, as well as from Odean (1998)'s measure of the disposition

[16]We also assessed whether there was a correlation between round number (1-8) and number of deviations, but find no pattern or statistically significant correlation. This held for all groups.

FIGURE 2.19: **Sophisticate: Strategy type vs actual play.**

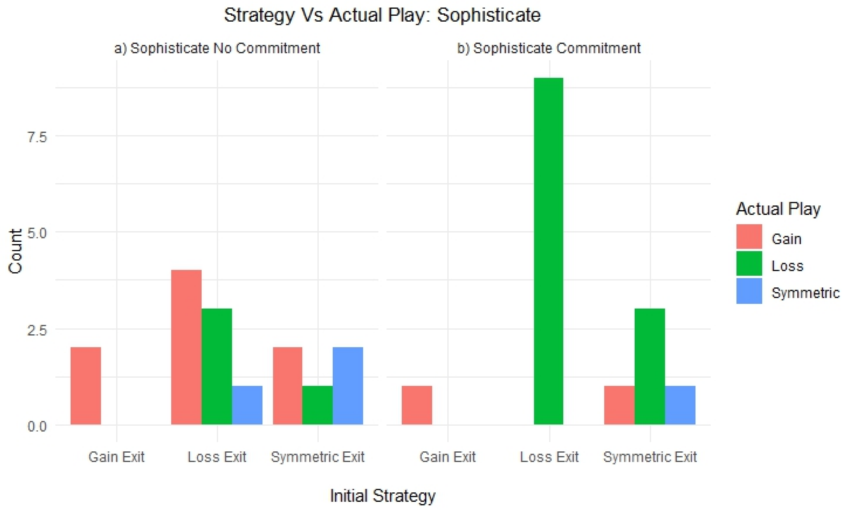

Notes: Panel A on the left represents the No Commitment condition, and Panel B on the right represents the Commitment condition.

effect. They define the consistency C_i of a subject as:

$$C_i = \frac{\text{Plan executions}_i}{\text{Deviation opportunities}_i} = 1 - \frac{\text{Deviations}_i}{\text{Deviation opportunities}_i}$$

Where $C_i = 1$ if subject i follows their plan in every game, and $C_i = 0$ if they never follow their plan. Ebert and Voigt (2023) find the distributions of dynamic consistency in their study are concentrated above 0.7, with a median of 0.89 over all studies. Where subjects follow their unconstrained plan for about 86% of sequential risk-taking actions, implying relatively consistent behaviour. However, in our setting, we find that the distributions of dynamic consistency have a wider concentration level range. We split our analysis into four subgroups. In the first two, we divide our sample into the two types, that being naive and sophisticate. As aforementioned, we do not find statistically significant differences between these groups. For the second two, we instead split our sample into two subgroups: those with, and those without, demand for commitment. This is to analyse the behaviour of subjects in the NC condition, but to segment subjects into those that will commit and those who will not. In Figure 2.21, we provide a violin plot to illustrate the distribution of the dynamic consistency measure, whilst providing a measure of the density's per group.

We can see there is substantially more heterogeneity in dynamic consistency than

FIGURE 2.20: **Deviations from ex-ante strategies: No commitment.**

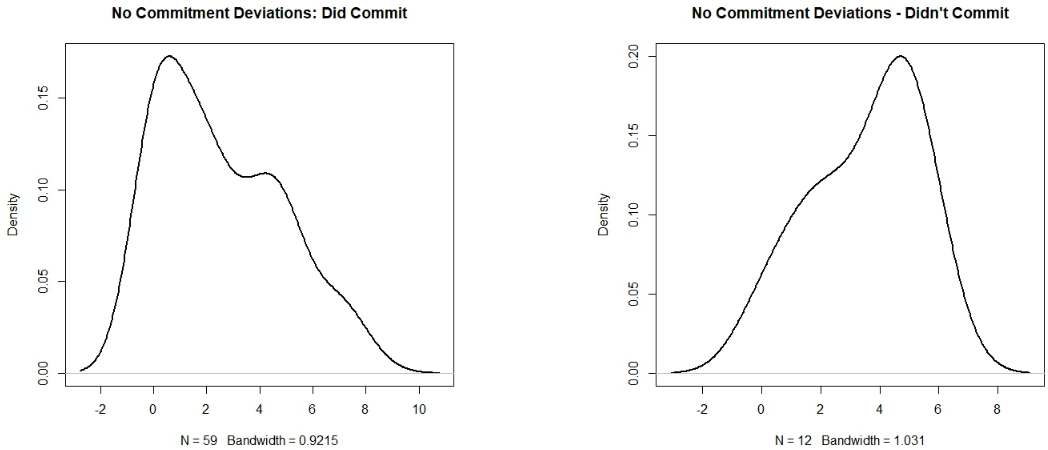

(A) Demand for Commitment (B) No Demand for Commitment

Notes: Kernel Density Estimates for the number of subjects who deviated from their initial strategies in the no-commitment setting. Separated into those who commit in the commitment setting (Panel A), and those who opt to play without commitment (Panel B).

FIGURE 2.21: **Distribution of dynamic consistency measure: No Commitment**

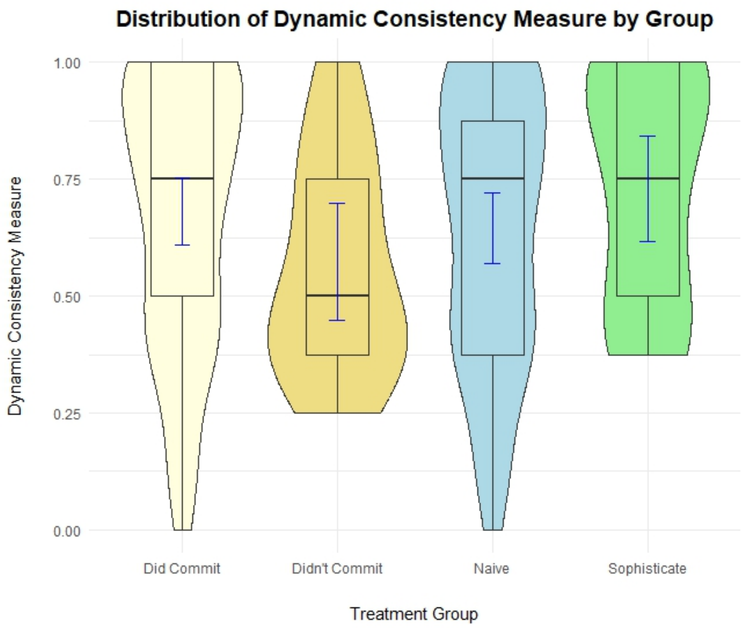

Notes: Distribution of dynamic consistency of four subgroups in the **No Commitment condition**. In the first two groups, subjects are separated into those who commit and those who do not. The second two groups separate's subjects into those classified as naive, and those classified as sophisticate.

in the Ebert and Voigt (2023) study. For those with demand for commitment, and those classified as sophisticate, we find their consistency measure is concentrated above 0.5, with median values of 0.75, and the largest density skewed towards dynamic consistency. For those classified as naive there is the largest heterogeneity, with a median of 0.75 that does not lie within the confidence levels of the mean. For these subjects, the dynamic consistency measure is concentrated between 0.375 and 0.875. The density of the measure is split between the two bounds of the concentration level, suggesting there are potentially two types of individuals within the original classification of "naive". Interestingly, for the subgroup who did not commit, we find the measure is concentrated between 0.375 and 0.75, with the density skewed towards the dynamically inconsistent bounds. Given that there are naives who committed, and naives who did not, it is likely that this has caused the heterogeneity in the dynamic consistency measure. Regardless, we find on average our subjects were less dynamically consistent than those in the Ebert and Voigt (2023) examination.

Ebert and Voigt (2023) suggest that offering flexible strategies may increase individuals' ability to stick to their plan beyond commitment devices that focus on restricted plans. However, in our setting, where flexible, trailing strategies were available, subjects still proceeded to deviate. We suggest that the reason subjects were more dynamically consistent in their setting is because they were able to see their plan throughout the task. This imposes a form of subconscious commitment. In other words, if an agent elicits a plan, and they can see the plan, they feel more obliged to stick to it. In our setting, where subjects were unable to see their plan following elicitation, we find dynamically inconsistent behaviour. We are not disregarding the experimental design of Ebert and Voigt (2023), in fact, combined with the current study, the two examinations have interesting implications for commitment devise design. As you will see shortly, we still observe relatively dynamically inconsistent behaviour in our condition in which semi-binding commitment was available. Whereas in the Ebert and Voigt (2023) assessment, where they indirectly impose a non-restrictive form of commitment by allowing subjects to visualise their plans during play, they find dynamically consistent behaviour. Perhaps the subconscious awareness of plans is more powerful than imposing semi-binding limits.

FIGURE 2.22: **Length of play.**

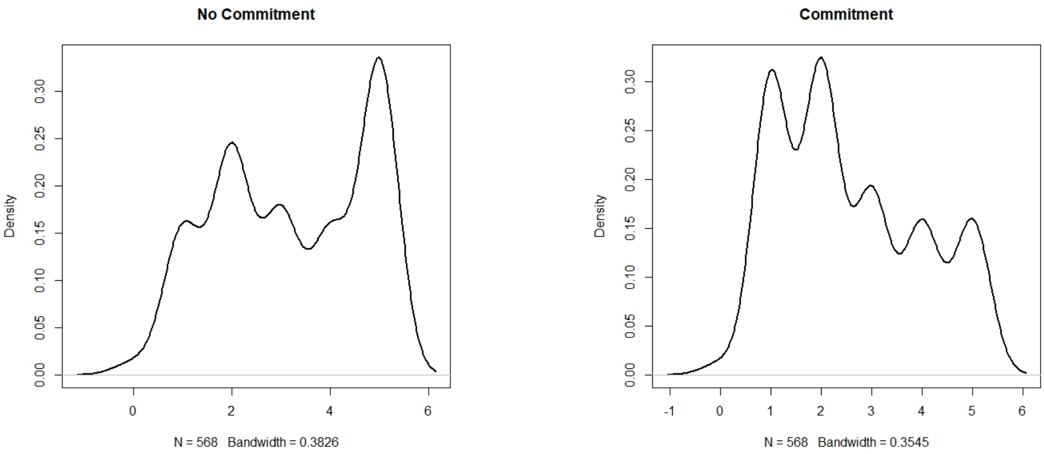

(A) No commitment

(B) Commitment

Notes: Density plot representing the horizon over which subjects' continued playing in both conditions: aggregated over all subjects. Panel A is for the No Commitment condition, and Panel B is the Commitment condition.

2.6.12 Role of Commitment

The existing literature has found conflicting results regarding the availability of commitment and whether this impacts a subjects' willingness to take on more risk. We now examine whether there is a relationship between commitment and playing for longer. Given that we find the majority of subjects in fact enter the casino (98% occurrence), and 83% of subjects opt for the commitment device, it is not surprising that subjects played for longer in the no-commitment rounds than the commitment rounds. We find that in the no-commitment rounds, the mean length of play per round was 3.26, whilst in the commitment condition, the mean length of play is 2.53, where the differences in means and medians per subject, per round, are statistically significant (p=0.0022, p<0.001). We find no difference between types in their mean risk-taking behaviour in the no commitment condition (p=0.915) or in the commitment condition (p=0.526). Figure 2.22 plots the kernel density estimates of length of play for both the no commitment condition and the commitment condition, for all subjects.

2.6.13 Dynamic Inconsistencies: C

In the commitment condition we now wish to identify whether subjects are exiting earlier because they are reaching their commitment bounds, or if they are deviating from their strategy by exiting before reaching these bounds.

Subjects played in n=8 rounds, where 59% (33/56) of naives, and 80% (12/15) of sophisticates deviated from their commitment device at least once. That being, they left the casino before they reached an exit-node from their initial strategy. Again, this is intuitive, given that the majority of subjects' elicited loss-exit strategies. With a loss-exit strategy in place, and with no-commitment, subjects' are likely to continue gambling in the loss domain in the hope that they can recover their losses over the longer time horizon. As the probability of recovering when already in the loss domain is already small, this probability is likely to be overweighted, resulting in optimism and riskier behaviour. However, with a commitment device restricting this horizon, if the low probability of a steady recovery has disappeared, and instead one is left with moderate probabilities of gains or losses, with maybe a single period left to recover their losses, this may lead to the cessation of gambling behaviour earlier on. Across all participants who opted for the commitment device, subjects deviated on a mean of 2.42 and median of 2 tasks out of a possible 8. However, we find that for those subjects who did not opt for the commitment device (all naive), they deviated on a mean of 4.25 and median of 4 tasks. Given that it is incredible hard to impose a hard-commitment device in a gambling setting, a semi-hard device was implemented, and was shown, on average, to reduce the average gambling time significantly (mean: $p=0.01$, median: $p<0.01$). Replicating the analysis of Figure 2.21, but in the commitment condition, we generate the violin plot of the Ebert and Voigt (2023) measure of dynamic consistency in Figure 2.23.

We can see there are still significant levels of deviation across all groups, where, inevitably, those who did not commit had dynamic consistency measure's concentrated between 0.3 and 0.675. For those who do commit, whilst the measure has the highest density around 1, we still observe a concentration of the measure between 0.5 and 0.875. This suggests dynamic inconsistency still prevailed, and therefore subjects deviated from their commitment devices, as aforementioned. Again, for

FIGURE 2.23: **Distribution of dynamic consistency measure: Commitment**

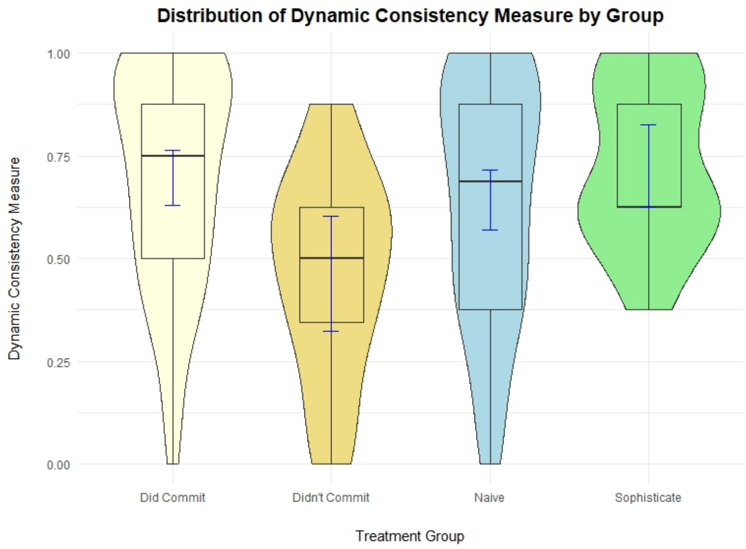

Notes: Distribution of dynamic consistency of four subgroups in the **Commitment condition**. In the first two groups, subjects are separated into those who commit and those who do not. The second two groups separate's subjects into those classified as naive, and those classified as sophisticate

our type classifications, there is substantial heterogeneity in the level of dynamic consistency observed. We observe large concentration levels and relatively uniform density's for our naive group between 0.3 and 1, and scattered densities for our sophisticate group.[17]

Our results have shown that differentiation between naives and sophisticates in this environment is ambiguous, in that for many comparisons, the two types, on average, act in a similar manner, with relatively large standard deviations in mean differences. We assume that the observed classification of types from our prediction exercise stems from the risk-preferences elicited coinciding either with their subsequent elicitation of strategies, **OR** their actual play given these parameters. Given that the design of the casino gambling model does not permit the direct elicitation of parameters, the true identification of types based on our parameter estimates could

[17]Additionally, we explore whether there is a link between subject-level risk-preferences and dynamic inconsistencies. We find little correlation for each preference parameter in each domain $-0.1 < cor < 0.2$, with p values greater than 0.4 for each. However, this is likely due to the heterogeneity in preference parameters elicited. Nonetheless we can conclude that we find large degrees of probability distortion, as well as large deviations from strategies. Additionally, the independence of the risk preference elicitation task from the casino task is a limitation we acknowledge, which may be contributing to the lack of correlation between risk preferences and dynamic consistency. Given it was not feasible to estimate parameters directly from the casino task, the analysis is somewhat restricted in this domain.

be restrictive. However, given the identified behaviour of subjects, we can make interesting insights into type-identification. Most notably, we propose the development of a new type: Quasi-naivety.

2.7 Quasi-naive

Evidently there are cases where the two types exhibit explicitly different behaviour, there are cases where there are no statistically significant differences, and in many cases, the theoretical assumptions of the model do not hold. Notably, so far we have identified that many naive subjects opt for commitment, sophisticates used backward induction for strategy elicitation or actual play, but did not tend to use it in both, and both types, but more importantly the naive agents, play for longer in the no-commitment round than the commitment round. Finally, we find that most subjects deviate from their initial strategies in the no-commitment setting, and in fact continue to deviate from their strategies, even when commitment is somewhat binding. Is it that the theoretical assumptions of the model are inaccurate, or do they merely need re-evaluation.

In this section we propose the development of a new classification, that provides an explanation for the results observed in the current study, as well those from existing studies (Heimer et al., 2023). This type stands as a middle-ground agent between the naive and a sophisticate, and is able to explain why in many cases, the theoretical predictions of the models attained large error rates. We call this type quasi-naive. We are not saying that purely naive subjects, and purely sophisticate subjects do not prevail. In fact we have found explicit support for these assumptions, in that there are those who did not opt for commitment, generated the same strategy in both conditions, and deviated from these strategies in actual play, thereby elicited very small prediction errors (0.14) with the naive model. Similarly, there were those who elicited strategies using backward induction in the no commitment condition, followed this initial strategy, opted for commitment, and elicited a forward-looking strategy which was followed ex-post, thereby eliciting small prediction errors (0.21) with the sophisticate model. However, we propose that "strictly naive" and "strictly sophisticate" classifications fall on the extreme ends of our type spectrum, with the majority of participants resembling a merger of the two types.

Quasi-naive subjects appear to be aware of their dynamic inconsistencies. They opt for commitment, yet still manage to deviate in light of updated preferences. In fact, they are likely to generate a strategy with commitment that leads to deviations from initial strategies that may not have occurred without commitment.

Even thought they are aware of their dynamic inconsistencies, a quasi-naive subject will devise a plan using a forward looking approach. They will enter the casino if they are able to generate a strategy that yields a positive expected return, given their preference parameter triples. Without the availability of commitment, quasi naive agents believe that "this time will be different" in that they are aware of their dynamic inconsistencies, but believe they have found a way to overcome them. Dai, Milkman, and Riis (2014) introduce how subjective temporal landmarks create new mental accounting periods and foster a "fresh start" effect, which is likely influencing the behaviour of quasi-naive agents. Inevitably, with the adrenaline of the risky task, their impulses take over and they proceed with a forward looking approach, in which the outcome distribution, as well as the probabilities of receiving the most desired outcome, change, leading to deviations from initial strategies. This is because subjects are acting on present time preferences, rather than their t=0 preferences. Up until this point, the strategies and actions will be the same as the original naive classification.

However, when commitment is available, quasi naive subjects do not need to take the risk of believing "this time will be different", and therefore opt for the commitment device. Bettega et al. (2023) and Kurth-Nelson and Redish (2012) have shown that in many cases, subjects have opted for commitment, even when it is costly, implying acknowledgement of dynamically inconsistent behaviour. Again, they generate a strategy using a forward looking approach, that they are now confident they will abide by, as they have a semi-binding commitment device to enforce this. This device provides the subjects with an illusion of safety, where they are reluctant to acknowledge deviations from strategies that are within the constraints of the device. They also fail to acknowledge that imposing a commitment device will alter the outcome distribution of the gamble, which means, as $t \to 5$, the outcome

FIGURE 2.24: **Example Commitment Strategy.**

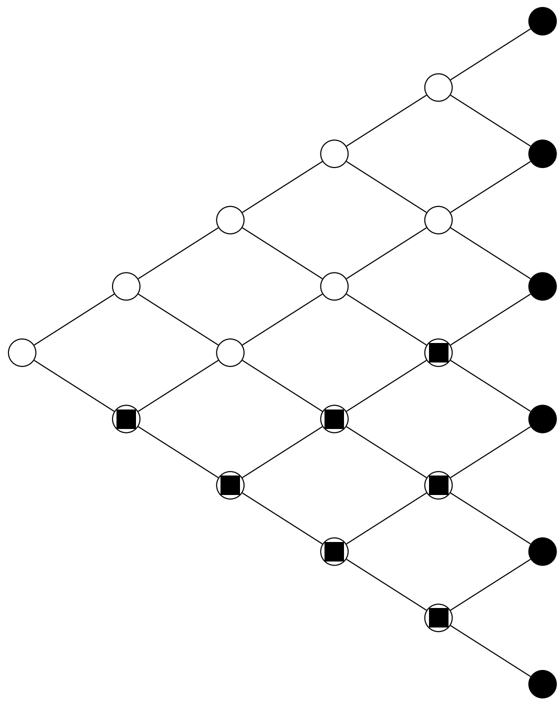

distribution of the game is not the same as it would be without the device. They have altered the outcome distribution of their gamble without knowing. For example, let us work with the strategy in Figure 2.24.

If our quasi naive subject elicits this strategy based on her parameter triple, then the expected outcome distribution, at a time of t=0 is (£-2, 0.625; £0 ,0.125; £2, 0.09375, £6, 0.125, £10, 0.03125). The quasi naive subject opts for the commitment device, and therefore these bounds are binding. As this strategy generates the highest utility given their parameter triple, they believe there will be no deviations from this plan, as the commitment device is there to restrict them from gambling in the loss domain for too long. However, as they still visualise the 5 period game in its original form, they neglect deviations from newly generated outcome distributions. In reality, the game has become Figure 2.25a.

Of course this is the intention of the quasi-naive, as they seek to generate an outcome distribution that yields a positive CPT value. However, if for example, the quasi-naive agent was to enter and win in t=0, and then lose in t=1, then their new position would be the red node in Figure 2.25b. At this point in t=2, the quasi naive agent looks to see if they can generate a strategy that provides a positive outcome distribution, given their parameter triples, and given the commitment bounds that

FIGURE 2.25: **Pruned Decision Tree.**

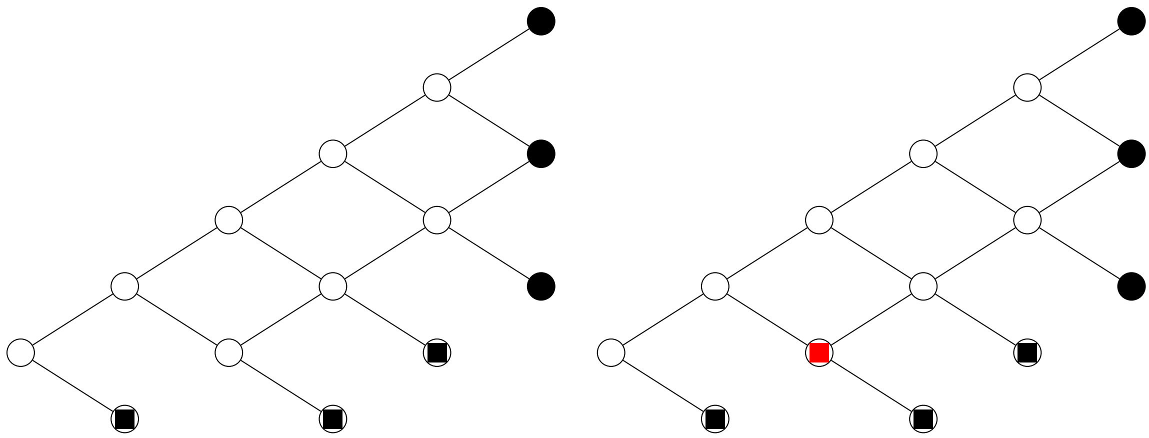

(A) Pruned Commitment Strategy

(B) Position after win in t=0, and loss in t=1.

Notes: Panel A represents the a visual representation of the pruned tree given a specific commitment strategy. Panel B illustrates where a subject would land following a win and then a loss.

restrict the full game. If they find a new strategy that generates a positive CPT value, then they continue playing. If they cannot, then they exit, even if their initial strategy with commitment was to continue. In their new position, they solve for an optimal strategy in the game represented in 2.26. The red node in Figure 2.25b is at the same position as the red node in Figure 2.26, only the latter has pruned the game further based on moves by nature.

At t=0, their optimal strategy suggests that they should play at this point in t=2, however, they now re-evaluate the new outcome distribution (-£2, 0.5; £0, 0.25: £2, 0.125, £6, 0.125), and in many cases will find their optimal strategy is now to exit. If for example we assume parameter triple (α, δ, λ) = (1.6,0.92,2). A quasi-naive agent with this parameter triple would elicit the optimal strategy illustrated in Figure 2.24. The utility of entering the casino in t=0 for this agent is CPT = 1.105. Therefore the quasi-naive agent enters with this strategy and opts for commitment. However, following a win and then a loss, they find themselves at the red node in figure 2.26. With a newly generated outcome distribution, they assess their CPT utility of continuing against their CPT utility of exiting, based on their parameter triple and their

FIGURE 2.26: **Pruned decision tree at t=2.**

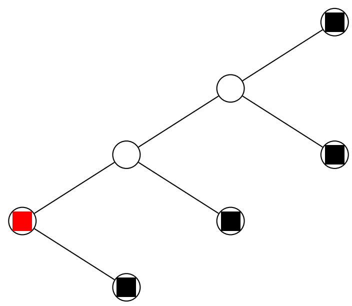

Notes: Following a pruning given the initial strategy, and then a win in t=0 and loss in t=1, the game now resembles a new decision tree.

current position. They find that the utility of playing, which their initial strategy advises them to do, is now CPT = -0.03. Whilst the utility of exiting at this point is 0. Therefore, the quasi-naive agent deviates from their initial commitment strategy and exits the casino. We can see this holds for a wide range of parameter triples over a wide range of commitment strategies. That being, where a subjects preference parameters generates an optimal strategy, commitment is imposed, but as the horizon shortens, subjects deviate from this strategy. This provides an intuitive explanation for the results found in our experiment, where subjects' classification into strictly naive or strictly sophisticate types was enigmatic. Similarly, it provides an explanation as to why subjects deviated from their soft commitment devices in the Heimer et al. (2023) study. Thaler and Johnson (1990) implicitly support this claim when they point out that people take on more risk after a loss, but only if the upside of the gamble allows them to recover from it and get back to the reference point. They refer to this behaviour as the "break even" effect.

Of course this is only relevant when there is a soft or semi-binding commitment device available, rather than a hard one. However, given most real-world applications of commitment devices are soft, this classification is pivotal. It is possible to impose a hard commitment device, but this would require a computer playing on ones behalf, and sticking to the elicited strategy strictly. However, this may change a quasi-naive subjects' mind regarding whether they wish to commit. The literature has provided research suggesting individuals require some level of control over financial choices, with Bettega et al. (2023) highlighting that many may feel uncomfortable with the idea of a commitment device that restricts future control. A semi-binding or soft commitment device provides this element of control, and the freedom to make ones own choices within the constraints of the device.

2.7.1 Welfare Implications

It is extremely difficult to devise a binding hard commitment plan in a gambling or financial setting, as most devises will still allow for deviations within the constraints of the commitment plan. Dependent on the structure of the task and the commitment device, semi-binding pre-commitment could result in increased risk taking or reduced risk-taking. In our setting, we find in fact that it reduced risky behaviour, in that subjects who opted for commitment exited the casino earlier in the commitment condition than the no-commitment condition. However, we also found that all subjects played in all tasks, so perhaps the inherent risky behaviour of subjects generates this result.

Of course, increased risk-taking and decreased risk-taking does not always correlate to higher or lower monetary gains or losses. Table 2.8 provides the summary statistics of the monetary outcome distribution of our subjects. Subjects completed eight rounds of play in each condition, whereby for each round, there was a monetary outcome. For each subject, we sum their cumulative winnings over each round and provide four realised moments of the monetary outcome distribution. The maximum cumulative payoff possible over 8 rounds is £80, the minimum was -£80, and the cumulative expected value over each gamble was 0. We analyse outcome distributions by segmenting subjects into different groups under various settings. Our first setting is in column 1, where we aggregate over the whole sample space, and provide the four moments in both the No Commitment condition (rounds 1-8), and the Commitment Condition (rounds 9-16). In the second setting, in columns 2 and 3, we segment subjects based on the Barberis (2012) classification of types, and again provide the four moments in both conditions. In the final setting, in columns 4 and

TABLE 2.8: **Four realized moments of the monetary outcome distribution.**

	All		Naïve		Soph		Didn't C		Did C	
Condition	NC	C	NC	C	NC	C	NC	C	NC	C
Mean[.]	-2.761	-0.958	-2.821	-0.714	-2.533	-1.867	-1	-0.833	-3.119	-0.983
Std[.]	5.538	3.077	5.756	2.735	4.809	4.103	3.954	1.337	5.769	3.329
Skew[.]	-0.019	-0.457	0.06	-0.274	-0.456	-0.269	0.647	-1.111	0.035	-0.4
Median[.]	-2	0	-2	0	-2	-2	-1	0	-2	0

Notes: Column 1 (ALL) is the whole sample space in the two conditions, column 2 and 3 are segmented into types, and column 4 and 5 are segmented based on their demand for commitment.

5, instead we segment subjects into two groups: those who did not opt for commitment in Part 2, and those who did, and provide the four moments of their outcome distributions in both conditions.

Again, for the differentiation between naive and sophisticate agents, we do not find any meaningful statistically significant differences in their outcome distributions, which is expected given their average lengths of play were similar. We are particularly interested in the welfare implications of commitment devices that are not strictly binding, so we now focus on the final two columns of Table 2.8. Interestingly, for those who decide not to commit, their losses are relatively low, and inevitably do not change much over the two conditions. However, for those who do opt for commitment, subjects are become better off in the commitment condition than in the no-commitment condition. Panels b and c of Figure 2.27 illustrate how the semi-binding commitment device shifted the concentration towards the positive domain and subsequently altered the direction of the skew. We are unable to identify whether or not the deviations from commitment in fact made subjects better or worse off, but this is perhaps something future considerations could look into. Additionally, it would be interesting to observe deviation behaviour in a setting where subjects may be less inclined to participate without a pre-commitment option. This would highlight whether semi-binding commitment hinders or benefits

FIGURE 2.27: **Monetary Outcome Distribution Concentrations.**

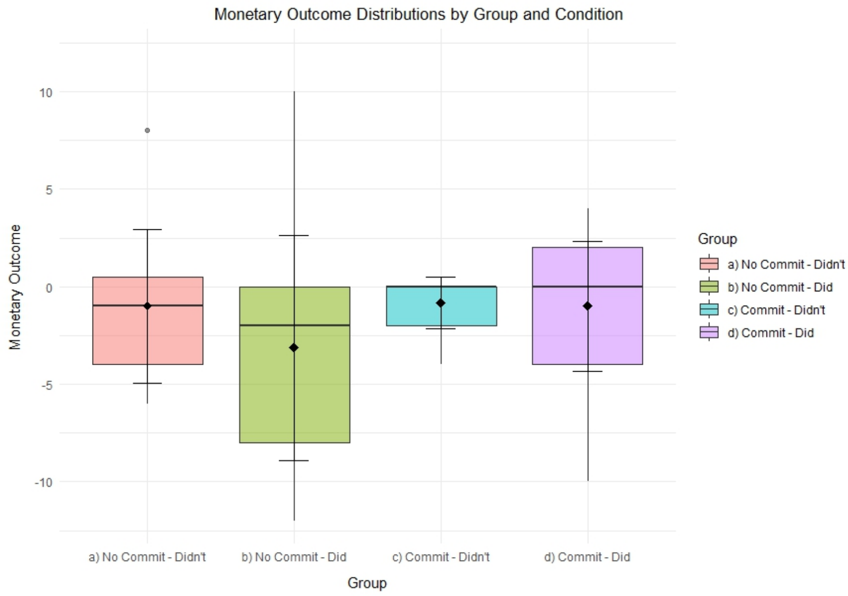

Notes: Subjects are segmented into those who have a demand for commitment (Did), and those who do not (Didn't). Results for both the no-commitment condition (No Commit), and the commitment condition (Commit).

financial welfare.

Nonetheless, regardless of the observed deviations from commitment within the constraints of the bounds, the commitment device still managed to shift the concentration of the monetary outcome distribution towards the gain domain.

Heimer et al. (2023) highlight that non-binding limits create an "illusion of commitment" whereby subjects overestimate the efficacy of the device. We observe this behaviour by the number of deviations from ex-ante commitment strategies. In our setting, these deviations did not hinder financial welfare, largely due to the risk-seeking attitudes adopted by our sample and that subjects were willing to gamble regardless of the availability of commitment. However these types of strategies could lead to increased risk-taking in environments where one may not usually feel comfortable engaging in the task. In these environments, it is likely that non-binding or semi-binding commitments could reduce financial welfare. For example, in our examination, aggregating over all participants, in the no-commitment condition we find the mean of the monetary outcome distribution is -2.761. Whilst not participating at all would yield a payoff of 0, subjectively benefiting the subjects financial welfare. The efficacy of commitment devices is therefore largely dependent on whether the individual would have engaged in the risky action anyway.

2.8 Discussion

Dynamic inconsistencies in decision-making under risk represent the inability to adhere, ex-post, to an ex-ante plan. This chapter sought to provide an explanation for dynamically inconsistent choices by being the first to directly test the empirical predictions of the Casino Gambling Model (Barberis, 2012). We show that direct parameter estimation from the theoretical design is not feasible, so we proceed with an experimental design that incorporates an independent risk-elicitation task, such that we can still predict the models primary assumptions. Of course, the static and independent environment in which parameter recovery takes place calls for questions concerning the heterogeneity of preferences in different tasks. Nonetheless, we are able to reveal fascinating insights regarding the structure of dynamically inconsistent choices, the generated distribution and preferences of ex-ante strategies, demand for pre-commitment, and how semi binding commitment devices affect financial welfare.

Regarding strategy type elicitation, we identify a preference for loss-exit strategies (sometimes referred to as right-biased rules), as has generally been found in the literature (Antler and Arad 2023; Ebert and Voigt 2023). When commitment is not available, we find significant deviations from these strategies ex-post, both in terms of strategy type deviations, as well as node level decisions. We find the majority of subjects acknowledged they acquired some form of dynamically inconsistent behaviour, in that they opted for a commitment device when available. The device was semi-binding, whereby subjects could not exceed the limits imposed by their strategy, but deviations within the bounds of the strategy was possible. We find that the majority of subjects deviate from their pre-commitment strategy at least once. We propose that "pure naivety" and "pure sophistication", defined as those who comply with the theoretical predictions of the CGM, do exist, although these types seem to represent the more extreme ends of the behavioural classification spectrum. The majority of subjects, rather, represent a merger of the two types.

The structural patterns observed from our investigation motivate the generation of theoretical predictions for a new behavioural type: Quasi-naivety. The theoretical assumptions imposed for a quasi naive agent are able to rationalise the behaviour observed by subjects in our experimental setting, as well as the experimental results from existing studies (Heimer et al. 2023; Bettega et al. 2023). Similarly, these assumptions provide support for the break-even effect of Thaler and Johnson (1990), whereby subjects take on more risk after a loss, but only if the upside of the gamble allows them to recover and converge back to the reference point.

We would also like to touch on the role of reinforcement learning and belief updating in experimental investigations. Existing studies have proven that there is evidence of learning throughout experimental investigations (Chen and Hsieh, 2011), which in turn changes subjects preferences over time. We find that a significant number of subjects changed their initial strategies across the two conditions. However, as we impose a within-subject design, we cannot directly assume this a by-product of commitment or sophistication. It could in fact be that subjects learn further from the game, and identify instances where they can optimise their strategies based on updated preferences. Additionally, we find that the majority of subjects opt for commitment, but again, we cannot impose that this is a result of inherent sophistication. As the commitment condition came after the no-commitment condition for all subjects, the substantial demand for commitment may have come from subjects updating their beliefs about their intrinsic dynamic inconsistencies, after experiencing deviations in prior stages. This in turn may have induced classification switches for many subjects, where a new demand for pre-commitment emerged. In future investigations it would be useful to control for ordering effects by counterbalancing the sample size, and allowing half of the subjects to engage in the commitment condition first, such that one can identify whether the demand for commitment was exogenous or endogenous. Similarly, it would be interesting to observe the entry rates for subjects who first had commitment available, but later had it taken away.

As aforementioned, there are various restrictions imposed by eliciting parameters via an elicitation task that is independent from the dynamic choice-task. For future research that identifies similar restrictions[18], we suggest the use of risk preference elicitation tasks that are nested within the core experimental task (reduction of compound lotteries). In our example, results may have been more robust if instead of using the certainty equivalent risk preference elicitation task, we had selected gambles from the outcome distributions of all feasible strategies. To elaborate, when determining the theoretical predictions of the CGM, we were required to generate the outcome distributions of all 801 feasible strategies. Selecting a sub-sample of these outcome distributions, and presenting them as static gambles, would increase the interdependence between parameter recovery and type identification.[19]

Our findings show that the availability of a commitment device, even when only semi-binding, increased financial welfare. However, we highlight that the efficacy of a commitment device is largely dependent on whether the individual would have engaged in the risky action in the first place. If the commitment device was the determinant in the tie-breaker between entering into the risky decision and not, then deviations from commitment could in fact reduce financial welfare with respect to the reference point. Finally, our results, combined with those of Ebert and Voigt (2023), provide interesting implications for the design of commitment devices. The combined results suggest that commitment devices that foster an intrinsic motivation may be more effective than extrinsically imposed commitments. This could also be an interesting avenue for future experimental investigations, with a between subject design, and treatment groups with commitment devices that vary based on their intrinsic or extrinsic nature, rather than varying solely on their level of restrictiveness.

[18] For example where parameter identification is not feasible

[19] Unfortunately, due to budget constraints, we were unable to re-run the experiment with the new design.

Chapter 3

Heuristics Unveiled: A Comparative Analysis of Toolbox Models and Prospect Theory in Risky Choice

Abstract

In an attempt to elucidate the classic violations of expected utility theory, the behavioural economics literature heavily relies on the influential work of Tversky and Kahneman (1992) and Tversky and Kahneman (1975), specifically the Cumulative Prospect Theory (CPT) model and the Heuristics-and-Biases program. While both approaches have significantly contributed to our understanding of decision-making under uncertainty, empirical evidence on the competing approaches remains inconclusive. In this study, we investigate the performance of each approach across a wide range of choice environments and increased cognitive load domains, encompassing gains, losses, time pressure, and complexity. Utilising data from various studies and employing Bayesian inference, we assess the performance of CPT in comparison to an adaptive cognitive toolbox model of heuristics. For subjects classified as toolbox decision makers, we examine the content (i.e., which heuristics) and the size of the toolbox (i.e., how many heuristics). Our findings reveal that as the choice environment objectively increases in complexity, individuals transition from using sophisticated expectation-based utility models to relying on a set of simplification heuristics for decision-making. We quantify the relationship between toolbox usage and complexity, showing a significant and positive correlation between the two. Furthermore, our results indicate that as task complexity rises, individuals tend to employ smaller toolboxes with fewer heuristics for decision-making.

3.1 Introduction

You are in a casino, and we ask you to choose red or black on a roulette wheel, what would your strategy for making this choice look like? What about if instead you are deciding on an insurance product to take out, where you need to consider the premium cost, coverage, excess, trust, and more. Do we use the same strategies to make decisions in both situations. With the former, it is fairly straightforward to calculate the expected return and objective probabilities associated with each choice. With the latter, it may be more complicated to evaluate the utility received from each component and assess it alongside the cost. So would we use simplification strategies to reduce the complexity of the decision, or would we still attempt to evaluate each component individually?

The behavioural economics literature on decision-making under risk and uncertainty relies heavily on Tversky and Kahneman (1992)'s Cumulative Prospect Theory to explain human choices. It has provided explanations to economic paradoxes (Allais, 1953), and can flexibly account for individuals' economic decisions. However, it remains inconclusive as to whether its assumptions remain as the most descriptive when the environment or characteristics of a task increase the complexity of the decision. This chapter seeks to determine whether CPT's explanatory dominance remains when subjects are presented with tasks that may overwhelm our cognitive loads. More specifically, we wish to determine whether subjects better explained by algebraic, expectation-based compensatory models like CPT (Hilbert 2011; Kahneman and Tversky 1979; Zindel, Zindel, and Quirino 2014), or simple rules of thumb (heuristics) (Tversky and Kahneman 1975; Gigerenzer and Gaissmaier 2011), when under time pressure and increased complexity.

The compensatory models predict overt decisions, combining components of risk attitudes such as utility curvature, probability weighting and loss aversion. Whilst heuristics are considered in order to capture the underlying cognitive process. These two approaches radically differ in their assumptions and the way they model decision making, and both have been extensively used in the literature, to explain a number of paradoxes, such as the Allais paradox, the four-fold pattern, the certainty effect, the possibility effect, and intransitivities.

Regarding the compensatory models, it is important to note that within CPT lies an indirect assumption that individuals have the time and capacity to subconsciously transform probabilities into decision weights, transform outcomes into utilities, and mentally endure the future effects of a potential loss. Many real-world economic choices do not provide the luxury of infinite time horizons, and some may not have the capacity to evaluate the utility received from all potential outcomes.

On the other hand, there are theories of human cognition assuming that people are equipped with a repertoire of heuristics and simplifying processes to solve the tasks they face in daily life. The literature has modelled this behaviour with the aid of a cognitive toolbox, from which people might adaptively choose their respective strategies. In the field of judgement and decision making, this concept was pioneered by Payne, Bettman, and Johnson (1993) arguing that the decision makers are equipped with a set of strategies and select among them when faced with a decision; an approach which was later extended in Gigerenzer (2002), modelling decision making as probabilistic draws from a toolbox of heuristic rules. While this modelling approach has been extensively investigated in the field of psychology and related studies in various domains, such as resource allocation, estimation and judgment of frequencies, skill acquisition, and learning processes, there is a notable absence of empirical evidence regarding the performance of these toolbox models of cognition, specifically in the context of risky choice. Scheibehenne, Rieskamp, and Wagenmakers (2013) highlight that while this theoretical framework of a cognitive toolbox provides a plausible account of intra- and interindividual differences in human behaviour, it is often unclear how to rigorously test the toolbox framework, how to quantitatively specify such a model, how to limit the number of toolbox strategies in the model to avoid the so called *strategy sprawl*, and how to formally test against alternative theories. To address these issues comprehensively, the authors propose the utilisation of Bayesian inference techniques. In this chapter, we adopt the statistical framework proposed by Scheibehenne, Rieskamp, and Wagenmakers (2013) to rigorously assess the performance of an adaptive toolbox model of heuristics in the field of risky choice compared to the benchmark CPT model. We aim to explore how individuals utilise either compensatory models or simple rules of thumb in various decision-making domains that may increase ones cognitive load, specifically in the gain domain, the loss domain, under time pressure, and when subject to overly complex tasks.

Interestingly, the the use of heuristics strategies is implicity supported by existing evidence in the eye-tracking data literature. Multiple studies have used process data to investigate and characterise decision strategies (Venkatraman, Payne, and Huettel 2014, Harrison and Swarthout 2019, Fiedler and Glöckner 2012). Noteably, Arieli, Ben-Ami, and Rubinstein (2011) find that the eye patterns of individuals suggest we compare prizes and probabilities separately, which opposes more holistic cognitive bias model approaches like CPT. This evidence coincides with more heuristic-type behaviour. It was also found that this effect was more pronounced when the suggested weighting of probabilities and utilities are more laborious to compute. Nonetheless, there are studies that say the eye tracking data opposes certain heuristics (Glöckner and Herbold, 2011), but these studies tend to also rule out the idea that individuals look to maximise any form of expectations models. However, they impose the assumption that individuals rely on a single heurstic rather than a repetoire of heuristics - an assumption we relax.

3.1.1 Research Questions

Our assessment of the explanatory performance of CPT and an adaptive toolbox of heuristics in tasks of varying complexity will answer the following research questions.

Research Question 1. *Can a toolbox model of simple heuristics explain lottery choices better relative to a sophisticated compensatory utility model (i.e. CPT)?*

Recent literature has primarily focused on one-to-one comparisons between a single heuristic and a flexible model with free parameters (most commonly CPT), providing overwhelming support in favour of the latter (Brandstätter, Gigerenzer, and Hertwig 2006; Rieskamp 2008; Glöckner and Pachur 2012; Balcombe and Fraser 2015; Peterson et al. 2021). Nevertheless, this approach ignores the concept of *ecological rationality*, that is, the fit between a heuristic or decision strategy and a choice environment which gives the agent the flexibility to adapt her strategy according to the decision task at hand. Various heuristics have been shown to take place simultaneously, so to better understand the behaviour of economic agents, it is imperative to determine which heuristics are used in which circumstances (Campo et al., 2016).

In the field of risky choice, there is a notable scarcity of research examining the performance of toolbox models. Stahl (2018) investigates whether a toolbox model of simple heuristic rules can help explain choices under risk relative to expected utility theory (EUT). The study concludes that *[$\cdots$] if we want to forecast the future lottery choices of humans and have limited prior data on which to make those forecasts, then our analysis suggests it would be better to use the Expected Utility Theory-Only model even if we believe it is not the true data generating process rather than using an overfitted toolbox model*. Mohnert, Pachur, and Lieder (2019) develop a model according to which the decision maker selects a decision strategy for a given choice problem rationally from a toolbox of strategies and they estimate the content of the toolbox at the individual level. Their adaptive toolbox model predicted people's risky choices better than single strategies and non-adaptive toolbox models, but performed worse than CPT. Finally, Olschewski and Rieskamp (2021) explore whether time pressure motivates subjects to use simple, non-compensatory strategies in a risky choice experiment. They assume 3 potential heuristics, and they find a slight but insignificant increase in the number of participants who resort to heuristics under time pressure, attributing this to increased noise in subjects' behaviour[1]. Nevertheless, prior efforts to assess toolbox models face obstacles stemming from simplification assumptions and econometric challenges in accurately characterising and identifying the model. These simplification issues arise in two primary forms. Firstly, with the exception of Mohnert, Pachur, and Lieder (2019), existing studies have not allowed for individual heterogeneity in toolboxes, implying that all subjects employ the same limited set of 3 or 4 heuristics. The second concern, raised by Stahl (2018), revolves around overfitting induced by the estimation method. Notably, all the aforementioned studies employ Maximum Likelihood Estimation (MLE) techniques, which are known to

[1] A relevant study is He, Analytis, and Bhatia (2022) focusing on the recent literature of collective model wisdom (model crowds) in decision analysis. It conducts a large scale comparison of 58 prominent models of risky choice and they find that crowds of risky choice models perform better than individual models. While they include a large number of popular heuristics in the candidate models, they do not explicitly test toolbox models of heuristics and their approach differs from ours in several aspects (i.e.type of lotteries, model selection, estimation method, CPT specifications, and number and type of heuristics).

produce estimates that are extreme, relatively noisy, and less reliable when compared to more flexible estimation methods, such as Bayesian Hierarchical modelling or Simulated Maximum Likelihood, that account for behavior at both individual and population levels. This leads us to our second question:

Research Question 2. *What is the best way to quantitatively specify and robustly estimate a toolbox model of cognition?*

Stahl (2018) using both simulated and actual data, suggests that it may be preferable to use the EUT-Only model, even when it is not believed to be the true data generating process, rather than employing an overfitted toolbox model. Building on this insight, we re-analyse the same datasets and investigate two potential explanations for this result (1) the choice of estimation method and (2) the nature of the data. Regarding the former, Bishop (2006, pp. 166) cautions that "the use of maximum likelihood, or equivalently least squares, can lead to severe over-fitting if complex models are trained using data sets of limited size", and goes on to suggest "the phenomenon of over-fitting is really an unfortunate property of maximum likelihood and does not arise when we marginalize over parameters in a Bayesian setting". Using Bayesian inference, we demonstrate that the overfitting issue is significantly mitigated, resulting in more robust and reliable estimates. With regards to the second point, we explore how the nature of the dataset, and in particular the experimental design, can explain the poor forecasting performance of the toolbox model. The next question we explore can be summarised as:

Research Question 3. *What is the relationship between increased cognitive load and the utilisation of cognitive toolboxes?*

Stahl (2018) analyses the data of Hey and Orme (1994) and Harrison and Ruström (2008) which consist of 3-outcome binary lotteries, in the gain domain, with only four potential monetary outcomes and varying probabilities. We argue that this design and decision environment, may be more "user friendly" for expectation based models (i.e. Expected Utility). We extend our comparison of CPT and Toolbox models in various domains and conditions, specifically looking to explore whether decision-makers resort to the use of heuristics when the tasks at hand are more cognitively demanding. We therefore, explore the performance of the toolbox model, in a number of domains/environments that require more cognitive effort, including losses and mixed gambles, time pressure and complexity. Using features of a dataset such as the number of alternatives in a choice set, the formatting of probabilities and outcomes, or the distribution moments, we set up an index of complexity which enables us to quantify the relationship between increased cognitive load and the utilisation of cognitive toolboxes, and show how the two are correlated.

Given the flexibility that our econometric approach offers, we are able to estimate a large set of models per subject, and identify the combination of number and type of heuristics that best describe the behaviour of each individual subject. The next two questions focus on the size and the content of the toolbox model.

Research Question 4. *What is the optimal number of heuristic strategies that one should include in a toolbox?*

As aforementioned, the previous literature on toolbox modelling makes the simplification assumption that all the subjects share the same limited and fixed set of heuristics (3 or 4). We relax this hypothesis and allow the toolbox size to vary between 2 and 5 heuristics, accommodating individual-specific combinations of size and heuristic number. Furthermore, the toolboxes can comprise various combinations drawn from a comprehensive set of 11 heuristics extensively studied in the literature. Next, we inquire into the contents of these toolboxes, identifying which heuristics are most frequently employed and what insights they provide into the subjects' risk preferences.

Research Question 5. *What is inside the toolbox?*

After identifying all the subjects for whom the toolbox model is the best specification, we can not only assess the size of the toolboxes (i.e., the number of heuristics) but also delve into the content of these toolboxes. This allows us to determine which heuristics are most frequently employed in different environments and gain insights regarding the risk preferences of the subjects.

To summarise, we estimate toolbox models of cognition using data from prominent studies in various domains and contexts (i.e., gain, loss and mixed domains, time pressure, and complexity). These toolbox models comprise a comprehensive set of 11 heuristics, and we compare them with four different CPT model specifications. We identify the domains and environments in which decision makers resort to the use of simplifying strategies, with the percentage of subjects classified as Toolbox decision makers, ranging from 2 to 67%. We find that for the subjects classified as Toolbox decision makers, the majority are using 3 or 4 strategies. Furthermore, we observe extensive heterogeneity regarding the type of heuristics subjects employ, with the vast majority resorting to heuristics that provide a safety net (e.g. Minimax or Least Likely), while risk seeking strategies are least preferred (e.g. Maximax). We introduce a metric to gauge the complexity of experimental stimuli, and based on this index, we observe a positive and significant correlation between complexity and the utilisation of heuristics in addressing risky choice problems. Finally, our analysis allows us to identify the domains/environments where CPT can explain behaviour better, which allows us to run a racehorse between different probability weighting functions, showing that a two-parameter weighting function always performs better compared to the one-parameter family of weighting functions.

The rest of the chapter is organised as follows. Section 4.2 outlines the theoretical frameworks of the the two competing models, section 3.3 provides the details of our econometric approach, section 3.4 briefly presents the datasets we employ, while section 3.5 presents the results. We then conclude. Please see Appendix B.1 for a comprehensive review of the literature on heuristics, the four domains of consideration, and the use of adaptive toolboxes.

3.2 Theoretical Framework

In this section, we present the underlying assumptions of the two decision making models that we estimate. In our analysis, the value of the outcomes can be positive, zero or negative, depending on the context and the environment that we explore. We present the full specification (i.e. the one that accounts for both gains and losses) which we adapt accordingly. We present both the deterministic assumptions of each model, as well as the stochastic assumptions we make, in order to take into account noise and heterogeneity in decision making.

3.2.1 Cumulative Prospect Theory

The decision maker faces pairs of n-outcome lotteries with outcomes $x_1 \leq \ldots \leq x_k \leq 0 \leq x_{k+1} \leq \ldots \leq x_n$ and corresponding probabilities $p_1 \ldots p_n$. Following Tversky and Kahneman (1992) we assume that a decision maker is endowed with a utility function $u(.)$, on monetary outcomes with $u(0) = 0$, and a probability weighting function $w(p)$, that transforms the objective probabilities into subjective decision weights. The overall evaluation of a lottery L is given by:

$$V(L) = \sum_{i=1}^{k} u(x_i)\pi_i^- + \sum_{j=k+1}^{n} u(x_j)\pi_j^+ \tag{3.1}$$

where π^+ and π^- are the decision weights for gains and losses, respectively. The decision weights are defined as:

$$\begin{aligned}
\pi_1^- &= w^-(p_1) \\
\pi_n^+ &= w^+(p_n) \\
\pi_i^- &= w^-(p_1 + \ldots + p_i) - w^-(p_1 + \ldots + p_{i-1}) \text{ for } 1 < i \leq k \\
\pi_j^+ &= w^+(p_j + \ldots + p_n) - w^+(p_{j+1} + \ldots + p_n) \text{ for } k < j < n
\end{aligned}$$

We assume a CRRA utility function over monetary outcomes of the following form:

$$u(x) = \begin{cases} \frac{x^r}{r} & \text{if } x \geq 0 \\ -\lambda\frac{(-x)^r}{r} & \text{if } x < 0 \end{cases} \tag{3.2}$$

where $r \geq 0$ is a parameter governing the utility curvature, and $\lambda \geq 1$ is the parameter of loss aversion. Previous studies have shown that the power function fits well experimental data, for the level of monetary payoffs usually used in experimental studies (see for example Stott 2006; Balcombe and Fraser 2015; Baillon, Bleichrodt, and Spinu 2020). When we consider losses, in order to avoid the scaling issues that a domain specific *power* (CRRA) function causes to the estimation of the loss aversion parameter, we follow Nilsson, Rieskamp, and Wagenmakers (2011) and assume the same power coefficient for both gains and losses (see Köbberling and Wakker 2005;

Wakker 2010; Harrison and Swarthout 2021). For the probability weighting function, we consider four specifications, the Prelec (1998) one-parameter function:

$$w(p) = \exp(-(-\log(p))^{\gamma}) \tag{3.3}$$

the Prelec (1998) two-parameter function:

$$w(p) = \exp(-(-\log(p))^{\gamma})^{\delta} \tag{3.4}$$

the Tversky and Kahneman (1992) function:

$$w(p) = \frac{p^{\gamma}}{(p^{\gamma} + (1-p)^{\gamma})^{1/\gamma}} \tag{3.5}$$

and the Goldstein and Einhorn (1987) probability weighting function:

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

All of the aforementioned specifications allow for an inverse-S shape of the weighting function, with overweighting of low probabilities and underweighting of moderate to high probabilities. The two-parameter families of the weighting function have the advantage of decomposing probability weighting to both its degree of curvature and its elevation. For instance, in Equation 3.6, $\delta > 0$ measures the elevation while $\gamma > 0$ measures the degree of curvature of the weighting function (likelihood insensitivity). As δ increases, the function becomes more elevated (exhibiting less overall risk aversion for gains and more for losses). On the other hand, the smaller $\gamma < 1$, the more curved the probability function, which means that the range of intermediate probabilities becomes flatter and therefore, exhibiting more rapidly diminishing sensitivity to probabilities close to the boundaries 0 and 1. As mentioned before, we allow the parameters of the weighting function to differ between domains, while we allow for both S and inverse-S shapes of the function.

To model stochastic choice we assume a logistic Luce (1959) choice rule such that the probability of choosing lottery A is given by:

$$P(A > B) = \frac{1}{1 + \exp(s(V(B) - V(A)))} \tag{3.7}$$

with $s > 0$ a choice sensitivity parameter, indicating how sensitively the predicted choice probability reacts to differences in the utility valuations of the two lotteries on a cardinal scale, and $V(.)$ the CPT value of the respective lottery.

In the case of lotteries defined purely in the gains domain we take into consideration the *Contextual Utility* Wilcox (2011) and the predicted probability becomes:

$$P(A > B) = \frac{1}{1 + \exp(s(\frac{V(B)-V(A)}{\nu}))} \tag{3.8}$$

where the normalising term ν is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair. That is, the difference between the two lotteries is relative to the range of outcomes found in the lottery pair. We also note that in the gains domain the model is equivalent to the Quiggin (1982) Rank Dependent Utility model.

3.2.2 Cognitive Toolbox

Many theories of human cognition assume that people are equipped with a repertoire of heuristics and simplifying processes to solve the tasks they face in daily life. In the literature, this idea has been theoretically modelled with the aid of a cognitive toolbox, from which people might adaptively choose their respective strategies. In the field of judgement and decision making, this idea was pioneered by Payne, Bettman, and Johnson (1993) arguing that the decision makers are equipped with a set of strategies and select among them when faced with a decision, approach which was later extended in Gigerenzer (2002), modelling decision making as probabilistic draws from a toolbox of heuristic rules.

Following Scheibehenne, Rieskamp, and Wagenmakers (2013), a toolbox model can be represented by a set of different psychological processes or strategies f, and each strategy predicts a particular course of action, depending on the environment or the *ecology* of the domain upon which decisions are made. Independent of the mechanism behind the strategy selection, the outcome of this process can be modelled with the aid of a mixture proportion parameter β which indicates the probability of choosing each strategy in the toolbox. For instance, for a particular toolbox TB consisting of J strategies, each strategy f_j will be selected with probability β_j, with $\sum_{j=1}^{J} \beta_j = 1$. For instance, a potential toolbox with 4 strategies would be defined as:

- Pick the lottery with the highest payoff (MAXIMIN) with probability β_1
- Avoid the lottery with the lowest payoff (MINIMAX) with probability β_2
- Pick the lottery with the highest most likely payoff (MOST LIKELY) with probability β_3
- Pick the lottery with the highest probability of the highest possible payoff (MOST PROBABLE) with probability $1 - \sum_{i=1}^{3} \beta_i$

This modelling specification allows for the underlying cognitive process of strategy selection to remain unspecified, given that the value of the parameter vector β will be estimated by the data, providing the empirical validation of the latent strategy mix. Given this mixture specification, the compound probability of choosing lottery A can be specified based on the sum of the individual likelihoods of each f_j, weighted by the mixture probability β_j:

$$p(A|TB) = \sum_{j=1}^{J} [\beta_j \times P(A|f_j)] \tag{3.9}$$

where $P(A|f_j)$ is the individual predicted probability of each strategy. Since the heuristics generate ordinal choice propensities (i.e. deterministic), we follow Balcombe and Fraser (2015), Rieskamp (2008) and He, Analytis, and Bhatia (2022) and assume a *constant-error* choice rule to capture stochastic choice in the data[2]. The constant-error specification, has been widely used in the game-theoretical literature (i.e. trembling hand) and it has been populated by Harless and Camerer (1994) in the

[2]Since a heuristic choice rule predicts only deterministically, there is lack of clarity on how the deterministic prediction of the theory translates into a probability of observing one choice or the other. Andersen et al. (2010) discuss ways of how one can modify the Priority Heuristic to make it worth testing against any real data, but they conclude that any modification of this kind would be contrary to main purpose of the model. We therefore feel that the constant-error stochastic rule is the most natural one can assume.

context of risky choice. Since the heuristics generate ordinal choice propensities (i.e. deterministic), we assume a *constant-error* choice rule to capture stochastic choice in the data, where the decision maker chooses with constant probability $1 - \varepsilon$, the option that the heuristic prescribes, and with probability ε she makes a mistake[3]. The overall likelihood for a given subject is therefore the product, across all the tasks, of the weighted sum of predicted probabilities across the number of strategies in a given toolbox. The next modelling choice we need to make, is how many heuristics to include in a toolbox. The previous literature has assumed an arbitrary, fixed number of heuristics, which is the same for all the subjects. Nevertheless, Scheibehenne, Rieskamp, and Wagenmakers (2013) discuss how restricting the repertoire to only a few strategies would ignore any *intra* and *inter*-individual differences through qualitatively different processes, and how it can lead to the strategy sprawl problem, if one assumes too many strategies for a particular subject. In our analysis, we aim to identify the optimal toolbox for each subject, both in terms of size (how many strategies) and in terms of content (which strategies). The process we adopt is as follows. First, we adhere to Glöckner and Pachur (2012) and investigate the performance of 11 heuristics as potential components of a cognitive toolbox[4]. Then, we calculate all the potential combinations of heuristics. A set of n elements has $2^n - 1$ potential subsets when the null subset is not taken into consideration. This means that if one considers all the toolboxes of any size (ranging from toolboxes with only 2 heuristics to toolboxes with all 11 available) it gives in total 2036 potential toolboxes (excluding the null toolbox and the toolboxes with only one heuristic available). In order to reduce the number of models to estimate, we consider toolboxes of size up to 5, giving a total of 1012 toolbox models[5]. There are different types of heuristics focusing on one or multiple attributes. Some of the heuristics focus exclusively on the monetary payoffs, such as the Minimax, the Maximin or the Better than Average heuristic (outcome heuristics), while others focus on a combination of payoffs

[3] This is the part $P(A|f_j)$ in Equation 4.4.

[4] The full list of the heuristics along with a description of the choice mechanism behind each heuristic is provided in Table C.1.

[5] In particular, with 11 available heuristics, there are 55 potential combinations for a toolbox of size 2, 165 for a toolbox of size 3, 330 for a toolbox of size 4, and, 462 for a toolbox of size 5. For example, with 4 available alternatives J with $j \in \{A, B, C, D\}$ one can form the following toolboxes: (1) nothing, (2) A or B or C or D, if the size of the toolbox is 1, (3) AB or AC or AD or BC or BD or CD, if size is 2, (4) BCD or ACD or ABD or ABC, if the size is 3, and; (4) ABCD, if the size is 4.

and probabilities (dual heuristics), such as the Least Likely, the Most Likely or the Probable (Brandstätter, Gigerenzer, and Hertwig, 2006). Finally, there are heuristics for multiple-attribute choice which include the Lexicographic (Gigerenzer and Goldstein 1996: Tversky 1972), the Priority (Hill, Raacke, and Park 2017; Todd et al. 1999) and the Tallying heuristic (Parpart, Jones, and Love 2018; Czerlinski, Gigerenzer, and Goldstein 1999; Dawes 1979). The latter follow Rubinstein (1988) three-step model, where the agent applies an algorithmic process of decision making, going through various degrees of reason, and if two options are similar in terms of one reason (e.g. dominance) attention is shifted to other reasons (e.g. similarity). *A priori*, we expect that decision makers who are using a toolbox to decide, will delegate a small number of heuristics due to cognitive, time or other limitations. Mohnert, Pachur, and Lieder (2019) provide support in favour of this modelling choice as they find that the majority of the estimated toolboxes are of size 4, while research on model crowds indicates that a number close to five models is optimal (see Makridakis and Winkler 1983; Ashton and Ashton 1985; He, Analytis, and Bhatia 2022.) and adding further models diminishes the prediction capacity of the select crowd. In addition, we do not take into consideration any toolboxes of size one, given that the aim of this chapter is to relax the assumption that subjects are using a single heuristic and to avoid this kind of one-to-one comparison that previous research has investigated. In the next section we describe both the estimation and the model selection method that we adopt.

3.3 Bayesian Hierarchical Modelling

There are various ways to estimate structural decision making models. The most common approach is the use of subject level Maximum Likelihood Estimation techniques (MLE). Nevertheless, MLE may generate noisy and unreliable estimates, and can therefore produce extreme estimates for some subjects if number of observations is limited. In addition, MLE is susceptible to overfitting and may adjust mostly noise rather than the actual preferences of the subject, leading to very poor predictive performance of the models. An alternative way is to pool all the data together and estimate a model for the representative agent, assuming a particular preference functional. Nevertheless, several studies have provided evidence against the assumption of a single data generating process and proposed the use of finite mixture models instead (see Harrison and Rustrom 2008; Fehr-Duda et al. 2010; Conte, Hey, and Moffatt 2011; Alam, Georgalos, and Rolls 2022). While it is a useful approach in that it allows one to test the presence of more than one preference functional, they rely on the extreme assumption of the presence of n representative agents, one for each assumed preference functional (that is for instance, all EUT subjects share the same behavioural parameters, all CPT the same and so on).

Scheibehenne, Rieskamp, and Wagenmakers (2013) provide arguments of how the Bayesian formalism can allow toolbox approaches to be rigorously tested via the use of the *Bayes Factor* (Kass and Raftery, 1995), a unifying comparison metric that quantifies the extend to which data support one model over another, taking model complexity into account. To mitigate the drawbacks of MLE, we adopt hierarchical Bayesian estimation techniques (see Balcombe and Fraser 2015; Ferecatu and Önçüler 2016; Baillon, Bleichrodt, and Spinu 2020; Alam, Georgalos, and Rolls 2022 and Gao, Harrison, and Tchernis 2022 for some recent applications of hierarchical Bayesian models for choice models under risk and Stahl (2014) for ambiguity models). The key aspect of hierarchical modelling is that even though it recognises individual variation, it also assumes that there is a distribution governing this variation (individual parameter estimates originate from a group-level distribution). A hierarchical Bayesian model simultaneously estimates the individual level parameters, along with the hyper-parameters of the group level distributions. In typical hierarchical models, the estimates of the low level parameters are pulled closer together than they would in the absence of a higher-level distribution, leading to the so called *shrinkage* of the estimates.

As Baillon, Bleichrodt, and Spinu (2020) highlight, Bayesian Hierarchical modelling is a compromise between a representative agent and subject-level type estimation. It estimates the model parameters for each subject separately, but it assumes that subjects share similarities and draw their individual parameters from a common, population level distribution. In that way, individual parameter estimates inform each other and lead to a *shrinkage* towards the group mean that reduces biases in parameter estimates.

We follow the Rouder and Lu (2005) and Nilsson, Rieskamp, and Wagenmakers (2011) set-up and we estimate all the specifications using BHM. Each subject i made a series of N binary choices in a given dataset and the observed choices vector is denoted by $D_i = (D_{i1} \cdots D_{iN})$. Every subject is characterised by its own parameter vector B_i and we assume that all the parameters are normally distributed ($b_i \sim N(\mu_b, \sigma_b)$), while for the hyper-parameters we assume normal priors for the mean μ_b and uninformative priors (uniform) for σ_b. For the mixture parameter vector β in the toolbox model, since it represents a probability distribution and the parameters in β are not independent from each other, we assume that it is a J-dimension categorical variable which follows a Dirichlet distribution $\beta \sim Dirichlet(\pi)$ with π a diffuse hyper prior parameter for the distribution. We also follow the standard procedure and transform all the parameters to their exponential form to ensure that they lie within the appropriate bounds.

The likelihood of subject's i choices is given by:

$$P(D_i|B_i) = \prod_{n=1}^{N} P(D_{i,n}|B_i)$$

where $P(D_{i,n}|B_i)$ is the predicted probability for each lottery pair n, as this was presented in the previous section. Combining the likelihood of the observed choices and the probability distribution of all the behavioural parameters, the posterior distribution of the parameters is given by:

$$P(B|D) \propto P(D|B) \times P(B)$$

with $P(D|B)$ being the likelihood of observed choices over all the subjects and $P(B)$ the priors for all parameters in the set B. Monte Carlo Markov Chains (MCMC) were used to estimate all the specifications. The estimation was implemented in JAGS (Plummer, 2017). The posterior distribution of the parameters is based on draws from two independent chains, with 50,000 MCMC draws each. Due to the high level of non-linearity of the models, there was a burn-in period of 25,000 draws, while to reduce autocorrelation on the parameters, the samples were thinned by 10 (every tenth draw was recorded). Convergence of the chains was confirmed by computing the $\hat{R}$ statistic (Gelman and Rubin, 1992).

All the inference and the subsequent comparison of the models is based on the log *Bayes Factor* measure (Kass and Raftery, 1995). Bayes factors penalise models with a large number of parameters, prevent over-fitting, and are a good measure of the forecasting capacity of each model. The Bayes Factor is defined as $\exp(LML_{FULL} - LML_{EUT})$, where LML_i denotes the log-marginal likelihood of model i. To estimate the log-marginal likelihoods we use the Newton-Raftery estimator (harmonic mean of the log-likelihood over all draws after the burn-in period, Newton and Raftery 1994).

3.4 Data Sets

We estimated the models using a wide range of data sets from experimental studies covering all domains of gains, losses and mixed gambles, as well as cognitive loaded environments such as time pressed decision making and complexity. Table 3.1 provides a summary of all the data sets we use for our analysis. All data sets were chosen in a way that they would satisfy the following criteria:

- The experimental designs have been developed with the objective to estimate structural econometric models and involve a substantial number of tasks per participant.
- All studies have been incentivised in monetary terms.
- The studies incorporate a wide range of probabilities and outcomes, introducing various levels of complexity.

TABLE 3.1: **Summary of the Data Sets.**

Data Set	Abbreviation	Domain	Cognitive Load	No. of Subjects	No. of Tasks	Payoff Structure	Probability Structure	Complexity Levels
Baillon, Bleichrodt, and Spinu (2020)	BBS20	G	N	139	70	1-4 outcome Random	Random	Medium
Hey and Orme (1994)	HO94	G	N	80	100	3 outcome Fixed	Constant Multiples of 0.125	Low
Glöckner and Pachur (2012)	GP12	G/M/L	N	66	138	2 outcome Random	Random	Medium
Harrison and Swarthout (2021)	HS21	G/M/L	N	175	100	1-3 outcome Fixed	Based on the Marschak-Machina triangle	Low
Olschewski and Rieskamp (2021)	$OR21_{TP}$	TP	Y	60	150	2-4 outcome Random	Random	High
Moffatt, Sitzia, and Zizzo (2015)	MSZ15	Complexity	Y	80	54	3-27 outcome Fixed	Combination of simple lotteries with $p \in \{0.2, 0.3, 0.5\}$	High

Notes: G indicates the gains, L the losses and M the mixed domain. TP stands for the time pressure condition.

In the gain domain, we analyse experimental data from two studies; Hey and Orme (1994) and Baillon, Bleichrodt, and Spinu (2020). Hey and Orme (1994) involves 80 subjects deciding over 100 pairwise choice questions. There are 3 outcomes in each of the lotteries, and these outcomes are held fixed at either £0, £10, £20 and £30, whilst the associated probabilities are all multiples of 0.125. Baillon, Bleichrodt, and Spinu (2020) data include the choices of 139 participants from 70 binary lotteries. Each option in a lottery (Option A or B) has between one and four possible outcomes (all framed as gains). The payoffs and probabilities in their experiment have been carefully chosen to maximise statistical efficiency and minimise redundancy. On top of allowing for a wide range of number of outcomes and magnitude of probabilities, the authors ensured that within each choice pair there were non-matching maximal or minimal outcomes, questions had similar expected value, and finally, questions were orthogonal to maximise statistical efficiency. It can be argued that the tasks in Hey and Orme (1994) are relatively less complex than those in Baillon, Bleichrodt, and Spinu (2020), as the fixed probabilities and payoffs permit for a simpler design and thus a less computationally challenging cognitive process. Furthermore, the way the maximum and minimum outcomes have been chosen in Baillon, Bleichrodt, and Spinu (2020)[6] that may generate the appropriate *ecology* for an adaptive toolbox.

In the loss domain, we use the data from Glöckner and Pachur (2012) and Harrison and Swarthout (2021) for our analysis. Glöckner and Pachur (2012) include the choices of 66 subjects in 138 pairwise choice problems, where 70 are gains, 30 are losses, and the rest are mixed. The problems involve binary two-outcome lotteries. The tasks were a combination of lottery pairs that have been used in previous studies to capture various decision making phenomena with tasks being either randomly generated, designed to differentiate between the priority heuristic and CPT, designed to measure risk attitudes using the Holt and Laury (2002) task, or designed to measure loss aversion. They generated two sets of 138 tasks and participants faced either of each sets in two separate sessions. We are using the data from session 1. Harrison and Swarthout (2021) include data from 175 undergraduate students in

[6]See Appendix B in Baillon, Bleichrodt, and Spinu (2020) for a description of how subsets of the task set differ in terms of the pairs' maximum and minimum outcomes.

100 binary lotteries framed as gains, losses, and mixed outcomes. 16 of the tasks are mixed and 16 are losses, the rest are gains. Each task consists of lottery pairs with either 1, 2 or 3 outcomes, where the outcomes are fixed and the probabilities have been chosen using the indifference curves in the Marschak-Machina triangle, ensuring maximal discrimination between Expected Utility Theory and CPT.

We now discuss the analysis data from environments that involve increased levels of cognitive load. For our first cognitive loaded domain, that of complexity, we analyse data from Moffatt, Sitzia, and Zizzo (2015). Their experiment involves 80 subjects making pairwise choices in 54 binary lottery tasks designed to separate between complexity aversion and risk aversion. The lotteries vary in complexity, where some involve as little as 3 outcomes, whilst the maximum includes 27 outcomes, labelling them as simple, complex and very complex. Adopting a structured procedure, they transformed 3-outcome simple lotteries to complex ones, and from these complex lotteries they could generate very complex lotteries, based on a similar procedure. They generate in total 27 lottery pairs (combinations of simple, complex and very complex lotteries) and they present them to their subjects twice, to test consistency. Here we use the first 27 tasks. All lotteries have the same expected value. For our second cognitive loaded domain, that being decision making under time pressure, we make use of Experiment 2 from Olschewski and Rieskamp (2021). 60 subjects faced a battery of binary lotteries in two treatments, no time pressure (NTP) and time pressure (TP). We analyse data from 150 tasks, where 75 of these tasks fall in the NTP domain, and the other 75 were under time pressure. The set of lotteries was identical for the two treatments, but both the order of tasks and the positioning (left-right) were shown in a randomised order, different for each participant. The experimental design controls for the level of *complexity* of the lotteries by manipulating the number of outcomes of a gamble. There were three conditions in total: complex, where both gambles consisted of four outcomes each, safe-easy: where the safer gamble with lower variance had only two outcomes; and risky-easy, where the riskier gamble with higher variance had only two outcomes. Both outcomes and probabilities varied with a random structure. In the time pressure condition, the average time constraint was set to 4.12 seconds based on reaction times of participants from a practice experiment. Whereas during the NTP condition, participants had 30 seconds per lottery to make their decision

3.4.1 Complexity Index

In this subsection, we devise an index to explore the relationship between the degree of complexity of the tasks in a particular study and the percentage of subjects classified as Toolbox decision makers in that study. To this end, we develop a measure that will help us characterise the degree of *complexity* of a given dataset. In the literature, complexity of a task is mostly characterised by the number of alternatives on the decision maker's choice set, or the number of payoff outcomes in a particular lottery (see among others Sonsino, Benzion, and Mador 2002, Moffatt 2016, Zilker, Hertwig, and Pachur 2020, Fudenberg and Puri 2022). Nevertheless, this measure ignores the role of the set of attributes of each lottery (or pair of lotteries) that may affect the degree of complexity a decision maker perceives, and complexity can be seen as increasing with the number of alternatives and the number of attributes. Recently, researchers have started to include further attributes of a lottery as indicators of complexity. For instance, Diecidue, Levy, and Ven (2015) accounts for focusing on both the number of outcomes and the format of the probabilities in a task, distinguishing between simple (rounded) and complex (non-rounded) probabilities. Huck and Weizsäcker (1999) examine how several features of a decision task lead subjects to deviate from expected value maximisation, including the number of outcomes, the format of probabilities and outcomes, as well as distribution moments of the lotteries such as the mean and the variance. While in a similar context, Enke and Shubatt (2023) construct complexity indices by evaluating the effectiveness of numerous features in forecasting the error rate when identifying the lottery with the highest expected value in their experiment.

In our index we aim to take into consideration all the features of the decision tasks that may increase the decision maker's perception of complexity, along with the amount of cognitive load. On top of the number of outcomes and probabilities, we want to take into consideration the presentation format of probabilities and the degree of similarity between lotteries. Following Huck and Weizsäcker (1999) we include in the index the following features:

- Average number of outcomes across all lotteries ($avg\ \#outc$).
- Average number of outcomes with non-rounded probabilities (non-divisible by 0.05, $avg\ \#probs$).
- Average expected value difference between lottery pairs ($avg(EV_{diff})$).
- Average standard deviation difference between lottery pairs ($avg(st.dev_{diff})$).
- Ratio of unique outcomes over total number of tasks ($\#\ outcomes/\#tasks$).

We define the index I_{cplx} as:

$$I_{cplx} = avg\ \#outc + avg\ \#probs - avg(EV_{diff}) - avg(st.dev_{diff}) + \frac{\#outcomes}{\#tasks} \quad (3.10)$$

Note that the expected value and the standard deviation differences enter the index with a negative sign. The closer the expected value of two lotteries, the harder is to make a decision. Likewise, comparing lotteries with similar variance makes the task of identifying the riskier lottery harder. Therefore, large differences in expected value and variance reduce the overall complexity score. Finally, experimental designs that involve a large number of tasks, with non-fixed multiple outcomes are expected to be more cognitively demanding. The index captures this aspect via the ratio of the total number of unique outcome values in a dataset over the total number of tasks in this dataset. We normalise the index in the interval $[0, 1]$, with 1 indicating the highest level of complexity[7].

It is possible to assign different weights to each of the attributes, allowing some to exert a more substantial influence on the overall measure of complexity. Nevertheless, estimating these weights would necessitate data on the perceived complexity of a lottery (or a set of lotteries) based on various attributes. Gathering such data might

[7] While time pressure increases subjects' cognitive load, there is no seamless way to capture this feature in the index. A potential way would be to use the median response time as an indicator, expecting that low response times correspond to high time pressure. Nevertheless, given that there is no straightforward relationship between the two, we instead extrapolate the value of the index for the time pressure data. More specifically, since the tasks in Olschewski and Rieskamp (2021) are identical for both the TP and the NTP conditions, and since time pressure is expected to put high cognitive demands, we set the index for the TP condition to be equal to the mean between the NTP index and 1.

be challenging due to its highly subjective nature[8]. As for our approach, we prefer to treat the effect of the various attributes to complexity perception as exogenous. By doing so, we employ a metric that assigns equal weight to all the attributes.

3.5 Results

This section presents the results of our analysis across three primary areas. Firstly, we determine whether subjects are better characterised as CPT or adaptive Toolbox decision makers. Next, among those classified as toolbox users, we ascertain the number of strategies present in their adaptive toolbox. Finally, we identify the specific heuristics that dominate the strategic portfolio of our toolbox users.

3.5.1 CPT Vs Heuristics

For the Hey and Orme (1994) data, we find that only 1 out of 80 (1.25%) subjects were characterised by an adaptive toolbox. The nature of the tasks involved may provide an explanation to this extreme result. Due to the way the tasks where chosen (fixed outcomes) and the nature of heuristics, there are several ties between two lotteries, where the heuristic predicts indifference, which dramatically decreases the model's predictive capacity. This may also explain the contradictory results of (Stahl, 2018)[9].

Of the remaining 79 subjects (98.75%) who are characterised by CPT, we find that the GE weighting specification was the best performing for the majority of the CPT subjects (46%). For our the second gain-domain dataset (Baillon, Bleichrodt, and Spinu, 2020), we find the number of participants utilising an adaptive toolbox is substantially larger than in the Hey and Orme (1994) data, with 58 out of 136

[8]Enke and Shubatt (2023) adopt a similar approach and develop an index for choice complexity using data from an experiment in which subjects are asked to identify the lottery with the highest expected value. Their index captures the predicted error rate in identifying the lottery with the highest expected value where the predictions are computed as convex combinations of choice set features. They find that the most important features include the excess dissimilarity of the lotteries, the number of outcomes, the lack of dominance, the presence of compound probabilities and the expected value difference.

[9]In fact, for 40/100; 73/100 and 35/100 of the tasks for the 1st, 2nd and 3rd tool respectively, for the HO dataset, there is a predicted probability of 0.5, as the highest, lowest or highest most likely payoffs of the two lotteries coincide. Furthermore, this dataset is known in the literature to favour Expected Utility. As Harrison and Swarthout (2021) highlight, the battery of lotteries deliberately avoided sets of lottery pairs that had generated"knife-edge" tests of EUT. Their design mantra was to be agnostic about choice patterns, and see which models best characterized the data, rather than selecting lottery pairs designed to be hard for EUT per se.

(43%) being characterised by a toolbox. Although CPT still dominates the decision-processes of these individuals, the complexity, and subsequent increased cognitive load, associated with the tasks involved may explain the increase in the toolbox percentage.

The results from the loss domain are fairly consistent across data sets, with 26% of subjects in the Harrison and Swarthout (2021), and 20% of subjects in the Glöckner and Pachur (2012) study fitting to an adaptive toolbox over CPT. Interestingly, of the remaining subjects in both studies who were better characterised by CPT, 80% in the Glöckner and Pachur (2012) were best fit by a GE weighting function, and 72% in the Harrison and Swarthout (2021) were best fit by the PRL_2 weighting function. There is some imbalance regarding the number of the tasks in the gains and losses domain in the two datasets we are using. In particular, for the Harrison and Swarthout (2021) dataset, there are 16 tasks on the losses, compared to 68 on the gains domain, while in the Glöckner and Pachur (2012) data there are 30 losses lotteries against 70 in the gains domain. Therefore, it is not clear whether the difference in the performance of the two models is because of the presence of losses, or due to other reasons. To disentangle the effect of negative payoffs (losses) on behaviour, we repeat the same exercise estimating the models solely on the gain and the loss domain, and compare them within domains. For the Harrison and Swarthout (2021) data, we find that for 21.7% (30.3%) of the subjects, the cognitive toolbox best explains their behaviour in the gains (losses) domain, while for the Glöckner and Pachur (2012) data, 9.3% (10.1%) are best characterised by a toolbox model in the gain (loss) domain.

For our more cognitively demanding domains, we find a more pronounced leap towards the use of adaptive toolboxes. From the results of Olschewski and Rieskamp (2021), we see the number of individuals using a toolbox of heuristics rises from 27% in the NTP control, to 48% when they are restricted by time constraints. Therefore in the NTP condition, 44 subjects (73.3%) are characterised by a CPT specification, while in the TP condition, CPT can only successfully describe the behaviour of 31 subjects (51.7%). Increased time pressure may have hindered the ability of individuals to subjectively weight utilities and probabilities, thus incentivising the use of heuristics to make their decisions. Finally, from the results of Moffatt, Sitzia, and Zizzo (2015), we see a significant switch in decision-processes, as this experimental

TABLE 3.2: **Subject Classifications.**

Dataset	Domain/Environment	Toolbox	PRL1	PRL2	TK	GE	TOTAL
Hey and Orme (1994)	Gains	2	11	14	16	37	80
%		0.025	0.140	0.180	0.200	0.463	
Baillon, Bleichrodt, and Spinu (2020)	Gains	60	22	39	5	10	136
%		0.441	0.162	0.287	0.037	0.074	
Harrison and Swarthout (2021)	Losses/Mixed	46	9	93	14	13	175
%		0.263	0.051	0.531	0.080	0.074	
Glöckner and Pachur (2012)	Losses/Mixed	13	2	8	0	41	64
%		0.203	0.031	0.125	0.000	0.641	
Olschewski and Rieskamp (2021)	Gains	16	15	19	9	1	60
%		0.267	0.250	0.317	0.150	0.017	
Olschewski and Rieskamp (2021)	Time Pressure	29	9	12	9	1	60
%		0.483	0.150	0.200	0.150	0.017	
Moffatt, Sitzia, and Zizzo (2015)	Complexity	54	1	4	19	2	80
%		0.675	0.013	0.050	0.238	0.025	

Notes: Percentage of subjects classified as Toolbox or CPT per dataset.

design, consisting of extremely complex decision-tasks, meant 68% of individuals are better characterised by an adaptive toolbox than any of the CPT specifications. Of the remaining 32% who are better fit to CPT, we find that the TK specification dominates, fitting the behaviour of 73% of these individuals.

Table 3.2 summarises the results for all the datasets while figure 3.1 illustrates the percentage of subjects classified as CPT or Toolbox decision maker in ascending order. It is apparent from the Figure that the two most cognitively demanding datasets (time pressure and complexity) have the highest frequency of subjects that resort to heuristics. This result indicates a correlation between the increased cognitive load and subjects using simple rules of thumb to make their decision-making process easier. We explore this relationship in section 3.5.4.

FIGURE 3.1: **Subject Classifications.**

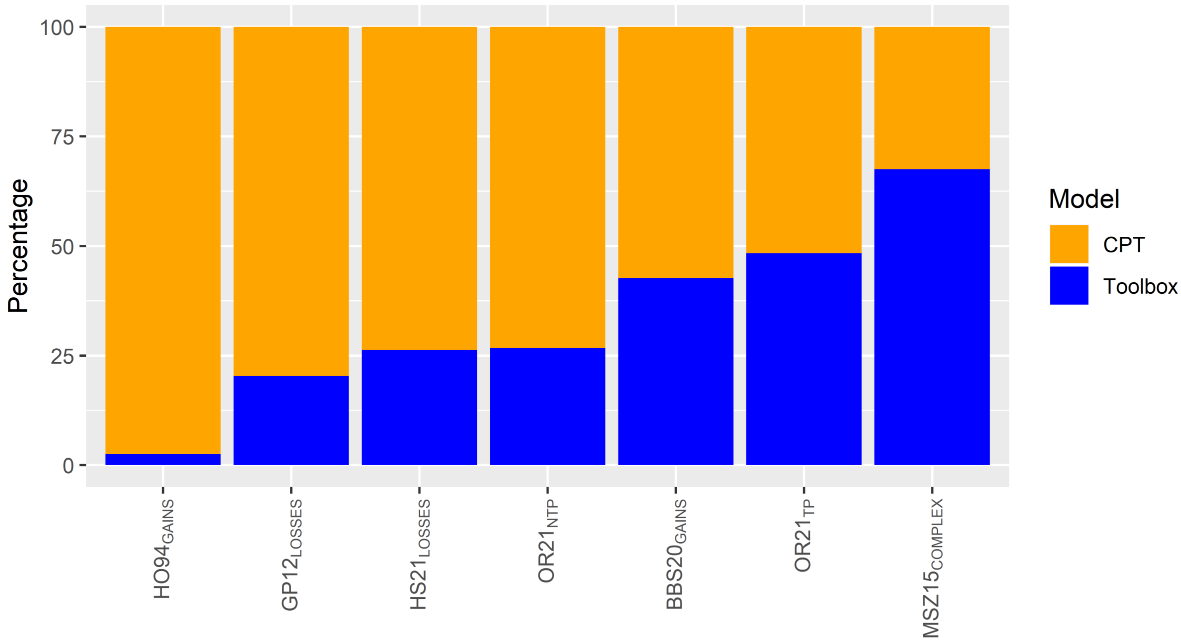

Notes: Percentage of subjects classified as CPT or Toolbox decision makers per dataset.

After classifying our data sets into two subdomains based on whether they are associated with an increase in cognitive load (CL) or not (NCL), we find in the NCL domain that 71% of individuals are better characterised by a CPT specification, and only 29% are better described by an adaptive toolbox. Whilst in our CL subdomain, this result changes to 50% for CPT and 50% for the toolboxes, further emphasising the increased reliance on, or switch to, toolboxes of heuristics when individuals are constrained by an increased cognitive load.

3.5.2 Number of strategies in a Toolbox

With regards to the optimal number of strategies to include in a toolbox, Figure 3.2 shows how many tools were used by each individual over all data sets.

We can see that, of the 217 participants, over all data sets, who were characterised as using a cognitive toolbox, 29% had a toolbox of 4 strategies, 29% used 3 strategies, 22% used 2 strategies, and 20% used 5. We separate the analysis into our CL and NCL domains to disentangle the effect that increased cognitive pressure has on the

FIGURE 3.2: **Number of tools in the toolboxes.**

(A) All

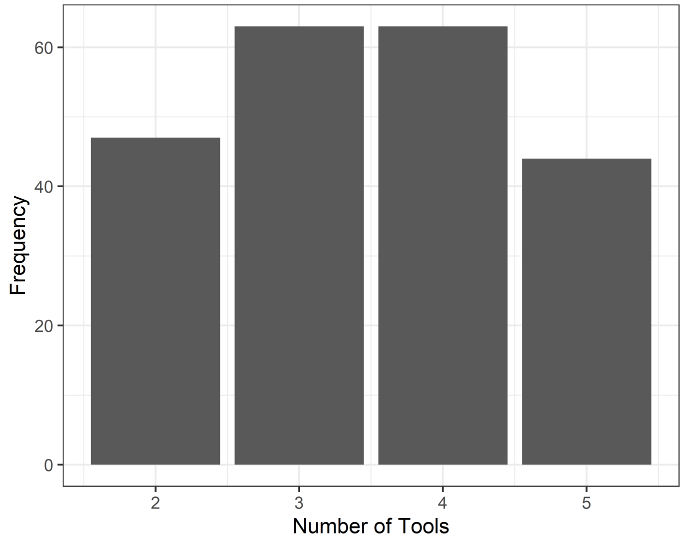

(B) High cognitive load

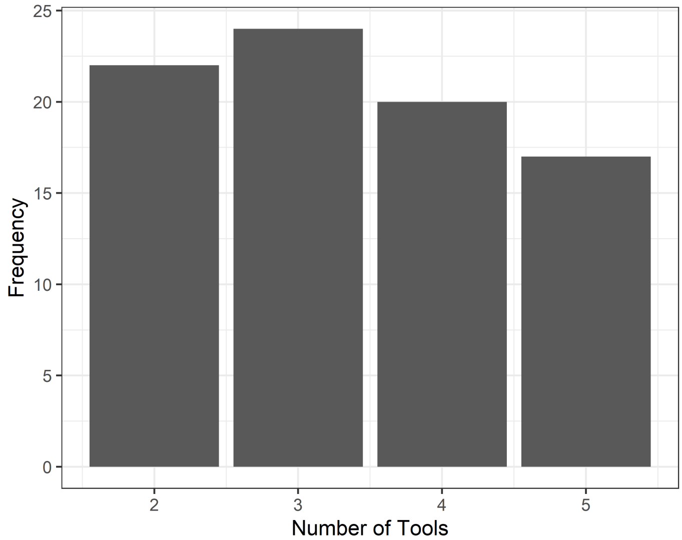

(C) Low cognitive load

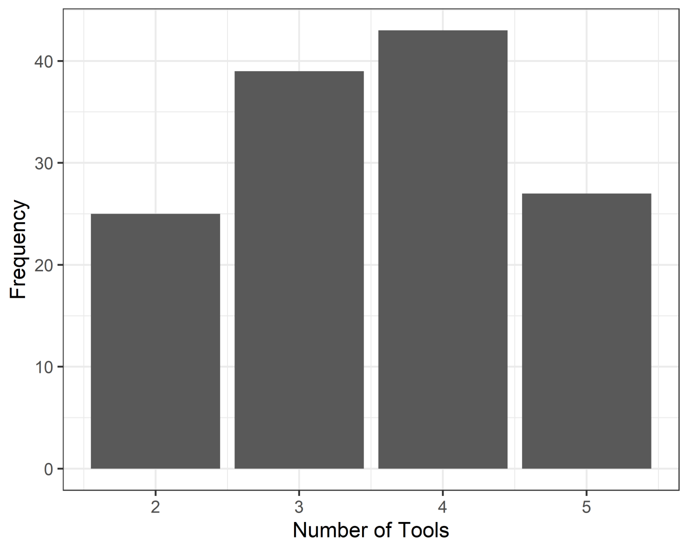

Notes: The top panel illustrates the frequency of heuristics in all toolboxes, the middle for the case of high cognitive load (time pressure and complexity) and the bottom panel for the case of low cognitive load (gains, losses and no time pressure).

number of strategies individuals use over a series of tasks. We find that, as cognitive load increases, the number of strategies used falls. When the subjects cognitive resources are restricted, 29% of subjects use 3 strategies and 27% use 2 strategies. Only 24% and 20% of subjects use 4 and 5 strategies respectively. On the other hand, when the tasks at hand and respective environment are not cognitively demanding, 32% of subjects use 4 strategies, 29% use 3, 20% use 5, and 19% use 2. This suggests that when tasks overwhelm ones cognitive capacity, the number of strategies they can process at once declines, and they are forced to comply with a mere couple of strategies for all decision-tasks.

From a methodological perspective, our results support the existing literature in that individual toolboxes tend to hold, on average, around 3-4 strategies, and that anything over 5 would diminish the predictive capacity of the model (Mohnert, Pachur, and Lieder 2019, Makridakis and Winkler 1983, Ashton and Ashton 1985, He, Analytis, and Bhatia 2022).

3.5.3 Which Heuristics are used

Finally, for all participants that are characterised as toolbox users, we extracted which specific heuristics formed each of their toolboxes to decipher which strategies were relied on in different enviroments. A graphical representation of our findings are represented in Figure 3.3. Of the 755 times that heuristics are used over decision-tasks and subjects, we find that the Minimax was used the most frequently accounting for 13% of these decisions. The Least Likely heuristic, Better-than-Average, and Equiprobable, also accounted for a substantial amount of individual decision processes. On the other hand, the Most likely heuristic, Maximin and Tallying were relied upon the least amount of times, with the latter only showing up in 6% of toolboxes. The observed extensive heterogeneity in heuristics used, however, follows a logical pattern, with the vast majority resorting to heuristics that provide a safety net, highlighting risk aversion, and with the strategies associated with risk seeking behaviour being avoided. These results highlight the fact that adaptive toolboxes of heuristics are still able to capture complex risk preferences.

Figure 3.3 divides these results into two mains: those concerning an increase in cognitive load and those associated with low cognitive load. This is to depict the effect an increase on ones' cognitive load, through time pressure or increased complexity, has on the strategies adopted. There are two specific results in this domain we would like to highlight, regarding The Minimax heuristic and the Priority Heuristic. When pooling the data from the experiments involving less cognitively demanding tasks, we see that the Minimax heuristic prevailed, accounting for 69 (15%) of decisions. In this subcategory we also see that the Priority heuristic is only the 6th/11 most commonly used heuristic, accounting for only 39 (8%) of heuristic strategies in this domain. However, as time pressure is introduced and the complexity of decision-tasks increase, we find that the use of these strategies switch places with regards to their dominance in individual decision-processes. In the high cognitive load domain, we find that the use of the Priority heuristic rises to 12% (34 out of 281 cases), making it the most used strategy in this domain, whilst the use of the Minimax heuristic falls to almost 10%.

Finally, we separate our results into the "types" of heuristics used in the various domains, namely into three domains. The first being monetary payoff heuristics, which consist of the equiprobable, equal-weight, better-than-average, minimax, and Maximin heuristics. The second being dual heuristics, consisting of the least likely, most likely, and probable heuristics. The final domain being heuristics with a lexicographic nature, consisting of the priority heuristic, the lexicographic heuristic (multiple attribute choice), and the tallying heuristics.

Overall 48% of heuristics used were monetary payoff heuristics, 28% are dual heuristics, and 24% are the multiple attribute choice heuristics. We find that this result remains the same across the cognitive load and non-cognitive load domains. Finally, once we account for the fact that we include more monetary payoff heuristics than its counterparts and adjust our results so they are directly comparable, these results change to 36% for monetary payoff heuristics, 34% for dual heuristics, and 30% for heuristics of multiple attribute choice. This result strengthens our hypothesis that in the analysis of individual decision-making using adaptive toolboxes, it is crucial to accommodate toolbox heterogeneity and consider a diverse range of strategies. It is evident that individuals rely on various heuristics, spanning different categories, when making economic decisions.

FIGURE 3.3: **Frequency of Heuristics in the toolboxes.**

(A) All

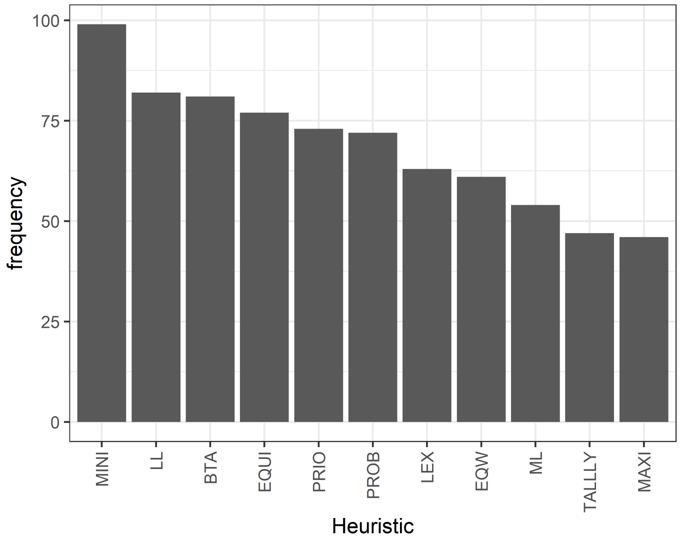

(B) High cognitive load

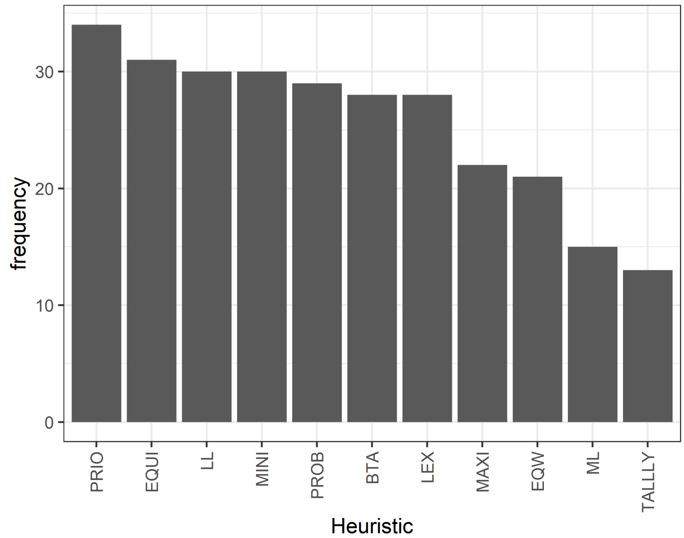

(C) Low cognitive load

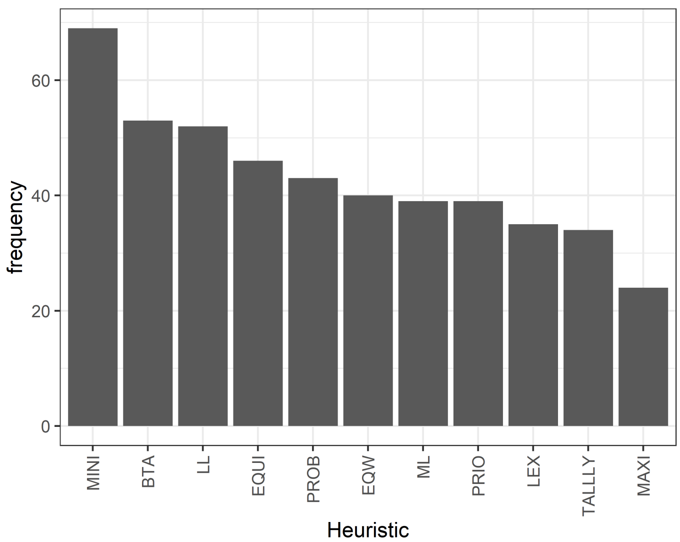

Notes: The top panel illustrates the frequency of heuristics in all toolboxes, the middle for the case of high cognitive load (time pressure and complexity) and the bottom panel for the case of low cognitive load (gains, losses and no time pressure).

3.5.4 Complexity Index

Figure 3.4 displays the relationship between the complexity index for each dataset and the percentage of subjects classified as Toolbox decision makers in that dataset. As expected, datasets with fixed outcomes, rounded probabilities and high differences in expected values such as Hey and Orme (1994) or Harrison and Swarthout (2023) generate the lowest complexity score, while datasets with a large number of outcomes and zero difference in the expected value of the lottery pairs, are the most complex (i.e. Moffatt, Sitzia, and Zizzo 2015).

FIGURE 3.4: **Complexity index per dataset.**

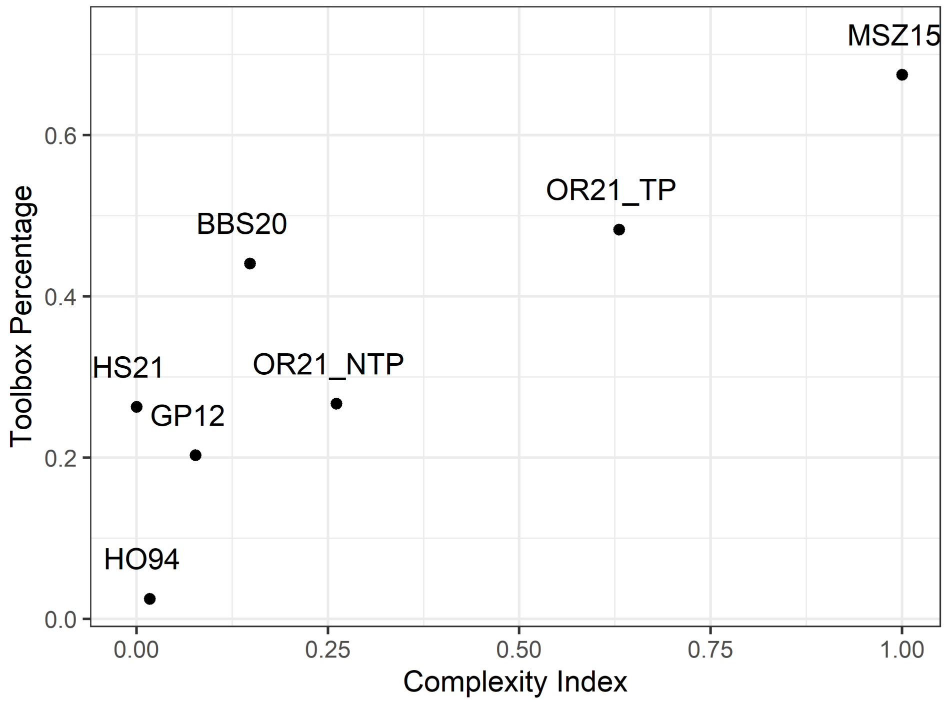

The Figure illustrates the positive relationship between complexity and use of heuristics. Using a Pearson product-moment correlation test, we find that the correlation is positive ($\rho = 0.859$) and significant (p=0.013). The latter has implications on the experimental design and stimuli used in various studies, to identify and estimate preference functionals. While there is a large literature aiming to develop methods that increase the informative content of experimental data (see for example adaptive experimental designs) there is a risk that over-complicated designs encourage participants to resort to heuristic decision making and therefore, lead to failed efforts to identify and estimate the underlying assumed preferences.

3.6 Conclusion

Incorporating environmental dependence into decision-making analysis is crucial for a comprehensive understanding of individual decision-making processes. Given that most day-to-day financial and economic decisions involve varying contexts and environments, it might be unreasonable to suggest that individual strategies and preferences remain constant. Specifically, when examining contextual environments demanding greater cognitive effort—whether due to increased complexity or time pressure in decision making, we propose that the nature of these decisions may overload an individual's cognitive capacity, potentially impeding the optimality and rationality of their choices. Using data from a wide range of studies encompassing diverse decision domains and environments, we investigate the efficacy of heuristics by estimating and rigorously testing toolbox models of cognition to explain behavior. Based on our complexity index, our analysis affirms that as tasks become objectively more complex, leading to an increase in choice-related cognitive load, subjects transition from the use of sophisticated expectation-based utility models to depending on a set of simplification heuristics when making decisions. Our results also indicate that with increasing task complexity, individuals use smaller toolboxes with fewer heuristics to make decisions. This is likely the result of a mental shortcut strategy to mitigate complexity, restricting the number of processable strategies an individual can handle. Lastly, we demonstrate that individuals employ a combination of various types of heuristics with extensive heterogeneity between subjects. Therefore, when assessing the capacity of toolbox models of cognition, we recommend incorporating varying types of heuristics, including monetary payoff heuristics, dual heuristics, and multiple attribute choice heuristics. Our analyses further support the findings of Stahl (2018): when one is agnostic on what is inside the toolbox, then she would be better off by assuming an expectation utility model, rather than a toolbox model with arbitrarily chosen heuristics. While we concentrate on just two domains that may constrain an individual's cognitive capacity, everyday decisions can be affected by numerous factors. These factors include increased time pressure, heightened uncertainty, vulnerable emotional states, interruptions or distractions, or simply a lack of experience. All of these elements can modify an individual's strategic process, ultimately leading to changes in their decisions and outcomes.

From a methodological point of view, we provide the tools on how to efficiently estimate toolbox models via Bayesian Hierarchical Modelling. This approach permits us to combine the beneficial elements of Maximum Likelihood estimation and representative agent assumptions (pooling), whilst eliminating the potential of overfitting and allowing for heterogeneity in individual preferences and strategies. It also accounts for the fact that humans share behavioural similarities, which should be exploited to enhance our understanding of human behaviour. As a by-product of this research, we estimate four full specifications of CPT and run a horse race comparison between different weighting functions in a wide range of decision environments. Our results indicate that the two-parameter family of probability weighting functions was always ranked best, with the PRL_2 having the best performance, followed by the GE function[10]. A final contribution is that we provide a measure of complexity of a dataset and we quantify the relationship between usage of toolboxes and complexity, finding a significant and positive relation between the two. The implications for the design of economic experiments appear to be multifaceted. Depending on the specific experimental design, it is possible that heuristics may be encouraged. It is not entirely clear-cut whether simpler designs inherently discourage heuristic-based problem-solving, nevertheless, it seems that when confronted with complexity, potential losses, or increased cognitive effort, subjects may tend to simplify their decision-making process and resort to heuristics.

Our study also adds to the burgeoning literature on collective model wisdom (model crowds) in decision analysis. As He, Analytis, and Bhatia (2022) highlight, a successful application that harnesses collective model wisdom in decision analysis is Scheibehenne, Rieskamp, and Wagenmakers (2013) who formulate the metaphor

[10]While there has been some limited research trying to identify the best combination of utility, probability weighting and stochastic link functions in the gains domain (see Stott 2006; and Balcombe and Fraser 2015 using the same dataset), there is a lack of a similar comparison in the domain of losses. Here we provide some initial evidence.

of the heuristic toolbox in a hierarchical Bayesian framework and show that by incorporating multiple heuristics, the toolbox explains behavioral data better than a single heuristic. Our study is the first to apply this framework to the field of risky choice.

This raises intriguing avenues for future research, including the exploration of decision environments beyond those considered here, such as decisions from experience, or choice under ambiguity, to gain a more comprehensive understanding of how the interplay between experimental design, cognitive load, and heuristics influencing decision-making.

Chapter 4

Testing Models of Complexity Aversion

Abstract

In this chapter we aim to investigate how the complexity of a decision-task may change an agent's strategic behaviour as a result of increased cognitive fatigue. In this framework, complexity is defined as a function of the number of outcomes in a lottery. Using Bayesian inference techniques, we re-analyse data from a lottery-choice experiment. We quantitatively specify and estimate adaptive toolbox models of cognition, which we rigorously test against popular expectation-based models; modified to account for complexity aversion. We find that for the majority of the subjects, a toolbox model performs best both in-sample, and with regards to its predictive capacity out-of-sample, suggesting that individuals resort to heuristics when the complexity of a task overwhelms their cognitive load.

Keywords: Complexity aversion · Toolbox models · Heuristics · Risky choice · Bayesian modelling

JEL codes: C91 · D81 · D91

4.1 Introduction

In recent years, the economic environment has witnessed a noticeable surge in complexity, driven by a confluence of interconnected factors. Technological advancements and globalization have expanded choices and convenience, while at the same time they have introduced overwhelming options that demand more of the consumers' attention and time. Mortgages, financial products, investment decisions and cryptocurrencies, all come with a plethora of options and features, that can exacerbate consumer decision-making, contributing to their increased cognitive fatigue.

In the field of choice under risk, complexity is represented by the number of payoff outcomes in a particular lottery. Early research on this topic has found that *complexity aversion* is a common attribute in subjects' behaviour, that is they reveal a strong preference for simple lotteries over complex ones (lotteries with higher number of outcomes). Huck and Weizsäcker (1999) and Sonsino, Benzion, and Mador (2002) were among the first to provide evidence that individuals discriminate heavily against complicated lotteries, such that even when the expected value was fixed, they still prefer the lotteries with fewer outcomes even when these lotteries have a higher variance. Moffatt, Sitzia, and Zizzo (2015) estimate the distribution of attitudes towards complexity, finding that 50% are complexity-averse, 33% complexity-neutral, and only 17% complexity-loving. They also find that this rate of responsiveness to complexity reduces with experience to the extent that the average subject becomes almost complexity neutral by the end of the experiment. This convergence to complexity neutrality does not necessarily mean that the subjects no longer have a distaste for complex tasks, as it could be that they merely adopted a different strategy to make their decision, one which meant the complexity of the task was no longer hindering their decision process (i.e. heuristics).

From a theoretical modelling point of view, various expectation-based utility models (e.g. mean-variance, Expected Utility, Cumulative Prospect Theory) have been modified to capture complexity aversion. Moffatt, Sitzia, and Zizzo (2015) test versions of the mean-variance model, and expected utility, while Fudenberg and

Puri (2022), propose a model that combines the standard cumulative prospect theory (CPT) model with a complexity cost. This model captured preferences for lotteries with smaller number of outcomes and show that both probability weighting and complexity costs have an important role to play in predicting these risky alternatives. Diecidue, Levy, and Ven (2015) find that their results are consistent with prospect theory, but can also be explained by a population with heterogeneous aspiration levels. On the other hand, Bernheim and Sprenger (2020) find that PT and CPT fail rigorous tests that they design, and conclude that there is a possibility the observed behaviour reflects a combination of standard CPT and a form of complexity aversion linked to heuristics. While Georgalos and Nabil (2023) show that the descriptive capacity of CPT is decreasing on the level of complexity in a dataset.

Previous research has also suggested that when decisions are more complex, individuals may avoid making a decision altogether, they might procrastinate, but more often than not they decide to stick with a default option or strategy (Iyengar and Lepper 2001; Thaler and Sunstein 2009). This lead a strand of the literature to associate the distaste for complexity with an increase in one's cognitive load and therefore, an increase in reliance on simplified strategies or heuristics (rules of thumb). Venkatraman, Payne, and Huettel (2014) shows that when faced with multiple-outcome gambles involving probabilities of both gains and losses, people often use simple heuristics that maximise the overall probability of winning. Coricelli, Diecidue, and Zaffuto (2018) find that subjects may employ both a simplifying strategy and a compensatory strategy, providing evidence in support of a multiple-strategy approach to decision making.(Oberholzer, Olschewski, and Scheibehenne, 2021) report evidence of complexity aversion, suggesting a tendency to avoid cognitive effort as a potential explanation. Zeisberger (2022) suggest that the more complex the decision problem, the more likely it is the decision-maker will apply heuristics. Further studies have also supported the idea that complexity induces the use of heuristics with a focus on gain and loss probabilities (Erev et al. 2010; Payne 2005).

While this literature hints towards the increased use of heuristics and simplification strategies as a response to the increased cognitive load, to the best of our knowledge, the relationship between heuristic decision making and complexity has not been thoroughly investigated. This is a gap in the literature that we aspire to bridge. In this short chapter, we aim to study the effects of complexity on decision making and whether the increased complexity, and therefore the increased cognitive fatigue, lead agents to resort to heuristic decision making (following simple rules of thumb) rather than using complicated expectation utility models that account for the level of complexity. The heuristics literature assumes that people are equipped with a repertoire of heuristics (strategies) and simplifying processes (rules of thumb) to solve the tasks they face in daily life. This idea has been theoretically modelled with the aid of a *cognitive toolbox*, from which people might adaptively choose their respective strategies. Payne, Bettman, and Johnson (1993) argued that the decision makers are equipped with a set of strategies and select among them when faced with a decision; an approach which was later extended in Gigerenzer (2002) who models decision making as probabilistic draws from a toolbox of heuristic rules. Scheibehenne, Rieskamp, and Wagenmakers (2013) propose a model of strategy selection. More specifically, they suggest a framework on how to quantitatively specify a toolbox model of cognition, and how to rigorously test it using Bayesian inference techniques. Using data from an experiment designed to elicit preferences towards risk and complexity aversion, we implement the methodology suggested in Scheibehenne, Rieskamp, and Wagenmakers (2013) to estimate cognitive toolbox models. We then test these models against popular expectation-based utility models, modified to account for complexity aversion. We compare the models based on both their in-sample and out-of-sample (predictive) capacity. We find that for the majority of the subjects, a toolbox model of simple heuristics has better descriptive and prescriptive capacity than competing compensatory models.

4.2 Theoretical Framework

In this section we present the theoretical models designed to capture preferences towards complexity and risk. The subjective complexity of a choice task is generally characterised in the literature by the number of alternatives on the decision maker's choice set, or the number of payoff outcomes in a particular lottery (see among others Sonsino, Benzion, and Mador 2002, Moffatt, Sitzia, and Zizzo 2015, Zilker, Hertwig, and Pachur 2020, Fudenberg and Puri 2022). In our comparison, we include three expectation-based utility models that have been developed or modified to account for this type of complexity, as well as a cognitive toolbox of heuristics. We include the two models tested in Moffatt, Sitzia, and Zizzo (2015), namely the *mean-variance* and the Viscusi (1989) *Prospective Reference Theory*, the *Simplicity Theory*, a recent Cumulative Prospect Theory specification to account for complexity, as proposed in Fudenberg and Puri (2022), and a toolbox model of simple heuristic rules, as proposed in Scheibehenne, Rieskamp, and Wagenmakers (2013) and implemented in Stahl (2018).

4.2.1 Mean-Variance

This model assumes that the utility function of the decision maker takes into consideration the expected value of the lottery (mean), the variance (exposure to risk), and its complexity (measured by the number of outcomes). The utility function for an individual i is given by:

$$U(p,x) = \mu_{(p,x)} - \alpha_i \sigma^2_{(p,x)} - \gamma_i C_{(p,x)} \tag{4.1}$$

where $\mu_{(p,x)}$ is the expected value of the J-outcome lottery $\mathcal{L} = \{p_1, x_1; \cdots ; p_J, x_J\}$ defined as:

$$\sum_{j=1}^{J} p_j x_j$$

$\sigma^2_{(p,x)}$ is the variance of the lottery defined as:

$$\sum_{j=1}^{J} p_j \left(x_j - \mu_{(p,x)}\right)^2$$

and $C_{(p,x)}$ is the measure of complexity of the lottery, operationalised as C=0 for a sure payoff, C=1 for a simple lottery, C=2 for a complex, and C=3 for a very complex lottery. The parameter α is closely related to the coefficient of absolute risk aversion, while γ represents the degree of complexity aversion when $\gamma > 0$.

4.2.2 Prospective Reference Theory

This model assumes that the decision makers do not take the stated probabilities at face value, but act as Bayesians, and view the prior probability of each outcome of the lottery $\mathcal{L}$ as $1/J$. The model follows the same specification as above but replaces the objective probabilities in the expected value formula with transformed ones of the form:

$$\tilde{p}_j = \frac{\delta \frac{1}{J} + p_j}{\delta + 1}, \; j = 1, \ldots, J; J > 1 \tag{4.2}$$

The parameter δ defines the degree of probability distortion. When $\delta \rightarrow 0$ the transformed probabilities coincide with the objective ones. On the contrary, as $\delta \rightarrow \infty$, $\tilde{p}_j \rightarrow 1/J$.

4.2.3 Simplicity Theory

Simplicity theory, introduced in Fudenberg and Puri (2022), modifies the CPT model to account for complexity aversion by introducing a complexity cost that captures a preference for lotteries with fewer number of outcomes. The CPT-simplicity model is defined as:

$$U(p, x) = \sum_{j=1}^{J} u(x_j) \left[w\left(\sum_{k=1}^{j} p_k \right) - w\left(\sum_{k=1}^{j-1} p_k \right) \right] - C(|support(p)|)$$

where $C(x)$ is a three-parameter sigmoid cost function to account for complexity, specified as:

$$C(x) = \frac{\iota}{1 + e^{-\kappa(x-\rho)}} - \frac{\iota}{1 + e^{-\kappa(1-\rho)}}$$

with x being the number of outcomes of a lottery, ι the height of the function, ρ the midpoint of the rise, and κ the slope, with larger values of κ indicating a steeper slope[1]. The function satisfies the condition $C(1) = 0$, while $w(.)$ is the Tversky and Kahneman (1992) probability weighting function[2]:

$$w(p) = \frac{p^{\gamma}}{(p^{\gamma} + (1-p)^{\gamma})^{1/\gamma}} \tag{4.3}$$

[1]Sigmoid functions have been extensively used in the artificial neural networks literature.

[2]We also tried different specifications of the probability weighting function (both one and two-parameter functionals) with TK being the best performing specification.

Finally, a power (CRRA) utility function is assumed for the monetary payoffs transformation.

4.2.4 Cognitive Toolbox

Following Scheibehenne, Rieskamp, and Wagenmakers (2013), a toolbox model can be represented by a set of different psychological processes or strategies f, and each strategy predicts a particular course of action, depending on the *ecology* of the decision environment. The outcome of this process can be modelled with the aid of a mixture proportion parameter β, which indicates the probability of choosing each strategy in the toolbox, where β a vector. For instance, for a particular toolbox TB consisting of J strategies, each strategy f_j will be selected with probability β_j, with $\sum_{j=1}^{J} \beta_j = 1$. For instance, a potential toolbox with 4 strategies would be defined as:

- Pick the lottery with the highest payoff (MAXIMAX) with probability β_1
- Pick the lottery with the highest minimum payoff(MINIMAX) with probability β_2
- Pick the lottery with the highest most likely payoff (MOST LIKELY) with probability β_3
- Pick the lottery with the highest probability of the highest possible payoff (MOST PROBABLE) with probability $1 - \sum_{i=1}^{3} \beta_i$

This modelling specification allows for the underlying cognitive process of strategy selection to remain unspecified, given that the value of the parameter vector β will be estimated by the data, providing the empirical validation of the latent strategy mix. Given this mixture specification, the compound probability of choosing lottery A can be specified based on the sum of the individual likelihoods of each f_j, weighted by the mixture probability β_j:

$$p(A|TB) = \sum_{j=1}^{J} [\beta_j \times P(A|f_j)] \tag{4.4}$$

where $P(A|f_j)$ is the individual predicted probability of each strategy. Since the most distinguishable feature of a toolbox model is its adaptive nature (each individual adopts their chosen strategies depending on the choice environment), we deviate from the standard practice of fixing a pre-determined set of strategies, same for all the subjects, and allow for heterogeneity between subjects, both in terms of size (how many strategies) and in terms of content (which strategies). The toolbox models we investigate can accommodate a variety of heuristics (out of a total of 10 heuristics extensively utilised in the literature[3]) and sizes (ranging from 2 to 5 strategies per toolbox[4]). We achieve so by estimating, subject-by-subject, every potential toolbox of size up to 5, that is formed as a combination of a subset of the available 10 heuristics. This gives in total 627 toolbox models.

4.3 Data

We re-analyse the data from Moffatt, Sitzia, and Zizzo (2015). This dataset involves 80 subjects participating in a 2-phase experiment, where in each phase subjects faced 27 tasks in which they were asked to choose between two lotteries with the same expected value, but with differing degrees of complexity and risk (phase 2 consisted of the same 27 tasks presented in a different order). The experiment was incentivised using the random lottery incentive mechanism. The experimental design builds on Sonsino, Benzion, and Mador (2002) and Sitzia and Zizzo (2011) single period tasks. The construction of the lotteries is based on the two tasks presented below. The first task involves the choice between a sure win (SW) and a simple 3-outcome lottery (S_3).

$$SW = \begin{cases} 107, & \text{with probability } 1 \end{cases} \quad S_3 = \begin{cases} 80, & \text{with probability } 0.40 \\ 100, & \text{with probability } 0.30 \\ 150, & \text{with probability } 0.30 \end{cases}$$

[3] We use the heuristics studied in Glöckner and Pachur (2012). In the Appendix there is the full list of heuristics along with a description of the choice they prescribe.

[4] Scheibehenne, Rieskamp, and Wagenmakers (2013) discuss how including too many strategies can lead to the strategy sprawl problem.

Using the S_3 lottery and following a particular procedure [5], it is then possible to generate a *complex* lottery, with nine outcomes, and a *very complex* lottery with 27 outcomes. The new lottery will be more complex, but at the same time *safer*, since it will be characterised by lower variance. On top of the SW lottery, they generated six simple, six complex and six very complex lotteries. Using three simple lotteries, they first generated three complex and three very complex lotteries. Then, using the so-called *safe* version of the simple lotteries, which has decreased spread of the extreme outcomes and unchanged the middle outcome, they constructed three further complex and three very complex *safe* lotteries. The pairwise combinations between a subset of these lotteries, along with the SW lottery, gives the total of the 27 tasks (see Moffatt, Sitzia, and Zizzo 2015, Table 2a, pp. 152-153 for the full set of tasks). All lotteries have the same expected value which also contributes to the complexity of the task.

4.4 Econometric Analysis and Results

We estimate all the models using Hierarchical Bayesian econometric techniques, which allow for the simultaneous estimation of individual level parameters and the hyper-parameters of the group level distributions (see Balcombe and Fraser 2015; Ferecatu and Önçüler 2016; Baillon, Bleichrodt, and Spinu 2020; Alam, Georgalos, and Rolls 2022 and Gao, Harrison, and Tchernis 2022 for some recent applications of Bayesian econometrics in risky choice). We compare models both *in-sample*, and *out-of-sample*. In particular, we first compare the models in-sample, based on the value of the Bayes Factor, using the data from phase 1 of the experiment. We then compare the models based on their out-of-sample predictive capacity (predicted log-likelihood) on the phase 2 tasks, using the estimates from phase 1. To capture stochasticity in choice, we model the error structure assuming a logit link function. The probability of choosing lottery A is given by:

$$p(A, B) = \frac{\exp(\phi U_A(p, x))}{\exp(\phi U_A(p, x)) + \exp(\phi U_B(p, x))}$$

[5]To save on space, we briefly describe the process in the online appendix and we refer the interested reader to the original study (Moffatt, Sitzia, and Zizzo 2015, p.151).

where $U(p, x)$ is the utility as defined in section 4.2, and ϕ an index of the sensitivity to differences in utility, to be estimated. The overall likelihood is a Bernoulli distribution that can be expressed as $P(D) = \prod p(A, B)^I \times (1 - p(A, B))^{(1-I)}$, where I is an indicator function, taking the value 1 when the subject chose A, otherwise 0.

For the toolbox model, since the heuristics generate ordinal choice propensities (i.e. deterministic), we assume a *constant-error* choice rule to capture stochastic choice in the data, where the decision maker chooses with constant probability $1 - \varepsilon$, the option that the heuristic prescribes, and with probability ε she makes a mistake[6]. The overall likelihood for a given subject is therefore the product, across all the tasks, of the weighted sum of predicted probabilities across the number of strategies in a given toolbox.

Table 4.1 reports the results of the classification. The first column classifies subjects to models based on the value of the Bayes Factor, while the second column, according to the models' predictive capacity. In-sample, the toolbox model has the best performance for 56.3% of the subjects, followed by the mean variance (26.3%), the simplicity theory (16.3%) and only one subject is characterised by the Prospective Reference model. A similar pattern is also observed in our out-of-sample prediction exercise. The toolbox model is best for 60% of the subjects, followed by the mean variance (16.3%), the Prospective Theory model (13.8%) and the Simplicity Theory model (10%).

[6] This is the part $P(A|f_j)$ in Equation 4.4.

TABLE 4.1: **Subject Classifications.**

Model	In-sample	Out-of-sample
Toolbox	45	48
%	0.563	0.600
Mean-variance	21	13
%	0.263	0.163
Prospective Reference Theory	1	11
%	0.013	0.138
Simplicity Theory	13	8
%	0.163	0.100
TOTAL	80	80

Notes: Number of subjects for which a model is classified as best, based on the in-sample fit (Bayes Factor) and the out-of-sample fit (predicted log-likelihood).

Given the performance of the toolbox model, we next focus on the size and the content of each toolbox. Figures 4.1 and 4.2, illustrate the distribution of the different sized toolboxes, both in and out-of-sample. In both cases, the majority of the subjects (who is classified as toolbox decision makers) uses 4 or 5 heuristics, while very few use only 2. This size is in line with previous results in the literature (see Makridakis and Winkler 1983; Ashton and Ashton 1985; He, Analytis, and Bhatia 2022). There seems to be a slight drop in the size of toolboxes, out-of-sample, which could be the effect of learning and increased familiarity with the task.

FIGURE 4.1: **Frequency of toolbox sizes (in-sample).**

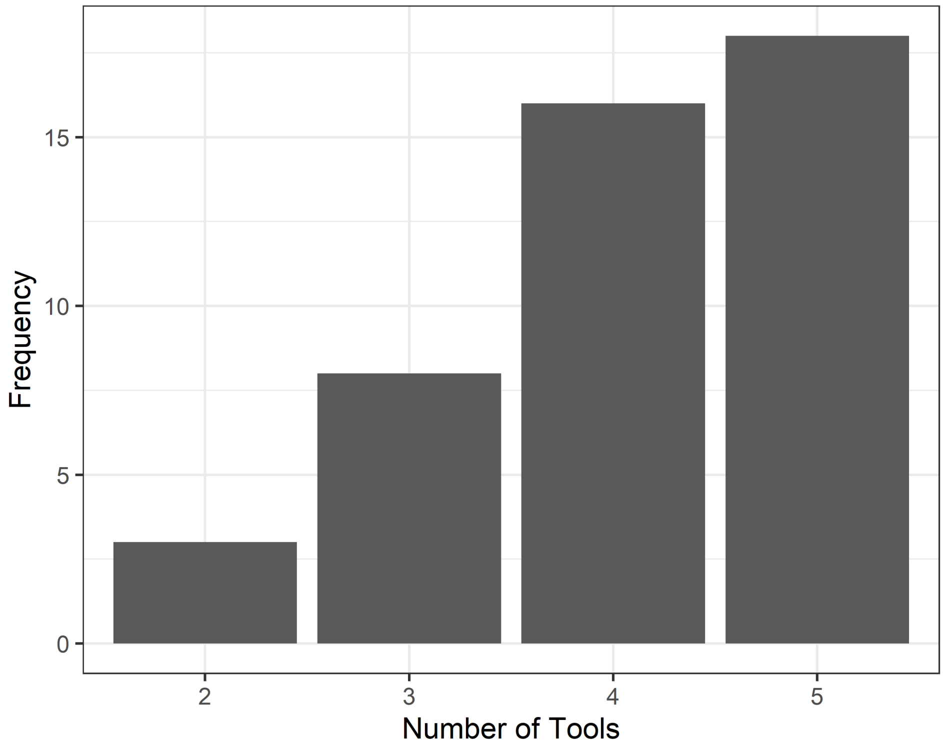

FIGURE 4.2: **Frequency of toolbox sizes (out-of-sample).**

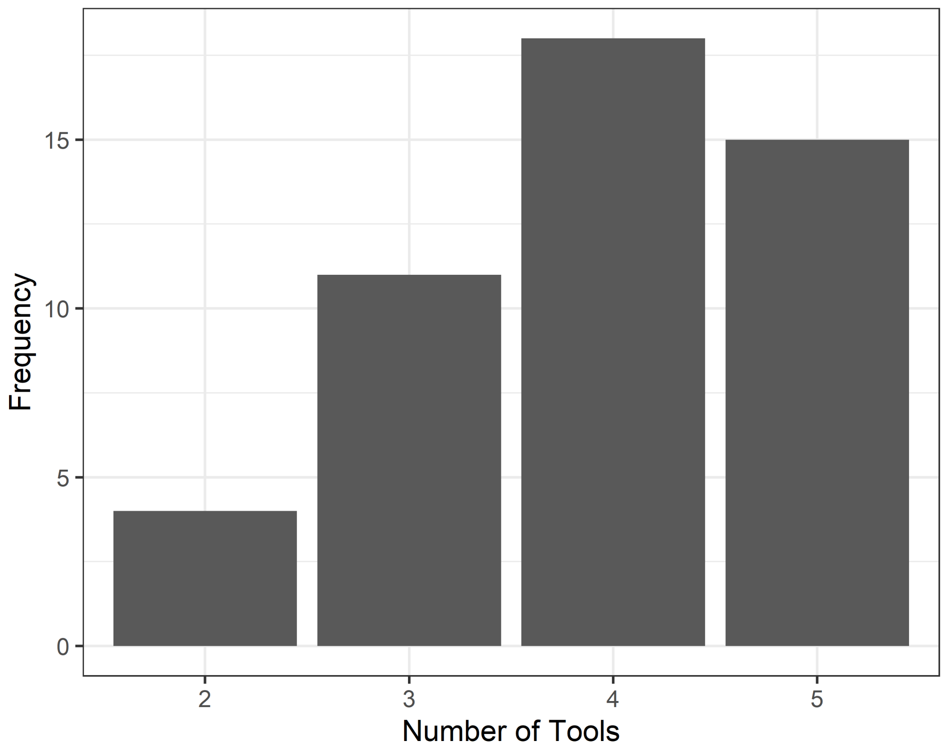

Regarding the content of these toolboxes, Figures 4.3 and 4.4 illustrate the distribution of heuristics across all toolboxes, in and out-of-sample, respectively. Three heuristics outperformed all others, both in and out-of-sample as they were present in the majority of the toolboxes, namely, the Minimax (MINI), the Least Likely (LL) and the Equal Weight (EW). Similarly, the three worst performing heuristics, both in and out-of-sample were the Maximax (MAXI), the Equiprobable (EQUI) and the Most Likely (ML). Given the nature of these heuristics, it is easy to infer that subjects tend to resort to strategies that they will protect them from the worst case scenario (i.e. worst outcome), while avoid strategies that would expose them to higher levels of complexity. When we compare in and out-of-sample differences, there are two points worth mentioning: (1) we find strong evidence in favour of the Priority Heuristic (PRIO), in-sample, a heuristic that has received much attention in the literature because of its capacity to explain risky choice, and (2) the performance of PRIO falls massively in the out-of-sample prediction, which can be seen as an indicator a change in the strategy set that subjects adopt to tackle similar tasks. The PRIO is a lexicographic strategy that requires several rounds of reason comparing payoffs and probabilities and is therefore more cognitively demanding compared to simpler heuristics. This may be a potential explanation of the drop of complexity averse and seeking subjects that Moffatt, Sitzia, and Zizzo (2015) find in the phase 2 data.

FIGURE 4.3: **Frequency of heuristics (in-sample).**

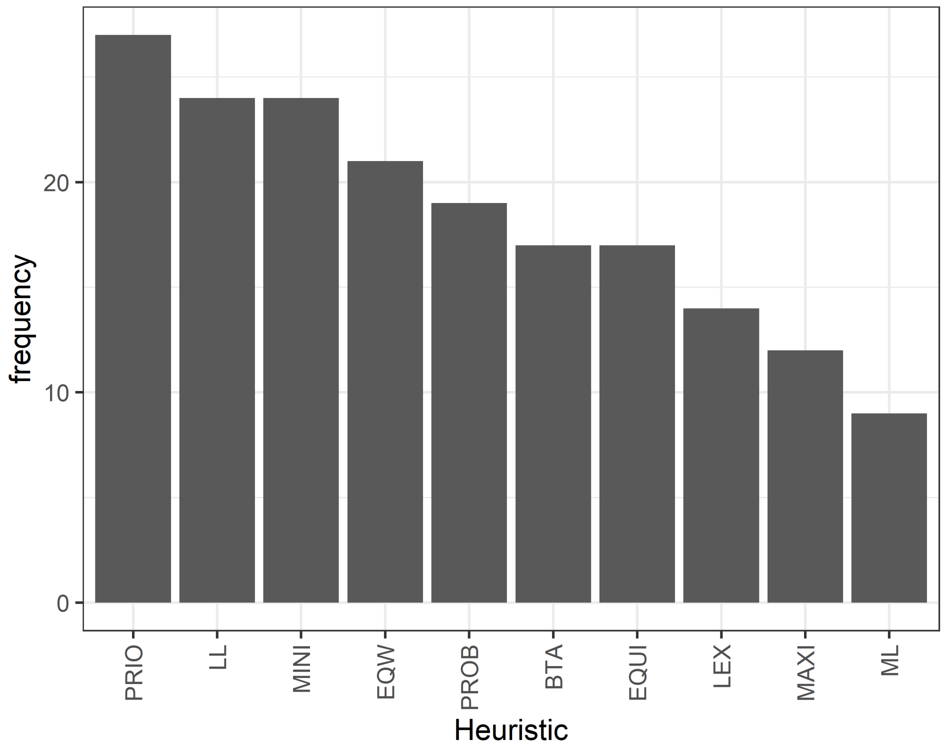

FIGURE 4.4: **Frequency of heuristics (out-of-sample).**

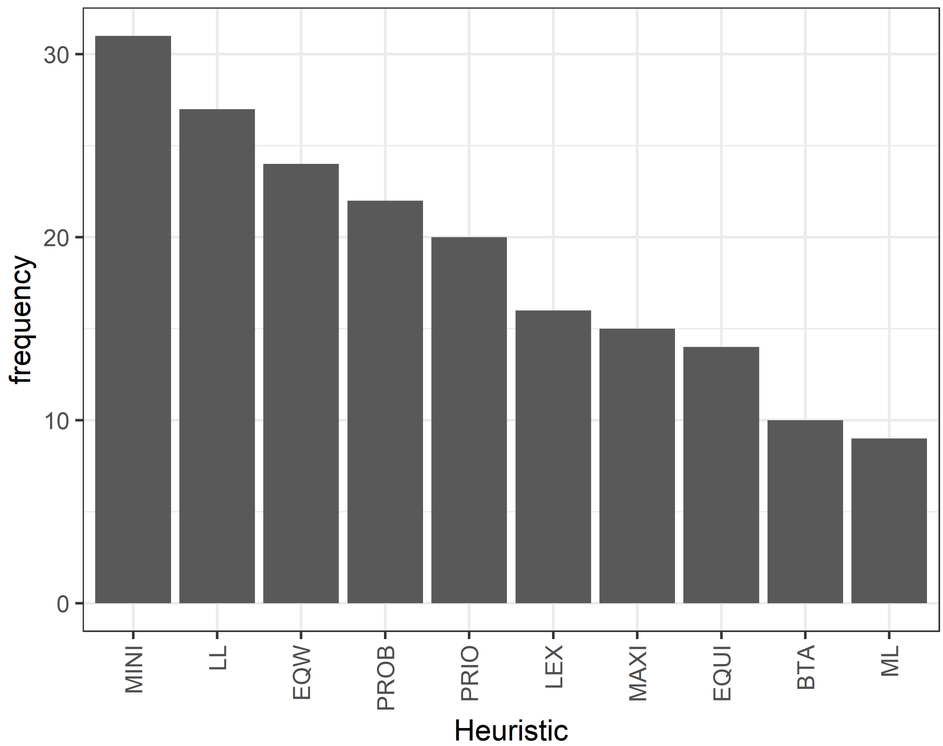

4.5 Conclusion

Our analysis highlights the importance of accounting for complexity when deciding on which explanatory model to adopt to describe individual behaviour. We have shown that with overly complex tasks comes increased cognitive fatigue in decision-making; a characteristic which heightens one's reliance on simple rules of thumb to make decisions. This results in an adaptive toolbox of heuristics outperforming other expectation-based models of decision-making, even when complexity aversion is captured within these competing parametric models. We provide a means of efficiently estimating structural models of decision-making, including a toolbox model, in-sample via the use of Bayesian Hierarchical modelling, and illustrate the robustness of these results in their alignment with our out-of-sample prediction results.

Ironically, analysing strategic processes and preferences in the face of increased complexity is a complex matter in itself, and it is easy to neglect vital attributes of complexity. Whilst most studies use the number of alternatives in a choice set as the key metric, we would urge future research to consider the works of Diecidue, Levy, and Ven (2015), Huck and Weizsäcker (1999) and Georgalos and Nabil, 2023 who discuss how the formatting of probabilities and outcomes, the distribution moments (e.g. variance and mean) and other factors may well fall into the complexity function. The latter design a metric as a benchmark to determine a data sets complexity levels.

Finally we would urge future studies to expand decision-tasks beyond binary lotteries, as research has suggested the impact of complexity on risk taking is largely dependent on the decision format (Oberholzer, Olschewski, and Scheibehenne, 2021). Before we jump to conclusions on complexity's effect on risky decision-making, we must ensure that numerous tasks of varying contexts and characteristics are examined, as it may be that the nature of certain tasks lead people to specific solutions.

Conclusion

We have looked at four chapters focusing on decision-making under risk and uncertainty. We have shown that parametric methods work well in capturing the risk preferences of individuals with inherently complex decision-making processes. Additionally, disentangling probability distortion from probability elevation in the weighting function changes the narrative of subjective perceptions of probabilities. We have shown that imposing restrictive assumptions, such as representative agent or linear utility assumptions, may restrict robust identification of preferences, especially when the number of data points is low. Additionally, we have generated interesting insights into dynamic inconsistencies in risky choices, and how probability distortion plays a role in driving dynamically inconsistent choices. We have identified that semi binding commitment devices work well in improving financial welfare, but only when the subject would have taken on the risk anyway. Finally we have shown that when the complexity of a task increases, whether that be a result of increased time-pressure, more options, more complex options, or something else, then individuals switch from using sophisticated compensatory models, to simplification strategies, such as to ease the decision-making process. Our complexity index illustrates a robust relationship between adaptive toolbox use and tasks of objectively higher complexity. Similarly we show that a toolbox model works better than existing decision-theory models, even when adjusted to include a complexity component.

The research in this thesis has created various pathways for future research. Regarding Chapter 1, it would be interesting to extend the analysis to the mixed domain, such as to capture loss aversion in problem gambling. Dependent on ethical considerations, a new experimental investigation in a dynamic setting could generate important insights, such as to capture how problem gamblers preferences change in light of prior wins or losses. For Chapter 2, it would be useful to re-run the experiment, but with a risk-elicitation task that resembles the generated outcome distributions of the dynamic choice task. Similarly, although it would be computationally challenging, extending the analysis to the domain of ambiguity could be fascinating, such as to replicate more of a trading/stock market environment in which probabilities are unknown. For chapter 2, we propose two more avenues for future research. The first being a similar 2-stage dynamic decision-task, but where one counterbalances the sample to control for ordering effects, such as to capture the role of learning, and whether awareness of dynamic inconsistencies comes from experience. Secondly, we recommend testing various control groups with different commitment devices that differ in their intrinsic/extrinsic nature rather than their level of restrictiveness. For chapter 3 and 4, one could delve deeper into the characteristics of complexity and the effect each characteristic has on complexity aversion and subsequent decision-making processes. Alternatively, given we are one of the first to identify the explanatory power of an adaptive toolbox model of heuristics in a risk setting, largely due to the efficacy of Bayesian inference methods, as well as the heterogeneity of our specification, it is evident that there is a need for more work on the composition of adaptive toolboxes in various settings.

A key takeaway from this thesis is that decisions made under risk and uncertainty are not homogeneous, and dependent on the nature, domain, and environment in which decisions take place, we need to account for varying strategic processes, as well as the effects of extreme risk preferences, and how preferences evolve over time. As a starting point, assessing the fit of various models prior to recovering structural parameters gives one an idea of how subjects process the decision-task. Whilst there are flexible models that have traditionally worked well in many settings, we live in an evolving and inherently complex economy, and tailoring decision-theory analysis to the evolving characteristics of decisions will allow economists to provide higher quality, more relevant, and robust, results and insights.

Appendix A

Chapter 2 Appendix

A.1 Literature Review

This section provides a comprehensive literature review related to dynamic consistency violations in decision-making under risk, disentangling the existing research into three primary domains: 1. Theoretical contributions to the dynamic decision-making under risk literature, 2. Empirical and experimental contributions, 3. The literature regarding overcoming dynamic consistency violations via the implementation of commitment devices to optimal stopping problems.

A.1.1 Dynamic Prospect Theory DPT

Traditionally, research into human choices and decision-making under risk has focused on isolated, static choices that are made independently from one another. Whilst in reality, most economic choices have some element of inter-connectivity. Traditionally, game theory posited that in dynamic choice problems, individuals would devise their solution based on backward induction (Zermelo, 1913), whereby one maximises their utility by determining their optimal choice at the final stage of a finite game, and iterate backwards determining optimal solutions at each stage by eliminating future strategies that will not be played, and forming sub-games (Selten, 1965). This approach, however, assumes dynamic consistency and sequential rationality (Kreps and Wilson 1982; Sarin and Wakker 1998; Machina 1989), a notion that's assumptions have since been faulted theoretically (Reny, 1992), and experimentally (Binmore et al., 2002). Nonetheless, elements of backward induction have been cited in experimental settings since, with Gneezy, Rustichini, and Vostroknutov (2007) finding that individuals initiated with a forward-looking approach (often referred to as the strategy approach), but after experiencing early losses, they switched their mode of analysis to backward induction. Similarly, Carbone and Hey (2001) replayed participants actions to determine how individuals solved dynamic choice problems and found that individuals used some version of backward induction, whereby they looked ahead to the final nodes of a decision-tree, but neglected the principle of optimality. On the other hand, Hey and Lotito (2009) find that the majority of subjects follow the strategy method over backward induction in their experimental analysis. The differentiation of behavioural types in the framework we adopt will further explore whether individuals adopt a strategy approach or backward induction in their dynamic decisions.

Identifying why individuals deviated from their optimal ex-ante strategies became a priority for the behavioural economics literature. The most well-known explanation comes from the discounting literature, in that individuals discount future periods and exhibit a favourable bias for the present (Harris and Laibson, 2001). Strotz (1955-56) was the first to suggest that people are more impatient in the short-run than they are in the long-run, and since then, economists have attempted to model this behaviour with hyperbolic discounting functions (Ainslie and Haslam 1992, Loewenstein and Prelec 1992, Laibson 1997). This function has been able to mirror the well documented psychological bias where individuals prioritise short-term immediate rewards over potentially larger future rewards, with a discount rate that declines as the horizon increases.

Whilst discounting models provide some explanations for dynamic choices, due to their focus being on how individuals value rewards over time, it does not inherently incorporate probabilities into its core components, and is restricting its capacity to explain decision-making inconsistencies under risk (Shoji and Kanehiro 2016; O'Donoghue and Rabin 1999). It therefore struggles to fully explain various behavioural phenomena such as the disposition effect or casino gambling. Heimer et al. (2023) highlight: *"the proposed mechanism for time inconsistency hyperbolic discounting is conceptually distinct from the driver of dynamic inconsistency in risky choice."*

Whilst this area of dynamic literature provides a focus on preferences towards monetary payoffs, a large branch of the static literature provides evidence that many deviations from "rationality" come as a result of probability distortion (Bruhin, Fehr-Duda, and Epper, 2010). Similarly, Nebout and Dubois (2014)'s experimental study into violations of dynamic axioms found that probability levels played an important role in violations of dynamic consistency. We therefore wish to focus on risky decision-making without discounting in order to identify the role probability distortion and loss aversion play on individual dynamic inconsistencies.

Structural static decision-theory models have been extended to dynamic settings to account for inconsistent dynamic decision-making under risk. Hotaling and Busemeyer (2012) extend the basic Decision Field Theory (DFT) framework to the domain of dynamic decision making, whereby a decision-maker deliberates by thinking about the possible outcomes of each action, and updates their preferences over time as new information is presented to them, leading to attention switching between various options until a threshold is reached and one option dominates. This model proposed a new alternative to backward induction and belongs to a class of sequential sampling models. Their analysis involved fitting over 100 model specifications to experimental data, and whilst they were unable to identify a universal classification scheme, they found that there are clear differences in strategies that individuals use to solve dynamic decision-making, where 3 specific behavioural types emerged (near-optimal planners, myopic planners, non-planners).

The identification of behavioural types was first formalised by McClennen (1990) and has been a theme of ongoing importance since, whereby he classifies three behavioural types (naive, resolute, sophisticated). Hey and Lotito (2009) seek to identify these types based on sequential choice problems and found that 50% were naive (unaware of their dynamic inconsistencies), 40% were resolute (find a way of committing to their initial strategies), and 10% were sophisticated (aware of their inconsistencies but struggle to find a way to commit). Nebout and Willinger (2014) extend the design of Hey and Lotito (2009) by additionally categorising the different types based on the dynamic axiom that is being violated (dynamic consistency, consequentialism, reduction of compound lottery). The theme of behavioural type identification has also been extended beyond these traditional classifications, with Hey and

Knoll (2011) recording entire sequences of actions, as in Hotaling and Busemeyer (2012), and identifying distinct groups in terms of broader qualitative characteristics (those who ignored information to minimise effort and those who approximated an optimal strategy). The Alaoui and Fons-Rosen (2021) behavioural classification separates the personality trait of grit into tenacity and diligence. They find that tenacity in itself captures the tendency to deviate from ex-ante preferences when this means accepting defeat. Given the heterogeneity of our population, accounting for heterogeneity in behavioural characteristics may provide further insights into where complex deviations from dynamic axioms originate.

The well renowned Cumulative Prospect Theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992), provides intuition into why certain behavioural types are increasingly dynamically inconsistent. Notably, Barberis (2012) provides a theoretical model of casino gambling, where he extents CPT to a dynamic setting, and shows how a combination of probability weighting and loss aversion leads to plausible time inconsistencies dependent on the classification of a behavioural type, which in turn provides an explanation as to why some individuals engage in casino gambling, even given the negative expected value of such gambles. In doing so, he defines 3 types, who differ in terms of their awareness of their dynamic inconsistencies, and their ability to commit to their initial strategies (naive, sophisticate with commitment, sophisticate without commitment). Barberis' 5-period model forms the framework for our experimental analysis, and will therefore will be expanded on in our theoretical framework. This dynamic version of CPT predicts that the same individual will reject a single fair gamble while accepting the same gamble as part of a dynamic sequence.

As with any emerging theoretical framework, it has received its setbacks, with Ebert and Strack (2015) faulting CPT's predictions in a dynamic framework, and suggesting that a combination of naivety (an individual who is unaware of their dynamic inconsistencies) in tandem with probability distortion, leads to absurd and unreliable predictions. Their theoretical analysis of the model suggests that these individuals would gamble forever. Nonetheless, Henderson, Hobson, and Alex (2017) provide support for the framework, and show that allowing for randomisation can significantly alter the predictions of their model, and leads to the voluntary cessation of gambling in many instances. Additionally, the Ebert and Strack (2015) framework consists of a continuous-time, infinite horizon model, whereas in real-world settings, most gambles and financial decisions are finite (we run out of money, time-constraint on an investment etc.). Ebert (2020) later discusses how an agent will accept gambles even if they are negatively skewed, as long as the horizon is long enough for them to generate positive skewness through their plans. Heimer et al. (2023) highlight that a finite planning horizon restricts the potential skewness of a dynamic loss-exit strategy. Again, this highlights why, as the horizon shortens, individuals may stop gambling. Barberis (2012)'s solution is numerical, in which the model is solved using exhaustive search across Markovian stopping strategies, and due to the probability weighting embedded in CPT's value function, Barberis suggests there is no known analytical solution. However, Hu, Obloj, and Zhou (2022) later turn the problem into a computationally tractable mathematical program by allowing for randomisation, and when solving the model analytically, are able to provide support for Barberis' theoretical predictions for longer time horizons.

The CGM provides an intuitive and applicable account of dynamic decision-making, and whilst it has received positive affirmation, and researchers have since tested various elements of the model, its theoretical predictions have never been directly tested. Considering CPT's contribution to explanations of other financial phenomena such as the equity premium puzzle (Benartzi and Thaler, 1995), the low average return on IPO's (Barberis and Huang, 2008), and the disposition effect (Barberis and Xiong, 2009), testing the CGM's key assumptions in an experimental or empirical setting could provide emerging insights into what characterises our dynamic inconsistencies in financial decision-making.

A.1.2 Empirical and Experimental Evidence

Although the Model of Casino Gambling's theoretical predictions have never been directly tested, in that no-one as of yet has attempted to estimate CPT parameters and identify types of individuals based on these parameters in a dynamic setting of more than 2 periods, features of the model have prevailed to from experimental analysis in the literature.

Hey and Panaccione (2011) estimate a Rank Dependant Expected utility model with a CRRA function and a Quiggin weighting function in a two-period dynamic choice problem. The examine 4 types: Naive, Sophisticate, Resolute and Myopic, where the resolute are those who somehow impose their first period preferences on their future selves (which could be perceived as a sophisticate with commitment) and myopic individuals are those who act as if each period is their last. Their experimental analysis finds evidence that most individuals are resolute. Their design only accounts for 2 periods, thus the extent to which individuals overweight extremely low probabilities in the future is restricted, and sticking to initial strategies over a shorter horizon may be perceived as simpler. Similarly, Johnson and Busemeyer (2001) in two laboratory experiments find that dynamic inconsistencies increase as the tree length increases, suggesting that increasing the horizon in the Hey and Panaccione (2011) study may lead to an increase in the number of individuals being classified as naive. Additionally, their use of RDEU over CPT restricts their analysis to the gain domain only. Gambling and stock market investments are associated with significant losses, and therefore to gain a better understanding of dynamically inconsistent economic choices, incorporating losses and a longer time-horizon is required.

Heimer et al. (2023) delve deeper into dynamic inconsistencies between planned and actual choices in both an experimental and empirical setting. Their empirical investigation, focusing on trading data from a large international online brokerage, and their experimental investigation using fair symmetrical gambles over multiple periods, find that there is a substantial discrepancy between people's risk-taking intentions and actual behaviour in a dynamic setting. Their results implied that a dynamic framework featuring probability weighting, reference dependence, and diminishing sensitivity, like the CGM, is most consistent with observed behavioural patterns. Whilst they do not test the theoretical predictions of the casino gambling model directly, in that they do not seek to identify behavioural types and recover CPT parameters, their results support the models underlying assumptions, and their investigation into different commitment plans provides interesting insights for the model. Similarly, in departure from the Barberis (2012) model, they restrict the subjects' initial plans to a pair of gain and loss limits. This approach allows subjects to define their strategies based on upper or lower bounds in which they wish to stop playing if these bounds are reached. Ebert and Voigt (2023) utilise the visual and operational design of the CGM to elicit stopping times and identify individual preferences for risk-taking strategies. They expand on the research by Heimer et al. (2023) by allowing subjects to use trailing stopping strategies, which refer to strategies that are path dependent. Their results suggest that subjects use trailing stop-loss strategies 3.5 times more than threshold stop-loss strategies. Additionally, they implement randomisation into the framework, allowing subjects to randomise their choices, and find that 80% of subjects use randomisation at least once. They highlight the demand for such randomisation and path dependence by showing how these mechanisms remained popular even when they incurred an additional cost. Their results showed that subjects exhibited dynamic consistency, in that they largely followed their initial pre-specified plans, and suggest that offering flexible strategies may increase individuals' ability to stick to their plan beyond commitment devices that focus on restricted plans.

Similarly, Andrade and Iyer (2009) and Barkan and Busemeyer (1999) analyse individual stopping strategies and how these differ from actual choices, with the former identifying dynamic inconsistencies in that individuals, contradicting their initial strategies, ended up betting more after an initial loss. The latter identify the same pattern in the loss domain, in which individuals become more risk seeking after a loss, but also that subjects are generally more risk averse after an initial gain. Again, this opposed the subjects' initial strategies. Both findings coincide with the predictions of cumulative prospect theory in a dynamic setting. Thaler and Johnson (1990) point out that people take on more risk after a loss, but only if the upside of the gamble allows them to recover from it and get back to the reference point (the "break even" effect).

Alaoui and Fons-Rosen (2021) explore how grit, tenacity, and self-discipline may drive deviations from initial strategies, where they elicit strategies by asking subjects for the minimum and maximum limit with which they would not like to surpass, and find that tenacity, where individuals are not willing to accept defeat, influences individual deviations from ex-ante strategies. Ploner (2017) use dynamic choice tasks to provide support for the disposition effect when choices are taken sequentially, however they find that when choices are planned (pre-commitment) there is a reversed disposition effect. Imas (2016), instead, looks to disentangle the effects of changing attitudes over periods dependant on whether an experienced loss is a realised loss or a paper loss. They show that following a realised loss, individuals avoid risk, however if the same loss is instead a paper loss, individuals are more likely to take on greater risk in the next period.

Barkan (2003) expand on how prior experiences affect dynamic choices in sequential games, and suggest that changes in preferences could be explained by changes in subjective probabilities rather than by the utility associated with a gamble. This insight suggests that after experiencing a gain in one period, there will be a decrease in the subjective probability associated with receiving another gain in the proceeding period, providing an explanation for why individuals plan to bet more after a loss than a gain.

A.1.3 Commitment devices

As of yet, we have discussed the existing evidence on what drives dynamic inconsistencies in risky choice, as well as what characterises these choices, but a co-existing branch of literature has focused on how one can overcome these discrepancies and measure individual preferences for pre-commitment to initial strategies. We use the term "commitment device" to define a means of pre-commiting to one's initial plan. The UK government's Behavioural Insights Team defines a commitment device as an arrangement entered into voluntarily by an individual that sees them put in place measures designed to dissuade them from breaking their intentions (BIT, 2021). Strotz (1955-56) was the first economist to formalize a theory of commitment and to show that commitment mechanisms could be potentially important determinants of economic outcomes.

In decision making under risk, this now focuses on how individuals find a way to commit to optimal stopping strategies, and identifying preferences for various levels of commitment. Bettega et al. (2023) suggest that these commitments that restrict one's action space are able to minimise deviations from initial strategies that come as a result of instant gratification.

Many experimental designs have played around with various types and levels of commitment devices to gain a better understanding of how this helps overcome dynamic inconsistencies, as well as how they increase an individual's willingness to engage in a risky activity. The literature can be divided into two broader domains: those who assess the effect of commitment on the discrepancy between initial strategies and actual choices (Trope and Fishbach 2000; Ariely 2002; Houser et al. 2017; Heimer et al. 2023), and those who focus on preferences for commitment and changes in risk preferences when commitment is available (Antler and Arad 2023; Hey and Paradiso 2006; Bettega et al. 2023).

Antler and Arad (2023) explore how subjects commit to a cutoff stopping rule when facing a sequence of lotteries. Their stopping rules involve the individual selection of an upper and lower bound for which their endowment can reach, with a left biased stopping rule referring to one in which the upper bound (gain) is smaller than the lower bound (loss), and a right-biased rule as the opposite. They find that subjects either consistently choose rules with higher upper bounds or consistently choose rules with lower lower-bounds. They classify subjects into types based on their tendencies to choose left bias, right bias, moderate, and extreme stopping rules, and find that the most common stopping rule was left bias and extreme. As their analysis did not focus on the comparison of ex-ante strategies and ex-post decisions, their results are in line with the existing literature on actual choices made in a sequential process.

On the other hand, Heimer et al. (2023) account for ex-ante strategies in their analysis of commitment devices, and find that individuals choose ex-ante stopping rule strategies that are right-biased, but the subject's actual choices reversed ex-post and were more highly correlated with left-biased strategies. Additionally, they assess the difference between "soft" commitment plans, and "hard" commitment plans, where hard plans are binding and soft plans are breakable. An interesting result highlighted how soft plans can have negative implications, in that the availability of a soft commitment plan encouraged individuals to engage in a risky decision, where they would not have otherwise. They identify welfare implications in that subjects deviated from their soft-commitment plans in real play.

This finding is implicitly supported by Hey and Paradiso (2006) who suggest that individuals are likely to engage in riskier activities, and prefer problems where an element of pre-commitment is available. However, Bettega et al. (2023), who implement a BART experimental design, find the opposite, in that the availability of soft and hard commitment devices led to subjects reducing their risk-taking behaviour, even when the device was non-binding. Psychologically, this result is also intuitive, in that the decision to deviate from a form of commitment may cause the individual to reflect on their willpower, which is associated with discomfort and other costs such as shame or guilt (Benabou and Tirole 2004; Kast, Meier, and Pomeranz 2016).

Whilst the available literature provides a comprehensive examination of the role commitment devices play in risky choice, they have not yet been used to differentiate between behavioural types, more specifically, to differentiate between how individuals plan and execute choices dependent on whether they are aware of their dynamic inconsistencies and if they are able to find a means of committing to their initial plans. Our focus on commitment falls in the demand side, in identifying if individuals have a desire to commit, and how the availability of commitment devices affects initial strategies and demand for risk.

Appendix B

Chapter 3 Appendix

B.1 Literature Review

This section provides an extensive summary and literature review on heuristics, existing evidence on the four key domains of consideration, and the development of adaptive toolboxes.

B.1.1 Heuristics

A heuristic can be thought of as a rule of thumb; a cognitive problem-solving approach that ignores part of the information with the goal of increasing the speed of decision-making (Gigerenzer and Gaissmaier, 2011). They simplify our decision-making process, allowing us to accomplish complex tasks or achieve short-term goals. How we make decisions is largely dependent on the environment and social context of the decision-task, therefore when examining an array of choices, it may be that more than heuristic played a part determining ones decision-strategies. Multiple heuristics have been shown to operate simultaneously, so to better understand the behavior of economic agents, we need to determine which heuristics are at play in various choice domains(Campo et al., 2016). See Brandstätter, Gigerenzer, and Hertwig (2006) for a summary of evidence in favor of heuristics ability to model cognitive processes.

Glöckner and Pachur (2012) compare different implementations of CPT to 11 different heuristics taken from the works of Brandstätter, Gigerenzer, and Hertwig (2006) and Payne, Bettman, and Johnson (1993). Table C.1 provides a comprehensive list of heuristics (note, the use of the word “reasons” or “cues” refer to a specific mental strategy taken within a heuristic). The current study follows Glöckner and Pachur (2012) and investigates the performance of 11 heuristics as potential components of a cognitive toolbox[1].

There are different types of heuristics focusing on one or multiple attributes. Some of the heuristics focus exclusively on the monetary payoffs, such as the Minimax, the Maximin or the Better than Average heuristic (outcome heuristics), while others focus on a combination of payoffs and probabilities (dual heuristics), such as the Least Likely, the Most Likely or the Probable. Finally, there are heuristics for multiple-attribute choice which include the Lexicographic, the Priority and the Tallying heuristic. The latter follow Rubinstein (1988) three-step model, where the agent applies an algorithmic process of decision making, going through various degrees of reason, and if two options are similar in terms of one reason (e.g. dominance) attention is shifted to other reasons (e.g. similarity). This kind of lexicographic type heuristics, have been quite popular in the literature, given that they are simple strategies that correctly predict classic violations of expected utility theory.

The most widely used heuristic in the literature is the Priority Heuristic. It only considers a few pieces of information at a time, it assumes that the decision maker cannot combine information during the process, and it is non-compensatory (Hill, Raacke, and Park, 2017). This heuristic has received considerate recognition as a result of its ability to facilitate a wide array of behaviours: risk aversion for gains (losses) with high (low) probabilities, risk seeking for gains (losses) with low (high) probabilities, and various other empirically proven decision-making preferences. The cues within the priority heuristic are ordered by the individual, such that when a dominant cue is found (going through the order sequentially) the individual stops and makes their decision based on that specific cue. Todd et al. (1999) proposes three building blocks (or rules) within the priority heuristic. The first is a search rule, whereby the individual specifies the direction (or order) in which the cues will take place (they will go through the order until one option dominates another under a specific cue). The second is the stopping rule, which specifies at what cue the search stops, and the third is the decision rule, which specifies how the final decision

[1]The full list of the heuristics along with a description of the choice mechanism behind each heuristic is provided in Table C.1.

is reached. For the Priority rule, in a lottery type setting, the decision maker is given four reasons/cues and goes through them in the following order: Minimum gain, probability of minimum gain, maximum gain, probability of maximum gain. If the decision-maker is indifferent at the first cue, they move to the second, and continue until one option can be discriminated.

There are countless heuristics that have been examined throughout the psychology and economics literature, and many individuals might use more than one heuristic for different problems depending on the nature of the decision. The lexicographic heuristic is similar to the priority heuristic in that they both adopt a lexicographic nature; however, they differ in the order that each attribute is being checked, or in different criteria that are applied in order to stop the search. Here the individual may compare the values (or most likely outcomes) of two options, one value at a time until, again, they are able to discriminate (Gigerenzer and Goldstein 1996; Tversky 1972). The ordering of the cues is dependent on the goal of the choice being made (Hill, Raacke, and Park, 2017). Similarly, the Tallying heuristic adopts the same lexicographic nature, but in this case differences of the magnitude, or predicted value of the cues are ignored, and the decision is based purely on tally marks (Parpart, Jones, and Love 2018; Czerlinski, Gigerenzer, and Goldstein 1999; Dawes 1979).

There are also heuristics that require even less cognitive attention. For example, the Minimax (Maximax) heuristics rule is to merely choose the gamble with the highest minimum outcome (highest outcome). Similarly, the "least likely" ("most likely") heuristic identifies each gambles worst outcome (most likely outcome) and chooses the gamble with the lowest probability of the worst outcome (highest, most likely outcome) (Glöckner and Pachur 2012;Brandstätter, Gigerenzer, and Hertwig 2006).

B.1.2 Gain Domain

In the field of behavioural economics there are two prevalent ways to model irrational behaviour. The one is to assume that agents suffer from cognitive biases (Hilbert 2011; Kahneman and Tversky 1979; Zindel, Zindel, and Quirino 2014).

While they aim to maximise utility, their cognitive biases may lead them to sub-optimal outcomes and therefore, deviations from rationality. For example, a decision maker who distorts objective probabilities, may be prone to violations of the Expected Utility axioms (Machina 1983: Tversky 1975). On the other hand, there is the literature that assumes that agents are not able to calculate probabilities or expectations and they therefore resort to simple rules of thumb (heuristics) to inform their decisions (Gigerenzer and Gaissmaier, 2011). We wish to determine what best characterises deviations from EU in various contexts and domains.

The gain domain provides an exemplary overview of the comparison of the two prevailing behavioural approaches, which is why we use the environment of gains as our baseline domain. Firstly, the gain domain provides the simplest and least cognitively demanding environment in which individuals can make decisions. Additionally, many of the violations of EU have been observed in the gains domain (Allais and Hagen 1979: Avineri and Prashker 2004), and the majority of studies in the vast behavioural literature focus mostly on gains.

CPTs traditional S-shaped utility function implying diminishing marginal utility (Dacey 2003; Fishburn and Kochenberger 1979), combined with its probability distortion component (Preston and Baratta, 1948) has seen a stream of evidence supporting its explanatory power in the gain domain (Stott 2006; Fennema and Wakker 1997; Hey and Orme 1994; Loomes, Moffat, and Sudgen 2002). However, others have argued that CPT's flexibility overpowers its descriptive capacity in some contexts, and that individuals, dependent on the specific environment/context of choice, do not always rely on compensatory models to make decisions, but rather they use general rules of thumb and simplification strategies (Brandstätter, Gigerenzer, and Hertwig, 2006).

B.1.3 Loss Domain

We postulate that individuals behave in a different way when losses are present, therefore we separate the analysis for gains and losses to account for variability in behaviour and decision strategies across the two domains. The change in environment transforms the framing of the decision-task, which in turn re-frames the benefits and consequences of our decisions and has different psychological implications to the gain domain.

The additional mental strain of enduring a potential loss may effect ones decision-strategy (Tom et al. 2007; Thaler 2000; Gal and Rucker 2018), where the increased stress in this domain may leave subjects relying on simpler strategies to reduce the pressure and attempt to avoid future losses (Payne, 1976). Zeisberger (2022) find that decision-makers, in fact, pay explicit attention to loss probabilities in a loss context. They emphasise that it is crucial to separate this hypothesis from the pure effect of loss aversion, as their hypothesis regarding a focus on loss probabilities has strategy implications, whilst loss aversion stems from a preference based cognitive bias model. This result is supported by experimental evidence on individuals providing a focus on loss probabilities in lottery tasks (Payne, Laughhunn, and Crum 1980; Lopes and Oden 1999; Sokolowska 2006), implying heuristics are informing individuals decisions.

On the other hand, there are the arguments that CPT indirectly implements the assumption of a more pronounced focus on loss probabilities into its specification. Whilst CPT captures the concept of loss aversion either in its loss aversion parameter, or by providing different free parameters for risk aversion in the gain and loss domains, it has been noted that loss aversion can also be captured by the weighting function in CPT (Schmidt and Zank, 2005). This idea allows CPT to provide a focus on loss probabilities, as existing evidence has suggested. CPT and its loss aversion assumption has successfully explained deviations from Expected utility, and provided explanations for loss based paradoxes, such as the endowment effect (Thaler, 1980), the equity premium puzzle (Benartzi and Thaler, 1995), and the status quo bias (Samuelson and Zeckhauser, 1988).

B.1.4 Complexity

An influential component in economic decision-making that can increase an individual's cognitive load is the complexity of the task at hand. Within a binary lottery setting, there are various characteristics of a decision-task that can influence its complexity. For example, the number of outcomes in a single binary lottery, the number of binary lotteries to complete, the variance in outcomes within a lottery, the ease of probabilistic calculations (e.g. 0.5 vs 0.634), etc.

There is limited literature on how complexity has may alter risky decision-making strategies in order to simplify the task at hand, with most complexity based studies focusing on individual preferences (e.g. avoid/neglect the overly complex) (Oberholzer, Olschewski, and Scheibehenne, 2021). Whilst this chapter does not focus on preferences for complexity, insights on preferences has implications for strategy development. Evidence suggests that the population is divided between subjects who respond to complexity and those who are complexity neutral. Moffatt, Sitzia, and Zizzo (2015) find 67% of their subjects respond to complexity, and the majority of these display moderate levels of complexity aversion. The distaste for complexity is largely explained by cognitive overload, which may facilitate simplification strategies if complexity cannot be avoided. Moffatt, Sitzia, and Zizzo (2015) find that this rate of responsiveness to complexity (67%) falls to 0% by the end of the experiment (complete complexity neutrality). This convergence to complexity neutrality does not necessarily mean that the subjects no longer have this distaste for complex tasks, but rather they found a simplification strategy that made the tasks less complex.

Sonsino, Benzion, and Mador (2002) find that individuals discriminate heavily against complicated lotteries, such that even when the expected value was fixed, and this information was disclosed to participants, they still prefer the lotteries with fewer outcomes even when these lotteries have a higher variance. In many cases the expected value of the complicated lotteries were higher, yet the simpler lotteries were still preferred. Fudenberg and Puri (2022) expand on this by discussing how a uniform distribution over 20 outcomes is seen as more appealing than a lottery over 19 outcomes that has a higher mean and lower variance if the distribution of the 19-outcomes lottery is very irregular.

It has also been suggested that when decisions are more complex, individuals may avoid making a decision altogether, they might procrastinate, but more often than not they decide to stick with a default option or strategy (Iyengar and Lepper 2001; Thaler and Sunstein 2009). A default option provides an indirect assumption than an individual has thus resulted to a rule of thumb to inform their decisions, providing further evidence to support the explanatory power of heuristics.

Fudenberg and Puri (2022), who capture a preference for lotteries with a smaller number of outcomes, conclude that both probability weighting has an important role to play in predicting these risky alternatives. On the other hand, Bernheim and Sprenger (2020) find that PT and CPT fail rigorous tests that they design. Their design involved lotteries with varying degrees of complexity and they provide suggestions that there is a possibility the observed behaviour reflects a combination of standard PT and a form of complexity aversion.

Zeisberger (2022) suggest that the more complex the decision problem, the more likely it is the decision-maker will apply heuristics. Their results find that individuals provide a primary focus on loss probabilities, which provides support for the heuristics with lexicographic rules and potentially the "least likely" heuristic. They find behaviour to be influences by the probabilities of gaining and losing all their tasks. Further studies have also supported the idea that complexity induces the use of heuristics with a focus on gain and loss probabilities (Erev et al. 2010; Payne 2005).

Although existing work has provided suggestions of possible strategies adopted under complexity, there is limited work on the role these strategies play and, more notably, how many strategies take place over a series of tasks. Venkatraman, Payne, and Huettel (2014) provide support for the assumption that, when faced with multiple-outcome gambles involving probabilities of both gains and losses, people often use simple heuristics that maximises the overall probability of winning. However, from this they expand on the concept of needing multiple strategies over a single strategy in risky choice. Not only do their results oppose the popular prediction of CPT, as they find a greater preference for choices that maximise the overall probability of winning, but they provide arguments that favour an adaptive toolbox of heuristics when tasks become increasingly more complex. They find that the results from single strategy models (i.e. those that predict risk aversion and loss aversion across subjects) fail to generalise to more complex gambles.

B.1.5 Time Pressure

Most economic decisions are made within a finite time horizon (e.g. cashing out a bet, buy/sell an asset, buy/sell a currency). The existing literature examining this domain of decision-making suggests that individuals, due to a lack of time and commitment, will rely on shortcuts when making choices (Campo et al. 2016; Payne, Bettman, and Johnson 1993. This increase in an individual's cognitive load may facilitate that switch from the thorough analysis of a decision based on cognition (the use of cognitive biases) to making a choice based on a simple heuristic. Payne, Bettman, and Johnson (1988) showed that heuristics, in particular the lexicographic rule, were more accurate than normative procedures (e.g. additive utility) when individuals had severe time deadlines. Similarly, Payne, Bettman, and Luce (1996) showed that opportunity cost time pressure meant the subjects in their study adapted their decision-making procedure by processing less information, being more selective in their processing, and thus processing more by attribute; all characteristics that provide support for the use of heuristics in decision-making. Time pressure has also been manipulated specifically to induce heuristic choice behaviour in consumer decision experiments (Langner and Krengel, 2013).

Olschewski and Rieskamp (2021) study decision-making under time pressure and the use of heuristic toolboxes. They find that risk preferences or strategy selection failed to explain the behaviour presented, and that is was down to pure choice inconsistency. However, in their theoretical framework they predetermine the set of strategies, such that all subjects have the same toolbox. We relax this hypothesis in our work, as a less restricted and more heterogeneous framework may provide more conclusive results regarding the identification of strategies under time pressure.

Saqib and Chan (2015) point out that, under time pressure, individuals see the maximal possible outcome (e.g., best gain in gain domain and worst loss in loss domain) as more likely to happen. In turn, they use that as their reference point and not the status quo to evaluate decisions. This tendency reverses traditional CPT risk preferences, as they see only see mediocre gain relative to the maximum gain, thus framing that outcome as a loss, which results in risk seeking behaviour (and vice versa for losses). On the other hand, the experimental evidence provided by Nursimulu and Bossaerts (2014) shows that time pressure leads to purchase impulsiveness, decreased aversion to variance and sensitivity to expected reward. These time varying sensitives can result in increased probability distortion and decreased risk aversion for gains, thus indicating that risk preferences do in fact change when individuals have a limited time to make decisions. The change in decisions that time pressure provokes could also be a consequence of the new, strict, stress induced environment. Cahlíková and Cingl (2013) examine this effect of stress on risk attitudes and found that acute psychological stress significantly increased risk aversion. Seen as time pressure can have a direct effect on stress levels, this is an important result to note for the proceeding analysis.

B.1.6 Toolbox

As of yet we have discussed heuristics and have mentioned regularly the idea of multiple heuristics working simultaneously. We will now expand on the concept of Adaptive Toolboxes of heuristics. In the literature there have been suggested various heuristics that aim to explain deviations from rationality. However, it is not clear whether agents resort to the use of heuristics under particular circumstances and whether they use a unique heuristic or a repertoire of heuristics depending the nature of the problem at hand. The latter is known as the use of an adaptive toolbox of heuristics.

As pointed out by Scheibehenne, Rieskamp, and Wagenmakers (2013), a single model approach does not assume qualitatively different cognitive processes or strategies. An adaptive toolbox of heuristics provides a solution to this particular problem, by pooling together multiple heuristics simultaneously to form a single, unified model, such that if an individual exhibits more than one heuristic throughout a decision-making process, the toolbox will be able to fit all heuristics to the individual's decisions, rather than choosing a single heuristic with substantial noise. He, Analytis, and Bhatia (2022), who introduce the idea of crowd models to capture heuristic behaviour, suggest that each individual model, or heuristic, captures unique aspects of the decision process. We should view the various heuristic strategies as complementary rather than competing accounts of behaviour.

Since the toolbox model is econometrically a mixture model (Stahl, 2018), it provides a percentage accuracy of each heuristic for all decisions. Obviously, with the inclusion of more strategies (a larger toolbox) comes the potential dilemma of overfitting due to an overly flexible toolbox; a problem that is likely to have caused ambiguity in existing results. In simple terms, if a toolbox is unable to capture behaviour accurately, one could keep adding strategies until the noise is virtually non-existent, which in turn does not provide us with reliable insights. Scheibehenne, Rieskamp, and Wagenmakers (2013), however, provide a solution to this problem by using a Bayesian Hierarchical Approach.

While toolbox models have been widely developed throughout the psychology literature, the use of toolbox models is quite restricted in the economics literature; particularly in the field of risky choice. Most of the literature focuses on the comparison between a single heuristic against compensatory expectation based models like CPT, providing overwhelming support in favour of the latter (Glöckner and Pachur 2012; Peterson et al. 2021; Balcombe and Fraser 2015; Brandstätter, Gigerenzer, and Hertwig 2006).

There have been a few attempts to estimate toolbox models in a risk context but most suffer of simplification assumptions that hinder the proper identification of the model (Stahl 2018; Olschewski and Rieskamp 2021). These simplification issues arise in 2 main forms. The first being that existing studies have not allowed for individual heterogeneity in toolboxes, where they assume each subject selects strategies from the same limited toolbox. Secondly, the structure of experimental designs have generated ties for heuristics in the toolbox. Therefore, we wish to assess the explanatory power of toolboxes with more holistic designs and heterogeneous assumptions. There is experimental evidence in the strategy selection literature that highlights an individual's tendency to alter their strategies in correspondence with the structure of the choice problem, and whether the characteristics of the task require speedy decision-making, or increased attention to detail (Mohnert, Pachur, and Lieder 2019; Payne, Bettman, and Johnson 1988). A less restricted approach allowing for heterogeneity, and adopting an econometric framework that can effectively recover strategies (e.g. (Scheibehenne, Rieskamp, and Wagenmakers, 2013)) will increase the likelihood of model identification.

B.2 Table of Heuristics

TABLE B.1: **Table of heuristics.**

	Heuristic	Description
1.	Priority Heuristic	Go through reasons in the order of: minimum gain, probability of minimum gain, and maximum gain. Stop examination if the minimum gains differs by 1/10 (or more) of the maximum gain; otherwise, stop examination if probabilities differ by 1/10 (or more) of the probability scale. Choose the gamble with the more attractive gain (probability). For loss gambles, the heuristic remains the same except that "gains" are replaced by "losses". For mixed gambles, the heuristic remains the same except that "gains" are replaced by "outcomes".
2.	Equiprobable	Calculate the arithmetic mean of all outcomes for each gamble. Choose the gamble with the highest mean.
3.	Equal-weight	Calculate the sum of all outcomes for each gamble. Choose the gamble with the highest sum.
4.	Better than average	Calculate the grand average of all outcomes from all gambles. For each gamble, count the number of outcomes equal to or above the grand average. Then choose the gamble with the highest number of such outcomes.
5.	Tallying	Give a tally mark to the gamble with (a) the higher minimum gain, (b) the higher maximum gain, (c) the lower probability of the minimum gain, and (d) the higher probability of the maximum gain. For losses, replace "gain" with "loss" and "higher" with "lower" (and vice versa). Choose the gamble with the highest number of tally marks.
6.	Probable	Categorize probabilities as probable (i.e., $\geq 1/2$ for a two-outcome gamble, $\geq 1/3$ for a three-outcome gamble, etc.) or improbable. Cancel improbable outcomes. Then calculate the arithmetic mean of the probable outcomes for each gamble. Finally, choose the gamble with the highest mean.
7.	Minimax	Choose the gamble with highest minimum outcome.
8.	Maximin	Choose the gamble with the highest outcome.
9.	Lexicographic	Determine the most likely outcome of each gamble and choose the gamble with the better outcome. If both outcomes are equal, determine the second most likely outcome of each gamble, and choose the gamble with the better (second most likely) outcome. Proceed until a decision is reached.
10.	Least likely	Identify each gamble's worst outcome. Then choose the gamble with the lowest probability of the worst outcome.
11.	Most likely	Identify each gamble's most likely outcome. Then choose the gamble with the highest, most likely outcome.

Appendix C

Chapter 4 Appendix

C.1 Tasks

This Appendix briefly describes the procedure Moffatt, Sitzia, and Zizzo (2015) are using to generate the lotteries for their experiment. Consider the following lottery S_α:

$$S_3 = \begin{cases} 80, & \text{with probability } 0.40 \\ 100, & \text{with probability } 0.30 \\ 150, & \text{with probability } 0.30 \end{cases}$$

This simple lottery with 3 outcomes, can generate a complex lottery with 9 outcomes, and a very complex lottery with 9 outcomes. In vector form, this lottery can be written as $S_\alpha = (p, x) = \left((p_1 p_2 p_3)', (x_1, x_2, x_3)'\right)$. A complex lottery C_α can be generated from S_α using the formula:

$$C_\alpha = \left(vec(pp'); vec\left(\frac{1}{2}xi_3' + \frac{1}{2}i_3 x'\right)\right)$$

where i_3 is a vector of size 3 consisting of ones and $vec(A)$ is the function that transforms a $n \times n$ matrix A into a $n^2 \times 1$ (column) vector consisting of the elements of A. This lottery is equivalent to playing S_α twice and using the arithmetic mean outcome from the two plays as the outcome.

Applying this to the above lottery, we get:

$$vec(p \times p') = \begin{bmatrix} 0.16 & 0.12 & 0.12 \\ 0.12 & 0.09 & 0.09 \\ 0.12 & 0.09 & 0.09 \end{bmatrix} = \begin{bmatrix} 0.16 & 0.12 & 0.12 & 0.12 & 0.09 & 0.09 & 0.12 & 0.09 & 0.09 \end{bmatrix}$$

and $vec(pp'$ generates a vector of size eleven with the element of the $p \times p'$ matrix. Then, for the payoffs:

$$vec\left(\frac{1}{2}xi_3' + \frac{1}{2}i_3x'\right) = \begin{bmatrix} 80 & 90 & 115 & 90 & 100 & 125 & 115 & 125 & 150 \end{bmatrix}$$

which gives the lottery

$$C_3 = \begin{cases} 80, & \text{with probability } 0.16 \\ 90, & \text{with probability } 0.24 \\ 100, & \text{with probability } 0.09 \\ 115, & \text{with probability } 0.24 \\ 125, & \text{with probability } 0.18 \\ 150, & \text{with probability } 0.09 \end{cases} \quad S_3 = \begin{cases} 80, & \text{with probability } 0.40 \\ 100, & \text{with probability } 0.30 \\ 150, & \text{with probability } 0.30 \end{cases}$$

Using a similar procedure, it is possible to create a very complex lottery with 27 outcomes. For the full set of tasks please see Moffatt, Sitzia, and Zizzo (2015, Table 2a, pp. 152-153).

C.2 List of Heuristics

TABLE C.1: **Table of heuristics**

	Heuristic	Description
1.	Priority Heuristic (PRIO)	Go through reasons in the order of: minimum gain, probability of minimum gain, and maximum gain. Stop examination if the minimum gains differs by 1/10 (or more) of the maximum gain; otherwise, stop examination if probabilities differ by 1/10 (or more) of the probability scale. Choose the gamble with the more attractive gain (probability).
2.	Equiprobable (EQUI)	Calculate the arithmetic mean of all outcomes for each gamble. Choose the gamble with the highest mean.
3.	Equal-weight (EQW)	Calculate the sum of all outcomes for each gamble. Choose the gamble with the highest sum.
4.	Better than average (BTA)	Calculate the grand average of all outcomes from all gambles. For each gamble, count the number of outcomes equal to or above the grand average. Then choose the gamble with the highest number of such outcomes.
5.	Probable (PROB)	Categorize probabilities as probable (i.e., $\geq 1/2$ for a two-outcome gamble, $\geq 1/3$ for a three-outcome gamble, etc.) or improbable. Cancel improbable outcomes. Then calculate the arithmetic mean of the probable outcomes for each gamble. Finally, choose the gamble with the highest mean.
6.	Minimax (MINI)	Choose the gamble with highest minimum outcome.
7.	Maximax (MAXI)	Choose the gamble with the highest outcome.
8.	Lexicographic (LEX)	Determine the most likely outcome of each gamble and choose the gamble with the better outcome. If both outcomes are equal, determine the second most likely outcome of each gamble, and choose the gamble with the better (second most likely) outcome. Proceed until a decision is reached.
9.	Least likely (LL)	Identify each gamble's worst outcome. Then choose the gamble with the lowest probability of the worst outcome.
10.	Most likely (ML)	Identify each gamble's most likely outcome. Then choose the gamble with the highest, most likely outcome.

Notes: Heuristics are from Thorngate (1980) and Payne, Bettman, and Johnson (1993), later used in Brandstätter, Gigerenzer, and Hertwig (2006) and Glöckner and Pachur (2012).

Bibliography

Abdellaoui, M., C. Barrios, and P. Wakker (2007). "Reconciling introspective utility with revealed preference: Experimental arguments based on prospect theory". In: *Journal of Econometrics* 138, pp. 356–378.

Abdellaoui, M., F. Vossmann, and M. Weber (2005). "Choice-based Elicitation and Decomposition of Decision Weights for Gains and Losses under Uncertainty". In: *Management Science* 5, pp. 1384–1399.

Abeler, Johannes et al. (2011). "Reference Points and Effort Provision". In: *American Economic Review* 101.2, pp. 470–92. DOI: 10.1257/aer.101.2.470. URL: https://www.aeaweb.org/articles?id=10.1257/aer.101.2.470.

Adinoff, Bryon (2004). "Neurobiological processes in drug reward and addiction". In: *Harv Rev Psychiatry*.

Ainslie, George and Nick Haslam (1992). *Hyperbolic discounting.* New York, NY, US.

Akaike, Hirotugu (1973). "Information Theory and an Extension of the Maximum Likelihood Principle". In: URL: https://api.semanticscholar.org/CorpusID:64903870.

Alam, J., K. Georgalos, and H. Rolls (2022). "Risk Preferences, Gender Effects and Bayesian Econometrics". In: *Journal of Economic Behavior & Organization* 202, pp. 168–183.

Alaoui, Larbi and Christian Fons-Rosen (2021). *Know when to fold 'em: The grit factor*. Economics Working Papers. Department of Economics and Business, Universitat Pompeu Fabra. URL: https://EconPapers.repec.org/RePEc:upf:upfgen:1521.

Allais, Maurice (1953). "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine". In: *Econometrica: Journal of the econometric society*, pp. 503–546.

Allais, Maurice and Ole Hagen (1979). *Expected Utility Hypotheses and the Allais Paradox*. D. Reidel.

Andersen, S. et al. (2010). "Behavioral Econometrics for Psychologists". In: *Journal of Economic Psychology* 31.4, pp. 553–576.

Andrade, Eduardo B. and Ganesh Iyer (2009). *Planned versus actual betting in sequential gambles.* US. DOI: 10.1509/jmkr.46.3.372.

Antler, Yair and Ayala Arad (Sept. 2023). "An Experimental Analysis of the Prize-Probability Tradeoff in Stopping Problems". In: *Review of Economics and Statistics*, pp. 1–45. DOI: 10.1162/rest_a_01366.

Arieli, Amos, Yaniv Ben-Ami, and Ariel Rubinstein (2011). "Tracking Decision Makers under Uncertainty". In: *American Economic Journal: Microeconomics* 3.4, pp. 68–76. DOI: 10.1257/mic.3.4.68. URL: https://www.aeaweb.org/articles?id=10.1257/mic.3.4.68.

Ariely Dan, Wertenbroch Klaus (June 2002). "Procrastination, Deadlines, and Performance: Self-Control by Precommitment". In: *Psychological science* 13, pp. 219–24. DOI: 10.1111/1467-9280.00441.

Ashton, A. and R. Ashton (1985). "Aggregating Subjective Forecasts: Some Empirical Results". In: *Management Science* 31.12, pp. 1499–1508.

Avineri, Erel and Joseph N. Prashker (2004). "Violations of Expected Utility Theory in Route-Choice Stated Preferences: Certainty Effect and Inflation of Small Probabilities". In: *Transportation Research Record* 1894.1, pp. 222–229. DOI: 10.3141/1894-23. eprint: https://doi.org/10.3141/1894-23. URL: https://doi.org/10.3141/1894-23.

Baillon, A., H. Bleichrodt, and V. Spinu (2020). "Searching for the Reference Point". In: *Management Science* 66.1, pp. 93–112.

Balcombe, K. and I. Fraser (2015). "Parametric Preference Functionals under Risk in the Gain Domain: A Bayesian Analysis". In: *Journal of Risk and Uncertainty* 50.2, pp. 161–187. ISSN: 1573-0476.

Baltussen, Guido et al. (Sept. 2012). "Random Incentive Systems in a Dynamic Choice Experiment". In: *Experimental Economics* 15, pp. 418–443. DOI: 10.1007/s10683-011-9306-4.

Barberis, N. (2012). "A Model of Casino Gambling". In: *Management Science* 58.1, pp. 35–51. eprint: http://dx.doi.org/10.1287/mnsc.1110.1435.

Barberis, Nicholas and Ming Huang (Nov. 2008). "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices". In: *American Economic Review* 98, pp. 2066–2100. DOI: 10.2139/ssrn.649421.

Barberis, Nicholas and Wei Xiong (2009). "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation". In: *The Journal of Finance* 64.2, pp. 751–784. DOI: https://doi.org/10.1111/j.1540-6261.2009.01448.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2009.01448.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2009.01448.x.

Barkan, Rachel (Oct. 2003). "Modeling Dynamic Inconsistency with a Dynamic Reference Point". In: *Journal of Behavioral Decision Making* 16, pp. 235 –255. DOI: 10.1002/bdm.444.

Barkan, Rachel and Jerome R. Busemeyer (1999). “Changing plans: Dynamic inconsistency and the effect of experience on the reference point”. In: *Psychonomic Bulletin Review* 6.4, pp. 547–554. ISSN: 1531-5320. DOI: 10.3758/BF03212962. URL: https://doi.org/10.3758/BF03212962.

Bechara, Antoine et al. (1994). “Insensitivity to future consequences following damage to human prefrontal cortex”. In: *Cognition* 50.1, pp. 7–15. ISSN: 0010-0277. DOI: https://doi.org/10.1016/0010-0277(94)90018-3. URL: https://www.sciencedirect.com/science/article/pii/0010027794900183.

Benabou, Roland and Jean Tirole (2004). “Willpower and Personal Rules”. In: *Journal of Political Economy* 112.4, pp. 848–886. URL: https://EconPapers.repec.org/RePEc:ucp:jpolec:v:112:y:2004:i:4:p:848-886.

Benartzi, Shlomo and Richard H. Thaler (1995). “Myopic Loss Aversion and the Equity Premium Puzzle”. In: *The Quarterly Journal of Economics* 110.1, pp. 73–92. ISSN: 00335533, 15314650. URL: http://www.jstor.org/stable/2118511 (visited on 03/22/2023).

Bernheim, B. Douglas and Charles Sprenger (2020). “On the Empirical Validity of Cumulative Prospect Theory: Experimental Evidence of Rank-Independent Probability Weighting”. In: *Econometrica* 88.4, pp. 1363–1409. DOI: https://doi.org/10.3982/ECTA16646. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA16646. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA16646.

Bernoulli, Daniel (1738). “Specimen Theoriae Novae de Mensura Sortis.” In: *Commentarii academiae scientiarum imperialis Petropolitanae, 5, 175-192.*

Berns, Gregory, C. Monica Capra, and Charles Noussair (June 2007). “Receptor Theory and Biological Constraints on Value”. In: *Annals of the New York Academy of Sciences* 1104, pp. 301–9. DOI: 10.1196/annals.1390.013.

Bettega, Paul et al. (2023). *Hard vs. soft commitments: Experimental evidence from a sample of French gamblers*. Working Papers. Grenoble Applied Economics Laboratory (GAEL). URL: https://EconPapers.repec.org/RePEc:gbl:wpaper:2023-05.

Bibby, Peter A. (2016). “Loss-Chasing, Alexithymia, and Impulsivity in a Gambling Task: Alexithymia as a Precursor to Loss-Chasing Behavior When Gambling.” In: *Front Psychol*.

Binmore, Ken et al. (May 2002). “A Backward Induction Experiment”. In: *Journal of Economic Theory* 104, pp. 48–88. DOI: 10.1006/jeth.2001.2910.

Birnbaum, Michael H. and Alfredo Chavez (1997). “Tests of Theories of Decision Making: Violations of Branch Independence and Distribution Independence”. In: *Organizational Behavior and Human Decision Processes* 71.2, pp. 161–194. URL: https://EconPapers.repec.org/RePEc:eee:jobhdp:v:71:y:1997:i:2:p:161-194.

Bishop, C. (2006). *Pattern Recognition and Machine Learning*. Springer-Verlag New York.

BIT (2021). "Applying behaviour insights to design better safer gambling tools". In: *Part 2: Commitment devices*.

Bleichrodt, Han. and Jose Luis. Pinto (Nov. 2000). "A Parameter-Free Elicitation of the Probability Weighting Function in Medical Decision Analysis". In: *Management Science* 46, pp. 1485–1496. DOI: 10.1287/mnsc.46.11.1485.12086.

Blondel, Serge (May 2002). "Testing Theories of Choice Under Risk: Estimation of Individual Functionals". In: *Journal of Risk and Uncertainty* 24, pp. 251–265. DOI: 10.1023/A:1015687502895.

Bookstein, Abraham, Vladimir Kulyukin, and Timo Raita (Oct. 2002). "Generalized Hamming Distance". In: *Information Retrieval* 5. DOI: 10.1023/A:1020499411651.

Bouchouicha, R. and F. Vieider (2017). "Accommodating Stake Effects under Prospect Theory". In: *Journal of Risk and Uncertainty* 55, pp. 1–28.

Bouchouicha, R. et al. (2019). "Gender effects for loss aversion: Yes, no, maybe?" In: *Journal of Risk and Uncertainty* 59, pp. 171–184.

Brand, Matthias et al. (2005). "Decision-making impairments in patients with pathological gambling". In: *Psychiatry Research* 133.1, pp. 91–99. ISSN: 0165-1781. DOI: https://doi.org/10.1016/j.psychres.2004.10.003. URL: https://www.sciencedirect.com/science/article/pii/S0165178104002537.

Brandstätter, E., G. Gigerenzer, and R. Hertwig (2006). "The Priority Heuristic: Making Choices without Trade-Offs". In: *Psychological Review* 113.2, pp. 409–432.

Breen, Robert B and Marvin Zuckerman (1999). "'Chasing' in gambling behavior: personality and cognitive determinants". In: *Personality and Individual Differences* 27.6, pp. 1097–1111. ISSN: 0191-8869. DOI: https://doi.org/10.1016/S0191-8869(99)00052-5. URL: https://www.sciencedirect.com/science/article/pii/S0191886999000525.

Bruhin, A., H. Fehr-Duda, and T. Epper (2010). "Risk and Rationality: Uncovering Heterogeneity in Probability Distortion". In: *Econometrica* 78.4, pp. 1375–1412.

Busemeyer, J. et al. (2000). "Dynamic and Consequential Consistency of Choices Between Paths of Decision Trees". In: *Journal of Experimental Psychology: General* 129, pp. 530–545.

Cahlíková, Jana and Lubomír Cingl (2013). "Risk preferences under acute stress". In: *Experimental Economics* 20.1, pp. 209–236. DOI: 10.1007/s10683-016-9482-3. URL: https://ideas.repec.org/a/kap/expeco/v20y2017i1d10.1007_s10683-016-9482-3.html.

Camerer, C. and H. Ho (1994). "Violations of the Betweenness Axiom and Nonlinearity in Probability". In: *Journal of Risk and Uncertainty* 8.2, pp. 167–196. ISSN: 08955646, 15730476. URL: http://www.jstor.org/stable/41760723.

Campo, Cristina del et al. (May 2016). "Decision making styles and the use of heuristics in decision making". In: *Journal of Business Economics* 86. DOI: 10.1007/s11573-016-0811-y.

Carbone, Enrica (1997). "Investigation of stochastic preference theory using experimental data". In: *Economics Letters* 57.3, pp. 305–311. ISSN: 0165-1765. DOI: https://doi.org/10.1016/S0165-1765(97)00244-9. URL: https://www.sciencedirect.com/science/article/pii/S0165176597002449.

Carbone, Enrica and John Hey (2000). "Which Error Story Is Best?" In: *Journal of Risk and Uncertainty* 20.2, pp. 161–76. URL: https://EconPapers.repec.org/RePEc:kap:jrisku:v:20:y:2000:i:2:p:161-76.

Carbone, Enrica and John D. Hey (2001). "A Test of the Principle of Optimality". In: *Theory and Decision* 50.3, pp. 263–281. ISSN: 1573-7187. DOI: 10.1023/A:1010342908638. URL: https://doi.org/10.1023/A:1010342908638.

Carpenter, Jeffrey and Erick Gong (Oct. 2015). "Motivating Agents: How Much Does the Mission Matter?" In: *Journal of Labor Economics* 34, pp. 000–000. DOI: 10.1086/682345.

Carpenter, Jeffrey Paul and Emiliano Huet-Vaughn (2019). "Real-effort tasks". In: *Handbook of Research Methods and Applications in Experimental Economics*. URL: https://api.semanticscholar.org/CorpusID:199337053.

Cavedini, Paolo et al. (Mar. 2002). "Frontal lobe dysfunction in gambling disorder patients". In: *Biological psychiatry* 51, pp. 334–41. DOI: 10.1016/S0006-3223(01)01227-6.

Charness, Gary, Uri Gneezy, and Austin Henderson (2018). "Experimental methods: Measuring effort in economics experiments". In: *Journal of Economic Behavior amp; Organization* 149.C, pp. 74–87. URL: https://EconPapers.repec.org/RePEc:eee:jeborg:v:149:y:2018:i:c:p:74-87.

Chen, Shu-Heng and Yi-Lin Hsieh (2011). "Reinforcement Learning in Experimental Asset Markets". In: *Eastern Economic Journal* 37.1, pp. 109–133. ISSN: 00945056, 19394632. URL: http://www.jstor.org.ezproxy.lancs.ac.uk/stable/41239496.

Clark, Crystal A. and Alain Dagher (2014). "The role of dopamine in risk taking: a specific look at Parkinsons disease and gambling". In: *Frontiers in Behavioral Neuroscience* 8. ISSN: 1662-5153. DOI: 10.3389/fnbeh.2014.00196. URL: https://www.frontiersin.org/articles/10.3389/fnbeh.2014.00196.

Clark, Luke (Jan. 2010). "Clark L. Decision-making during gambling: an integration of cognitive and psychobiological approaches. Philos Trans R Soc Lond B Biol Sci 365: 319-330". In: *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 365, pp. 319–30. DOI: 10.1098/rstb.2009.0147.

Conte, Anna, John D. Hey, and Peter G. Moffatt (2011). "Mixture models of choice under risk". In: *Journal of Econometrics* 162.1, pp. 79 –88. ISSN: 0304-4076.

Coricelli, Giorgio, Enrico Diecidue, and Francesco D. Zaffuto (2018). "Evidence for multiple strategies in choice under risk". In: *Journal of Risk and Uncertainty* 56.3, pp. 193–210. DOI: 10.1007/s11166-018-9278-2.

Cubitt R., P., C. Starmer, and R. Sugden (1998). "Dynamic Choice and the Common Ratio Effect: An Experimental Investigation". In: *Economic Journal* 108.450, pp. 1362–80.

Czerlinski, Jean, Gerd Gigerenzer, and Daniel Goldstein (Jan. 1999). "How good are simple heuristics?" In: pp. 97–118. ISBN: 978-0195143812.

Dacey, R. (2003). "The s-Shaped Utility Function". In: *Synthese* 135.2, pp. 243–272. DOI: 10.1023/a:1023465024536.

Dai, Hengchen, Katherine Milkman, and Jason Riis (2014). "The Fresh Start Effect: Temporal Landmarks Motivate Aspirational Behavior". In: *Management Science* 60.10, pp. 2563–2582. URL: https://EconPapers.repec.org/RePEc:inm:ormnsc:v:60:y:2014:i:10:p:2563-2582.

Davies, Rob (2020). "UK gambling addiction much worse than though". In: *The Guardian*.

Dawes, Robyn M. (1979). "The Robust Beauty of Improper Linear Models in Decision Making". In: *American Psychologist* 34.7, pp. 571–582. DOI: 10.1037/0003-066x.34.7.571.

Delfabbro, Paul, Julie Lahn, and Peter Grabosky (Aug. 2006). "Its Not What You Know, but How You Use It: Statistical Knowledge and Adolescent Problem Gambling". In: *Journal of gambling studies / co-sponsored by the National Council on Problem Gambling and Institute for the Study of Gambling and Commercial Gaming* 22, pp. 179–93. DOI: 10.1007/s10899-006-9009-5.

Diecidue, E, M Levy, and J van de Ven (2015). "No aspiration to win? An experimental test of the aspiration level model". In: *Journal of Risk and Uncertainty* 51, pp. 245–266.

Ebert, S. and P. Strack (2015). "Until the Bitter End: On Prospect Theory in a Dynamic Context". In: *American Economic Review* 105.4, pp. 1618–33.

Ebert, Sebastian (2020). "On taking a skewed risk more than once," in: *Available at SSRN: https://ssrn.com/abstract=3731565*.

Ebert, Sebastian and Maximilian Voigt (Jan. 2023). "Eliciting Stopping Times". In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4526931.

Enke, Benjamin and Cassidy Shubatt (Jan. 2023). "Quantifying Lottery Choice Complexity". In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4568763.

Erev, Ido et al. (2010). "A choice prediction competition: Choices from experience and from description". In: *Journal of Behavioral Decision Making* 23.1, pp. 15–47.

DOI: https://doi.org/10.1002/bdm.683. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/bdm.683. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/bdm.683.

Fehr-Duda, H. et al. (2010). "Rationality on the Rise: Why Relative Risk Aversion Increases with Stake Size". In: *Journal of Risk & Uncertainty* 40, pp. 147–180.

Fehr-Duda, Helga and Thomas Epper (2012). "Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences". In: *Annual Review of Economics* 4.1, pp. 567–593. URL: https://EconPapers.repec.org/RePEc:anr:reveco:v:4:y:2012:p:567-593.

Fennema, Hein and Peter Wakker (1997). "Original and cumulative prospect theory: a discussion of empirical differences". In: *Journal of Behavioral Decision Making* 10.1, pp. 53–64. DOI: https://doi.org/10.1002/(SICI)1099-0771(199703)10:1<53::AID-BDM245>3.0.CO;2-1.

Ferecatu, A. and A. Önçüler (2016). "Heterogeneous Risk and Time Preferences". In: *Journal of Risk and Uncertainty* 53.1, pp. 1–28.

Fiedler, Susann and Andreas Glöckner (Oct. 2012). "The Dynamics of Decision Making in Risky Choice: An Eye-Tracking Analysis". In: *Frontiers in psychology* 3, p. 335. DOI: 10.3389/fpsyg.2012.00335.

Fischbacher, Urs, Gerson Hoffmann, and Simeon Schudy (Feb. 2017). "The Causal Effect of Stop-Loss and Take-Gain Orders on the Disposition Effect". In: *The Review of Financial Studies* 30. DOI: 10.1093/rfs/hhx016.

Fishburn, Peter C. and Gary A. Kochenberger (1979). "TWO-PIECE VON NEUMANN-MORGENSTERN UTILITY FUNCTIONS*". In: *Decision Sciences* 10.4, pp. 503–518. DOI: https://doi.org/10.1111/j.1540-5915.1979.tb00043.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-5915.1979.tb00043.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-5915.1979.tb00043.x.

Fudenberg, D. and I. Puri (2022). "Simplicity and Probability Weighting in Choice under Risk". In: *American Economic Review Papers & Proceedings* 112, pp. 421–425.

Gal, David and Derek Rucker (Mar. 2018). "The Loss of Loss Aversion: Will It Loom Larger than Its Gain?" In: *Journal of Consumer Psychology* 28. DOI: 10.1002/jcpy.1047.

Gao, X., G. Harrison, and R. Tchernis (2022). "Behavioral Welfare Economics and Risk Preferences: a Bayesian Approach". In: *Experimental Economics*.

Gelman, A. and D. Rubin (1992). "Inference from Iterative Simulation Using Multiple Sequences". In: *Statistical Science* 7.4, pp. 457–472. ISSN: 08834237.

Gelskov, Sofie V. et al. (2016). "Aberrant neural signatures of decision-making: Pathological gamblers display cortico-striatal hypersensitivity to extreme gambles". In: *NeuroImage* 128, pp. 342–352. ISSN: 1053-8119. DOI: https://doi.org/10.

1016/j.neuroimage.2016.01.002. URL: https://www.sciencedirect.com/science/article/pii/S1053811916000070.

Georgalos, Konstantinos and Nathan Nabil (Nov. 2023). *Heuristics Unveiled: A Comparative Analysis of Toolbox Models and Prospect Theory in Risky Choice*. English. WorkingPaper. Lancaster University, Department of Economics.

Georgalos, Konstantinos, Ivan Paya, and David A. Peel (2021). "On the contribution of the Markowitz model of utility to explain risky choice in experimental research". In: *Journal of Economic Behavior Organization* 182, pp. 527–543. ISSN: 0167-2681. DOI: https://doi.org/10.1016/j.jebo.2018.11.010. URL: https://www.sciencedirect.com/science/article/pii/S0167268118303068.

Gigerenzer, G. (2002). "The Adaptive Toolbox". In: *Bounded Rationality: The Adpative Toolbox*. Ed. by G. Gigerenzer and R. Selten. Cambridge, MA: MIT Press, pp. 37–50.

Gigerenzer, Gerd and Wolfgang Gaissmaier (Jan. 2011). "Heuristic Decision Making". In: *Annual review of psychology* 62, pp. 451–82. DOI: 10.1146/annurev-psych-120709-145346.

Gigerenzer, Gerd and Daniel Goldstein (Oct. 1996). "Goldstein, D.G.: Reasoning the fast and frugal way: models of bounded rationality. Psychological Review 103(4), 650". In: *Psychological Review* 103, pp. 650–669. DOI: 10.1037/0033-295X.103.4.650.

Glöckner, A. and T. Pachur (2012). "Cognitive Models of Risky Choice: Parameter Stability and Predictive Accuracy of Prospect Theory". In: *Cognition* 123, pp. 21–32.

Glöckner, Andreas and Ann-Katrin Herbold (2011). "An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes". In: *Journal of Behavioral Decision Making* 24.1, pp. 71–98. DOI: https://doi.org/10.1002/bdm.684. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/bdm.684. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/bdm.684.

Gneezy, Uri, Aldo Rustichini, and Alexander Vostroknutov (Oct. 2007). "I'll cross that Bridge when I Come to It: Backward Induction as a Cognitive Process". In.

Goldstein, W. and H. Einhorn (1987). "Expression Theory and the Preference Reversal Phenomena". In: *Psychological Review* 94, pp. 236–254.

Gonzalez, R. and G. Wu (1999). "On the Shape of the Probability Weighting Function". In: *Cognitive Psychology* 38.1, pp. 129–166.

Goudriaan, Anna E et al. (2004). "Pathological gambling: a comprehensive review of biobehavioral findings". In: *Neuroscience Biobehavioral Reviews* 28.2, pp. 123–141. ISSN: 0149-7634. DOI: https://doi.org/10.1016/j.neubiorev.2004.

03.001. URL: https://www.sciencedirect.com/science/article/pii/S0149763404000247.

Greene, William (Jan. 2003). "Econometric Analysis Fifth Edition". In: *Chapter 14*.

Gul, Faruk and Wolfgang Pesendorfer (2001). "Temptation and Self-Control". In: *Econometrica* 69.6, pp. 1403–1435. DOI: https://doi.org/10.1111/1468-0262.00252. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00252. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00252.

Harless, D. and C. Camerer (1994). "The Predictive Utility of Generalised Expected Utility Theories". In: *Econometrica* 62.6, pp. 1251–1289.

Harris, Christopher and David Laibson (2001). "Dynamic Choices of Hyperbolic Consumers". In: *Econometrica* 69.4, pp. 935–57. URL: https://EconPapers.repec.org/RePEc:ecm:emetrp:v:69:y:2001:i:4:p:935-57.

Harrison, G. and E. Rustrom (2008). "Risk Aversion in the Laboratory". In: *Research in Experimental Economics*. Ed. by J. Cox and G. Harrison. Vol. 12. Emerald Group Publishing Limited, pp. 41–196.

Harrison, G. and E. Rutström (2008). "Expected Utility Theory and Prospect Theory: one Wedding and a Decent Funeral". In: *Experimental Economics* 12.2, p. 133. ISSN: 1573-6938.

Harrison, G. and T. Swarthout (2021). "Cumulative Prospect Theory in the Laboratory: A Reconsideration". In: *Models of Risk Preferences: Descriptive and Normative Challenges*. Ed. by G. Harrison and D. Ross. UK: Emerald, Research in Experimental Economics.

Harrison, Glenn and J. Swarthout (May 2019). "Eye-tracking and economic theories of choice under risk". In: *Journal of the Economic Science Association* 5. DOI: 10.1007/s40881-019-00063-3.

Harrison, Glenn W. and J. Todd Swarthout (Jan. 2023). "Cumulative Prospect Theory in the Laboratory: A Reconsideration". In: ed. by Glenn W. Harrison and Don Ross. Vol. 22. Research in Experimental Economics. Emerald Publishing Limited, pp. 107–192. DOI: 10.1108/S0193-230620230000022003. URL: https://doi.org/10.1108/S0193-230620230000022003.

He, L., P. Analytis, and S. Bhatia (2022). "The Wisdom of Model Crowds". In: *Management Science* 68.5, pp. 3175–3973.

Heimer, Rawley et al. (Jan. 2023). "Dynamic Inconsistency in Risky Choice: Evidence from the Lab and Field". In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4351897.

Henderson, Vicky (2012). "Prospect Theory, Liquidation, and the Disposition Effect". In: *Management Science* 58.2, pp. 445–460. ISSN: 00251909, 15265501. URL: http://www.jstor.org/stable/41406399.

Henderson, Vicky, David Hobson, and SL Alex (2017). "Randomized strategies and prospect theory in a dynamic context". In: *Journal of Economic Theory* 168, pp. 287–300.

Hey, J. and G. Lotito (2009). "Naive, Resolute or Sophisticated? A Study of Dynamic Decision Making". In: *Journal of Risk and Uncertainty* 38.1, pp. 1–25.

Hey, J. and C. Orme (1994). "Investigating Generalizations of Expected Utility Theory Using Experimental Data". In: *Econometrica* 62.6, pp. 1291–1326.

Hey, J. and L. Panaccione (2011). "Dynamic Decision Making: What Do People Do?" In: *Journal of Risk and Uncertainty* 42.2, pp. 85–123.

Hey, J. and M. Paradiso (Mar. 2006). "Preferences Over Temporal Frames In Dynamic Decision Problems: An Experimental Investigation". In: *Manchester School* 74.2, pp. 123–137.

Hey, John and Julia Knoll (June 2011). "Strategies in dynamic decision making - An experimental investigation of the rationality of decision behaviour". In: *Journal of Economic Psychology* 32, pp. 399–409. DOI: 10.1016/j.joep.2011.02.011.

Hilbert, Martin (Nov. 2011). "Toward a Synthesis of Cognitive Biases: How Noisy Information Processing Can Bias Human Decision Making". In: *Psychological bulletin* 138, pp. 211–37. DOI: 10.1037/a0025940.

Hill, W. Trey, John D. Raacke, and J. April Park (2017). "Examining the Priority Heuristic in Conditions of Resource Need Levels". In: *Psychological Reports* 120.5. PMID: 28558547, pp. 824–845. DOI: 10.1177/0033294117709786. eprint: https://doi.org/10.1177/0033294117709786. URL: https://doi.org/10.1177/0033294117709786.

Holt, C. and S. Laury (2002). "Risk Aversion and Incentive Effects". In: *American Economic Review* 92.5, pp. 1644–1655.

Hotaling, Jared M. and Jerome R. Busemeyer (2012). "DFT-D: a cognitive-dynamical model of dynamic decision making". In: *Synthese* 189.1, pp. 67–80. ISSN: 1573-0964. DOI: 10.1007/s11229-012-0157-0. URL: https://doi.org/10.1007/s11229-012-0157-0.

Houser, Daniel et al. (Nov. 2017). "Temptation and Commitment in the Laboratory". In: *Games and Economic Behavior* 107. DOI: 10.1016/j.geb.2017.10.025.

Hu, Sang, Jan Obloj, and Xun Zhou (May 2022). "A Casino Gambling Model Under Cumulative Prospect Theory: Analysis and Algorithm". In: *Management Science* 69. DOI: 10.1287/mnsc.2022.4414.

Huck, S and G Weizsäcker (1999). "Risk, complexity, and deviations from expected-value maximization: Results of a lottery choice experiment". In: *Journal of Economic Psychology* 20, pp. 699–715.

Imas, Alex (Aug. 2016). "The Realization Effect: Risk-Taking After Realized Versus Paper Losses". In: *American Economic Review* 106, pp. 2086–2109. DOI: 10.1257/aer.20140386.

Iyengar, Sheena and Mark Lepper (Jan. 2001). "When Choice is Demotivating: Can One Desire Too Much of a Good Thing?" In: *Journal of personality and social psychology* 79, pp. 995–1006. DOI: 10.1037/0022-3514.79.6.995.

Johnson, Cathleen et al. (2021). "Prince: An improved method for measuring incentivized preferences". In: *Journal of Risk and Uncertainty* 62.1, pp. 1–28. ISSN: 1573-0476. DOI: 10.1007/s11166-021-09346-9. URL: https://doi.org/10.1007/s11166-021-09346-9.

Johnson, J. and R. Busemeyer J. (2001). "Multiple-Stage Decision-Making: The Effect of Planning Horizon Length on Dynamic Consistency". English. In: *Theory and Decision* 51.2-4, pp. 217–246. ISSN: 0040-5833. DOI: 10.1023/A:1015507110940.

Kahneman, D. and A. Tversky (1979). "Prospect Theory: An Analysis of Decision Under Risk". In: *Econometrica* 47.2, pp. 263–291.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler (1991). "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias". In: *Journal of Economic Perspectives* 5.1, pp. 193–206. DOI: 10.1257/jep.5.1.193. URL: https://www.aeaweb.org/articles?id=10.1257/jep.5.1.193.

Kass, R. and A. Raftery (1995). "Bayes Factors". In: *Journal of the American Statistical Association* 90, pp. 773–795.

Kast, Felipe, Stephan Meier, and Dina Pomeranz (Apr. 2016). "Saving more in groups: Field experimental evidence from Chile". In: *Journal of Development Economics* 133. DOI: 10.1016/j.jdeveco.2018.01.006.

Kirman, Alan P. (1992). "Whom or What Does the Representative Individual Represent?" In: *Journal of Economic Perspectives* 6.2, pp. 117–136. DOI: 10.1257/jep.6.2.117. URL: https://www.aeaweb.org/articles?id=10.1257/jep.6.2.117.

Köbberling, V. and P. Wakker (2005). "An Index of Loss Aversion". In: *Journal of Economic Theory* 122, pp. 119–131.

Kreps, D. and R. Wilson (1982). "Reputation and Imperfect Information". In: *Journal of Economic Theory* 27, pp. 253–279.

Kurth-Nelson, Zeb and A. Redish (Oct. 2012). "Dont Let Me Do That! Models of Precommitment". In: *Frontiers in Neuroscience* 6, p. 138. DOI: 10.3389/fnins.2012.00138.

Labudda, Kirsten et al. (2007). "Decision-making and neuroendocrine responses in pathological gamblers". In: *Psychiatry Research* 153.3, pp. 233–243. ISSN: 0165-1781. DOI: https://doi.org/10.1016/j.psychres.2007.02.002. URL: https://www.sciencedirect.com/science/article/pii/S0165178107000315.

Laibson, David (1997). "Golden Eggs and Hyperbolic Discounting". In: *The Quarterly Journal of Economics* 112.2, pp. 443–478. URL: https://EconPapers.repec.org/RePEc:oup:qjecon:v:112:y:1997:i:2:p:443-478..

Lambos, Chrisi and Paul Delfabbro (Aug. 2007). "Numerical Reasoning Ability and Irrational Beliefs in Problem Gambling". In: *International Gambling Studies* 7, pp. 157–171. DOI: 10.1080/14459790701387428.

Langner, Tobias and Martin Krengel (2013). "The mere categorization effect for complex products: The moderating role of expertise and affect". In: *Journal of Business Research* 66.7. Scientific Advancements in Consumer-Retailer Relationships Internationalization of Marketing Communication and Consumer Research, pp. 924–932. ISSN: 0148-2963. DOI: https://doi.org/10.1016/j.jbusres.2011.12.012. URL: https://www.sciencedirect.com/science/article/pii/S014829631100419X.

Lezzi, Emanuela, Piers Fleming, and Daniel Zizzo (Jan. 2015). "Does it Matter Which Effort Task You Use? A Comparison of Four Effort Tasks When Agents Compete for a Prize". In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.2594659.

l'Haridon, Olivier et al. (Aug. 2010). "Separating curvature and elevation: A parametric probability weighting function". In: *Journal of Risk and Uncertainty* 41, pp. 39–65. DOI: 10.1007/s11166-010-9097-6.

Ligneul, Romain et al. (Aug. 2012). "Shifted risk preferences in pathological gambling". In: *Psychological medicine* 43, pp. 1–10. DOI: 10.1017/S0033291712001900.

Limbrick-Oldfield, Eve et al. (Nov. 2020). "Gambling disorder is associated with reduced sensitivity to expected value during risky choice". In: *Journal of Behavioral Addictions* 9. DOI: 10.1556/2006.2020.00088.

Loewenstein, George and Drazen Prelec (1992). "Anomalies in Intertemporal Choice: Evidence and an Interpretation". In: *The Quarterly Journal of Economics* 107.2, pp. 573–597. URL: https://EconPapers.repec.org/RePEc:oup:qjecon:v:107:y:1992:i:2:p:573-597..

Loomes, G., P. Moffat, and R. Sudgen (2002). "A Microeconometric Test of Alternative Stochastic Theories and Risky Choice". In: *Journal of Risk and Uncertainty* 2, pp. 103–130.

Lopes, Lola L and Gregg C Oden (1999). "The Role of Aspiration Level in Risky Choice: A Comparison of Cumulative Prospect Theory and SP/A Theory". In: *Journal of Mathematical Psychology* 43.2, pp. 286–313. ISSN: 0022-2496. DOI: https://doi.org/10.1006/jmps.1999.1259. URL: https://www.sciencedirect.com/science/article/pii/S0022249699912592.

Lopez-Gonzalez, Hibai, Mark D. Griffiths, and Ana Estévez (2020). "Why Some Sports Bettors Think Gambling Addiction Prevented Them from Becoming Winners? A Qualitative Approach to Understanding the Role of Knowledge in Sports Betting

Products”. In: *Journal of Gambling Studies* 36.3, pp. 903–920. ISSN: 1573-3602. DOI: 10.1007/s10899-020-09944-3. URL: https://doi.org/10.1007/s10899-020-09944-3.

Luce, R (1959). “On the Possible Psychophysical Laws”. In: *Psychological Review* 66, pp. 81–95.

Luce, R.Duncan (1991). “Rank- and sign-dependent linear utility models for binary gambles”. In: *Journal of Economic Theory* 53.1, pp. 75–100. ISSN: 0022-0531. DOI: https://doi.org/10.1016/0022-0531(91)90143-R. URL: https://www.sciencedirect.com/science/article/pii/002205319190143R.

Machina, M. (1989). “Dynamic Consistency and Non-expected Utility Models of Choice under Uncertainty”. In: *Journal of Economic Literature* 27.4, pp. 1622–68.

Machina, Mark J. (1983). “Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom”. In: *Foundations of Utility and Risk Theory with Applications*. Ed. by Bernt P. Stigum and Fred Wenstøp. Dordrecht: Springer Netherlands, pp. 263–293. ISBN: 978-94-017-1590-4. DOI: 10.1007/978-94-017-1590-4_15. URL: https://doi.org/10.1007/978-94-017-1590-4_15.

Makridakis, S. and R. Winkler (1983). “Averages of Forecasts: Some Empirical Results”. In: *Management Science* 29.9, pp. 987–996.

Markowitz, H. (1952). “The Utility of Wealth”. In: *Journal of Political Economy* 60.2, pp. 152–158.

Mccabe, Kevin et al. (Nov. 1994). “Preferences, Property Rights, and Anonymity in Bargaining Games”. In: *Games and Economic Behavior* 7, pp. 346–380. DOI: 10.1006/game.1994.1056.

McClennen, E.F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Paperback re-issue. Cambridge University Press. ISBN: 9780521360470.

Mishra, Akanksha, Sonu Singh, and Shubha Shukla (May 2018). “Physiological and Functional Basis of Dopamine Receptors and Their Role in Neurogenesis: Possible Implication for Parkinsons disease”. In: *Journal of Experimental Neuroscience* 12, p. 117906951877982. DOI: 10.1177/1179069518779829.

Moffatt, P. (2016). *Experimetrics*. Macmillan Palgrave.

Moffatt, P., S. Sitzia, and D. Zizzo (2015). “Heterogeneity in Preferences Towards Complexity”. In: *Journal of Risk and Uncertainty* 51, pp. 147–170.

Moffatt, Peter and Simon Peters (2001). “Testing for the Presence of a Tremble in Economic Experiments”. In: *Experimental Economics* 4.3, pp. 221–228. URL: https://EconPapers.repec.org/RePEc:kap:expeco:v:4:y:2001:i:3:p:221-228.

Mohnert, F., T. Pachur, and F. Lieder (2019). “What’s in the Adaptive Toolbox and How Do People Choose From It? Rational Models of Strategy Selection in Risky Choice”. In.

Moloi, Tankiso and Tshilidzi Marwala (May 2020). "Dynamic Inconsistency Theory". In: pp. 43–52. ISBN: 978-3-030-42961-4. DOI: 10.1007/978-3-030-42962-1_5.

National Research Council, US (1999). "Committee on the Social and Economic Impact of Pathological Gambling. Pathological Gambling: A Critical Review". In: US. Washington (DC): National Academies Press (US).

Navarro, Daniel J. et al. (2006). "Modeling individual differences using Dirichlet processes". In: *Journal of Mathematical Psychology* 50.2. Special Issue on Model Selection: Theoretical Developments and Applications, pp. 101–122. ISSN: 0022-2496. DOI: https://doi.org/10.1016/j.jmp.2005.11.006. URL: https://www.sciencedirect.com/science/article/pii/S0022249605000969.

Nebout, A. and D. Dubois (2014). "When Allais Meets Ulysses: Dynamic Axioms and the Common Ratio Effect". In: *Journal of Risk and Uncertainty* 48.1, pp. 19–49. ISSN: 1573-0476.

Nebout, A. and M. Willinger (2014). *Are Non-Expected Utility Individual Really Dynamic Inconsistent? Experimental Evidence*. Discussion Papers 99/26. LAMETA-CNRS, University of Montpellier.

Newton M, A. and A. Raftery (1994). "Approximate Bayesian Inference with the Weighted Likelihood Bootstrap". In: *Journal of the Royal Statistical Society, Series B* 56, pp. 3–48.

Nilsson, H., J. Rieskamp, and E-J. Wagenmakers (2011). "Hierarchical Bayesian Parameter Estimation for Cumulative Prospect Theory". In: *Journal of Mathematical Psychology* 55.1. Special Issue on Hierarchical Bayesian Models, pp. 84–93. ISSN: 0022-2496. DOI: https://doi.org/10.1016/j.jmp.2010.08.006. URL: https://www.sciencedirect.com/science/article/pii/S0022249610001070.

Nursimulu, Anjali D. and Peter Bossaerts (2014). "Risk and Reward Preferences under Time Pressure". In: *Review of Finance* 18.3, pp. 999–1022. URL: https://ideas.repec.org/a/oup/revfin/v18y2014i3p999-1022..html.

Oberholzer, Yvonne, Sebastian Olschewski, and Benjamin Scheibehenne (2021). *Complexity Aversion in Risky Choices and Valuations: Moderators and Possible Causes*. DOI: 10.31234/osf.io/tnvzr. URL: psyarxiv.com/tnvzr.

Odean, Terrance (1998). "Are Investors Reluctant to Realize Their Losses?" In: *The Journal of Finance* 53.5, pp. 1775–1798. DOI: https://doi.org/10.1111/0022-1082.00072. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/0022-1082.00072. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/0022-1082.00072.

O'Donoghue, T. and M. Rabin (1999). "Doing It Now or Later". In: *American Economic Review* 89.1, pp. 103–124.

Olschewski, S and J. Rieskamp (2021). "Distinguishing three effects of time pressure on risk taking: Choice consistency, risk preference and strategy selection". In: *Journal of Behavioral Decision Making* 34.4, pp. 541–554.

Osberghaus, Daniel (2017). "Prospect theory, mitigation and adaptation to climate change". In: *Journal of Risk Research* 20.7, pp. 909–930. DOI: 10.1080/13669877.2015.1121907. eprint: https://doi.org/10.1080/13669877.2015.1121907. URL: https://doi.org/10.1080/13669877.2015.1121907.

Pachur, Thorsten, Renata S. Suter, and Ralph Hertwig (2017). "How the twain can meet: Prospect theory and models of heuristics in risky choice". In: *Cognitive Psychology* 93, pp. 44–73. ISSN: 0010-0285. DOI: https://doi.org/10.1016/j.cogpsych.2017.01.001. URL: https://www.sciencedirect.com/science/article/pii/S0010028516300093.

Parpart, Paula, Matt Jones, and Bradley C. Love (2018). "Heuristics as Bayesian inference under extreme priors". In: *Cognitive Psychology* 102, pp. 127–144. ISSN: 0010-0285. DOI: https://doi.org/10.1016/j.cogpsych.2017.11.006. URL: https://www.sciencedirect.com/science/article/pii/S0010028517303286.

Patmanidis, Spyridon et al. (Sept. 2019). "A comparison between Nonlinear Least Squares and Maximum Likelihood estimation for the prediction of tumor growth on experimental data of human and rat origin". In: *Biomedical Signal Processing and Control* 54, p. 101639. DOI: 10.1016/j.bspc.2019.101639.

Payne, J., J. Bettman, and E. Johnson (1993). *The Adaptive Decision Maker*. Cambridge University Press.

Payne, John, James Bettman, and Eric Johnson (July 1988). "Adaptive Strategy Selection in Decision Making". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 14, pp. 534–552. DOI: 10.1037/0278-7393.14.3.534.

Payne, John W (1976). "Heuristic Search Processes in Decision Making". In: *Advances in Consumer Research*.

Payne, John W. (2005). "It is Whether You Win or Lose: The Importance of the Overall Probabilities of Winning or Losing in Risky Choice". In: *Journal of Risk and Uncertainty* 30.1, pp. 5–19. ISSN: 1573-0476. DOI: 10.1007/s11166-005-5831-x. URL: https://doi.org/10.1007/s11166-005-5831-x.

Payne, John W., James R. Bettman, and Mary Frances Luce (1996). "When Time Is Money: Decision Behavior under Opportunity-Cost Time Pressure". In: *Organizational Behavior and Human Decision Processes* 66.2, pp. 131–152. URL: https://EconPapers.repec.org/RePEc:eee:jobhdp:v:66:y:1996:i:2:p:131-152.

Payne, John W., Dan J. Laughhunn, and Roy Crum (1980). "Translation of Gambles and Aspiration Level Effects in Risky Choice Behavior". In: *Management Science* 26.10, pp. 1039–1060. DOI: 10.1287/mnsc.26.10.1039. eprint: https://doi.org/10.1287/mnsc.26.10.1039. URL: https://doi.org/10.1287/mnsc.26.10.1039.

Peel, D.A. and Jie Zhang (2012). "On the potential for observational equivalence in experiments on risky choice when a power value function is assumed". In: *Economics Letters* 116.1, pp. 8–10. ISSN: 0165-1765. DOI: https://doi.org/10.1016/j.econlet.2011.12.116. URL: https://www.sciencedirect.com/science/article/pii/S0165176511006392.

Pelletier, Marie-France and Robert Ladouceur (Apr. 2007). "The effect of knowledge of mathematics on gambling behaviours and erroneous perceptions". In: *International Journal of Psychology* 42, pp. 134–140. DOI: 10.1080/00207590600788047.

Peterson, J. et al. (2021). "Using Large-scale experiments and Machine Learning to Discover Theories of Human Decision-Making". In: *Science* 372, pp. 1209–1214.

Ploner, Matteo (Mar. 2017). "Hold on to it? An experimental analysis of the disposition effect". In: *Judgment and Decision Making* 12, pp. 118–127. DOI: 10.1017/S1930297500005660.

Plummer, M. (2017). *JAGS Version 4.3.0 User Manual*. Tech. rep.

Prelec, D. (1998). "The Probability Weighting Function". English. In: *Econometrica* 66.3, pp. 497–527. ISSN: 00129682.

Preston, Malcolm G. and Paolo Baratta (1948). "An experimental study of the auction-value of an uncertain outcome." In: *The American journal of psychology* 61 2, pp. 183–93.

Quiggin, J. (1982). "A Theory of Anticipated Utility". In: *Journal of Economic Behavior and Organization* 3.4, pp. 323–343.

R Core Team (2023). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: https://www.R-project.org/.

Reny, Philip (1992). "Backward Induction, Normal Form Perfection and Explicable Equilibria". In: *Econometrica* 60.3, pp. 627–49. URL: https://EconPapers.repec.org/RePEc:ecm:emetrp:v:60:y:1992:i:3:p:627-49.

Rieskamp, J. (2008). "The Probabilistic Nature of Preferential Choice". In: *Experimental Psychology: Learning, Memory and Cognition* 44, p. 1.

Ring, Patrick et al. (June 2018). "It's All About Gains: Risk Preferences in Problem Gambling". In: *Journal of Experimental Psychology: General* 147. DOI: 10.1037/xge0000418.

Rouder, J. and J. Lu (2005). "An Introduction to Bayesian Hierarchical Models with an Application in the Theory of Signal Detection". In: *Psychonomic Bulletin & Review* 55, pp. 84–93.

Rubinstein, A. (1988). "Similarity and Decision-making under Risk (Is Htere a Utility Theory Resolution to the Allais Paradox?)" In: *Journal of Economic Theory* 46, pp. 145–153.

Saha, Atanu (1993). "Expo-Power Utility: A Flexible Form for Absolute and Relative Risk Aversion". In: *American Journal of Agricultural Economics* 75.4, pp. 905–913. URL: https://ideas.repec.org/a/oup/ajagec/v75y1993i4p905-913..html.

Samuelson, William and Richard Zeckhauser (1988). "Status quo bias in decision making". In: *Journal of Risk and Uncertainty* 1.1, pp. 7–59. ISSN: 1573-0476. DOI: 10.1007/BF00055564. URL: https://doi.org/10.1007/BF00055564.

Saqib, Najam U. and Eugene Y. Chan (2015). "Time pressure reverses risk preferences". In: *Organizational Behavior and Human Decision Processes* 130, pp. 58–68. ISSN: 0749-5978. DOI: https://doi.org/10.1016/j.obhdp.2015.06.004. URL: https://www.sciencedirect.com/science/article/pii/S0749597815000722.

Sarin, Rakesh. and Peter. Wakker (1998). "Dynamic Choice and NonExpected Utility". In: *Journal of Risk and Uncertainty* 17.2, pp. 87–119. ISSN: 08955646, 15730476. URL: http://www.jstor.org/stable/41760907.

Scheibehenne, B, J Rieskamp, and E Wagenmakers (2013). "Testing Adaptive Toolbox Models: A Bayesian Hierarchical Approach". In: *Phychological Review* 120.1, pp. 39–64.

Schmidt, U and H. Zank (2005). "What is Loss Aversion". In: *Journal of Risk &Uncertainty* 30.2, pp. 157–167.

Scholten, M. and D. Read (2014). "Prospect Theory and the Forgotten Fourfold Pattern of Risk Preferences". In: *Journal of Risk and Uncertainty* 48.1, pp. 67–83.

Schwartz, Alan, Julie Goldberg, and Gordon Hazen (Feb. 2008). "Prospect theory, reference points, and health decisions". In: *Judgment and Decision Making* 3, pp. 174–180. DOI: 10.1017/S1930297500001510.

Selten, R. (1975). "Rexamination of the Perfectness Concept for Equilibrium Points in Extensive Games". In: *International Journal of Game Theory* 4.1, pp. 25–55. ISSN: 1432-1270. DOI: 10.1007/BF01766400.

Selten, Reinhard (1965). "SPIELTHEORETISCHE BEHANDLUNG EINES OLIGOPOLMODELLS MIT NACHFRAGETRÄGHEIT. TEIL II: EIGENSCHAFTEN DES DYNAMISCHEN PREISGLEICHGEWICHTS". In: *Zeitschrift für die gesamte Staatswissenschaft / Journal of Institutional and Theoretical Economics* 121.4, pp. 667–689. ISSN: 00442550. URL: http://www.jstor.org/stable/40748908.

Shead, N., Mitchell Callan, and David Hodgins (Aug. 2008). "Probability discounting among gamblers: Differences across problem gambling severity and affect-regulation expectancies". In: *Personality and Individual Differences* 45. DOI: 10.1016/j.paid.2008.06.008.

Shefrin, Hersh and Meir Statman (1985). "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence". In: *The Journal of Finance* 40.3, pp. 777–790. ISSN: 00221082, 15406261. DOI: 10.2307/2327802. URL: http://www.jstor.org/stable/2327802.

Shoji, Isao and Sumei Kanehiro (2016). "Disposition effect as a behavioral trading activity elicited by investors' different risk preferences". In: *International Review of Financial Analysis* 46.C, pp. 104–112. URL: https://EconPapers.repec.org/RePEc:eee:finana:v:46:y:2016:i:c:p:104-112.

Sitzia, Stefania and Daniel J. Zizzo (2011). "Does product complexity matter for competition in experimental retail markets?" In: *Theory and Decision* 70.1, pp. 65–82.

Slovic, Paul (1969). *Differential effects of real versus hypothetical payoffs on choices among gambles.* US. DOI: 10.1037/h0027489.

Smith, Donald (2001). "Gambling Behaviour and Understanding of Probability Concepts Among University Students". In: *Victoria University of Technology*.

Sneddon, Robert and R.Duncan Luce (2001). "Empirical Comparisons of Bilinear and Nonbilinear Utility Theories". In: *Organizational Behavior and Human Decision Processes* 84.1, pp. 71–94. ISSN: 0749-5978. DOI: https://doi.org/10.1006/obhd.2000.2919. URL: https://www.sciencedirect.com/science/article/pii/S0749597800929194.

Sokolowska, Joanna (2006). "Risk Perception and Acceptance: One Process or Two?" In: *Experimental Psychology* 53.4. PMID: 17176656, pp. 247–259. DOI: 10.1027/1618-3169.53.4.247. eprint: https://doi.org/10.1027/1618-3169.53.4.247. URL: https://doi.org/10.1027/1618-3169.53.4.247.

Sonsino, D., U. Benzion, and G. Mador (2002). "The Complexity Effects on Choice under Uncertainty". In: *The Economic Journal* 112.482, p. 9.

Stahl, D. (2014). "Heterogeneity of Ambiguity Preferences". In: *The Review of Economics and Statistics* 96.5, pp. 609–617.

— (2018). "An Empirical Evaluation of the Toolbox Model of Lottery Choices". In: *The Review of Economics and Statistics* 100.3, pp. 528–534.

Stone, M. (1974). "Cross-Validatory Choice and Assessment of Statistical Predictions". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 36.2, pp. 111–133. DOI: https://doi.org/10.1111/j.2517-6161.1974.tb00994.x. eprint: https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1974.tb00994.x. URL: https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.2517-6161.1974.tb00994.x.

Stott, H. (2006). "Cumulative Prospect Theory's Functional Menagerie". English. In: *Journal of Risk and Uncertainty* 32.2, pp. 101–130. ISSN: 0895-5646. DOI: 10.1007/s11166-006-8289-6.

Strotz, R. (1955-56). "Myopia and Inconsistency in Dynamic Utility Maximization". In: *The Review of Economic Studies* 23, pp. 165–180.

Takahashi, Hidehiko et al. (Dec. 2010). "Dopamine D-1 Receptors and Nonlinear Probability Weighting in Risky Choice". In: *The Journal of neuroscience : the official* *journal of the Society for Neuroscience* 30, pp. 16567–72. DOI: 10.1523/JNEUROSCI.3933-10.2010.

Takahashi, Taiki (Sept. 2008). "Biophysics of Risk Aversion Based on Neurotransmitter Receptor Theory". In: *Neuro endocrinology letters* 29, pp. 399–404. DOI: 10.2139/ssrn.1878915.

Thaler, Richard (Mar. 1980). "Toward a Positive Theory of Consumer Choice". In: *Journal of Economic Behavior & Organization* 1, pp. 39–60. DOI: 10.1016/0167-2681(80)90051-7.

Thaler, Richard and Eric Johnson (June 1990). "Gambling With the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice". In: *Management Science* 36, pp. 643–660. DOI: 10.1287/mnsc.36.6.643.

Thaler, Richard and C. Sunstein (June 2009). *NUDGE: Improving Decisions About Health, Wealth, and Happiness*. Vol. 47. ISBN: 9780141040011.

Thaler, Richard H. (2000). "From Homo Economicus to Homo Sapiens". In: *Journal of Economic Perspectives* 14.1, pp. 133–141. DOI: 10.1257/jep.14.1.133. URL: https://www.aeaweb.org/articles?id=10.1257/jep.14.1.133.

Thorngate, W. (1980). "Efficient decision heuristics". In: *Behavioral Science* 25, pp. 219–225.

Todd, Peter et al. (Jan. 1999). *Simple Heuristics That Make Us Smart*.

Tom, Sabrina et al. (Feb. 2007). "The Neural Basis of Loss Aversion in Decision-Making Under Risk". In: *Science (New York, N.Y.)* 315, pp. 515–8. DOI: 10.1126/science.1134239.

Trope, Yaacov and Ayelet Fishbach (Nov. 2000). "Counteractive self-control in overcoming tempation". In: *Journal of personality and social psychology* 79, pp. 493–506. DOI: 10.1037//0022-3514.79.4.493.

Tversky, A. and D. Kahneman (1992). " Advances in Prospect Theory: Cumulative Representation of Uncertainty". In: *Journal of Risk and Uncertainty* 5.4, pp. 297–323.

Tversky, Amos (1972). "Elimination by Aspects: A Theory of Choice". In: *Psychological Review* 79.4, pp. 281–299. DOI: 10.1037/h0032955.

— (1975). "A Critique of Expected Utility Theory: Descriptive and Normative Considerations". In: *Erkenntnis* 9.2, pp. 163–173.

Tversky, Amos and Daniel Kahneman (1975). "Judgment under Uncertainty: Heuristics and Biases". In: *Utility, Probability, and Human Decision Making: Selected Proceedings of an Interdisciplinary Research Conference, Rome, 3-6 September, 1973*. Ed. by Dirk Wendt and Charles Vlek. Dordrecht: Springer Netherlands, pp. 141–162. DOI: 10.1007/978-94-010-1834-0_8. URL: https://doi.org/10.1007/978-94-010-1834-0_8.

Venkatraman, Vinod, John W. Payne, and Scott A. Huettel (2014). "An overall probability of winning heuristic for complex risky decisions: Choice and eye fixation evidence". In: *Organizational Behavior and Human Decision Processes* 125.2, pp. 73–87. ISSN: 0749-5978. DOI: https://doi.org/10.1016/j.obhdp.2014.06.003. URL: https://www.sciencedirect.com/science/article/pii/S0749597814000491.

Vieider, Ferdinand (Dec. 2012). "Moderate Stake Variations for Risk and Uncertainty, Gains and Losses: Methodological Implications for Comparative Studies". In: *Economics Letters* 117, 718721. DOI: 10.1016/j.econlet.2011.12.028.

Vieider, Ferdinand M. et al. (2015). "Common Componetns of Risk and Uncertainty Attitudes Across Contexts and Domains: Evidence from 30 Countries". In: *Journal of the European Economic Association* 13.3, pp. 421–452. DOI: https://doi.org/10.1111/jeea.12102. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jeea.12102. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jeea.12102.

Viscusi, W. Kip (1989). "Prospective reference theory: Toward an explanation of the paradoxes". In: *Journal of Risk and Uncertainty* 2.3, pp. 235–263.

Von Neumann, John and Oskar Morgenstern (1947). *Theory of games and economic behavior, 2nd rev. ed.* Princeton, NJ, US: Princeton University Press, pp. xviii, 641–xviii, 641.

Vuong Q., H. (1989). "Likelihood Ratio Tests for Model Selection and non-nested Hypothesis". In: *Econometrica* 57, pp. 307–33.

Wakker, P. (2010). *Prospect Theory*. Cambridge University Press.

Wakker, Peter, Ido Erev, and Elke U Weber (1994). "Comonotonic Independence: The Critical Test between Classical and Rank-Dependent Utility Theories". In: *Journal of Risk and Uncertainty* 9.3, pp. 195–230. URL: https://EconPapers.repec.org/RePEc:kap:jrisku:v:9:y:1994:i:3:p:195-230.

Wakker, Peter and Amos Tversky (1993). "An Axiomatization of Cumulative Prospect Theory". In: *Journal of Risk and Uncertainty* 7.2, pp. 147–75. URL: https://EconPapers.repec.org/RePEc:kap:jrisku:v:7:y:1993:i:2:p:147-75.

Weber, Martin and Colin F. Camerer (1998). "The disposition effect in securities trading: an experimental analysis". In: *Journal of Economic Behavior Organization* 33.2, pp. 167–184. ISSN: 0167-2681. DOI: https://doi.org/10.1016/S0167-2681(97)00089-9. URL: https://www.sciencedirect.com/science/article/pii/S0167268197000899.

Weglarczyk, Stanislaw (Jan. 2018). "Kernel density estimation and its application". In: *ITM Web of Conferences* 23, p. 00037. DOI: 10.1051/itmconf/20182300037.

Wilcox, N. (2011). "Stochastically more Risk Averse: A Contextual Theory of Stochastic Discrete Choice under Risk". In: *Journal of Econometrics* 162.1, pp. 89 –104.

Wilcoxon, Frank (1992). "Individual Comparisons by Ranking Methods". In: *Breakthroughs in Statistics: Methodology and Distribution*. Ed. by Samuel Kotz and Norman L. Johnson. New York, NY: Springer New York, pp. 196–202. ISBN: 978-1-4612-4380-9. DOI: 10.1007/978-1-4612-4380-9_16. URL: https://doi.org/10.1007/978-1-4612-4380-9_16.

Wu, George and Richard Gonzalez (Dec. 1996). "Curvature of the Probability Weighting Function". In: *Management Science* 42, pp. 1676–1690. DOI: 10.1287/mnsc.42.12.1676.

Zeisberger, Stefan (2022). "Do people care about loss probabilities?" In: *Journal of Risk and Uncertainty* 65.2, pp. 185–213. ISSN: 1573-0476. DOI: 10.1007/s11166-022-09391-y. URL: https://doi.org/10.1007/s11166-022-09391-y.

Zermelo, Ernst (1913). "Uber eine Anwendung der Mengenlehre auf die Theorie ¨des Schachspiels". In: *Fifth Congress Mathematicians, (Cambridge 1912),Cambridge University Press 1913*.

Zhong, Songfa et al. (Sept. 2009). "A Neurochemical Approach to Valuation Sensitivity over Gains and Losses". In: *Proceedings. Biological sciences / The Royal Society* 276, pp. 4181–8. DOI: 10.1098/rspb.2009.1312.

Zilker, V, R Hertwig, and T Pachur (2020). "Age Differences in Risk Attitude Are Shaped by Option Complexity". In: *Journal of Experimental Psychology: General* 149.9, pp. 1644–1683.

Zindel, Márcia Longen, Thilo Zindel, and Marcelo Grangeiro Quirino (2014). "Cognitive bias and their implications on the financial market". In: *International Journal of Engineering and Technology* 14.3, pp. 11–17.