THE ECONOMICS OF ADDICTION:
AN EXPERIMENTAL INVESTIGATION

ANDRÉ HOFMEYR

Thesis Presented for the Degree of

DOCTOR OF PHILOSOPHY

in the School of Economics

UNIVERSITY OF CAPE TOWN

FEBRUARY 2015

Abstract

Addiction is an ideal puzzle for economic theory: why do most addicts expend resources to acquire their targets of addiction but then incur real costs to try and reduce or limit their consumption of these goods? Furthermore, why is the typical course of addiction characterised by repeated unsuccessful attempts to quit prior to final abstention? From the standpoint of standard consumer theory in economics these patterns of behaviour are difficult to rationalise. There is a rich theoretical literature in economics which models habit-forming behaviours, of which addiction is the exemplar, but there is a paucity of experimental economic studies eliciting and comparing the preferences – specifically, risk and time preferences – that economic theory suggests may differ between addicts and non-addicts. The experimental research that has been conducted has been dominated by psychologists, and some economists have begun to follow their methodological lead. However, detailed reviews of the experimental literature on addiction highlight a number of methodological and statistical flaws in the ways these data have been collected and analysed. This thesis is primarily concerned with methodological and statistical issues at the boundary between economics and psychology as these bear upon developing a general, consistent explanation of addiction. An incentive-compatible experimental design is formulated which lends itself to the estimation of several different theories of choice under risk and over time. In addition, a full information maximum likelihood statistical framework, which is consistent with the data generating processes proposed by structural theories and accounts for subject errors in decision making, is used to explore the relationship between risk preferences, time preferences and addiction. This thesis challenges some of the maintained assumptions in the addiction literature; e.g., that the probability discounting and hyperbolic discounting models best characterise choice under risk and over time, respectively. But it also replicates a previous finding; i.e., that smokers tend to discount the future more heavily than non-smokers. It shows, therefore, that some results withstand careful methodological and statistical scrutiny, whereas other results do not. Ultimately, this thesis argues that tools from experimental economics and econometrics, which have been under-used in addiction studies, contribute to a more accurate and reliable characterisation of this phenomenon.

Acknowledgements

As a young research assistant to Professor Don Ross in 2006 I was exposed to the use of experimental economic tools to investigate the behavioural correlates of addiction. That experimental economics could be used to study addiction fascinated me and with the encouragement of Professor Ross I registered for my PhD in 2008 to conduct research in this area. Generous funding from the National Responsible Gambling Programme (NRGP) in South Africa allowed a research team, of which I was a member, to design an ambitious longitudinal study focussing on the risk and time preferences of gamblers. In late 2008, Professor Ross introduced me to Professor Glenn Harrison who had just published an influential paper in *Econometrica* focussing on the elicitation and estimation of risk and time preferences. Professor Harrison kindly offered his time and expertise to the research team and played a crucial role in the development of the project's experimental design. During this process Professor Harrison agreed to act as co-supervisor on my PhD.

The longitudinal study of risk and time preferences and gambling behaviour, which was to form the basis of my PhD, took 18 months to plan and 18 months to run. Halfway through the study, the research team uncovered widespread fieldworker fraud which had contaminated the entire sample: some fieldworkers had paid subjects the highest prizes on offer in the experiments and not on the basis of their choices in the experiments. Needless to say, this destroyed the incentive compatibility of the experiments and meant that I had to find new data to use in my thesis.

Professors Ross and Harrison put me in touch with Professor John Monterosso who had collected some experimental data on the discounting behaviour of smokers in California. Professor Monterosso then contacted Professor Edythe London who had also collected experimental data on the discounting behaviour of smokers in California. I am deeply indebted to them for allowing me to use these data for Chapter 3 of this thesis and for many helpful discussions.

To conduct experiments investigating the risk and time preferences of a sample of smokers and non-smokers at the University of Cape Town (UCT), I obtained a

Research Development grant from the Emerging Researcher Programme (ERP) at UCT. Through the kind assistance of Professor Corne van Walbeek, who contacted the American Cancer Society on my behalf, I was granted some additional funding which allowed me to expand the sample. Professor Todd Swarthout developed the software I used to elicit risk and time preferences and generously adapted it for the South African context. The data that were collected are analysed in Chapter 5. For their invaluable assistance in making this component of my research possible, I am most grateful to the ERP, Professors van Walbeek and Swarthout, and the American Cancer Society.

I would like to thank Professors Justine Burns and Martine Visser for originally introducing me to experimental economics, for allowing me to run my first experiments in 2005, and for the many insightful conversations we have had. For their unfailing encouragement, I would also like to thank Doctor George Ainslie and Professor Harold Kincaid. As mentors and collaborators they have taught me much of what I know about experimental psychology and the philosophy of science.

I am profoundly grateful and indebted to my supervisors, Professors Ross and Harrison. They exemplify the highest standards of scholarship and they have encouraged, nudged, and supported me throughout my PhD. Their tireless energy and insistence on "doing it right" have inspired me as a researcher and I will always remember their wit and intellect.

My mother Jane, father Karl, and sister Kate have been a constant source of encouragement and support. They have generously allowed me to take writing retreats at their homes, without which this thesis would not have been finished, and they have always been there when I needed them.

I dedicate this work to my wife Sarah. She met me before I started my PhD, married me halfway through it, and her love and support has made the completion of this thesis possible.

TABLE OF CONTENTS

1. INTRODUCTION

This thesis is about addiction and the ways in which economists and researchers from other disciplines model it theoretically, investigate it experimentally, and analyse it statistically. While there is a rich theoretical literature in economics which models habit-forming behaviours, of which addiction is the exemplar, there is a paucity of experimental economic studies eliciting and comparing the preferences – specifically, instantaneous[1] risk and time preferences, intertemporal risk preferences, and subjective beliefs – that economic theory suggests may differ between addicts and non-addicts.

The experimental research that has been conducted has been dominated by psychologists, and some economists have begun to follow their methodological lead. This thesis shows that tools from experimental economics and econometrics, which have been under-used in addiction studies, contribute to a more accurate and reliable characterisation of this phenomenon. Thus, this thesis is primarily concerned with methodological and statistical issues at the boundary between economics and psychology as these bear upon developing a general, consistent explanation of addiction.

Chapter 2 provides a motivation for studying addiction by chronicling the history of substance use and abuse, discussing current trends in drug consumption and dependence, and highlighting the costs that addiction imposes on society. This broad overview makes it clear that addiction is a pressing public health concern that warrants the attention that economists have increasingly been devoting to it.

Psychologists, with clinical concerns in mind, have worked to devise a diagnostic classification system which distinguishes between addicts and non-addicts. Chapter 2 traces the history of the diagnostic classification of addiction and highlights the view

[1] The prefix "instantaneous" is used to differentiate instantaneous risk and time preferences from intertemporal risk preferences. Intertemporal risk preferences refer to preferences over intertemporal lotteries, the outcomes of which may be serially correlated. By contrast, instantaneous risk and time preferences define, respectively, atemporal attitudes to risk and uncertainty, and the valuation of goods which are available at different points in time.

that classification systems do not solely serve scientific ends but that they develop in concert with wider political, social, and economic forces. The chapter draws attention to a number of issues with the measurement and interpretation of addiction. It also recognises, though, that no classification system is ever perfect, and that research into substance dependence has made progress over the last 50 years by refining the definition and measurement of addiction.

The disease model of addiction is the primary understanding of substance dependence among the medical fraternity, the public, and the media. Chapter 2 shows why the disease model of addiction gained general acceptance in the 20th century. It then outlines Heyman's (2009) critique of this model and his argument that addiction is more usefully characterised as a disorder of choice and thereby amenable to economic analysis.

Ardent proponents of the disease model and Heyman's (2009) choice-based account of addiction appear to lie at two extreme ends of a continuum of explanations for substance dependence. But Ross, Sharp, Vuchinich and Spurrett (2008) explain that these apparently contradictory hypotheses are ultimately compatible if one attends to distinctions between different levels of scientific analysis. Chapter 2 therefore outlines the reconciliation which Ross, Sharp, Vuchinich and Spurrett (2008) formulate because this resolves the tension between these two views and motivates research, at different but inextricably linked levels, into the various causes of substance dependence. The chapter then discusses economic models of choice which are relevant to the study of addiction.

In many ways addiction is an ideal puzzle for economic theory. Why do most addicts expend resources to acquire their targets of addiction while maintaining concurrent commitments, involving real costs, to personal policies and programmes aimed at reducing or limiting their consumption of these goods? Furthermore, why is the typical course of addiction characterised by repeated unsuccessful attempts to quit prior to final abstention? From the standpoint of standard consumer theory in economics these patterns of behaviour are difficult to rationalise.

The economic models reviewed in Chapter 2, which seek to explain these behavioural puzzles, are organised into three categories: models of rational addiction; dual self and dual system models; and exogenous addiction models. The chapter outlines the models' similarities and differences and a deep philosophical issue on which the comparisons among these models force attention: the nature of agency in economic theory. The discussion of these models is used to provide theoretical justification for the empirical analyses conducted in later chapters. These empirical analyses centre on the instantaneous risk and time preferences of addicts and non-addicts, so particular attention is given to the role of these preferences in the different theoretical models.

Instantaneous risk and time preferences are important components of economic models as they define, respectively, atemporal attitudes toward risk and uncertainty, and the valuation of goods which are available at different points in time. It makes intuitive sense that instantaneous risk and time preferences may be an important factor in addiction. For example, people who discount the future heavily will tend to focus on the short-term benefits of addictive consumption while putting relatively less weight on the future costs, thereby inclining them to consume the addictive good. Similarly, the decision to consume an addictive good involves clear risks, such as the potential for negative health consequences, so it is presumably linked to an individual's attitude towards risk.

The economic models in Chapter 2 adopt different assumptions about agents' time preferences and they place more or less emphasis on the role that risk plays in driving choice behaviour. Given these differences across the economic models of addiction, and the plethora of different theories of choice under risk and over time, the question, from methodological and statistical perspectives, turns to how best to elicit and analyse the instantaneous risk and time preferences of addicts and non-addicts.

The three empirical chapters of this thesis (Chapters 3, 4 and 5) provide detailed reviews of the experimental and statistical tools which have been used to compare the instantaneous risk and time preferences of addicts and non-addicts. Chapter 3 reviews studies investigating the instantaneous time preferences of smokers and non-smokers and finds that in the majority of these studies there is a positive relationship between smoking and time discounting: smokers indeed tend to discount the future more

heavily than non-smokers. Chapter 4 introduces and critiques the probability discounting model which has become a popular framework in psychology and addiction studies for exploring people's atemporal attitudes to risk. Chapter 5, building on the exposition in Chapter 4, reviews studies investigating the instantaneous risk preferences of smokers and non-smokers and finds that the relationship between attitudes to risk and smoking is equivocal.

Unfortunately, these reviews paint a depressing picture of the methodological and statistical practices commonly used in experimental studies of addiction. First, the bulk of this empirical research makes a number of strong, and unwarranted, assumptions about the nature of instantaneous risk and time preferences (e.g., that discounting is hyperbolic, utility is linear in money, and risk attitudes are determined solely by the shape of a probability weighting function). Second, the experimental tools used in these studies typically do not satisfy Smith's (1982) precepts for controlled microeconomic experiments (e.g., the studies often employ hypothetical incentives and use elicitation mechanisms which are susceptible to being "gamed" by subjects). Finally, the statistical tools used to analyse the experimental data are potentially mis-specified (e.g., subjects' "inconsistent" choices are assumed away rather than modelled statistically and point estimates are used as data in statistical models).

Given these problems with experimental design and statistical analysis, how should one go about eliciting and comparing the instantaneous risk and time preferences of addicts and non-addicts so as to draw accurate and robust inferences? Harrison, Lau and Rutström (2014) argue that theory, experimental design, and econometrics are complementary. In other words, it is crucial to review and understand theory because this informs experimental design and analysis. And analysis itself should be constrained by, and interpreted jointly with, theoretical considerations, prior empirical work, complementary data, econometric methodology, and intended applications.

Thus, it is crucial to have an intimate theoretical knowledge of different models of instantaneous risk and time preferences and their implications for choice behaviour generally, and addiction, specifically. This knowledge should then be used to design experiments to elicit instantaneous risk and time preference data that are amenable to

rigorous statistical analysis. Clearly, the statistical models which are developed to analyse these data must be structured by the theory and informed by the experimental design so that valid inferences can be drawn. And the results from these models must be interpreted jointly with the theory and experimental design which motivated the models' development.

Consequently, this thesis advocates, and lends more credence to, experimental designs which are incentive compatible, incorporate transparent payment schemes and offer salient rewards, and include tasks that are easily understood. Only experimental designs which satisfy these criteria promote the truthful revelation of preferences. In addition, an experimental design should lend itself to the estimation of several different theories of choice under risk and over time because these have different implications for our understanding of addiction.

For example, the quasi-hyperbolic or β-δ discounting model has become a popular theory of instantaneous time preferences because it captures the potential for time inconsistency, an apparent hallmark of addiction, in a mathematically tractable framework. To facilitate estimation of this model it is helpful to vary the front end delay (FED) to the smaller, sooner (SS) reward in a discounting task. Specifically, a zero day FED allows one to pin down the estimate of β, which captures a "passion for the present" or "present-bias" in decision making, whereas positive FEDs (e.g., a delay of 7 days to the SS reward) allow one to estimate the long-term discounting parameter δ. Experimental designs which do not incorporate a zero day FED must assume that a subject's *present*-bias persists beyond the *present* to estimate β. Thus, theory must inform experimental design so that features of the theory can be identified by the experiment and subsequently estimated statistically.

As another example of how theory should inform experimental design and statistical analysis consider the work of Andersen, Harrison, Lau and Rutström (2008) (AHLR). As a matter of theory, time preferences are defined over time-dated utility flows, not flows of money. These are equivalent if a utility function is linear but AHLR showed that if a utility function is concave then the assumption of linearity will, for the same observed choices, bias the estimation of discounting parameters upwards. Thus, to draw accurate inferences about discounting behaviour it is important to incorporate

utility function curvature in the estimation of discounting models. AHLR used a risk preference task to elicit the concavity of the utility function, under the assumption that expected utility theory characterises choice under risk, and a time preference task to elicit discounting behaviour. The researchers then adopted a full information maximum likelihood statistical framework to estimate the parameters of time preference models jointly with the curvature of the utility function. This statistical approach is adopted in Chapter 5 because it represents a consistent union between theory, experimental design, and econometrics.

At a more general statistical level, this thesis stresses the importance of understanding and respecting the type of data generating processes (DGP) encountered in addiction research, so that valid inferences can be drawn. A practice which is far too common in the experimental literature on addiction is a two-step approach to statistical analysis. In the case of time preferences, this two-step approach entails using non-linear least squares (NLLS), or some other technique, to estimate discounting parameters at the level of the individual and then using these point *estimates* as *data* in subsequent statistical models. The problem with the two-step approach, other than that it often uses tiny samples to estimate discounting parameters at the level of the individual, is that estimated discounting parameters are estimates, and not data. Such estimates comprise both a point estimate (of the mean) and a standard error, and to use only the point estimate is to throw away information on the uncertainty of that estimate. Moreover, using an estimated discounting parameter as a datum violates the statistical assumptions of the second-stage models: specifically, that the covariates are measured without error. Thus, the statistical inferences drawn from this approach are simply not valid.

The full information maximum likelihood statistical approach which this thesis adopts estimates instantaneous risk and time preference parameters as a linear function of observable characteristics (e.g., age, gender, and addiction) so that the statistical uncertainty of the estimates propagates into the inferences which are drawn from the data. This approach allows one to make relatively strong claims about differences in the instantaneous risk and time preferences of addicts and non-addicts.

Another issue which is typically ignored, or assumed away, in the analysis of experimental data in psychology and addiction studies is the presence of "inconsistent" choices. For example, a person's pattern of choices on a discounting task may not be deterministically consistent with a particular discounting parameter. In this case, researchers typically calculate the discounting parameter which is consistent with the highest proportion of the subject's choices and use this discounting parameter thereafter; if multiple discounting parameters yield the same level of consistency, an average of these values is used instead.

The problem with this approach is that the uncertainty surrounding a particular parameter estimate is ignored when the estimate is treated as a datum which is measured without error. This thesis adopts the view that people's choice behaviour in experiments is inherently stochastic and one should be mindful of this fact and model it statistically. Consequently, the statistical models in this thesis formally incorporate the possibility that subjects make mistakes in experiments by adopting a behavioural error specification which is estimated jointly with the parameters of instantaneous risk and time preference models. This behavioural error specification becomes part of the structural theory being tested and applied, rather than being included as a statistical "after thought."

Determining which theories of choice under risk and over time best characterise the choices of people in experiments has important implications for our understanding of, and policy response to, addiction. For example, time inconsistency is an important feature of many economic models of addiction. Under the assumption of an additively-separable intertemporal utility function, only the exponential discounting model is time consistent whereas hyperbolic, quasi-hyperbolic, and Weibull discounting functions may yield time inconsistency. When agents are time inconsistent they can benefit from interventions which lock in their current preferences if they are apt to change in future; time-consistent agents do not benefit from these interventions. Thus, it is important to know whether people discount exponentially or, say, hyperbolically because this sharpens our understanding of, and influences our policy response to, addiction.

Unfortunately what few tests have been conducted in the addiction literature comparing the various theories of choice under risk and over time are not valid statistically. And these invalid statistical tests have been used to justify the near-universal adoption of the hyperbolic discounting and probability discounting models in experimental studies of addiction. This thesis employs statistically-valid Clarke (2007) non-nested model selection tests to adjudicate between the various models of risk and time preference. Importantly, Chapter 3 shows that the exponential discounting model better characterises the data than the hyperbolic discounting model. In addition, Chapter 4 shows that the probability weighting function (PWF) popularised by Tversky and Kahneman (1992) and the PWF developed by Prelec (1998) better characterise the data than the PWF implicit in the probability discounting model.

Clearly these findings are specific to the datasets used in these chapters, but they should give researchers pause when they are inclined to assume that one model of choice under risk or over time best characterises the choices of all subjects under all circumstances. Consequently, this thesis advocates the estimation of several theories of instantaneous risk and time preferences so that the ability of these theories to explain the data can be tested formally and so that one can check the sensitivity and robustness of the results across different specifications.

A related issue is whether the assumption that only one theory of choice under risk or over time can adequately characterise all of the choices in a dataset is justified. A recent development in the estimation of instantaneous risk and time preference models is to employ a mixture of different specifications. When one separately estimates particular models of choice under risk or over time one is implicitly assuming that only one DGP characterises all of the choices in the data. But what if one theory better explains some choices in a dataset while another theory better explains other choices in that same dataset? In this case, forcing all the data to fit one model will yield biased estimates of the underlying parameters. Mixture models allow one to combine different theories of choice under risk or over time and estimate the weight accorded to the different theories in the data. These models therefore directly incorporate the potential for multiple DGPs and provide an alternative measure of how much support each theory finds in the data. Mixture models are used extensively

in the three empirical chapters of this thesis because they formally recognise the heterogeneity in responses to instantaneous risk and time preference tasks.

With this background to the theoretical, methodological and statistical approach advocated in this thesis, some specific details on the three empirical chapters can be provided. Chapter 3 focuses solely on instantaneous time preferences and takes a step back from the ideal world in which theory, experimental design, and econometrics coalesce perfectly. The fundamental question which this chapter addresses is: if you have a dataset of discounting choices, which was generated by an experiment that does *not* satisfy Smith's (1982) precepts, and which does *not* include information that makes it possible to jointly estimate utility function curvature and discounting behaviour, what is the appropriate way to proceed statistically? In other words, this chapter takes the data that are typically collected in psychology experiments of time preferences and addiction, and provides an approach to data analysis which is statistically coherent.

If one delves into the literature on instantaneous risk preferences and addiction, one inevitably encounters the probability discounting model of Rachlin and colleagues (Rachlin, Logue, Gibbon and Frankel (1986), Rachlin, Castrogiovanni and Cross (1987), Rachlin, Raineri and Cross (1991)) which attempts to tie probabilistic choice to a temporal framework by reinterpreting the probability of a reward as the delay to, or rate of reinforcement of, that reward. The probability discounting model has become a popular framework in psychology and addiction studies for eliciting and analysing people's atemporal attitudes to risk. Despite the model's popularity in these fields it is virtually unheard of in economics. Consequently, Chapter 4 traces the development of the model and reinterprets it in language familiar to economists, statisticians, and decision theorists. The chapter then conducts an empirical analysis of the efficacy of the model's PWF in comparison to other PWFs commonly used in empirical studies of instantaneous risk preferences.

The complementarity between theory, experimental design, and econometrics is exemplified in Chapter 5 of this thesis. The chapter discusses a set of incentive-compatible experiments which were designed to elicit the instantaneous risk and time preferences of smokers and non-smokers, and to jointly estimate utility function

curvature and discounting behaviour. The experimental design lends itself to the estimation of different models of choice under risk and over time so as to rigorously explore the relationship, if any, between smoking status and instantaneous risk and time preferences.

In sum, this thesis discusses the experimental elicitation and estimation of preferences which may be relevant to our understanding of addiction. It stresses the importance of the theory, experimental design, and econometric trinity which is necessary to draw accurate and reliable inferences about potential differences between addicts and non-addicts. It also can serve as a constructive guide to others who aim to conduct research into the behavioural correlates of addiction.

REFERENCES

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583-618.

CLARKE, K. A. (2007): "A Simple Distribution-Free Test for Nonnested Model Selection," *Political Analysis*, 15, 347-363.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2014): "Theory, Experimental Design and Econometrics Are Complementary (and So Are Lab and Field Experiments)." in *The Methods of Modern Experimental Economics*, ed. by G. Frechette, and A. Schotter. New York, NY: Oxford University Press.

HEYMAN, G. M. (2009): *Addiction: A Disorder of Choice*. Cambridge, MA: Harvard University Press.

PRELEC, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497-527.

RACHLIN, H., A. CASTROGIOVANNI, AND D. CROSS (1987): "Probability and Delay in Commitment," *Journal of the Experimental Analysis of Behavior*, 48, 347-53.

RACHLIN, H., A. W. LOGUE, J. GIBBON, AND M. FRANKEL (1986): "Cognition and Behavior in Studies of Choice," *Psychological Review*, 93, 33-45.

RACHLIN, H., A. RAINERI, AND D. CROSS (1991): "Subjective Probability and Delay," *Journal of the Experimental Analysis of Behavior*, 55, 233-44.

ROSS, D., C. SHARP, R. VUCHINICH, AND D. SPURRETT (2008): *Midbrain Mutiny: The Behavioural Economics and Neuroeconomics of Disordered Gambling*. Cambridge, MA: MIT Press.

SMITH, V. L. (1982): "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, 923-955.

TVERSKY, A., AND D. KAHNEMAN (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297-323.

2. ECONOMIC MODELS OF ADDICTION

I. INTRODUCTION

Economic models of addiction have a long lineage and in many ways addiction is an ideal puzzle for economic theory. Why do most addicts expend resources to acquire their targets of addiction while maintaining concurrent commitments, involving real costs, to personal policies and programmes aimed at reducing or limiting their consumption of these goods? Furthermore, why is the typical course of addiction characterised by repeated unsuccessful attempts to quit prior to final abstention? From the standpoint of standard consumer theory in economics these patterns of behaviour are difficult to rationalise.

This thesis is primarily concerned with methodological and statistical issues at the boundary between economics and psychology in the context of addiction. Research into the behavioural aspects of addiction has been dominated by psychologists but I will argue that there should be a greater reliance on the tools from experimental economics and econometrics so as to draw accurate inferences about potential differences between addicts and non-addicts. In this chapter I discuss addiction generally and then focus on economic models of this phenomenon to highlight the role that instantaneous risk and time preferences[1], intertemporal risk preferences[2], and subjective beliefs play in the initiation, maintenance, and resolution of substance dependence[3].

[1] Instantaneous risk and time preferences are important components of economic models as they define, respectively, atemporal attitudes toward risk and uncertainty, and the valuation of goods which are available at different points in time. Formally, instantaneous risk preferences are defined over lotteries, where a lottery is a probability distribution over outcomes. Under expected utility theory, the concavity or convexity of a utility function determines attitudes to risk. Time preferences are defined over time-dated utility flows and are captured by a discounting function, which implies discounting behaviour. Although time preferences imply discounting behaviour, these terms (i.e., "time preferences" and "discounting behaviour") will be used interchangeably, and context will clarify the distinction.

[2] Intertemporal risk preferences refer to preferences over intertemporal lotteries, the outcomes of which may be serially correlated. Intertemporal risk preferences will be discussed in greater detail later in this chapter.

[3] The terms "dependence" and "addiction" refer to the most severe form of problematic substance use and will be used interchangeably. This is in keeping with the history of diagnostic classification systems for addiction, which I discuss later.

The chapter proceeds as follows. I provide a brief overview of the history of substance use and abuse, discuss current trends in drug consumption and dependence, and highlight the costs that addiction imposes on society. I discuss the way in which addiction has been diagnosed over time and then chronicle the history of the disease model of addiction, which is the primary understanding of substance dependence among the medical fraternity, the public, and the media. I provide a brief outline of Heyman's (2009) critique of the disease interpretation and then discuss the reconciliation provided by Ross, Sharp, Vuchinich and Spurrett (2008) (RSVS) which resolves the tension between these apparently contradictory views. Finally, I review a number of economic models which are relevant to the study of addiction and emphasise the roles that instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs play in substance dependence.

II. BACKGROUND

Substance Use and Abuse: A Brief History and Current Trends

Psychoactive substances have been used by different cultures for millennia. Wine vessels, dating back to late prehistorical and early historical times, have been found in shipwrecks in the Mediterranean and burial sites in China (Westermeyer (2005)). Ancient Greek archaeological artefacts from the Bronze Age (1600 – 1200 BC) include representations of the poppy seed capsule and poppy seed caches have been found in a prehistoric site in northern Turkey (Heyman (2009)).

Anthropological studies of preliterate societies suggest that psychoactive substance use was prevalent in a variety of different cultures. Tribes in North and South America cultivated numerous stimulant (e.g., coca leaf, tobacco leaf, and coffee bean) and hallucinogenic (e.g., peyote) drugs while different ethnic groups in Africa and the Middle East produced qat, a stimulant, and cannabis. Societies on the Eurasian land mass and Africa extracted alcohol from a number of sources including honey, grains, fruits, and milk, and the stimulant betel nut was used from Oceania to South Asia (Westermeyer (2005)).

Early Egyptian and Chinese writings make reference to alcohol, opium, and other substances whereas Mayan, Aztec, and Incan statues and glyphs depict rituals in which drugs were consumed. Alcohol became an integral part of the Jewish Passover festival and the Roman Catholic Mass while opium was used in certain Hindu marriages (Westermeyer (2005)). The early history of psychoactive substances suggests that they were used medicinally, in specific ceremonies or for particular purposes, e.g., manual labour or long hunts, and possibilities for abuse were severely circumscribed (Westermeyer (2005), United Nations Office on Drugs and Crime (UNODC) (2012)).

Over the last few centuries, economic and technological changes have had a marked impact on the cost and distribution of drugs, usage patterns, and methods of administration (Westermeyer (1987)). Industrialisation and international trade led to large increases in substance production and distribution, and rising incomes fostered recreational drug use. New modes of administration like snorting, smoking, and injecting, delivered drug doses to the brain far more rapidly and efficiently than ingestion, and as the rate and efficiency of delivery increases so too does the drug's concentration at active sites in the brain. New methods for purifying and modifying plant compounds along with laboratory synthesis has led to the production of drugs that are more potent and more easily smuggled (Westermeyer (2005)). Some of these changes coincided with, and were responsible for, the first large-scale instances of substance abuse.

The first two large drug abuse epidemics occurred at roughly the same time but in different regions and with respect to different substances. The English gin epidemic or gin plague began in the late 1600s and lasted several decades. The proximate causes were: the importation of vast quantities of rum and gin which were sold very cheaply; a lack of traditions and social sanctions to deter over-consumption of these beverages; and the start of the Industrial Revolution which led to mass urbanisation and the erosion of traditional rural values. The gin plague was eventually "cured" through a combination of import taxes on alcohol, anti-alcohol propaganda, and the formation of Protestant sects which were opposed to the consumption of alcohol (Westermeyer (2005)).

The Asian opium epidemic affected numerous countries in East and Southeast Asia and lasted for centuries. Opium had been used in Asia for millennia with few adverse consequences until Europeans introduced tobacco smoking to the region. Asians began to add opium to their tobacco pipes and this fast became a common practice. Smoking opium delivers the active ingredient, morphine, to the brain far more quickly and efficiently than ingestion and, as discussed earlier, this magnifies its effect. A stronger pharmacological effect coupled with surplus wealth, available leisure time, and a genetic aversion to alcohol, precipitated an opium crisis (Heyman (2009)). According to a report prepared by the International Opium Commission (1909), approximately one quarter of the population of China used opium at the start of the 20th century. As Westermeyer (2005, p. 28) notes, some Asian countries (e.g., China, Japan, and South Korea) managed to arrest the epidemic in the 20th century, while others are still plagued by it (e.g., Laos, Afghanistan, and Myanmar).

Current patterns of substance use and abuse have been linked to the emergence of "youth culture" in the US in the 1960s, which spread rapidly to Western Europe and other parts of the world (UNODC (2012)). Cannabis and hallucinogenic drugs, like lysergic acid diethylamide (LSD), became anti-authority symbols and were used by increasing numbers of people to explore altered states of consciousness (Westermeyer (2005)). Heroin use also became more widespread during this period, and was particularly high among veterans of the Vietnam war. Cocaine use began to expand in the 1960s but was mostly confined to the richer segments of society until the invention of "crack," a cheaper form of cocaine that is typically smoked, led to an explosion in use in the early 1980s. Towards the end of this decade, cocaine use started to stabilise but this coincided with the start of the "rave" scene which became synonymous with the use of 3,4-methylenedioxymethamphetamine (MDMA or "ecstasy") and other amphetamine-type stimulants (UNODC (2012)).

The World Health Organisation (WHO) estimates that approximately 42 percent of the world's population, aged 15 years and older[4], consumed alcohol in 2004 and that 11.5 percent of these people (4.8 percent of the world's population) engaged in heavy episodic drinking on a weekly basis (WHO (2011)). With regard to past-month

[4] All of the estimates presented in this section relate to people aged 15 years and older.

prevalence, approximately 25 percent of the world's population smoked tobacco in 2006, whereas 19 percent smoked tobacco daily in 2009 (WHO (2010)). The United Nations Office on Drugs and Crime (UNODC) estimates that approximately 5 percent of the world's population used an illicit drug in 2010. The most widely-used illicit substance is cannabis, with approximately 119 – 224 million users worldwide. Amphetamine-type stimulants (including "ecstasy") have the second highest prevalence rate with between 14.3 and 52.5 million users. Opioids and cocaine round out the top four most used illicit substances, with approximately 39.6 – 55.5 million users in 2010 (UNODC (2012)).

Illicit drug use is typically initiated in the teens or early twenties and is largely a youth phenomenon in many countries (UNODC (2012)). Men in urban areas are particularly susceptible and there is a large gender gap in use. Prevalence rates for substance use peak among people aged 18 – 25, and then decline gradually until reaching very low levels in people older than 65. Approximately 27 million people, or 0.6 percent of the world's population, meet the criteria for substance use disorders; these criteria will be discussed in more detail later. In addition, in 2008, there were an estimated 16 million injecting drug users, of which approximately 19 percent were living with the human immunodeficiency virus (HIV), 46.7 percent were infected with hepatitis C, and 14.4 percent had hepatitis B. Less than 20 percent of people who require treatment for substance use disorders receive it but those seeking treatment are typically in their late 20s or early 30s. Drug-related deaths are also concentrated among the younger members of society, with an average age in the mid-30s (UNODC (2012)).

The Costs of Substance Use and Abuse

Estimating the costs to society of substance use and abuse is a difficult task, not least because usage tends to be underreported, many of the drugs are illegal, and the extent to which they independently affect crime, productivity, mortality, and morbidity is difficult to gauge. These issues notwithstanding, the UNODC and WHO, drawing on numerous sources, compile estimates of the impact of alcohol, tobacco, and illicit drugs on society.

Tobacco use has been identified as the leading cause of preventable death on earth (WHO (2013)).[5] Globally, approximately 12 percent of all deaths among people over 30 can be attributed to tobacco. In 2004 alone, about 5 million people died from the direct use of tobacco, while another 600,000 died as a result of second-hand smoke. In addition, among people over the age of 30, 5 percent of all deaths from communicable diseases, and 14 percent of all deaths from non-communicable diseases have been linked to tobacco (WHO (2012)). Finally, estimates from 2009 suggest that 56.9 million disability-adjusted life-years (DALYs) will be lost to the use of tobacco (WHO (2009)).

A WHO (2009) *Global Health Risks* report identified alcohol use – along with childhood underweight, unsafe sex, and unsafe water, sanitation and hygiene – as one of the leading global risks for burden of disease as measured by DALYs. Specifically, the report estimates that 69.4 million DALYs will be lost due to alcohol use. In addition, alcohol use, high blood pressure and the other factors mentioned previously (i.e., unsafe sex etc.), account for one quarter of all deaths in the world and one fifth of all DALYs. Alcohol use has also been identified as a causal factor in at least 60 major types of diseases and injuries and it leads to approximately 2.5 million deaths each year. This implies that 4% of all deaths globally are attributable to alcohol. In addition, alcohol is responsible for between 20 and 50 percent of all cases of cirrhosis of the liver, epilepsy, poisonings, road accidents, violence, and numerous cancers (WHO (2011)).

Illicit drug use leads to approximately 245,000 deaths each year worldwide and about half of these deaths are due to fatal overdoses (WHO (2009)). Some 4.5 million people are in treatment for drug-related problems globally but, as mentioned earlier, this is less than 20% of people with substance use disorders. Billions of dollars are spent on treatment each year but the UNODC estimates that, at current prevalence

[5] While the harm caused by tobacco use is undoubtedly huge, whether it is the leading cause of preventable death globally is questionable. For example, providing clean water and sanitation to people who live without them my be more cost-effective, and have a far greater impact, than trying to curb tobacco use, particularly in less developed countries, where the bulk of the global population resides. The provision of clean water and sanitation was one of the key themes of the World Bank's (1992) *World Development Report 1992: Development and the Environment* and numerous subsequent World Bank (1993, 2003, 2010) and WHO (2009) publications stress the importance of clean water and sanitation for reducing the global burden of disease.

rates, $200 - $250 billion would have been required to cover global costs associated with treatment for illicit drug use in 2010; these amounts dwarf current expenditure.

Productivity losses due to substance abuse are linked to a number of factors including labour force non-participation, unemployment, absenteeism, workplace accidents, and incarcerations. Deriving precise estimates of these losses is difficult, but a US study by the National Drug Intelligence Center (2011) estimated that $120 billion (0.9 percent of GDP) in productivity losses were due to the use of illicit drugs in 2007. A 2002 study in Canada, covering alcohol, tobacco, and illicit drugs, estimated that productivity losses totalling 4.7 billion Canadian dollars (0.4 percent of GDP) could be attributed to substance abuse (Rehm et al. (2006)). Finally, an Australian study, focussing on legal and illegal substances, calculated productivity losses at 2.1 billion Australian dollars (0.3 percent of GDP) in 2004/2005 (Collins and Lapsley (2008)).

Psychoactive substances affect crime through a number of different channels: the use of drugs can promote criminal acts by lowering inhibitions; crimes are often committed to finance drug-taking; competition between drug traffickers can lead to "turf wars;" and the sale of drugs is often closely linked to fraud and corruption. Studies linking drugs to crime have found that the costs can be substantial. In the UK, the cost of drug-related crime in England and Wales was estimated at £13.9 billion in 2003/2004, which amounted to almost 90 percent of all social and economic costs associated with drug abuse (Gordon, Tinsley, Godfrey and Parrott (2006)). Similarly, an Austrian study calculated the cost of drug-related crime at €2.6 billion in 2002 (UNODC (2012)).

Diagnosing Addiction

The two primary diagnostic classification systems for substance use disorders are *The Diagnostic and Statistical Manual of Mental Disorders* (DSM), published by the American Psychiatric Association (APA), and *The International Classification of Diseases* (ICD), published by the WHO. These two systems differed markedly when they were first published but a sustained standardisation effort by the APA, WHO, and scientists working in the area of mental disorders, has ensured that the diagnostic systems are now compatible (WHO (1994)). The latest version of the ICD (ICD-10)

was published in 1994, whereas the most recent version of the DSM (DSM-5) was published in 2013. Given the massive growth in addiction-related research over the last 20 years, I will focus on the DSM criteria in what follows and trace the history of the diagnostic classification of addiction.

The first edition of the DSM (APA (1952)) included categories for "Alcoholism" and "Drug addiction" under the more general category of "Sociopathic Personality Disturbances." The diagnosis of alcoholism was confined to cases "in which there is a well established addiction to alcohol without recognizable underlying disorder" (APA (1952, p. 39)). Presumably clinical judgement and experience was needed to determine whether a "well established addiction to alcohol" was present. Drug addiction, by contrast, was a symptom of an underlying personality disorder and was reserved for cases where "the individual is actually addicted" (APA (1952, p. 39)).

The second edition of the DSM (DSM-II) (APA (1968)) included categories for "Alcoholism" and "Drug dependence" under the more general category of "Personality Disorders and Certain Other Non-psychotic Mental Disorders." Alcoholism itself was split into three categories, representing greater levels of severity: episodic excessive drinking, habitual excessive drinking, and alcohol addiction. Alcohol addiction, the most severe form, was confined to people who were "dependent on" alcohol, although the term "dependent on" was not defined (APA (1968, p. 45)). The best evidence for a diagnosis of alcohol addiction was the presence of withdrawal symptoms, but failing that, heavy drinking for a period of three months was taken as presumptive evidence of dependence.

A diagnosis of drug dependence was reserved for people who were "addicted to or dependent on" drugs other than alcohol, tobacco, and caffeine, but like alcohol addiction, the key terms "addicted to" and "dependent on" were not defined. Diagnosis required evidence of habitual use and the manual recognised that withdrawal only occurs with certain substances (e.g., opioids) and not others (e.g., marijuana).

DSM-III (APA (1980)) marked a fundamental shift in the classification of mental disorders by incorporating specific criteria for diagnosis so as to improve reliability.

"Substance Use Disorders" were given their own category in the manual and "Substance abuse" was distinguished from "Substance dependence." Substance abuse was diagnosed when there was: 1) a pattern of pathological use which 2) lasted for at least one month and 3) led to impairment in social or occupational functioning. A detailed explanation was given for each criterion so as to aid diagnosis.

Substance dependence, a more severe form of substance abuse, was diagnosed when there was evidence of physiological dependence, specifically, tolerance or withdrawal. The only exceptions were alcohol and cannabis dependence where a pattern of pathological use and impairment in social or occupational functioning was also required. An interesting change in DSM-III, other than specific criteria for diagnosis, was the decision to drop the term "addiction" in favour of "dependence." This was done to reduce the stigma associated with the term "addiction" and in the hope that "dependence" would provide a more objective and precise definition (RSVS (2008)).

DSM-IV (APA (1994)) kept the abuse/dependence distinction but provided a far more detailed set of criteria to diagnose either mental disorder. Substance dependence was defined as, "… a cluster of cognitive, behavioral, and physiological symptoms indicating that the individual continues use of the substance despite significant substance-related problems. There is a pattern of repeated self-administration that usually results in tolerance, withdrawal, and compulsive drug-taking" (APA (1994, p. 176)).

A diagnosis of substance dependence required the presence of at least three of the following criteria within a 12-month period: 1) tolerance; 2) withdrawal; 3) the substance is taken in larger amounts or for a longer period than was intended; 4) there is a persistent desire or unsuccessful efforts to cut down or control substance use; 5) a large amount of time is spent in activities to obtain the substance, use the substance, or recover from its effects; 6) important social, occupational or recreational activities are given up or reduced due to substance use; and 7) there is continued use of the substance despite persistent or recurrent problems linked to use (APA (1994, p. 181)). Note that tolerance and withdrawal were no longer mandatory criteria for a diagnosis of substance dependence.

Substance abuse, by contrast, was defined as, "… a maladaptive pattern of substance use manifested by recurrent and significant adverse consequences related to the repeated use of substances" (APA (1994, p. 182)). Only one of the following criteria needed to be present within a 12-month period for a diagnosis of substance abuse: 1) failure to fulfil major role obligations; 2) recurrent substance use in situations in which it is physically hazardous; 3) recurrent substance-related legal problems; and 4) continued substance use despite problems caused or exacerbated by the substance (APA (1994, p. 183)).

DSM-5 (APA (2013)) created a new overarching category called "Substance-Related and Addictive Disorders" which incorporates a behavioural addiction, "Gambling Disorder," in addition to substance use disorders. DSM-5 also dropped the distinction between substance abuse and dependence in favour of the single term "substance use disorder" which has three levels of severity: mild, moderate, and severe. The definition of a substance use disorder is identical to the definition of substance dependence in DSM-IV, "… a cluster of cognitive, behavioral, and physiological symptoms indicating that the individual continues using the substance despite significant substance-related problems" (APA (2013, p. 483)).[6] A new set of criteria, drawing heavily on DSM-IV, and new thresholds for diagnosis were also adopted. Table I lists the criteria for a substance use disorder in DSM-5.

The criteria in Table I are broken down into four main categories: impaired control, social impairment, risky use, and pharmacological criteria (APA (2013, p. 483)). The list of symptoms is a mixture of the criteria for substance abuse and dependence in DSM-IV, other than criterion (4). Criterion (4), craving, is an important addition to the list because it was mentioned in DSM-III and DSM-IV but was not included as a symptom even though it is "… likely to be experienced by most (if not all) individuals with Substance Dependence" (APA (1994, p.176)).

[6] The Oxford English Dictionary (OED), Third Edition (2010), defines addiction as, "Immoderate or compulsive consumption of a drug or other substance; *spec*. a condition characterized by regular or poorly controlled use of a psychoactive substance despite adverse physical, psychological, or social consequences, often with the development of physiological tolerance and withdrawal symptoms; an instance of this." Thus, the definition of addiction in the OED and the definition of substance dependence in the DSM are very similar.

TABLE I
CRITERIA FOR SUBSTANCE USE DISORDERS IN DSM-5

A problematic pattern of substance use leading to clinically significant impairment or distress, as manifested by at least two of the following, occurring within a 12-month period:

Impaired control

(1) The substance is often taken in larger amounts or over a longer period of time than was intended.

(2) There is a persistent desire or unsuccessful efforts to cut down or control substance use.

(3) A great deal of time is spent in activities necessary to obtain the substance, use the substance, or recover from its effects.

(4) Craving, or a strong desire or urge to use the substance.

Social impairment

(5) Recurrent substance use resulting in failure to fulfil major role obligations at work, school, or home.

(6) Continued substance use despite having persistent or recurrent social or interpersonal problems caused or exacerbated by the effects of the substance.

(7) Important social, occupational, or recreational activities are given up or reduced because of substance use.

Risky use

(8) Recurrent substance use in situations in which it is physically hazardous

(9) Substance use is continued despite knowledge of having a persistent or recurrent physical or psychological problem that is likely to have been caused or exacerbated by the substance.

Pharmacological criteria

(10) Tolerance, as defined by either of the following:
  (a) A need for markedly increased amounts of the substance to achieve intoxication or desired effect.
  (b) A markedly diminished effect with continued use of the same amount of the substance.

(11) Withdrawal, as manifested by either of the following:
  (a) a substance-specific problematic behavioral change, with physiological and cognitive concomitants, that is due to the cessation of, or reduction in, heavy and prolonged substance use.[1]
  (b) the same (or a closely related) substance is taken to relieve or avoid withdrawal symptoms.

Source: DSM-5 (APA (2013, p 483-484)).
[1]Criterion (a) is specific to each substance in the DSM because the nature of withdrawal tends to differ across substances; I have included criterion (a) in its most general form (APA (2013, p. 486)).

An important change in DSM-5 is that only two criteria need to be present for a diagnosis of substance use disorder; in DSM-IV, one criterion needed to be present for substance abuse while three had to be present for substance dependence. As mentioned previously, DSM-5 also includes a severity specifier: 2 – 3 criteria

represent mild substance use disorder, 4 – 5 criteria are indicative of moderate substance use disorder, whereas 6 or more symptoms imply severe substance use disorder. DSM-5 makes reference to the term "addiction" in the title of the mental disorder category (i.e., "Substance-Related and Addictive Disorders") but then explains why the term is not used anywhere else in the manual, "… because of its uncertain definition and its potentially negative connotation" (APA (2013, p.485)).

The evolution of the diagnostic classification of addiction raises some important issues. First, iterations of the DSM show that diagnostic criteria and classification systems do not solely serve scientific ends but that they develop in concert with wider political, social, and economic forces (RSVS (2008)). For example, a diagnosis of substance dependence in DSM-III required evidence of physiological dependence, regardless of other forms of harm that a person may inflict on herself[7] through substance use, because this was regarded as a clear sign of a disease state and the disease model of addiction had gained general acceptance.

Second, the terms "addiction" and "dependence" have changed from being descriptors of behaviour into explanations of that behaviour. RSVS (2008, p. 4) state this point clearly, "This appears to have happened over the years with the term 'addiction,' which has gone from being a metaphorical description of devotion to a particular activity, to being an hypothesized internal state, the existence of which is revealed by that devotion, and, finally, to being the cause of that devotion." Akers (1991, p. 779) makes a similar point and then argues, "However precisely it is defined, addiction is a label, a term applied to behavior. It cannot, itself, provide an explanation for that behaviour."

Third, West (2006) recognises that DSM diagnostic criteria leave much room for interpretation. For example, how strong do cravings need to be to be classified as such? Similarly, how painful or pronounced does withdrawal have to be to tick off this symptom? Either we need to rely on clinical judgement to adjudicate, recognising that clinicians will have different views, or let the individual decide when she fills out a diagnostic questionnaire.

[7] In this thesis I adopt the convention of using the female gender in even-numbered chapters and the male gender in odd-numbered chapters.

Finally, given that only two criteria need to be present for a diagnosis of substance use disorder, we can have up to 5 "addicts" with non-overlapping symptoms. In addition, every criterion is weighted equally in the diagnosis of addiction even though certain criteria (e.g., withdrawal) tend to be more indicative than others (e.g., interpersonal problems).

These issues notwithstanding, research into addiction has made steady progress over the last 50 years by using the criteria set out in the DSM. Clearly no classification system is perfect but addiction research has improved, and will continue to improve, as DSM criteria and the definition of substance dependence evolve.

The Disease Model of Addiction

Warner (1994) traces the history of the disease model of addiction back to the early 17th century. As alcoholism or "habitual drunkenness," as it was referred to at the time, became increasingly widespread, members of the clergy began to focus on its causes, course, and consequences. They argued that alcoholism was caused by the consumption of spirituous alcohol, as opposed to fermented liquors[8], which subsequently led to a loss of control over the use of alcohol. This loss of control implied that habitual drunkenness was a disease, the only cure for which was complete abstinence.

For example, in 1609, John Downame referred to people "addicted to drunkennese" who turned "delight into necessitie" (cited in Warner (1994, p. 687)). In 1619, Robert Harris described alcoholism as a "dropsilike disease" (cited in Warner (1994, p. 688)). In 1677, Edward Bury argued that habitual drunkenness was a disease "so epidemical" that "all the Physicians in England know not how to set a stop to it" (cited in Warner (1994, p. 688)). As Heyman (2009, p. 98) notes, members of the clergy reasoned that they were dealing with a disease because their parishioners were

[8] "Spirituous alcohol" refers to distilled spirits like whiskey, rum, gin and brandy, whereas "fermented liquors" refer to types of alcohol - like beer, wine and cider - that are produced through a process of fermentation.

still pious but had lost control over the use of alcohol and continued drinking despite alcohol-related problems.

This view, that addiction is a disease, won few adherents in the 1600s but, a century or so later, was embraced by Benjamin Rush, who is often credited as the father of American psychiatry (e.g., Westermeyer (2005, p. 23)). Rush believed that consumption of alcohol over a period of time could turn a voluntary drinker into an involuntary one. In other words, drinking became a compulsion or necessity merely as a function of prior consumption. Rush's work coupled with a marked increase in alcohol consumption in the US post-1780, and the subsequent rise in opium and morphine abuse, brought addiction to the attention of the wider medical community (Weiner and White (2007)).

In the 1850s and 1860s a number of asylums and sanatoria were established in the US and Europe to treat addiction, and greater awareness of addiction as a problem which required medical treatment culminated in the establishment of the American Association for the Cure of Inebriety (AACI) in 1870. The AACI was based on four founding principles: 1) addiction is a disease; 2) like other diseases, addiction is curable; 3) its primary cause is a "constitutional susceptibility" to addiction; 4) and this susceptibility can be either acquired or inherited (Weiner and White (2007)).

The AACI published the *Quarterly Journal of Inebriety* in 1876, which was the first scientific journal focussing specifically on addiction, and the first to explicitly embrace the disease interpretation of this disorder. Weiner and White (2007) trace the history of this journal and the terms that were used to describe dependence on psychoactive substances. Although "inebriety" remained the professional term of choice throughout the journal's history, "alcoholism" and "addiction" took on greater prominence over the years.

Despite the growing acceptance of the disease model of addiction among the set of physicians studying substance dependence, lay people and religious organisations vehemently opposed this conception (Weiner and White (2007)). Levine (1978) explains that the public were of the opinion that addicts *chose* to abuse their targets of addiction, rather than being *compelled* to do so, as the disease interpretation implies.

The American Medical Temperance Movement (AMTA) was formed in 1891 and it published the *Bulletin of the American Temperance Association* to expound this alternative view. AMTA advocated the prohibition of alcohol and the legal regulation of opiates and cocaine, policy suggestions which were ultimately codified in the Harrison Narcotics Tax Act of 1914 (Weiner and White (2007)).

The Harrison Act did not explicitly cover recreational drug use and addiction, but according to Heyman (2009, p. 9) it was enforced as though its goal was to eradicate substance use and abuse. This shift in policy from treatment to prosecution of addiction led to the demise of the American Medical Society for the Study of Alcohol and Other Narcotics (AMSSAON), the successor to the AACI, and all but a handful of the sanatoria and asylums which had been established in the 19th century (Weiner and White (2007)).

However, the failure of prohibition in the US and the growing awareness of substance abuse and addiction following World War II, led to renewed interest in the etiology, treatment, and prevention of addiction, and the establishment of alcohol and drug abuse divisions under the National Institute of Mental Health (NIMH). In 1952, the first edition of *The Diagnostic and Statistical Manual of Mental Disorders* (DSM) was published and, as discussed earlier, it included addiction to alcohol and drugs as recognised mental disorders. The hope at the time was that the scientific recognition of addiction as a mental disorder would lead to better research into, and treatment and prevention of, addiction.

However, post-World War II increases in alcohol abuse, and the cannabis, LSD, and opiate epidemics of the 1960s, made it clear that the alcohol and drug divisions of the NIMH were not responding adequately to substance use problems in the US. This prompted the formation of the National Institute on Alcohol Abuse and Alcoholism (NIAAA) in 1970 and the National Institute on Drug Abuse (NIDA) in 1974, both of which fell under the NIMH. The formation of these institutes provided the funding and impetus for large-scale research into, and treatment of, addiction (Westermeyer (2005)). They also became two of the most important proponents of the disease model of addiction, which Heyman (2009, p. 17) argues is the primary understanding of addiction among the public, the medical fraternity, and the media.

That NIDA and the NIAAA still embrace the disease model of addiction is evident in their press releases, research reports, and publicly disseminated information. For example, NIDA's webpage "DrugFacts: Understanding Drug Abuse and Addiction" provides the following definition of addiction:

> "Addiction is a chronic, often relapsing brain disease that causes compulsive drug seeking and use, despite harmful consequences to the addicted individual and to those around him or her. Although the initial decision to take drugs is voluntary for most people, the brain changes that occur over time challenge an addicted person's self control and hamper his or her ability to resist intense impulses to take drugs" (NIDA (2012)).

Similarly, the NIAAA's webpage on "Alcohol Use Disorders" states that, "Like many other diseases, alcoholism is typically considered chronic, meaning that it lasts a person's lifetime" (NIAAA (2014)). In the next section I discuss Heyman's (2009) critique of the disease model of addiction and outline his argument for understanding addiction as a disorder of choice. I then present the reconciliation provided by RSVS which resolves the traditional tension between these hypotheses.

The Disease Model of Addiction: A Critique and Reconciliation

Heyman (2009, p. 91) argues that proponents of the disease interpretation of addiction typically rely on at least one of the following justifications for their adoption of this model: 1) addiction has a biological basis; 2) psychoactive substances can convert a voluntary drug user into an involuntary one; and 3) the disease model is the humane interpretation and leads to better treatment for addicts. Heyman critiques each of these justifications and I will discuss them in turn.

Numerous research studies have found that genes influence substance dependence (for recent reviews consult Li (2003), Agrawal and Lynskey (2008), Li and Burmeister (2009)) and this has been used as evidence in favour of the disease model of addiction. The logic is as follows: people do not choose their genes, if genes

influence behaviour then that behaviour cannot be voluntary, genes influence addiction, thus addiction is not voluntary, it is a disease.

While genes certainly influence addiction, this does not imply that addiction is a disease because the correlation between genes and addiction is far from perfect. A recent review of twin and adoption studies (Agrawal and Lynskey (2008)) finds that genetic influences on addiction range from 0.3 – 0.7. In other words, a genetic predisposition to addiction raises the likelihood of addiction without making it certain. In addition, the fact that genes influence behaviour does not imply that behaviour is involuntary; genes affect many behaviours which are clearly voluntary. For example, Heyman (2009, p. 94) cites numerous studies showing that a range of attitudes and beliefs (e.g., support for the death penalty and religious beliefs) are heritable. In other words, just because something is heritable does not mean it is involuntary: attitudes, beliefs, and behaviours which are voluntary have genetic as well as cultural, social, and psychological antecedents. Thus, the link between genes and substance dependence does not imply that addiction is a disease.

Another line of reasoning which forms part of the biological justification for the disease interpretation of addiction is that drug use leads to neuroadaptation. As drug use changes brain structure and function, addiction must be a disease, specifically a brain disease, or so the argument goes (see, for example, Leshner (1997), Hanson, Leshner and Tai (2002)). The major flaw with this hypothesis is that voluntary behaviours lead to neuroadaptation too. In other words, the fact that drug use leads to neuroadaptation does not mean that addiction is a disease. Changes in voluntary behaviour lead to changes in the brain and it is this plasticity which makes voluntary behaviour possible (Heyman (2009, p. 95-96)). Thus, using neuroadaptation as the basis for the disease model of addiction runs the risk of classifying everything which leads to changes in the brain as a disease.

In sum, genes are related to addiction and addiction causes changes in the brain, but genes also influence a host of voluntary behaviours and these voluntary behaviours lead to neuroadaptation. Thus, the biological basis of addiction, in and of itself, does not imply that addiction is a disease. For addiction to be classified in this way, drugs

must lead to a loss of control over substance use and thereby transform a voluntary drug user into an involuntary one. It is to this point that I now turn.

Heyman (2009, p. 97-98) recognises that to determine whether drug use turns a voluntary user into an involuntary one, an impartial understanding of voluntary and involuntary behaviour is crucial. The history of the disease model of addiction shows that the disease concept is closely linked to the idea that addicts are *compelled* to consume drugs of abuse despite the deleterious effects this has on their welfare. In other words, the disease model of addiction is tied to the idea that people do not *voluntarily* engage in self-destructive acts. Under this conception, voluntary behaviour is optimal behaviour, in the sense of maximising benefits and minimising costs, and as addiction is not optimal it cannot be voluntary.

Heyman (2009, p. 104) provides an alternative definition of voluntary behaviour, "… the degree to which an activity is voluntary is the degree to which it systematically varies as a function of its consequences, and the degree to which it is feasible to apply such consequences." This definition distinguishes blinking and sneezing, which are involuntary acts because they do not vary as a function of their consequences, from winking and spitting, which do. The second component of this definition accommodates cases where it is not feasible to adjust consequences so as to alter a voluntary act; in this case, voluntary behaviour is functionally involuntary. For example, eating is a voluntary act but, as there are no substitutes for calories, it is not feasible to apply consequences to completely deter eating. In the context of addiction, drug use may become involuntary if it is not sensitive to its consequences or if it is not legitimate or practical to apply these consequences.

Studies of addicts have found that there is a clear inverse relationship between price and consumption of drugs, and between access to alternative, substitutable activities and substance use (Vuchinich and Heather (2003) provide an overview of the literature on this topic). In other words, addiction is responsive to changes in its consequences. Furthermore, according to a set of addicts' self-reports, many decided to quit using drugs in response to fears of arrest, financial problems, and the desire to be good parents (Heyman (2009, p. 44-65)).

Heyman draws attention to the four largest epidemiological studies of addiction ever conducted in the US: the Epidemiologic Catchment Area (ECA) study (Anthony and Helzer (1991)); the National Comorbidity Survey (NCS) (Warner et al. (1995)); the National Comorbidity Survey Replication (Kessler et al. (2005a, 2005b)); and the National Epidemiologic Survey on Alcohol and Related Conditions (NESARC) (Stinson et al. (2005)). He finds that the vast majority of addicts quit by their early thirties and do so without seeking formal treatment, which echoes the findings of the 2012 *World Drug Report* (UNODC (2012)). In other words, as the demands of adulthood take on greater significance, many addicts successfully manage to quit their targets of addiction, which clearly implies that addiction varies as a function of its consequences and thereby involves voluntary behaviour.

Heyman (2009, p. 86-87, 105-108) also focuses on contingency management (CM) as a treatment for addiction. CM aims to modify behaviour by incentivising and reinforcing changes in behaviour through the provision of vouchers, modest financial rewards, and prizes. If addiction is voluntary and responds to changes in its consequences then CM may be an effective form of treatment. Heyman finds that a large number of studies successfully employed CM techniques to promote abstinence and prevent relapse.

This finding suggests that the disease interpretation of addiction may not yield the best treatment for addicts, which is a common justification for the model. If addiction is a chronic, relapsing disorder, characterised by involuntary drug use, then CM treatment should have little effect. However, a recent meta-analysis of CM techniques concluded that it is one of the most effective approaches for the treatment of addiction (Prendergast et al. (2006)).

The preceding discussion is a brief outline of Heyman's thesis and his critique of the three primary justifications for the disease model of addiction. His argument casts doubt on the validity of the disease model and it suggests that addiction involves voluntary behaviour insofar as it varies systematically as a function of its costs and benefits.

RSVS argue that these two apparently contradictory hypotheses (i.e., the disease model and Heyman's voluntary choice model) can be reconciled by focussing on recent developments in behavioural science and neuroscience and by adopting the appropriate levels or scales of scientific analysis. Addiction manifests itself and can be understood at both molar (i.e., behavioural) and molecular levels; these levels put constraints on each other but neither can be reduced to the other. Heyman provides a molar account of addiction whereas the disease model is inherently molecular in that neuroadaptation produces a disease state typified by compulsive drug-taking.

RSVS adopt Ainslie's (1992, 2001) "picoeonomic" framework to provide a molar scale account of addiction. Ainslie argues that Samuelson's (1937) discounted utility model – which assumes an additively-separable intertemporal utility function and exponential (E) discounting, and thereby implies time-consistent preferences – is an inaccurate description of how people and nonhuman animals discount delayed rewards. Numerous studies in psychology (see Green and Myerson (2004) for a review) suggest that people discount rewards hyperbolically[9] which, in contrast to E discounting with an additively-separable intertemporal utility function, can lead to preference reversals (viz., time-inconsistent choices): preference for a larger, later (LL) reward switches to a smaller, sooner (SS) reward temporarily when the SS reward is imminent.

Hyperbolic (H) discounting provides an explanation for the pattern of behaviour described at the start of this chapter: addicts consume drugs while taking active, and costly, steps to prevent or limit their consumption of these goods. It also explains the common cycle of quit attempts and relapse observed in addicts. For example, a smoker who wants to quit may take active steps to stop smoking so as to enjoy the long-term health benefits which result (i.e., the smoker has a preference for the LL reward) but then cave in when offered a cigarette by a friend (i.e., switch to the SS reward when it is imminent). Thus, H discounting provides an explicit rationale for the time-inconsistency which we observe in addicts. But if all people discount

[9] As this thesis focuses on statistical and experimental methodology at the boundary of economics and psychology, I will rigorously interrogate the claim that hyperbolic discounting more accurately characterises time preferences in later chapters.

hyperbolically then what prevents the vast majority of people who use drugs from becoming addicted?

To explain this apparent tendency to make time-inconsistent choices while accounting for the fact that people successfully avoid the temptation of SS rewards most of the time, Ainslie models molar choice behaviour as the Nash equilibrium of a bargaining game among subpersonal interests (hence "picoeconomics"). These subpersonal interests are a succession of selves who bargain with one another across time and ultimately determine whether SS or LL rewards are consumed.

Ainslie lists four mechanisms which are used by long-term interests to supplant short-term interests in these intertemporal bargaining games. First, external commitment devices, like the drug disulfiram which makes consumption of alcohol extremely unpleasant, can commit the agent to a course of action when there is a temptation to renege. Second, long-term interests can control the attention of the agent so as to prevent the processing of information about the availability of a SS reward; this is a form of internal, as opposed to external, commitment. Third, long-range interests can inhibit the agent's emotional response to "visceral rewards" (Loewenstein (1996, 1999)) and thereby make consumption of the visceral SS reward less likely.

Finally, and most importantly, "personal rules" can be used to *bundle* together whole series of LL rewards so as to bolster the agent's self-control and subvert short-term interests which favour immediate consumption. Note that preference reversal is more likely when an agent chooses between a discrete SS-LL reward pair as opposed to two series of SS and LL rewards. As Ainslie (2005, p. 640) explains, H discount curves level off at longer delays which means that when a set of LL rewards are bundled together in a series, the aggregate value of the LL series increases, relative to the SS series, as both series lengthen. With a long enough series, the total value of the LL series may not only exceed the value of the associated SS series but also the value of a lone SS reward which is available immediately. If so, then the agent may choose to forgo both the single SS reward and the series of SS rewards in favour of the series of LL rewards and thereby achieve self-control.

This framing of choice as between series of rewards rather than individual rewards is more likely if the agent views her current choice as predictive of future choices. If the current choice is a test case for future choices, this puts the whole bundle of LL rewards at stake in the current choice. What this means is that if the agent chooses the SS reward, she not only loses out on the LL reward but also the expectation of choosing LL rewards in future. Ainslie argues that people often adopt personal rules to promote the bundling of rewards. These rules specify the choice to be made within particular classes or categories of decision problems and they are strengthened by adherence and weakened by transgression. Ainslie (1992, 2001) provides detailed discussions of the conscious and unconscious formation of personal rules, their likely efficacy, and the role they play in circumscribing and promoting patterns of behaviour. In the context of addiction, addicts are understood as people who have lost the ability to make and "enforce" personal rules which prevent or curtail their consumption (see Monterosso and Ainslie (2009) for more details). Heyman's view that addiction is a disorder of choice fits squarely within this framework.

At a molecular level, RSVS argue that a pattern of drug consumption affects the way in which the brain's dopamine reward system learns. The dopaminergic circuit of the brain, which projects from the ventral tegmental area through the ventral striatum and into the orbitofrontal cortex, releases dopamine spikes when rewards turn out better than expected, and it learns to predict the cues that precede these spikes. The purpose of this learning process is to prime animals to exploit stochastic consumption opportunities and, as such, the dopamine system integrates valuation, attention, and motor preparation (Ross (2012)).

Addictive substances influence the brain in different ways, but all ultimately affect the dopamine reward system and the way in which it learns. Specifically, drugs influence reward timing signals in the dopamine system which makes it "surprised" and hyper-stimulated, implying that a reward-rich environment has been found. Consequently, the brain reduces GABA and serotonin inhibitory signals that typically prevent the dopamine system from directly controlling behaviour. That the use of drugs is paired with reliable cues for reward-rich consumption, makes the reward system learn to predict and respond to cues that are linked to drug use, which ultimately leads to the

relentless pursuit of addictive consumption. As RSVS (2008, p. 16) argue, the reward system "commits mutiny against the normal personal control apparatus."

To provide a bridge between these molar and molecular scale accounts, RSVS explain that choice behaviour at the molar level is what allows the reward system to learn. Similarly, changes in the reward system limit available options but do not preclude choice. As discussed previously, the majority of addicts quit by their early thirties and do so without seeking formal treatment, implying that a reward system which has adapted to seek out and use drugs does not completely predetermine behaviour. But from the other direction, recovered addicts who relapse tend to escalate drug use far faster than people in the early stages of addiction who have no history of dependence. This implies that reward system learning is not completely forgotten. Thus, the molar and molecular scales place constraints on each other but do not absolve the need to understand addiction at both levels.

The reconciliation provided by RSVS resolves the tension between choice-based and disease models of addiction by recognising that both of them contribute to our understanding of substance dependence and should be judged on their own respective levels of analysis. To borrow the analogy from Ross (2012, p. 336), there is nothing contradictory about the following two statements, "… in consequence of an inherited condition, lactose tolerant people are more likely to quench their thirst with milk than are lactose-intolerant people; and everyone who drinks milk chooses to do so."

In what follows, I will discuss economic models of choice which are relevant to the study of addiction and thereby highlight some of the molar scale factors which can precipitate substance dependence. I do this in full knowledge though that choice-based accounts of addiction provide only a partial understanding of this phenomenon, but, fortunately, one which is not at odds with a molecular scale interpretation.

III. ECONOMIC MODELS OF ADDICTION

Deep philosophical issues which models of addiction bring to light are the notions of choice, agency, and rationality in economic theory. Ross (2005, 2012, 2014) argues that any entity which behaviourally responds to incentives, specifically changes in relative opportunity costs, is an economic agent. This definition includes people as well as firms, industries, countries, nonhuman animals, and clusters of brain cells.

Economic agents, faced with various incentives and constraints, make choices. To make a choice, in the economic sense, agents must process changes in real or expected costs and benefits somewhere, but as Ross (2012, p. 325) argues, this processing is not necessarily "internal" to the agent. For example, changes in the external environment may limit the supply of certain goods (e.g., typewriters and fax machines) which leads to a reduction in their consumption without the agent internally processing this change. Instead the agent merely responds to the incentives and constraints created by the lack of supply (e.g., purchases a computer and a wireless router). As this example shows, choices need not involve any conscious processing or deliberation. Clearly some choices do involve this sort of reasoning but others are triggered by changes in the external environment and some are simply conditioned responses to stimuli; these conditioned responses still conform to the economic notion of *choice* because they are sensitive to their consequences.

An issue which will become apparent in my discussion of economic models of addiction is whether agency is tied to the whole biological life of a person. Some of the models (e.g., Becker and Murphy (1988), Orphanides and Zervos (1995), Laibson (2001)), either implicitly or explicitly, make this assumption while others (e.g., Gruber and Köszegi (2001), Benhabib and Bisin (2004), Fudenberg and Levine (2006, 2011, 2012)) adopt a more circumscribed notion of agency. The issue with models which assume that people are the prototypical economic agents is that people's tastes naturally change over time. This can be accommodated by indexing preferences to life-stages or particular life events which reliably track these changes (e.g., Becker and Mulligan (1997)). But in cases where preferences are not indexed in this way, we must assume that the agent (viz., person) maximises a single *lifetime* utility function which represents a stable set of preferences. Clearly this is a very

strong assumption and one which creates serious problems, discussed later, for the economic models of addiction which adopt it.

Ross (2005, 2012, 2014) argues that there is nothing in the axioms of economic theory which makes agency synonymous with a whole biological person. Thus, agency can instead be ascribed to temporal slices of a person, under the assumption that in each temporal slice a person's preferences are locally stable. Under this formulation, a biological person is made up of a sequence of economic agents and these agents interact strategically because later agents inherit the consequences (i.e., debts, contracts, etc.) of earlier agents' choices. A major strength of this approach is that changes in tastes are easily modelled as the Nash equilibria of intertemporal strategic interactions among a sequence of agents (recall the discussion of Ainslie's picoeconomic framework). Thus, models which adopt this formulation account for taste change at the level of the person (i.e., the molar level) by associating agency with a temporal slice of the person.

Ross (2014, p. 34) identifies two forms of rationality which have been emphasised in economics: 1) preference consistency over outcomes; and 2) the formation of accurate statistical predictions given available information. I will briefly focus on the first form of rationality (i.e., consistency of preferences over outcomes) because it is intimately linked to conceptions of agency, and is particularly relevant to economic models of addiction given that addicts often appear to make inconsistent (viz., irrational) choices, e.g., seek out and use drugs while paying real costs to prevent or limit the consumption of these goods.

In a deterministic intertemporal choice context, preference consistency means that an agent's valuation of different consumption bundles is invariant through time. In other words, if the agent prefers bundle A to bundle B at time point *t*, she will also prefer A to B at all other time points. Strotz (1955-1956) famously showed that, in the class of additively-separable intertemporal utility functions, such preference consistency only arises under the assumption of E discounting. Under this formulation, the agent compares the exponentially discounted costs and benefits of different consumption streams and chooses the one which maximises her utility through time.

Becker and Murphy (1988) modify this standard framework by incorporating intertemporal complementarities in consumption of the addictive good (viz., they abandon the assumption of complete additive separability of the intertemporal utility function). Their model is one of "rational addiction" where "rationality means a consistent plan to maximize utility over time" (Becker and Murphy (1998, p. 675)). While this model captures some of the salient features of addiction, it cannot explain the typical course of addiction (i.e., multiple unsuccessful quit attempts before eventual abstinence) because it conflates the agent with a whole biological person who must implement a consistent plan. As I will discuss later, the model also cannot account for the use of commitment devices and rehabilitation, and other seemingly inconsistent choices that are the hallmarks of addiction.

The models which I discuss in the next section have been grouped into three categories: 1) models of rational addiction; 2) dual self and dual system models; and 3) exogenous addiction models. I will provide an overview of these models and discuss their similarities and differences. The discussion of these models is used to provide a theoretical justification for the empirical analyses conducted in later chapters. These empirical analyses centre on the instantaneous risk and time preferences of addicts and non-addicts so particular attention will be given to the role of these preferences in the different theoretical models. As discussed in Chapter 1, this approach is consistent with Harrison, Lau and Rutström (2014) who argue that theory, experiments, and econometrics are complementary. In other words, it is crucial to review and understand theory because this informs experimental design and analysis. And analysis itself should be constrained by, and interpreted jointly with, theoretical considerations, prior empirical work, complementary data, econometric methodology, and intended applications.

III.I RATIONAL ADDICTION MODELS

Although many researchers made important contributions to economic theories of addiction (see Stigler and Becker (1977), Elster (1979), Spinnewyn (1981), Thaler and Shefrin (1981), Schelling (1984)), Becker and Murphy (1988) (BM)[10], building on the work of Stigler and Becker (1977), developed a theory of rational addiction which largely supplanted previous models and became the cornerstone of the neoclassical perspective on addiction (Bernheim and Rangel (2007)). The stated purpose of the model is to show that addiction can be explained as the outcome of forward-looking rational choice which does not rely on dividing the economic agent into multiple selves.

As mentioned previously, the BM model abandons the standard additively-separable intertemporal utility function so as to incorporate intertemporal dependencies in consumption of the addictive good. Specifically, BM assume that utility is additively separable over time in the stock of addictive capital, the addictive good, and the non-addictive good altogether, but not in the addictive and non-addictive goods alone because their marginal utilities are influenced by the stock of addictive capital. In other words, current consumption of an addictive good is affected by past consumption of the addictive good through changes in a person's stock of addictive capital.[11] Despite the fact that the intertemporal utility function is not completely additively separable in all variables, agents in the BM model have perfect foresight, discount exponentially, and therefore make time-consistent choices. As Johnsen and Donaldson (1985) explain, as long as intertemporal complementarities are properly encoded in current preferences, agents who discount exponentially will formulate and follow time-consistent plans.

[10] When I use the abbreviation BM without a date reference I am referring to Becker and Murphy's seminal article on addiction (Becker and Murphy (1988)). If I make reference to other papers that they have co-authored I will still use the abbreviation BM but include the relevant date reference. The same applies to all the author abbreviations that I use in this chapter.

[11] Elster (1997, p. 750-751) argues that BM use a person's stock of addictive capital as an analogy for a firm's stock of physical capital. Physical capital has three features: 1) it is produced by the investment decisions of firms; 2) investment lowers profits now but promises to increase profits in future; and 3) it decays or depreciates over time. The stock of addictive capital displays feature 1) but it does not typically display feature 2) in that consumption of an addictive good increases present utility at the cost of a decrease in future utility. Thus, consumption of an addictive good is akin to disinvesting. Furthermore, it is unclear whether it retains feature 3) because, as discussed earlier, relapse and escalation in usage after prolonged periods of abstinence is quite common. This implies that the effects of previous consumption have not depreciated or worn off completely.

The BM model tries to capture some of the salient features of addiction (e.g., tolerance, reinforcement, and withdrawal) through restrictions on the instantaneous utility function which agents maximise. To capture the effect of tolerance, BM assume that an increase in consumption of the addictive good leads to a decline in future utility through an increase in the stock of addictive capital. This is the sense in which the addiction is harmful: repeated use lowers welfare.[12]

This restriction also implies that abstention or a decline in consumption of the addictive good will lead to an increase in welfare over time due to a fall in the stock of addictive capital. However, a decrease in consumption of the addictive good leads to an immediate fall in utility which represents the physiological process of withdrawal.

The assumption that substance use leads to a decline in future welfare is one conception of tolerance to an addictive good but another, arguably more plausible one, would be that greater past consumption of an addictive good lowers the marginal utility of present consumption, implying that one needs to consume more of the good over time to enjoy the same overall level of utility. The DSM defines tolerance as, "… requiring a markedly increased dose of the substance to achieve the desired effect or a markedly reduced effect when the usual dose is consumed" (APA (2013, p. 484)). Thus, tolerance would be better captured by the assumption that past consumption lowers the utility of present consumption.[13]

However, the BM model assumes that higher consumption of the addictive good in the past raises the marginal utility of present consumption. In other words, the more

[12] It is also harmful in that greater consumption of the addictive good decreases the marginal utility from consumption of the non-addictive good. This can be interpreted as the crowding out effect that drug use often has on the enjoyment of other goods and activities.

[13] Ferguson (2000, p. 588) also recognises that the BM tolerance assumption is strange because addiction typically involves some form of habituation where greater prior use yields lower satisfaction from each successive unit of the addictive good (Goldstein (2001, p. 86)). He argues that this can be incorporated easily in the model by adding a new variable to the instantaneous utility function which represents "kicks" from consumption of the addictive good. These kicks would be produced by a production function which takes consumption of the addictive good as its argument and where the marginal utility of kicks increases as the stock of addictive capital rises. However, as consumption of the addictive good increases, the marginal productivity of addictive consumption in the production of kicks declines, implying that higher levels of consumption are necessary to satisfy the demand for kicks.

you have consumed, the greater the benefit from consumption now.[14] This is a necessary condition for the addictive good to capture the property of reinforcement (i.e., greater past consumption leads to greater present consumption). The sufficient condition for reinforcement is that the benefits from consumption now must offset the harmful effects which accumulate over time. When these necessary and sufficient conditions are satisfied, "adjacent complementarity" holds, which means that consumption of the addictive good is a complement, rather than a substitute, across time periods.[15,16] This assumption is central to most of the economic models of addiction reviewed in this section (an important exception is Bernheim and Rangel (2004)) because it provides the rationale for why agents continue to consume their targets of addiction despite the decline in welfare associated with increases in the stock of addictive capital.

The intuition of the BM model is relatively straightforward. People inherit a stock of addictive capital in time period 0. They then make intertemporal consumption decisions pertaining to addictive and non-addictive goods while taking into account the effect of these choices on their current and future welfare as determined by the restrictions on the instantaneous utility function and the addiction dynamics captured by the equation of motion for the stock of addictive capital. This is the sense in which agents are rational: they take into account all past, present and future implications of their consumption choices and implement a consistent plan to maximise their welfare through time.

The BM model admits multiple possible consumption patterns depending on the relationship between the immediate benefits of consumption of the addictive good and the long-term consequences which result. This is the tension inherent in the model.

[14] This resembles the notion of "cravings" which is defined as, "… an intense desire or urge for the drug" (APA (2013, p. 483)).

[15] Adjacent complementarity implies that anything which affects consumption of the addictive good at one point in time will affect consumption of the addictive good at all points in time. For example, an expected increase in future prices will not only decrease consumption when the change comes into effect but will also decrease consumption in every period leading up to that date. This is an important testable implication of the BM model and has led to a cottage industry of econometric models which attempt to show that consumers respond to future price changes by adjusting current consumption. For a review of this literature in relation to tobacco smoking consult Chaloupka and Warner (2001). For critiques of this literature consult Ferguson (2000) and Baltagi (2007).

[16] Adjacent complementarity also has important implications for intertemporal risk preferences. Specifically, in *stochastic* economic models of addiction, it yields correlation seeking behaviour. As the BM model is deterministic though, I defer this discussion until later.

Consumption of the addictive good raises welfare now and makes addictive consumption more attractive in future while simultaneously lowering future welfare through increases in the stock of addictive capital. So, do you consume now and enjoy the short-term benefit of addictive consumption while suffering the long-term costs or do you abstain now and enjoy the long-term benefits which arise due to a decrease in the stock of addictive capital?

Given the deterministic nature of the BM model, there is no place for subjective beliefs, instantaneous risk preferences nor intertemporal risk preferences in the analysis of choice behaviour.[17] The role of time preferences, however, is readily apparent. As BM (p. 682) assert, "our analysis implies the common view that present-oriented individuals are potentially more addicted to harmful goods than future-oriented individuals." In other words, people who discount the future heavily will tend to focus on the short-term benefits of consumption of the addictive good while putting relatively little weight on the future costs, thereby inclining them to consume the addictive good. This is a testable implication of the BM model which has been studied extensively in psychology but only in a handful of studies in economics (e.g., Chabris et al. (2008), Harrison, Lau and Rutström (2010), Kang and Ikeda (2014)). In Chapter 3, I provide a detailed review of experimental studies focussing on the relationship between time preferences and smoking behaviour and highlight the methodological and statistical issues that plague the bulk of studies in this literature.

The fact that consumers choose to become addicts and are "happy," in a utility-maximising sense, with the consequences of their choices is one of the most controversial implications of rational addiction because the common view of addicts is that they regret their choices and would be better off if they had never consumed the addictive good. As Gul and Pesendorfer (2007, p. 150) state, "the Becker-Murphy formulation entails an *a priori* rejection of 'the problem' of addiction." According to the BM model, even though a person's welfare declines over time through consumption of the addictive good, this was taken into account when the good was originally consumed (i.e., choices were made with full cognisance of their repercussions and as such there is no place for regret). This has clear implications for

[17] The impact of subjective beliefs, and instantaneous and intertemporal risk preferences come to the fore in the model of Orphanides and Zervos (1995), which is discussed next.

policy: any restrictions placed on consumption of the addictive good, other than those to eliminate negative externalities, are necessarily welfare reducing. As BM (p. 691) assert, "[people] would be even more unhappy if they were prevented from consuming the addictive goods."

Orphanides and Zervos (1995) (OZ) take the BM model and relax the assumption of perfect foresight so as to incorporate the roles of instantaneous risk preferences, intertemporal risk preferences, and subjective beliefs in the explanation of addiction. Note that the OZ model is still rational in the sense that people form consistent plans to maximize *expected* utility through time. However, these plans are based on agents' subjective beliefs about their addiction potential and these beliefs, although updated optimally via a Bayesian learning process, can lead to addiction and regret.

The OZ model is based on three fundamental assertions: consumption of the addictive good affects people in different ways (i.e., it may be very harmful to some but not at all harmful to others); people have subjective beliefs concerning their addiction potential; and these beliefs are updated through the information that consumption of the addictive good imparts.

To capture the idea that people have different susceptibilities to addiction, OZ assume that the population consists of two groups: *non-addicts* and *potential addicts*. Increases in the stock of addictive capital have no effect on non-addicts but can produce harmful side effects in potential addicts. An agent formulates her consumption plan by trading off the costs and benefits of consumption of addictive and non-addictive goods. Whether a particular consumption plan leads to addiction critically depends on a person's vulnerability to addiction and the speed with which uncertainty about the person's true type (i.e., non-addict or potential addict) is resolved.

If a potential addict consumes the addictive good and experiences no negative side effects then she will revise the probability of being a non-addict upwards over time. When the uncertainty regarding her true type is finally resolved and she realises she is a potential addict, her subsequent consumption decisions will be determined by the stock of addictive capital she has accumulated over time. The longer it takes for this

uncertainty to be resolved and the larger the stock of addictive capital, the more likely it is that the person will become addicted.[18]

The preceding discussion highlights one of the crucial insights of the model, which is that addiction is voluntary but unintentional. People in this model do not want to become addicts because they recognise the harmful consequences of addiction. But they may decide to risk the possibility of addiction because they are unsure of their addiction potential and consumption of the addictive good is attractive. As OZ (p. 741) argue, "[addiction] is the unintended occasional outcome of experimenting with an addictive good known to provide certain instant pleasure and only probabilistic future harm." What this implies is that addicts in the OZ model are not "happy addicts" because they did not want to become addicted ex ante. Owing to their erroneous initial beliefs and/or the slow resolution of uncertainty, they were "hooked" into an addiction and now suffer regret.

Thus the BM and OZ models differ markedly in their explanation of addiction. In the BM model, people choose to become addicts after comparing all the costs and benefits of consumption of the addictive good. In the OZ model, people do not want to be addicts but may choose to experiment with the addictive good, which over time leads to a harmful addiction. In addition, the BM and OZ models have notably different policy implications. As discussed previously, there is no room for government intervention in the BM model except to combat externalities associated with consumption of the addictive good. The OZ model, by contrast, prescribes a more interventionist role for government policy, specifically in the provision of accurate information about the true probability of addiction P^*. Furthermore, there may be a role for insurance against people's idiosyncratic risks of addiction.

[18] A person's initial beliefs P_0 play an important role in the decision to experiment with the addictive good. Let P^* represent the true probability that the person is a *non-addict*. In the case where someone knows P^* with certainty, she will set her initial belief $P_0 = P^*$. In general this is not the case and the person will need to form her initial beliefs by using information from different sources. The weight that people give to particular sources can have a large impact on the formation of their initial beliefs. For example, if a person puts a lot of weight on the information provided by a non-addict, this could yield an overoptimistic prior, $P_0 > P^*$, which inclines the individual to experiment with the addictive good. By contrast, if someone puts a high weight on information from an anti-drug campaign this may yield a conservative prior, $P_0 < P^*$, and prevent that person from experimenting with the good.

The OZ model makes it clear how instantaneous risk and time preferences, intertemporal risk preferences and subjective beliefs interact to influence addiction.[19] In the case of time preferences, if the future costs of consumption of the addictive good are discounted at a high rate then potential addicts may be inclined to experiment with the addictive good which provides immediate benefits. With only moderate discounting of the future though, both the present benefits and future costs of consumption of the addictive good will loom large in the mind of a potential addict and may prevent her from consuming the good.

In the case of instantaneous or atemporal risk preferences, the stochastic nature of the OZ model implies that every decision is affected by people's attitudes toward risk. Someone who is very risk averse with respect to atemporal payoffs may never choose to consume the addictive good as she fears the risk of harmful consequences and/or the realisation of her true type. Someone who is less risk averse with respect to atemporal payoffs may be willing to gamble on the possibility of harmful consequences and/or the realisation of her addiction potential. Thus, instantaneous risk preferences are crucial for explaining the initial choice to use an addictive good and any subsequent decisions to increase or decrease consumption of the good.

Intertemporal risk preferences[20] play an important role in the OZ model, and in the other stochastic models which I discuss in this chapter, and thereby deserve further comment. Richard (1975) is credited with introducing intertemporal risk preferences to the economic literature (see Eeckhoudt, Rey and Schlesinger (2007)), albeit with different terminology, but the concept apparently first appeared in de Finetti (1952). Richard (1975) basically extended the notion of risk preferences over one variable to

[19] The role of subjective beliefs has been discussed in detail so I will only focus on instantaneous risk and time preferences and intertemporal risk preferences in the remainder of this section.

[20] The literature on intertemporal risk preferences emerged from the literature on multi-attribute utility theory (see Keeney and Raiffa (1976) and Andersen, Harrison, Lau and Rutström (2014b) for reviews). A multi-attribute utility function captures the idea that agents may take into account multiple characteristics or attributes of a good when making choices. For example, suppose someone wants to purchase a dishwasher and cares both about the speed with which it finishes its cycle and its energy efficiency. One machine may be very fast but energy inefficient while another is slower but more efficient. To represent the person's preferences over these different attributes one could employ a multi-attribute utility function. In the context of intertemporal consumption streams, both the amounts of different goods and the times at which they are received can be regarded as two distinct attributes or characteristics of the consumption stream. Viewed in this way, preferences over intertemporal consumption streams are modelled naturally using a multi-attribute utility function.

risk preferences over multiple variables and referred to the latter as multivariate risk aversion.[21]

To understand this idea, consider the following two lotteries which yield outcomes in two time periods. Lottery L_1 gives an even chance of either high consumption in period 1 and low consumption in period 2 or low consumption in period 1 and high consumption in period 2. Lottery L_2, by contrast, gives an even chance of either high consumption in both periods or low consumption in both periods. Note that the outcomes in Lottery L_1 are negatively serially correlated whereas the outcomes in lottery L_2 are positively serially correlated. If an agent chooses lottery L_1 over lottery L_2, she displays intertemporal risk aversion or correlation aversion because, as Richard (1975, p. 12) remarked, "[she] prefers getting some of the 'best' and some of the 'worst' to taking a chance on all of the 'best' or all of the 'worst.'"[22] If the agent is indifferent between the lotteries she is intertemporally risk neutral or correlation neutral, and if she prefers lottery L_2 to lottery L_1 she is intertemporally risk seeking or correlation seeking.

Richard (1975) showed that the sign of the cross partial derivatives of an agent's intertemporal utility function determines her preferences towards serially correlated lotteries. Specifically, if $\partial^2 U(\mathbf{c})/(\partial c_i \partial c_j) \leq 0$ the agent is intertemporally risk averse. In words, if the intertemporal utility function's cross partial derivative is non-positive then the agent prefers lotteries where the outcomes are negatively serially correlated because the marginal utility of current consumption is decreasing in past consumption. By contrast, if $\partial^2 U(\mathbf{c})/(\partial c_i \partial c_j) = 0$ the agent is intertemporally risk neutral, whereas if $\partial^2 U(\mathbf{c})/(\partial c_i \partial c_j) \geq 0$ the agent is intertemporally risk seeking. Thus,

[21] Researchers in the field of intertemporal risk preferences typically employ the risk averse component of these preferences in their terminology. Keeney (1973) used the term "conditional risk aversion," Richard (1975) referred to "multivariate risk aversion," Epstein and Tanny (1980) defined "correlation aversion," Strzalecki (2013) employed "long-run risk aversion," and Andersen, Harrison, Lau and Rutström (2014b) used "intertemporal risk aversion" or "intertemporal correlation aversion." I prefer the term "intertemporal risk preferences" because it does not presuppose an aversion to lotteries with positive serial correlation. I think this is particularly important in the context of economics models of addiction because adjacent complementarity is linked to intertemporal risk seeking behaviour. Note that when I refer to a particular type of intertemporal risk preference (i.e., intertemporal risk aversion) I will use the "intertemporal risk" and "correlation" prefixes (i.e., correlation aversion) interchangeably.

[22] Andersen, Harrison, Lau and Rutström (2014b, p. 5) provide another intuitive definition for correlation aversion when drawing an analogy between instantaneous risk aversion and intertemporal risk aversion: "The correlation averse individual prefers to have non-extreme payoffs *across* periods, just as the risk averse individual prefers to have non-extreme payoffs *within* periods."

the form of an agent's intertemporal utility function determines her intertemporal risk preferences, just as the form of an agent's instantaneous utility function determines her instantaneous risk preferences under expected utility theory.[23]

The standard additively-separable intertemporal utility function exhibits intertemporal risk neutrality because consumption at different points in time are independent so the cross partial derivatives of the intertemporal utility function are necessarily zero. Thus, even though people may be risk averse, risk neutral or risk seeking over instantaneous or atemporal lotteries, an additively-separable intertemporal utility function yields correlation neutrality.[24,25] The natural question that emerges in the context of this chapter is: what does the form of the intertemporal utility function in economic models of addiction imply for agents' intertemporal risk preferences?

Recall that adjacent complementarity, which drives the addiction dynamics in most economic models of addiction, means that consumption of the addictive good is a complement across time periods; this clearly has implications for intertemporal risk preferences. As Bommier and Rochet (2006, p. 725-726) recognise, models of habit formation, of which economic models of addiction are some of the exemplars, place strong restrictions on the intertemporal utility function. Specifically, as consumption of the addictive good increases the stock of addictive capital and increases in the stock of addictive capital increase the marginal utility of addictive consumption, $\partial^2 U(\mathbf{c})/(\partial c_i \partial c_j) > 0$, implying that agents in these models are *typically* correlation seeking.

That agents in these stochastic models tend to be correlation seeking makes intuitive sense in the context of addiction. Referring back to the lotteries L_1 and L_2, the

[23] Much as the literature on instantaneous choice under risk has evolved to incorporate rank and sign dependence (which I discuss in more detail in Chapter 4), so too has the literature on intertemporal risk preferences (see Fishburn (1984), Miyamoto and Wakker (1996)).

[24] The same is true for other combinations of instantaneous and intertemporal risk preferences. In other words, agents can be risk averse, risk neutral or risk seeking with respect to atemporal payoffs and yet exhibit correlation aversion, correlation seeking or correlation neutrality with respect to intertemporal lotteries.

[25] That an additively-separable intertemporal utility function yields correlation neutrality has the unfortunate implication that the intertemporal elasticity of substitution simply equals the inverse of instantaneous risk attitudes. In economic models of addiction, where the intertemporal utility function is not completely additively separable, this link between instantaneous risk attitudes and the intertemporal elasticity of substitution is broken (see Bommier (2007), Andersen, Harrison, Lau and Rutström (2014b)).

correlation seeking agent prefers to gamble on all of the best or all of the worst, rather than accept some of the best and some of the worst. With regard to addictive consumption, this preference for lotteries with positive serial correlation seems to characterise the intertemporal risk preferences of addicts quite well: they chase the positive serial correlation in "good" realisations of addictive lotteries while running the risk of positive serial correlation in "bad" realisations of addictive lotteries. Correlation averse agents, by contrast, prefer non-extreme payoffs across periods (e.g., a binge in one period followed by abstention in the next), which potentially, but only partially, immunises them against the progressive, deleterious effects of addiction.

Although stochastic economic models of addiction typically assume, sometimes implicitly, that agents are correlation seeking, Bommier and Rochet (2006) and Lichtendahl, Chao and Bodily (2012) show that correlation averse preferences are not incompatible with habit formation. Lichtendahl, Chao and Bodily (2012) develop two models, one of which is a generalised exponential utility model and another which is a generalised power utility model, which assume correlation aversion *a priori* and yet still generate intertemporal dependencies in consumption. This leads the researchers to conclude that habits should be understood from two perspectives. The first, which they refer to as "habit formation," can be viewed as an input to a decision-making model, and they show that habit formation is consistent with both correlation seeking and correlation averse preferences. The second, which they refer to as "habit following," can be understood as an output of a behavioural model: specifically, that current behaviour positively influences future behaviour. Thus, while economic models of addiction tend to assume correlation seeking preferences, one can still generate some of the qualitative features of addiction (e.g., reinforcement in consumption of the addictive good) with correlation averse preferences. As I will discuss later, this is a feature of the model by Bernheim and Rangel (2004).

Andersen, Harrison, Lau and Rutström (2014b) conducted a set of instantaneous risk and time preference experiments, coupled with an intertemporal risk preference experiment, on a representative sample of 413 people in Denmark. The researchers used a full information maximum likelihood statistical framework to jointly estimate instantaneous risk and time preference parameters and an intertemporal risk

preference parameter. They found that, on average, the sample displayed intertemporal risk aversion but that about 5% of the sample was intertemporally risk neutral or risk seeking. Although not a focus of their research, it would be interesting to use this robust experimental design and statistical framework to explore potential differences in the intertemporal risk preferences of addicts and non-addicts. It may be that addicts are more likely to be correlation seeking than non-addicts, which would provide support for the basic structure of most economic models of addiction, but if not, the theoretical work of Lichtendahl, Chao and Bodily (2012) is particularly important because correlation seeking is not a prerequisite for habit formation.

Given the importance of instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs in the determination of choice behaviour in the OZ model, it is worthwhile to empirically test the extent to which these preferences and beliefs differ among addicts and non-addicts. In Chapter 5, I will present results from a study which was designed to elicit people's instantaneous risk and time preferences so as to explore rigorously potential differences between smokers and non-smokers; investigation of potential differences in the intertemporal risk preferences and subjective beliefs of smokers and non-smokers is deferred for future research.

In sum, the OZ model is an important extension to the theory of rational addiction for a number of reasons. First, it deals with the objection to "happy addicts" in the BM model; addicts in the OZ model are not happy, did not choose to be addicts ex ante and regret experimenting with an addictive good. Second, the model incorporates the uncertainty inherent in consumption of the addictive good. That people can perfectly foresee the long-term negative consequences of consumption of an addictive good when first deciding whether to experiment with it, is a very strong assumption. OZ relax it to show how the slow resolution of uncertainty can hook someone into a harmful addiction. The model also provides an explanation for the success of public awareness campaigns which provide people with information on the true probability of addiction and the negative side effects which result from consumption of addictive goods. Finally, the stochastic nature of the model highlights the role that instantaneous and intertemporal risk preferences play in addiction.

However, as discussed previously, theories of rational addiction assume time consistency and tie agency to the whole biological life of a person. These assumptions are so restrictive that the BM and OZ models fail to account for the typical course of addiction (i.e., multiple unsuccessful quit attempts, before final abstention), a point recognised by Ross (2010). People who are substance dependent do not usually consume drugs until they have exhausted their budgets, as these models imply. Rather, they tend to expend resources to prevent or limit their consumption of these goods.

Ross (2010) explains that the models of rational addiction could incorporate this by assuming that addicts limit their consumption so as to reduce their tolerance to drugs; this would allow them to get high cheaply, at least for a time. But if this is the rationale for limiting use, the models must then explain why most addicts eventually quit. To accommodate this fact, the models must assume that addicts only come to learn how bad addiction is through experience and this prompts a quit attempt.[26] While this behaviour can be rationalised, it is difficult to reconcile with the tendency for addicts to relapse. To do so, we must assume that when an addict quits and then relapses, she was merely trying to recalibrate her tolerance levels. But, when she finally quits for good, we must then assume that she eventually learnt that she was doing herself more harm than good. If she never quits then we must assume that she was not doing herself more harm than good, at least in her own terms.

As Ross (2010, p. 133) argues, "… to capture addiction in a traditional economic model, it seems we must characterise the most common course of addiction as involving a radical psychological discontinuity between unsuccessful and successful attempts at quitting: the former are for the sake of enhanced enjoyment of the addiction, and the latter result from a decision to abandon this enjoyment." This

[26] In the OZ model, the agent understands the consequences of addiction but may nevertheless decide to consume the addictive good so as to reveal information about her true type. Thus, one would still need to revise the OZ model such that agents are ignorant about the consequences of addiction. If the model was adapted in this way (i.e., agents are ignorant of their addiction potential and unaware of the consequences of addiction) it would be reminiscent of the literature in social psychology (e.g., Kruger and Dunning (1999)) which ties competence in a particular domain to the ability to evaluate competence in that domain. The analogue for the refined OZ model would be that if a person is ignorant of their addiction potential they may also be unable to accurately assess the consequences of addiction.

psychological explanation is unpersuasive and there is no empirical evidence to support it.

The BM and OZ models of rational addiction have generated a remarkable amount of controversy since their publication (for a particularly scathing critique consult Rogeberg (2004)). Despite their shortcomings, it is crucial to understand the importance of these theories in modelling addiction. One the BM model's foremost critics recognised this when he wrote:

> "Although I disagree sharply with much of it, it has raised the level of discussion enormously. Before Becker, most explanations of addiction did not involve choice at all, much less rational choice. By arguing that addiction is a form of rational behavior, Becker offers other scholars the choice between agreeing with him or trying to identify exactly where he goes wrong. Whatever option we take (I'm going to take the second), our understanding of addiction will be sharpened and focused." Elster (1997, p. 758).

As Elster recognises, the BM and OZ models of rational addiction provide some important insights on this phenomenon. The BM model captures the tension between consuming addictive goods and abstaining from them, and it highlights the role that time preferences play in this decision. The stochastic nature of the OZ model shows how instantaneous and intertemporal risk preferences, and subjective beliefs, may foster experimentation with drugs which can ultimately hook someone into a harmful addiction. However, the models misinterpret the physiological process of tolerance so as to incorporate reinforcement in consumption of the addictive good which ultimately drives the addiction dynamics in the models. More importantly, both models are hamstrung by the assumption of time consistency and their decision to link economic agency to the whole biological life of a person. This means that the models cannot account for addicts' time-inconsistent choices nor quit attempts and relapse and, thus, they fail to capture the typical course of addiction. The models in the next section dispense with this agency assumption so as to better capture the salient features of substance dependence.

III.II DUAL SELF AND DUAL SYSTEM MODELS

Dual self and dual system models are subpersonal in nature which means they do not link agency to the whole biological life of a person. These models are either diachronic or synchronic or both.[27] Dual system models are typically synchronic in that multiple systems or processes jostle for control of behaviour at the same time. Dual self models, by contrast, are usually diachronic which means that each self has full control of the person's cognitive and other capacities at a single point in time but that there is a succession of selves over time. One dual self model (Fudenberg and Levine (2006, 2011, 2012)) in this section is actually both synchronic and diachronic because in each time period a short-run (SR) and long-run (LR) agent compete for control of the person, but there is a succession of short-run selves over time. As mentioned above, a real strength of the subpersonal agency assumption is that it allows one to model time-inconsistent choice behaviour. All of the models in this section incorporate this possibility but do so in different ways.

Diachronic Models

In the spirit of early economic models of addiction which divided agents into multiple selves (see Winston (1980), Thaler and Shefrin (1981), Schelling (1984)), Gruber and Köszegi (2001) (GK) extend the model of rational addiction to incorporate time-inconsistent preferences. When people are time-inconsistent, the preferences they have now over options available in the future can differ to the preferences they have

[27] There is rich tradition of dual self and dual system models in psychology (see Chaiken and Trope (1999), Stanovich and West (2000) for reviews). As Andersen, Harrison, Lau and Rutström (2014a) explain, most models of choice in economics posit multiple psychological processes but are ultimately characterised in terms of single decision criteria. By contrast, models in psychology often assume that multiple decision criteria simultaneously affect choice; a good example is the SP/A theory of Lopes (1984, 1995). Unlike expected utility theory, rank-dependent utility theory, and cumulative prospect theory, which characterise choice in terms of a single decision criterion, SP/A theory evaluates lotteries in terms of dual criteria. Specifically, the first criterion is based on the "security" and "potential" (i.e., SP) of a lottery, where security relates to the smallest outcomes in a lottery and potential to the highest. Thus, the first criterion gives special weight to the extreme outcomes in a lottery and less weight to the intermediate outcomes. The second criterion is the "aspiration" (i.e., A) level of the decision maker and serves as a reference point which determines the extent to which a lottery satisfies the subject's aspirations. Conflict between the SP and A criteria ultimately determine choice behaviour according to the theory and Andersen, Harrison, Lau and Rutström (2014a) show how one can model the conflict between these different criteria using finite mixture models; these models will be used extensively in later chapters.

over these same options as the future becomes the present.[28] For example, someone may prefer to study hard next week for a test rather than procrastinate but when next week arrives they may prefer to procrastinate rather than study hard. In other words, the mere passage of time can lead to changes or reversals in preference thereby, "raising the spectre of intrapersonal conflict over decisions that have implications for the future" (GK, p. 1277). To model time-inconsistent preferences, GK assume that each person is made up of multiple selves and that each intertemporal incarnation of the agent (i.e., each self) makes choices according to the preferences prevailing at the time (i.e., the model is diachronic).

As discussed previously, the BM model incorporates standard time-consistent preferences by assuming that people discount future utility with a constant E discount factor $D^{E}(t) = 1 / (1 + \delta)^{t}$ and that they correctly encode the effect of current choices on future welfare (i.e., they have perfect foresight). Phelps and Pollak (1968) developed a mathematically tractable quasi-hyperbolic (QH) discounting function which incorporates time inconsistency. This specification was popularised by Laibson (1997) and it is the one that GK employ. The QH discounting function has a discount factor:

$$D^{QH}(t) = 1 \qquad \text{if } t = 0 \qquad (1a)$$

$$D^{QH}(t) = \beta / (1 + \delta)^{t} \qquad \text{if } t > 0 \qquad (1b)$$

Note that when $\beta < 1$, the discount factor between the current period and the next one ($\beta / (1 + \delta)$) is lower than the discount factor for consecutive periods in the future ($1 / (1 + \delta)$); hence, non-constant discounting. Thus, when $\beta < 1$, people have a "present-bias"[29] or a "passion for the present" and may fail therefore to carry out plans they make for the future. In other words, someone's current preferences over future

[28] Stated formally, under an additively-separable intertemporal utility function, preferences are time-inconsistent when the marginal rate of substitution between consumption at two future points in time varies according to the date at which it is evaluated (Caillaud and Jullien (2000)).

[29] O'Donoghue and Rabin (1999, p. 106) define present-biased preferences in the following way, "When considering trade-offs between two future moments, present-biased preferences give stronger relative weight to the earlier moment as it gets closer."

options may diverge from the choices they make over these options in the future (i.e., they may act in a time-inconsistent manner).

GK cite experimental studies from psychology as one reason for incorporating time-inconsistent preferences in their model; as these studies typically employ hypothetical incentives, elicitation methods which are not incentive compatible, and statistical tools which are not appropriate for the data that are collected, one should treat the results with caution. In a review of this literature, Green and Myerson (2004) argue that H discounting functions[30], almost without fail, fit discounting data better than E functions.[31] In other words, in the experiments which Green and Myerson (2004) review, which predominately use hypothetical rewards and lack incentive compatibility, people tend to make choices which reveal a higher discount factor for options available in the future and a lower discount factor when these options are shifted closer to the present (i.e., non-constant discounting).

Time-inconsistent preferences also provide an explanation for the use of commitment devices by addicts. Note that in the BM model addicts have no need for commitment devices because they form consistent plans through time. Consequently, any constraints on their behaviour will be welfare reducing. By contrast, if someone has time-inconsistent preferences then she can benefit from the use of commitment devices to lock in her current preferences if they are apt to change in future. Thus, the GK model provides a rationale for the use of rehabilitation (viz., a costly commitment device) by addicts.

Bénabou and Tirole (2004) (BT) develop a formal model of willpower and personal rules which is based on the picoeconomic framework of Ainslie (1992, 2001), discussed earlier. While the GK model emphasises the importance of external commitment devices for locking in long-term preferences and controlling behaviour,

[30] The studies which Green and Myerson (2004) review typically make use of Mazur's (1984) H discounting function which has a discount factor: $D^{H}(t) = 1 / (1 + \delta t)$. Ainslie (2012) discusses the important qualitative differences between H and QH discounting functions.

[31] This is a very strong claim and one which will be interrogated in Chapters 3, 4, and 5. As noted above, experimental studies in psychology typically use hypothetical incentives and elicitation methods which are not incentive compatible, so people have little reason to respond truthfully. In addition, the statistical tools that have been applied to these data do not take into account all the information that people's choices impart and they do not allow multiple data generating processes to account for these data. This makes it difficult to draw robust inferences from these studies.

Ainslie (1992, 2001) and BT focus on internal commitment devices (i.e., personal rules) for achieving the same ends.

A major stumbling block with the notion of internal commitment mechanisms is that they lack enforceability and hence, credibility. In other words, why would someone's resolution not to eat chocolate have any effect on her behaviour if her resolution does not circumscribe her behaviour in any way and if she can renege on the promise without any tangible costs or punishments? According to Ainslie (1992, 2001), and formalised by BT, personal rules acquire their force through their impact on self-reputation. Thus, sticking to your personal rules serves a self-signalling function (see Bodner and Prelec (2001)) about the type of person you are and if this identity is valuable, in that it allows you to overcome temptations, you may take costly actions to support it. What this implies is that if you succumb to temptation in any one instance, this lapse may set a precedent for future lapses and thereby undermine your reputation and your long-term goals.

There are two key features of the BT model: imperfect willpower, which is captured by a QH discounting function; and imperfect recall of past actions or the circumstances in which past actions were taken. This imperfect recall means that agents have a limited understanding of their willpower and propensity for self-control. Such understanding can only be acquired through direct experience but direct experience does not perfectly reveal an agent's willpower because she is apt to forget the exercise of willpower, or lack thereof, and the extraneous circumstances involved when her willpower was last put to the test.

The development and maintenance of personal rules in the context of uncertainty about the agent's preferences (i.e., willpower) is modelled as a game of imperfect information among the agent's selves in each period (i.e., the model posits diachronic subpersonal agents). Ultimately this is a signalling game which admits separating equilibria (i.e., where weak willed and strong willed agents take different actions) and pooling equilibria (i.e., where weak and strong willed agents take the same actions). But unlike signalling games involving multiple players, this game is about signals sent between different intertemporal incarnations of the same player.

With regard to the literature on addiction, agents in the BT model have a strong desire for external commitment devices that make consumption less likely and informational aids to improve their recall of past motives and actions. But the real strength of the model is that it provides a rationale for internal commitment mechanisms which allow agents to bundle rewards and achieve self-control. Where these personal rules are eroded, either through transgression or ex post rationalisation, excessive consumption is likely. As a corollary to this result, BT show that external constraints on behaviour can make an individual less likely to formulate their own personal rules and thus less likely to exercise self-control when these external constraints are lifted or relaxed.

The BT model formalises many of Ainslie's (1992, 2001) insights but does so by assuming a QH specification as opposed to Ainslie's preferred H discounting function (see Mazur (1984)). As Ainslie (2012) explains, H discount curves provide a natural explanation for the formation of personal rules while the QH function does not. Consequently, BT rely on the imperfect recall of past actions, which means agents must draw inferences about their past behaviour, to provide a rationale for the creation and maintenance of personal rules. Ainslie (2012, p. 27) critiques this "roundabout route" which is used to preserve the implications of his theory.

Note that both the GK and BT models explain the pattern of quit attempts and relapse which typically characterises addiction. The present self may decide to quit her target of addiction, possibly with the help of a commitment device, only to have her plan undone by a future self who is tempted by the immediate benefit of addictive consumption. In the GK model this occurs when the agent fails to recognise her future self-control problem (i.e., the agent is "naïve" in the terminology of O'Donoghue and Rabin (1999)). In the BT model, a personal rule which bundles the LL rewards of quitting but which lacks credibility, gets hijacked by the immediate temptation of addictive consumption. That the models incorporate the typical course of addiction, highlights the benefit of divorcing agency from biology.

As the preceding discussion shows, the role of time preferences take centre stage in the GK and BT frameworks. But, rather than simply being an issue of how heavily the future is discounted, time preferences in these models also determine the extent to which people's behaviour is time-inconsistent. For example, suppose that with $\beta = 1$,

like in the BM model, a person's discount rate δ is low enough that she decides not to consume the addictive good and thus does not become addicted. An otherwise identical person with a present-bias (i.e., $\beta < 1$)[32] may choose to consume the addictive good now with every intention of limiting consumption in the future so as to avoid addiction. However, this noble intention may be undone by the person's self-control problem, thereby leading to a cycle of good intentions and lapses which ultimately leads to addiction. Thus, it is worthwhile to examine empirically the extent to which addicts and non-addicts discount the future according to the QH discounting function. It may be the case that addicts and non-addicts discount the distant future similarly but that addicts have a passion for the present which non-addicts do not.

This discussion highlights a potential shortcoming of the GK and BT models: they assumes that a present-bias *always* influences intertemporal choice behaviour. In other words, a passion for the present is not specific to addictive goods, which means that anything which provides immediate benefits but entails future costs will be consumed excessively from the perspective of former selves unless internal or external commitment devices are in place.

By adopting the QH discounting specification, the GK and BT models account for addicts time-inconsistent choices, their use of commitment devices, and the pattern of quit attempts and relapse. It also promotes empirical tests of the discounting behaviour of addicts and non-addicts, not only with respect to the steepness of discounting, but also the extent to which discounting is not constant.

Synchronic Models

The models in this section posit synchronic subpersonal agents that jostle for control of behaviour at each moment in time. Both models assume E discounting but manage to generate inconsistent choice behaviour through the presence of stochastic temptations. These temptations potentially overwhelm cognitive processes and drive consumption in the direction of immediate gratification. As these temptations are

[32] O'Donoghue and Rabin (1999) assume that time-consistent and time-inconsistent agents do not discount the future at all ($\delta = 0$) to focus specifically on the role of present-bias ($\beta < 1$) in decision-making.

realised stochastically, the role of instantaneous and intertemporal risk preferences come to the fore in these models.[33] This stands in stark contrast to the deterministic, diachronic models reviewed previously, where time preferences are responsible for inconsistent choice behaviour.

Bernheim and Rangel (2004) (BR) develop an economic theory of addiction based on the assumption that the brain has a Hedonic Forecasting Mechanism (HFM) which uses environmental cues to predict the short-term hedonic effects of consumption. The HFM works quickly and, at least in the case of non-addictive goods, gives relatively accurate predictions and tends to produce optimal outcomes (i.e., ones where preferences and choices coincide). However, addictive goods act directly on the HFM and exaggerate the predicted short-term benefits of consumption. If people made all of their choices on the basis of HFM predictions, minimal exposure to addictive goods would rapidly trigger addiction.[34]

But people also make use of higher cognitive processes that are responsible for functions like deliberation and future planning which can override exaggerated HFM predictions. Competition between these systems often yields outcomes where preferences and choices coincide but can, particularly in the case of addictive goods, lead to outcomes where preferences and choices diverge.[35] Thus, the BR model is a dual system approach to inconsistent choice which posits synchronic subpersonal agents.

To incorporate the roles of the HFM and higher cognitive processes in decision-making BR develop the following model. They assume that people make decisions in one of two modes: a deliberative "cold" mode where higher cognitive processes take

[33] This does not imply that time preferences play no role in these models. Time preferences always matter in a dynamic choice context because they affect a person's valuation of rewards available at different points in time. However, the role of time preferences in these synchronic models is no different to their role in theories of rational addiction so I will not dwell on them in any detail.

[34] BR cite evidence from incentivised laboratory experiments with rats (e.g., Gardner and David (1999)) showing that these nonhuman animals rely exclusively on HFM predictions and, if given the chance, will compulsively self-administer cocaine until they die. The rats therefore ignore all other drives like hunger, thirst, and reproduction due to the exaggerated HFM predictions associated with cocaine use.

[35] The notion that preferences and choices can diverge is heretical to revealed preference theory because choices reveal preferences. In other words, preferences are understood as a summary of choices (Ross (2014, p. 251).

precedence and ensure that choices and preferences coincide; and a dysfunctional "hot" mode where the HFM assumes control and choices and preferences may diverge.

The agent enters each period in the cold mode and chooses an activity after taking into account all the present and future costs and benefits associated with that activity. This choice, in combination with the person's addictive state and other environmental factors, determine the likelihood that the person observes cues related to past consumption of the addictive good which trigger the hot mode. If the hot mode is triggered, the person always consumes the addictive good regardless of underlying preferences. If not triggered, the person rationally chooses whether to consume the addictive good or not. BR assume that initial use of an addictive good is intentional but that through repeated use, people become more susceptible to stochastic addictive cues which trigger mistaken usage.

The BR model's distinctive features arise due to the inclusion of state-contingent stochastic shocks which can trigger the hot mode and cause preferences and choices to diverge; without these shocks the model reduces to BM. These taste shocks explain cue-triggered relapse, cue-avoidance through cognitive and behavioural therapies, and rehabilitation which commits the agent to abstention.

While other economic models of addiction rely on adjacent complementarity to generate addictive consumption patterns, the BR model can do without it as long as the probability of entering the hot mode increases with the addictive state.[36] This is an important departure from other theories because, as discussed previously, repeated exposure to an addictive good often leads to hedonic tolerance which implies that the marginal benefit of consumption is lower, not higher, at more advanced addictive states. Thus, only the BR model is able to incorporate the physiological process of tolerance.

[36] In the context of intertemporal risk preferences, this feature of the BR model (i.e., that it does not rely on adjacent complementarity) implies that addictive consumption can be generated with correlation averse preferences, as the discussion earlier made clear (see Bommier and Rochet (2006) and Lichtendahl, Chao and Bodily (2012)). To the extent that correlation aversion characterises most people's intertemporal risk preferences (see Andersen, Harrison, Lau and Rutström (2014b)), this is a noteworthy feature of the BR model.

The BR model generates different consumption patterns depending on characteristics of the individual, the addictive good and the environment; this is something which is neglected by all other theories. BR use simulations to show that consumption trajectories differ markedly for goods like caffeine, tobacco, alcohol and heroin and that these trajectories produce very different addiction prevalence rates in the population. For example, BR (2005, p. 120-125) show that plausible assumptions concerning the characteristics of tobacco and heroin yield far higher population prevalence rates of tobacco addiction than heroin addiction.

Benhabib and Bisin (2004) (BB) develop a model of internal commitment mechanisms and self-control which has a number of similar features to the BR model (i.e., temptations are stochastic and decisions are influenced by multiple decision-making processes: controlled and automatic) but is more general in scope. The BB model is of a single agent who faces temptation shocks and responds to them according to the interplay of expected future regret and present attention costs which are incurred when controlled processing is activated. Note that these temptation shocks are exogenous and undefined, and in this way, the BB model shares similarities with the exogenous addiction models discussed in the next section. In the BR model, by contrast, temptation shocks are specific to addictive goods and are more likely with a larger stock of addictive capital.

BB take as their starting point the literature on time-inconsistent choice behaviour and argue that to understand the choices of agents with time-inconsistent preferences it is imperative to model internal commitment mechanisms. This is similar in spirit to Ainslie (1992, 2001) and BT although BB's model does not assume H, nor QH, discounting as a primitive. As explained earlier, BB's model employs an E discount factor but the presence of stochastic temptations can produce time-inconsistent choices (viz., a present-bias).

An agent's choice behaviour in the BB model is driven by two processes: automatic processes which can be hijacked by temptations and controlled processes which are immune to these temptations. Automatic processes are essentially conditioned responses which are sensitive to environmental cues whereas controlled processes are responsible for long-term goals. Controlled processes either invoke cognitive

processes or inhibit automatic processes to fulfil these goals. An overarching executive function, or supervisory attention system, determines the extent to which these processes, automatic and controlled, are activated and influence behaviour. As alluded to previously, the executive system will only activate controlled processing if the expected future regret from excessive consumption exceeds the attention costs needed to suppress it.

An important feature of the model is that if an agent gives in to temptation, she will anticipate that this temptation will be followed by others in future, and the agent will adjust her long-term goals to incorporate the presence of these temptations; the model, therefore, derives similar predictions, albeit from a different starting point, to Ainslie (1992, 2001) and BT. In the context of addiction, an agent who succumbs to the temptation of consuming drugs may revise her long-term goals to incorporate the possibility that she will consume drugs in future which may lead to a self-fulfilling pattern of repeated drug use.

BB focus on the effect of increasing the presence of temptations on dynamic choice behaviour; this is accomplished through a first order stochastic dominance increase in the distribution of temptations. With a high volume of temptations, agents exercise self-control more often so as to inhibit automatic processes. However, they also tend to set less ambitious long-term goals because they expect, and take into account, the higher volume of temptations in future. Thus, even though automatic processes are inhibited more often, agents consume more under controlled processes than they would have if there was no increase in temptations. This example has an interesting parallel with enlistees in the Vietnam war, many of whom went to Vietnam never having tried opiates but then became addicted to heroin due to the widespread availability of the drug and a high level of drug-related temptations. As Robbins, Davis and Goodwin (1973) showed, the vast majority of heroin addicts managed to quit when they returned home to the US, one reason for which was the large decrease in the volume of temptations.

The notion of competing systems which jostle for control of behaviour has been the focus of neuroimaging studies conducted by McClure et al. (2004, 2007). These researchers claim to have found that parts of the brain which humans share with other

nonhumans (i.e., the reward system) respond only to the availability of SS rewards, whereas the pre-frontal cortex, which is uniquely developed in humans, responds to both SS and LL rewards. McClure et al. (2004, 2007) interpret this neuroimaging data as evidence of an internal tug of war between subpersonal agents which generates non-constant discounting at the behavioural level.[37] Although certainly a compelling idea, and one which gives credence to these synchronic models, the evidence for this hypothesis has not withstood rigorous replication (e.g., Glimcher, Kable and Louie (2007), Pine et al. (2009)).

Instantaneous and intertemporal risk preferences assume crucial roles in the BR and BB models. In the BB model, even though stochastic temptations are realised at the start of the decision making period, all future costs and benefits depend on the future stochastic realisation of temptations, and therefore hinge on the agent's intertemporal risk preferences. In the BR model, the presence of stochastic shocks which can trigger the hot mode may make someone who is risk averse with respect to atemporal payoffs avoid consumption of the addictive good altogether whereas someone who is less risk averse with respect to atemporal payoffs may be willing to gamble on experimentation. Clearly an agent's instantaneous risk preferences combine with their intertemporal risk preferences to affect every consumption decision, not just initial use, because deciding to consume the addictive good now increases the risk of entering the hot mode in future.

A Synchronic and Diachronic Model

Fudenberg and Levine (2006, 2011, 2012) (FL) develop a dual self model of impulse control which has numerous economic applications, many of which are relevant to the study of addiction. In the FL (2006, 2011) model, a patient, less risk averse LR self (i.e., a "planner") and a sequence of completely myopic, more risk averse SR selves (i.e., "doers") interact strategically to determine behaviour.

[37] This assertion is a strong form of reductionism which implies that behaviour at the molar level can be explained by dynamics at the molecular level. RSVS and Ross (2014) provide detailed critiques of reductionism in science generally, and with regard to the McClure et al. (2004, 2007) hypothesis, specifically.

The stage-game between the LR self and one intertemporal incarnation of the SR self plays out across two phases. In the first phase, the LR self chooses a self-control action that affects the utility function of the SR self. Specifically, the LR self can alter the baseline preferences of the SR self by incurring a self-control cost which is borne by both selves. In the second phase, the SR self makes a decision based on the preferences determined in the first phase. Although this interaction plays out across two hypothetical phases, it determines behaviour at a single moment in time and is best considered synchronic. The model is also diachronic because the LR self interacts with a sequence of SR selves over time.

FL (2006) apply their model to a set of economic problems: a simple savings model; a model of banking, commitment, and risk aversion; a stationary stopping-time problem; and a model of cognitive load and self-control. In the simple savings model, the saving rate is determined by the interplay of the LR self's time preference and the cost of self-control. As the agent becomes more present-oriented or as the cost of self-control rises, so the savings rate falls. The analogues with addictive consumption are immediate: as a person becomes more present-oriented, consumption of addictive goods will tend to rise, implying that investment in future health and welfare declines. Moreover, as the cost of self-control rises, perhaps due to a larger stock of addictive capital, exercising self-control becomes less likely, which then leads to an increase in addictive consumption and a reduction in welfare over time.

The model of banking, commitment, and risk aversion shows how the interaction between LR and SR selves can yield time-inconsistent choices. For example, suppose an agent is offered the choice between one unit of a good next period or two units of the good in two periods. As the SR self is completely myopic and the options are only available in the future, the LR self can influence behaviour at no cost and select the more rewarding utility stream. By contrast, if the agent is offered a choice between one unit of a good now, or two units of a good next period, the SR self will be inclined to consume one unit of the good now. To prevent this, the LR self would have to exercise self-control and this may be too costly. Thus, the agent, who prefers the LL reward when both options are delayed, switches to choosing the SS reward when it is available immediately, and thereby exhibits time inconsistency. In this

model, time inconsistency is driven by the myopia of the SR self, and is very similar, therefore, to models of QH discounting.

FL use a stopping-time problem to show how self-control problems can lead to procrastination and delay. In every period, the SR self must choose whether to act or wait: waiting yields a stochastic amount of utility, the value of which is known at the start of the period. Acting, by contrast, leads to a flow of utility starting next period, the present value of which is larger than the present value of waiting forever. If the agent waits, the problem repeats itself next period. Note that waiting confers a payoff in the present whereas acting only yields a payoff next period; this pits the interests of the SR selves in direct conflict with the LR self.

FL show that the solution to this problem is a cutoff rule: wait if the immediate payoff is above a threshold, act if the immediate payoff is below the threshold. Clearly this model is relevant to addiction, and, specifically, attempts to quit. To exercise self-control and act (e.g., stop using addictive goods) is costly in that one forgoes the payoff from waiting (e.g., consuming addictive goods). This is why there will be a tendency to procrastinate and delay. However, when the payoff from waiting drops below a threshold (perhaps due to the decline in welfare associated with advanced addictive states), acting becomes optimal. FL show that it is always eventually optimal to act but that waiting can occur for long periods of time. Thus, the FL model provides one explanation for why addicts, who often express a sincere desire to quit immediately and unconditionally, often procrastinate and delay before attempting to quit.

Finally, FL show how a heavy cognitive load can defeat attempts at self-control. To do so they adopt a nonlinear cost of self-control specification to capture the idea that self-control is a limited resource. As an agent's cognitive load gets heavier, it becomes increasingly costly to exercise self-control and, thus, self-control becomes less likely. This model could easily be reformulated or reinterpreted such that the cost of self-control is linked to the addictive state, rather than the level of cognitive load, and where higher addictive states make it more costly to exercise self-control. With a high enough addictive state, self-control can be so costly that it is not used, thereby perpetuating the tendency to consume addictive goods.

Andersen, Harrison, Lau and Rutström (2008) apply the FL (2006) model when jointly estimating the instantaneous risk and time preferences of a sample of 253 people in Denmark. The researchers' motivation for using FL's theory is to model the extent to which money earned in experiments is integrated with lifetime wealth. Rabin (2000) showed that if people integrate their experimental income with their lifetime wealth, then the risk aversion that is observed with small stakes in the lab yields implausible levels of risk aversion with larger stakes.[38] But to what extent is experimental income integrated with lifetime wealth? Applying the logic of the FL model, the extent to which this occurs will be determined by the interaction between an impulsive SR self and the more patient LR self. While the LR self wants to perfectly integrate experimental income with lifetime wealth, the SR self has diametrically opposed preferences, and Andersen, Harrison, Lau and Rutström (2008) show how to incorporate this conflict in the analysis of experimental data.

[38] Formally, Rabin's (2000) so-called "calibration critique" is based on four propositions: 1) economic agents are risk averse expected utility maximisers; 2) expected utility theory assumes full asset integration (i.e., the agent's utility function is defined over terminal wealth); 3) agents turn down small-stakes gambles at all wealth levels, or, at least, over a range of wealth levels; and 4) agents turn-down large-stakes gambles that have extremely favourable expected values. Rabin's (2000) primary claim is that if 1), 2) and 3) hold, then 4) follows. In other words, if risk averse expected utility maximisers, with utility functions defined over terminal wealth, display risk aversion over small-stakes gambles, then this yields implausible levels of risk aversion over large-stakes gambles. The implication that Rabin (2000) draws from this implausible large-stakes risk aversion is that 1) must, therefore, be false (i.e., economic agents do not satisfy expected utility theory). But, clearly, 4) will not hold if any of the three propositions which jointly imply 4) do not hold. Note there is nothing in the axioms of expected utility theory which implies that proposition 2) must hold. In other words, utility need not be defined over terminal wealth. Andersen et al. (2014) elicited the instantaneous risk preferences of a nationally-representative sample of people in Denmark to explore the asset integration hypothesis: specifically, whether the subjects integrate their experimental earnings with their wealth fully (i.e., utility is defined over terminal wealth), partially (i.e., utility is defined over wealth and experimental income and these two arguments of the utility function are not perfect substitutes), or not at all (i.e., utility is defined purely over experimental income). They then matched each subject's experimental data with detailed information on their individual wealth provided by Statistics Denmark. Using this uniquely rich combination of data sources, Andersen et al. (2014) find that partial, rather than full, asset integration best characterises their data, which thereby challenges proposition 2). Furthermore, with partial asset integration, the small-stakes risk aversion which the researchers observe is consistent with plausible levels of large-stakes risk aversion. Proposition 3) is also open to question: do people display small-stakes risk aversion over a relatively large range of wealth? Cox and Sadiraj (2008, p. 33) argue that proposition 3) could be tested using a within-subject experimental design where subjects are given some initial wealth w at the start of the experiment and then endowed, sequentially, with increasing amounts of income x_t such that the range over which the subjects' wealth varies during the course of the experiment is relatively large. Each time subjects are endowed with x_t they are asked whether they are willing to play a 50:50 gamble for, say, a gain of 11 and a loss of 10, where lab wealth is greater or equal to 10. Only if subjects reject the gamble for four *sequential* values of x_t, and thereby display risk aversion over a relatively large range of wealth, can one argue that proposition 3) holds. Until such an experiment is conducted, the validity of this proposition remains untested.

FL (2011) focus specifically on the implications of non-linear self-control costs in their model. The model still predicts time inconsistency but it also leads to an additional implication: preference reversals are less likely when the probability of receiving the SS and LL rewards declines. This result is based on the dichotomy between the SR and LR selves: the SR self is less patient and more risk averse than the LR self. When temptations are high, the cost of self-control is high, and more weight is given to the impatient, more risk averse SR self. When temptations are low, by contrast, the cost of self-control is low and little weight is given to the SR self. Thus, if the probability of receiving delayed rewards declines, the SS reward becomes less tempting, the cost of self-control declines, and the likelihood of selecting the SS reward falls. This implies that time-inconsistent preference reversals are less likely when the probability of receiving the SS and LL rewards declines; tentative evidence for this hypothesis is provided by the experiments of Keren and Roelofsma (1995).

FL (2012) extend the dual self model to account for some of their original model's implausible predictions regarding the timing of decisions. First, they assume that the SR self is not completely myopic but values future utility less than the LR self.[39] Second, FL formulate their model in discrete time but investigate how the solutions change with the length of decision-making periods. Finally, FL incorporate a stock of willpower which is depleted by using self-control. The evolution of this stock means that the use of self-control in one period can spill over to the next and thereby affect decision making.

FL (2012) explore the implications of their new model through a series of examples, a number of which are relevant to the study of addiction. To flesh out these examples they make use of an important conceptual tool: a *simple temptation*. A simple temptation is a choice between either utility 0 in every period or a flow of utility $u_g > 0$ for a set number of periods N, followed by the flow of utility $u_b < 0$ every period thereafter. A simple temptation pits the interests of the LR and SR selves against each other and is easy to extend to other types of temptation. For example, a *persistent temptation* is a simple temptation which is present each period unless or until it is accepted. Similarly, an *intermediate temptation* [my terminology] is intermediate

[39] To justify this assumption, the SR selves can be regarded as successive intertemporal incarnations with random lifetimes or as a single SR self who has a lower discount factor than the LR self.

between a simple and persistent temptation: resisting the temptation in period n removes the temptation until period $n + j$ where $j \in (1, \infty)$.

As discussed previously, Ainslie argues that temptations may be easier to resist when they are "bundled" together rather than tackled individually. To incorporate this idea in the dual self model, consider an agent who faces a set of simple temptations over time. If the agent is indifferent between accepting or rejecting the first simple temptation, she will strictly prefer to reject the second simple temptation, if that choice can be made now. The explanation for this result is that the LR self always prefers to resist the temptation and, if the temptation is far enough in the future, so too does the SR self. If we extend this logic, the agent who would accept a set of simple temptations presented one at a time, may choose to reject a bundle of these temptations. This bundle is easier to resist because the SR self is less tempted by future rewards than rewards available immediately. Thus, the FL (2012) model directly incorporates the potential for reward bundling.

Ainslie (2012) argues that H, but not E nor QH, discounting functions provide an obvious rationale for the formation of personal rules so as to promote reward bundling. The reason for this difference is that a H discounting function declines gradually at long delays and thus has higher tails than E or QH functions. Consequently, the value of a series of LL rewards grows relative to a series of SS rewards as the length of the series increases, which motivates the formation of personal rules to bundle rewards. The original dual self model (i.e., FL (2006, 2011)), and the QH discounting model, imply that after a one-time decline in value after the first period, incremental discount rates are time-invariant. By dropping the assumption of completely myopic SR selves, FL's dual self model incorporates the qualitative properties of a H discounting function (i.e., discount rates decline gradually at long delays), which provides a further impetus, over and above the one discussed previously, to bundle rewards.

In sum, FL model choice behaviour as the outcome of a strategic interaction between a patient, less risk averse LR self and a sequence of impatient, more risk averse SR selves. The SR selves are present-oriented and care most about immediate

consumption but the LR self can manipulate the preferences of the SR selves by exercising costly self-control. The interaction between the LR and SR selves can promote reward bundling to serve the agent's long-term interests but it can also lead to overconsumption, time inconsistency, and an inability to act. The FL (2012) model is arguably the most flexible of those reviewed in this section and it highlights the importance of uncertainty and the timing of decisions in choice behaviour.

The dual self and dual system models that I have covered divorce agency from biology and explain molar choice behaviour as the outcome of strategic interactions between subpersonal agents. All of the models incorporate time-inconsistent behaviour and thereby accommodate the empirical observation of addicts expending resources to consume their targets of addiction while simultaneously paying real costs to reduce or limit their consumption of these goods. The models also predict the typical course of addiction which involves a series of quit attempts and relapses prior to final abstention. The diachronic models emphasise the importance of time preferences in these dynamics whereas the synchronic models highlight the role which instantaneous and intertemporal risk preferences play. Thus, one should empirically investigate the instantaneous risk and time preferences, and intertemporal risk preferences, of addicts and non-addicts to determine whether they differ in the ways that these models suggest.

III.III EXOGENOUS ADDICTION MODELS

The economic models discussed in this section share the features that agency is tied to the whole biological life of a person and temptations are *exogenous*[40] to the agent and present a challenge to welfare maximisation (Ross (2010, 2012)). Ross (2012, p. 330) explains that these models, sometimes implicitly, make neural processing exogenous to the agent which, "… involves treating the person's 'true' agency as logically prior to all neural processing, that is, regarding parts of their brain as generating exogenous impacts on their choice and budget sets, just like features of the environment outside their skulls."

[40] Ainslie (2012, p. 21) uses slightly different terminology, arguing that these temptations arise "autonomously."

This family of models draws its inspiration from Loewenstein's (1996, 1999) *visceral* account of choice behaviour. The basic argument here is that only some types of rewards (i.e., visceral rewards like drugs and desserts) are discounted hyperbolically while others are not (i.e., non-visceral rewards like petrol and dishwashing liquid).[41] In other words, these visceral rewards can produce a strong drive or impulse for immediate consumption which overpowers other goals. And if utility maximisation entails "cold," deliberative processing, then visceral rewards are an "autonomous" impediment to this goal; in Ross' (2010, p. 147) words, these visceral rewards are, "… a threat to successful maximization lying in ambush in the 'external' environment within the brain." In the context of addiction, a genetic predisposition to substance dependence is archetypical of an exogenous impediment to utility maximisation.

Laibson (2001) develops a cue-theory of consumption which assumes that cues influence habit-forming behaviours by altering the marginal utility of a good which was previously consumed in the presence of the cue. In other words, cues associated with past consumption of a good are complementary with current consumption (i.e., cues raise the marginal utility of current consumption).

The model admits four steady states, two of which coincide with the "addiction" and "no addiction" steady states of the BM model in which cues play no role. The other two steady states are characterised by cue-contingent consumption. Specifically, choices are directly influenced by cues (e.g., cues for visceral rewards) even though the cues are just white noise which are not correlated with any other variables in the decision maker's choice problem. This captures an important feature of addiction: cue-triggered recidivism.

The model also implies that consumers are willing to spend significant resources to manage the cues that they receive so as to manage their susceptibility to cue-contingent consumption which can have deleterious effects on future welfare (e.g., addictive consumption). Thus, the model can reconcile the empirical phenomenon of

[41] Loewenstein (1996, p. 273) argues that visceral factors have two characteristics: a direct hedonic impact; and an effect on the relative valuation of other goods and actions.

people simultaneously expending money on a good and taking active steps to try and control or prevent their consumption of this good.

Despite the fact that consumers have time-consistent preferences in Laibson's model, there is a role for commitment devices which limit the choice set of agents. This result arises when cues and consumption opportunities are linked such that to consume a good, a cue for consumption of that good must be present. When cues and consumption possibilities are linked, a large choice set implies a large number of cues and consumers who want to limit their exposure to cue-contingent consumption may choose to limit their choice set.

Laibson's model also incorporates the possibility of non-constant discounting (viz., time inconsistency) through the linkage of cues and rewards. The idea is that if a cue is paired with a reward, this may make the reward irresistible even though the consumer could receive a larger quantity of the reward by delaying consumption. In other words, the consumer has a passion for the present which is driven by the cue. By contrast, when a cue is present but the reward is only available in the future, the consumer can trade off the benefits and costs of consumption without the cue driving choice in a particular direction. Thus, people may simultaneously display a passion for the present and patience over consumption opportunities available in the future (viz., non-constant discounting).

Laibson's model is similar to BB and BR in that stochastic shocks affect choice behaviour. But unlike BB and BR, these shocks do not produce an internal tug of war between competing systems. Instead, a unified agent faces external shocks that impede successful utility maximisation. As Laibson (2001, p. 109) argues, a storm of visceral temptations implies that, "… behavior changes rapidly from moment to moment, temptation can/should be actively avoided, and public consumption can be a negative externality."

Loewenstein, O'Donoghue and Rabin (2003) (LOR) develop a model of projection bias which assumes that people tend to understand the way in which their tastes will change over time but struggle to accurately predict the magnitude of these changes.

Consequently people tend to project their current preferences into the future without adequately accounting for the extent to which tastes will change.

These errors in prediction, which cause predicted and realised utilities to diverge, imply that behaviour may not correspond to intertemporal utility maximisation with perfect foresight; projection bias is like a cognitive deficit which is "external" to a person's "true" agency. This failure of intertemporal utility maximisation can have particularly problematic effects in the case of habit forming behaviours like addiction where current consumption of the addictive good raises the marginal utility of future consumption while simultaneously leading to a fall in future welfare through increases in the stock of addictive capital. With projection bias, people may fail to realise, or may underappreciate, the extent to which current consumption affects future welfare, leading them to consume more of the good in the present than they would with perfect foresight. Furthermore, if tastes change over time but people fail to predict the magnitude of these changes, they may make plans which they fail to carry out. Thus, projection bias can lead to time inconsistency. Note that LOR assume that decision makers are completely unaware of this potential for time inconsistency and, thus, the agents are like the naïve QH discounters in O'Donoghue and Rabin (1999).

LOR argue that economic models of addiction capture many of the properties of substance dependence but often fail to explain how people became addicted to goods in the first place. Projection bias provides two possible and related explanations. If people fail to predict the future negative consequences of addictive consumption they may consume too much of the addictive good in the present. In addition, if people fail to predict the extent to which current consumption alters the marginal utility of future consumption they will underappreciate the power of habit formation. Consequently people may intend to consume an addictive good now without doing so in future, but without accurately predicting the influence of habit formation may then revise their plans in the future and continue to consume the addictive good.

LOR also highlight the role that day-to-day fluctuations in tastes have on a particular aspect of addiction: craving. For example, if someone's craving for a drug is particularly high on one day and this is projected on to the future, they may be

unwilling to quit because they incorrectly predict how difficult it will be to do so. Conversely, when a person's craving for an addictive good is low, they may overestimate the ease of quitting and make repeated attempts to quit which ultimately fail when cravings return to normal levels.

The model of projection bias adds to our understanding of addiction by highlighting the role that expectations of future utility have on current choices and how prediction errors can make addiction more likely and more intractable. An interesting difference between the models of LOR and BT is that the former focuses on errors in prediction while the latter focuses on errors in recall. In LOR, agents struggle to correctly forecast the effect of current actions on future welfare and may therefore take actions with immediate benefits without appropriately internalising their future costs. In BT, imperfect recall means that agents may have to continually "prove" to themselves that they have and can exercise willpower, a process that can unravel through transgression and precipitate addiction.

Gul and Pesendorfer (2007) (GP) develop an axiomatic theory of addiction which builds on their earlier work on models of temptation and self-control (GP (2001, 2004)). In this framework a person's welfare depends on the choices that she makes as well as the set of options from which choices are made. The key idea here is that the set of options may contain tempting alternatives that lower an agents' utility either by skewing her choice, requiring the use of costly self-control, or both. With regard to addiction, consumption of an addictive good is a tempting option which undermines self-control in the future. Thus, people are more likely to give in to temptation, the greater their past consumption of the addictive good.

This logic is formalised through the notion of "compulsive consumption" (GP, p. 154). A choice is compulsive if it differs to what would have been chosen if commitment were available. A good is addictive if consumption now leads to more compulsive consumption in future. Thus, addiction is defined as an increasing divergence between what is chosen and what would have been chosen absent temptation. Note that temptations here are ever-present, external shocks to the agent's welfare which must be combatted, if possible, through the use of commitment devices or costly self-control.

In the GP model, an agent's preferences are defined over the set of decision problems rather than the set of consumption bundles. This relatively large domain of preferences allows GP to incorporate the ideas of commitment, self-control and temptation in their model. For example, a person may strictly prefer decision problem *A* to decision problem *B* because the latter includes tempting alternatives that are difficult to resist. Even though this person may make the identical choice under each decision problem, she may strictly prefer decision problem *A* because making the same choice under *B* requires costly self-control. Thus, the person can benefit by committing to decision problem *A* at the outset.

Unlike GP's earlier work (GP (2001, 2004)), preferences in this model depend on the person's consumption history. Specifically, only consumption of the addictive good in the last period affects the person's current preferences and this is captured by the addictive state: the greater the agent's past consumption of the addictive good the higher the addictive state. Note that the cost of self-control increases at higher addictive states, implying that it becomes increasingly difficult to resist temptation.

Thus, if someone resists temptation now, it will be easier to do so in the next period because the addictive state and hence the cost of self-control will decline. By contrast, if someone succumbs to temptation in this period it will be harder to resist temptation next period because the addictive state rises and with it, the cost of self-control. Clearly, these intertemporal trade-offs are influenced by the person's time preferences. Someone who discounts future consequences heavily will downplay the effect of an increase in addictive consumption now on the difficulty of resisting temptation in future. Such a person will tend to consume more of the addictive good now which implies more compulsive consumption in future and ultimately, addiction.[42]

A shortcoming of the GP model is that there is no place for inconsistent choice behaviour. Choice in this model represents the ideal trade-off between the utility

[42] GP's earlier models (GP (2001, 2004)) incorporate risk and uncertainty but their model of addiction is defined solely over deterministic decision problems. Consequently, instantaneous and intertemporal risk preferences play no role in this framework.

gained from consumption of non-addictive goods and the self-control costs which must be incurred to limit consumption of addictive goods. Although agents desire commitment devices, this is not the product of time inconsistency. Instead, agents demand commitment devices so as to remove temptations, decrease the cost of self-control, and increase welfare.

Another issue with the model is that it assumes only consumption of the addictive good in the previous period affects current preferences. What this implies is that if someone abstains successfully, perhaps with the help of a commitment device, then all previous effects of consumption of the addictive good wear off completely by the subsequent period. In other words, even with extreme exposure to an addictive good a person *will* return to a "virgin" state through successful abstention where it is as if they never consumed the good. This is a difficult assumption to justify given that the effects from extreme exposure to an addictive good tend to persist.

The models discussed in this section link agency to the whole biological life of a person and regard deficits in human intelligence or reasoning as separate to, and distinct from, the person's "true" agency, and which should be modelled as exogenous influences on the agent's decision-making. Cues in Laibson's framework are external taste shocks which must be resisted where possible; projection bias is a cognitive deficit which can have problematic effects by causing predicted and realised utilities to diverge; and the exogenous temptations in GP's model are a threat to utility maximisation because they necessitate the use of costly self-control. These models capture many of the salient features of addiction and do so from a completely different starting point to the dual self and dual system models discussed previously.

In sum, the economic models of addiction reviewed in this chapter find common ground on the importance of instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs in the explanation of addiction. Thus, these preferences are the logical starting point for empirical analyses of addiction.

IV. CONCLUSION

In this chapter I provided an overview of the history of substance use and abuse, discussed current trends in drug consumption and dependence, and focussed on the costs that addiction imposes on society. This broad overview made it clear that addiction is a pressing public health concern and motivates research into the causes, course, and consequences of substance dependence.

I then focussed on the diagnostic classification of addiction to highlight the view that classification systems do not solely serve scientific ends but that they develop in concert with wider political, social, and economic forces. Despite some issues with measurement and interpretation, research into substance dependence has made progress over the last 50 years by refining the definition of, and diagnostic criteria for, addiction.

As a prelude to my discussion of economic models of addiction, I traced the history of the disease concept of substance dependence. I also outlined Heyman's (2009) critique of the disease model to provide an alternative account of the nature of addiction but ultimately relied on the reconciliation of RSVS which shows that these apparently contradictory hypotheses are ultimately compatible if one adopts the appropriate levels of scientific analysis.

In the second half of the chapter I reviewed a set of economic models which are relevant to the study of addiction so as to highlight the role that instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs play in substance dependence. The models that I covered were organised into three categories: models of rational addiction; dual self and dual system models; and exogenous addiction models.

The models of rational addiction tie agency to the whole biological life of a person and assume time consistency. These models show how the interplay of instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs can precipitate addiction. OZ (p. 741) provide a succinct statement of this point when they argue that people often experiment with drugs because they, "… provide certain

instant pleasure and only probabilistic future harm." While these models manage to capture some of the salient features of addiction they ultimately founder by mis-predicting the typical course of addiction; the models cannot account for addicts' inconsistent choice behaviour and the pattern of quit attempts and relapse.

Dual self and dual system models assume that strategic interactions among subpersonal agents determine molar level choice behaviour. When these subpersonal agents are divided synchronically, they become a community of agents. When they are divided diachronically, they become a sequence of agents (Ross (2012, p. 31)). All of the dual self models which I covered incorporate time-inconsistent choice behaviour but do so through different mechanisms. GK and BT assume QH discounting and thereby link inconsistency directly to discounting behaviour. BB and BR, by contrast, make use of standard E discounting but assume that stochastic shocks lead to an internal tug of war between competing systems which often produces time-inconsistent choices at the molar level. FL adopt a synchronic and diachronic subpersonal framework and model inconsistent choice as the outcome of interactions between a patient, less risk averse LR self and an impatient, more risk averse SR self. The FL (2012) model captures the qualitative features of H discounting which provides a rationale for Ainslie's notion of reward bundling.

The exogenous addiction models link agency to the whole biological life of a person and view temptations as exogenous to the agent and an impediment to successful utility maximisation. These models are inspired by Loewenstein's *visceral* factors which can overwhelm "cold," deliberative reasoning and drive choice in the direction of immediate gratification. These models show how cues for consumption, prediction errors, and temptations can affect an agent's intertemporal trade-offs and generate behavioural choice patterns that share a number of similar features of addiction.

The models as a group highlight the roles that instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs play in addiction.[43] From a methodological standpoint the question now turns to how best to measure

[43] As discussed previously, I do not explore potential differences in the intertemporal risk preferences and subjective beliefs of addicts and non-addicts in this thesis, but defer these important areas of enquiry for future research.

instantaneous risk and time preferences, particularly considering that the models incorporate different discounting specifications and place more or less emphasis on the roles that risk and uncertainty play. In Chapter 3, I provide a comprehensive review of the literature on time preferences and smoking behaviour to discuss the different experimental tools which have been used to elicit these preferences, and to draw attention to their benefits and drawbacks. In Chapter 4, I critique the probability discounting model which has been used extensively in psychology and addiction studies to measure instantaneous risk attitudes. In Chapter 5, I review alternative approaches to the experimental elicitation of instantaneous risk preferences and discuss a battery of risk preference questions which were developed to allow both utility function curvature and probability weighting to determine atemporal attitudes toward risk. I also present a set of time preference questions which were carefully constructed to reliably estimate alternative discounting models.

From a statistical perspective, economic models of addiction suggest that one should adopt a flexible estimation framework which can incorporate different instantaneous risk and time preference specifications. In other words, one should not be wedded to a particular discounting function or atemporal risk preference theory because these not only differ across the economic models that I reviewed but may also differ across people. Thus, it is crucial to employ a flexible statistical framework to draw accurate inferences about potential differences between addicts and non-addicts. In Chapter 3, I present a statistical approach to the analysis of experimental data which incorporates E, H, QH, and Weibull discounting specifications and which allows multiple data generating processes to account for the data. In Chapter 4, I extend this approach to the study of instantaneous risk preferences and in Chapter 5, I combine the insights from the previous chapters in a joint estimation framework of instantaneous risk and time preferences. It is these methodological and statistical issues to which I now turn.

REFERENCES

AGRAWAL, A., AND M. T. LYNSKEY (2008): "Are There Genetic Influences on Addiction: Evidence from Family, Adoption and Twin Studies," *Addiction*, 103, 1069-81.

AINSLIE, G. (1992): *Picoeconomics: The Interaction of Successive Motivational States within the Individual*. Cambridge: Cambridge University Press.

AINSLIE, G. (2001): *Breakdown of Will*. Cambridge: Cambridge University Press.

AINSLIE, G. (2005): "Precis of Breakdown of Will," *Behavioral and Brain Sciences*, 28, 635-50; discussion 650-73.

AINSLIE, G. (2012): "Pure Hyperbolic Discount Curves Predict "Eyes Open" Self-Control," *Theory and Decision*, 73, 3-34.

AKERS, R. (1991): "Addiction: The Troublesome Concept," *Journal of Drug Issues*, 21, 777-793.

AMERICAN PSYCHIATRIC ASSOCIATION (1952): *Diagnostic and Statistical Manual of Mental Disorders*. Washington, DC: American Psychiatric Association.

AMERICAN PSYCHIATRIC ASSOCIATION (1968): *Diagnostic and Statistical Manual of Mental Disorders, Second Edition (DSM-II)*. Washington, DC: American Psychiatric Association.

AMERICAN PSYCHIATRIC ASSOCIATION (1980): *Diagnostic and Statistical Manual of Mental Disorders, Third Edition (DSM-III)*. Washington, DC: American Psychiatric Association.

AMERICAN PSYCHIATRIC ASSOCIATION (1994): *Diagnostic and Statistical Manual of Mental Disorders, Fourth Edition (DSM-IV)*. Washington, DC: American Psychiatric Association.

AMERICAN PSYCHIATRIC ASSOCIATION (2013): *Diagnostic and Statistical Manual of Mental Disorders, Fifth Edition (DSM-5)*. Washington, DC: American Psychiatric Association.

ANDERSEN, S., J. C. COX, G. W. HARRISON, M. I. LAU, E. E. RUTSTRÖM, AND V. SADIRAJ (2014): "Asset Integration and Attitudes to Risk: Theory and Evidence," Working Paper WP 2011-17, Center for the Economic Analysis of Risk (CEAR) Georgia State University.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583-618.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2014a): "Dual Criteria Decisions," *Journal of Economic Psychology*, 41, 101-113.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2014b): "Multiattribute Utility Theory, Intertemporal Utility and Correlation Aversion," Working Paper WP 2011-03, Center for the Economic Analysis of Risk (CEAR) Georgia State Univerity.

ANTHONY, J. C., AND J. E. HELZER (1991): "Syndromes of Drug Abuse and Dependence," in *Psychiatric Disorders in America: The Epidemiologic Catchment Area Study*, ed. by L. N. Robins, and D. A. Regier. New York: Free Press, 116-154.

BALTAGI, B. H. (2007): "On the Use of Panel Data Methods to Estimate Rational Addiction Models," *Scottish Journal of Political Economy*, 54, 1-18.

BECKER, G. S., AND C. B. MULLIGAN (1997): "The Endogenous Determination of Time Preference," *Quarterly Journal of Economics*, 112, 729-758.

BECKER, G. S., AND K. M. MURPHY (1988): "A Theory of Rational Addiction," *Journal of Political Economy*, 96, 675-700.

BÉNABOU, R., AND J. TIROLE (2004): "Willpower and Personal Rules," *Journal of Political Economy*, 112, 848-886.

BENHABIB, J., AND A. BISIN (2004): "Modeling Internal Commitment Mechanisms and Self-Control: A Neuroeconomics Approach to Consumption-Saving Decisions," *Games and Economic Behavior*, 52, 460-492.

BERNHEIM, B. D., AND A. RANGEL (2004): "Addiction and Cue-Triggered Decision Processes," *American Economic Review*, 94, 1558 - 1590.

BERNHEIM, B. D., AND A. RANGEL (2005): "From Neuroscience to Public Policy: A New Economic View of Addiction," *Swedish Economic Policy Review*, 12, 99-144.

BERNHEIM, B. D., AND A. RANGEL (2007): "Behavioral Public Economics: Welfare and Policy Analysis with Non-Standard Decision-Makers," in *Behavioral Economics and Its Applications*, ed. by P. Diamond, and H. Vartiainen. Princeton: Princeton University Press, 7 - 81.

BODNER, R., AND D. PRELEC (2001): "The Diagnostic Value of Actions in a Self-Signaling Model," in *Collected Essays in Psychology and Economics*, ed. by I. Brocas, and J. D. Carillo. Oxford: Oxford University Press.

BOMMIER, A. (2007): "Risk Aversion, Intertemporal Elasticity of Substitution and Correlation Aversion," *Economics Bulletin*, 4, 1-8.

BOMMIER, A., AND J. ROCHET (2006): "Risk Aversion and Planning Horizons," *Journal of the European Economic Association*, 4, 708-734.

CAILLAUD, B., AND B. JULLIEN (2000): "Modelling Time-Inconsistent Preferences," *European Economic Review*, 44, 1116-1124.

CHABRIS, C., D. LAIBSON, C. L. MORRIS, J. P. SCHULDT, AND D. TAUBINSKY (2008): "Individual Laboratory-Measured Discount Rates Predict Field Behavior," *Journal of Risk and Uncertainty*, 37, 237-269.

CHAIKEN, S., AND Y. TROPE (1999): *Dual Process Theories in Social Psychology*. New York, NY: Guilford Press.

CHALOUPKA, F., AND K. WARNER (2001): "The Economics of Smoking," in *Handbook of Health Economics*, ed. by J. Newhouse, and D. Cutler. Amsterdam: North-Holland.

COLLINS, D. J., AND H. M. LAPSLEY (2008): *The Costs of Tobacco, Alcohol and Illicit Drug Abuse to Australian Society in 2004/2005*. Australia: The Department of Health.

COX, J. C., AND V. SADIRAJ (2008): "Risky Decisions in the Large and in the Small: Theory and Experiment," in *Research in Experimental Economics: Volume 12. Risk Aversion in Experiments.*, ed. by J. C. Cox, and G. W. Harrison. Bingley: Emerald, 9-46.

DE FINETTI, B. (1952): "Sulla Preferibilitia," *Giornal degli Economisti e Annali di Economia*, 11, 685-709.

EECKHOUDT, L., B. REY, AND H. SCHLESINGER (2007): "A Good Sign for Multivariate Risk Taking," *Management Science*, 53, 117-124.

ELSTER, J. (1979): *Ulysses and the Sirens: Studies in Rationality and Irrationality*. Cambridge: University of Cambridge Press.

ELSTER, J. (1997): "More Than Enough," *The University of Chicago Law Review*, 64, 749-764.

EPSTEIN, L. G., AND S. M. TANNY (1980): "Increasing Generalized Correlation: A Definition and Some Economic Consequences," *Canadian Journal of Economics*, 13, 16-34.

FERGUSON, B. S. (2000): "Interpreting the Rational Addiction Model," *Health Economics*, 9, 587-598.

FISHBURN, P. C. (1984): "Multiattribute Nonlinear Utility Theory," *Management Science*, 30, 1301-1310.

FUDENBERG, D., AND D. K. LEVINE (2006): "A Dual-Self Model of Impulse Control," *American Economic Review*, 96, 1449-1476.

FUDENBERG, D., AND D. K. LEVINE (2011): "Risk, Delay and Convex Self-Control Costs," *American Economic Journal: Microeconomics*, 3, 34-68.

FUDENBERG, D., AND D. K. LEVINE (2012): "Timing and Self-Control," *Econometrica*, 80, 1-42.

GARDNER, E., AND J. DAVID (1999): "The Neurobiology of Chemical Addiction," in *Getting Hooked: Rationality and Addiction*, ed. by J. Elster, and O. Skog. Cambridge: Cambridge University Press, 93-136.

GLIMCHER, P. W., J. KABLE, AND K. LOUIE (2007): "Neuroeconomic Studies of Impulsivity: Now or Just as Soon as Possible?," *American Economic Review*, 97, 142-147.

GOLDSTEIN, A. (2001): *Addiction: From Biology to Drug Policy*. New York: Oxford University Press.

GORDON, L., L. TINSLEY, C. GODFREY, AND S. PARROTT (2006): "The Economic and Social Costs of Class A Drug Use in England and Wales, 2003/04," in *Measuring Different Aspects of Problem Drug Use: Methodological Developments*, ed. by N. Singleton, R. Murray, and L. Tinsley. London: Home Office.

GREEN, L., AND J. MYERSON (2004): "A Discounting Framework for Choice with Delayed and Probabilistic Rewards," *Psychological Bulletin*, 130, 769-92.

GRUBER, J., AND B. KOSZEGI (2001): "Is Addiction "Rational"? Theory and Evidence," *Quarterly Journal of Economics*, 116, 1261 - 1303.

GUL, F., AND W. PESENDORFER (2001): "Temptation and Self-Control," *Econometrica*, 69, 1403-1435.

GUL, F., AND W. PESENDORFER (2004): "Self-Control and the Theory of Consumption," *Econometrica*, 72, 119-158.

GUL, F., AND W. PESENDORFER (2007): "Harmful Addiction," *The Review of Economic Studies*, 74, 147-172.

HANSON, G. R., A. I. LESHNER, AND B. TAI (2002): "Putting Drug Abuse Research to Use in Real-Life Settings," *Journal of Substance Abuse Treatment*, 23, 69-70.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2010): "Individual Discount Rates and Smoking: Evidence from a Field Experiment in Denmark," *Journal of Health Economics*, 29, 708-717.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2014): "Theory, Experimental Design and Econometrics Are Complementary (and So Are Lab and Field Experiments)." in *The Methods of Modern Experimental Economics*, ed. by G. Frechette, and A. Schotter. New York, NY: Oxford University Press.

HEYMAN, G. M. (2009): *Addiction: A Disorder of Choice*. Cambridge, MA: Harvard University Press.

INTERNATIONAL OPIUM COMMITTEE (1909): *Report of the International Opium Commission, Shanghai, China, February 1 to February 26 1909, Vol. II, Reports of the Delegations*. Shanghai: North-China Daily News & Herald Ltd.

JOHNSEN, T. H., AND J. B. DONALDSON (1985): "The Structure of Intertemporal Preferences under Uncertainty and Time Consistent Plans," *Econometrica*, 53, 1451-1458.

KANG, M. I., AND S. IKEDA (2014): "Time Discounting and Smoking Behavior: Evidence from a Panel Survey," *Health Economics*, 23, 1443-1464.

KEENEY, R. L. (1973): "Risk Independence and Multiattributed Utility Functions," *Econometrica*, 41, 27-34.

KEENEY, R. L., AND H. RAIFFA (1976): *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York, NY: Wiley.

KEREN, G., AND P. ROELOFSMA (1995): "Immediacy and Certainty in Intertemporal Choice," *Organizational Behavior and Human Decision Processes*, 63, 287-297.

KESSLER, R. C., P. BERGLUND, O. DEMLER, R. JIN, K. R. MERIKANGAS, AND E. E. WALTERS (2005a): "Lifetime Prevalence and Age-of-Onset Distributions of DSM-IV Disorders in the National Comorbidity Survey Replication," *Archives of General Psychiatry*, 62, 593-602.

KESSLER, R. C., W. T. CHIU, O. DEMLER, K. R. MERIKANGAS, AND E. E. WALTERS (2005b): "Prevalence, Severity, and Comorbidity of 12-Month DSM-IV Disorders in the National Comorbidity Survey Replication," *Archives of General Psychiatry*, 62, 617-27.

KRUGER, J., AND D. DUNNING (1999): "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments," *Journal of Personality and Social Psychology*, 77, 1121-1134.

LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443-477.

LAIBSON, D. (2001): "A Cue-Theory of Consumption," *Quarterly Journal of Economics*, 116, 81-119.

LESHNER, A. I. (1997): "Addiction Is a Brain Disease, and It Matters," *Science*, 278, 45-7.

LEVINE, H. G. (1978): "The Discovery of Addiction. Changing Conceptions of Habitual Drunkenness in America," *Journal of Studies on Alcohol*, 39, 143-74.

LI, M. D. (2003): "The Genetics of Smoking Related Behavior: A Brief Review," *The American Journal of the Medical Sciences*, 326, 168-73.

LI, M. D., AND M. BURMEISTER (2009): "New Insights into the Genetics of Addiction," *Nature Reviews Genetics*, 10, 225-31.

LICHTENDAHL, K. C., R. O. CHAO, AND S. E. BODILY (2012): "Habit Formation from Correlation Aversion," *Operations Research*, 60, 625-637.

LOEWENSTEIN, G. (1996): "Out of Control: Visceral Influences on Behavior," *Organizational Behavior and Human Decision Processes*, 65, 272-292.

LOEWENSTEIN, G. (1999): "A Visceral Account of Addiction," in *Getting Hooked: Rationality and Addiction*, ed. by J. Elster, and O. Skog. Cambridge: Cambridge University Press.

LOEWENSTEIN, G., T. O'DONOGHUE, AND M. RABIN (2003): "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics*, 118, 1209-1248.

LOPES, L. L. (1984): "Risk and Distributional Inequality," *Journal of Experimental Psychology: Human Perception and Performance*, 10, 465-485.

LOPES, L. L. (1995): "Algebra and Process in the Modeling of Risky Choice," in *Decision Making from a Cognitive Perspective*, ed. by J. R. Busemeyer, R. Hastie, and D. L. Medin. San Diego, CA: Academic Press.

MAZUR, J. E. (1984): "Tests of an Equivalence Rule for Fixed and Variable Reinforcer Delays," *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 426-437.

MCCLURE, S. M., K. M. ERICSON, D. I. LAIBSON, G. LOEWENSTEIN, AND J. D. COHEN (2007): "Time Discounting for Primary Rewards," *Journal of Neuroscience*, 27, 5796-804.

MCCLURE, S. M., D. I. LAIBSON, G. LOEWENSTEIN, AND J. D. COHEN (2004): "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 306, 503-7.

MIYAMOTO, J. M., AND P. P. WAKKER (1996): "Multiattribute Utility Theory without Expected Utility Foundations," *Operations Research*, 44, 313-326.

MONTEROSSO, J., AND G. AINSLIE (2009): "The Picoeconomic Approach to Addictions: Analysing the Conflict of Successive Motivational States," *Addiction Research and Theory*, 17, 115-134.

NATIONAL DRUG INTELLIGENCE CENTER (2011): *The Economic Impact of Illicit Drug Use on American Society*. Washington, DC: United States Department of Justice.

NATIONAL INSTITUTE ON ALCOHOL ABUSE AND ALCOHOLISM. (2014): *Alcohol Use Disorders*. Retrieved 6 April, 2014, from http://www.niaaa.nih.gov/alcohol-health/overview-alcohol-consumption/alcohol-use-disorders

NATIONAL INSTITUTE ON DRUG ABUSE. (2012): *Drugfacts: Understanding Drug Abuse and Addiction*. Retrieved 16 April, 2014, from http://www.drugabuse.gov/publications/drugfacts/understanding-drug-abuse-addiction

O'DONOGHUE, T., AND R. RABIN (1999): "Doing It Now or Later," *American Economic Review*, 89, 103-124.

ORPHANIDES, A., AND D. ZERVOS (1995): "Rational Addiction with Learning and Regret," *Journal of Political Economy*, 103, 739-758.

PHELPS, E. S., AND R. A. POLLAK (1968): "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 35, 185-199.

PINE, A., B. SEYMOUR, J. P. ROISER, P. BOSSAERTS, K. J. FRISTON, H. V. CURRAN, AND R. J. DOLAN (2009): "Encoding of Marginal Utility across Time in the Human Brain," *Journal of Neuroscience*, 29, 9575-81.

PRENDERGAST, M., D. PODUS, J. FINNEY, L. GREENWELL, AND J. ROLL (2006): "Contingency Management for Treatment of Substance Use Disorders: A Meta-Analysis," *Addiction*, 101, 1546-60.

RABIN, M. (2000): "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68, 1281-1292.

REHM, J., D. BALIUNAS, S. BROCHU, B. FISCHER, W. GNAM, J. PATRA, S. POPOVA, A. SARNOCINSKA-HART, AND B. TAYLOR (2006): *The Costs of Substance Abuse in Canada 2002: Highlights*. Ottawa, ON: Canadian Centre on Substance Abuse.

RICHARD, S. F. (1975): "Multivariate Risk Aversion, Utility Independence and Separable Utility Functions," *Management Science*, 22, 12-21.

ROBINS, L. N., D. H. DAVIS, AND D. W. GOODWIN (1973): "Drug Use by U.S. Army Enlisted Men in Vietnam: A Follow-up on Their Return Home," *American Journal of Epidemiology*, 99, 235-249.

ROGEBERG, O. (2004): "Taking Absurd Theories Seriously: Economics and the Case of Rational Addiction Theories," *Philosophy of Science*, 71, 263-285.

ROSS, D. (2005): *Economic Theory and Cognitive Science: Microexplanation*. Cambridge, MA: MIT Press.

ROSS, D. (2010): "Economic Models of Pathological Gambling," in *What Is Addiction?*, ed. by D. Ross, H. Kincaid, D. Spurrett, and P. Collins. Cambridge, MA: MIT Press, 131-158.

ROSS, D. (2012): "Addictive, Impulsive and Other Counter-Normative Consumption," in *Handbook of Developments in Consumer Behaviour*, ed. by V. K. Wells, and G. R. Foxall. Cheltenham: Edward Elgar Publishing.

ROSS, D. (2014): *Philosophy of Economics*. Basingstoke: Palgrave Macmillan.

ROSS, D., C. SHARP, R. VUCHINICH, AND D. SPURRETT (2008): *Midbrain Mutiny: The Behavioural Economics and Neuroeconomics of Disordered Gambling*. Cambridge, MA: MIT Press.

SAMUELSON, P. (1937): "A Note on Measurement of Utility," *Review of Economic Studies*, 4, 155-161.

SCHELLING, T. (1984): *Choice and Consequence*. Cambridge, MA: Harvard Univerity Press.

SPINNEWYN, F. (1981): "Rational Habit Formation," *European Economic Review*, 15, 91-109.

STANOVICH, K. E., AND R. F. WEST (2000): "Individual Differences in Reasoning: Implications for the Rationality Debate?," *Behavioral and Brain Sciences*, 23, 645-726.

STIGLER, G. J., AND G. S. BECKER (1977): "De Gustibus Non Est Disputandum," *American Economic Review*, 67, 76-90.

STINSON, F. S., B. F. GRANT, D. A. DAWSON, W. J. RUAN, B. HUANG, AND T. SAHA (2005): "Comorbidity between DSM-IV Alcohol and Specific Drug Use Disorders in the United States: Results from the National Epidemiologic Survey on Alcohol and Related Conditions," *Drug and Alcohol Dependence*, 80, 105-16.

STROTZ, R. H. (1955-1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165-180.

STRZALECKI, T. (2013): "Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion," *Econometrica*, 81, 1039-1074.

THALER, R. H., AND H. M. SHEFRIN (1981): "An Economic Theory of Self-Control," *Journal of Political Economy*, 89, 392-406.

UNITED NATIONS OFFICE ON DRUGS AND CRIME (UNODC) (2012): *World Drug Report 2012*. United Nations publication: Sales No. E.12.XI.1.

VUCHINICH, R., AND N. HEATHER (2003): *Choice, Behavioral Economics, and Addiction*. Oxford: Pergamon/Elsevier.

WARNER, J. (1994): ""Resolv'd to Drink No More": Addiction as a Preindustrial Construct," *Journal of Studies on Alcohol*, 55, 685-91.

WARNER, L. A., R. C. KESSLER, M. HUGHES, J. C. ANTHONY, AND C. B. NELSON (1995): "Prevalence and Correlates of Drug Use and Dependence in the United States. Results from the National Comorbidity Survey," *Archives of General Psychiatry*, 52, 219-29.

WEINER, B., AND W. WHITE (2007): "The Journal of Inebriety (1876-1914): History, Topical Analysis, and Photographic Images," *Addiction*, 102, 15-23.

WEST, R. (2006): *Theory of Addiction*. Oxford: Blackwell Publishing.

WESTERMEYER, J. (1987): "Cultural Patterns of Drug and Alcohol Use: An Analysis of Host and Agent in the Cultural Environment," *Bulletin on Narcotics*, 39, 11-27.

WESTERMEYER, J. (2005): "Historical and Social Context of Psyhoactive Substance Use Disorders," in *Clinical Textbook of Addictive Disorders*, ed. by R. J. Frances, S. I. Miller, and A. H. Mack. New York, NY: The Guilford Press, 16-34.

WINSTON, G. C. (1980): "Addiction and Backsliding: A Theory of Compulsive Consumption," *Journal of Economic Behavior and Organization*, 1, 295-324.

WORLD BANK (1992): *World Development Report 1992: Development and the Environment*. New York, NY: Oxford University Press.

WORLD BANK (1993): *World Development Report 1993: Investing in Health*. New York, NY: Oxford University Press.

WORLD BANK (2003): *World Development Report 2003: Sustainable Development in a Dynamic World - Transforming Institutions, Growth, and Quality of Life*. World Bank: World Bank.

WORLD BANK (2010): *World Development Report 2010: Development and Climate Change*. Washington, DC: World Bank.

WORLD HEALTH ORGANIZATION (WHO) (1994): *The ICD-10 Classification of Mental and Behavioural Disorders: Clinical Descriptions and Diagnostic Guidelines*. Geneva: World Health Organisation.

WORLD HEALTH ORGANIZATION (WHO) (2009): *Global Health Risks: Mortality and Burden of Disease Attributable to Selected Major Risks*. Geneva: World Health Organization.

WORLD HEALTH ORGANIZATION (WHO) (2010): *World Health Statistics 2010*. Geneva: World Health Organization.

WORLD HEALTH ORGANIZATION (WHO) (2011): *Global Status Report on Alcohol and Health*. Geneva: World Health Organization.

WORLD HEALTH ORGANIZATION (WHO) (2012): *WHO Global Report: Mortality Attributable to Tobacco*. Geneva: World Health Organisation.

WORLD HEALTH ORGANIZATION (WHO) (2013): *WHO Report on the Global Tobacco Epidemic, 2013: Enforcing Bans on Tobacco Advertising, Promotion and Sponsorship*. Geneva: World Health Organization.

3. TIME PREFERENCES AND SMOKING BEHAVIOUR

I. INTRODUCTION

Cigarette smoking involves a clear intertemporal trade off: any short-term benefits from smoking are coupled with the potential for large long-term costs. The intertemporal implications of the decision to smoke have stimulated research into the relationship between instantaneous time preferences[1] and smoking behaviour. There is a large literature in psychology and a burgeoning literature in economics which analyses this relationship. The most common, although not universal, finding is that smokers tend to discount future rewards more heavily than non-smokers[2], implying that smokers place greater emphasis on the present relative to the future.

However, most of the studies in this literature rely on small samples, hypothetical rewards, restrictive assumptions about the form that discounting takes, and statistical tools which are not appropriate to the data that are collected. In this chapter, I remedy some of these shortcomings: I use two relatively large samples, four distinct discounting functions, and a full information maximum likelihood statistical approach which allows one to draw robust inferences from discounting data.

I find that, in contrast to the maintained assumption of most studies in this literature, simple hyperbolic (H) discounting, such as the function proposed by Mazur (1984), does not provide the best fit to the data. Furthermore, I find that there is a need to allow multiple decision processes to characterise discounting choices. Finally, I find a significant relationship between smoking status and discounting behaviour which is relatively robust across the different discounting specifications.

[1] Chapter 2 highlighted the distinction between instantaneous risk and time preferences and intertemporal risk preferences. However, as this thesis only empirically examines instantaneous risk and time preferences, all subsequent references to "time preferences" will dispense with the "instantaneous" prefix.
[2] The term "non-smokers" refers to both ex-smokers and never-smokers.

The chapter is structured as follows. Section II reviews previous research on the relationship between discounting and smoking behaviour. Section III introduces the different discounting models which will be estimated. Section IV discusses the studies that will be used to investigate the smoking-discounting relationship and Section V formulates the statistical approach to data analysis. Section VI presents the results and Section VII concludes.

II. A REVIEW OF THE SMOKING AND DISCOUNTING LITERATURE

Table I provides a detailed summary of experimental studies investigating the relationship between smoking and time preferences. Online searches of PubMed and Econlit, employing the search criteria "smoking" and "discounting" and their variants (e.g., "smoke", "discount", and "time preference"), were used to locate these papers. An initial list of over 50 studies was trimmed according to the following rules: the study had to include a clear smoker, non-smoker comparison[3]; study participants had to make choices between amounts of money, rather than cigarettes or quality-adjusted life years, available at different points in time[4]; and the instrument used to assess discounting had to include at least 20 questions.[5] The 30 studies satisfying these criteria are listed in Table I.

Mitchell (1999) and Bickel, Odum and Madden (1999) conducted the first experiments investigating the relationship between smoking and discounting behaviour. Mitchell (1999) presented 20 relatively heavy[6], current smokers ($N_S = 20$) and 20 never-smokers ($N_{NS} = 20$) with 137 choice questions between a real larger, later (LL) reward of $10 available after one of six delays (0, 7, 30, 90, 180, or 365 days, i.e., the temporal horizon ranged from 0 to 365 days) and a real smaller, sooner

[3] A number of studies (e.g., Field et al. (2006), Dallery and Raiff (2007), Epstein et al. (2003)) focus purely on discounting among smokers and were excluded due to the lack of non-smokers in the sample.
[4] Odum, Madden and Bickel (2002) and van der Pol and Ruggeri (2008) focus on the discounting of health outcomes and Field et al. (2006) and Odum and Baumann (2007) focus on the discounting of hypothetical cigarette rewards.
[5] Some panel studies (e.g., the Health and Retirement Study (HRS)) include a module to assess discounting behaviour but the limited number of questions (e.g., three questions in the HRS, see Bradford (2010)) makes precise estimation and inference difficult, so these studies were excluded.
[6] The smokers in Mitchell's (1999) study stated that they smoked at least 15 cigarettes per day and provided a breath sample to verify their smoking status.

(SS) reward, which varied between \$0.01 and \$10.50, available immediately.[7] The questions were drawn randomly from this battery, without replacement, and presented to subjects sequentially. At the end of the experiment, one of a subject's choices was selected randomly for payment.

Mitchell used each subject's choices to determine an indifference point between the LL reward (i.e., \$10) available after a particular delay (e.g., 7 days) and an SS reward available immediately. For example, if a subject chose \$10 in 7 days over \$6.50 immediately but then chose \$7 immediately over \$10 in 7 days, the subject was assigned an indifference point of \$6.75. Taking the average of these two values is arbitrary and doing so throws away information about the uncertainty of this estimate; all that one can infer from this pattern of choices is that a subject's indifference point lies in the open interval (\$6.50, \$7). Interval data of this form is analysed appropriately using interval regression methods but Mitchell (1999) used the *estimated* indifference points as *data* to construct Mann-Whitney tests of differences in the indifference points of heavy smokers and never-smokers. Mitchell (1999) found that current smokers' indifference points were significantly lower than never-smokers' indifference points for the 7 day ($p < 0.05$), 30 day ($p < 0.01$), and 90 day ($p < 0.06$) delays.

In addition, Mitchell (1999) fitted Mazur's (1984) H discounting function to the indifference points for each subject and then compared the estimated discounting parameters of current smokers and never-smokers using a Mann-Whitney test; she found that current smokers discounted significantly more than never-smokers ($p < 0.06$). Note that using the point estimate of a discounting parameter as a datum ignores the uncertainty of this estimate and, thus, should not be used for inferential purposes.

[7] One would expect people to choose \$10.50 now over \$10 available after a delay so incorporating this question may be a test of subject comprehension; Mitchell (1999) provided no justification for the question's inclusion.

TABLE I: REVIEW OF EXPERIMENTAL LITERATURE ON SMOKING AND DISCOUNTING BEHAVIOUR

Study	Sample (size)	Elicitation method	Task-related incentives (max LL)	Horizon	Front end delay (FED)	Correct for non-linear utility	Models (estimated rates)	Statistical method	Hyperbolicky discounting?	Significant relationship with smoking?
Bickel, Odum & Madden (1999)	Adults in Burlington, VT, USA ($N_S = 23$, $N_{NS} = 22$, $N_{ES} = 21$)	Choice (ordered)	No ($1000)	7 - 9131 days	No	No	H and E ($\delta^H_S = 0.054$) ($\delta^H_{NS} = 0.007$) ($\delta^H_{ES} = 0.007$)	NLLS for discounting, ANOVA and non-parametric tests for analysis	Yes (compared to E) based on R^2 comparisons	Yes, positive for S relative to NS ($p<0.01$) and ES ($p<0.01$); No for NS relative to ES.
Mitchell (1999)	Adults in Durham, NH, USA ($N_S = 20$, $N_{NS} = 20$)	Choice (random)	Yes ($10)	0 - 365 days	No	No	H ($\delta_S = 0.012$) ($\delta_{NS} = 0.006$)	NLLS for discounting, non-parametric tests for analysis	By assumption	Yes ($p<0.06$), positive.
Baker, Johnson & Bickel (2003)	Adults in Burlington, VT, USA ($N_S = 30$, $N_{NS} = 30$)	Titration (random - Richards et al. (1999))	Yes ($100)	Real: 1 - 183 days	No	No	H NRD but from Figure 2: ($10: $\delta_S = 0.008$, $\delta_{NS} = 0.001$) ($100: $\delta_S = 0.005$, $\delta_{NS} = 0.001$)	NLLS for discounting, ANOVA for analysis	By assumption	Real: Yes ($p<0.01$), positive.
			No ($1000)	Hypothetical: 1 - 9131 days	No	No	H NRD but from Figure 2: ($10: $\delta_S = 0.008$, $\delta_{NS} = 0.003$) ($100: $\delta_S = 0.006$, $\delta_{NS} = 0.0008$) ($1000: $\delta_S = 0.004$, $\delta_{NS} = 0.0005$)			Hypothetical: Yes ($p<0.01$), positive.
Reynolds, Karraker, Horn & Richards (2003)	Adolescents in Morgantown, WV, USA ($N_S = 19$, $N_{NS} = 19$, $N_T = 17$)	Titration (random - Richards et al. (1999))	1-out-of-2-tasks ($10)	1 - 365 days	No	No	H ($\delta_S = 0.010$) ($\delta_{NS} = 0.007$) ($\delta_T = 0.016$)	NLLS for discounting, ANOVA for analysis	By assumption	No.
Reynolds (2004)	Adolescents and young adults in Morgantown, WV, USA ($N_{S(adolescent)} = 19$, $N_{S(adult)} = 25$, $N_{NS} = 29$)	Titration (random - Richards et al. (1999))	1-out-of-2-tasks ($10)	1 - 365 days	No	No	H ($\delta_{S(adolescent)} = 0.016$) ($\delta_{S(adult)} = 0.075$) ($\delta_{NS(adult)} = 0.012$)	NLLS for discounting, ANOVA, correlations and post hoc tests for analysis	By assumption	Yes, positive for $S_{(adult)}$ relative to $S_{(adolescent)}$ ($p<0.05$) and $NS_{(adult)}$ ($p<0.05$). No for $S_{(adolescent)}$ relative to $NS_{(adult)}$.
Reynolds, Richards, Horn & Karraker (2004)	Mostly students in Morgantown, WV, USA ($N_S = 25$, $N_{NS} = 29$)	Titration (random - Richards et al. (1999))	1-out-of-2-tasks ($10)	1 - 365 days	No	No	H ($\delta_S = 0.066$) ($\delta_{NS} = 0.015$)	NLLS for discounting, ANOVA for analysis	By assumption	Yes ($p<0.05$), positive.
Ohmura, Takahashi & Kitamura (2005)	Students in Sapporo, Japan ($N_S = 27$, $N_{NS} = 23$)	Titration (random - Richards et al. (1999))	No (¥100,000 = $1000)	7 - 1826 days)	No	No	H, E and AUC. ($AUC_S = 0.54$) ($AUC_{NS} = 0.58$)	AUC and NLLS for discounting, correlations and t tests for analysis	Yes (compared to E) based on R^2 comparisons	No.

Source: Author's construction.

Notes: S = smoker; NS = non-smoker/never-smoker; ES = ex-smoker; LS = light smoker; T = trier; FS = fast smoking adopter; SS = slow smoking progressor.

H = hyperbolic; E = exponential; QH = quasi-hyperbolic; AUC = area under the curve; NRD = not reported directly; [a] = annual rate; [b] = weekly rate; NLLS = non-linear least squares; ML = maximum likelihood.

ANOVA = analysis of variance; ANCOVA = analysis of covariance

TABLE I: REVIEW OF EXPERIMENTAL LITERATURE ON SMOKING AND DISCOUNTING BEHAVIOUR (CONTINUED)

Study	Sample (size)	Elicitation method	Task-related incentives (max LL)	Horizon	Front end delay (FED)	Correct for non-linear utility	Models (estimated rates)	Statistical method	Hyperbolicky discounting?	Significant relationship with smoking?
Heyman & Gibb (2006)	Students in Cambridge, MA, USA ($N_S = 19, N_{NS} = 31, N_{LS} = 21$)	Choice (ordered)	Yes ($29)	Real: 1 – 30 days	No	No	H Real: ($\delta_S = 0.074$) ($\delta_{NS} = 0.036$) ($\delta_{LS} = 0.045$)	Algebra and averaging for discounting, F-test and post-hoc tests for analysis	By assumption	Real: Yes, positive for S relative to NS ($p<0.01$) and LS ($p<0.05$); No for LS relative to NS.
			No ($1000)	Hypothetical: 7 – 3650 days	No	No	H Hypothetical: ($\delta_S = 0.007$) ($\delta_{NS} = 0.009$) ($\delta_{LS} = 0.004$)			Hypothetical: No.
Reynolds (2006)	Adults in Buffalo, NY, USA ($N_S = 15, N_{NS} = 15$)	Titration (random – Richards et al. (1999))	No ($10)	1 – 365 days	No	No	H ($\delta_S = 0.088$) ($\delta_{NS} = 0.020$)	NLLS for discounting, non-parametric tests for analysis	By assumption	Yes ($p<0.01$), positive.
Johnson, Bickel & Baker (2007)	Adults in Burlington, VT, USA ($N_S = 30, N_{NS} = 30, N_{LS} = 30$)	Titration (random – Richards et al. (1999))	Yes ($100)	Real: 1 – 183 days	No	No	H Real: ($10: $\delta_S = 0.006$, $\delta_{LS} = 0.003$, $\delta_{NS} = 0.0009$) ($100: $\delta_S = 0.003$, $\delta_{LS} = 0.001$, $\delta_{NS} = 0.0008$)	NLLS for discounting, ANOVA for analysis	By assumption	Real: Yes, positive for S ($p<0.05$) and LS ($p<0.05$) relative to NS; No for S relative to LS.
			No ($1000)	Hypothetical: 1 – 9131 days	No	No	H Hypothetical: ($10: $\delta_S = 0.006$, $\delta_{LS} = 0.007$, $\delta_{NS} = 0.002$) ($100: $\delta_S = 0.004$, $\delta_{LS} = 0.002$, $\delta_{NS} = 0.0005$) ($1000: $\delta_S = 0.002$, $\delta_{LS} = 0.0008$, $\delta_{NS} = 0.0003$)			Hypothetical: Yes, positive for S ($p<0.01$) and LS ($p<0.05$) relative to NS; No for S relative to LS.
Reynolds et al. (2007)	Adolescents in Columbus, OH, USA ($N_S = 25, N_{NS} = 26$)	Titration (random – Richards et al. (1999))	Yes ($10)	1 – 365 days	No	No	AUC NRD but from Figure 1: ($AUC_S = 0.129$) $AUC_{NS} = 0.234$)	AUC for discounting, ANOVA and ANCOVA for analysis	AUC, but dropped subjects that had poor H fit	Yes ($p<0.05$), positive.
Bickel, Yi, Kowal & Gatchalian (2008)	Adults in Little Rock, AR, USA ($N_S = 30, N_{NS} = 29$)	Titration (random - Richards et al. (1999))	No ($1000)	1 - 9131 days	No	No	H and E ($\delta^H_S = 0.007$) ($\delta^H_{NS} = 0.001$)	NLLS for discounting, ANCOVA for analysis	Yes (compared to E) based on R^2 comparisons	Yes ($p<0.05$), positive.

Source: Author's construction.

Notes: S = smoker; NS = non-smoker/never-smoker; ES = ex-smoker; LS = light smoker; T = trier; FS = fast smoking adopter; SS = slow smoking progressor.

H = hyperbolic; E = exponential; QH = quasi-hyperbolic; AUC = area under the curve; NRD = not reported directly; [a] = annual rate; [b] = weekly rate; NLLS = non-linear least squares; ML = maximum likelihood.

ANOVA = analysis of variance; ANCOVA = analysis of covariance

TABLE I: REVIEW OF EXPERIMENTAL LITERATURE ON SMOKING AND DISCOUNTING BEHAVIOUR (CONTINUED)

Study	Sample (size)	Elicitation method	Task-related incentives (max LL)	Horizon	Front end delay (FED)	Correct for non-linear utility	Models (estimated rates)	Statistical method	Hyperbolicky discounting?	Significant relationship with smoking?
Chabris et al. (2008)	1. Adults in Boston, MA, USA (N = 126)	Choice (random – Kirby et al. (1999))	1-in-6-chance ($85)	7 – 186 days	No	No	H ($\delta = 0.015$, SD = 0.02)	ML for discounting, OLS, Tobit, Probit for analysis	By assumption	Yes ($p<0.05$), positive.
	2. Adults in the USA (recruited online) (N = 326)		1-in-6-chance ($85)	7 – 186 days	No	No	H ($\delta = 0.008$, SD = 0.009)	ML for discounting, OLS, Tobit, Probit for analysis	By assumption	No
Sweitzer et al. (2008)	Adults in Allegheny County, PA, USA ($N_S = 101$, $N_{NS} = 145$, $N_T = 279$, $N_{ES} = 185$)	Choice (random)	No ($100)	7 - 1825 days	No	No	H ($\delta_S = 0.120$) ($\delta_{NS} = 0.079$) ($\delta_T = 0.090$) ($\delta_{ES} = 0.086$)	NLLS for discounting, ANCOVA for analysis	By assumption	Yes, positive for S relative to NS ($p<0.01$), ES ($p<0.01$) and T ($p<0.01$); No for all other comparisons.
Adams & Nettle (2009)	Adults in 15 major urban areas in the USA (recruited online) ($N_S = 70$, $N_{NS} = 346$)	Choice (ordered)	No ($1000)	30 - 3652 days	No	No	H ($\delta = 1.3$)[a]	NLLS for discounting, logistic regression for analysis	By assumption	No.
Audrain-McGovern et al. (2009)	High school students in northern Virginia, USA ($N_{NS} = 556$, $N_{FS} = 112$, $N_{SS} = 241$)	Choice (random - Kirby et al. (1999))	Not reported ($85)	7 - 186 days	No	No	H Assuming ln transformation: ($\delta_{FS} = 0.023$) ($\delta_{SS} = 0.016$) ($\delta_{NS} = 0.010$)	Algebra and averaging for discounting, latent growth curve modeling (LGCM) and growth mixture modeling (GMM) for analysis	By assumption	LCGM: Yes ($p<0.05$), positive. GMM: Yes, positive for FS ($p<0.05$) and SS ($p<0.05$) relative to NS; No for FS relative to SS.
Jones, Landes, Yi & Bickel (2009)	Adults in Little Rock, AR, USA ($N_S = 86$, $N_{NS} = 141$)	Titration (ordered or random)	No ($1000)	1 - 9131 days	No	No	H NRD but from Figure 3: ($100: $\delta_{S(men)} = 0.012$, $\delta_{NS(men)} = 0.001$, $\delta_{S(women)} = 0.0015$, $\delta_{NS(women)} = 0.002$) ($1000: $\delta_{S(men)} = 0.0075$, $\delta_{NS(men)} = 0.0005$, $\delta_{S(women)} = 0.001$, $\delta_{NS(women)} = 0.001$)	NLLS for discounting, ANCOVA for analysis	By assumption	Yes, positive for $S_{(men)}$ ($p<0.01$) relative $NS_{(men)}$ at $100 and $1000 magnitudes; No for $S_{(women)}$ relative to $NS_{(women)}$ at both magnitudes.
Melanko et al. (2009)	Adolescents in central Ohio, USA ($N_S = 50$, $N_{NS} = 25$). Smokers were split into high and low pyschopathology groups.	Titration (random - Richards et al. (1999))	Yes ($10)	1 - 365 days	No	No	AUC NRD but from Figure 1: ($AUC_{S(low)} = 0.126$) ($AUC_{S(high)} = 0.214$) ($AUC_{NS} = 0.275$)	AUC for discounting, ANOVA for analysis	AUC, no assumption about form of discounting	Yes, positive for $S_{(low)}$ relative to NS ($p=0.01$); No for all other comparisons.
Businelle, McVay, Kendzor & Copeland (2010)	Adults in southern USA ($N_S = 20$, $N_{NS} = 34$)	Choice (ordered)	No ($1000)	0.25 - 9131 days	No	No	H and AUC ($\delta_S = 0.077$) ($\delta_{NS} = 0.039$)	NLLS and AUC for discounting, ANCOVA for analysis	By assumption (but also used AUC)	Yes ($p=0.01$), positive.

Source: Author's construction.

Notes: S = smoker; NS = non-smoker/never-smoker; ES = ex-smoker; LS = light smoker; T = trier; FS = fast smoking adopter; SS = slow smoking progressor.

H = hyperbolic; E = exponential; QH = quasi-hyperbolic; AUC = area under the curve; NRD = not reported directly; [a] = annual rate; [b] = weekly rate; NLLS = non-linear least squares; ML = maximum likelihood.

ANOVA = analysis of variance; ANCOVA = analysis of covariance

TABLE I: REVIEW OF EXPERIMENTAL LITERATURE ON SMOKING AND DISCOUNTING BEHAVIOUR (CONTINUED)

Study	Sample (size)	Elicitation method	Task-related incentives (max LL)	Horizon	Front end delay (FED)	Correct for non-linear utility	Models (estimated rates)	Statistical method	Hyperbolicky discounting?	Significant relationship with smoking?
Harrison, Lau & Rutström (2010)	Adults in Denmark ($N_S = 71$, $N_{NS} = 181$)	Choice (ordered)	1-in-10-chance ($1175)	30 - 730 days	Yes	Yes	H and E Linear utility: ($\delta^H_{S(men)} = 0.341$)[a] ($\delta^H_{NS(men)} = 0.240$)[a] ($\delta^H_{S(women)} = 0.329$)[a] ($\delta^H_{NS(women)} = 0.250$)[a]	ML for discounting and analysis	25% - 40% of choices by smokers and non-smokers best characterised by H	Linear utility: Men: Yes ($p<0.05$), positive; Women: Yes ($p<0.10$), positive.
							H and E Concave utility: ($\delta^H_{S(men)} = 0.127$)[a] ($\delta^H_{NS(men)} = 0.093$)[a] ($\delta^H_{S(women)} = 0.109$)[a] $\delta^H_{NS(women)} = 0.095$)[a]			Concave utility: Men: Yes ($p<0.05$), positive; Women: No.
Bickel et al. (2012)	Adults in the USA (recruited online) ($N_S = 182$, $N_{NS} = 614$)	Choice (random)	No ($85)	10 - 75 days	No	No	H (Not reported)	Algebra and averaging for discounting, ANCOVA for analysis	By assumption	Yes ($p<0.01$), positive.
Mitchell & Wilson (2012)	1. Adults in Portland, OR, USA ($N_S = 20$, $N_{NS} = 20$)	Choice (random)	Yes ($50)	14 - 154 days	Yes	No	H and QH (0 FED: $\delta^H_S = 0.230$, $\delta^H_{NS} = 0.020$) (+ FED: $\delta^H_S = 0.070$, $\delta^H_{NS} = 0.010$)	NLLS and ML for discounting, ANOVA for analysis	By assumption (but also estimated QH)	Yes ($p<0.01$), positive.
	2. Adults in Portland, OR, USA ($N_S = 16$, $N_{NS} = 16$)		No ($50)	14 - 154 days	Yes	No	H and QH (0 FED: $\delta^H_S = 0.120$), $\delta^H_{NS} = 0.020$) (+ FED: $\delta^H_S = 0.050$, $\delta^H_{NS} = 0.010$)			Yes ($p<0.01$), positive.
Reynolds & Fields (2012)	Adolescents in Columbus, OH, USA ($N_S = 50$, $N_{NS} = 50$, $N_T = 41$)	Titration (random - Richards et al. (1999))	Yes ($10)	1 - 365 days	No	No	AUC NRD but from Figure 1: ($AUC_S = 0.166$) ($AUC_T = 0.224$) ($AUC_{NS} = 0.347$)	AUC for discounting, ANOVA and ANCOVA for analysis	AUC, no assumption about form of discounting	Yes, positive for S ($p<0.01$) and T ($p<0.05$) relative to NS; No for S relative to T.
Stillwell & Tunney (2012)	International online study ($N_S = 1592$, $N_{LS} = 669$, $N_{NS} = 6777$)	Choice (ordered or random)	No ($1000)	7 - 1826 days	No	No	H NRD but from Figure 3: ($\delta_S = 0.437$) ($\delta_{LS} = 0.397$) ($\delta_{NS} = 0.369$)	NLLS for discounting, ANOVA for analysis	Yes (compared to E) based RSS comparisons	Yes, positive for S relative to LS ($p<0.01$) and NS ($p<0.01$) and positive for LS relative to NS ($p<0.01$).
Wing, Moss, Rabin & George (2012)	Adults in the greater Toronto area, Canada ($N_S = 23$, $N_{NS} = 37$)	Choice (random - Kirby et al. (1999))	No ($85)	7 - 186 days	No	No	H NRD but from Figure 1: ($\delta_S = 0.017$) ($\delta_{NS} = 0.011$)	Algebra and averaging for discounting, ANCOVA for analysis	By assumption	No.

Source: Author's construction.

Notes: S = smoker; NS = non-smoker/never-smoker; ES = ex-smoker; LS = light smoker; T = trier; FS = fast smoking adopter; SS = slow smoking progressor.

H = hyperbolic; E = exponential; QH = quasi-hyperbolic; AUC = area under the curve; NRD = not reported directly; [a] = annual rate; [b] = weekly rate; NLLS = non-linear least squares; ML = maximum likelihood.

ANOVA = analysis of variance; ANCOVA = analysis of covariance

TABLE I: REVIEW OF EXPERIMENTAL LITERATURE ON SMOKING AND DISCOUNTING BEHAVIOUR (CONTINUED)

Study	Sample (size)	Elicitation method	Task-related incentives (max LL)	Horizon	Front end delay (FED)	Correct for non-linear utility	Models (estimated rates)	Statistical method	Hyperbolicky discounting?	Significant relationship with smoking?
Balevich, Wein & Flory (2013)	Students in Flushing, NY, USA ($N_S = 50$, $N_{NS} = 102$, $N_T = 91$)	Choice (random) or titration (random)	No ($100)	1 - 1825 days	No	No	H ($\delta_S = 0.126$) ($\delta_{NS} = 0.135$) ($\delta_T = 0.138$)	NLLS for discounting, ANOVA for analysis	By assumption	No.
Poltavski & Weatherly (2013)	Students in Grand Forks, ND, USA ($N_S = 16$, $N_{LS} = 74$, $N_{NS} = 92$)	Choice (random)	No ($100,000)	183 - 3652 days	No	No	H and AUC ($1000: $\delta_S = 0.010$, $\delta_{LS} = 0.010$, $\delta_{NS} = 0.007$) ($100,000: $\delta_S = 0.008$, $\delta_{LS} = 0.008$, $\delta_{NS} = 0.007$)	NLLS and AUC for discounting, ANOVA for analysis	By assumption (but also used AUC)	No.
Sheffer et al. (2013)	Adults in Little Rock, AR, USA ($N_S = 47$, $N_{NS} = 19$)	Titration (random - Richards et al. (1999))	No ($1000)	1 - 9131 days	No	No	H NRD but from Figure 1: ($\delta_S = 0.020$) ($\delta_{NS} = 0.004$	NLLS for discounting, ANCOVA for analysis	By assumption	Yes ($p<0.05$), positive.
Kang & Ikeda (2014)	Adults in Japan ($N_S \approx 862$, $N_{NS} \approx 2588$)	Choice (ordered)	No (¥1000,000 = $10000)	7 - 365 days	Yes	No	E and proxies for H (See Table III, the mean of δ^E ranges from 0.022 to 1.904)[a]	ML for discounting, hurdle model for analysis	Assumes E but constructs H proxies	Yes ($p<0.01$), positive.
Kobiella et al. (2014)	Adults in Mannheim, Germany ($N_S = 27$, $N_{NS} = 31$)	Choice (random)	Yes (€41.32)	14 - 28 days	Yes	No	H ($\delta_S = 0.055$)[b] ($\delta_{NS} = 0.038$)[b]	NLLS for discounting, t-tests for analysis	By assumption	Yes ($p<0.05$), positive.

Source: Author's construction.

Notes: S = smoker; NS = non-smoker/never-smoker; ES = ex-smoker; LS = light smoker; T = trier; FS = fast smoking adopter; SS = slow smoking progressor.

H = hyperbolic; E = exponential; QH = quasi-hyperbolic; AUC = area under the curve; NRD = not reported directly; [a] = annual rate; [b] = weekly rate; NLLS = non-linear least squares; ML = maximum likelihood.

ANOVA = analysis of variance; ANCOVA = analysis of covariance

Bickel, Odum and Madden (1999) (BOM) presented 23 heavy, current smokers, 22 never-smokers, and 21 ex-smokers (N_S = 23, N_{NS} = 22, and N_{ES} = 21)[8] with 189 choice questions between a hypothetical LL reward of $1000 available after one of seven delays (the temporal horizon ranged from 7 days to 25 years) and a hypothetical SS reward available immediately. For each delay, the SS rewards were presented sequentially in descending and then ascending order and subjects were asked to indicate their preference between each SS reward and the LL reward of $1000. A simple average of the last SS reward chosen in descending order and the first SS reward chosen in ascending order was used to define a person's indifference point for the $1000 LL reward at a particular delay; this method was used to derive 7 indifference points, for the seven delays in the task, for each subject even though taking the average of these two values is arbitrary and doing so throws away information about the uncertainty of the estimate.

BOM then fitted an exponential (E) function and Mazur's (1984) H function, using non-linear least squares (NLLS) estimation, to each subject's 7 derived indifference points. For each subject, BOM then compared[9] the fit of these two functions using the coefficient of determination, R^2, and found that for most subjects, H provided a better fit than E. They then used each subject's *estimated* R^2 value for the E and H functions as *data* to construct tests of whether the E or H functions provided a better fit to the discounting data across all subjects in the different smoking groups. Using the point estimate of a statistic (i.e., the value of R^2) as a datum ignores the uncertainty of this estimate and, thus, does not produce a valid test of one function's ability to better explain discounting data. Nevertheless, BOM state that the H function provided a better fit than the E function among current smokers ($p < 0.01$), never-smokers ($p < 0.01$), and ex-smokers ($p < 0.01$).

Finally, BOM compared the point *estimates* of the H discounting function across the current smoker, never-smoker, and ex-smoker groups by estimating an analysis of variance (ANOVA) model which included a smoking status covariate; they found a

[8] The smokers in BOM (1999) reported smoking at least 20 cigarettes per day for 5 years and had a Fagerström Test for Nicotine Dependence score of at least 6. Never-smokers reported never smoking and ex-smokers reported abstinence for at least one year following 5 years of smoking at least 20 cigarettes a day.

[9] Note that this was a simple comparison of the point estimates of R^2 for the E and H functions and not a formal statistical test.

significant overall effect of smoking status but estimates should not be used as data for inferential purposes. In addition, planned Mann-Whitney pairwise comparisons of the H discounting function *estimates* showed that current smokers discounted significantly more than never-smokers ($p < 0.01$), and ex-smokers ($p < 0.01$); there was no significant difference between never-smokers and ex-smokers.

Thus, the first two studies analysing the relationship between smoking and discounting behaviour suggested that current smokers discount more heavily than never-smokers. In addition, it appeared that this result was robust to different subjects pools, real as opposed to hypothetical rewards, different LL reward magnitudes ($10 versus $1000), and different elicitation mechanisms (random versus ordered choice). Although not confirmed by both studies, there was some evidence that the H function provided a better fit to the discounting data than the E function.

Reynolds, Karraker, Horn and Richards (2003) (RKHR), in a study with adolescent smokers ($N_S = 19$), adolescent never-smokers ($N_{NS} = 19$), and adolescent "triers" ($N_T = 17$)[10], provided the first null result in this literature. They used the titration procedure of Richards, Zhang, Mitchell and de Wit (1999) (RZMW), dubbed "Titration (random) – Richards et al. (1999)" in Table I, to derive indifference points for real $10 LL rewards available at different points in time (the temporal horizon ranged from 1 to 365 days). This titration procedure has been used extensively in the smoking and discounting literature and deserves further comment.

A titration procedure uses a subject's choices to determine the next set of choices that the subject faces. As a simple example, if someone chooses $10 after 7 days over $5 now, the titration algorithm assumes that $10 after 7 days will be chosen over all amounts of money less than $5 available now (e.g., $4 or $3 available now). Consequently, the titration algorithm will narrow the search for a subject's indifference point for that delay period (i.e., 7 days in our example) to the open interval ($5, $10). Some titration procedures take the average of the two values defining that interval to determine the next SS reward presented to the subject, $7.50

[10] "Triers" had smoked cigarettes for the first time in the 6 months prior to the study and they smoked an average of 3.76 cigarettes in total over this time span. Smokers, by contrast, had smoked every week for at least 6 months prior to the study and they smoked 46.42 cigarettes, on average, per week.

in this case. If the subject chooses $7.50 now over $10 in 7 days, then the range for indifference points is narrowed to ($5, $7.50). If, by contrast, the subject chooses $10 in 7 days over $7.50 now, then the range for indifference points is narrowed to ($7.50, $10). By continually splitting the difference of an interval, the titration algorithm converges to a subject's indifference point for a particular delay.[11,12]

An issue with this titration procedure is that if a subject makes a mistake (e.g., chooses $10 after 7 days when he meant to choose $5 now), it becomes impossible to recover the subject's "true" indifference point because the algorithm uses that mistake to refine the subsequent set of choices presented to the subject. The algorithm developed by RZMW is more sophisticated and uses two top and two bottom limits, rather than one top limit (e.g., $10 in our previous example) and one bottom limit (e.g., $5 in our previous example), to alleviate this issue. By employing multiple top and bottom limits, the algorithm of RZMW can recover a subject's indifference point even after a mistake.

Another issue with a simple titration algorithm which splits the difference of an interval is that the adjusting nature of the algorithm is evident to the subject and he can deduce that his future choices depend on his current choices. This raises an obvious incentive-compatibility problem because the subject can "game" the algorithm so as to be presented with higher SS rewards on subsequent decisions. The algorithm of RZMW attempts to mitigate this problem by randomly drawing SS amounts from within an interval, rather than simply splitting the difference, and by randomly selecting LL reward delays, rather than determining the indifference point for one delay before moving on to the next.

RKHR used this algorithm to investigate the discounting behaviour of adolescent smokers, never-smokers, and "triers." Subjects also completed a probability discounting task, which will be discussed in more detail in Chapter 4, and they were paid for one choice across both tasks; this payment scheme is referred to as 1-out-of-

[11] Clearly the algorithm must terminate at some point, lest it continue indefinitely. In studies with $10 LL rewards, as in RKHR, the algorithm stopped when the difference between the rewards in the interval had declined to $0.50.

[12] As the interval within which a person's indifference point lies gets smaller and smaller, it is questionable whether that person is willing or able to make increasingly fine-grained choices.

2-tasks in Table I. The H discounting function was estimated for each subject, using NLLS, and the estimated discounting parameters were log transformed to normalise their distribution.[13] These transformed discounting parameters were used as data and fed into an ANOVA model so as to compare the discounting behaviour of the 3 smoking groups: there were no significant differences between smokers, never-smokers and "triers."

This two-step approach to analysis (i.e., using NLLS, or some other technique, to estimate discounting parameters and then using the, typically log-transformed, point estimates as data in subsequent statistical models) is remarkably common in this literature. Harrison, Lau and Rutström (2010) (HLR) is the only study in Table I which does not use this method, and for good reason. The problem with the two-step approach, other than that it often uses tiny samples to estimate discounting parameters at the level of the individual, is that estimated discounting parameters are estimates, and not data. Such estimates comprise both a point estimate (of the mean) and a standard error, and to use only the point estimate is to throw away information on the uncertainty of that estimate.[14] Moreover, using an estimated discounting parameter as data violates the statistical assumptions of the second-stage models: specifically, that the covariates are measured without error. Thus, the statistical inferences drawn from this approach are simply not valid. HLR estimate discounting parameters as a linear function of observable characteristics (e.g., age, gender, and smoking status) so that the uncertainty of the discounting parameter estimates propagates into the inferences which are drawn from the data. This valid statistical approach will be used in this chapter.

Table I collates the results from the other studies and, on inspection, a number of interesting patterns emerge. The vast majority (i.e., 24) of the studies investigating

[13] RKHR used the common (i.e., base 10) logarithm to transform their data. A number of studies (e.g., Reynolds, Richards, Horn and Karraker (2004), Reynolds (2004), Heyman and Gibb (2006), Johnson, Bickel and Baker (2007) also adopt the common logarithmic transformation while others (e.g., Baker, Johnson and Bickel (2003), Bickel, Yi, Kowal and Gatchalian (2008), Jones, Landes, Yi and Bickel (2009), Sheffer et al. (2013)) use the natural logarithm to transform discounting parameters.

[14] This problem is compounded when indifference points are computed by taking an average of the interval within which a person's indifference point lies (i.e., taking the average of an interval derived by a titration mechanism). In this case, there are two levels of uncertainty (i.e., uncertainty about the indifference points and uncertainty about the parameter estimates) which are ignored when the final point estimate of a discounting parameter is used as data.

smoking and discounting behaviour were conducted in the US, 2 took place in Japan, 1 in Denmark, 1 in Canada, 1 in Germany, and 1 recruited subjects internationally over the internet. An important feature of these studies is that they have relatively diverse subject pools (i.e., they do not typically rely on convenient student samples but rather recruit from the community at large) which thereby bolsters the external validity of the results.

However, most of the studies have small sample sizes: 19 of the 30 studies recruited less than 100 people and 15 of these studies had samples of less than 70 people. Fortunately, since 2008, the trend has been towards larger and larger samples (e.g., Sweitzer et al. (2008) recruited 710 subjects, Audrain-McGovern et al. (2009) recruited 909 subjects, Kang and Ikeda (2014) used a sample of 3450 people, and Stillwell and Tunney (2012) recruited 9038 individuals).

With regards to elicitation mechanisms, 16 studies used choice procedures, 13 used titration, and 1 employed both methods (see Balevich, Wein and Flory (2013)). The task designed by Kirby, Petry and Bickel (1999) (KPB) was used in 3 of the choice procedure studies and this task will be discussed in more detail in Section IV because it was used in the research projects reported in this chapter.

A perennial issue in the interpretation of experimental results is whether real or hypothetical rewards were used in a study. If a study uses hypothetical rewards, all it really elicits is the choices a person thinks he would make when faced with those contingencies, or the choices he thinks the experimenter wants him to make. If real rewards are used, by contrast, a subject's choices ultimately determine the payment he receives and this – coupled with a task that is easily understood, a transparent payment scheme, salient rewards, and an incentive-compatible experimental design – promotes truthful revelation of preferences. Thus, one should give far more credence to studies using real as opposed to hypothetical rewards because in the former instance one can analyse what people actually did rather than what they think they would do or what they want the experimenter to think they would do.

Of the studies in Table I, 5 only used real rewards (Mitchell (1999), Reynolds et al. (2007), Melanko et al. (2009), Reynolds and Fields (2012), Kobiella et al. (2014))

whereas 4 used a combination of real and hypothetical rewards (Baker, Johnson and Bickel (2003), Heyman and Gibb (2006), Johnson, Bickel and Baker (2007), Mitchell and Wilson (2012)). Entirely hypothetical rewards were used in 15 studies (BOM, Ohmura, Takahashi and Kitamura (2005), Reynolds (2006), Bickel, Yi, Kowal and Gatchalian (2008), Sweitzer et al. (2008), Adams and Nettle (2009), Jones, Landes, Yi and Bickel (2009), Businelle, McVay, Kendzor and Copeland (2010), Bickel et al. (2012), Stillwell and Tunney (2012), Wing, Moss, Rabin and George (2012), Balevich, Wein and Flory (2013), Poltavski and Weatherly (2013), Sheffer et al. (2013), Kang and Ikeda (2014)), 1 study did not report whether real or hypothetical rewards were used (Audrain-McGovern et al. (2009)), and 5 studies used probabilistic payment schemes (RKHR, Reynolds (2004), Reynolds, Richards, Horn and Karraker (2004), Chabris et al. (2008), HLR).[15] Thus, approximately half of the studies in Table I used entirely hypothetical rewards and this should be taken into account when drawing conclusions about the relationship between smoking and discounting behaviour.

The temporal horizon (i.e., the time delay between the SS and LL rewards) of the studies reported in Table I ranges from 6 hours to 25 years. Studies using real rewards or probabilistic payment schemes tend to employ far shorter temporal horizons than studies using hypothetical rewards; this makes sense because the credibility of payments in the distant future would be open to question. Of the studies using real rewards or probabilistic payment schemes, only 1 had a temporal horizon as long as 2 years (i.e., HLR) whereas 7 studies using hypothetical rewards had temporal horizons extending out to 25 years (BOM, Baker, Johnson and Bickel (2003), Johnson, Bickel and Baker (2007), Bickel, Yi, Kowal and Gatchalian (2008), Jones, Landes, Yi and Bickel (2009), Businelle, McVay, Kendzor and Copeland (2010), Sheffer et al. (2013)).

[15] Studies employing real rewards typically make use of the random lottery incentive mechanism (RLIM) to determine subject payment. RLIM randomly selects one of a subject's choices on a task and, in a study with real rewards, pays out this choice with certainty. A probabilistic payment scheme also makes use of RLIM but subjects are only given some chance of being paid for the randomly selected choice (i.e., subjects are not paid with certainty). In Chabris et al. (2008) subjects were given a 1-in-6 chance of being paid for one of their choices while in HLR subjects were given a 1-in-10 chance of being paid for one of their choices. By contrast, RKHR, Reynolds, Richards, Horn and Karraker (2004) and Reynolds (2004) paid subjects for 1 choice across two different tasks, implying that subjects had roughly a 50% chance of being paid for one of their choices on the discounting task.

Time preferences are represented mathematically using a discounting function. There are a number of discounting functions which have been proposed (see Section III) but the majority of studies in Table I (i.e., 22 out of 30) adopted the assumption that people discount hyperbolically and, thus, only used Mazur's (1984) H function in their analyses.[16] There are 4 studies in Table I which directly compared the E and H discounting functions, using either R^2 or the residual sum of squares (RSS) to adjudicate between them, and all of the studies found that the H function better explains discounting data (BOM, Ohmura, Takahashi and Kitamura (2005), Bickel, Yi, Kowal and Gatchalian (2008), Stillwell and Tunney (2012)). There are 3 studies (Reynolds et al. (2007), Melanko et al. (2009), Reynolds and Fields (2012)) which only used the "theoretically neutral" area under the curve (AUC) method of Myerson, Green and Warusawitharana (2001) to compare the discounting of smokers and non-smokers[17], and 1 study (HLR) estimated a statistical model that allows both E and H discounting functions to characterise the data.

HLR's approach is based on the sensible idea that some discounting choices may be better explained by an E function whereas others may be better explained by an H function and that the data should be used to determine the proportion of choices best explained by each model. Using this so-called "mixture model" approach, HLR found that approximately 25% - 40% of discounting choices were best characterised by the H function. This suggests that researchers investigating the link between smoking and discounting behaviour may have relied too heavily on the H function because it does not explain all discounting choices all of the time.

Mixture models also address a deeper issue of bias in the estimation of discounting models. Suppose, for example, that at least 50% of the choices in a dataset are best characterised by the H function whereas the remaining fraction is best characterised

[16] Businelle, McVay, Kendzor and Copeland (2010) and Poltavski and Weatherly (2013) assumed hyperbolic discounting but also used the area under the curve (AUC) method of Myerson, Green and Warusawitharana (2001) to draw inferences about the relationship between smoking and discounting. Mitchell and Wilson (2012) assumed hyperbolic discounting but also estimated a quasi-hyperbolic discounting function.

[17] The AUC method is "theoretically neutral" because it does not assume that discounting takes a particular form (e.g., E or H). Instead, when using the AUC method, one calculates the area under a subject's derived indifference points and normalises this to lie in the closed unit interval [0, 1]. Larger AUCs imply shallower discounting and, thus, the AUCs of smokers and non-smokers can be compared to determine whether the groups differ in their discounting behaviour.

by the E function. If one just estimates the H model on the whole dataset then one will reject the E model in favour of the H model because the estimate obtained from the H model will be halfway between the "true" H estimate and the estimate one would obtain from the E model. Thus, if one just estimates the H model then this biases against the E model. Mixture models remove this source of bias by allowing both discounting models to account for the data, and by estimating the proportion of the data which each model explains. In this chapter I will adopt the approach of HLR and estimate mixture models of a number of different discounting specifications so as not to be wedded to any particular discounting framework and to determine the proportion of discounting choices that is explained by each specification.

HLR deserves further comment because it was the first study in this literature to use a front end delay (FED) to the SS reward and it is the only study which incorporates utility function curvature when estimating discounting models. Prior to the work of Coller and Williams (1999) it was common to make receipt of the SS reward immediate, as is the case in most of the studies in Table I. An issue with an experimental design where the SS reward is immediate (i.e., a design with no FED) is that it may increase preference for the SS reward due to the additional transaction costs and uncertainty associated with receipt of the LL reward. A FED is used to hold these transactions costs constant across the two rewards. Following the work of HLR, Mitchell and Wilson (2012), Kang and Ikeda (2014) and Kobiella et al. (2014) used a FED for some of the choices they presented to subjects.

Time preferences are defined over time-dated utility flows, not flows of money. These are equivalent if a utility function is linear but Andersen, Harrison, Lau and Rutström (2008) showed that if a utility function is concave then the assumption of linearity will, for the same observed choices, bias the estimation of discounting parameters upwards. Thus, to draw accurate inferences about discounting behaviour it is important to incorporate utility function curvature in the estimation of discounting models. To my knowledge, HLR are the only researchers to incorporate the shape of the utility function when analysing the relationship between smoking and discounting behaviour. They used a risk preference task to determine the shape of the utility function, under the assumption that expected utility theory characterised choices over

risky prospects, which they then estimated jointly with the parameters of discounting models.

Incorporating utility function curvature had a marked effect on their results. Assuming linear utility, HLR found that both male and female smokers discount significantly more than their non-smoking counterparts. However, when the discounting models were estimated jointly with the curvature of the utility function, only male smokers discounted significantly more than male non-smokers; there was no statistically significant difference in the discounting behaviour of female smokers and non-smokers. The null result for women, under joint estimation, was driven by the fact that female smokers were significantly more risk averse (i.e., had significantly more curvature in their utility functions) than female non-smokers. By assuming linear utility, this difference in utility function curvature among women showed up as a difference in their discounting behaviour.

Thus, to draw accurate inferences about smoking and discounting, it is crucial to jointly estimate utility function curvature and discounting parameters. Unfortunately the studies used in this chapter did not collect data on risk preferences so I cannot incorporate the shape of the utility function in this way. Nevertheless, I will briefly discuss a parametric approach to this issue later.

An important feature of the estimates presented in Table I is that 26 of the studies computed daily discount rates. This is common in the behavioural psychology literature but not in economics where annual discount rates are the norm. Of the remaining studies, 3 estimated annual rates, and 1 estimated weekly rates. I estimate daily discount rates in this chapter due to a puzzling inferential issue which emerges when estimating or inferring annual discount rates with the KPB discounting task. In Appendix A I focus directly on this inferential issue by conducting 125,000 replications of a simulation which generates KPB discounting data. The weight of the simulation evidence in Appendix A suggests that daily estimates are more reliable numerically than annual estimates, irrespective of whether I assume a daily or annual discount rate data generating process (DGP).

The last column of Table I reports whether the researchers found a significant statistical relationship between smoking and discounting behaviour. To understand this column, note that a "positive" relationship between smoking and discounting means that smokers discount more heavily than non-smokers. Some of the papers in Table I report findings from different studies or from different treatments in the same study. For example, Baker, Johnson and Bickel (2003) report results from real and hypothetical experimental treatments whereas Chabris et al. (2008) report findings from multiple studies. In some cases (e.g., Baker, Johnson and Bickel (2003)) results were the same across studies and treatments, while in others (e.g., Chabris et al. (2008), Heyman and Gibb (2006)) they differed. Of the 36 reported findings in Table I, 28 were positive and significant while the remaining 8 were null results.[18] Thus, the bulk of findings in this literature – irrespective of whether real or hypothetical rewards, long or short temporal horizons, choice or titration elicitation mechanisms, small or large samples, and simple or complex statistical procedures were used – point to a positive relationship between smoking and discounting behaviour.

In this chapter I will examine the smoking-discounting relationship using two relatively large samples, four distinct discounting models, and a statistical framework which allows one to draw robust inferences from data generated by the KPB discounting task. I will also estimate a mixture of the different discounting specifications to determine what proportion of choices is best characterised by the different models, and to explore the possibility that smokers are more prone to time-inconsistency than non-smokers in that, under the assumption of an additively-separable intertemporal utility function, they are more likely to discount according to the H model than the E model. In the next section I discuss the different discounting models which will be used.

[18] Some studies classified smokers using more than one category (e.g., heavy and light smokers in Stillwell and Tunney (2012)), others classified non-smokers using more than one category (e.g., never-smokers and ex-smokers in BOM), and still others separated male and female smokers and non-smokers (e.g., Jones, Landes, Yi and Bickel (2009) and HLR). In a few of these cases, comparisons between some of the groups were significant while others were not, which makes coding the study problematic. To resolve this issue, studies were coded as having found a significant result if at least one smoker, non-smoker comparison was statistically significant. Clearly this is not a perfect system but it is arguably preferable to coding a study as having found no significant results just because one comparison (between, say, light smokers and non-smokers) was not significant even though another comparison (between, say, heavy smokers and non-smokers) was significant.

III. DISCOUNTING FUNCTIONS

Samuelson (1937) developed the discounted utility model which has been the workhorse of theoretical and empirical economics since its conception.[19] From about the mid-1950s though, both economists (e.g., Strotz (1955-1956) and Phelps and Pollak (1968)) and psychologists (e.g., Herrnstein (1961) and Ainslie (1975)) have challenged the validity of the model on theoretical and empirical grounds, which has led to the development of numerous alternative models of discounting behaviour. In this section, I will introduce the discounted utility model and some of the major alternatives by drawing on the exposition in Andersen, Harrison, Lau and Rutström (2014, p. 16-17) (AHLR), which provides a general framework for discussing the various models of time preference.

The discount factor D is the scalar, for a particular time horizon τ, which equates the *utility* of income received at time t with the *utility* of income received at time $t + \tau$:

$$U(y_t) = DU(y_{t+\tau}), \tag{1}$$

for some utility function $U(\cdot)$. A common assumption in the empirical literature is that $U(\cdot)$ is linear in y which means that D is then the discount factor that equalises *income* received at time t with *income* received at time $t + \tau$. Other than HLR, every study in Table I makes this assumption.

Samuelson's (1937) discounted utility model employs the exponential (E) discount factor:

$$D^{E}(t) = 1 / (1 + \delta)^{t}, \tag{2}$$

for $t \geq 0$, and where the discount rate d is:

$$d^{E}(t) = \delta \tag{3}$$

[19] Samuelson (1937) was explicit that he did not expect the model to describe the discounting behavior of real people but the model was nevertheless adopted for this purpose.

Two important features of this model are that it is mathematically tractable (i.e., the geometric series $\sum_t D^E(t)$ converges in the limit), and the discount rate $d^E(t)$ is a constant over time which, when coupled with an additively-separable intertemporal utility function, implies time-consistent preferences.[20]

Phelps and Pollak (1968) developed the quasi-hyperbolic (QH) discounting function, in the context of a social planning problem, which has a discount factor:

$$D^{QH}(t) = 1 \qquad \text{if } t = 0 \tag{4a}$$

$$D^{QH}(t) = \beta / (1 + \delta)^t \qquad \text{if } t > 0 \tag{4b}$$

If $\beta = 1$ the QH specification collapses to the E model, whereas if $\beta < 1$ discounting is quasi-hyperbolic. I use β and δ to represent the parameters in all of the discounting models in this chapter even though there is nothing which implies that they should be equal across the different specifications; this choice was made for notational simplicity. The QH discounting function was first used to model individual behaviour by Elster (1979, p. 71) and was subsequently popularised by Laibson (1997). It plays a prominent role in behavioural economics because it is mathematically tractable (i.e., $\sum_t D^{QH}(t)$ converges in the limit) and it incorporates the potential for time-inconsistent preferences (i.e., preferences over future options can change depending on the time point at which these future options are evaluated) when $\beta < 1$ and the intertemporal utility function is additively separable.

The discount rate in the QH model is the value of $d^{QH}(t)$ which solves $D^{QH}(t) = 1 / (1 + d^{QH})^t$. Thus,

$$d^{QH}(t) = [\beta / (1 + \delta)^t]^{(-1/t)} - 1, \tag{5}$$

for $t > 0$. Clearly the discount rate under the QH model is not constant over time if $\beta < 1$. When $\beta < 1$, there is a sharp drop in the value of a reward if it is not available immediately, but the extent of this drop declines over time as the discount rate

[20] As discussed in Chapter 2, time consistency, or the lack thereof, is central to economic models of addiction.

asymptotes towards δ. As AHLR (p. 17) note, "the drop 1 – β can be viewed as a fixed utility cost of discounting anything relative to the present …." Thus, the QH model can account for a "present-bias" or a "passion for the present" in discounting behaviour.

Mazur (1984, p. 427) developed the H discounting function to account for pigeons' preferences over fixed and variable schedules of reinforcement. The H specification has a discount factor:

$$D^{H}(t) = 1 / (1 + \delta t) \tag{6}$$

This function has been used extensively in the psychology literature and in 27 of the 30 studies in Table I. Unlike the E and QH discounting specifications, the harmonic series $\sum_t D^{H}(t)$ does not converge and the H model has not been widely used therefore in the theoretical economics literature. Mazur's (1984) H function forms part of a whole family of hyperbolic discounting models, but I will use (6) due to its importance in the literature on time preferences and smoking. The discount rate $d^{H}(t)$ for this model is:

$$d^{H}(t) = (1 + \delta t)^{(1/t)} - 1, \tag{7}$$

which declines over time and therefore admits time-inconsistent preferences under the assumption of an additively-separable intertemporal utility function.

The final discounting function which will be used in this chapter was originally proposed by Read (2001, equation 16, p. 25) and has been dubbed the "Weibull" (WB) discounting function by Jamison and Jamison (2011, p. 5) because it has an associated Weibull probability density function. The discount factor for the WB specification is:

$$D^{WB}(t) = \exp(-\delta t^{(1/\beta)}), \tag{8}$$

for $\delta > 0$ and $\beta > 0$. Note that when $\beta = 1$, (8) collapses to the E specification so, in the terminology of Jamison and Jamison (2011, p. 25), the parameter β either "expands"

or "contracts" time. When $\beta > 1$, it is as if time has contracted or is perceived to be "slowing down" by the individual. By contrast, when $\beta < 1$, it is as if time has expanded or is "speeding up" as perceived by the individual.[21] The discount rate for the WB model is:

$$d^{WB}(t) = \exp(\delta t^{(1-\beta)/\beta}) - 1, \qquad (9)$$

which is constant over time when $\beta = 1$, declines over time when $\beta > 1$, and increases over time when $\beta < 1$. AHLR (Appendix D, p. A35) draw the analogy between subjective distortions of probabilities in the literature on choice under risk and subjective distortions of time as captured by β in the WB model. Just as people may perceive probabilities as different to their objective values when making choices under risk, they may perceive time to be faster or slower when making choices over time-dated utility flows. The WB model incorporates this possibility in a flexible specification which may shed further light on the relationship between smoking and discounting behaviour.

In this section I have discussed four models which have been proposed to explain the way in which people discount future utility flows. These models are but a small subset of those which exist in the literature (see, for example, Harvey (1986), Loewenstein and Prelec (1992), Bleichrodt, Rohde and Wakker (2009), Benhabib, Bisin and Schotter (2010)) but they were chosen for good reason. The E, H, and QH models are the most prominent in the economics, psychology, and addiction literatures so they should be included for the sake of comparison with previous studies. The WB model was included due to its flexibility (i.e., it admits constant, increasing, and decreasing discount rates) and because it incorporates the notion of subjective time perception which may be an important driver of behaviour and may shed light on the relationship between smoking and discounting.

[21] Read (2001, p. 25) credits Green, Fry and Myerson (1994) with the notion that subjective time perception influences discounting behaviour.

IV. THE UCLA AND USC DISCOUNTING STUDIES

Researchers at the University of California Los Angeles (UCLA) and the University of Southern California (USC) collected the discounting data which will be analysed in this chapter.

The UCLA Consortium for Neuropsychiatric Phenomics (CNP) is a large, collaborative research project focussing on the genetic and environmental bases of variation in psychological and neural system phenotypes. It was one of nine interdisciplinary research consortia supported by the National Institutes of Health (NIH) Roadmap Initiative starting in 2007 (Bilder et al. (2009)). Discounting data was obtained as part of this ambitious project.

Study participants were recruited from the Los Angeles metropolitan area using print advertisements, flyers, internet postings (e.g., craigslist.org) and presentations by study investigators. All potential participants underwent telephone screening followed by additional in-person screening to ensure that they satisfied study inclusion/exclusion criteria. Subjects were required to be men or women between the ages of 21 and 50; speak English or Spanish; have at least 8 years of formal education; have no significant medical illness; test negative for drugs of abuse; have visual acuity of 20/60 or better; and have no Axis I or II disorders as defined by the DSM-IV.

Following a thorough explanation of the research project, participants gave written informed consent and completed a number of tasks and questionnaires (consult Bilder et al. (2009) for more information on the other tasks and questionnaires). All procedures were approved by the Institutional Review Boards at UCLA and the Los Angeles County Medical Department of Mental Health. Smoking status was assessed by self-report; specifically, subjects answered the question, “Have you ever smoked cigarettes on a daily basis?” by choosing “no” “yes, current” or “yes, past.” Subjects who self-identified as current smokers were also asked how many cigarettes they smoked per day, on average, and the length of time they had been smokers. The mean number of cigarettes smoked per day is 11.13 with a standard deviation of 18.99. The

mean smoking duration is 9.77 years with a standard deviation of 7.34.[22] Thus, the UCLA sample provides information on current and former smoking status, as well as the severity and duration of smoking among current smokers.

For the USC study, recruitment took place through internet postings on craigslist.org and was confined to the Los Angeles metropolitan area. The postings attracted 380 individuals who signed informed consent to be interviewed for inclusion in the study. Of these people, 229 self-identified as current cigarette smokers and 109 self-identified as never-smokers. As with the UCLA study, participants were screened for Axis I and II pathologies, and subjects meeting diagnostic criteria for Axis I or II disorders were excluded. From the original recruitment pool, 88 smokers and 86 never-smokers met all inclusion/exclusion criteria and completed a demographic questionnaire, a discounting task, and a neuroimaging task. One of the subject's choices across both tasks was randomly selected for payment using gift cards. If the randomly selected choice had an associated delay, that gift card was activated on the corresponding day. The neuroimaging task will not be discussed further.

In the USC study, all smokers were nicotine-dependent according to DSM-IV criteria, as assessed using the Mini-International Neuropsychiatric Interview (M.I.N.I.) (Sheehan et al. (1998)). Additionally, smoker inclusion required: 1) self-reported smoking of at least 15 cigarettes per day for at least two years; and 2) biochemical confirmation of smoking by either carbon monoxide in expired breath of at least 15 ppm, or a positive urinalysis (cutoff level = 200ng/ml). Thus, subjects in the USC sample were either heavy, current smokers or never-smokers. The mean number of cigarettes smoked per day is 19.50, with a standard deviation of 4.27, and the mean smoking duration is 11.78 years, with a standard deviation of 7.41. Mann-Whitney tests confirm that smokers in the USC study smoked significantly more cigarettes per day ($p < 0.001$) and smoked for a significantly longer period of time ($p = 0.055$) than smokers in the UCLA sample.

[22] In the health sciences literature, smoking status is often defined on the basis of a 100-cigarettes-in-a-lifetime threshold. In the UCLA sample, 71 out of the 75 self-identified current smokers satisfy this criterion, 2 did not provide information on smoking duration so they cannot be classified, and 2 do not satisfy this criterion as they only started smoking in the month prior to the study.

TABLE II
SUMMARY STATISTICS OF UCLA AND USC STUDIES

Variable	UCLA (N = 1031) Mean (Std deviation)	USC (N = 174) Mean (Std deviation)	Combined (N = 1205) Mean (Std deviation)
Age[a]	31.451	33.810	31.792
	(8.526)	(8.574)	(8.570)
Hispanic[b]	0.403	0.121	0.363
	(0.491)	(0.327)	(0.481)
Male[d]	0.455	0.552	0.469
	(0.498)	(0.499)	(0.499)
Years of education[c]	15.045	14.884	15.022
	(1.992)	(2.111)	(2.009)
Smoker: past	0.202	0	0.173
	(0.401)	(0)	(0.378)
Smoker: present[b]	0.073	0.506	0.135
	(0.260)	(0.501)	(0.342)

[a]Significantly different across samples at the 1% level according to a Mann-Whitney test
[b]Significantly different across samples at the 1% level according to a χ^2 test
[c]Significantly different across samples at the 5% level according to a Mann-Whitney test
[d]Significantly different across samples at the 5% level according to a χ^2 test

Summary statistics for the two samples are presented in Table II. In the combined sample, approximately 36% of subjects are Hispanic, 47% are male, 17% are past or ex-smokers, and 14% are current smokers. In addition, subjects are approximately 32 years old and have 15 years of education on average. The UCLA and USC samples differ markedly according to these demographic characteristics. The USC sample is significantly older, has a higher proportion of men, a lower proportion of Hispanics, fewer years of education on average, and a far greater proportion of current smokers. Clearly it will be important to control for these differences in statistical models of the relationship between smoking and discounting behaviour.

Discounting data in both studies was collected using a questionnaire originally developed by Kirby and Maraković (1996) and later refined by KPB. Subjects had to make 27 choices between SS and LL monetary rewards. The monetary rewards were entirely hypothetical in the UCLA study whereas subjects had a chance of receiving payment for one of their choices in the USC study. To accommodate this difference in payment scheme and other potential differences across studies, a fixed effect dummy variable indicating study site (i.e., UCLA or USC) will be included in the statistical analyses.

TABLE III
DELAY DISCOUNTING TASK WITH FRACTION OF SUBJECTS CHOOSING LL AMOUNT

	Reward values			δ_H at	δ_H	δ_E at	LL	Fraction choosing LL	
Order	SS	LL	Delay[1]	indiff.	rank	indiff.	size	Smokers	Non-smokers[2]
13	$34	$35	186	0.00016	1	0.00016	S	0.01	0.03
1	$54	$55	117	0.00016	1	0.00016	M	0.08	0.11
9	$78	$80	162	0.00016	1	0.00016	L	0.03	0.05
20	$28	$30	179	0.0004	2	0.0004	S	0.02	0.08
6	$47	$50	160	0.0004	2	0.0004	M	0.06	0.07
17	$80	$85	157	0.0004	2	0.0004	L	0.02	0.05
26	$22	$25	136	0.001	3	0.0009	S	0.04	0.08
24	$54	$60	111	0.001	3	0.0009	M	0.05	0.09
12	$67	$75	119	0.001	3	0.0009	L	0.03	0.10
22	$25	$30	80	0.0025	4	0.0023	S	0.06	0.13
16	$49	$60	89	0.0025	4	0.0023	M	0.15	0.21
15	$69	$85	91	0.0025	4	0.0023	L	0.17	0.27
3	$19	$25	53	0.006	5	0.005	S	0.18	0.29
10	$40	$55	62	0.006	5	0.005	M	0.13	0.23
2	$55	$75	61	0.006	5	0.005	L	0.28	0.43
18	$24	$35	29	0.016	6	0.013	S	0.29	0.36
21	$34	$50	30	0.016	6	0.013	M	0.42	0.55
25	$54	$80	30	0.016	6	0.013	L	0.59	0.67
5	$14	$25	19	0.041	7	0.03	S	0.53	0.60
14	$27	$50	21	0.041	7	0.03	M	0.61	0.73
23	$41	$75	20	0.041	7	0.03	L	0.75	0.79
7	$15	$35	13	0.10	8	0.07	S	0.71	0.75
8	$25	$60	14	0.10	8	0.06	M	0.86	0.90
19	$33	$80	14	0.10	8	0.07	L	0.84	0.88
11	$11	$30	7	0.25	9	0.15	S	0.91	0.94
27	$20	$55	7	0.25	9	0.16	M	0.88	0.88
4	$31	$85	7	0.25	9	0.15	L	0.90	0.93

Source: Author's construction but based on Kirby, Petry and Bickel (1999)
Notes: The "Order" column lists the order of each question as presented to subjects. SS refers to the smaller, sooner reward. LL refers to the larger, later reward. "δ_E at indiff." and "δ_H at indiff." list the parameter values at which the SS and LL amounts are of equal value as determined by the exponential (2) and hyperbolic (6) discounting functions, respectively. "δ_H rank" groups the discounting questions according to their implied δ_H indifference values. S, M, and L refer to the small, medium and large LL reward sizes, respectively.

[1]Delays are in days
[2]Non-smokers include both never-smokers (n = 834) and ex-smokers (n = 208).

In the KPB task, the SS reward is always available immediately whereas the delay to the LL reward varies between 7 and 186 days. Of the 27 questions, 9 use small reward magnitudes (LL reward: $25 - $35), 9 use medium reward magnitudes (LL reward: $50 - $60), and 9 use large reward magnitudes ($75 - $85). Table III lists the 27 binary choice questions.

Under the assumption that the utility function $U(\cdot)$ is linear in income y, one can use a participant's choices in Table III to determine a range for the discounting parameters

in the E and H models.[23] Table III includes the values of δ_E and δ_H for each choice question at which the SS and LL rewards are of equal value as determined by the E (2) and H (6) discounting functions, respectively. These values and a subject's pattern of choices can be used to determine his discounting parameters under either specification. For example, suppose that someone chooses the SS reward on the first 3 rows of the table. This implies that the person's discounting parameter, under either the E or H model, is greater than 0.00016. Suppose further that this person switches to the LL reward on row 4 of the table and continues to select the LL reward on every subsequent row. This implies that the person's discounting parameter, under either the E or H model, is less than 0.0004. Thus, this pattern of choices reveals a discounting parameter which lies in the open interval (0.00016, 0.0004). KPB used the geometric mean[24] of these two values (i.e., $(0.00016 \times 0.0004)^{1/2} = 0.00025$) as the person's discounting parameter.

One shortcoming of this approach is that it discards information on the precision of the parameter estimate. In other words, the pattern of choices described above reveals a discounting parameter which lies in the open interval (0.00016, 0.0004). By calculating the geometric mean of this interval (i.e., 0.00025) and using it as the person's discounting parameter, one ignores the uncertainty surrounding this estimate. In other words, a person's "true" discounting parameter may be 0.00036 or it may be 0.00017 and one cannot tell which value it is based on the choices that the person made (i.e., either value is consistent with the DGP). In effect, KPB change interval data into point estimates by implicitly assuming a uniform distribution within the discount rate intervals.

Unfortunately this method is even less tenable in the presence of "inconsistent" choices. As Kirby and Maraković (1996, p. 102) and KPB (p. 81) point out, a person's choices on this task are not always deterministically consistent with a particular

[23] The QH and WB models have two parameters so a participant's choices in Table III cannot pin down a range of values for both parameters. This can be achieved statistically though, as I will discuss in the next section of this chapter.

[24] The stated rationale for using a geometric, rather than an arithmetic, mean was, "… to avoid underweighting the smaller of the two rate parameters" (KPB, p. 80). A geometric mean is useful when calculating a measure of central tendency for two values which have different ranges (i.e., one value falls within the closed interval [0, 3] while the other falls within the closed interval [0, 100]). However, in the context of the KPB task it is unclear how an arithmetic mean underweights the smaller value because the smaller and larger values are given equal weight.

discounting parameter. Building on the example from earlier, suppose that instead of someone switching to the LL reward on row 4 of Table III and then consistently choosing the LL reward on every subsequent row, the person switches back and chooses the SS reward on rows 5, 6, and 7 and then chooses the LL reward for every subsequent row. In such a case, Kirby and Maraković (1996) and KPB calculated the discounting parameter which was consistent with the highest proportion of a subject's choices and used this discounting parameter thereafter. If multiple discounting parameters yielded the same level of consistency, the geometric mean of these values was used as the person's discounting parameter.

From the preceding discussion the problem with this method should be clear. In this case there is even less certainty about a particular parameter estimate and yet it is treated as a datum which is measured without error.[25] The statistical approach which I adopt remedies these problems by recognising the uncertainty of parameter estimates and by incorporating a behavioural error specification which allows for the possibility that subjects make mistakes in their discounting choices.

V. STATISTICAL SPECIFICATION

The statistical method I employ is direct estimation by maximum likelihood of a structural model of a latent choice process. The latent choice process in question is captured by the various discounting models discussed earlier. These models provide the structure necessary to estimate people's time preferences using their choice data. One of the major benefits of the maximum likelihood approach is that it uses all of the available information which the participants' data impart to estimate the discounting parameters and the precision of these estimates. Following Andersen, Harrison, Lau and Rutström (2008), I review the basic logic of the estimation strategy below, focussing on the canonical case of the E model. I will then briefly discuss the straightforward extension to the other discounting models.

[25] As discussed previously, interval data is analysed appropriately using interval regression methods. For a discussion of this point, in the context of experimental studies of choice under risk, consult Harrison and Rutström (2008, §2.1)

Under the E model, δ is the discounting parameter which equalises the *utility* of income received at time t with the *utility* of income received at time $t + \tau$:

$$[1 / (1 + \delta)^{t}]U(y_t) = [1 / (1 + \delta)^{t+\tau}]U(y_{t+\tau}) \quad (10)$$

for some utility function $U(\cdot)$. In the KPB task, the SS reward is available immediately, so (10) simplifies to:

$$U(y_0) = [1 / (1 + \delta)^{\tau}]U(y_\tau) \quad (11)$$

Without data on risk attitudes, it is not possible to jointly estimate the curvature of the utility function $U(\cdot)$ and the discounting parameter δ but one can make parametric assumptions about the form and shape of $U(\cdot)$ and then estimate "profile likelihoods" to determine the optimal shape of $U(\cdot)$ given the data. This will be taken up in Appendix B but for what follows I will assume that $U(\cdot)$ is linear in income y so that (11) simplifies to:

$$y_0 = [1 / (1 + \delta)^{\tau}]y_\tau \quad (12)$$

Note that y_0 in (12) represents the SS reward in the KPB task whereas y_τ represents the LL reward in this task. Thus, δ is the discounting parameter which equalises the *value* of the SS and LL rewards under the assumptions of the E model and linear utility. Stated another way, the left hand side (LHS) of (12) represents the present value (PV) of the SS reward:

$$PV_{SS} = y_0 \quad (13)$$

The right hand side (RHS) of (12) represents the PV of the LL reward:

$$PV_{LL} = [1 / (1 + \delta)^{\tau}]y_\tau \quad (14)$$

To determine the value of δ, the PV of the SS and LL rewards are calculated for an initial estimate of δ and the index below is formed:

$$\nabla PV = PV_{SS} - PV_{LL} \quad (15)$$

Note that this is a latent index, based on latent preferences, which captures the difference in present values of the two rewards presented to subjects. This index is then linked to the subjects' observed choices using the cumulative logistic distribution function $\Lambda(\nabla PV)$. This function takes any argument (∇PV) between $\pm\infty$ and transforms it into a number between 0 and 1. Thus, we have the so-called "logit" link function:

$$\Pr(\text{Choose SS reward}) = \Lambda(\nabla PV) \quad (16)$$

The latent index in (15) is linked to subjects' observed choices by specifying that the SS reward is chosen when $\Lambda(\nabla PV) > ½$, which is precisely what (16) tells us. To see this note that if $PV_{SS} = PV_{LL}$ then $\nabla PV = 0$. Plugging 0 into the cumulative logistic distribution function (i.e., $\Lambda(0)$) yields a value of 0.5. In other words, when the PV of the SS and LL rewards are equal, the probability of choosing the SS reward is 0.5 (i.e., there is an equal chance of choosing the SS reward or the LL reward). By contrast, if $\nabla PV > 0$ then $\Lambda(\nabla PV) > 0.5$ (i.e., the probability of selecting the SS reward is greater than 0.5 when its PV exceeds the PV of the LL reward) and if $\nabla PV < 0$ then $\Lambda(\nabla PV) < 0.5$ (i.e., the probability of selecting the SS reward is less than 0.5 when its PV is less than the PV of the LL reward).

Thus, the likelihood of the observed responses, conditional on the E model being true, depends on the estimates of δ given the statistical model above and the choices of subjects in the discounting task. The conditional log-likelihood of the model is:

$$\ln L_i(\delta; z, X) = \sum_i [(\ln \Lambda(\nabla PV) \times I(z_i = 1)) + (\ln (1 - \Lambda(\nabla PV)) \times I(z_i = 0))], \quad (17)$$

where $I(\cdot)$ is the indicator function, $z_i = 1(0)$ denotes the choice of the SS (LL) reward in choice pair i, and X is a vector of individual characteristics capturing age, gender, education etc.

One of the advantages of structural maximum likelihood estimation is that it is a straightforward extension to make the parameter of interest, the discounting parameter δ, a linear function of individual characteristics. In this case, one estimates $\delta = \delta_0 + \delta_\beta \times X$, where δ_0 is a fixed parameter and δ_β is a coefficient vector linked to the variable vector X of individual characteristics. If no individual characteristics are included in the model we estimate $\delta = \delta_0$, which is the discounting parameter estimated at the level of the sample without taking into account observed, individual heterogeneity (i.e., assuming homogenous preferences). Note that every estimate of δ includes a standard error which reflects our ignorance and uncertainty as to the "true" value of δ. This stands in stark contrast to the deterministic method of calculating a person's discounting parameter which Kirby and Maraković (1996) and KPB employ.

Another important extension to the simple model defined above is to allow for some errors on the part of subjects when they make choices between the SS and LL rewards. This error could be as simple as a "tremble," where, say, a subject wants to choose the SS reward but mistakenly selects the LL reward instead. A different behavioural error specification introduced by Fechner (1966) uses the latent index:

$$\nabla PV = (PV_{SS} - PV_{LL}) / \mu, \qquad (18)$$

rather than (15). Note how different values of μ affect our ∇PV index. As $\mu \rightarrow 0$ our specification collapses to a deterministic choice model where the choice is strictly determined by the PV of the two rewards. However, as $\mu \rightarrow \infty$, $\nabla PV \rightarrow 0$, and a subject's choice is essentially random (i.e., the probability of selecting either reward is 0.5). When $\mu = 1$ we are back to specification (15), so the Fechner error term is a parameter which basically flattens the logit link function as its value increases. In this chapter I will estimate the E model with and without the Fechner error term to highlight the importance of incorporating it; I will always include the error term in the other discounting models. Note that by using a behavioural error term one formally accommodates subjects' inconsistent choices on the KPB task, rather than letting this inconsistency be soaked up by the standard errors on the discounting parameter estimates. Thus, the Fechner error term allows one to model behavioural errors directly instead of capturing these errors as part of a statistical error term.

Fortunately it is a simple matter to incorporate other discounting models in this statistical framework. Assuming linear utility, one simply replaces the E model's discount factor in (12) with the discount factors for the other specifications. In the case of the WB model, for example, (12) becomes:

$$y_0 = [\exp(-\delta\tau^{(1/\beta)})]y_\tau \qquad (19)$$

One then adjusts (13) and (14) appropriately to incorporate this new expression, and to form the latent index in (18), so as to proceed as before.

VI. RESULTS

The primary focus of this chapter is to explore differences in discounting behaviour between smokers and non-smokers when applying the appropriate statistical tools to the data generated by the KPB discounting task. The second aim of this chapter is to highlight the methodological issues involved in drawing accurate inferences about the smoking-discounting relationship.

A. Baseline Estimates

The final columns of Table III show the fraction of smokers and non-smokers[26] who selected the LL reward for each choice question. As is evident, the fraction of non-smokers selecting the LL reward is greater than the fraction of smokers selecting the LL reward for every question other than number 27, which is in the second last row of the table, where the fractions are equal. This provides preliminary evidence of differences in the discounting behaviour of smokers and non-smokers. Before any definitive conclusions are reached though, the data must be subjected to appropriate statistical analysis.

[26] As discussed previously, non-smokers include both never-smokers and ex-smokers. In subsequent analyses I distinguish between never-smokers and ex-smokers.

Table IV presents baseline estimates of the E, H, QH, and WB discounting models. Note that the results pool choices across all individuals, which means I am estimating the value of δ_0 for the E and H models, and δ_0 and β_0 in the case of the QH and WB models, for the sample as a whole. In other words, I am assuming homogenous preferences. Note further that the results account for clustering at the individual level which adjusts the standard errors of the estimates to take into account the fact that each respondent made multiple choices across 27 questions. Finally, note that I have estimated daily discount rates so as to conform to the bulk of studies in the smoking-discounting literature, and due to a puzzling inferential issue which I investigate in Appendix A.

TABLE IV: DISCOUNTING FUNCTION ML ESTIMATES
HOMOGENOUS PREFERENCES AND HOMOSKEDASTIC ERRORS

	Model 1	**Model 2**	**Model 3**	**Model 4**	**Model 5**
	Exponential No error	Exponential Fechner error	Hyperbolic Fechner error	Quasi-Hyperbolic Fechner error	Weibull Fechner error
Discounting parameter (δ)	0.00510***	0.00955***	0.01283***	0.00432***	0.02160***
	(0.00008)	(0.00044)	(0.00066)	(0.00037)	(0.00441)
Discounting parameter (β)				0.80607***	1.26329***
				(0.01384)	(0.08136)
Error (μ)		15.04102***	13.54160***	10.85905***	13.55935***
		(0.38471)	(0.33044)	(0.40535)	(0.49900)
N	32535	32535	32535	32535	32535
log-likelihood	-87637.774	-14722.578	-14737.615	-14694.619	-14714.908

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The E model results are sensible and fall well within the range of estimates implied by the questions on the task: 0.00016 - 0.16 (see Table III). Note that the estimate of the discounting parameter δ is almost twice as large when we incorporate the Fechner error term in the model. Note further that the estimate of the error term μ is large and statistically significant. Finally, the log-likelihood of the model which includes the Fechner error term is far higher than the log-likelihood of the model which does not. These results highlight the importance of incorporating a behavioural error specification when estimating models of discounting behaviour and thus, every subsequent analysis includes this error term.

The estimate of δ in the H model also falls well within the range of estimates implied by the task: 0.00016 – 0.25. The log-likelihood of the H model is less than the log-likelihood of the E model, which provides preliminary evidence that the E model better characterises these data. This is an important finding given that 22 of the studies in Table I just assumed that discounting was hyperbolic and did not estimate the E model at all. This result, that E outperforms H, will be subjected to closer scrutiny later when I conduct non-nested model selection tests and estimate mixture models of discounting behaviour.

The QH results reveal significant evidence of quasi-hyperbolic discounting and, thus, declining discount rates. Specifically, the estimate of $\beta = 0.806$ is less than 1 at any standard level of significance ($p < 0.001$). The estimate of $\delta = 0.004$ is approximately half the estimate of δ in the E model, which suggests that part of the discounting of future rewards is attributable to a present-bias in decision making. Finally, the log-likelihood for the QH model is the highest of all the models that were estimated, which suggests that it provides the best fit to these data.

The WB results suggest that subjective time perception plays an important role in discounting behaviour. The estimate of $\beta = 1.263$, which is significantly greater than 1 ($p < 0.01$), implies that subjects tend to perceive time as "slowing down" and this generates declining discount rates over time. Thus, both the QH and WB results suggest that discount rates decline over time which, under the assumption of an additively-separable intertemporal utility function, raises the spectre of time-inconsistent choices. However, the two models provide competing explanations for time inconsistency: a present-bias in the case of the QH model and subjective time perception in the case of the WB model.

B. Non-Nested Model Selection Tests

A comparison of the E, H, QH, and WB models' log-likelihoods is a somewhat crude measure of their ability to explain the data. The hypothesis that one model better characterises the data than another can be tested formally using Vuong (1989) and Clarke (2007) non-nested model selection tests.[27]

Vuong (1989) developed a model selection test for non-nested models which can be used to adjudicate between the various discounting functions. The Vuong (1989) test uses the likelihoods for each observation under each model and forms the ratio of these values. Thus, for observation *i* one forms the ratio of the likelihood from, say, the E model with the likelihood from the H model. One then calculates the natural logarithm of this ratio and tests whether the expected value of these log-ratios is zero over the entire sample. This produces a test statistic which is normally distributed under fairly general conditions and allows one therefore to compare the models and determine which model better characterises the data.

Clarke (2007) developed an alternative test for non-nested model selection. He showed that this test is asymptotically more efficient and has greater power in discriminating between models than the Vuong (1989) test when the distribution of the individual log-ratios is highly peaked. In other words, when the distribution of these log-ratios is leptokurtic[28], the Clarke (2007) test is superior, from both statistical efficiency and power perspectives, to the Vuong (1989) test.

Figure I plots the distribution of the individual log-ratios, with a normal density overlay, for the six discounting model comparisons. The distribution of these log-ratios is leptokurtic which suggests that the Clarke (2007) test is more appropriate for these data. The Clarke (2007) test yields a test statistic based on the binomial distribution which must be compared to a critical value to determine which model, in

[27] The QH and WB models are nested in the E model, but not in each other. We know from earlier that if $\beta = 1$ in either the QH or WB models, these models collapse to the E model. So an obvious test of whether discounting is, say, E or WB is to test whether $\beta = 1$ in the WB model, as I did earlier. Nevertheless, it is still worthwhile to conduct formal model selection tests even if models are nested in each other, as will become clear later.

[28] The normal distribution is the quintessential mesokurtic distribution. A distribution which has positive excess kurtosis (i.e., a highly peaked distribution) is leptokurtic.

a pairwise comparison, receives the most support in the data. A Clarke (2007) test comparing the E and H models yields a test statistic of 17737, which is above the critical value of 16267.5, implying that the E model better characterises the data ($p < 0.001$). This result provides formal confirmation of the earlier finding that the E model provides a better fit to these data than the H model, which was based on a comparison of their log-likelihoods. To my knowledge, this is the first valid statistical test of the E and H models in the smoking-discounting literature.

Figure I: Distribution of individual log-ratios

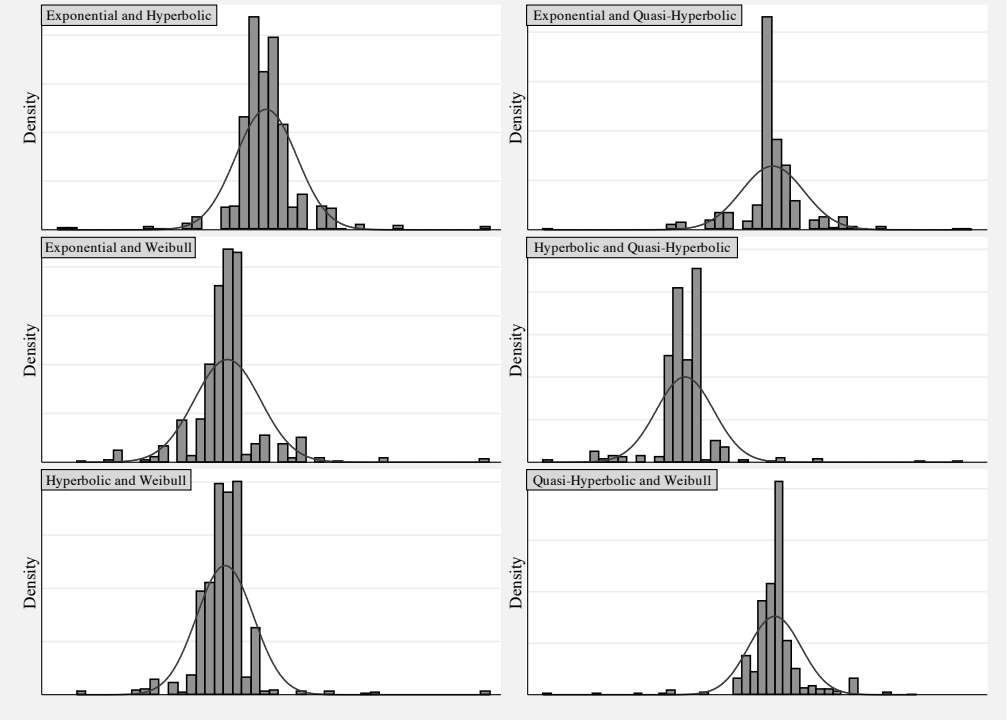

A Clarke (2007) test comparing the E and QH models finds in favour of the QH model ($p < 0.001$) whereas the test comparing the E and WB models finds in favour of the E model ($p < 0.001$). This last result, that E outperforms WB, is revealing because when estimating the WB model previously we could easily reject the hypothesis that $\beta = 1$. Thus, even though there was significant evidence of subjective time perception in the WB model, the E model better characterises the data, at least according to the Clarke (2007) test. This highlights the value of conducting model selection tests prior to claiming that one model is the "winner" on the basis of a

parameter test or a comparison of the models' log-likelihoods[29]; the log-likelihood of the WB model was higher than the E model.

A Clarke (2007) test comparing the H and QH models suggests that the H model finds more support in the data ($p < 0.001$) and the test of the H and WB models also finds in favour of the H model ($p < 0.001$). Finally, a Clarke (2007) test of the QH and WB models suggests that the QH model better characterises the data ($p < 0.001$). The preceding results show quite clearly that the WB model has the least support in the data because the E, H, and QH models outperform it on the Clarke (2007) test. Finding a transitive ranking of the other discounting models though, is not possible. We found that the E model outperforms H, that H outperforms QH, but that QH outperforms E. This yields an intransitive or cyclic ranking: E > H > QH > E. Thus, although the Clarke (2007) test provides a formal method for comparing models it is not immune to cyclical rankings which clearly makes it difficult to claim that one model is the obvious winner in terms of goodness of fit.

C. Smoking and Discounting Behaviour

To analyse the link between smoking and discounting behaviour, one can make the parameters of interest in the four discounting models a linear function of smoking status. By estimating these models one captures the "total effect" of smoking status on discounting behaviour without controlling for any potential differences between smokers and non-smokers, like age, education, and gender, and any potential differences between the UCLA and USC study sites. In other words, these models provide a first pass at exploring this relationship.

Table V presents the results from the E, H, QH, and WB models where the discounting parameters are estimated as a linear function of smoking status. In the E model, current smokers discount significantly more heavily than both never-smokers ($p < 0.01$), the omitted base category, and ex-smokers ($p < 0.01$). To see the economic significance of this result, note that the coefficient estimate for current smokers of 0.006 means that they have a daily discount rate approximately 0.5% higher than

[29] As I will discuss later, there is also little value in making strong claims about the superiority of one model over another when multiple decision-making processes may be present in the data.

never-smokers. Note that there is no significant difference in the discounting behaviour of ex-smokers and never-smokers, which echoes the results of BOM and Sweitzer et al. (2008).

TABLE V: DISCOUNTING FUNCTION ML ESTIMATES
SMOKING TOTAL EFFECT

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Discounting parameter (δ)				
Smoker: past	0.00064	0.00114	-0.00025	0.01004
	(0.00103)	(0.00164)	(0.00068)	(0.01301)
Smoker: present	0.00585***	0.00887***	0.00225**	0.01191
	(0.00168)	(0.00261)	(0.00112)	(0.01291)
Constant	0.00882***	0.01168***	0.00420***	0.01857***
	(0.00048)	(0.00071)	(0.00041)	(0.00444)
Discounting parameter (β)				
Smoker: past			-0.03434	0.14817
			(0.03202)	(0.19990)
Smoker: present			-0.0441	0.02298
			(0.03807)	(0.17206)
Constant			0.82074***	1.23345***
			(0.01641)	(0.08854)
Error (μ)				
Constant	14.91054***	13.42915***	10.85760***	13.46171***
	(0.37865)	(0.32563)	(0.40395)	(0.49746)
N	32535	32535	32535	32535
log-likelihood	-14667.145	-14683.557	-14637.089	-14658.484

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In the H model, current smokers discount significantly more heavily than never-smokers ($p < 0.01$) and ex-smokers ($p < 0.01$) but there is no significant difference between ex-smokers and never-smokers. In the QH model, the estimate of δ, sometimes referred to as the "long-term discount rate," is significantly higher for current smokers than never-smokers ($p < 0.05$) and ex-smokers ($p < 0.05$). The estimate of β though, which reflects a present-bias, does not differ significantly among current smokers, never-smokers and ex-smokers. Finally, in the WB model, there are no significant differences between current smokers, never-smokers, and ex-smokers in the estimate of δ. Similarly, there are no significant differences in subjective time perception (i.e., in the estimate of β) between these groups. However, a test of the joint hypothesis that current smoking status has no effect on discounting across both equations can be rejected ($p = 0.015$). These results imply that the effect of current smoking status in the WB model cannot be confined to particular parameter

estimates but, rather, that current smoking status has a statistically significant effect on the model as a whole.

These results provide a simplistic - in the sense that the models only include the smoking status covariate - but nuanced view of the relationship between smoking and discounting behaviour. While the estimate of δ is significantly higher among current smokers than never-smokers and ex-smokers in the E, H, and QH models, this is not the case in the WB model. Furthermore, current smokers do not differ significantly from never-smokers and ex-smokers in terms of present-bias in the QH model, nor subjective time perception in the WB model. Clearly these results are only preliminary because we need to control for a number of factors which may mediate the relationship between smoking and discounting behaviour.

Table VI presents the results from the E, H, QH, and WB models that take into account observed, individual heterogeneity by conditioning the discounting parameter estimates on a set of covariates and task parameters. Specifically, the models include the demographic variables from Table II as well as a fixed effect dummy variable for study site and two dummy variables for LL reward size. In addition, the models incorporate heteroskedasticity by making the Fechner error term a linear function of study site[30] and the subject's decision time on each question in the task.[31]

In the E model, current smokers discount significantly more heavily than never-smokers ($p < 0.05$), even after controlling for other potential differences between these groups. However, a Wald test of the hypothesis that current smokers discount more heavily than ex-smokers fails to reach significance ($p = 0.103$). There is also evidence of a magnitude effect in discounting behaviour in that higher LL reward sizes are associated with significantly lower discounting. Years of education has a negative impact on discounting whereas age and Hispanic ethnicity tend to increase

[30] As discussed previously, the UCLA and USC study sites used different payment schemes for the discounting task. In addition, the study sites employed different researchers. It is natural, therefore, to allow for heteroskedasticity by study site.

[31] Hey (1995) found that a subject's decision time on each question of a risk preference task was negatively and significantly associated with the estimate of a Fechner error term. In other words, people who spent longer on each question in the task made fewer behavioural errors than people who went through each question at a faster pace. I use the natural logarithm of decision time because this variable was measured in milliseconds and, thus, it has an enormous range.

the discounting of delayed rewards. Finally, the USC study site is associated with significantly fewer behavioural errors; this is not surprising given that subjects were incentivised at this site.

TABLE VI: DISCOUNTING FUNCTION ML ESTIMATES
HETEROGENOUS PREFERENCES AND HETEROSKEDASTIC ERRORS

	Model 1		**Model 2**	
	Exponential		Hyperbolic	
	Estimate	Std error	Estimate	Std error
Discounting parameter (δ)				
Age	0.00007*	0.00004	0.00012**	0.00006
Hispanic	0.00121*	0.00071	0.00222**	0.00109
Male	0.00017	0.00056	0.00027	0.0008
Years of education	-0.00051***	0.00015	-0.00080***	0.00019
Smoker: past	0.00033	0.00074	0.00057	0.00109
Smoker: present	0.00278**	0.00136	0.00418**	0.00211
LL size: medium	-0.01327***	0.00069	-0.01925***	0.00117
LL size: large	-0.01697***	0.00084	-0.02450***	0.00141
USC sample	0.00122	0.00118	0.0017	0.0017
Constant	0.02804***	0.00262	0.03971***	0.00358
Error (μ)				
Decision time (ln)	-0.27509	0.35758	-0.20374	0.32448
USC sample	-1.88827**	0.9091	-1.93433**	0.81023
Constant	16.51280***	2.8729	14.63170***	2.61076
N	32508		32508	
log-likelihood	-14064.931		-14121.63	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The pattern of results in the H model is qualitatively identical to the E model. Specifically, current smokers discount significantly more than never-smokers ($p < 0.05$) but not ex-smokers ($p = 0.116$). LL reward size and years of education have a negative impact on discounting whereas age and Hispanic ethnicity tend to increase discounting. The USC study site is also associated with fewer behavioural errors.

In the QH model, with a full set of covariates, smoking status is not significantly associated with the long-term discount rate δ nor the extent of present-bias β. The point estimates are in the expected direction (i.e., positive for δ signifying higher long-term discounting, and negative for β signifying more present-bias) but they fail to reach significance. Furthermore, a test of the joint hypothesis that smoking has no effect on discounting across both equations, cannot be rejected ($p = 0.189$).

In the WB model, by contrast, the estimate of δ is significantly higher among current smokers than never-smokers ($p < 0.001$) and ex-smokers ($p < 0.01$). In addition, the

estimate of β is significantly higher among current smokers than never-smokers ($p < 0.001$) and ex-smokers ($p < 0.01$). Thus, current smokers in the WB model discount the future more heavily than never-smokers and ex-smokers and current smokers are more likely to perceive time as "slowing down" than never-smokers and ex-smokers.

TABLE VI: DISCOUNTING FUNCTION ML ESTIMATES (CONTINUED)
HETEROGENOUS PREFERENCES AND HETEROSKEDASTIC ERRORS

	Model 3		**Model 4**	
	Quasi-Hyperbolic		Weibull	
	Estimate	Std error	Estimate	Std error
Discounting parameter (δ)				
Age	-0.00006***	0.00002	-0.00013***	0.00004
Hispanic	-0.00118**	0.00047	0.04797***	0.01214
Male	0.00031	0.0004	0.00152	0.00124
Years of education	0.00024**	0.00011	0.00019***	0.00007
Smoker: past	0.00008	0.00055	0.00397	0.0068
Smoker: present	0.00083	0.00067	0.06284***	0.01751
LL size: medium	-0.03089***	0.00255	0.07314***	0.02294
LL size: large	-0.03395***	0.00268	0.19204***	0.0454
USC sample	-0.00194***	0.00057	0.00031	0.00178
Constant	0.03475***	0.0035	0.00360**	0.00169
Discounting parameter (β)				
Age	-0.00518***	0.0013	-0.00630***	0.00154
Hispanic	-0.10241***	0.0249	0.64905***	0.11729
Male	0.0141	0.02286	0.03721	0.04399
Years of education	0.02881***	0.00617	0.03523***	0.00916
Smoker: past	-0.00827	0.03327	0.09659	0.14436
Smoker: present	-0.0056	0.0354	0.68462***	0.12416
LL size: medium	-0.59887***	0.06211	1.53727***	0.35252
LL size: large	-0.64385***	0.06668	6.12850**	2.67538
USC sample	-0.12031***	0.03149	-0.0498	0.05762
Constant	1.20748***	0.1384	0.31942**	0.14713
Error (μ)				
Decision time (ln)	-0.16724	0.26206	-0.2609	0.24016
USC sample	-3.03018***	0.66895	-1.60585***	0.55084
Constant	12.80384***	2.1619	12.05736***	2.0211
N	32508		32508	
log-likelihood	-13582.172		-13764.512	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In sum, the preceding results provide a complex perspective on the relationship between smoking and discounting behaviour. In the E, H, and WB models, current smokers tend to discount the future more heavily than never-smokers, but current smokers only discount more heavily than ex-smokers in the WB model. Current smokers are also more likely to perceive time as "slowing down" in the WB model than ex-smokers and never-smokers. By contrast, there are no significant differences between current smokers, ex-smokers, and never-smokers in the QH model.

D. Mixture Models of Discounting Behaviour

The analyses conducted thus far have been based on the implicit assumption that the observations are produced by only one DGP (i.e., E, H, QH, or WB) when more may be present in the data. In other words, the E model may explain some discounting choices better than the H model whereas the H model may explain other choices better than the E model. The assumption that only one DGP characterises all of the data, clearly precludes such a possibility.

Finite mixture models[32] allow two or more DGPs to account for the data and also provide a measure of the proportion of the data which is explained by each process. In the current context, one can estimate a mixture model of, say, the E and H discounting functions and then ask the data to determine how much support each function has. To do so one specifies a "grand likelihood" function which is just a probability-weighted average of the likelihoods of the two models.

Letting π^E represent the probability that the E model is correct, and $\pi^H = (1 - \pi^E)$ the probability that the H model is correct, the grand likelihood is the probability-weighted average of the two conditional likelihoods L^E and L^H for the E and H models, respectively. Thus, the likelihood for the mixture model is given by:

$$\ln L_i(\delta_E, \delta_H, \mu, \kappa; z, X) = \Sigma_i \ln [(\pi^E \times L^E) + (\pi^H \times L^H)] \qquad (20)$$

where κ is a parameter which defines the log odds of the probability of the E model: $\pi^E = 1 / (1 + \exp(\kappa))$. Note that this transformation allows the parameter κ to take on any value during the maximisation process but it constrains the probabilities π^E and π^H to lie within the unit interval. The grand likelihood in (20) is maximised to estimate the parameters of each model and the weight accorded to each model in the data.

[32] For detailed discussions of mixture models consult McLachlan and Peel (2000), Harrison and Rutström (2009), and Conte, Hey and Moffat (2011).

Table VII presents estimates of the mixture model of the E and H discounting functions, assuming homogenous preferences and homoskedastic errors.[33] The estimate of $\pi^E = 0.537$ implies that the E model accounts for approximately 54% of the choices in the data; the H model therefore accounts for roughly 46% of the choices. A hypothesis test that $\pi^E = 0.5$ cannot be rejected at any regular level of significance ($p = 0.204$ as noted in the table).[34] Thus, the E and H discounting models both account for approximately half of the choices in the dataset. This result implies that it is a mistake to assume that only one DGP characterises the data.

TABLE VII: MIXTURE MODEL ML ESTIMATES
HOMOGENOUS PREFERENCES AND HOMOSKEDASTIC ERRORS

	Estimate	Std error	*p*-value	95% Confidence interval	
		Exponential discounting model			
Discounting parameter (δ_E^{mix})	0.00317***	0.00021	<0.001	0.00277	0.00357
Mixture probability (π^E)	0.53692***	0.02909	<0.001	0.4799	0.59394
		Hyperbolic discounting model			
Discounting parameter (δ_H^{mix})	0.06232***	0.00738	<0.001	0.04784	0.07679
Mixture probability (π^H)	0.46308***	0.02909	<0.001	0.40606	0.5201
		Fechner error term			
Error (μ)	8.37233***	0.17598	<0.001	8.02742	8.71724
N	32535				
log-likelihood	-14427.64				
H_0: $\pi^E = 0.5$, *p*-value = 0.20440					

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In addition, it suggests that model selection tests are fraught with difficulty because each model has to account for all the data. To use an analogy, assuming that the E or H model accounts for every discounting choice is like embarking on a DIY project

[33] Appendix C contains the results from all of the two process mixture models that can be estimated from the four discounting models used in this chapter. I only present the results from the E and H mixture model in this section because these are the most commonly used discounting functions in the addiction literature and they are representative of the results from the other mixture models.
[34] The mixture probabilities in all of the mixture models estimated in Appendix C are close to 0.5. However, in the E and QH mixture model, the QH mixture probability is significantly greater than 0.5 ($p < 0.01$). Similarly, in the H and QH mixture model, the QH mixture probability is significantly greater than 0.5 ($p < 0.05$).

armed only with a hammer. If a model selection test favours the E model over the H model, all that this tells us is that the E model is better at explaining all the data. This is similar to the finding that a hammer is more useful than a wrench for most DIY jobs. But if there are multiple DGPs then we should respect this feature of the data and estimate a mixture of these processes, just like one should use a wrench to tighten a bolt and a hammer to knock in a nail.

The mixture model in Table VII also shows how discounting parameter estimates are distorted when the E or H models have to account for all the data. Figure II plots the discounting parameter estimates from the E and H models in Table IV (i.e., when they are assumed to account for all the data) and the discounting parameter estimates from the mixture model in Table VII.

Figure II: Exponential and hyperbolic discount factors

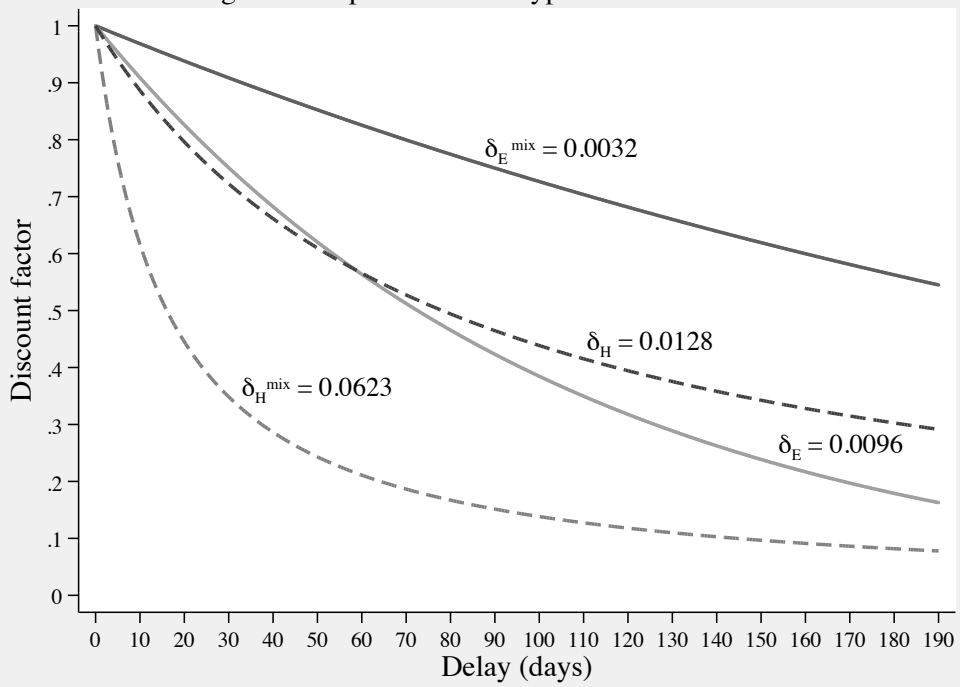

In Model 2 of Table IV, where the E model was assumed to be the sole DGP, the estimate of $\delta_E = 0.0096$. In the mixture model, the estimate of δ_E, which I refer to as δ_E^{mix} in Table VII and Figure II, is far lower at 0.0032. This implies that when one tries to make all the data fit the E model, one inflates the estimate of the discounting parameter. Similarly, in Model 3 of Table IV, where the H model was assumed to be

the sole DGP, the estimate of $\delta_H = 0.0128$. In the mixture model in Table VII, the estimate of δ_H^{mix} is far higher at 0.0623. Thus, by assuming one DGP we are averaging the estimates that we derive when allowing multiple DGPs to characterise the data.

Finally, notice that the estimate of the Fechner error term $\mu = 8.372$ in the mixture model in Table VII is far lower than the estimates of μ for the E and H models in Table IV. Thus, what was being captured as subject errors in decision making when estimating the E and H models separately is partly the product of forcing the data to fit one DGP.

Mixture models also allow one to explore the hypothesis that current smokers are more prone to time inconsistency than ex-smokers and never-smokers, under the assumption of an additively-separable intertemporal utility function, by making the mixture probability a function of observable characteristics and task parameters. Recall that with an additively-separable intertemporal utility function the E model implies time-consistent preferences whereas the H model may yield time-inconsistent choices. If one finds that current smokers are more likely to discount according to the H model as opposed to the E model, this implies that they may be more prone to time inconsistency.

Table VIII presents estimates of the mixture model of the E and H discounting functions with a heterogenous mixture probability, heterogenous preferences, and heteroskedastic errors. The results show that there are no significant differences in the likelihood that current smokers, ex-smokers and never-smokers discount according to the E model. Thus, under the assumption of an additively-separable intertemporal utility function, current smokers are no more prone to time inconsistency than ex-smokers and never-smokers. However, large LL rewards and more years of education increase the likelihood of discounting according to the E model, and thus decrease the likelihood of discounting according to the H model. By contrast, people who took part in the USC study are more likely to discount hyperbolically.

TABLE VIII: MIXTURE MODEL ML ESTIMATES
HETEROGENOUS PREFERENCES AND MIXTURE PROBABILITY AND
HETEROSKEDASTIC ERRORS

	Estimate	Std error	*p*-value	95% Confidence Interval	
Discounting parameter (δ_E^{mix})					
Age	-0.00002	0.00002	0.26122	-0.00006	0.00002
Hispanic	-0.00049	0.00039	0.20272	-0.00125	0.00027
Male	0.00021	0.00034	0.52564	-0.00045	0.00087
Years of education	0.00007	0.00008	0.39723	-0.00009	0.00023
Smoker: past	0.00018	0.00044	0.67261	-0.00067	0.00104
Smoker: present	0.00110*	0.00058	0.05813	-0.00004	0.00225
LL size: medium	-0.01484***	0.00148	0.00000	-0.01774	-0.01194
LL size: large	-0.01487***	0.00126	0.00000	-0.01734	-0.01241
USC sample	-0.00030	0.00056	0.58801	-0.00141	0.00080
Constant	0.01773***	0.00201	0.00000	0.01380	0.02166
Discounting parameter (δ_H^{mix})					
Age	0.00051	0.00036	0.15314	-0.00019	0.00121
Hispanic	0.01100*	0.00592	0.06300	-0.00060	0.02260
Male	0.00158	0.00584	0.78628	-0.00985	0.01302
Years of education	-0.00363***	0.00089	0.00005	-0.00538	-0.00188
Smoker: past	0.00619	0.00810	0.44452	-0.00968	0.02207
Smoker: present	0.01756	0.01190	0.13993	-0.00576	0.04088
LL size: medium	0.00671*	0.00355	0.05883	-0.00025	0.01366
LL size: large	0.08974***	0.01725	0.00000	0.05593	0.12356
USC sample	0.00674	0.00885	0.44647	-0.01061	0.02408
Constant	0.06166***	0.01562	0.00008	0.03104	0.09228
Mixture probability (π^E)					
Age	-0.00498	0.00313	0.11148	-0.01111	0.00115
Hispanic	-0.08435	0.05209	0.10537	-0.18644	0.01774
Male	0.02903	0.03526	0.41038	-0.04008	0.09814
Years of education	0.03239***	0.00647	0.00000	0.01972	0.04507
Smoker: past	0.00888	0.04498	0.84350	-0.07928	0.09704
Smoker: present	0.03490	0.04750	0.46247	-0.05819	0.12799
LL size: medium	-0.09518	0.08167	0.24384	-0.25526	0.06489
LL size: large	0.15659**	0.06940	0.02405	0.02057	0.29261
USC sample	-0.08818*	0.05207	0.09038	-0.19024	0.01388
Constant	0.26554	0.16868	0.11544	-0.06507	0.59615
Error (μ)					
Decision time (ln)	-0.42812	0.27743	0.12279	-0.97187	0.11563
USC sample	-2.19142***	0.57184	0.00013	-3.31221	-1.07064
Constant	13.00284***	2.22874	0.00000	8.63459	17.37110
N	32508				
log-likelihood	-13652.552				

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In sum, the results from the mixture model analyses suggest that it is mistaken to assume that only one model of discounting behaviour accurately characterises all discounting choices by all subjects on the KPB task. In other words, by forcing the data to conform to one DGP one may draw incorrect inferences about the type and extent of discounting behaviour. Finally, the results suggest that, under the

assumption of an additively-separable intertemporal utility function, current smokers are no more prone to time inconsistency than ex-smokers and never-smokers.

VII. DISCUSSION AND CONCLUSIONS

This chapter provided a detailed summary of the studies which have been conducted to explore the relationship between smoking and discounting behaviour. This review highlighted a number of patterns in the literature: 1) most of the studies have been conducted in the US; 2) approximately half of the studies used hypothetical rewards; 3) there was a roughly equal split between choice and titration elicitation mechanisms; 4) every study except HLR used a two-step approach to data analysis; 5) a majority of studies assumed hyperbolic discounting; 6) the subject pools, although somewhat small, were fairly representative of the general population; and 7) the vast majority of studies found a positive and significant relationship between smoking and discounting behaviour.

I sought to replicate this finding of a positive and significant relationship between smoking and discounting behaviour by using two relatively large datasets, four distinct discounting models, and a statistical framework which respects the process(es) generating the data.

The baseline estimates presented in Section VI highlighted the importance of incorporating a behavioural error specification to account for subjects' inconsistent choices, rather than assume them away, on the KPB discounting task. Thus, this error specification was included in all subsequent analyses. A comparison of the models' log-likelihoods provided preliminary evidence that the E model had more support in the data than the H model. This is a striking finding given that the bulk of studies in this literature assume hyperbolic discounting. Finally, the QH and WB model estimates showed that discount rates decline over time, thereby raising the spectre of time-inconsistent choice behaviour under the assumption of an additively-separable intertemporal utility function.

To formally adjudicate between the discounting models, Clarke (2007) non-nested model selection tests were conducted. These tests confirmed that the E model better characterised the discounting data than the H model. In addition, I found that the WB model had the least support in the data but that it was not possible to find a transitive ranking of the remaining specifications.

The relationship between smoking and discounting was explored in two ways. The first simply added a smoking status covariate to the discounting models so as to capture the "total effect" of smoking. In the E and H models I found that current smokers discounted significantly more than both ex-smokers and never-smokers. In the QH model, the long-term discount rate δ was significantly higher among current smokers than both ex-smokers and never-smokers but there were no significant differences between these groups in terms of present-bias (i.e., the estimate of β). Finally, in the WB model, there were no significant differences between current smokers, ex-smokers, and never-smokers in the separate estimates of β and δ, but the estimates of β and δ were jointly significant for current smokers.

The second approach to exploring this relationship was to include a full set of covariates and task parameters in the models. In so doing, the focus switched to the "marginal effect" of smoking, over and above what was accounted for by other observable individual characteristics. In the E and H models, current smokers discounted significantly more than never-smokers. In the QH model, by contrast, there were no significant differences in the estimates of β and δ, either individually or jointly, across the smoking status groups. Finally, in the WB model, the estimate of δ was significantly higher for current smokers than both never-smokers and ex-smokers. Furthermore, current smokers were significantly more likely to perceive time as slowing down than both never-smokers and ex-smokers.

The final set of statistical analyses was the estimation of mixture models of the different discounting specifications. These models showed that multiple decision processes characterise the discounting of delayed rewards and that it is a mistake to force all discounting choices to fit one particular model. In addition, the models allowed me to explore the hypothesis that smokers are more prone to time inconsistency than non-smokers, under the assumption of an additively-separable

intertemporal utility function, by making the mixture probability a function of smoking status. No significant differences were found in the extent to which smokers, never-smokers, and ex-smokers discount hyperbolically.

This research makes a number of contributions to the literature. This is the first study of the smoking-discounting relationship which estimates four distinct models of discounting behaviour. These estimates were produced using a statistical framework which recognises, and takes into account, the uncertainty of discounting parameter estimates and the potential for subject errors in decision making (i.e., this framework makes it possible to draw accurate statistical inferences). That the smoking-discounting relationship was robust to different discounting models lends further credence to the hypothesis that smokers discount the future more heavily than non-smokers.

This was also the first study in this literature to provide a valid statistical test of whether the E or H models better characterise all of the discounting data. According to the Clarke (2007) test, the E model outperformed the H model in this dataset, and this result should give researchers pause when deciding whether to assume hyperbolic discounting.

The two samples used in the analyses were relatively large and subjects were carefully screened to ensure that they did not suffer from other mental disorders which may have biased the smoking-discounting results. While careful screening is common in this literature (e.g., Epstein et al. (2003), Dallery and Raiff (2007), Sweitzer et al. (2008)), this is the largest sample (N = 1205) where potential participants were screened for a wide range of mental disorders.

Perhaps the most important result in this chapter, which echoes the findings of HLR and Coller, Harrison and Rutström (2012), is that multiple decision processes characterise discounting data and that one should adopt statistical models which can accommodate this. This turns the argument about the superiority of E or H into one which focuses on the relative contributions of each model to the explanation of discounting behaviour. It also allows one to determine whether different subsamples are more or less likely to discount according to the E or H specifications.

This research also suffers from a number of shortcomings. That only a small subset (N = 174) of the people in this study were incentivised, and only probabilistically, means that, to a large extent, I am drawing inferences about the choices people think they would make when faced with those contingencies, or the choices they think the experimenter wants them to make, rather than the choices they made and experienced.

In addition, the KPB discounting task is a relatively blunt instrument for making precise inferences about discounting behaviour. A task employing more choices, different reward magnitudes, a larger range of implied interest rates, and front end delays to the SS reward, would lend itself more to the precise estimation of discounting parameters.

Finally, the lack of a risk preference task in these studies makes it difficult to estimate discounting parameters jointly with the curvature of the utility function. As HLR show, this can be very important for the conclusions that one draws. In Appendix B I discuss a parametric approach to this problem which involves the estimation of "profile likelihoods."

These issues notwithstanding, this chapter can serve as a constructive guide to researchers who want to explore the relationship between smoking and discounting behaviour in a statistical framework suitable for this task.

REFERENCES

ADAMS, J., AND D. NETTLE (2009): "Time Perspective, Personality and Smoking, Body Mass, and Physical Activity: An Empirical Study," *British Journal of Health Psychology*, 14, 83-105.

AINSLIE, G. (1975): "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, 82, 463-96.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583-618.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2014): "Discounting Behavior: A Reconsideration," *European Economic Review*, 71, 15-33.

AUDRAIN-MCGOVERN, J., D. RODRIGUEZ, L. H. EPSTEIN, J. CUEVAS, K. RODGERS, AND E. P. WILEYTO (2009): "Does Delay Discounting Play an Etiological Role in Smoking or Is It a Consequence of Smoking?," *Drug and Alcohol Dependence*, 103, 99-106.

BAKER, F., M. W. JOHNSON, AND W. K. BICKEL (2003): "Delay Discounting in Current and Never-before Cigarette Smokers: Similarities and Differences across Commodity, Sign, and Magnitude," *Journal of Abnormal Psychology*, 112, 382-92.

BALEVICH, E. C., N. D. WEIN, AND J. D. FLORY (2013): "Cigarette Smoking and Measures of Impulsivity in a College Sample," *Substance Abuse*, 34, 256-62.

BENHABIB, J., A. BISIN, AND A. SCHOTTER (2010): "Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs," *Games and Economic Behavior*, 69, 205-223.

BICKEL, W. K., D. P. JARMOLOWICZ, E. T. MUELLER, C. T. FRANCK, C. CARRIN, AND K. M. GATCHALIAN (2012): "Altruism in Time: Social Temporal Discounting Differentiates Smokers from Problem Drinkers," *Psychopharmacology (Berlin)*, 224, 109-20.

BICKEL, W. K., A. L. ODUM, AND G. J. MADDEN (1999): "Impulsivity and Cigarette Smoking: Delay Discounting in Current, Never, and Ex-Smokers," *Psychopharmacology (Berlin)*, 146, 447-54.

BICKEL, W. K., R. YI, B. P. KOWAL, AND K. M. GATCHALIAN (2008): "Cigarette Smokers Discount Past and Future Rewards Symmetrically and More Than Controls: Is Discounting a Measure of Impulsivity?," *Drug and Alcohol Dependence*, 96, 256-62.

BILDER, R. M., F. W. SABB, T. D. CANNON, E. D. LONDON, J. D. JENTSCH, D. S. PARKER, R. A. POLDRACK, C. EVANS, AND N. B. FREIMER (2009): "Phenomics: The Systematic Study of Phenotypes on a Genome-Wide Scale," *Neuroscience*, 164, 30-42.

BLEICHRODT, H., K. I. M. ROHDE, AND P. P. WAKKER (2009): "Non-Hyperbolic Time Inconsistency," *Games and Economic Behavior*, 66, 27-38.

BRADFORD, W. D. (2010): "The Association between Individual Time Preferences and Health Maintenance Habits," *Medical Decision Making*, 30, 99-112.

BUSINELLE, M. S., M. A. MCVAY, D. KENDZOR, AND A. COPELAND (2010): "A Comparison of Delay Discounting among Smokers, Substance Abusers, and Non-Dependent Controls," *Drug and Alcohol Dependence*, 112, 247-50.

CHABRIS, C., D. LAIBSON, C. L. MORRIS, J. P. SCHULDT, AND D. TAUBINSKY (2008): "Individual Laboratory-Measured Discount Rates Predict Field Behavior," *Journal of Risk and Uncertainty*, 37, 237-269.

CLARKE, K. A. (2007): "A Simple Distribution-Free Test for Nonnested Model Selection," *Political Analysis*, 15, 347-363.

COLLER, M., G. W. HARRISON, AND E. E. RUTSTRÖM (2012): "Latent Process Heterogeneity in Discounting Behavior," *Oxford Economic Papers*, 64, 375-391.

COLLER, M., AND M. B. WILLIAMS (1999): "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 107-127.

CONTE, A., J. D. HEY, AND P. G. MOFFAT (2011): "Mixture Models of Choice under Risk," *Journal of Econometrics*, 162, 79-88.

DALLERY, J., AND B. R. RAIFF (2007): "Delay Discounting Predicts Cigarette Smoking in a Laboratory Model of Abstinence Reinforcement," *Psychopharmacology (Berlin)*, 190, 485-96.

ELSTER, J. (1979): *Ulysses and the Sirens: Studies in Rationality and Irrationality*. Cambridge: University of Cambridge Press.

EPSTEIN, L. H., J. B. RICHARDS, F. G. SAAD, R. A. PALUCH, J. N. ROEMMICH, AND C. LERMAN (2003): "Comparison between Two Measures of Delay Discounting in Smokers," *Experimental and Clinical Psychopharmacology*, 11, 131-8.

FECHNER, G. (1966): *Elements of Psychophysics*. New York, NY: Holt, Rinehart and Winston.

FIELD, M., M. SANTARCANGELO, H. SUMNALL, A. GOUDIE, AND J. COLE (2006): "Delay Discounting and the Behavioural Economics of Cigarette Purchases in Smokers: The Effects of Nicotine Deprivation," *Psychopharmacology (Berlin)*, 186, 255-63.

GREEN, L., A. FRY, AND J. MYERSON (1994): "Discounting of Delayed Rewards: A Life-Span Comparison," *Psychological Science*, 5, 33-36.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2010): "Individual Discount Rates and Smoking: Evidence from a Field Experiment in Denmark," *Journal of Health Economics*, 29, 708-717.

HARRISON, G. W., AND E. E. RUTSTRÖM (2009): "Expected Utility Theory and Prospect Theory: One Wedding and a Decent Funeral," *Experimental Economics*, 12, 133-158.

HARVEY, C. M. (1986): "Value Functions for Infinite-Period Planning," *Management Science*, 32, 1123-1139.

HERRNSTEIN, R. J. (1961): "Relative and Absolute Strength of Response as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Behavior*, 4, 267-72.

HEY, J. D. (1995): "Experimental Investigations of Errors in Decision Making under Risk," *European Economic Review*, 39, 633-640.

HEYMAN, G. M., AND S. P. GIBB (2006): "Delay Discounting in College Cigarette Chippers," *Behavioural Pharmacology*, 17, 669-79.

JAMISON, D. T., AND J. JAMISON (2011): "Characterizing the Amount and Speed of Discounting Procedures," *Journal of Benefit-Cost Analysis*, 2, 1-53.

JOHNSON, M. W., W. K. BICKEL, AND F. BAKER (2007): "Moderate Drug Use and Delay Discounting: A Comparison of Heavy, Light, and Never Smokers," *Experimental and Clinical Psychopharmacology*, 15, 187-94.

JONES, B. A., R. D. LANDES, R. YI, AND W. K. BICKEL (2009): "Temporal Horizon: Modulation by Smoking Status and Gender," *Drug and Alcohol Dependence*, 104 Supplement 1, S87-93.

KANG, M. I., AND S. IKEDA (2014): "Time Discounting and Smoking Behavior: Evidence from a Panel Survey," *Health Economics*, 23, 1443-1464.

KIRBY, K. N., AND N. N. MARAKOVIĆ (1996): "Delay-Discounting Probabilistic Rewards: Rates Decrease as Amounts Increase," *Psychonomic Bulletin & Review*, 3, 100-104.

KIRBY, K. N., N. M. PETRY, AND W. K. BICKEL (1999): "Heroin Addicts Have Higher Discount Rates for Delayed Rewards Than Non-Drug-Using Controls," *Journal of Experimental Psychology: General*, 128, 78-87.

KOBIELLA, A., S. RIPKE, N. B. KROEMER, C. VOLLMERT, S. VOLLSTADT-KLEIN, D. E. ULSHOFER, AND M. N. SMOLKA (2014): "Acute and Chronic Nicotine Effects on Behaviour and Brain Activation During Intertemporal Decision Making," *Addiction Biology*, 19, 918-930.

LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443-477.

LOEWENSTEIN, G., AND D. PRELEC (1992): "Anomalies in Intertemporal Choice: Evidence and Interpretation," *Quarterly Journal of Economics*, 107, 573-592.

MAZUR, J. E. (1984): "Tests of an Equivalence Rule for Fixed and Variable Reinforcer Delays," *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 426-437.

MCLACHLAN, G., AND D. PEEL (2000): *Finite Mixture Models*. New York, NY: Wiley.

MELANKO, S., K. LERAAS, C. COLLINS, S. FIELDS, AND B. REYNOLDS (2009): "Characteristics of Psychopathy in Adolescent Nonsmokers and Smokers: Relations to Delay Discounting and Self Reported Impulsivity," *Experimental and Clinical Psychopharmacology*, 17, 258-65.

MITCHELL, S. H. (1999): "Measures of Impulsivity in Cigarette Smokers and Non-Smokers," *Psychopharmacology (Berlin)*, 146, 455-64.

MITCHELL, S. H., AND V. B. WILSON (2012): "Differences in Delay Discounting between Smokers and Nonsmokers Remain When Both Rewards Are Delayed," *Psychopharmacology (Berlin)*, 219, 549-62.

MYERSON, J., L. GREEN, AND M. WARUSAWITHARANA (2001): "Area under the Curve as a Measure of Discounting," *Journal of the Experimental Analysis of Behavior*, 76, 235-43.

ODUM, A. L., AND A. A. BAUMANN (2007): "Cigarette Smokers Show Steeper Discounting of Both Food and Cigarettes Than Money," *Drug and Alcohol Dependence*, 91, 293-6.

ODUM, A. L., G. J. MADDEN, AND W. K. BICKEL (2002): "Discounting of Delayed Health Gains and Losses by Current, Never- and Ex-Smokers of Cigarettes," *Nicotine & Tobacco Research*, 4, 295-303.

OHMURA, Y., T. TAKAHASHI, AND N. KITAMURA (2005): "Discounting Delayed and Probabilistic Monetary Gains and Losses by Smokers of Cigarettes," *Psychopharmacology (Berlin)*, 182, 508-15.

PHELPS, E. S., AND R. A. POLLAK (1968): "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 35, 185-199.

POLTAVSKI, D. V., AND J. N. WEATHERLY (2013): "Delay and Probability Discounting of Multiple Commodities in Smokers and Never-Smokers Using Multiple-Choice Tasks," *Behavioural Pharmacology*, 24, 659-67.

READ, D. (2001): "Is Time-Discounting Hyperbolic or Sub-Additive?," *Journal of Risk and Uncertainty*, 23, 5-32.

REYNOLDS, B. (2004): "Do High Rates of Cigarette Consumption Increase Delay Discounting? A Cross-Sectional Comparison of Adolescent Smokers and Young-Adult Smokers and Nonsmokers," *Behavioural Processes*, 67, 545-9.

REYNOLDS, B. (2006): "The Experiential Discounting Task Is Sensitive to Cigarette-Smoking Status and Correlates with a Measure of Delay Discounting," *Behavioural Pharmacology*, 17, 133-42.

REYNOLDS, B., AND S. FIELDS (2012): "Delay Discounting by Adolescents Experimenting with Cigarette Smoking," *Addiction*, 107, 417-24.

REYNOLDS, B., K. KARRAKER, K. HORN, AND J. B. RICHARDS (2003): "Delay and Probability Discounting as Related to Different Stages of Adolescent Smoking and Non-Smoking," *Behavioural Processes*, 64, 333-344.

REYNOLDS, B., M. PATAK, P. SHROFF, R. B. PENFOLD, S. MELANKO, AND A. M. DUHIG (2007): "Laboratory and Self-Report Assessments of Impulsive Behavior in Adolescent Daily Smokers and Nonsmokers," *Experimental and Clinical Psychopharmacology*, 15, 264-71.

REYNOLDS, B., J. B. RICHARDS, K. HORN, AND K. KARRAKER (2004): "Delay Discounting and Probability Discounting as Related to Cigarette Smoking Status in Adults," *Behavioural Processes*, 65, 35-42.

RICHARDS, J. B., L. ZHANG, S. H. MITCHELL, AND H. DE WIT (1999): "Delay or Probability Discounting in a Model of Impulsive Behavior: Effect of Alcohol," *Journal of the Experimental Analysis of Behavior*, 71, 121-43.

SAMUELSON, P. (1937): "A Note on Measurement of Utility," *Review of Economic Studies*, 4, 155-161.

SHEEHAN, D. V., Y. LECRUBIER, K. H. SHEEHAN, P. AMORIM, J. JANAVS, E. WEILLER, T. HERGUETA, R. BAKER, AND G. C. DUNBAR (1998): "The Mini-International Neuropsychiatric Interview (M.I.N.I.): The Development and Validation of a Structured Diagnostic Psychiatric Interview for DSM-IV and ICD-10," *Journal of Clinical Psychiatry*, 59 Supplement 20, 22-33;quiz 34-57.

SHEFFER, C. E., M. MENNEMEIER, R. D. LANDES, W. K. BICKEL, S. BRACKMAN, J. DORNHOFFER, T. KIMBRELL, AND G. BROWN (2013): "Neuromodulation of Delay Discounting, the Reflection Effect, and Cigarette Consumption," *Journal of Substance Abuse Treatment*, 45, 206-14.

STILLWELL, D. J., AND R. J. TUNNEY (2012): "Effects of Measurement Methods on the Relationship between Smoking and Delay Reward Discounting," *Addiction*, 107, 1003-12.

STROTZ, R. H. (1955-1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165-180.

SWEITZER, M. M., E. C. DONNY, L. C. DIERKER, J. D. FLORY, AND S. B. MANUCK (2008): "Delay Discounting and Smoking: Association with the Fagerstrom Test for Nicotine Dependence but Not Cigarettes Smoked Per Day," *Nicotine & Tobacco Research*, 10, 1571-5.

VAN DER POL, M., AND M. RUGGERI (2008): "Is Risk Attitude Outcome Specific within the Health Domain?," *Journal of Health Economics*, 27, 706-17.

VUONG, Q. H. (1989): "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 57, 307-333.

WING, V. C., T. G. MOSS, R. A. RABIN, AND T. P. GEORGE (2012): "Effects of Cigarette Smoking Status on Delay Discounting in Schizophrenia and Healthy Controls," *Addictive Behaviors*, 37, 67-72.

4. THE PROBABILITY DISCOUNTING MODEL

I. INTRODUCTION

The probability discounting (PD) model is a popular framework for investigating people's instantaneous or atemporal attitudes toward risk[1] in experimental settings (see Green and Myerson (2004), Madden and Bickel (2010) for reviews). It is particularly common in studies of addiction where delay discounting data is also often obtained (see, for example, Mitchell (1999), Richards, Zhang, Mitchell and de Wit (1999), Reynolds, Richards, Horn and Karraker (2004)). This model was developed in a series of three papers by Rachlin, Logue, Gibbon and Frankel (1986) (RLGF), Rachlin, Castrogiovanni and Cross (1987) (RCC), and Rachlin, Raineri and Cross (1991) (RRC). As its name suggests, PD draws its inspiration from models of temporal or delay discounting, where the delay to a reward is replaced with the odds against receiving a reward. This chapter explains the PD model in language familiar to decision theorists, statisticians and economists, it shows how the model relates to standard theories of choice under risk like expected utility theory, prospect theory, rank-dependent utility theory, and rank-dependent expected value theory, and it highlights some of the shortcomings of the PD approach.

RLGF set themselves the task of translating Kahneman and Tversky's (1979) (KT) cognitive model of choice (i.e., prospect theory[2]) into a behavioural model of choice inspired by Herrnstein's (1961) pioneering work on the "matching law"[3] which relates behavioural outputs to environmental inputs. Prospect theory was developed to account for a number of purported anomalies in choice among lotteries[4] that the canonical model of choice under risk (i.e., expected utility theory) allegedly cannot

[1] The discussion in Chapter 2 highlighted the importance of both instantaneous and intertemporal risk preferences in the analysis of choice behaviour. However, as this thesis only empirically examines instantaneous risk preferences, all subsequent references to "risk preferences" refer to the instantaneous or atemporal variety, unless otherwise noted.

[2] Prospect theory has become one of the popular theories of choice under risk in addition to expected utility theory. This chapter will discuss prospect theory extensively, not to put it on a pedestal but rather to determine whether Rachlin and colleagues were able to develop a behavioural model of choice which indeed has the same features as prospect theory.

[3] Other early contributions to the literature on the matching law are Davenport (1962), Logan (1965), Chung (1965), Chung and Herrnstein (1967), and Herrnstein (1970).

[4] The terms "lottery," "gamble," and "prospect" refer to a probability distribution over outcomes and will be used interchangeably.

explain. It proposes that two phases take place during the choice process: an initial editing phase which typically yields a simpler representation of the gambles; and a subsequent evaluation phase where the edited gambles are evaluated and the one with the highest value is chosen. KT provide a descriptive explanation of the editing phase and develop a formal model of the evaluation phase. The model of the evaluation phase will be outlined briefly, along with other popular models of choice under risk, so as to highlight the link between these cognitive models and the behavioural model of Rachlin and colleagues.

II. THEORIES OF CHOICE UNDER RISK WITH SIMPLE, REGULAR PROSPECTS

Using the notation in KT, $(x, p; y, q)$ is a prospect yielding outcome x with probability p, outcome y with probability q and 0 with probability $1 - p - q$, where $p + q \leq 1$. KT define a "regular" prospect as one where: $p + q < 1$, or $x \geq 0 \geq y$, or $x \leq 0 \leq y$. For the present discussion I will be concerned with regular prospects of the following form $(x, p; 0, q)$: prospects that pay x with probability p and 0 with probability q, where $p + q = 1$. I make use of these simple, regular prospects because the theory of probability discounting developed by Rachlin and colleagues appears to be limited only to gambles involving the outcome 0 and one positive or negative outcome x. In the language of prospect theory, I will only consider lotteries in the gain frame (i.e., where $x \geq 0$) to keep the discussion simple.

During the evaluation phase of prospect theory (PT), simple, regular prospects are evaluated according to the following function:

$$V(x, p; y, q) = \pi(p)v(x) + \pi(q)v(y), \qquad (1)$$

where $v(0) = 0$, $\pi(0) = 0$, $\pi(1) = 1$, $\pi(\cdot)$ is an increasing function of p, and $\pi(p) \in [0, 1]$. According to KT, people respond differently to gains and losses. This matters in (1) because the outcomes x and y are evaluated relative to some reference point r which determines whether the outcomes are perceived as gains or losses. KT were

agnostic about the determination of r whereas their subsequent model, Cumulative Prospect Theory (CPT) (Tversky and Kahneman (1992)), assumes the reference point coincides with zero income (i.e., $r = 0$). I adopt the latter assumption for ease of exposition because it implies that when $x, y \geq 0$, the lottery is in the gain frame, and I need not worry about the outcomes being perceived as losses.

As $v(0) = 0$ and I am concerned with simple, regular prospects $(x, p; 0, q)$, (1) reduces to:

$$V(x, p; 0, q) = \pi(p)v(x) \tag{2}$$

In other words, the value V of prospect $(x, p; 0, q)$ is determined solely by the product of some function $\pi(\cdot)$ of the probability p assigned to outcome x and some utility function $v(\cdot)$ of the outcome x. The function $\pi(\cdot)$ is referred to as a probability weighting function (PWF) in the literature because it takes probabilities as its argument and returns so-called decision weights. In other words, the function $\pi(\cdot)$ incorporates the possibility of subjective distortions of objective probabilities (i.e., the objective probability 0.1 may be subjectively perceived as greater or less than 0.1). Under the PT formulation, risk preferences are determined both by the shape of the utility function $v(\cdot)$ and the shape of the PWF $\pi(\cdot)$.

Unlike PT, rank-dependent utility (RDU) theory (Quiggin (1982)) is not defined relative to some reference point r, it does not treat gains and losses differently, and it applies rank-dependent transformations of probabilities to lotteries involving more than two prizes. However, under the assumption that $r = 0$, and with simple, regular prospects of the form I am considering (i.e., gambles in the gain frame with only two prizes), RDU is equivalent to PT. Thus, risk preferences are determined both by the shape of the utility function $v(\cdot)$ and the shape of the PWF $\pi(\cdot)$ under RDU.

By contrast, expected utility (EU) theory assumes $\pi(p) = p$ and defines the value V of prospect $(x, p; 0, q)$ as:

$$V(x, p; 0, q) = pv(x) \tag{3}$$

EU is not defined relative to a reference point, it does not treat gains and losses differently, and probabilities are interpreted objectively. Thus, risk preferences are determined solely by the shape of the utility function $v(\cdot)$ because $V(\cdot)$ is linear in p.[5]

Finally, Yaari's (1987) dual theory of choice under risk, sometimes referred to as the rank-dependent expected value (RDEV) model (see Harrison and Rutström (2008)), defines the value of a simple, regular prospect as:

$$V(x, p; 0, q) = \pi(p)x \tag{4}$$

Like EU and RDU, RDEV is not defined relative to a reference point and it does not treat gains and losses differently. As V is linear in outcome x in (4), risk preferences in this model are determined solely by the shape of the PWF $\pi(p)$.

As I will discuss later, risk preferences in the PD model are determined solely by the shape of the PWF and it is therefore equivalent to RDEV, except that the PD model is limited to simple, regular prospects and it employs a specific PWF whereas the RDEV model is more general. KT (p. 280-284) argue empirically that people tend to overweight low probabilities (i.e., $\pi(p) > p$ for small p) and underweight moderate to high probabilities (i.e., $\pi(p) < p$ for moderate to high probabilities). This overweighting and underweighting yields an "inverse S-shaped" PWF $\pi(\cdot)$ with the following properties: subcertainty, subproportionality, and subadditivity. Note that these properties do not define the model, they are simply specific assumptions within the model. In Section V I will conduct empirical analyses to determine whether this overweighting and underweighting of probabilities is a common feature of data from PD experiments.

[5] Note that $V(\cdot)$ is unique up to a positive affine transformation so $U(\cdot) = a + bV(\cdot)$, for $b > 0$, represents the same preferences as (3).

III. THE PD MODEL: A BEHAVIOURAL THEORY OF CHOICE UNDER RISK

Rachlin and colleagues (RLGF (1986), RCC (1987), RRC (1991)) developed the PD model in a series of three papers. As each paper built on its predecessor, the papers will be discussed in chronological order. This will provide an historical account of the development of the model while highlighting some of the problems that arose during each stage.

RLGF (1986): Cognition and Behavior in Studies of Choice

RLGF (p. 35) explain that a behavioural model of choice is one where an animal's behaviour in a specific environment is described by two sets of rules. The first consists of reinforcement schedules that are independent of the animal's behaviour while the second describes an animal's behaviour as a function of exposure to environmental stimuli. The following behavioural model of choice has been used to explain a large proportion of animal choice data:

$$B_1/B_2 = (A_1/A_2)^a(R_1/R_2)^r(D_2/D_1)^d, \qquad (5)$$

where B_1 and B_2 are rates of responding on two levers, buttons etc., during a particular time interval, A_1 and A_2 are the respective reinforcement amounts that are delivered over that interval, R_1 and R_2 are the respective rates of reinforcement over that interval, D_1 and D_2 are the respective delays to the rewards, and a, r and d represent an animal's sensitivity to amount, rate and delay, respectively.

This function occupies a venerable space in the experimental literature on animals' choice behaviour between different schedules of reinforcement. It is a generalised version of Herrnstein's (1961) matching law and it has been used to explain the commonly found tendency of animals to sharply discount delayed rewards (RLGF (1986, p. 36)).[6]

[6] Kagel, Battalio and Green (1995) critique the matching law and its implications for the commodity-choice behaviour (p. 51-71), labour-supply behaviour (p. 110-128), and time discounting (p. 178-180) of animals, predominantly rats and pigeons, under experimental conditions. They argue in favour of an economic account of animal choice behaviour that relies on maximising (i.e., the comparison of marginal rates of return) rather than matching (i.e., the comparison of average rates of return).

Although this function was derived to provide an account of animals' delay discounting behaviour, RLGF employ it to explain how people value probabilistic rewards. The crucial step in this reformulation is to interpret the probability associated with a reward as the delay to, or the rate of reinforcement of, this reward. As RLGF (p. 35) note, delay and rate of reinforcement are typically interdependent in experiments, an issue to which I will return later.

To understand this reformulation of probability, consider the following thought experiment presented in RLGF (p. 36). A person plays a gamble with a probability p = ⅓ of winning x and a probability of $1 - p$ = ⅔ of receiving 0. A physical randomisation device is used to determine whether the subject wins x and the gamble is repeated periodically.[7] Suppose that the physical randomisation device takes c seconds to deliver a result and the intertrial interval (ITI) between gambles is t seconds. Then, timed from the start of the first gamble, the average or expected delay (D) to the person's first win is given by the following waiting-time function:

$$D = [(t + c) / p] - t \qquad (6)$$

Note that t is subtracted on the assumption that there is no delay to the first gamble. Expression (6) shows how, at least in the repeated-gambles case, probability and delay are linked. As probability increases so delay to reward falls, and vice versa. In addition, probability affects the rate of reinforcement of a reward in the following way. Using the parameters from the description above (i.e., p = ⅓), over a long series of repeated gambles, the rate at which the person would receive money is $x / [3(t + c)]$ per second.

Now suppose that the person is in one room and the physical randomisation device is in another. If the person wins x this is delivered through a trapdoor. Suppose further that there are two randomisation devices, for two different lotteries, and two trapdoors, and that the person must choose between them. According to RLGF (p.

[7] As the gamble is repeated and the person wins x on successful realisations of the gamble, payoffs in this thought experiment are cumulative. RLGF implicitly assume away the effect that cumulative payoffs have on choice behaviour. In other words, they assume that the person does not have an intertemporal utility function which is affected by changes in income.

36), the choice can be viewed as one between gambles or as one between rates of reinforcement. The behavioural model of choice developed by RLGF assumes that the person is choosing between different rates of reinforcer delivery and in so doing reinterprets probabilities as delays to, or rates of reinforcement of, rewards.

In experiments where people choose between lotteries, the two effects of probability (i.e., on the rate of reinforcement of a reward and the delay to a reward) are confounded. Following RLGF (p. 39), to see how this affects the generalised matching law equation (5), let $R_1 = 1/D_1$ and $R_2 = 1/D_2$ so that rate and delay are perfectly confounded. This yields the following equation:

$$B_1/B_2 = (A_1/A_2)^a(D_2/D_1)^{r+d} \tag{7}$$

Using (6) and assuming that $c_1 = c_2 = c$ and $t_1 = t_2 = t$, (7) becomes:

$$B_1/B_2 = (A_1/A_2)^a([p_1(t(1 - p_2) + c)] / [p_2(t(1 - p_1) + c)])^{r+d} \tag{8}$$

Expression (8) represents the matching law for probabilistic rewards when the rate of reinforcement of a reward is perfectly confounded with the delay to that reward.

RLGF derive this expression to provide a behavioural account of choice among gambles: preference for lottery 1 over lottery 2, as measured by rate of responding (B_1/B_2) over a time interval, is a function of the rewards in these lotteries (A_1 and A_2), the probabilities associated with these rewards (p_1 and p_2), the reward delivery period c, the ITI length t, and the sensitivity exponents a, r, and d. RLGF use (8) and set $p_2 = 1$, $t = 1$, $c = 0.2$, and $A_1 = A_2$, in an attempt to replicate KT's assertion that people overweight low probabilities and underweight moderate to high probabilities. Under these assumptions, (8) becomes:

$$B_1/B_2 = [(0.2p_1) / (1.2 - p_1)]^{r+d} \tag{9}$$

Plotting B_1/B_2 as a function of p_1/p_2 (where $p_2 = 1$) for different values of $r + d$, yields Figure I. In the figure, $s_R = r$ and $s_D = d$ according to the notation that I have used. RLGF (p. 39) argue that when $r + d = 0.5$, the plotted function is their model's

counterpart to the inverse S-shaped PWF of KT which is based on the overweighting of low probabilities and the underweighting of moderate to high probabilities.[8] Note that a PWF relates stated probabilities (i.e., p_1 and p_2) to subjective decision weights (i.e., $\pi(p_1)$ and $\pi(p_2)$). Figure I, by contrast, plots relative rate of responding (B_1/B_2) as a function of relative probabilities (p_1/p_2). While the shape of the curve in Figure I when $r + d = 0.5$ resembles an inverse S-shaped PWF, it does not relate probabilities to decision weights and hence it is inaccurate for RLGF to claim that they have translated KT's model into a behavioural model of choice.

Figure I

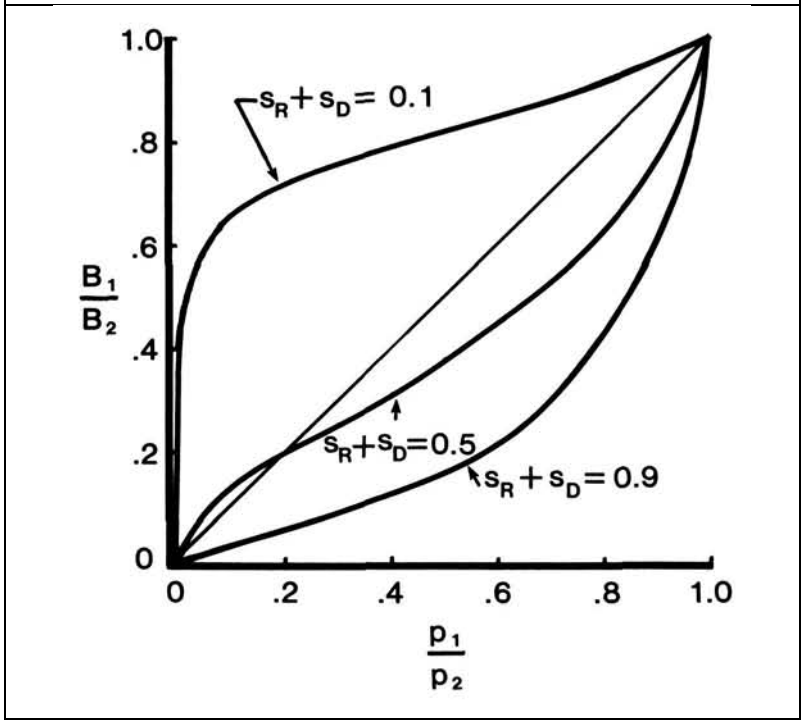

Source: RLGF (p. 40), Figure 4.

Furthermore, Figure I was plotted on the assumptions that $A_1 = A_2$ and $p_2 = 1$. In other words, the rewards under both lotteries are the same but reward A_2 is received with certainty under lottery 2. Over almost its entire range, Figure I has $p_1 < 1$ but the figure suggests that lottery 1 will still be chosen some of the time (i.e., $B_1/B_2 \neq 0$). This means that a lottery which pays, say, \$100 with a probability less than one and

[8] As stated previously under expression (1), $\pi(0) = 0$ and $\pi(1) = 1$, which, when coupled with the properties of subadditivity, subcertainty, and subproportionality, means the KT PWF has jump discontinuities at $p = 0$ and $p = 1$. KT (p. 282-283) suggest that these discontinuities may capture the distinction between certainty and uncertainty and conclude that their function is not well-behaved near the end points of the probability interval [0, 1]. Figure I does not incorporate these jump discontinuities.

$0 otherwise will be chosen over a degenerate lottery which pays $100 for sure. While KT (p. 284) recognise that their model can lead to violations of stochastic dominance, and subsequently developed CPT to account for this, they argued that such violations were unlikely because dominated prospects would be removed during the initial editing phase. Clearly this is not the case in the model of RLGF because Figure I implies that the stochastically dominated lottery will be chosen at least some of the time.

A testable implication of (8) is whether a longer ITI *t* will affect choice among probabilistic rewards. RLGF conducted an experiment to test this hypothesis, one which was replicated by Silberberg, Murray, Christensen and Asano (1988) (SMCA). Both studies used hypothetical rewards so their results should be treated with caution. In the experiment, subjects chose between two spinners, each of which rotated over a circle made up of 18 pie-shaped wedges (see Figure II). The wedges were black on one side and white on the other. If the spinner landed on a white wedge, the subject was told that she had won, if it landed on a black wedge, the subject was told that she had lost. One spinner, referred to as the "sure thing," had 17 white wedges and 1 black wedge whereas the other spinner, referred to as the "risky gamble," had 7 white wedges and 11 black wedges. The hypothetical payoff for the sure thing was $100 and the hypothetical payoff for the risky gamble was $250.

Figure II

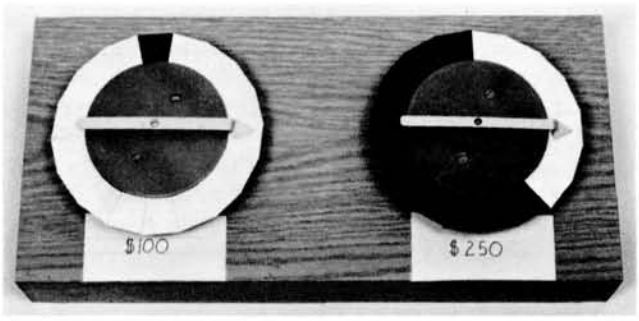

Source: RLGF (p. 40), Figure 4.

The number of white wedges on the sure thing spinner remained constant throughout the experiment. However, the number of white wedges on the risky gamble spinner was adjusted depending on the subject's previous choice: a choice of the sure thing

resulted in one more white wedge on the risky gamble (i.e., an increase in the odds of winning on the risky gamble) whereas a choice of the risky gamble led to one less white wedge on the risky gamble (i.e., a decrease in the odds of winning on that spinner). This titration procedure[9] was used so that, at equilibrium, the subject would be indifferent between the sure thing and the risky gamble (i.e., B_1 and B_2 in (8) would be equal at equilibrium). Subjects made 10 choices in total in the experiment.[10]

To see the effect of a longer ITI on choice between the two gambles, note that $A_1 = 100$, $A_2 = 250$, p_1 is approximately equal to 1[11], $c_1 = c_2 = 5s$, $t_1 = t_2 = t$ and, for simplicity, $a = r + d = 1$. Substituting these values into (8), remembering that at equilibrium $B_1 = B_2$, and solving for p_2:

$$p_2 = (t + 5) / (t + 12.5) \qquad (10)$$

When the ITI $t = 0$, $p_2 = 0.4$. In other words, a person choosing between the spinners would be indifferent between the sure thing and the risky gamble when the probability of receiving \$250 under the risky gamble is 0.4. However, as t increases so $p_2 \rightarrow 1$. Thus, with a long ITI, a subject is only indifferent between the sure thing and the risky gamble when the likelihood of receiving \$250 under the risky gamble is approximately 1. This is precisely the implication that RLGF and SMCA set out to test.[12]

---

[9] As discussed in Chapter 3, titration procedures are susceptible to being "gamed" by subjects and do not, therefore, promote truthful revelation of preferences (i.e., they lack incentive compatibility). For example, subjects may disproportionately choose the sure thing spinner on the first few trials so as to increase the odds of winning on the risky gamble in the final trials of the experiment. This point is moot for studies involving hypothetical rewards as these lack incentive compatibility to begin with, but should be taken into account for studies with titration procedures and real rewards.

[10] Whether 10 choices are enough to reach equilibrium is an open question and one which is not taken up by RLGF.

[11] KT (p. 265) argue that people respond differently to certain outcomes as opposed to near-certain outcomes. RLGF ignore this point by assuming that when $p = 17/18$ this is the same as $p = 1$. In other words, they treat an inherently risky prospect as one involving no risk.

[12] This implication is at odds with the literature on the matching law applied to choice among delayed rewards. This literature, cited previously, suggests that as the delay to all rewards increases, the likelihood of selecting the larger, more delayed reward - rather than the smaller, more immediate reward - increases. If people understand probability as delay then a low probability is equivalent to a long delay. So as the ITI increases (i.e., as the delay to both rewards increases), subjects should be more likely to select the larger, more uncertain reward (viz., the larger, more delayed reward). This works against the hypothesis which RLGF sought to test.

RLGF found evidence in support of this hypothesis in a between-subject experimental design involving 30 subjects, where one group's ($n_1 = 15$) ITI was 30*s* and the other group's ($n_2 = 15$) ITI was 90*s*.[13] Specifically, the group with the longer ITI selected the sure thing spinner more often, which means that the odds of winning on the risky gamble was higher, than the group with the shorter ITI. This comparison used the number of risky gamble choices of each subject over the course of the experiment as the data on which to conduct a *t*-test. Subjects in the 30*s* ITI group selected the risky gamble an average of 5.87 times, whereas the subjects in the 90*s* ITI group selected the risky gamble an average of 3.93 times, over the course of the experiment; this difference was significant, $t = 4.65$, df = 28, $p < 0.01$.

SMCA replicated RLGF's experiment, albeit with three procedural differences: SMCA used a computer, rather than an experimenter, to present and record subjects' choices; SMCA added an additional experimental treatment: some of the subjects began the experiment with a choice trial while others[14] started the experiment with an ITI; and one group had an ITI of 25*s* as compared to the RLGF study where the comparable group had an ITI of 30*s*.

SMCA recruited 101 subjects to take part in the experiment. They compared the number of white wedges on the risky gamble spinner in the final round of the experiment across the groups (choice trial vs. ITI × 25*s* vs. 90*s*), using a Kruskal-Wallis test, and found a significant between-group difference: $H(3) = 15.8$, $p < 0.05$. Using pairwise post hoc contrasts, they found no statistically significant differences[15] between subjects who began the experiment with a choice trial and subjects who began the experiment with an ITI. However, subjects in the short ITI group, regardless of whether they started the experiment with a choice trial or ITI, selected the risky gamble spinner significantly more often than subjects in the long ITI group. Thus, SMCA replicated the result of RLGF.

In a follow-up experiment, SMCA decided to test the robustness of this result by telling subjects how many choices (i.e., 10) they would make in the experiment; this

[13] With different ITIs, temporal discounting behaviour could drive the results in the experiment. RLGF implicitly assume that their subjects did not discount delayed rewards.
[14] Unfortunately SMCA do not provide the exact number of subjects in each experimental treatment.
[15] SMCA do not report the test statistics nor *p*-values for the pairwise post hoc contrasts.

information was not provided in SMCA's original experiment nor the experiment of RLGF. Using a sample of 20 students, SMCA found that there was no significant difference in the number of risky gamble choices between the group with the short ITI and the group with the long ITI, although they did not provide test statistics nor *p*-values for their comparisons.[16]

Finally, SMCA conducted another experiment with 40 students where, in addition to the different ITIs, they told one group of subjects that they had been endowed with $10 of hypothetical money and the other group was told that they had been endowed with $10,000 of hypothetical money. This information was given to subjects at the start of the experiment and they were also told that the experiment consisted of 10 choice trials. They found differences between the group endowed with $10 and the group endowed with $10,000 but no differences between the short ITI and the long ITI groups. In sum, SMCA replicated the result in RLGF but found that the ITI effect disappeared as soon as subjects were told how many choices they would have to make in the experiment. Thus, the results reported by RLGF appear to be very sensitive to the information provided to subjects.

Ignoring the sensitivity of these results for the moment, what does the RLGF model imply when these different ITIs are used? Using (10), when $t = 30s$, $p_2 \approx 0.82 \approx 15/18$ white wedges on the risky gamble. When $t = 90s$, $p_2 \approx 0.84 \approx 15/18$ white wedges on the risky gamble. So, for these parameters, the RLGF model implies very little difference in the choice behaviour of the two groups.[17] RLGF (p. 41) found that, at

[16] SMCA dropped the treatment where one group of subjects started the experiment with a choice trial while the other group started with an ITI.

[17] Harrison (1989, 1992, 1994) provides a detailed critique of inferences drawn from experimental studies when the studies fail to satisfy Smith's (1982) precepts for controlled microeconomic experiments. Harrison (1989, 1992, 1994) shows that deviations from "optimal" behaviour in many experiments often entail such a low cost in terms of forgone expected income that experimental subjects have almost no incentive to find the theoretical optimum nor reveal their preferences with any real precision. In other words, some experimental tasks do not adequately incentivise subjects to reveal their "true" preferences because the expected increase in earnings from doing so does not offset the required cognitive effort. Thus, many experiments fail to satisfy Smith's (1982) precept of "dominance," which requires that the experimental reward medium dominates the subjective costs of decision making in the experiment; Harrison refers to this as the "payoff-dominance critique" whereas Bardsley et al. (2010) refer to this as the "flat maximum critique." Morgan and Tustin (1992, p. 1142-1143) make a similar point in the context of labour supply choices by pigeons: to locate the optimum in some decision problems requires the comparison of margins which are so small that only a fully informed, optimising agent could conceivably evaluate. In these cases, the decision problems do not provide a valid test of the theory. The experiment of RLGF is a case in point.

the end of the experiment, the group with the longer ITI had an average of 9 white wedges on the risky gamble whereas the group with the shorter ITI had an average of 6 white wedges on the risky gamble. Thus, the model of RLGF not only predicts very little difference between the groups it also vastly overestimates the number of white wedges on the risky gamble spinner.

A fundamental problem with this model is the assumption that people interpret probabilities as delays to reward or as rates of reinforcement of reward. The thought experiment described earlier makes a case for this view, but it assumes that individuals are exposed to repeated gambles. How then do people behave when faced with a one-shot gamble? This is when the logic of the model becomes tenuous. RLGF (p. 38) argue that, "… the behavioral model must infer the existence of past external events (events that had been paired with the stated probability)." In other words, RLGF are forced to assume that people still interpret the probability of winning a reward as a delay to that reward which is linked to some past, external event, even in a one-shot setting. They must maintain this assumption even though, in a one-shot gamble with a prize of 0 and a prize of $x > 0$ (i.e., a simple, regular prospect), the probability of receiving x represents the likelihood that one receives x and the complementary probability represents the likelihood of receiving nothing. This is a problem because if one does not receive x then it will not be received, regardless of the length of time that one waits.

Another issue with the thought experiment described earlier is that people only observe the outcome of a gamble when it is favourable. In other words, there is an implicit assumption that a subject chooses one gamble and then simply waits for it to pay out rather than choosing a gamble and observing when it does and does not pay out. These procedural differences likely have a large impact on how the experiment is perceived by the subject. Choosing a gamble and then waiting until it pays out without observing each trial, frames the task as one involving amounts and delays. Watching a gamble pay out or not across repeated trials, frames the experiment as one involving risk and allows the subject to estimate the probability of receiving x, and the complementary probability of receiving 0. In the former case, interpreting probability as delay may be valid because that is how the subject experiences the task. In the latter case, the link between probability and delay is tenuous to the point of vanishing.

Finally, note what the RLGF model implies for the relationship between risk aversion and the length of the ITI. A person is risk averse if she prefers the certain outcome x to a gamble with an expected value of x. From (10) it is clear that as the ITI increases, p_2 increases, which means that as the ITI increases, risk aversion increases too. Pushing this logic to its natural conclusion, a long enough ITI would generate an aversion to risk that would make a small, certain reward preferable to a far larger, near-certain reward. While this may apply to some agents in some circumstances, empirical evidence supporting the general validity of this prediction is not provided by RLGF.

RCC (1987): Probability and Delay in Commitment

RCC conducted an experiment with real rewards to bolster RLGF's claim that probabilistic choice can be tied to a temporal framework. The experiment was based on Rachlin and Green (1972) who studied the behaviour of pigeons in a delay-commitment paradigm. The major difference with this experiment, other than that RCC used human subjects, was that delays were replaced with probabilities: long delays with low probabilities and short delays with high probabilities.

To understand the RCC experiment, it helps to discuss the original experiment of Rachlin and Green (1972). Figure IIIa shows the experimental design of Rachlin and Green (1972) whereas Figure IIIb shows the experimental design of the RCC study. Rachlin and Green (1972) gave pigeons a choice, at point B, between the path leading to point A or the path leading to point C; pigeons pecked at illuminated keys to make this choice. If pigeons chose the path leading to point A, they made a subsequent choice between a smaller, sooner (SS) reward ($a = 2s$, $d = 0s$) and a larger, later (LL) reward ($a = 4s$, $d = 4s$).[18] Thus, choice of the path leading to point A gave pigeons *flexibility* in their choice in the second stage. If pigeons chose the path leading to point C, there was no subsequent choice and the pigeons automatically received the LL reward ($a = 4s$, $d = 4s$). Thus, choice of the path leading to point C *committed* pigeons to the LL reward.

[18] Note that the reward a is measured in seconds because this is the amount of time that pigeons were given access to food.

Rachlin and Green (1972) manipulated the delay t to points A and C from a choice at B to determine whether this delay affected choice behaviour at point B. They found that when t was short ($t < 4s$), pigeons at point B predominantly chose the path leading to point A and then subsequently chose the SS reward ($a = 2s$, $d = 0s$) over the LL reward ($a = 4s$, $d = 4s$) on almost every trial (i.e., more than 90% of the time). By contrast, when t was relatively long ($t > 4s$), pigeons at point B predominantly chose the path leading to point C where only the LL reward was available ($a = 4s$, $d = 4s$). This is a delay-commitment paradigm because it tests whether commitment (i.e., choice of the path leading to point C over the path leading to point A) increases as the delay to all rewards increases. Rachlin and Green (1972) found that longer delays lead to a preference for commitment.

Figure III

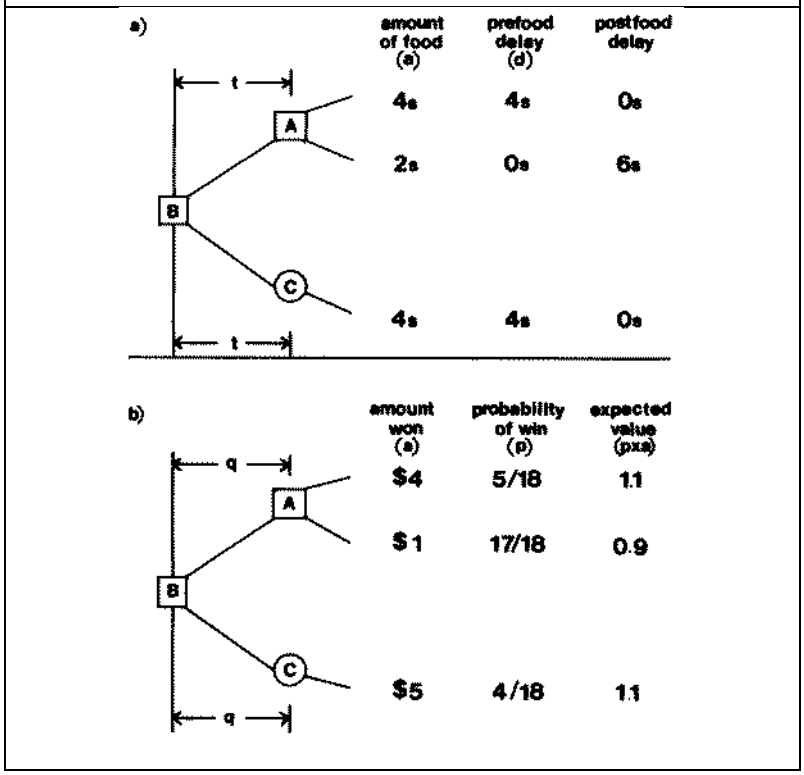

Source: RCC (p. 348), Figure 1.

As mentioned previously, RCC replaced delays with probabilities, and pigeons with humans. Their experiment played out across two stages. In the first stage, subjects had to choose whether to allocate a blue chip (probability $q = 3/18$ of advancing to stage 2) or red chip (probability $q = 15/18$ of advancing to stage 2) to one of two cards: X

or Y. Card X corresponds to point A and card Y corresponds to point C in Figure IIIb. After allocating a red or blue chip to card X or Y, a spinner, programmed with the relevant probability (i.e., $q = 3/18$ for the blue chip or $q = 15/18$ for the red chip), was used to determine whether the subject proceeded to stage 2.

If a subject allocated her chip to card X and was successful, she moved on to stage 2 (point A) where she had to choose whether to play a low reward, high probability gamble (i.e., \$1 reward with probability 17/18) or a high reward, low probability gamble (i.e., \$4 reward with probability 5/18). Thus, if the subject allocated a chip to card X and was successful, she had *flexibility* in her choice at stage 2. If, by contrast, she placed her chip on card Y during stage 1 and was successful, she then played a high reward, low probability gamble (i.e., \$5 reward with probability 4/18) in stage 2 (point C). Thus, by allocating a chip to card Y, the subject was *committed* to the high reward, low probability gamble in stage 2, if it was reached. As mentioned previously, the rewards used in this study were real, rather than hypothetical, and subjects were paid their winnings at the end of the experiment.

At the start of the experiment each subject was given 10 blue chips and 10 red chips which they could allocate, in any order, to card X or Y across 20 trials. This differs from the design in Rachlin and Green (1972) because pigeons could not choose the delays themselves; the pigeons were exposed to the different delay treatments. RCC provide a strange justification for this difference:

> "This method of having the subjects themselves select trial order was chosen because pilot experiments of ours as well as published accounts of human laboratory analogs to animal experiments … indicate that corresponding results are more likely when people's tasks are made more complicated and varied than corresponding animal tasks. For similar reasons, the two larger rewards (\$5 and \$4) were not identical (although their expected values were identical)." (RCC (p. 349)).

RCC argue that their experiment is the probabilistic choice analogue of the experiment in Rachlin and Green (1972). Rather than manipulate the delay t to points A and C from B, RCC manipulated the probability q of reaching points A and C from

B. Choice of the path leading to point A (i.e., allocating a token to card X) gives a subject flexibility in her choice in the second stage of the experiment, *if it is reached*. By contrast, choice of the path leading to point C (i.e., allocating a token to card Y) commits the subject to a high reward, low probability gamble in the second stage of the experiment, *if it is reached*. By varying q, RCC could test whether a low continuation probability, which RCC argue is analogous to a long delay, is associated with more commitment choices than a high continuation probability.

RCC found that at point A (i.e., after allocating a token to card X and successfully proceeding to stage 2), 28% of choices were for the high reward, low probability gamble, as represented by the bar on the far left of Figure IV. Note that this fraction of choices was significantly less ($t = 3.28$, df = 10, one-tailed test) than 50%. In other words, if point A was reached, there was a preference for the low reward, high probability gamble. Note that this result does not line up perfectly with Rachlin and Green (1972), who found an almost exclusive preference for the smaller, sooner reward at point A.

Figure IV

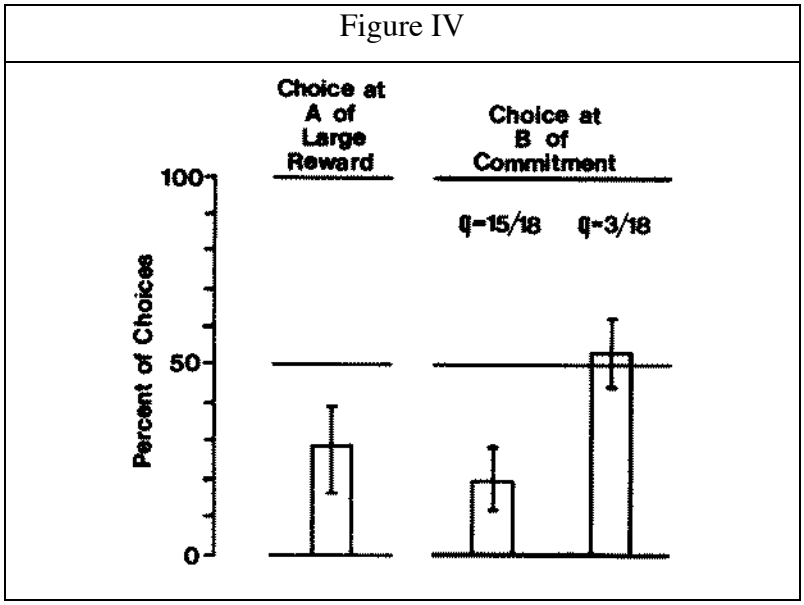

Source: RCC (p. 351), Figure 3.

In stage 1, approximately 18% of the high continuation probability chips (i.e., red chips where $q = 15/18$) and 53% of the low continuation probability chips (i.e., blue chips where $q = 3/18$) were placed on card Y, as represented by the bars on the right

of Figure IV.[19] Note that the allocation of red chips to card Y was significantly less than the allocation of blue chips to card Y (t = 5.42, df = 10, two-tailed test). RCC interpret the preceding set of results as evidence that a preference for the low reward, high probability gamble changed to indifference when the continuation probability q fell from 15/18 to 3/18.

The conclusion which RCC reached relies on a strange and dubious comparison: the allocation, approximately 53%, of low continuation probability blue chips to card Y during stage 1, as shown by the bar on the far right of Figure IV, and the choice of the $4 low probability gamble during stage 2 of card X (approximately 28%), as shown by the bar on the far left of Figure IV. In other words, the comparison is between the choice of gambles after the resolution of stage 1 uncertainty and the initial stage 1 choice between cards prior to the resolution of uncertainty.

A more appropriate comparison would be with the allocation of red and blue chips to cards X and Y in stage 1. Subjects allocated significantly more red chips than blue chips to card X (t = 5.42, df = 10, two-tailed test). This means that subjects had a preference for flexibility in stage 2, rather than commitment in stage 2, when using the high continuation probability red chips.[20] This does not imply the converse though: that subjects had a preference for commitment over flexibility when using low continuation probability blue chips. Subjects were practically and statistically indifferent between flexibility and commitment when allocating blue chips. RCC (p. 350) simply state that the 53% allocation of blue chips to card Y is not significantly different to 50% without providing test statistics, although this can be seen to be true by looking at the whiskers of the box on the far right of Figure IV.

Thus, RCC replicated the Rachlin and Green (1972) result of a preference for flexibility at short delays (viz., high probabilities), but failed to replicate the result of a *preference* for commitment at long delays (viz., low probabilities). While the fraction of commitment choices was greater with low probabilities than with high

[19] RCC do not provide standard deviations for these estimates but these are represented by the whiskers in Figure IV.

[20] Recall that allocating a chip to card X gives you freedom of choice (viz., flexibility) in stage 2, if it is reached. By contrast, allocating a chip to card Y commits you to playing the high reward, low probability gamble in stage 2, if it is reached.

probabilities, this fraction was not significantly different to 0.5. Thus, it is not valid for RCC (p. 350) to claim that, "these results parallel those obtained with pigeons choosing among rewards of various amounts and delays."

A major experimental design issue of the RCC study was the sequential allocation of tokens: subjects had to allocate a token, observe the result of the ensuing gamble, and then allocate another token, until all of their red and blue tokens were finished. Consequently, each person's idiosyncratic payoff history may have influenced her subsequent choices. In other words, this design is not immune to order or wealth effects across trials.[21] This is a point which RCC (p. 350) acknowledge, but their approach to the problem is not satisfactory.

RCC focussed on the last four trials of the experiment to see whether choices during these trials were markedly different to the choices made in previous trials. At the level of the sample as a whole, there was an equal number of red and blue tokens left for allocation over the last four trials. Thus, the last four trials, at least in terms of the proportion of red and blue tokens in the sample, were comparable to the trials at the start of the experiment. RCC found that the allocation of tokens to cards X and Y was very similar in the last four trials as in the experiment as a whole.[22] In addition, they found that choices in stage 2 of card X (i.e., between the high reward, low probability gamble and the low reward, high probability gamble) were very similar in the last four trials as in the experiment as a whole.

While these results suggest that order and wealth effects were unlikely to be driving RCC's findings, one should heed the warning of Harrison (2007) that appropriate statistical techniques need to be used to draw inferences from experimental data when there is the potential for correlation of responses at the level of the individual and over time. RCC ignore these possibilities by treating the *K* choices of each subject as independent and by not taking into account the time path of choices in the experiment.

[21] A cleaner experimental design would be to ask subjects to allocate their red and blue tokens across the cards and across the gambles at stage 2 at the outset of the experiment, as if they were constructing a portfolio of risky assets, and then play out all of the gambles.

[22] RCC compared the last four trials to the full twenty trials. Ideally they should have compared the last four trials to the first sixteen trials of the experiment.

In addition to the issues outlined above, only 11 subjects took part in the experiment so there was minimal power for the statistical tests that were conducted. Furthermore, RCC specifically asked their subjects not to make any mathematical calculations even though they were presented with options that had real financial consequences. Thus, the experiment of RCC does little to support their contention that probability is best interpreted as delay.

RRC (1991): Subjective Probability and Delay

The preceding two papers laid the foundations for the PD model, but it was the method in RRC that has been replicated numerous times (see, for example, Ostaszewski (1997), Mitchell (1999), Richards, Zhang, Mitchell and de Wit (1999), Reynolds, Richards, Horn and Karraker (2004)) and has defined the model as it is currently employed. RRC drew on the work of Mazur (1984), who ran a series of delay discounting experiments with pigeons, to further develop the purported link between probability and delay. Mazur ((1984, p. 427)) found that the pigeons' delay discounting data was best explained by a hyperbolic discounting function:

$$V = x / (1 + \delta D), \tag{11}$$

where V is the present or discounted value of the delayed reward x, D is the delay to the reward, and δ is a parameter which captures the extent to which future values are discounted; as δ increases so the present value of a delayed reward declines.[23]

RRC argue, once again, that probability is best interpreted as delay and then use the waiting-time function (6) to derive the following result. On the assumption that c, which is the time it takes for a physical randomisation device to deliver a result, is small relative to the ITI t, the waiting-time function (6) can be re-written as:

$$D = (t / p) - t = t[(1 / p) - 1] = t\Theta, \tag{12}$$

where $\Theta = ((1 / p) - 1)$ represents the odds against receiving a reward.

[23] Using the notation from Chapter 3, the hyperbolic (H) discounting function has a discount factor $D^{H}(t) = 1 / (1 + \delta t)$. RLGF distinguish between the delay to a reward D and the ITI t so, using their notation, the discount factor (DF) for the hyperbolic function is: $DF^{H}(D) = 1 / (1 + \delta D)$.

Thus, if the probability of receiving reward x under some gamble is 0.2, then the odds against receiving x is 4:1. As RRC (p. 235) note, in the context of repeated gambles, odds against is the average number of losses expected before a win.

The derivation in (12) relies on the assumptions that $c = 0$ and that there is no ITI prior to the outcome of the first gamble in a set of repeated gambles. This latter assumption is why t is subtracted in the waiting-time function (6) and why $\Theta = ((1 / p) - 1)$ represents the odds against receiving a reward; if t was not subtracted then $\Theta = (1 / p)$.

RRC used (11) and (12) to develop the following function, which they argue is analogous to the delay discounting function (11), to describe how people value or "discount" lotteries of the form $(x, p; 0, q)$:

$$V = x / (1 + \gamma\Theta), \tag{13}$$

where γ performs the same role as δ in (11): it captures the extent to which the probabilistic reward x is "discounted" as a function of the odds against receiving it. Expression (13) defines the PD model as it is currently employed.

This derivation appears quite sensible and (13) certainly looks just like (11), with γ taking the place of δ and Θ taking the place of D. However, if probability and delay are inextricably linked[24] then presumably the correct substitution for D in (11) is $t\Theta$ which is the result that was derived in (12). After all, Θ only represents odds against when we subtract t in the waiting-time function (6). Thus, to derive a probability discounting function which takes the same form as the delay discounting function

[24] Researchers have investigated the link between probability and delay by conducting discounting experiments where they vary the probability that subjects receive payment for one of their choices on the task. Keren and Roelofsma (1995) conducted an experiment, using hypothetical rewards, where subjects were offered a choice between 100 Dutch Guilders now or 110 Guilders in 4 weeks. When the probability of payment was 1, 82% of 60 subjects chose the SS reward. When the payment probability was reduced to 0.9, only 54% of 70 subjects chose the SS reward, and when the payment probability declined to 0.5, only 39% of 100 subjects chose the SS reward. These results suggest that probability and delay are linked, and that they have similar effects on behaviour. However, follow-up studies by Weber and Chapman (2005), using hypothetical rewards, and Andersen, Harrison, Lau and Rutström (2014), using real rewards, failed to replicate this result.

(11), RRC simply abandoned the link between probability and delay which they argued was so central to the way people interpret probability. This point should not be ignored because it was the stated rationale for performing the experiment with different ITIs in RLGF.

RRC set out to test the PD model (13) and the delay discounting model (11) by recruiting 80 undergraduate students to take part in experiments with hypothetical rewards: 40 subjects were used to obtain a PD function and 40 were used to obtain a delay discounting function. In the PD experiment, subjects made binary choices between $1,000 available with different probabilities (0.95, 0.9, 0.7, 0.5, 0.3, 0.1 and 0.05) and an amount of money to be received with certainty ($1,000, $990, $980, $960, $940, $920, $900, $850, $800, $750, $700, $650, $600, $550, $500, $450, $400, $350, $300, $250, $200, $150, $100, $80, $60, $40, $20, $10, $5 or $1). For each probability, half of the subjects in the PD experiment were presented with these certain amounts of money in decreasing and then increasing order, while the other subjects were presented with these amounts in increasing and then decreasing order. This elicitation procedure was used to find the certain amount of money that was subjectively equivalent (i.e., that made the subject indifferent) to $1,000 with the various probabilities listed above.

Note that points of indifference (viz., certainty equivalents) were obtained by averaging the amounts of money before and after a switch was made.[25] For example, if someone chose $1,000 with a probability of 0.9 over $850 with certainty, and then chose $900 with certainty over a 0.9 probability of receiving $1,000, then the indifference point was assumed to be $875 for a probability of 0.9. The procedure in the delay discounting experiment was identical except that probabilities were replaced with delays: 1 month, 6 months, 1 year, 5 years, 10 years, 25 years and 50 years.

[25] Taking the average of these amounts is arbitrary and doing so throws away information on the nature and uncertainty of the estimate. The correct approach for analysing interval data (e.g., where a person's indifference point lies in the interval ($850, $900)) is interval regression (see Harrison and Rutström (2008, p. 62-69) for a discussion of this point) but this is not the approach which RRC adopt. While interval regression is the appropriate statistical model for interval data it does not allow one to compare different PWFs which, as I discuss later, is an express purpose of this chapter. Consequently, I adopt a complementary full information maximum likelihood framework, which directly incorporates the uncertainty of the estimates and models it statistically, so that I can draw robust inferences about the ability of different PWFs to characterise PD data.

RRC fitted the hyperbolic delay (11) and probability (13) functions to the median[26] indifference points and certainty equivalents in the sample, respectively. They argued that these hyperbolic functions provided a better fit to the data than exponential functions, but they did not conduct formal statistical tests of this assertion. Nevertheless, they concluded that, "the corresponding form of Equations [11] and [13] implies that odds against in probabilistic discounting acts like delay in delay discounting and tends to confirm the speculation of Rachlin et al. (1986) that stated probability and stated delay have corresponding effects on behaviour" (RRC (p. 239-240)). In the remainder of this section, I will evaluate this claim to determine whether it is justified.

As discussed earlier, the "probability as delay" interpretation relies on a repeated gambles context. As RRC (p. 243) admit, the probabilities in their experiment were "one-shot" as opposed to repeated gambles. Furthermore, subjects did not experience any of the gambles, and hence did not experience any delay between choice and reward, because the experiment was entirely hypothetical. Thus, the original reinforcement rationale underlying the thought experiment in RLGF does not apply in this study and, thus, the probability as delay interpretation is tenuous in this context.

A more fundamental problem is that the link between probability and delay in (6) and (12) was severed by removing the parameter t from expression (13). RRC provide no justification for this omission and I have been unable to find any explanation in studies which employ the PD model.

Methodologically and statistically, the elicitation method suffers from notable drawbacks. To illustrate this point, suppose that someone is indifferent between \$850 with certainty and a lottery that pays \$1,000 with a probability of 0.9 and \$0 with a probability of 0.1. This implies that certain amounts greater than \$850 will be preferred to the gamble but certain amounts less than \$850 will not. Given that people cannot state indifference and that the amounts used in the procedure were typically in

[26] Using the median of the sample's indifference points for each delay, and the median of the sample's certainty equivalents for each probability, ignores the distribution of these indifference points and certainty equivalents (see Figure V). This approach, therefore, discards crucial information which should be used when drawing inferences from these data. In the analyses conducted in Section V I specifically use all of the information which the data imparts so as to estimate the parameters of PWFs and the uncertainty of these estimates.

increments of $50, the elicitation procedure will always over- or under-estimate this true value.

For example, suppose a subject with this indifference point is offered the choice between $1,000 with a probability of 0.9 and $850 for sure. Given that the subject is indifferent between these options we can assume, for the sake of argument, that she selects either one with equal probability. If the subject selects $850 for this pair but then switches to the gamble when offered $800 with certainty, as her preferences dictate, her derived indifference point will be $825. If, on the other hand, the subject selects $1,000 with probability 0.9 over $850 with certainty, having selected the certain $900 previously, then her derived indifference point will be $875. Thus, unless a person's true indifference point lies midway between two of the certain values used in the elicitation procedure, one will always over- or under-estimate her indifference point. This problem is magnified if people make mistakes in their decisions by, for example, selecting a less preferred option to a more preferred option.

Figure V

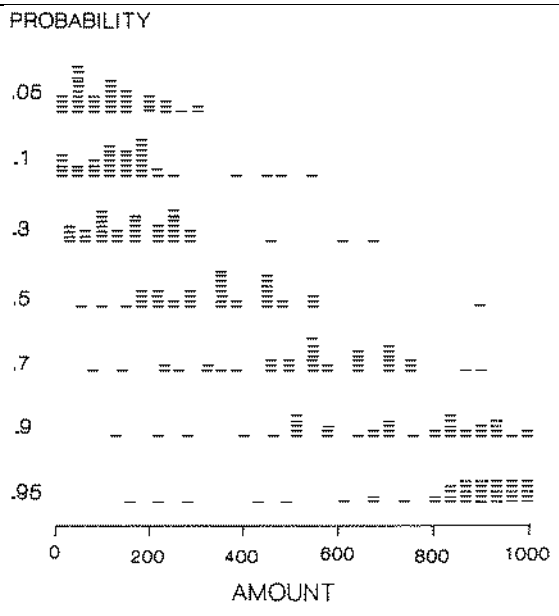

Source: RRC (p. 237), Figure 3.

This methodological issue raises an important statistical issue. Given that the estimate of an indifference point contains some error, one should be cognisant of the uncertainty of the estimate and model it statistically. In other words, every point

estimate has a distribution around it and to ignore this distribution is to assume it is degenerate. As the example above showed, this cannot be the case. Furthermore, to use some measure of central tendency, like the median as in RRC, to select the sample's indifference point for each probability is to ignore the distribution of indifference points in the sample. As Figure V shows, for every probability in the experimental task there is a distribution of indifference points. Statistical techniques, which treat each choice by every subject in the elicitation procedure as a datum, should be used so that important information is not thrown away and the uncertainty surrounding derived indifference points is modelled appropriately.

Thus, the PD model suffers from several theoretical shortcomings, and, in addition, the methodological and statistical approaches used to generate and analyse PD data have notable drawbacks. It seems premature then for Rachlin and colleagues to have claimed that probability is best interpreted as delay, and that they have provided a behavioural foundation for KT's cognitive theory of choice, PT.

In this section I traced the historical development of the PD model. The key feature of this model is that probabilities are interpreted as delays to, or rates of reinforcement of, rewards which therefore ties choice under risk to a temporal framework.[27] I have discussed a number of shortcomings of this model, not least that the link between probability and delay was severed so as to derive a form for the PD function which is similar to a hyperbolic delay discounting function. In the next section I will investigate the relationship between the PD model and other models of choice under risk.

[27] As discussed in Chapter 2, intertemporal risk preferences represent a person's preferences over intertemporal lotteries, the outcomes of which may be serially correlated. By assuming that probability is best interpreted as delay, Rachlin and colleagues not only confound instantaneous risk and time preferences but also fail to recognise the importance of intertemporal risk preferences.

IV. THE PD MODEL AND OTHER THEORIES OF CHOICE UNDER RISK

A natural question when first encountering the PD model is: how does it relate to other theories of choice under risk? To answer this question, one must first understand the lotteries to which the theory applies. The PD model only applies to a limited class of gambles: simple, regular prospects in the terminology used earlier (i.e., those that take the form $(x, p; 0, q)$). The model can incorporate gambles with a negative outcome ($x < 0$) but to keep the discussion simple I will focus exclusively on gambles with a positive outcome ($x > 0$). While the class of simple, regular prospects is undoubtedly interesting, the PD model does not address gambles involving two non-zero outcomes (i.e., gambles with rewards $x > 0$ and $y > 0$), nor gambles with mixed domains (i.e., gambles with rewards $x > 0$ and $y < 0$), nor gambles involving more than two outcomes (i.e., gambles of the form $(x, p; y, q; z, 1 - p - q)$. The model could be extended to incorporate these other types of lotteries but the form that this model would take, and whether the purported link between probability and delay would remain, is not clear.

An issue that has been neglected in the literature on the PD model is that it employs the implicit assumption that utility or value is linear in outcomes. To see this, recall that the PD function takes the following form: $V = x / (1 + \gamma\Theta) = x\pi(\Theta) = x\pi(p)$ as $\Theta = ((1 / p) - 1)$. Researchers use an elicitation procedure to find the value V which is subjectively equivalent to the gamble $(x, p; 0, q)$. Thus, the experimental procedure is being used to elicit the certainty equivalent for gamble $(x, p; 0, q)$. The certainty equivalent (CE) of a gamble is defined as the outcome or amount of money received with certainty that provides the same *utility* as the prospective utility of the gamble:

$$v(\text{CE}) = \pi(p)v(x) \tag{14}$$

By replacing $v(\text{CE})$ with V and $v(x)$ with x, one assumes that utility is linear in outcomes. A large body of empirical research (see, for example, Harless and Camerer (1994), Hey and Orme (1994), Holt and Laury (2002), Andersen, Harrison, Lau and Rutström (2008), Harrison and Rutström (2008)) suggests that v is typically concave in outcomes. To assume that v is linear in x implies that risk preferences in the PD

model are determined solely by the function $\pi(p)$ where $\pi(p) = 1 / (1 + \gamma\Theta)$ and $\Theta = ((1 / p) - 1)$.

Note that apart from the specific functional form for $\pi(p)$, the PD model is therefore exactly Yaari's (1987) dual theory of choice under risk (i.e., the RDEV model) limited to simple, regular prospects; recall that the RDEV model incorporates the potential for non-linear transformations of probabilities while assuming that utility is linear in outcomes. Yaari's model is more general than the PD approach to risk preferences though because it admits different PWFs, lotteries with more than two prizes[28], and lotteries which incorporate both positive and negative prizes (i.e., mixed domain gambles). Thus, the RDEV model is arguably the preferable theory if one wants to assume linear utility.

When one recognises that the PD model is just the dual theory of choice under risk with a specific functional form for the PWF, the question of interest changes to whether this function is useful in applied research. Writing the PWF $\pi(p)$ in terms of probabilities p rather than odds against Θ, it is clear what form this function takes:

$$\pi(p) = p / [p + \gamma(1 - p)] \qquad (15)$$

Figure VI plots this PWF for different values of γ. When $\gamma = 1$, $\pi(p)$ is linear. When $\gamma < 1$, $\pi(p)$ is concave, which represents probability optimism, and when $\gamma > 1$, $\pi(p)$ is convex, which represents probability pessimism and risk aversion. Given that this function is linear, concave or convex throughout its range, it cannot account for overweighting of low probabilities and underweighting of moderate to high probabilities. This function is very similar therefore to the common power PWF: $\pi(p) = p^{\gamma}$.

KT suggest that the PD model's PWF might be too restrictive in that it cannot account for overweighting of low probabilities and underweighting of moderate to high probabilities. I sought to test this by comparing the PD model's PWF to two functions

[28] In these cases, the RDEV model applies rank-dependent non-linear transformations of probabilities so that stochastic dominance is not violated.

commonly used in the literature.[29] Tversky and Kahneman (1992) (TK) popularised the following PWF:

$$\pi(p) = p^{\gamma} / [p^{\gamma} + (1 - p)^{\gamma}]^{1/\gamma}, \qquad (16)$$

which is defined for $1 > p > 0$. This function permits linear, inverse S-shaped and S-shaped forms. Gonzalez and Wu (1999) review the empirical evidence on this function and find that $1 > \gamma > 0$ in most studies. This gives the function an inverse S-shape with overweighting of low probabilities up to a crossover point where $\pi(p) = p$, and then underweighting of moderate to high probabilities.[30]

Figure VI: PD model probability weighting function

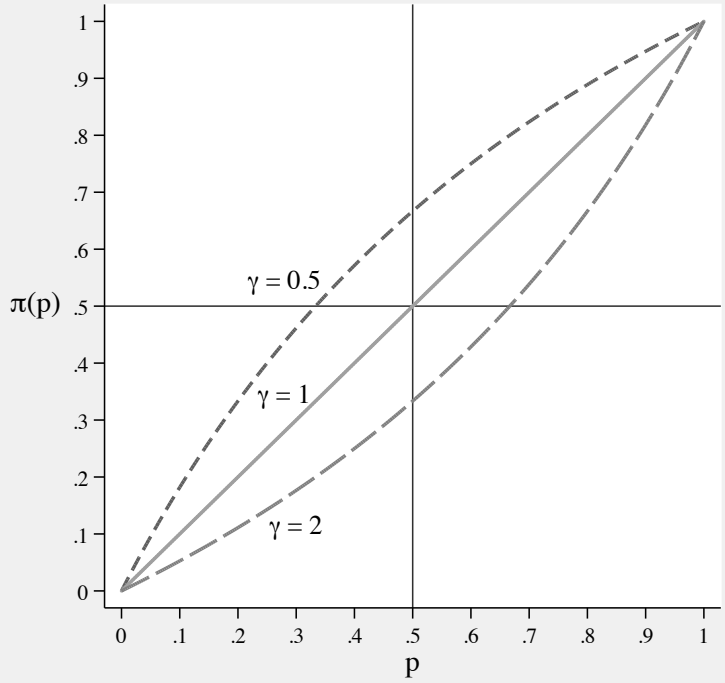

Prelec (1998) axiomatically derived the following two-parameter PWF:

$$\pi(p) = \exp[\, -\eta(-\ln p)^{\gamma}\,], \qquad (17)$$

[29] Stott (2006) reviews the "menagerie" of PWFs which have been developed for models (e.g., PT, RDU, and RDEV) that incorporate subjective distortions of probabilities.
[30] However, Ingersoll (2008) shows that this function is not monotonic at very small values of γ.

which is defined for $1 > p > 0$, $\eta > 0$, and $\gamma > 0$.[31,32] This function exhibits considerable flexibility in that it admits linear, inverse S-shaped and S-shaped forms when $\eta = 1$, and it incorporates objective weighting, underweighting, or overweighting of all probabilities when $\gamma = 1$. Thus, the Prelec function incorporates the qualitative properties of both the PD and TK functions for different parameter values.

The data from Richards, Zhang, Mitchell and de Wit (1999) (RZMW) will be used to test the efficacy of the PD model's PWF in applied research in comparison to those provided by TK and Prelec.

V. AN EMPIRICAL ANALYSIS OF THE PD MODEL'S PWF

RZMW recruited 24 subjects to take part in a within-subject experimental study, using real as opposed to hypothetical rewards, of the acute effects of moderate doses of alcohol on delay and probability discounting. They used the titration procedure discussed at length in Chapter 3, to elicit certainty equivalents for lotteries which paid out \$10 with various probabilities, and to determine the SS rewards which were subjectively equivalent to \$10 at different delays. Participants took part in four experimental sessions or treatments: a pre-placebo session, a post-placebo session, a pre-ethanol session, and a post-ethanol session. Subjects were blind to the treatments.

RZMW fitted hyperbolic and exponential delay discounting functions to each participant's derived indifference points, and fitted hyperbolic and exponential PD functions to each participant's elicited certainty equivalents. The researchers saved the R^2 values for each function and for each participant, and then used them as data to construct tests of whether the hyperbolic or exponential functions provided better fits to the risk and time preference data. Note that using the point estimate of a statistic like R^2 as a datum ignores the uncertainty of this estimate and, thus, does not produce

[31] Prelec (1998, proposition 1, part C, p. 503) provides these parameter restrictions. Prelec (1998, proposition 1, part B, p. 503) constrains $1 > \gamma > 0$, but this constraint can be quite restrictive in practice because it ensures that the PWF is either linear, S-shaped or inverse-S shaped.

[32] Note that I impose these constraints when estimating the models using nonlinear transformations of the parameters.

a valid test of one function's ability to better explain subject choices. Ignoring this issue, RZMW used Wilcoxon matched-pairs signed-rank tests on these *estimated* R^2 values and "found" that the hyperbolic delay and PD functions explained the subject's choices significantly better than the exponential delay and PD functions, respectively.

Given the apparent superiority of the hyperbolic delay and PD functions, only these functions were used in subsequent analyses. Appendix C of RZMW (p. 140) includes all of the elicited certainty equivalents from the probability discounting task in the four experimental sessions, and these data will be used to determine the efficacy of the PD model's PWF in empirical research.[33] These data are used because they are readily available and the study's experimental methodology has been replicated numerous times, both in incentivised (see, for example, Mitchell (1999), Reynolds, Karraker, Horn and Richards (2003), Reynolds, Richards, Horn and Karraker (2004)) and non-incentivised (see Ohmura, Takahashi and Kitamura (2005) and Sheffer et al. (2013)) studies.

RZMW elicited certainty equivalents for simple, regular prospects of the form ($10, p; 0, q) where p took on the values: 1, 0.9, 0.75, 0.5, and 0.25. These certainty equivalents were elicited in each of the four experimental choice sessions and, thus, each participant provides 20 data points for analysis. Note that each subject's elicited certainty equivalents are not independent observations and this should be taken into account in the analyses; RZMW did not cluster the standard errors of the estimates to accommodate this lack of independence.

RZMW converted the five probabilities listed above into odds against winning and then used non-linear least squares (NLLS) estimation to find the best fitting hyperbolic function (13) for these data. RZMW (see Appendix A, p. 138) estimated the value of γ in (13) for each subject in each experimental session, they averaged the certainty equivalents in the pre-placebo and pre-ethanol sessions and then estimated the value of γ for each subject (see Table 1, p. 131), and they used the median

[33] Unfortunately RZMW do not provide the choice data which was used to derive the certainty equivalents and I am forced, therefore, to analyse the certainty equivalents in this section rather than the choice data itself. Nevertheless, the statistical approach I adopt is the appropriate method for analysing these data.

certainty equivalent of the sample for each probability to estimate a grand value of γ for all subjects (see RZMW (p. 132), Figure 4).[34,35]

I adopt a different approach to data analysis which uses all of the data (i.e., elicited certainty equivalents and not statistics like the mean or median of elicited certainty equivalents) provided by all of the subjects across all of the sessions to estimate the parameters of PWFs *at the level of the sample*. In addition, I formally incorporate the fact that each subject made multiple choices in the task and across the sessions. By using all of the information that the data provides while accounting for the lack of independence of observations, I am able to draw accurate inferences from these data. Note that NLLS and maximum likelihood (ML) estimators are consistent as the sample size n tends toward infinity. All of RZMW's estimates rely on 5 observations so it is highly questionable whether any of the asymptotic properties of the estimators can be invoked.

To compare the PD model's PWF to those provided by TK and Prelec, I will use probabilities rather than odds against in my estimating equations. Expressions (13) and (15) can be used derive the non-linear equation for estimating γ when odds against has been transformed back into probability:

$$V_i = x_i(p_i / [p_i + \gamma(1 - p_i)]) + \varepsilon_i, \qquad (18)$$

where the subscript i denotes each observation, V is the elicited certainty equivalent of the reward x = \$10 for a probability p, and ε is the regression error term assumed to be a Normal distribution with zero mean and variance σ^2.

To estimate this model in a ML framework one must explicitly identify the log-likelihood by expanding (18):

$$\ln L_i(\gamma, \sigma; y, X) = \ln \Phi \{[V_i - x_i(p_i / [p_i + \gamma(1 - p_i)])] / \sigma_i\} - \ln \sigma_i, \qquad (19)$$

[34] Note that using the average and/or the median of elicited certainty equivalents discards useful information on the distribution of these data and does not allow one, therefore, to draw accurate inferences from these data.
[35] I have replicated these results but do not present them here because they are available in RZMW.

where Φ is the standard normal density with mean 0 and variance 1, y represents the data that is used to estimate γ and σ (i.e., V, x, and p), and X is a vector of individual characteristics and task parameters like age, gender, and experimental treatment.

An ML framework is attractive because it is straightforward to allow for multiple responses by the same subject (i.e., clustering), to perform non-nested model selection tests which rely on comparisons of the log-likelihoods of each observation in each model, to make the parameter of interest (i.e., γ in (19)) a linear function of observable characteristics, and to estimate a mixture model of the different PWFs.[36] It is also simple to adjust (19) to incorporate the TK and Prelec PWFs.

TABLE I: PROBABILITY WEIGHTING FUNCTION ML ESTIMATES
HOMOGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**
	PD	TK	Prelec
PWF parameter (γ)	1.303***	0.754***	0.739***
	(0.181)	(0.053)	(0.062)
PWF parameter (η)			1.039***
			(0.082)
Sigma (σ)	0.141***	0.136***	0.136***
	(0.016)	(0.014)	(0.014)
N	480	480	480
log-likelihood	257.853	274.844	278.315

Results account for clustering at the individual level
Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Table I presents estimates of the RDEV model which employs the PD, TK and Prelec PWFs and accounts for multiple responses by the same subject. The estimate of γ = 1.303 for the PD function, which is significantly greater than 1 (p = 0.094), implies underweighting of all probabilities. By contrast, the estimate of γ = 0.754 for the TK PWF, which is significantly less than 1 (p < 0.001), implies overweighting of low probabilities and underweighting of moderate to high probabilities. Similarly, the estimate of γ = 0.793 for the Prelec PWF also yields an inverse S-shaped function. The estimate of η = 1.039, however, is not significantly different to 1 (p = 0.637).

[36] With NLLS it is also straightforward to incorporate clustering and to make the parameter of interest a linear function of observable characteristics. Thus, one benefits from an ML approach if one wants to conduct non-nested model selection tests or estimate mixture models.

Figure VII plots the PWFs, and implied decision weights, for the estimates in Table I. The decision weights are graphed for an equi-probable (i.e., 50:50) reference lottery and show the decision weight applied to the worst outcome in the lottery ($0 in the PD task) and the decision weight applied to the best outcome in the lottery ($10 in the PD task). Focussing on the implied decision weights for the PD model, the decision weight applied to the worst outcome ($0) is 0.57 and the decision weight applied to the best outcome ($10) is 0.43. In other words, the PD PWF implies that the probability of the worst outcome is overweighted (0.57), and the probability of the best outcome is underweighted (0.43), relative to the objective probabilities of 0.5.

The same pattern emerges for the other PWFs but it is not as pronounced. In the case of the TK PWF, the decision weight applied to the worst outcome is 0.53 and the decision weight applied to the best outcome is 0.47. Finally, for the Prelec PWF, the decision weight applied to the worst outcome is 0.55 and the decision weight applied to the best outcome is 0.45. To reiterate, these decision weights are based on an equi-probable reference lottery and the decision weights will change as the probabilities in the PD task vary.

Figure VII: Estimated probability weighting functions and implied decision weights
Based on an equi-probable (i.e., 50:50) reference lottery

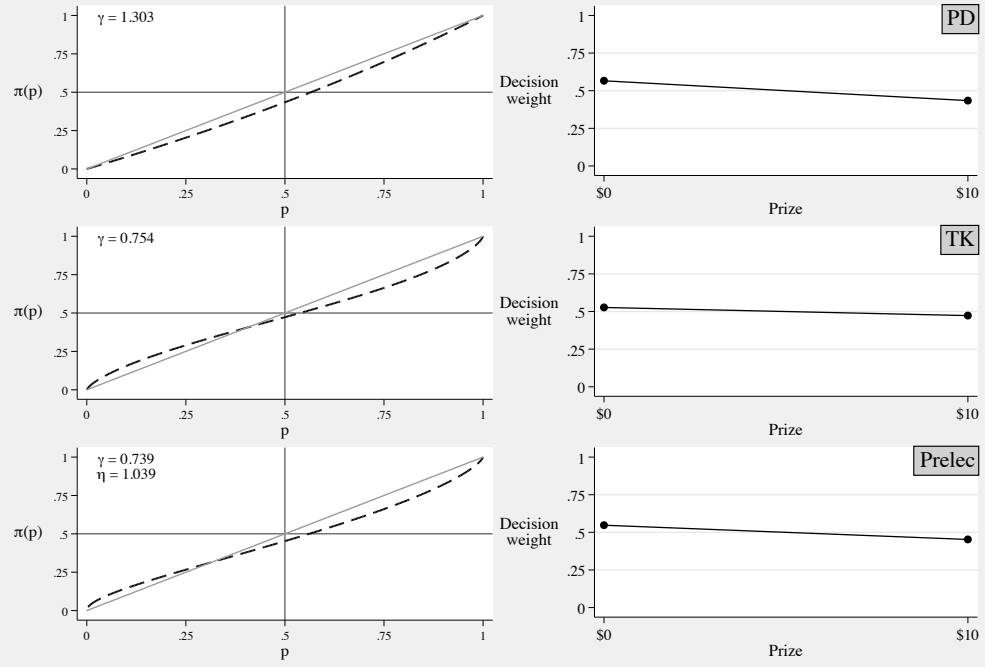

The log-likelihoods for the TK and Prelec functions exceed the log-likelihood for the PD function, suggesting that the TK and Prelec functions better characterise the data. This hypothesis can be tested formally using Vuong (1989) and Clarke (2007) non-nested model selection tests. As discussed in Chapter 3, the Clarke test is asymptotically more efficient and has greater power in discriminating between models than the Vuong test when the distribution of the models' individual log-ratios is highly peaked. In other words, when the distribution of these log-ratios is leptokurtic[37], the Clarke test is superior, from both statistical efficiency and power perspectives, to the Vuong test.

Figure VIII: Distribution of individual log-ratios

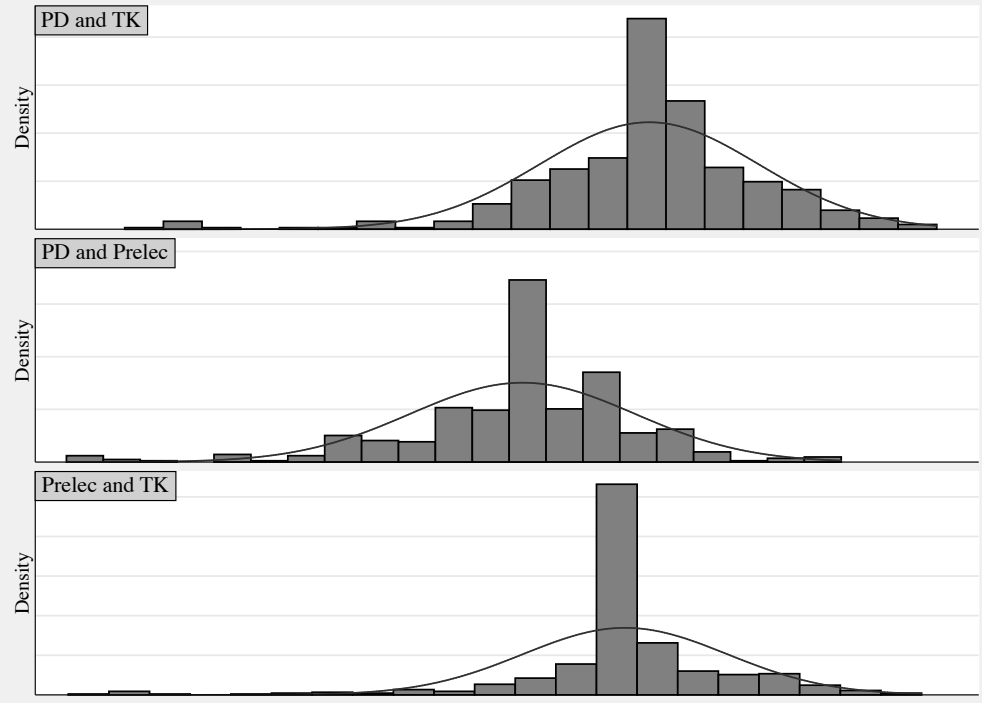

Figure VIII plots the distribution of the individual log-ratios, with a normal density overlay, for the three PWF comparisons. The distribution of these log-ratios is leptokurtic (i.e., highly peaked) which suggests that the Clarke test is more appropriate for these data. The Clarke test yields a test statistic based on the binomial distribution which must be compared to a critical value to determine which model, in a pairwise comparison, receives the most support in the data. A Clarke test comparing

[37] The normal distribution is the quintessential mesokurtic distribution. A distribution which has positive excess kurtosis (i.e., a highly peaked distribution) is leptokurtic.

the PD and TK functions yields a test statistic of 197, which is below the critical value of 240, implying that the TK function better characterises the data ($p < 0.001$).

A Clarke test comparing the PD PWF and the Prelec PWF finds in favour of the Prelec function ($p < 0.001$). Finally, a Clarke test of the TK and Prelec PWFs suggests that the Prelec function finds more support in the data ($p < 0.001$). Based on these tests the following transitive ranking of PWFs emerges: the Prelec function finds the most support in the data, followed by the TK function, and then the PD function.

As mentioned earlier, it is straightforward to make the parameter(s) of interest in our models a linear function of observable characteristics and thereby admit heterogeneity in the PWF estimates. RZMW collected each participant's gender and Appendix C (p. 140) groups the elicited certainty equivalents by experimental session and divides participants by whether they received a low or high dose of ethanol. Table II presents estimates of the RDEV model which incorporates these variables in the three models.

Gender and ethanol dose are not statistically significant in any of the models and none of the experimental treatment variables are statistically significant in the TK model. However, the estimate of γ in the PD model is significantly higher in the pre-ethanol session than in the pre-placebo session (the omitted base category). Wald tests show that the estimate in the pre-ethanol session is also significantly greater than the estimates in the post-placebo ($p = 0.038$) and the post-ethanol ($p = 0.034$) sessions. This result is contrary to the hypothesis that ethanol increases probabilistic discounting, which the researchers set out to test, and differs to RZMW who found no statistically significant difference between the estimates of γ in the pre- and post-ethanol sessions.[38] Note that RZMW used *estimates* of γ as *data* to conduct *t*-tests of potential differences across the pre- and post-ethanol sessions. The valid approach to analysis which I have adopted uses all of the information which a dataset imparts to estimate the parameters of a model and conduct hypothesis tests on these estimates. These differences in analysis likely explain the contradictory findings.

[38] RZMW did not compare the estimate of γ in the pre-ethanol session to the estimates of γ in the pre-placebo and post-placebo sessions.

The estimates of γ for the Prelec PWF do not differ significantly according to observable characteristics and task parameters but the estimate of η in the pre-ethanol session is significantly higher than in the pre-placebo session (the omitted base category). In addition, Wald tests show that the estimate of η in the pre-ethanol session is significantly higher than estimates in the post-placebo ($p = 0.017$) and post-ethanol ($p < 0.038$) sessions. These results mirror those for γ in the PD model.

TABLE II: PROBABILITY WEIGHTING FUNCTION ML ESTIMATES
HETEROGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**
	PD	TK	Prelec
PWF parameter (γ)			
Male	0.035	-0.011	0.007
	(0.341)	(0.096)	(0.130)
Ethanol - high dose	0.45	-0.08	0.023
	(0.378)	(0.107)	(0.130)
Post-placebo session	-0.006	-0.002	-0.002
	(0.048)	(0.018)	(0.027)
Pre-ethanol session	0.223**	-0.045	0.005
	(0.108)	(0.031)	(0.045)
Post-ethanol session	0.062	0.005	0.044
	(0.099)	(0.036)	(0.042)
Constant	1.003***	0.811***	0.715***
	(0.184)	(0.073)	(0.118)
PWF parameter (η)			
Male			0.027
			(0.172)
Ethanol - high dose			0.215
			(0.172)
Post-placebo session			0.001
			(0.021)
Pre-ethanol session			0.111**
			(0.045)
Post-ethanol session			0.056
			(0.052)
Constant			0.880***
			(0.116)
Sigma (σ)			
Constant	0.138***	0.135***	0.132***
	(0.016)	(0.014)	(0.014)
N	480	480	480
log-likelihood	269.434	278.331	289.727

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The Clarke tests conducted earlier are based on the implicit assumption that the observations are produced by only one data generating process (DGP) (i.e., the PD PWF or the TK PWF) when more may be present in the data. In other words, the PD

function may explain some choices in the data better than the TK function whereas the TK function may explain other choices in the data better than the PD function. The assumption that only one DGP characterises all of the data, clearly precludes such a possibility.

Mixture models[39] allow two or more DGPs to account for the data and also provide a measure of the proportion of the data which is explained by each process. In the current context, one can estimate a mixture model of, say, the PD and TK PWFs and then ask the data to tell one how much support each function has. To do so one specifies a "grand likelihood" function which is just a probability-weighted average of the likelihoods of the two models.

Letting π^{PD} represent the probability that the PD model's PWF is correct, and $\pi^{TK} = (1 - \pi^{PD})$ the probability that the TK function is correct, the grand likelihood is the probability-weighted average of the two conditional likelihoods L^{PD} and L^{TK} for the PD and TK models, respectively. Thus, the likelihood for the mixture model is given by:

$$\ln L_i(\gamma^{PD}, \gamma^{TK}, \kappa; y, X) = \Sigma_i \ln [(\pi^{PD} \times L^{PD}) + (\pi^{TK} \times L^{TK})], \qquad (20)$$

where κ is a parameter which defines the log odds of the probability of the PD model: $\pi^{PD} = 1 / (1 + \exp(\kappa))$. Note that this transformation allows the parameter κ to take on any value during the maximisation process but it constrains the probabilities π^{PD} and π^{TK} to lie within the unit interval. The grand likelihood in (20) is maximised to estimate the parameters of each model and the weight accorded to each model in the data.

Table III presents estimates of the mixture model of the PD and TK PWFs. The estimate of $\gamma^{PD} = 4.802$ is large and implies extreme underweighting of probabilities but the 95% confidence interval shows that it is estimated very imprecisely. The mixture probability $\pi^{PD} = 0.130$ implies that approximately 13% of the data is best characterised by the PD model's PWF but this estimate is not significantly different to

[39] For detailed discussions of mixture models consult McLachlan and Peel (2000), Harrison and Rutström (2009), and Conte, Hey and Moffatt (2011).

zero ($p = 0.146$).[40] The estimate of $\gamma^{TK} = 0.855$, which is significantly less than 1 ($p = 0.005$) implies overweighting of low probabilities and underweighting of moderate to high probabilities. Finally, the estimate of $\pi^{TK} = 0.870$ implies that approximately 87% of the data is best characterised by the TK PWF, although I cannot reject the hypothesis that this estimate is equal to 1 ($p = 0.146$). Thus, the PD model's PWF finds almost no support in the data, even when it is allowed to account for only a fraction of the choices in the experiment.

TABLE III: MIXTURE MODEL ML ESTIMATES
PD AND TK FUNCTIONS

	Estimate	Std error	*p*-value	95% Confidence interval	
		PD probability weighting function			
PWF parameter (γ^{PD})	4.802**	2.163	0.026	0.563	9.042
Mixture probability (π^{PD})	0.130	0.089	0.146	-0.045	0.304
		TK probability weighting function			
PWF parameter (γ^{TK})	0.855***	0.052	0.000	0.752	0.957
Mixture probability (π^{TK})	0.870***	0.089	0.000	0.696	1.045
		Sigma			
Constant (σ)	0.110***	0.016	0.000	0.079	0.141
N	480				
log-likelihood	298.681				

H_0: $\pi^{TK} = 1$, p-value = 0.146

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The preceding results suggest that the PD model's PWF is potentially too restrictive when trying to classify risk preferences using the RDEV model. According to Clarke tests, the TK and Prelec PWFs find more support in the data than the PD function. A

[40] The lower bound of the 95% confidence interval for the mixture probability π^{PD} is less than 0 and the upper bound of the 95% confidence interval for the mixture probability π^{TK} is greater than 1, even though the log odds transformation constrains these probabilities to lie within the unit interval. These values lie outside the unit interval due to the use of the delta method (see Oehlert (1992)), which is an approximation, much like a Taylor series, to transform κ into the mixture probabilities.

mixture model of the PD and TK PWFs confirms this result and finds almost no support for the PD model in the data.[41,42]

VI. DISCUSSION AND CONCLUSIONS

This chapter reviewed the development of the PD model so as to link it to other theories of choice under risk and to highlight its theoretical and empirical shortcomings. The PD model attempts to tie choice under risk to a temporal framework by reinterpreting probability as delay to, or rate of reinforcement of, reward. This interpretation is loose in the context of one-shot gambles, and RRC abandoned the purported link between probability and delay so as to develop a PWF which has a similar form to a hyperbolic delay discounting function when probabilities have been transformed into odds against winning. This chapter showed that the PD model is formally isomorphic to Yaari's (1987) dual theory of choice under risk, limited to a certain class of lotteries and with a specific functional form for the PWF.

The theoretical issues with the PD model are matched by the methodological problem of using an experimental procedure, which is not incentive compatible, to elicit certainty equivalents. Furthermore, the statistical tools applied to these data do not take into account the DGP and the use of estimates as data is invalid. Finally, estimators which only have attractive asymptotic properties should not be used with tiny samples.

An obvious critique of the analyses conducted in Section V is that I did not incorporate concave utility when comparing the virtues of different PWFs. This

[41] Appendix D presents estimates from a mixture model of the TK and Prelec PWFs and discusses problems I encountered when trying to estimate a mixture model of the Prelec and PD PWFs.

[42] Harrison and Rutström (2009) extend mixture models to incorporate *observable* heterogeneity in individual characteristics and treatments. Conte, Hey and Moffatt (2011), by contrast, show that one can estimate the mixture probability as a random coefficient, and thereby attain a better econometric characterisation of *unobserved* heterogeneity. I adopt the approach of Harrison and Rutström (2009) in Chapters 3 and 5 by making the mixture probability a linear function of observable characteristics and task parameters. I do not incorporate observable heterogeneity in estimates of the mixture probability in this chapter because the focus is on the ability of different PWFs to explain the data, under the assumption of homogenous preferences; incorporating observable heterogeneity is also not warranted given the small sample size of the RZMW data.

choice was deliberate because I wanted to judge the PD model as it was formulated rather than reformulate it to account for risk attitudes which are affected by the shape of the utility function.[43]

Another critique of the analyses conducted in Section V is that I focussed on only two PWFs when numerous others exist in the literature. I selected the TK and Prelec functions because they have been used extensively, they allow for overweighting and underweighting of probabilities, and the two-parameter Prelec function is remarkably flexible. The crucial finding is that the RZMW data is better characterised by functions which allow for overweighting and underweighting of probabilities, which implies that the PD function may be too restrictive in some circumstances.

In sum, the PD model adds little to our understanding of choice under risk and suffers from a number of theoretical and empirical limitations. Researchers in psychology and addiction studies would be well served by embracing other theories of choice under risk, using experimental methods which promote, rather than hinder, elicitation of preferences, and statistical techniques that are appropriate to the data generated by these experiments.

[43] In Chapter 5, I estimate RDU models which allow risk preferences to be determined both by the curvature of the utility function and probability weighting.

REFERENCES

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583-618.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2014): "Discounting Behavior: A Reconsideration," *European Economic Review*, 71, 15-33.

BARDSLEY, N., R. CUBITT, G. LOOMES, P. MOFFATT, C. STARMER, AND R. SUGDEN (2010): *Experimental Economics: Rethinking the Rules*. Princeton, NJ: Princeton University Press.

CHUNG, S. H. (1965): "Effects of Delayed Reinforcement in a Concurrent Situation," *Journal of the Experimental Analysis of Behavior*, 8, 439-44.

CHUNG, S. H., AND R. J. HERRNSTEIN (1967): "Choice and Delay of Reinforcement," *Journal of the Experimental Analysis of Behavior*, 10, 67-74.

CLARKE, K. A. (2007): "A Simple Distribution-Free Test for Nonnested Model Selection," *Political Analysis*, 15, 347-363.

CONTE, A., J. D. HEY, AND P. G. MOFFAT (2011): "Mixture Models of Choice under Risk," *Journal of Econometrics*, 162, 79-88.

DAVENPORT, J. W. (1962): "The Interaction of Magnitude and Delay of Reinforcement in Spatial Discrimination," *Journal of Comparative and Physiological Psychology*, 55, 267-73.

GONZALEZ, R., AND G. WU (1999): "On the Shape of the Probability Weighting Function," *Cognitive Psychology*, 38, 129-166.

GREEN, L., AND J. MYERSON (2004): "A Discounting Framework for Choice with Delayed and Probabilistic Rewards," *Psychological Bulletin*, 130, 769-92.

HARLESS, D. W., AND C. F. CAMERER (1994): "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica*, 62, 1251-1289.

HARRISON, G. W. (1989): "Theory and Misbehavior of First-Price Auctions," *American Economic Review*, 79, 749-762.

HARRISON, G. W. (1992): "Theory and Misbehavior of First-Price Auctions: Reply," *American Economic Review*, 82, 1426-1443.

HARRISON, G. W. (1994): "Expected Utility Theory and the Experimentalists," *Empirical Economics*, 19, 223-253.

HARRISON, G. W. (2007): "House Money Effects in Public Good Experiments: Comment," *Experimental Economics*, 10, 429-437.

HARRISON, G. W., AND E. E. RUTSTRÖM (2008): "Risk Aversion in the Laboratory," in *Research in Experimental Economics: Volume 12. Risk Aversion in Experiments*, ed. by J. C. Cox, and G. W. Harrison. Bingley: Emerald, 41-196.

HARRISON, G. W., AND E. E. RUTSTRÖM (2009): "Expected Utility Theory and Prospect Theory: One Wedding and a Decent Funeral," *Experimental Economics*, 12, 133-158.

HERRNSTEIN, R. J. (1961): "Relative and Absolute Strength of Response as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Behavior*, 4, 267-72.

HERRNSTEIN, R. J. (1970): "On the Law of Effect," *Journal of the Experimental Analysis of Behavior*, 13, 243-66.

HEY, J. D., AND C. ORME (1994): "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62, 1291-1326.

HOLT, C. A., AND S. K. LAURY (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92, 1644-1655.

INGERSOLL, J. (2008): "Non-Monotonicity of the Tversky-Kahneman Probability-Weighting Function: A Cautionary Note," *European Financial Management*, 14, 385-390.

KAGEL, J. H., R. C. BATTALIO, AND L. GREEN (1995): *Economic Choice Theory: An Experimental Analysis of Animal Behavior*. New York, NY: Cambridge University Press.

KAHNEMAN, D., AND A. TVERSKY (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263-291.

KEREN, G., AND P. ROELOFSMA (1995): "Immediacy and Certainty in Intertemporal Choice," *Organizational Behavior and Human Decision Processes*, 63, 287-297.

LOGAN, F. A. (1965): "Decision Making by Rats: Delay Versus Amount of Reward," *Journal of Comparative and Physiological Psychology*, 59, 1-12.

MADDEN, G. J., AND W. K. BICKEL (2010): *Impulsivity: The Behavioral and Neurological Science of Discounting*. Washington, DC: American Psychological Association.

MAZUR, J. E. (1984): "Tests of an Equivalence Rule for Fixed and Variable Reinforcer Delays," *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 426-437.

MCLACHLAN, G., AND D. PEEL (2000): *Finite Mixture Models*. New York, NY: Wiley.

MITCHELL, S. H. (1999): "Measures of Impulsivity in Cigarette Smokers and Non-Smokers," *Psychopharmacology (Berlin)*, 146, 455-64.

MORGAN, P. B., AND D. TUSTIN (1992): "The Perception and Efficiency of Labour Supply Choices by Pigeons," *The Economic Journal*, 102, 1134-1148.

OEHLERT, G. W. (1992): "A Note on the Delta Method," *The American Statistician*, 46, 27-29.

OHMURA, Y., T. TAKAHASHI, AND N. KITAMURA (2005): "Discounting Delayed and Probabilistic Monetary Gains and Losses by Smokers of Cigarettes," *Psychopharmacology (Berlin)*, 182, 508-15.

OSTASZEWSKI, P. (1997): "Temperament and the Discounting of Delayed and Probabilistic Rewards: Conjoining European and American Psychological Traditions," *European Psychologist*, 2, 35-43.

PRELEC, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497-527.

QUIGGIN, J. (1982): "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3, 323-343.

RACHLIN, H., A. CASTROGIOVANNI, AND D. CROSS (1987): "Probability and Delay in Commitment," *Journal of the Experimental Analysis of Behavior*, 48, 347-53.

RACHLIN, H., AND L. GREEN (1972): "Commitment, Choice and Self-Control," *Journal of the Experimental Analysis of Behavior*, 17, 15-22.

RACHLIN, H., A. W. LOGUE, J. GIBBON, AND M. FRANKEL (1986): "Cognition and Behavior in Studies of Choice," *Psychological Review*, 93, 33-45.

RACHLIN, H., A. RAINERI, AND D. CROSS (1991): "Subjective Probability and Delay," *Journal of the Experimental Analysis of Behavior*, 55, 233-44.

REYNOLDS, B., K. KARRAKER, K. HORN, AND J. B. RICHARDS (2003): "Delay and Probability Discounting as Related to Different Stages of Adolescent Smoking and Non-Smoking," *Behavioural Processes*, 64, 333-344.

REYNOLDS, B., J. B. RICHARDS, K. HORN, AND K. KARRAKER (2004): "Delay Discounting and Probability Discounting as Related to Cigarette Smoking Status in Adults," *Behavioural Processes*, 65, 35-42.

RICHARDS, J. B., L. ZHANG, S. H. MITCHELL, AND H. DE WIT (1999): "Delay or Probability Discounting in a Model of Impulsive Behavior: Effect of Alcohol," *Journal of the Experimental Analysis of Behavior*, 71, 121-43.

SHEFFER, C. E., M. MENNEMEIER, R. D. LANDES, W. K. BICKEL, S. BRACKMAN, J. DORNHOFFER, T. KIMBRELL, AND G. BROWN (2013): "Neuromodulation of Delay Discounting, the Reflection Effect, and Cigarette Consumption," *Journal of Substance Abuse Treatment*, 45, 206-14.

SILBERBERG, A., P. MURRAY, J. CHRISTENSEN, AND T. ASANO (1988): "Choice in the Repeated-Gambles Experiment," *Journal of the Experimental Analysis of Behavior*, 50, 187-195.

SMITH, V. L. (1982): "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, 923-955.

STOTT, H. P. (2006): "Cumulative Prospect Theory's Functional Menagerie," *Journal of Risk and Uncertainty*, 32, 101-130.

TVERSKY, A., AND D. KAHNEMAN (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297-323.

VUONG, Q. H. (1989): "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 57, 307-333.

WEBER, B., AND G. B. CHAPMAN (2005): "The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect?," *Organizational Behavior and Human Decision Processes*, 96, 104-118.

YAARI, M. E. (1987): "The Dual Theory of Choice under Risk," *Econometrica*, 55, 95-115.

5. RISK AND TIME PREFERENCES AND SMOKING BEHAVIOUR

I. INTRODUCTION

The literature review in Chapter 3 suggested that there is a positive relationship between smoking and discounting behaviour: smokers tend to discount the future more heavily than non-smokers. However, the bulk of studies in this literature rely on small samples, hypothetical rewards, restrictive assumptions about the form that discounting takes, and statistical tools which are not appropriate to the data that are collected. I sought to remedy some of these shortcomings by investigating the smoking-discounting relationship using two relatively large samples, four distinct discounting functions, and a maximum likelihood estimation framework which is consistent with the process(es) which generate time preference data.

However, the analyses in Chapter 3 did not directly incorporate the shape of people's utility functions in the estimation of discounting parameters.[1] This is a crucial issue because time preferences are defined over time-dated utility flows, not flows of money. These are equivalent if a utility function is linear but Andersen, Harrison, Lau and Rutström (2008) (AHLR) showed that if a utility function is concave then the assumption of linearity will bias the estimation of discounting parameters upwards. Thus, to draw accurate inferences about discounting behaviour it is important to incorporate utility function curvature in the estimation of discounting models.

In this chapter I report the results from a set of instantaneous risk[2] and time preference experiments conducted on a sample of student smokers and non-smokers at the University of Cape Town (UCT) in 2012. The experimental design allows me to

[1] In Appendix B I present a parametric approach to this identification problem which involves the estimation of "profile likelihoods" to determine the optimal shape of the sample's utility function given the data. This approach has its merits but is second-best to collecting experimental data which allows one to estimate the shape of the utility function jointly with the parameters of discounting models. I will adopt this latter approach in this chapter.

[2] The discussion in chapter 2 highlighted the importance of both instantaneous and intertemporal risk preferences in the analysis of choice behaviour. As mentioned previously, this thesis only empirically examines instantaneous risk preferences so all subsequent references to "risk preferences" refer to the instantaneous or atemporal variety, unless otherwise noted.

explore potential differences in the risk and time preferences of smokers and non-smokers and jointly estimate utility function curvature and discounting behaviour. I find no significant differences in the risk preferences of smokers and non-smokers but find that smokers discount significantly more heavily than non-smokers. These results are robust to different assumptions about the way people evaluate lotteries and the way they discount utility flows. In addition, I find that smokers may be more likely to discount hyperbolically than non-smokers, which, under the assumption of an additively-separable intertemporal utility function, means they may be more prone to time inconsistency.

The chapter is structured as follows. In Section II I review previous research on the relationship between risk preferences and smoking behaviour. This section complements the review in Chapter 3 of the relationship between time preferences and smoking behaviour and ties in the discussion of probability discounting in Chapter 4. In Section III I discuss the experimental design and present summary statistics for the sample and in Section IV I formulate the statistical approach to data analysis. Section V presents the results and Section VI concludes.

II. A REVIEW OF THE SMOKING AND RISK PREFERENCES LITERATURE

The decision to smoke involves clear risks, like the potential for negative health consequences, and is made under conditions of uncertainty, i.e., without knowing one's susceptibility to these risks. It is presumably linked, therefore, to an individual's attitudes toward risk and uncertainty. Specifically, smokers may be less risk averse than non-smokers and this greater penchant to take on risk may incline them to smoke. In this section I review the experimental literature on risk preferences and smoking behaviour.

Table I provides a detailed summary of studies investigating this relationship. Online searches of PubMed and Econlit, employing the search criteria "smoking" and "risk preferences" and their variants (e.g., "smoke," "risk", and "probability discounting"), were used to locate these papers. Unlike the literature on time preferences and smoking behaviour there is a dearth of studies analysing the risk preferences of

smokers and non-smokers. An initial list of studies was trimmed according to the following rules: the study had to include a clear smoker, non-smoker comparison[3]; and study participants had to make choices between lotteries[4] involving amounts of money, rather than cigarettes or quality-adjusted life years.[5] The 11 studies satisfying these criteria are listed in Table I.

Mitchell (1999) conducted the first experimental study investigating the risk preferences of 20 relatively heavy[6], current smokers ($N_S = 20$) and 20 never-smokers ($N_{NS} = 20$). She presented subjects with 137 choice questions between a lottery which paid out \$10 with specific probabilities ($p = 0.1, 0.25, 0.5, 0.75, 0.9$, and 1) and \$0 with the complementary probability (i.e., the lottery (\$10, p; \$0, $1 - p$) using the notation from Chapter 4), and a sure amount of money which varied between \$0.01 and \$10.50. The questions were drawn randomly from this battery, without replacement, and presented to subjects sequentially. At the end of the experiment, one of a subject's choices was selected randomly for payment.

Mitchell used each subject's choices to determine a certainty equivalent for the lottery (\$10, p; \$0, $1 - p$) at different values of p. For example, if a subject chose \$4 over the lottery (\$10, 0.5; \$0, 0.5) but then chose the lottery (\$10, 0.5; \$0, 0.5) over \$3.50, the subject was assigned a certainty equivalent of \$3.75. Taking the average of these two values is arbitrary and doing so throws away information about the uncertainty of this estimate; all that one can infer from this pattern of choices is that a subject's certainty equivalent lies in the open interval (\$3.50, \$4). Interval data of this form is analysed appropriately using interval regression methods but Mitchell used the *estimated* certainty equivalents (i.e., the point estimate \$3.75 in the example) as *data* to construct Mann-Whitney tests of differences in the certainty equivalents of heavy smokers and never-smokers; no significant differences between these groups were found.

[3] Lawyer, Schoepflin, Green and Jenks (2011) investigate whether the risk preferences of smokers and non-smokers differ when they make choices over hypothetical or real rewards. However, the researchers do not compare the risk preferences of smokers and non-smokers so this study is not included in Table I.
[4] A number of studies (e.g., Bradford (2010), Jusot and Khlat (2013)) use survey questions which try to elicit general attitudes toward risk and were excluded for this reason.
[5] van der Pol and Ruggeri (2008) investigate risk preferences over hypothetical health outcomes.
[6] The smokers in Mitchell's (1999) study stated that they smoked at least 15 cigarettes per day and provided a breath sample to verify their smoking status.

TABLE I: REVIEW OF EXPERIMENTAL LITERATURE ON SMOKING AND RISK PREFERENCES

Study	Sample (size)	Elicitation method	Incentives (max prize)	Probabilities	Models (estimated rates)	Statistical method (valid?)	Significant relationship with smoking?
Mitchell (1999)	Adults in Durham, NH, USA ($N_S = 20$, $N_{NS} = 20$)	Choice (random)	Yes ($10)	0.1, 0.25, 0.5, 0.75, 0.9, 1	PD ($\gamma_S = 1.328$) ($\gamma_{NS} = 1.371$)	NLLS for risk aversion, non-parametric tests for analysis. (not valid)	No.
Reynolds, Karraker, Horn & Richards (2003)	Adolescents in Morgantown, WV, USA ($N_S = 19$, $N_{NS} = 19$, $N_T = 17$)	Titration (random - Richards et al. (1999))	1-out-of-2-tasks ($10)	0.25, 0.5, 0.75, 0.9, 1	PD NRD but from Figure 2: ($\gamma_S = 1.610$) ($\gamma_{NS} = 1.110$) ($\gamma_T = 3.820$)	NLLS for risk aversion, ANOVA for analysis. (not valid)	Yes, positive for T relative S ($p<0.05$) and NS ($p<0.05$); No for S relative to NS.
Reynolds, Richards, Horn & Karraker (2004)	Mostly students in Morgantown, WV, USA ($N_S = 25$, $N_{NS} = 29$)	Titration (random - Richards et al. (1999))	1-out-of-2-tasks ($10)	0.25, 0.5, 0.75, 0.9, 1	PD ($\gamma_S = 1.910$) ($\gamma_{NS} = 1.470$)	NLLS for risk aversion, ANOVA for analysis. (not valid)	Yes ($p<0.05$), positive (smokers are more risk averse)
Ohmura, Takahashi & Kitamura (2005)	Students in Sapporo, Japan ($N_S = 27$, $N_{NS} = 23$)	Titration (random - Richards et al. (1999))	No (¥100,000 = $1000)	0.1, 0.3, 0.5, 0.7, 0.9	PD and AUC $AUC_S = 0.230$ $AUC_{NS} = 0.180$	AUC and NLLS for risk aversion, correlations and t-tests for analysis. (not valid)	Yes ($p=0.08$), negative (smokers are less risk averse)
Reynolds (2006)	Adults in Buffalo, NY, USA ($N_S = 15$, $N_{NS} = 15$)	Titration (random - Richards et al. (1999))	No ($10)	0.25, 0.5, 0.75, 0.9, 1	PD ($\gamma_S = 3.908$) ($\gamma_{NS} = 1.574$)	NLLS for risk aversion, non-parametric tests for analysis. (not valid)	Yes ($p<0.05$), positive (smokers are more risk averse)
Reynolds et al. (2007)	Adolescents in Columbus, OH, USA ($N_S = 25$, $N_{NS} = 26$)	Titration (random - Richards et al. (1999))	Yes ($10)	0.25, 0.5, 0.75, 0.9, 1	AUC and PD (Not reported)	AUC for risk aversion, ANOVA for analysis. (not valid)	No.
Yi, Chase & Bickel (2007)	Adults in Little Rock, AR, USA ($N_S = 30$, $N_{NS} = 29$)	Titration (ordered)	No ($1000)	0.01, 0.05, 0.1, 0.25, 0.5, 0.75, 0.95	PD and AUC (Not reported)	NLLS and AUC for risk aversion, ANOVA for analysis. (not valid)	No when analysing all the data; Yes ($p<0.05$), positive, when using only probabilities ≥ 0.5.
Anderson & Mellor (2008)	Adults subjects in Williamsburg, VA, USA ($N_S \approx 79$, $N_{NS} \approx 898$)	Choice (ordered - MPL)	Yes ($11.55)	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1	CRRA (1-r) ($r = 0.257$)	Algebra and averaging for risk aversion, probit model for analysis. (not valid)	Yes ($p<0.1$), negative (smokers are less risk averse).
Harrison, Lau & Rutström (2010)	Adults in Denmark ($N_S = 71$, $N_{NS} = 181$)	Choice (ordered - MPL)	1-in-10-chance ($687)	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1	CRRA (1-r) ($r_{S(men)} = 0.729$) ($r_{NS(men)} = 0.746$) ($r_{S(women)} = 0.811$) ($r_{NS(women)} = 0.755$)	ML for risk aversion and analysis. (valid)	Men: No; Women: Yes ($p<0.06$), positive (smokers are more risk averse)
Szrek, Chao, Ramlagan & Peltzer (2012)	Adults in Witbank, South Africa ($N_S \approx 59$, $N_{NS} \approx 292$)	Choice (ordered - MPL)	Yes (R48 ≈ $7)	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1	CRRA (1-r) ($r = 0.35$, SD = 0.62)	Algebra and averaging for risk aversion, logit model for analysis. (not valid)	No.
Poltavski & Weatherly (2013)	Students in Grand Forks, ND, USA ($N_S = 16$, $N_{LS} = 74$, $N_{NS} = 92$)	Choice (random)	No ($100,000)	0.01, 0.1, 0.5, 0.9, 0.99	PD and AUC ($1000: $\gamma_S = 0.118$, $\gamma_{LS} = 0.134$, $\gamma_{NS} = 0.307$) ($100,000: $\gamma_S = 0.031$, $\gamma_{LS} = 0.167$, $\gamma_{NS} = 0.181$)	NLLS and AUC for risk aversion, ANOVA for analysis. (not valid)	Yes, negative for S relative to NS ($p<0.05$); No for all other comparisons.

Source: Author's construction.

Notes: S = smoker; NS = non-smoker/never-smoker; LS = light smoker; T = trier; PD = probability discounting; AUC = area under the curve; NRD = not reported directly.

NLLS = non-linear least squares; ML = maximum likelihood; ANOVA = analysis of variance.

In addition, Mitchell fitted the probability discounting (PD) probability weighting function (PWF)[7] to the certainty equivalents for each subject and then compared the *estimated* PWF parameters of smokers and never-smokers. She found evidence of risk aversion in both groups (i.e., $\gamma > 1$) but no significant differences in the risk preferences of smokers and never-smokers. Echoing the issue raised earlier, using the point estimate of any parameter as a datum ignores the uncertainty of this estimate and should not be used for inferential purposes.[8]

Reynolds, Karraker, Horn and Richards (2003) (RKHR) found that risk preferences differ according to smoking status but perhaps not in the way that would be expected, and with a statistical approach which is not valid. RKHR used the titration algorithm of Richards, Zhang, Mitchell and de Wit (1999), which was discussed extensively in Chapter 3, to elicit certainty equivalents for the lottery ($10, p; $0, $1 - p$) at different values of p among adolescent smokers ($N_S = 19$), adolescent never-smokers ($N_{NS} = 19$), and adolescent "triers" ($N_T = 17$).[9,10] Subjects also completed a delay discounting task and they were paid for one choice across both tasks; this payment scheme is referred to as 1-out-of-2-tasks in Table I.

RKHR fitted the PD PWF to the *estimated* certainty equivalents, using non-linear least squares (NLLS) estimation, and then used the *estimated* PWF parameters as *data* in an ANOVA model so as to compare the three smoking status groups. For reasons

---

[7] Chapter 4 showed that the PD model is just Yaari's (1987) dual theory of choice under risk limited to a circumscribed class of lotteries and with a specific PWF: $\pi(p) = p / [p + \gamma(1 - p)]$; if $\gamma > 1$ this represents probability pessimism and risk aversion. Recall that to derive this form for the PWF, Rachlin, Raineri and Cross (1991) severed the purported link between probability and delay which they argued was so central to the way people interpret probabilities. Thus, the specific form of this PWF lacks consistent theoretical foundations.

[8] The seventh column of Table I lists the statistical method that was adopted in each study and provides a binary summary judgement (i.e., valid or not valid), in parentheses, of whether the statistical approach was valid given the data obtained in the experiment. This binary summary judgement does not imply that the estimates which the researchers obtained were "wrong" but rather that the method used to derive the estimates was not appropriate for the data.

[9] "Triers" had smoked cigarettes for the first time in the 6 months prior to the study and they smoked an average of 3.76 cigarettes in total over this time span. Smokers, by contrast, had smoked every week for at least 6 months prior to the study and they smoked 46.42 cigarettes, on average, per week.

[10] As discussed in Chapter 3, a titration algorithm is susceptible to being "gamed" by subjects because it narrows the search for the interval within which a person's certainty equivalent, for a particular value of p, lies by making the choices an experimental subject faces contingent on his prior choices. Thus, titration procedures lack incentive compatibility. Moreover, RKHR used the mid-point of the titration-derived interval as the person's certainty equivalent, even though any value within this interval is consistent with the data generating process (DGP). In other words, RKHR used an estimate as data, without taking into account the uncertainty of this estimate.

outlined earlier, this statistical approach is not valid but RKHR report that they found evidence of risk aversion in all groups ($\gamma > 1$) and they found that "triers" were more risk averse than smokers ($p < 0.05$) and never-smokers ($p < 0.05$); there were no significant differences between smokers and never-smokers.

This two-step approach to analysis (i.e., using NLLS, or some other technique, to estimate risk preference parameters and then using the point estimates as data in subsequent statistical models) is remarkably common in this literature. Harrison, Lau and Rutström (2010) (HLR) is the only study in Table I which does not use this method, and for good reason. The problem with the two-step approach, other than that it often uses tiny samples to estimate risk preference parameters at the level of the individual, is that estimated risk preference parameters are estimates, and not data. Such estimates comprise both a point estimate (of the mean) and a standard error, and to use only the point estimate is to throw away information on the uncertainty of that estimate.[11] Moreover, using an estimated risk preference parameter as data violates the statistical assumptions of the second-stage models: specifically, that the covariates are measured without error. Thus, the statistical inferences drawn from this approach are simply not valid. HLR estimate risk preference parameters as a linear function of observable characteristics (e.g., age, gender, and smoking status) so that the uncertainty of the risk preference parameter estimates propagates into the inferences which are drawn from the data. This valid statistical approach will be used in this chapter.

Table I collates the results from the other studies. A clear majority of the studies (8 out of 11) were conducted in the US, with only one study a piece taking place in Japan, Denmark, and South Africa. An important feature of these studies is that only 3 use student subject pools while the rest recruit from the community at large; diverse samples help to bolster the external validity of the results so it is unfortunate that the statistical analyses in every study except HLR hinder meaningful inferences.

[11] This problem is compounded when certainty equivalents are computed by taking an average of the interval within which a person's certainty equivalent lies (i.e., taking the average of an interval derived by a titration mechanism). In this case, there are two levels of uncertainty (i.e., uncertainty about the certainty equivalents and uncertainty about the parameter estimates) which are ignored when the final point estimate of a risk preference parameter is used as data.

The majority of studies on risk preferences and smoking behaviour have small sample sizes: the first 7 studies listed in Table I recruited less than 60 people. Fortunately, since 2008, 4 relatively large studies have taken place: Anderson and Mellor (2008) (AM) elicited risk preference data on 79 smokers and 898 non-smokers; HLR recruited 252 subjects; Szrek, Chao, Ramlagan and Peltzer (2012) (SCRP) used a sample of 351 individuals; and Poltavski and Weatherley (2013) recruited 182 people.

With regards to elicitation mechanisms, there is a roughly equal split between titration (6 out of 11 studies) and choice procedures. AM, HLR, and SCRP used an ordered choice elicitation mechanism, originally devised by Miller, Meyer and Lanzetta (1969) and refined by Holt and Laury (2002) (HL), which has been used extensively in the experimental economics literature on choice under risk, and thereby deserves further comment.[12] This elicitation procedure is referred to as a multiple price list (MPL).

In a MPL, subjects are given a table with 10 rows, and on each row they must choose between a "safe" and a "risky" lottery. In Table II, which is adapted from Table I in HL (p. 1645), Option A is the "safe" lottery because the range of the prizes is small (e.g., ($2.00, p; $1.60, $1 - p$)), and Option B is the "risky" lottery because the range of the prizes is large (e.g., ($3.85, p; $0.10, $1 - p$)). On row 1 of the table $p = 0.1$, and as you move down the table p increases by 0.1 on each row, implying that by row 10, $p = 1$. In the last 3 columns of the table I have included the expected value (EV) of Option A, the EV of Option B, and their difference, although this information is not usually provided to subjects.

In row 1 of Table II, the EV of Option A exceeds the EV of Option B but by row 5 the EV of Option B exceeds the EV of Option A. The logic behind this elicitation mechanism is that only a very risk loving subject would choose Option B (the "risky" lottery) on row 1, and only a very risk averse subject would choose Option A (the "safe" lottery) on row 9.[13] A risk neutral subject would switch from choosing Option

[12] Harrison and Rutström (2008, p. 44-61) provide a detailed discussion of different risk preference elicitation mechanisms.

[13] As row 10 involves sure outcomes (i.e., $p = 1$) it is not relevant to risk preferences at all but is a good test of whether subjects understood the experiment because one would expect them to choose the larger sure outcome (e.g., $3.85 from the example) over the smaller sure outcome (e.g., $2.00 from the

A to Option B as the EV difference first changes sign (i.e., on row 5 of the table). Thus, if a subject switches to Option B before row 5 he is risk loving, if he switches to Option B on row 5 he is risk neutral, and if he switches to Option B after row 5 he is risk averse.

TABLE II: LOTTERY CHOICES IN THE HL RISK PREFERENCE EXPERIMENT

	Option A				**Option B**				EV^A	EV^B	Difference
Row	p	$	p	$	p	$	p	$	($)	($)	($)
1	0.1	2.00	0.9	1.60	0.1	3.85	0.9	0.10	1.64	0.48	1.17
2	0.2	2.00	0.8	1.60	0.2	3.85	0.8	0.10	1.68	0.85	0.83
3	0.3	2.00	0.7	1.60	0.3	3.85	0.7	0.10	1.72	1.23	0.50
4	0.4	2.00	0.6	1.60	0.4	3.85	0.6	0.10	1.76	1.60	0.16
5	0.5	2.00	0.5	1.60	0.5	3.85	0.5	0.10	1.80	1.98	-0.18
6	0.6	2.00	0.4	1.60	0.6	3.85	0.4	0.10	1.84	2.35	-0.51
7	0.7	2.00	0.3	1.60	0.7	3.85	0.3	0.10	1.88	2.73	-0.85
8	0.8	2.00	0.2	1.60	0.8	3.85	0.2	0.10	1.92	3.10	-1.18
9	0.9	2.00	0.1	1.60	0.9	3.85	0.1	0.10	1.96	3.48	-1.52
10	1	2.00	0	1.60	1	3.85	0	0.10	2.00	3.85	-1.85

Source: HL (p. 1645)

We can say even more about risk preferences by putting some parametric structure on the subjects' utility functions. Specifically, if we assume that subjects employ a power utility function $U(y) = y^r$, which displays constant relative risk aversion (CRRA), and that they evaluate lotteries according to expected utility (EU) theory, then we can use a subject's pattern of choices on the MPL to define bounds on the risk preference parameter r.[14]

For example, suppose that a subject chose Option A on the first 5 rows of Table II and then switched to Option B on row 6. To calculate the upper bound on r we solve the following equation:

$$0.5(\$2.00)^r + 0.5(\$1.60)^r = 0.5(\$3.85)^r + 0.5(\$0.10)^r \Leftrightarrow r \approx 0.85$$

example). Harrison and Rutström (2009, p. 132) also advocate including a row 0 where the smaller outcome under each lottery (i.e., $1.60 under Option A and $0.10 under Option B) is received with certainty so as to "bracket" the MPL logic. In other words, if subjects can see that they should choose Option A on row 0 and Option B on row 10, then all they need to determine is the row on which they switch.

[14] Under EU theory the shape of a utility function determines attitudes toward risk. Using the power utility function above, $r > 1$ denotes risk loving behaviour, $r = 1$ denotes risk neutral behaviour, and $r < 1$ denotes risk aversion. If $r = 0$, $U(y) = \ln y$, and if $r < 0$, $U(y) = -y^r$, following Wakker (2008). AM, HLR and SCRP use a different parameterisation of the CRRA utility function: $U(y) = y^{(1-r)}/(1-r)$. Under this formulation, $r < 0$ denotes risk loving behaviour, $r = 0$ implies risk neutral behaviour, and $r > 0$ denotes risk aversion; if $r = 1$, $U(y) = \ln y$.

This equation defines the value of r which makes a subject indifferent between the two lotteries on row 5. To calculate the lower bound on r we solve the following equation:

$$0.6(\$2.00)^r + 0.4(\$1.60)^r = 0.6(\$3.85)^r + 0.4(\$0.10)^r \Leftrightarrow r \approx 0.59$$

This equation defines the value of r which makes a subject indifferent between the two lotteries on row 6. Thus, if a subject chooses the Option A lottery on the first 5 rows and then switches to the Option B lottery on row 6, this pattern of choices implies a risk preference parameter r which lies in the open interval (0.59, 0.85). Interval data of this form is analysed appropriately using interval regression methods but AM and SCRP used the mid-point of these intervals as data to compare the risk preferences of smokers and non-smokers. As discussed previously, this approach throws away useful information on the uncertainty of the parameter estimates and violates the statistical assumptions of the second-stage models. Thus, the inferences drawn from these data are not valid statistically.

In contrast to studies of smoking and discounting behaviour, there is a greater proportion of studies using real rewards or probabilistic payment schemes in the literature on smoking and risk preferences. Table I shows that 4 studies (Mitchell (1999), Reynolds et al. (2007), AM, SCRP) used only real rewards, whereas 3 studies used probabilistic payment schemes (RKHR, Reynolds, Richards, Horn and Karraker (2004), HLR).[15] The remaining 4 studies (Ohmura, Takahashi and Kitamura (2005), Reynolds (2006), Yi, Chase and Bickel (2007), Poltavski and Weatherly (2013)) used entirely hypothetical rewards. As discussed in Chapter 3, the use of real rewards or probabilistic payment schemes – coupled with a task that is easily understood, a transparent payment scheme, salient rewards, and an incentive-compatible

[15] Studies employing real rewards typically make use of the random lottery incentive mechanism (RLIM) to determine subject payment. RLIM randomly selects one of a subject's choices on a task and, in a study with real rewards, pays out this choice with certainty. A probabilistic payment scheme also makes use of RLIM but subjects are only given some chance of being paid for the randomly selected choice (i.e., subjects are not paid with certainty). In HLR subjects were given a 1-in-10 chance of being paid for one of their choices. By contrast, RKHR and Reynolds, Richards, Horn and Karraker (2004) paid subjects for 1 choice across 2 different tasks, implying that subjects had roughly a 50% chance of being paid for one of their choices on the risk preference task.

experimental design – promotes the truthful revelation of preferences and, thus, far more credence should be given to the results from these studies than those which employ hypothetical rewards.

A majority of the studies in Table I (8 out of 11) adopted the PD approach[16] to risk preferences, which defines risk aversion solely in terms of the shape of the PWF.[17] As subjective probability distortions drive risk preferences in the PD framework, it is surprising that 6 out of these 8 studies only used 5 probabilities in the elicitation task; the remaining two studies (Mitchell (1999) and Yi, Chase and Bickel (2007)) only used 6 and 7 probabilities, respectively. AM, HLR, and SCRP assumed that EU theory characterises choice under risk, so risk preferences are determined solely by the shape of the utility function. All of these studies use a MPL, which has 10 probabilities, and they assumed a CRRA utility function: specifically, $U(y) = y^{(1-r)}/(1-r)$. In this chapter, I allow risk preferences to be determined both by the shape of the utility function and the shape of the PWF so as to provide a bridge between prior studies in the literature. In addition, this allows me to explore whether smokers and non-smokers differ in the shape of their utility functions, the shape of their PWFs, or both.

The final column of Table I shows whether the studies found a significant statistical relationship between risk preferences and smoking behaviour: the results are equivocal and, other than HLR, the statistical analyses are not valid. To interpret this column, note that a positive relationship between smoking and risk preferences means that smokers are more risk averse than non-smokers whereas a negative relationship means that smokers are less risk averse than non-smokers. Null results were reported in 3 studies, positive results were reported in 5 studies, and negative results were

[16] Chapter 4 showed that the PD model is just Yaari's (1987) dual theory of choice under risk limited to a circumscribed class of lotteries and with a specific PWF: $\pi(p) = p / [p + \gamma(1 - p)]$; if $\gamma > 1$ this represents probability pessimism and risk aversion. Recall that to derive this form for the PWF, Rachlin, Raineri and Cross (1991) severed the purported link between probability and delay which they argued was so central to the way people interpret probabilities. Thus, the specific form of this PWF lacks consistent theoretical foundations.

[17] Of these studies, 3 also employed the area under the curve (AUC) method of Myerson, Green and Warusawitharana (2001). When using the AUC method, one calculates the area under a subject's derived certainty equivalents and normalizes this to lie in the closed unit interval. Larger AUCs imply less risk aversion and, thus, the AUCs of smokers and non-smokers can be compared to determine whether the groups differ in their risk preferences.

reported in 3 studies.[18] These conflicting results cut across different elicitation mechanisms, real and hypothetical rewards, different frameworks for choice under risk, and different methods of analysis. Thus, Table I shows that the relationship between risk preferences and smoking behaviour, or lack thereof, differs markedly across studies.

In this chapter I will add to the extant literature by investigating the smoking-risk preference relationship using an incentive-compatible experimental design, a relatively large sample of South African university students, and a statistical framework which allows one to draw robust inferences about smokers and non-smokers. Rather than define risk preferences solely in terms of utility function curvature or probability weighting, I allow both sources to affect attitudes toward risk.[19] In addition, I use choices over risky prospects to determine the shape of subjects' utility functions which I then estimate jointly with models of discounting behaviour to characterise time preferences over time-dated utility flows, not flows of money. This will allow me to draw strong conclusions about the relationship between discounting and smoking behaviour.

III. EXPERIMENTAL DESIGN AND SUMMARY STATISTICS

The study of risk and time preferences and smoking behaviour recruited 175 subjects from undergraduate classes at UCT. Given the focus on smoking behaviour, sign-up sheets included a simple screening question asking for the potential participant's smoking status. A large number of people applied to take part in the study and individuals from the smoking and non-smoking groups were randomly selected for inclusion in the project. Those that were selected were added to a website which

[18] Some of the studies classified smokers using more than one category (e.g., heavy smokers and light smokers in Poltavski and Weatherly (2013), and smokers and "triers" in RKHR), and HLR separated male and female smokers and non-smokers. I adopt the classification scheme from Chapter 3 which codes a study as having found a significant result if at least one smoker, non-smoker comparison was statistically significant, even if all comparisons were not.

[19] Several researchers (e.g., Harrison, Humphrey and Verschoor (2009), Harrison and Rutström (2009), Tanaka, Camerer and Nguyen (2010), Andersen, Harrison, Lau and Rutström (2014)) have conducted experimental studies where both utility function curvature and probability weighting affect attitudes toward risk. To my knowledge, this is the first study of smoking behaviour which incorporates both sources of risk preferences.

allowed them to sign up for an experimental session that worked with their academic timetable.

The experiments took place in a computer lab at UCT which had been set up to run the risk and time preference software developed by Todd Swarthout at Georgia State University. Subjects were separated by partitions and were not allowed to talk to each other during the session. Experiments were conducted in August 2012 and 10 sessions took place in total. The median group size was 17 participants and the author assumed the role of experimenter for every session; two research assistants (RAs) were also employed to help administer subject payments and answer questions.

Upon arrival at the lab, subjects were randomly allocated to computer terminals and given an overview of the tasks that they would complete. Subjects then signed informed consent before being taken through a detailed presentation of the risk or time preference task[20]; the order of these tasks was counter-balanced across sessions so subjects either performed the risk or time preference task first. Participants were given the opportunity to ask questions at any stage of the presentations or during the tasks. After questions had been addressed, subjects completed the first task.

Once all participants had completed the first task, the experimenter went through a detailed presentation of the other task. Subjects then completed this task before filling out a questionnaire which collected standard demographic characteristics and information on smoking behaviour. The experimenter or RAs then determined their earnings for the tasks. All subjects received a show-up fee of R20. Earnings for the risk preference task were paid out immediately in cash and earnings for the time preference task were paid out on the date corresponding to the subject's choice on the randomly selected discounting question. Delayed payments were done via electronic transfer and subjects received a payment notification on their cell phones as soon as the transfer took place. Experimental sessions lasted approximately an hour and

[20] The introductory presentation used in this study is included in Appendix E, the risk preference task presentation is included in Appendix F, and the time preference task presentation is included in Appendix G. The presentations, and explanations therein, were carefully developed to make the tasks transparent and easy to understand. The payment system is also discussed in detail so that subjects understand how their final earnings are determined. This attention to detail, coupled with salient rewards, promotes incentive compatibility and the truthful revelation of preferences.

subjects earned R370 ($66 at PPP at the time) on average. The risk and time preference tasks will be discussed in detail below.

Risk Preference Task

The risk preference task was based on the seminal contribution of Hey and Orme (1994). It presented subjects with a choice between two lotteries on each screen; these lotteries were displayed as pie charts with accompanying text that listed the probabilities of the prizes. Figure I shows a screenshot of the risk preference task.

The task used prize magnitudes between R0 and R280 ($0 - $50 at PPP at the time) and probabilities which varied in increments of 0.05 between 0 and 1. Thus, other than HLR, this study used larger lottery prizes than any of the studies in Table I which have incentive-compatible experimental designs. In addition, this study had more variation in the probability domain than every other study in Table I.

Figure I

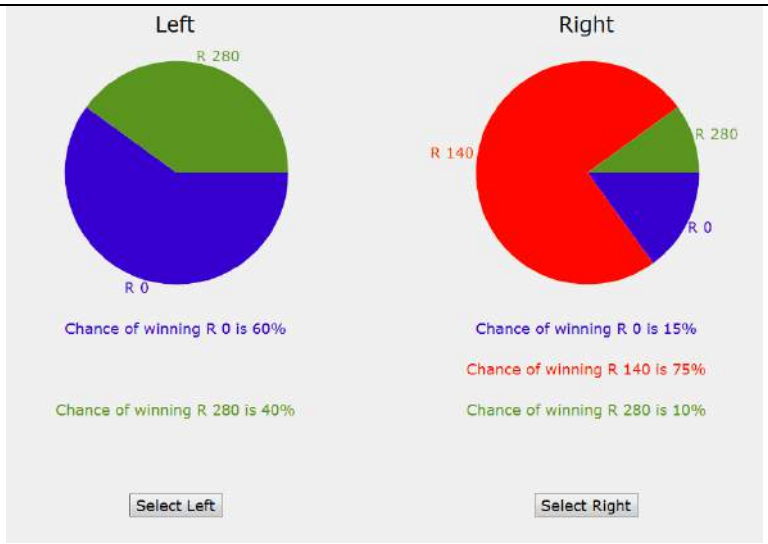

The lottery pairs in the task were based on the set developed by Loomes and Sugden (1998) (LS) to test different stochastic specifications of choice under risk. LS designed the lottery pairs to accommodate a wide range of risk preferences, to provide good coverage of the probability space, and to generate common ratio tests of EU

theory. However, all the lotteries over which each subject made choices had the same context (i.e., the same set of prizes).[21] By contrast, I used four prize contexts in my experiments: (R0, R140, R280), (R40, R80, R240), (R20, R100, R220), and (R60, R120, R180). Incorporating a number of different prizes and probabilities is helpful for the separate identification of the utility function and the PWF in models which admit both sources of risk preferences (e.g., rank-dependent utility theory).

Figure II: Marschak-Machina triangles of lotteries in the risk preference task

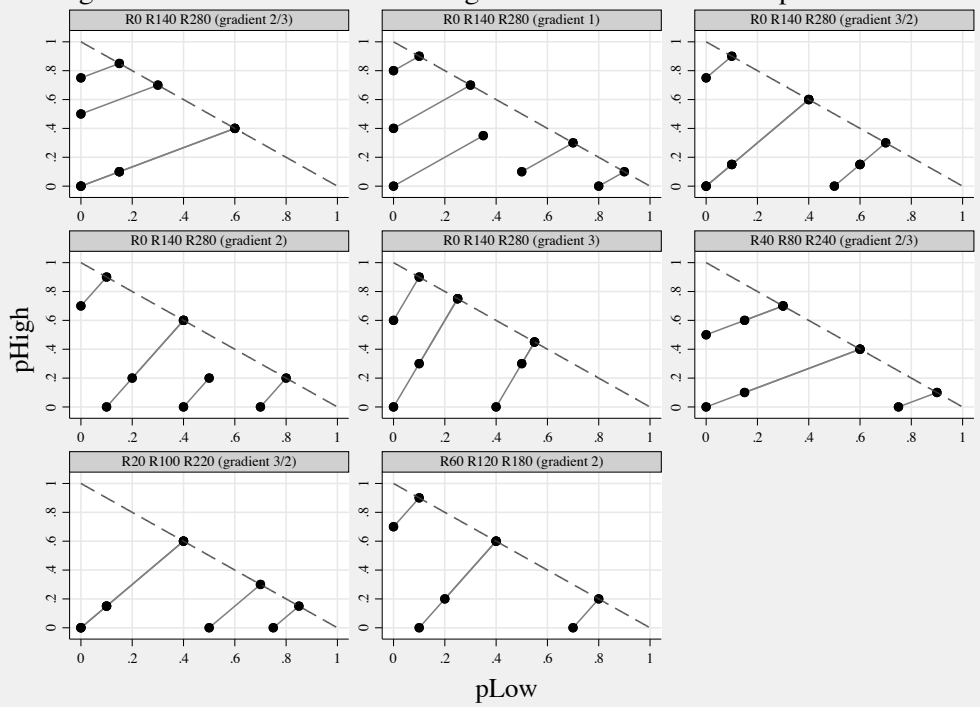

Figure II shows the set of Marschak-Machina (MM) triangles representing the lotteries, and lottery pairs, which were used in the risk preference task. At the top of each diagram I list the context of the lotteries (e.g., (R0, R140, R280)) and the gradient of the lines connecting lottery pairs. In a MM triangle the vertical axis represents the probability of the highest prize in a lottery (e.g., R280) and the horizontal axis represents the probability of the lowest prize in a lottery (e.g., R0). The probability of the intermediate prize (e.g., R140) can be determined by subtracting the sum of high and low prize probabilities from 1. Each point in the MM

[21] LS used two experimental treatments: one where subjects made choices over lotteries defined on the context ($0, $10, $20) and one where subjects made choices over lotteries defined on the context ($0, $10, $30). The probability distributions over these contexts were identical across the two groups except for 8 out of the 45 lotteries in the task.

triangle represents a lottery and the line connecting two, or more, points represents a lottery pair, or set of lottery pairs, on offer in the task. The figure shows that the risk preference task provided excellent coverage of the MM triangle and that it captures the full range of risk preferences: risk-loving (gradients less than 1); risk neutral (gradients equal to 1); and risk averse (gradients greater than 1).

Subjects made 40 choices in the risk preference task and one choice was selected at random at the end of the experimental session for payment.

Time Preference Task

The time preference task presented subjects with choices between smaller, sooner (SS) and larger, later (LL) monetary rewards. Figure III shows a screenshot of the time preference task. On each screen subjects had to make 4 choices before proceeding to the next screen. The principal (i.e., SS reward) and time horizon were fixed on each screen but varied across screens. A calendar was displayed on every screen so as to show the subjects when they would receive the amounts of money they chose.

Figure III

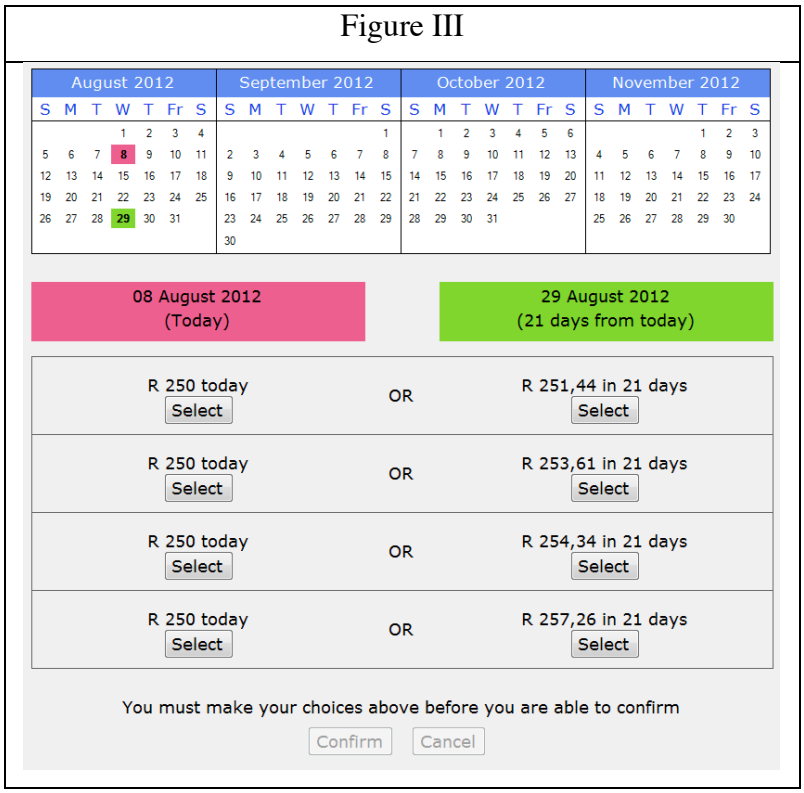

Following Coller and Williams (1999), three front end delays (FEDs) to the SS rewards were used: zero days, 7 days, and 14 days. As discussed in Chapter 3, this design allows one to hold subjective transaction costs constant for the SS and LL rewards at positive FEDs. It also facilitates estimation of the parameters of a quasi-hyperbolic discounting function because the zero day FED allows one to pin down the estimate of β, which captures a "passion for the present" or "present-bias" in decision making, whereas the positive FEDs allow one to estimate the long-term discounting parameter δ.[22] Subjects in an experimental session were only exposed to one of these FED treatments.

Two principals (R150 and R250: $27 and $45 at PPP at the time), 14 time horizons between the SS and LL rewards (7 – 98 days), and nominal annual interest rates between 5% and 250% were used in the time preference task. These parameters define a battery of 224 possible choice pairs. Each subject made 60 choices in the task which were drawn randomly, without replacement, from this battery. At the end of the experimental session, one of these choices was randomly selected for payment.

Summary Statistics

Table III presents summary statistics for the sample of 175 students. The average age in the sample is approximately 20 years old, 42% of the sample is white[23], two-thirds of the sample is enrolled in the Commerce faculty at UCT, and approximately one-third of the sample receives financial aid. Current smokers make up 62% of the sample[24] and this is the largest number of smokers (i.e., 108 smokers) ever recruited for a study exploring risk preferences and smoking behaviour.[25] They were deliberately oversampled to investigate whether intensity of smoking is related to risk

[22] Coller, Harrison and Rutström (2012) explain that to estimate the β parameter in the quasi-hyperbolic discounting function one must use an experimental design which incorporates a zero day FED because this allows one to identify a subject's passion for the present. In designs which do not incorporate a zero day FED, one must assume that a subject's *present*-bias persists beyond the *present* to estimate β; this assumption is highly questionable.

[23] Approximately 24% of the sample is black, 14% is Coloured ("Coloured" is an officially designated population group in South Africa which refers to individuals of mixed-race origin), and 17% is Asian/Indian. The remaining 3% preferred not to classify their race.

[24] The remaining 38% of the sample comprises both former-smokers and never-smokers who will be referred to collectively as non-smokers.

[25] Smokers were defined as those people who answered "yes" to the question: "Do you currently smoke cigarettes?"

and time preferences. The mean number of cigarettes smoked per day is 8.67 with a standard deviation of 5.81 and a range of 1 - 25.

Smokers also completed the Fagerström Test for Nicotine Dependence (FTND) (Heatherton, Kozlowski, Frecker and Fagerström (1991)) which is a measure of smoking severity that scores people on a scale of 0 - 10, where higher numbers indicate greater severity. The average FTND score among smokers is 2.22 with a standard deviation of 2.08. Thus, on average, the smokers in this sample are relatively mild smokers. In addition, given the young age of the sample, the smokers' lifetime exposure to cigarettes is relatively low. In the literature on risk and time preferences and smoking behaviour, researchers usually try to maximise the difference between smokers and non-smokers by selecting heavy smokers to take part in the study. I decided to recruit smokers across the entire spectrum of severity to determine whether being a smoker, irrespective of intensity, is associated with risk and time preferences. This also allows me to explore the relationship, if any, between smoking severity and risk and time preferences.

TABLE III
SUMMARY STATISTICS

Variable	Mean	Std Deviation
Demographics		
Age	19.789	1.815
White	0.417	0.495
Male	0.549	0.499
Commerce faculty	0.674	0.470
Financial aid	0.314	0.466
Smoke	0.617	0.487
Treatments		
Risk task first	0.514	0.501
FED: 0 days	0.343	0.475
FED: 1 week	0.326	0.469
FED: 2 weeks	0.331	0.471
High Principal	0.498	0.500

Table III shows that randomisation across experimental treatments ensured that approximately 50% of the sample completed the risk preference task prior to the time preference task. FED treatments were split evenly across the sample and 50% of choices in the time preference task involved the high principal of R250.

IV. STATISTICAL SPECIFICTION

The statistical method I employ is direct estimation by maximum likelihood of a structural model of a latent choice process; the explanation in this section closely follows AHLR. The latent choice process in question is captured by models of risk and time preferences. These models provide the structure necessary to estimate people's risk and time preferences using their choice data. One of the major benefits of the maximum likelihood approach is that it uses all of the available information which the participants' data impart to estimate discounting and risk preference parameters and the precision of these estimates. I review the basic logic of the estimation strategy below, focussing on the canonical cases of EU theory and exponential (E) discounting. I will then briefly discuss the straightforward extension to other risk and time preference models.

Assume that utility of income is defined by a power utility function which displays CRRA:

$$U(y) = y^r, \tag{1}$$

where y is a lottery prize in the risk preference task and r is a parameter to be estimated. If $r = 0$, $U(y) = \ln y$, and if $r < 0$, $U(y) = -y^r$, following Wakker (2008). We know that under EU theory risk preferences are determined by the shape of the utility function so with the power utility function parameterisation, $r > 1$ yields a convex utility function and risk loving behaviour, $r = 1$ implies a linear utility function and risk neutrality, and $r < 1$ yields a concave utility function and risk aversion.

Let there be three possible outcomes in a lottery, just like the risk preference task reported in this chapter. Under EU theory the probabilities for each outcome y_j, $p(y_j)$, are those that are used in the experimental task, so expected utility is simply the probability-weighted utility of each outcome in each lottery i:

$$\mathrm{EU}_i = \sum_{j=1,2,3} [\, p(y_j) \times U(y_j) \,] \tag{2}$$

To determine the value of r, the EU for each lottery pair (i.e., the Left and Right lotteries in Figure I) is calculated for a candidate estimate of r and the index below is formed:

$$\nabla EU = EU_R - EU_L \tag{3}$$

Note that this is a latent index, based on latent preferences, which captures the difference in EU of the Right and Left lotteries presented to subjects. This index is then linked to the subjects' observed choices using the cumulative normal distribution function $\Phi(\nabla EU)$. This function takes any argument (∇EU) between $\pm\infty$ and transforms it into a number between 0 and 1. Thus, we have the so-called "probit" link function:

$$\Pr(\text{Choose lottery R}) = \Phi(\nabla EU) \tag{4}$$

The latent index in (3) is linked to subjects' observed choices by specifying that lottery R is chosen when $\Phi(\nabla EU) > ½$, which is precisely what (4) tells us. To see this note that if $EU_R = EU_L$ then $\nabla EU = 0$. Plugging 0 into the cumulative normal distribution function (i.e., $\Phi(0)$) yields a value of 0.5. In other words, when the EU of the lotteries L and R are equal, the probability of choosing lottery R is 0.5 (i.e., there is an equal chance of choosing lottery L or R). By contrast, if $\nabla EU > 0$ then $\Phi(\nabla EU) > 0.5$ (i.e., the probability of selecting lottery R is greater than 0.5 when its EU exceeds the EU of lottery L) and if $\nabla EU < 0$ then $\Phi(\nabla EU) < 0.5$ (i.e., the probability of selecting lottery R is less than 0.5 when its EU is less than the EU of lottery L).

Thus, the likelihood of the observed responses, conditional on the power utility and EU model being true, depends on the estimates of r given the statistical model above and the choices of subjects in the risk preference task. The conditional log-likelihood for the risk preference responses is:

$$\ln L_i^{RP}(r; z, X) = \sum_i [(\ln \Phi(\nabla EU) \times I(z_i = 1)) + (\ln (1 - \Phi(\nabla EU)) \times I(z_i = 0))], \tag{5}$$

where I(·) is the indicator function, z_i = 1(0) denotes the choice of the R (L) lottery in choice pair i, and X is a vector of individual characteristics capturing age, gender, education etc.

One of the advantages of structural maximum likelihood estimation is that it is a straightforward extension to make the parameter of interest, the risk preference parameter r, a linear function of individual characteristics. In this case, one estimates $r = r_0 + r_\beta \times X$, where r_0 is a fixed parameter and r_β is a coefficient vector linked to the variable vector X of individual characteristics. If no individual characteristics are included in the model we estimate $r = r_0$, which is the risk preference parameter estimated at the level of the sample without taking into account observed, individual heterogeneity (i.e., assuming homogenous preferences). Note that every estimate of r includes a standard error which reflects our ignorance and uncertainty as to the "true" value of r. This stands in stark contrast to the bulk of studies in Table I which use risk preference estimates as data in subsequent statistical models.

Another important extension to the simple model defined above is to allow for some errors on the part of subjects when they make choices between lotteries L and R. This error could be as simple as a "tremble," where, say, a subject wants to choose lottery R but mistakenly selects lottery L instead. I adopt the "contextual utility" (CU) behavioural error specification of Wilcox (2011) to allow mistakes on the part of subjects from the perspective of the deterministic EU model and to draw robust inferences about the primitive "stochastically more risk averse than" relation.[26] The CU specification normalises the ∇EU index so that it falls within the closed unit interval [0, 1] and it incorporates the behavioural error term originally due to Fechner (1966). Thus, rather than adopt the simple ∇EU index in (3), I make use of the index below:

$$\nabla EU = [(EU_R - EU_L) / \lambda] / \mu, \qquad (6)$$

where λ is the normalising term and μ is the Fechner error term. Note how different values of μ affect our ∇EU index. As $\mu \rightarrow 0$ our specification collapses to a

[26] The "stochastically more risk averse than" relation is the stochastic choice counterpart to the "more risk averse than" relation (see Pratt (1964)) which is defined for the deterministic EU model.

deterministic choice model where the choice is strictly determined by the EU of the two lotteries. However, as $\mu \rightarrow \infty$, $\nabla EU \rightarrow 0$, and a subject's choice is essentially random (i.e., the probability of selecting either lottery is 0.5). When $\mu = 1$ we are back to specification (3), so the Fechner error term is a parameter which basically flattens the probit link function as its value increases. The new conditional log-likelihood is:

$$\ln L_i^{RP}(r, \mu; z, X) = \sum_i [(\ln \Phi(\nabla EU) \times I(z_i = 1)) + (\ln (1 - \Phi(\nabla EU)) \times I(z_i = 0))] \quad (7)$$

The expression in (7) can be maximised using standard numerical methods to estimate the power function parameter r, which defines risk preferences under EU theory, and the Fechner error term μ which determines the extent to which choices involve errors on the part of subjects.

Fortunately it is a simple matter to incorporate other theories of choice under risk in this statistical framework. Quiggin (1982) developed the rank-dependent utility (RDU) model which assumes that a decision maker transforms objective probabilities into subjective decision weights which are then used to evaluate lotteries. According to this theory, risk preferences are determined both by the shape of the utility function, like EU theory, and the shape of the PWF, like the PD approach[27] to risk attitudes. Under RDU we replace (2) with:

$$RDU_i = \sum_{j=1,\dots,n} [\, w(y_j) \times U(y_j) \,], \quad (8)$$

where

$$w_j = \pi(p_j + \dots + p_n) - \pi(p_{j+1} + \dots + p_n), \quad (9)$$

for $j = 1, \dots, n\text{-}1$, and

[27] Chapter 4 showed that the PD model is just Yaari's (1987) dual theory of choice under risk limited to a circumscribed class of lotteries and with a specific PWF: $\pi(p) = p / [p + \gamma(1 - p)]$; if $\gamma > 1$ this represents probability pessimism and risk aversion. Recall that to derive this form for the PWF, Rachlin, Raineri and Cross (1991) severed the purported link between probability and delay which they argued was so central to the way people interpret probabilities. Thus, the specific form of this PWF lacks consistent theoretical foundations.

$$w_j = \pi(p_j), \tag{10}$$

for $j = n$. Note that the subscript j represents outcomes ranked from worst to best, and $\pi(p)$ is a specific PWF. A number of different PWFs have been used in the literature and Stott (2006) provides a useful review. Tversky and Kahneman (1992) (TK) popularised the following PWF:

$$\pi(p) = p^{\gamma} / [\, p^{\gamma} + (1-p)^{\gamma} \,]^{1/\gamma}, \tag{11}$$

for $1 > p > 0$. This function permits linear, "inverse S-shaped" and "S-shaped" forms. Gonzalez and Wu (1999) review the empirical evidence on this function and find that $1 > \gamma > 0$ in most studies. This gives the function an inverse S-shape with overweighting of low probabilities up to a crossover point where $\pi(p) = p$, and then underweighting of moderate to high probabilities.[28] In this chapter I will estimate the TK PWF, amongst others, to see whether I replicate this inverse S-shaped result in this sample.

To estimate a RDU model, assuming power utility, the TK PWF, and the CU behavioural error specification, one simply forms the RDU index $\nabla RDU = [(RDU_R - RDU_L)/\lambda]/\mu$ and then links this to the subjects' observed choices using the cumulative normal distribution function. This defines the conditional log-likelihood for the model which is then used to estimate r, μ, and γ, where γ is the parameter defining the TK PWF. In this chapter I will estimate EU and RDU models to compare the risk preferences of smokers and non-smokers. In addition, I will estimate the parameters of a variety of PWFs to ensure that the results are robust across different specifications.

Shifting the focus to time preferences, recall from Chapter 3 that under the E model, δ is the discounting parameter which equalises the *utility* of income received at time t with the *utility* of income received at time $t + \tau$:

$$[1 / (1 + \delta)^{t}]U(y_t) = [1 / (1 + \delta)^{t+\tau}]U(y_{t+\tau}), \tag{12}$$

[28] However, Ingersoll (2008) shows that this function is not monotonic at very small values of γ.

for some utility function $U(\cdot)$. Under the assumptions that EU characterises choices over risky prospects and that subjects employ the power utility function, we can add more structure to this indifference condition. Specifically, (12) becomes:

$$[1 / (1 + \delta)^{t}](y_t)^{r} = [1 / (1 + \delta)^{t+\tau}](y_{t+\tau})^{r}, \tag{13}$$

where the general form of the utility function $U(\cdot)$ in (12) has been replaced with the specific power utility function $U(y) = y^{r}$ in (13). Note that the left hand side (LHS) of (13) represents the present value (PV) of the *utility* of the SS reward in the time preference task whereas the right hand side (RHS) of (13) represents the present value of the *utility* of the LL reward. Thus,

$$PV_{SS} = [1 / (1 + \delta)^{t}](y_t)^{r}, \tag{14}$$

and

$$PV_{LL} = [1 / (1 + \delta)^{t+\tau}](y_{t+\tau})^{r} \tag{15}$$

To estimate the parameters of our time preference model, conditional on EU theory, power utility, and the E model, we form the latent index below:

$$\nabla PV = (PV_{SS} - PV_{LL}) / \nu, \tag{16}$$

where ν is a Fechner error term for the time preference task, just as μ was the behavioural error term for the risk preference task. We could force $\mu = \nu$ but there is little sense in doing so if we think that one task may be more cognitively challenging than the other, and hence more prone to subject error. To remain open to this possibility, I allow μ and ν to vary independently.[29] The latent index (16) captures the difference in the present values of the utility of the SS and LL rewards. It is linked to

[29] My prior is that the risk preference task, which incorporated up to three prizes in each lottery and a host of different probabilities, is more cognitively challenging than the time preference task, where subjects simply had to make choices between two rewards available at different points in time.

subjects' observed choices using the cumulative normal distribution function $\Phi(\nabla PV)$. This defines our probit link function:

$$\Pr(\text{Choose SS reward}) = \Phi(\nabla PV) \tag{17}$$

The latent index in (16) is linked to subjects' observed choices by specifying that the SS reward is chosen when $\Phi(\nabla PV) > ½$, which is exactly what (17) tells us.

Thus, the likelihood of the observed time preference responses, conditional on the EU, power utility, and E models being true, depends on the estimates of r, δ, and ν, given the statistical model above. The conditional log-likelihood is:

$$\begin{aligned} \ln L_i^{TP}(r, \delta, \nu; z, X) = \sum_i [&(\ln \Phi(\nabla PV) \times I(z_i = 1)) \\ &+ (\ln (1 - \Phi(\nabla PV)) \times I(z_i = 0))] \end{aligned} \tag{18}$$

The joint likelihood of the risk and time preference responses can then be formed as:

$$\ln L_i(r, \delta, \mu, \nu; z, X) = \ln L_i^{RP} + \ln L_i^{TP} \tag{19}$$

This "joint estimation" approach uses subjects' choices in the risk preference task to pin down the parameters of the utility function, and subjects' choices in the time preference task to pin down the parameters of the E discounting model, conditional on the shape of the utility function. This approach ensures, therefore, that we estimate time preferences defined over utility flows, and not flows of money.

As explained in Chapter 3, it is straightforward to incorporate other discounting models in this statistical framework. In the case of Weibull (WB) discounting, (13) becomes:

$$[\exp(-\delta t^{(1/\beta)})](y_t)^r = [\exp(-\delta(t+\tau)^{(1/\beta)})](y_{t+\tau})^r \tag{20}$$

(14) and (15) are adjusted appropriately to incorporate this new expression and then one forms the latent index in (16) and proceeds as before.

V. RESULTS

In this section I present the results from a set of risk and time preference models so as to explore potential differences in the risk and time preferences of smokers and non-smokers. I begin with the risk preference results because they provide a natural segue to the time preference results which are conditional on the utility function curvature identified by the risk preference task.

A. Risk Preferences: EU Theory

Table IV presents baseline estimates of an EU model employing a power utility function and the CU behavioural error specification. Note that the results pool choices across all individuals, which means I am estimating the value of r_0 for the sample as a whole. In other words, I am assuming homogenous preferences. Note further that the results account for clustering at the individual level which adjusts the standard errors of the estimates to take into account the fact that each respondent made multiple choices across the 40 risk preference questions.

The estimate of $r = 0.306$ implies a relatively high level of risk aversion in the sample. The estimate of $\mu = 0.175$ is positive and statistically significant, implying that subjects make behavioural errors in the risk preference task and it is important, therefore, to take this into account so as to draw accurate inferences concerning the "stochastically more risk averse than" relation.

TABLE IV: EXPECTED UTILITY THEORY ML ESTIMATES
HOMOGENOUS PREFERENCES

	Model
	CU error
Power function parameter (r)	0.306***
	(0.028)
Error (μ)	0.175***
	(0.009)
N	7000
log-likelihood	-4198.932

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

To analyse the link between risk preferences and smoking behaviour, one can make the parameter of interest *r* a linear function of smoking status. This captures the "total effect" of smoking status on risk preferences without controlling for any potential differences between smokers and non-smokers like age, education, and gender. In other words, this model provides a first pass at exploring this relationship. The point estimate of the "smoker" variable in this model is -0.027 with a standard error of 0.048 which means there is not a statistically significant relationship between risk preferences and smoking status in this sample.[30]

Table V presents the results from a model that takes into account observed, individual heterogeneity by conditioning the power function parameter estimate on a set of covariates and task parameters. Specifically, the model includes the demographic variables from Table III and a variable specifying whether the risk preference task preceded the time preference task. This model captures the marginal effect of smoking status on risk preferences while controlling for other factors which may mediate this relationship.

TABLE V: EXPECTED UTILITY THEORY ML ESTIMATES
HETEROGENOUS PREFERENCES

	Model	
	Estimate	Std Error
Power function parameter (r)		
Age	0.005	0.012
White	0.045	0.05
Male	0.119**	0.047
Commerce faculty	0.081	0.055
Financial aid	-0.033	0.055
Risk task first	-0.031	0.046
Smoker	-0.036	0.058
Constant	0.123	0.223
Error (μ)		
Constant	0.173***	0.009
N	7000	
log-likelihood	-4180.528	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table V shows that the only variable which is significantly, individually related to risk preferences in this sample is gender: men are less risk averse than women. Thus,

[30] I also estimated a model which allows risk preferences to vary as a quadratic function of smoking intensity as measured by the average number of cigarettes smoked per day: risk preferences were not significantly related to smoking intensity.

estimates from the EU model with a power utility function and CU error specification, point to no statistically significant differences in the risk preferences of smokers and non-smokers.

To explore the possibility that the power utility function is too restrictive to accurately characterise choice under risk in this sample, I also estimated the expo-power (EP) utility function of Saha (1993) which admits increasing relative risk aversion (IRRA), decreasing relative risk aversion (DRRA), and CRRA. The EP utility function takes the following form:

$$U(y) = \theta - \exp(-\alpha y^{r}), \tag{21}$$

where $\theta > 1$ and $\alpha r > 0$.[31] Note that the parameter θ is simply included to make utility positive and monotonically increasing, and has no effect on risk preferences.[32] This is most easily seen by focussing on the coefficient of relative risk aversion R which for the EP utility function is: $R(y) = 1 - r + r\alpha(y^{r})$. Note that this function nests CRRA as $\alpha \rightarrow 0$.

Thus, the parameters of interest in the EP utility function are α and r because these determine whether relative risk aversion is increasing, decreasing, or constant relative to income. Baseline estimates (i.e., assuming homogenous preferences) of the EP utility function indicate IRRA in the sample: $\alpha = 0.002$ is significantly greater than 0 ($p = 0.020$) and $r = 1.373$ is significantly greater than 1 ($p < 0.001$). Figure IV shows the estimates of relative risk aversion, with a 95% confidence interval, for the range of prizes in the risk preference task; dashed lines in the figure represent the prizes in the task. Thus, it appears that the assumption of a power utility function, which displays CRRA, is too restrictive.

[31] Holt and Laury (2002) adopted a different specification of the EP utility function: $U(y) = [(1 - \exp(-\alpha y^{1-r})) / \alpha]$.
[32] Saha (1993, p. 908-909) actually sets $\theta = 0$ in his numerical analyses.

Figure IV: Relative risk aversion

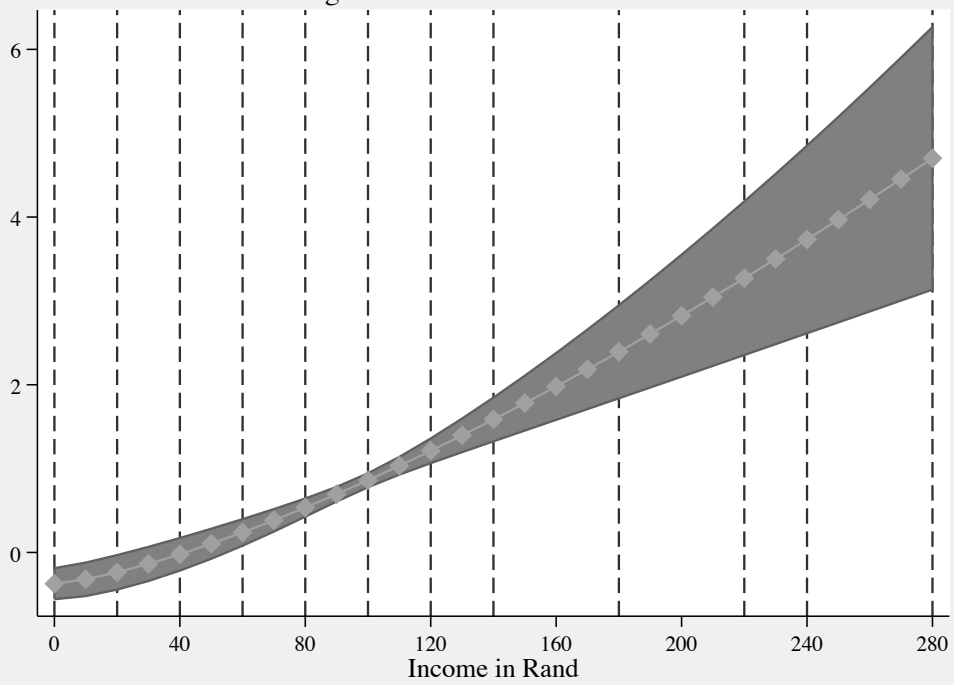

However, when I allow the parameters of the EP utility function to vary according to demographic characteristics and task parameters, a different picture emerges. Table VI presents the results and shows that, for the expo parameter α, none of the covariates, nor the constant term, is significantly different to zero. In addition, a test of the joint hypothesis that all of the covariates, including the constant term, are equal to zero, cannot be rejected ($p = 0.833$).[33] Thus, while there is some evidence of IRRA in the sample, it does not vary according to demographic characteristics and, specifically, smoking status. Consequently, the power utility function will be used in subsequent analyses because it adequately characterises choice under risk in this sample.

[33] Harrison, Lau and Rutström (2007, p. 358) used this approach to determine whether CRRA held over the range of prizes used in their experiments. They too found that a test of the joint hypothesis that all of the covariates, and the constant term, are equal to zero, could not be rejected, which lead them to conclude that CRRA was an appropriate characterization for their sample.

TABLE VI: EXPECTED UTILITY THEORY ML ESTIMATES
EP UTILITY FUNCTION AND HETEROGENOUS PREFERENCES

	Model	
	Estimate	Std Error
Power parameter (r)		
Age	0.004	0.021
White	0.214**	0.107
Male	-0.014	0.045
Commerce faculty	0.117	0.075
Financial aid	0.025	0.031
Risk task first	0.054	0.037
Smoker	0.020	0.055
Constant	1.078**	0.420
Expo parameter (α)		
Age	0.000	0.000
White	-0.002	0.001
Male	0.000	0.000
Commerce faculty	-0.001	0.001
Financial aid	0.000	0.000
Risk task first	0.000	0.000
Smoker	0.000	0.000
Constant	0.004	0.004
Error (μ)		
Constant	0.141***	0.008
N	7000	
log-likelihood	-4110.578	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

B. Risk Preferences: RDU Theory

The EU results in the previous subsection suggest that there are no significant differences in the risk preferences of smokers and non-smokers. However, this analysis, by assumption, ignored the role of probability weighting and it may be the case that smokers perceive probabilities differently to non-smokers. To explore this possibility, I estimate RDU models in this subsection.

One of the key components of a RDU model is the specification of the PWF. The TK PWF was presented in (11) and two other commonly used PWFs are: the power function, and the Prelec (1998) function. The power PWF is just like the power utility function except that prizes are replaced with probabilities. Specifically, the power PWF takes the following form:

$$\pi(p) = p^{\gamma} \tag{22}$$

An important feature of the power PWF is that it is concave, convex, or linear throughout its range. This means that interior probabilities are either viewed objectively (i.e., linear weighting), always overweighted, or always underweighted. Thus, the power PWF does not permit the inverse S-shaped or S-shaped forms of the TK PWF.

Prelec (1998) axiomatically derived a two-parameter PWF which exhibits considerable flexibility. The functional form for this PWF is:

$$\pi(p) = \exp[\, -\eta(-\ln p)^{\gamma}\,], \tag{23}$$

which is defined for $1 > p > 0$, $\eta > 0$, and $\gamma > 0$.[34,35] This function allows independent specification of location and curvature in probability weighting. Note that it nests the power PWF when $\gamma = 1$, and it nests a one-parameter function when $\eta = 1$, which is very similar to the TK function in that it admits linear, inverse S-shaped, and S-shaped forms.

Table VII presents baseline estimates of RDU models employing the power utility function, the CU behavioural error specification, and the three PWFs discussed above.[36] In Model 1, the power PWF parameter $\gamma = 0.953$, implying slight overweighting of all probabilities. However, this estimate is not significantly different to 1 ($p = 0.301$) so we cannot rule out a linear PWF where probabilities are viewed objectively.

In Model 2, the TK PWF parameter $\gamma = 0.868$ which yields an inverse S-shaped function implying overweighting of low probabilities and underweighting of moderate to high probabilities. The estimate of γ is significantly less than 1 at any regular level

[34] Prelec (1998, proposition 1, part C, p. 503) provides these parameter restrictions. Prelec (1998, proposition 1, part B, p. 503) constrains $1 > \gamma > 0$, but this constraint can be quite restrictive in practice because it ensures that the PWF is either linear, S-shaped or inverse-S shaped.

[35] Note that I impose these constraints when estimating the models using nonlinear transformations of the parameters.

[36] As mentioned previously, I impose the PWF constraints, using nonlinear transformations of the parameters, when estimating the models. To recover the core parameters I use the inverse of these nonlinear transformations, and then apply the "delta method" to derive standard errors and *p*-values for the estimates (see Oehlert (1992)).

of significance ($p < 0.001$)[37] and this has a marked effect on the estimate of the power function parameter r. To see this, compare the estimate of r under the EU model in Table IV (i.e., $r = 0.306$) to the estimate of r in Model 2 of Table VII (i.e., $r = 0.351$). A Wald test that the estimate of r under the EU model is equal to the estimate of r under the RDU model with the TK function, is easily rejected ($p < 0.001$), implying that the power utility function is more concave in the EU model.

TABLE VII: RANK-DEPENDENT UTILITY THEORY ML ESTIMATES
HOMOGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**
	Power	TK	Prelec
Power function parameter (r)	0.283***	0.351***	0.324***
	(0.024)	(0.032)	(0.026)
PWF parameter (γ)	0.953***	0.868***	0.797***
	(0.045)	(0.022)	(0.025)
PWF parameter (η)			0.882***
			(0.033)
Error (μ)	0.176***	0.170***	0.169***
	(0.009)	(0.009)	(0.008)
N	7000	7000	7000
log-likelihood	-4197.975	-4177.421	-4151.295

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Under RDU, both probability weighting and utility function curvature influence risk attitudes. The RDU model with the TK PWF shows that the EU estimate of r is biased downward, yielding a more concave utility function. This is an important issue because when jointly estimating utility function curvature and discounting behaviour, the extent of utility function curvature identified by the risk preference task propagates into estimates of discounting parameters. Thus, if one ignores probability weighting when it is present, this may lead to biased estimates of utility function curvature and, hence, biased estimates of discounting parameters. In effect, when probability weighting is present, one wants to apportion risk preferences into their concave utility and probability weighting components so that accurate inferences about discounting behaviour can be drawn.

[37] The presence of inverse S-shaped probability weighting explains why the estimate of γ is not significantly different to 1 in the model with the power PWF: the power PWF is "confused" because it has to be linear, concave, or convex throughout its range.

The Prelec estimates in Table VII replicate the inverse S-shaped PWF that I find with the TK function. The estimates of $\gamma = 0.797$ and $\eta = 0.882$ are significantly less than 1 ($p < 0.001$ in both cases) and the estimate of $r = 0.324$ is greater than in the EU model of Table IV, although I cannot reject the hypothesis that the estimate of r under the EU model is the same as the estimate of r under the RDU model with the Prelec PWF ($p = 0.358$). Thus, the estimates in Table VII show that probability weighting plays a role in the determination of risk attitudes in this sample. This will need to be taken into account when adopting the joint estimation approach to discounting behaviour.

Figure V: Estimated probability weighting functions and implied decision weights
Based on equi-probable reference lotteries

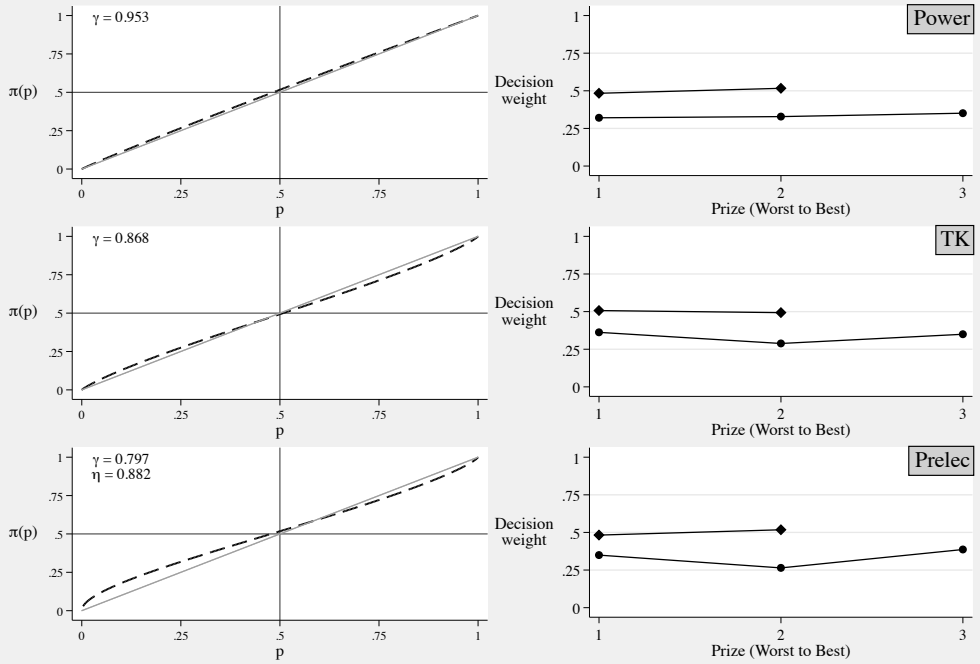

Figure V plots the PWFs, and implied decision weights, for the estimates in Table VII. The decision weights are graphed for equi-probable reference lotteries and show the decision weight applied to the worst outcome, the decision weight applied to the intermediate outcome, and the decision weight applied to the best outcome in a lottery. Focussing on the implied decision weights for a two outcome equi-probable (i.e., 50:50) reference lottery, the power PWF assigns a decision weight of 0.48 to the worst outcome and a decision weight of 0.52 to the best outcome in a lottery. The Prelec PWF also assigns a decision weight of 0.48 to the worst outcome and a decision weight of 0.52 to the best outcome in a lottery. Thus, the power and Prelec

PWFs imply that the probability of the worst outcome is underweighted (0.48) and the probability of the best outcome is overweighted (0.52), relative to the objective probabilities of 0.5. By contrast, the TK PWF assigns a decision weight of 0.51 to the worst outcome and a decision weight of 0.49 to the best outcome in a lottery. In this case, the worst outcome is overweighted and the best outcome is underweighted relative to the objective probabilities of 0.5.

Focussing on the implied decision weights for a three outcome equi-probable reference lottery, the power PWF assigns a decision weight of 0.32 to the worst outcome, a decision weight of 0.33 to the intermediate outcome, and a decision weight of 0.35 to the best outcome. Thus, the power PWF implies that the worst outcome is slightly underweighted, the intermediate outcome is viewed objectively, and the best outcome is overweighted, relative to the objective probabilities of 0.33. By contrast, the TK PWF assigns a decision weight of 0.36 to the worst outcome, a decision weight of 0.29 to the intermediate outcome, and a decision weight of 0.35 to the best outcome in a lottery. Similarly, the Prelec PWF assigns a decision weight of 0.35 to the worst outcome, a decision weight of 0.26 to the intermediate outcome, and a decision weight of 0.39 to the best outcome in a lottery. Thus, the TK and Prelec PWFs overweight the best and worst outcomes, and underweight the intermediate outcome, of a lottery, relative to the objective probabilities of 0.33.

To investigate the possibility that smokers perceive probabilities differently to non-smokers, even if their utility functions do not differ, I estimated the two models in Table VII which admit inverse S-shaped PWFs and allowed the parameters to vary as a function of observable characteristics and task parameters. Results are presented in Table VIII.[38]

[38] As mentioned previously, the experimental design of the risk preference task lends itself to common ratio tests of EU theory. To complement the analyses in this section, I conducted a set of common ratio tests for the lotteries represented in the MM triangles in Figure II to determine whether smokers were more or less EU-consistent than non-smokers. I adopted the non-parametric Cochran Q test for this purpose and found that both smokers and non-smokers violated EU theory in every MM triangle in Figure II ($p < 0.001$ in every test) except the MM triangle with a gradient of 3. In this latter MM triangle, I could not reject the hypothesis that non-smokers satisfy EU theory ($p = 0.111$) but I could reject this hypothesis for smokers ($p = 0.027$). Thus, in only 1 of the 8 MM triangles of Figure II were non-smokers more EU-consistent than smokers. The bulk of the evidence, therefore, suggests little difference in the extent to which smokers and non-smokers violate EU theory; one reaches the same conclusion from the estimates in Table VIII.

TABLE VIII: RANK-DEPENDENT UTILITY THEORY ML ESTIMATES
HETEROGENOUS PREFERENCES

	Model 1 TK		**Model 2** Prelec	
	Estimate	Std Error	Estimate	Std Error
Power function parameter (r)				
Age	0.005	0.013	-0.004	0.011
White	0.038	0.060	0.029	0.051
Male	0.114**	0.055	0.062	0.049
Commerce faculty	0.113*	0.060	0.030	0.062
Financial aid	-0.057	0.065	-0.051	0.058
Risk task first	-0.057	0.055	-0.015	0.050
Smoker	-0.048	0.068	-0.005	0.055
Constant	0.179	0.246	0.366	0.230
PWF parameter (γ)				
Age	-0.002	0.011	-0.003	0.006
White	0.021	0.054	0.001	0.047
Male	0.016	0.050	-0.009	0.044
Commerce faculty	-0.083	0.057	-0.084	0.120
Financial aid	0.061	0.059	0.034	0.056
Risk task first	0.055	0.051	0.054	0.080
Smoker	0.026	0.056	0.028	0.049
Constant	0.876***	0.228	0.871***	0.206
PWF parameter (η)				
Age			-0.027	0.046
White			-0.062	0.121
Male			-0.166	0.137
Commerce faculty			-0.216	0.184
Financial aid			-0.014	0.139
Risk task first			0.166	0.153
Smoker			0.146	0.153
Constant			1.425**	0.676
Error (μ)				
Constant	0.168***	0.008	0.166***	0.008
N	7000		7000	
log-likelihood	-4153.594		-4119.762	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In both of the models in Table VIII, smokers do not differ significantly to non-smokers in the shape of their utility functions (i.e., in the estimate of r) nor in the way they perceive probabilities (i.e., in the estimates of γ and η). In addition, tests of the joint hypothesis that the coefficients for smokers across r, γ, and η are equal to zero, cannot be rejected under either model ($p = 0.771$ for the TK model and $p = 0.823$ for the Prelec model).[39] Thus, at least in this sample, there are no significant differences in the risk preferences of smokers and non-smokers. This result is robust to different

[39] I also estimated a RDU model with the expo-power utility function, the CU error specification, the Prelec PWF, and the full set of covariates from Table VIII. The smoker variable was not significantly different to zero for any of the parameters in the model. In addition, a test of the joint hypothesis that the coefficients for smokers across r, α, γ, and η are equal to zero, cannot be rejected ($p = 0.967$).

theories of choice under risk, different PWFs, and a utility function that admits DRRA, IRRA, and CRRA.

C. Time Preferences: Baseline Results

In this subsection I will estimate the four time preference models introduced in Chapter 3 and discussed extensively in Andersen, Harrison, Lau and Rutström (2014): the E model, Mazur's (1984) hyperbolic (H) model, the quasi-hyperbolic (QH) model, and the Weibull (WB) model. Unlike Chapter 3 though, I will jointly estimate the parameters of these models with the curvature of the utility function to focus on the discounting of utility flows, not flows of money. In addition, I will allow both EU and RDU to characterise choices over risky prospects. This will ensure that the results are robust to different assumptions about how people evaluate lotteries and will determine whether assuming EU, rather than RDU, biases the estimates of discounting parameters.

Recall from Chapter 3 that the discount factor for the E model is:

$$D^{E}(t) = 1 / (1 + \delta)^{t}, \tag{24}$$

for $t \geq 0$. By contrast, the discount factor for the QH model is:

$$D^{QH}(t) = 1 \qquad \text{if } t = 0 \tag{25a}$$

$$D^{QH}(t) = \beta / (1 + \delta)^{t} \qquad \text{if } t > 0 \tag{25b}$$

Note that if $\beta = 1$, the QH specification collapses to the E model whereas if $\beta < 1$ discounting is quasi-hyperbolic. Recall that I use β and δ to represent the parameters in all of the discounting models even though there is nothing which implies that they should be equal across the different specifications; this choice was made for notational simplicity.

Mazur's (1984) H discounting function has a discount factor:

$$D^{H}(t) = 1 / (1 + \delta t) \quad (26)$$

Finally, the WB model has a discount factor:

$$D^{WB}(t) = \exp(-\delta t^{(1/\beta)}), \quad (27)$$

for $\delta > 0$ and $\beta > 0$. Note that when $\beta = 1$, (27) collapses to the E specification so, in the terminology of Jamison and Jamison (2011, p. 25), the parameter β either "expands" or "contracts" time. When $\beta > 1$, it is as if time has contracted or is perceived to be "slowing down" by the individual. By contrast, when $\beta < 1$, it is as if time has expanded or is "speeding up" as perceived by the individual.

TABLE IX: DISCOUNTING FUNCTION ML ESTIMATES
LINEAR UTILITY AND HOMOGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Discounting parameter (δ)	3.234***	1.715***	2.833***	0.890***
	(0.287)	(0.096)	(0.271)	(0.066)
Discounting parameter (β)			0.962***	1.518***
			(0.013)	(0.107)
Error (ν)	24.272***	24.043***	23.669***	23.573***
	(1.774)	(1.742)	(1.626)	(1.591)
N	10500	10500	10500	10500
log-likelihood	-5419.508	-5335.484	-5352.777	-5233.649

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table IX presents results from the four discounting models employing the Fechner error term, assuming linear utility, and using years, rather days, as the unit of measurement for the estimation of the parameters[40]; this table is included for comparative purposes. Note, from Chapter 3, that if no individual characteristics are included in the model we estimate $\delta = \delta_0$ and $\beta = \beta_0$, which are the discounting parameters estimated at the level of the sample without taking into account observed, individual heterogeneity (i.e., assuming homogenous preferences). Note further that the results account for clustering at the individual level which adjusts the standard

[40] All of the discounting results in this chapter refer to annual, rather than daily, time periods.

errors of the estimates to take into account the fact that each respondent made multiple choices across the 60 time preference questions.

Under the assumption of linear utility, estimated discount rates are huge and differ markedly across the different specifications. In the E model, the estimate of $\delta = 3.234$ implies an annual discount rate in excess of 320%. In Table X, by contrast, where I have estimated the discounting models jointly with the curvature of the utility function, discount rates are far lower, and more similar. In the E model in Table X, the estimate of $\delta = 0.507$ implies an annual discount rate of 50%. Similar declines are evident in the other discounting models in Table X, which highlights the point, now familiar from AHLR, that the assumption of linear utility biases estimates of δ upwards. Thus, to draw accurate inferences about discounting behaviour it is crucial to estimate time preferences jointly with the curvature of the utility function.

TABLE X: DISCOUNTING FUNCTION ML ESTIMATES
CONCAVE UTILITY AND HOMOGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Power function parameter (r)	0.283***	0.309***	0.273***	0.260***
	(0.032)	(0.028)	(0.034)	(0.037)
Discounting parameter (δ)	0.507***	0.472***	0.441***	0.223***
	(0.070)	(0.049)	(0.065)	(0.032)
Discounting parameter (β)			0.987***	1.608***
			(0.004)	(0.114)
Risk error (μ)	0.183***	0.174***	0.187***	0.192***
	(0.011)	(0.009)	(0.012)	(0.014)
Time error (ν)	0.159***	0.198***	0.145***	0.128***
	(0.047)	(0.050)	(0.045)	(0.044)
N	17500	17500	17500	17500
log-likelihood	-9519.026	-9488.92	-9430.8	-9282.495

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

With regard to utility curvature, note that in a joint estimation framework, the parameter r is estimated *jointly* with the parameters of the discounting models and, thus, its value inevitably varies across the different specifications. However, the point estimates of r fall within a narrow range, i.e., from 0.260 in the WB model to 0.309 in the H model, and are very similar to the estimate of $r = 0.306$ in Table IV.[41]

[41] Using a 10% level of significance, one-sided Wald tests show that the estimates of r, across the different discounting models, fall within the interval (0.21, 0.35).

In the QH model in Table X, the estimate of $\beta = 0.987$ is significantly less than 1 ($p = 0.003$), which provides evidence of quasi-hyperbolic discounting and declining discount rates. The same is true in the WB results: the estimate of $\beta = 1.608$ is significantly greater than 1 ($p < 0.001$) which means people perceive time as "slowing down" and this generates declining discount rates. Thus, both the QH and WB results suggest that discount rates decline over time, which, when coupled with an additively-separable intertemporal utility function, raises the spectre of time-inconsistent choices. However, the two discounting functions provide competing explanations for this result: a present-bias in the case of the QH model and subjective time perception in the case of the WB model. Note, finally, that the WB model has the highest log-likelihood of all the specifications in Table X, which provides preliminary evidence that it best characterises all of the data.

To estimate the four discounting models under the assumption that RDU characterises choices over risky prospects, a specific PWF had to be selected, lest I present an endless series of results.[42] The analyses in the previous subsection showed that probability weighting takes an inverse S-shaped form which rules out the power PWF as a candidate. Of the remaining PWFs, the Prelec function is easily the most flexible so it will be used to incorporate probability weighting in the estimation of discounting models.[43]

Table XI presents the results from the four discounting models assuming RDU and the Prelec PWF. Despite the presence of probability weighting in choices over risky prospects (i.e., γ and η are significantly less than 1 across all models), estimates of the power function parameter r are very similar in the models in Table X and Table XI. Consequently, the estimates of β and δ do not differ significantly across EU and RDU specifications. Specifically, for the E discounting model, the estimate of δ assuming EU in Table X does not differ significantly to the estimate of δ assuming RDU in Table XI ($p = 0.731$). Similarly, the estimates of δ for the H, QH, and WB models in

[42] Four discounting functions and three PWFs yield a total of 12 models. Note that I have estimated all of these models but the results across the PWFs are very similar so a representative set of results, employing one PWF, will be reported here.

[43] I also attempted to estimate the different discounting models, assuming that either EU or RDU characterises choice under risk, at the individual level. Unfortunately I did not get numerical convergence in enough subjects to be able to provide reliable estimates at this level of analysis.

Table X do not differ significantly to the estimates of δ for the H, QH, and WB models in Table XI (p = 0.322, p = 0.471, p = 0.253, respectively). Finally, the estimates of β in the QH and WB models in Table X do not differ significantly to the estimates of β in the QH and WB models in Table XI (p = 0.500, p = 0.269, respectively). Thus, while probability weighting is certainly a feature of the subjects' choices over risky prospects, it does not significantly affect the estimates of concavity in our discounting models and, thus, it does not significantly affect the estimates of our discounting parameters

TABLE XI: DISCOUNTING FUNCTION ML ESTIMATES
RANK-DEPENDENT UTILITY AND HOMOGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential Prelec	Hyperbolic Prelec	Quasi-Hyperbolic Prelec	Weibull Prelec
Power function parameter (r)	0.277***	0.327***	0.260***	0.238***
	(0.028)	(0.027)	(0.029)	(0.030)
PWF parameter (γ)	0.797***	0.797***	0.796***	0.795***
	(0.025)	(0.025)	(0.025)	(0.026)
PWF parameter (η)	0.838***	0.884***	0.823***	0.804***
	(0.032)	(0.034)	(0.032)	(0.031)
Discounting parameter (δ)	0.493***	0.502***	0.415***	0.204***
	(0.062)	(0.050)	(0.057)	(0.028)
Discounting parameter (β)			0.988***	1.611***
			(0.004)	(0.115)
Risk error (μ)	0.178***	0.169***	0.181***	0.186***
	(0.009)	(0.008)	(0.010)	(0.010)
Time error (ν)	0.151***	0.231***	0.128***	0.104***
	(0.041)	(0.055)	(0.036)	(0.031)
N	17500	17500	17500	17500
log-likelihood	-9471.828	-9441.151	-9383.297	-9234.32

Results account for clustering at the individual level
Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Consequently, on a qualitative level, the results in tables X and XI are identical: incorporating utility function curvature leads to a marked decline in estimates of δ; and there is evidence of declining discount rates in both the QH ($\beta < 1, p = 0.002$) and WB ($\beta > 1, p < 0.001$) models.

While the presence of probability weighting in choices over risky prospects does not significantly affect the discounting parameter estimates, relative to the EU case, I make use of a RDU model and the Prelec PWF in my subsequent analyses to incorporate this feature (i.e., probability weighting) of the data. To highlight the

robustness of the results though, in Appendix H I report discounting estimates which assume EU characterises choice under risk.

Figure VI: Distribution of individual log-ratios

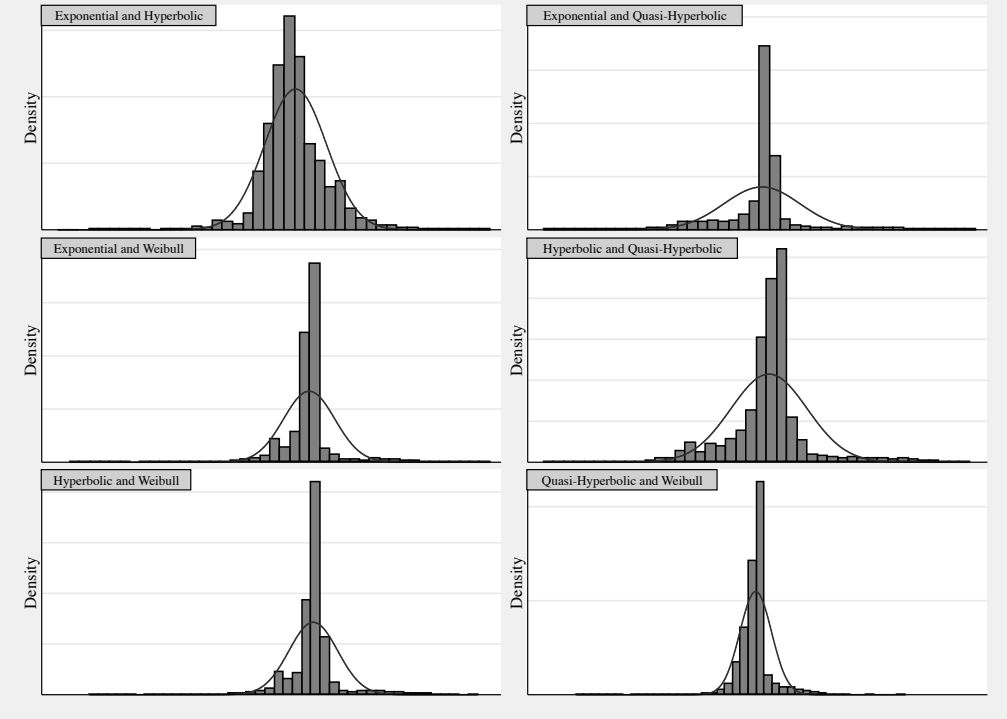

A comparison of the models' log-likelihoods in Table XI suggests that the WB model best characterises the data, followed by the QH model, the H model, and the E model. As discussed in chapters 3 and 4, the hypothesis that one model outperforms another can be tested formally using Vuong (1989) and Clarke (2007) non-nested model selection tests.[44] The choice between these tests is based on the distributions of individual log-ratios of the various models. Figure VI plots the distribution of individual log-ratios, with a normal density overlay, for the six discounting model comparisons. The distribution of these log-ratios is leptokurtic[45] which suggests that the Clarke (2007) test is more appropriate for these data.

[44] The E model is nested in the QH and WB models, but the QH and WB models are not nested in each other. We know from earlier that if $\beta = 1$ in either the QH or WB models, these models collapse to the E model. So an obvious test of whether discounting is, say, E or WB is to test whether $\beta = 1$ in the WB model, as I did earlier. Nevertheless, it is still worthwhile to conduct formal model selection tests even if models are nested in each other, as will become clear later.

[45] As explained in earlier chapters, the normal distribution is the quintessential mesokurtic distribution. A distribution which has positive excess kurtosis (i.e., a highly peaked distribution) is leptokurtic.

The Clarke (2007) test yields a test statistic based on the binomial distribution which must be compared to a critical value to determine which model, in a pairwise comparison, receives the most support in the data. A Clarke (2007) test comparing the E and H models yields a test statistic of 6488, which is below the critical value of 8750, implying that the H model better characterises the data ($p < 0.001$). Thus, in this sample, the H model outperforms the E model. This stands in direct contrast to the finding from Chapter 3 where the E model outperformed the H model and suggests that one should estimate a range of discounting models rather than assume that one model always best characterises the data.

A Clarke (2007) test comparing the E and QH models finds in favour of the E model ($p < 0.001$) whereas a test comparing the E and WB models finds in favour of the WB model ($p < 0.001$). A test of the H and QH models suggests that the H model finds more support in the data ($p < 0.001$) while a test of the H and WB models finds in favour of the WB model ($p < 0.001$). Finally, a Clarke (2007) test of the QH and WB models suggests that the WB model better characterises the data ($p < 0.001$).

The preceding results provide the following transitive ranking of discounting models: WB > H > E > QH. In words, the WB model finds the most support in the data, the QH model finds the least support in the data, and the E and H models are intermediate to these. Thus, if one had to select a single model that best characterises the time preferences of this sample, the WB model would be the obvious choice. However, when multiple time preference processes are present in a dataset, it is preferable to estimate mixture models which allow these different processes to explain discounting choices; this will be taken up in the final subsection.

D. Smoking and Discounting Behaviour

As a prelude to the results in this subsection, Figure VII shows a kernel-weighted local polynomial regression, with a 95% confidence interval, of the fraction of LL choices by smokers and non-smokers for the nominal annual interest rates on offer in the time preference task. At each interest rate, the fraction of LL choices by smokers is less than the fraction of LL choices by non-smokers. This suggests that smokers

discount more heavily than non-smokers, but clearly this result must be subjected to closer scrutiny before any definitive conclusions are reached.[46]

Figure VII: Fraction of LL choices and interest rate offered

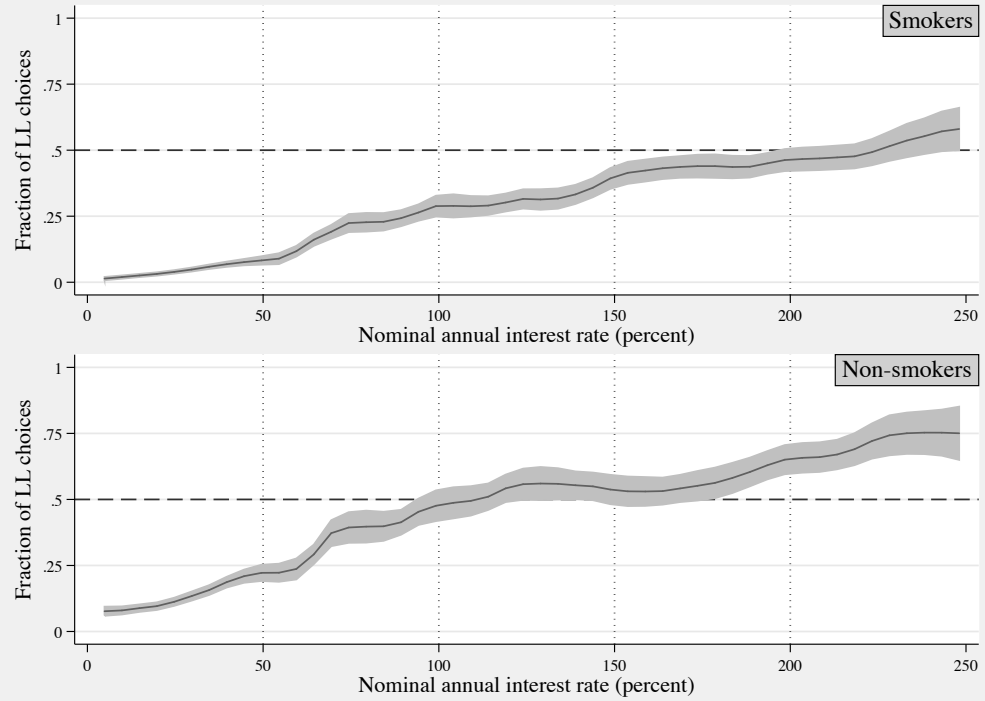

Table XII presents results from the four time preference models, assuming RDU and the Prelec PWF, where risk and discounting parameters are allowed to vary by smoking status. These models, therefore, capture the "total effect" of smoking on discounting behaviour without controlling for any potential differences between smokers and non-smokers like age, gender, etc. Across all specifications, the effect of smoking on the estimate of δ is positive and statistically significant, implying that smokers tend to discount the future more heavily than non-smokers. The magnitude of this difference in discounting behaviour is substantial. In the E model, for example, smokers have an annual discount rate which is 20% higher than non-smokers. Thus, the positive relationship between smoking and discounting identified in Chapter 3 has

[46] A simplistic approach to this endeavour, which can be referred to as "descriptive modelling," is to estimate a probit model of LL choices on the demographic characteristics and task parameters in Table III. Given that subjects made multiple choices in the time preference task, it is imperative to cluster the standard errors of the estimates by subject ID. Estimating this model I find that the coefficient on the "smoker" variable is -0.449 (std error = 0.138) which means smokers are significantly less likely to choose the LL reward than non-smokers. Although suggestive of a relationship between discounting and smoking, a full structural model of discounting behaviour, which takes into account all of the information which the participants' data imparts, is necessary to determine whether the time preferences of smokers and non-smokers differ.

been replicated using a completely different subject pool, different elicitation mechanisms, and a joint estimation approach to time preferences which controls for utility function curvature and probability weighting.[47]

TABLE XII: DISCOUNTING FUNCTION ML ESTIMATES
RANK-DEPENDENT UTILITY AND HETEROGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Power function parameter (r)				
Smoker	0.020	0.023	0.019	0.018
	(0.017)	(0.018)	(0.016)	(0.015)
Constant	0.259***	0.310***	0.245***	0.224***
	(0.028)	(0.027)	(0.030)	(0.030)
PWF parameter (γ)				
Smoker	0.041	0.042	0.040	0.039
	(0.055)	(0.052)	(0.055)	(0.056)
Constant	0.774***	0.774***	0.773***	0.772***
	(0.041)	(0.039)	(0.041)	(0.042)
PWF parameter (η)				
Smoker	0.099	0.105	0.096	0.094
	(0.079)	(0.079)	(0.079)	(0.078)
Constant	0.776***	0.820***	0.764***	0.747***
	(0.049)	(0.049)	(0.049)	(0.049)
Discounting parameter (δ)				
Smoker	0.200***	0.178***	0.176***	0.081**
	(0.060)	(0.051)	(0.059)	(0.032)
Constant	0.359***	0.385***	0.306***	0.156***
	(0.053)	(0.047)	(0.052)	(0.029)
Discounting parameter (β)				
Smoker			0.001	-0.149
			(0.007)	(0.268)
Constant			0.989***	1.670***
			(0.006)	(0.247)
Risk error (μ)				
Constant	0.179***	0.169***	0.182***	0.187***
	(0.009)	(0.008)	(0.010)	(0.010)
Time error (ν)				
Constant	0.136***	0.213***	0.118***	0.097***
	(0.035)	(0.048)	(0.032)	(0.028)
N	17500	17500	17500	17500
log-likelihood	-9338.422	-9306.809	-9263.355	-9118.667

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The estimates of β in the QH and WB models, by contrast, do not vary according to smoking status. Thus, smokers are no more present-biased than non-smokers in the QH model nor are they more likely to perceive time as slowing down in the WB

[47] Appendix H presents results from the four time preference models where EU is assumed to characterise choice under risk: the results are virtually identical to the models in Table XII.

model. It is only the long-term discount rate δ which differs between smokers and non-smokers in these models.

TABLE XIII: DISCOUNTING FUNCTION ML ESTIMATES
RANK-DEPENDENT UTILITY AND HETEROGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Power function parameter (r)				
Number of cigarettes	0.001	0.002	0.001	0.002
	(0.003)	(0.004)	(0.003)	(0.003)
(Number of cigarettes)^2	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	0.262***	0.312***	0.251***	0.219***
	(0.031)	(0.027)	(0.031)	(0.034)
PWF parameter (γ)				
Number of cigarettes	-0.012	-0.011	-0.012	-0.013
	(0.012)	(0.011)	(0.012)	(0.012)
(Number of cigarettes)^2	0.001	0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.809***	0.807***	0.809***	0.810***
	(0.037)	(0.036)	(0.038)	(0.039)
PWF parameter (η)				
Number of cigarettes	0.002	0.003	0.001	0.002
	(0.018)	(0.018)	(0.018)	(0.018)
(Number of cigarettes)^2	0.001	0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.782***	0.826***	0.773***	0.747***
	(0.047)	(0.047)	(0.047)	(0.047)
Discounting parameter (δ)				
Number of cigarettes	0.041***	0.036***	0.038***	0.026***
	(0.013)	(0.011)	(0.012)	(0.008)
(Number of cigarettes)^2	-0.002***	-0.002***	-0.002***	-0.001***
	(0.001)	(0.001)	(0.001)	(0.000)
Constant	0.372***	0.396***	0.317***	0.137***
	(0.053)	(0.046)	(0.050)	(0.027)
Discounting parameter (β)				
Number of cigarettes			0.001	-0.107*
			(0.001)	(0.056)
(Number of cigarettes)^2			0.000	0.004*
			(0.000)	(0.002)
Constant			0.987***	1.948***
			(0.005)	(0.261)
Risk error (μ)				
Constant	0.179***	0.170***	0.182***	0.188***
	(0.009)	(0.008)	(0.010)	(0.011)
Time error (v)				
Constant	0.129***	0.199***	0.116***	0.088***
	(0.035)	(0.046)	(0.033)	(0.028)
N	17500	17500	17500	17500
log-likelihood	-9337.392	-9308.739	-9264.074	-9106.204

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

To investigate whether smoking intensity and discounting behaviour are related, I estimated the four time preference models and allowed the parameters of interest to vary as a quadratic function of number of cigarettes smoked per day; results are presented in Table XIII. In all models, both the linear and quadratic terms are statistically significant in the estimate of δ: the linear term is positive and significant whereas the quadratic term is negative and significant. Thus, there is a concave relationship between discounting behaviour and number of cigarettes smoked per day: every additional cigarette leads to an increase in discounting, but at a decreasing rate until a maximum is reached, after which every additional cigarette leads to a decrease in discounting.

Table XIV maps out the response surface for estimates of δ in the four time preference models evaluated at different values of number of cigarettes smoked per day.[48] At low values of number of cigarettes, the conditional marginal effect of additional cigarettes is positive. By 15 cigarettes though, the conditional marginal effect of additional cigarettes is negative. Thus, Table XIV highlights the nonlinear effect of smoking intensity on discounting behaviour. To my knowledge, this is the first study of time preferences and smoking behaviour which has identified this effect.

TABLE XIV: NUMBER OF CIGARETTES CONDITIONAL MARGINAL EFFECTS FOR δ

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Number of cigarettes				
0	0.041 (0.013)	0.036 (0.011)	0.038 (0.012)	0.026 (0.008)
5	0.023 (0.007)	0.020 (0.006)	0.022 (0.007)	0.016 (0.005)
10	0.005 (0.004)	0.004 (0.003)	0.005 (0.003)	0.006 (0.002)
15	-0.013 (0.006)	-0.013 (0.006)	-0.011 (0.005)	-0.004 (0.002)
20	-0.031 (0.011)	-0.029 (0.011)	-0.028 (0.009)	-0.014 (0.004)
25	-0.049 (0.017)	-0.045 (0.016)	-0.044 (0.014)	-0.025 (0.007)

Standard errors in parentheses

[48] I map out the response surface by calculating the conditional marginal effect of "number of cigarettes" on estimates of δ at different levels of cigarette consumption per day. Note that the conditional marginal effect evaluated at zero cigarettes in Table XIV is identical to the point estimate of the "number of cigarettes" linear term in Table XIII. In other words, the estimate of the linear term of δ in Table XIII is the marginal effect of an infinitesimal increase in cigarette consumption around zero cigarettes, and not the effect evaluated at the average number of cigarettes. To see this, note that $\partial\delta/\partial$numberofcigarettes = $B_1 + 2B_2$(numberofcigarettes), where B_1 and B_2 represent the coefficient estimates of the linear and quadratic terms, respectively. When numberofcigarettes = 0, $\partial\delta/\partial$numberofcigarettes = B_1 and this is the coefficient estimate of the linear term of δ in Table XIII. At positive levels of cigarette consumption, both the linear and quadratic terms affect the estimate of δ. For example, at 5 cigarettes per day in the E model in Table XIV, an infinitesimal increase in number of cigarettes smoked increases the estimate of δ by 0.023 ($p = 0.002$).

In the QH model in Table XIII, smoking intensity is not significantly related to the extent of present-bias (i.e, the estimate of β). In the WB model, by contrast, the number of cigarettes' linear term is negative and significant in the estimate of β, albeit at the 10% level. Thus, the more cigarettes smoked per day, the less likely people are to perceive time as slowing down.

The preceding results are only preliminary because we need to control for a number of factors which may mediate the relationship between smoking and discounting behaviour. Table XV:A presents the results from the E and H models, and Table XV:B presents the results from the QH and WB models, that take into account observed, individual heterogeneity by conditioning the discounting and risk preference parameter estimates on the set of covariates and task parameters from Table III.[49]

In the E model in Table XV:A, smokers discount significantly more than non-smokers. There is also evidence of a magnitude effect in that the higher principal in the time preference task (i.e., R250 as opposed to R150) is associated with significantly lower discounting.[50] Finally, men tend to have higher discount rates than women, although this result is only significant at the 10% level.

[49] In Appendix I I estimate the marginal effect of smoking status across the four discounting specifications using the demographic characteristics and task parameters from Table III but under the assumption of linear utility. On a qualitative level, the results across linear and concave utility specifications are very similar, except in the WB and, to a lesser extent, the QH models. For example, in the WB model I find that the smoker variable is not statistically significant in the equations for β and δ. Recall that the WB model better characterises the data in this chapter than the other discounting specifications. Thus, if one adopted a "naïve" approach to data analysis, which entailed selecting the one function which best characterises the data and then using it for multivariate analysis, one could conclude that smoking and discounting behaviour are not related. I do not advocate this approach but the example highlights the fact that one can reach importantly different conclusions about the relationship between time preferences and smoking without incorporating the concavity of the utility function when estimating discounting parameters.

[50] Andersen, Harrison, Lau and Rutström (2013) provide a comprehensive review of studies investigating the magnitude effect. They find evidence of a magnitude effect, albeit slight, in a study of risk and time preferences conducted on a representative sample of the Danish population in 2009. The size of the magnitude effect identified in this chapter is far larger than that identified by Andersen, Harrison, Lau and Rutström (2013).

TABLE XV:A: DISCOUNTING FUNCTION ML ESTIMATES
RANK-DEPENDENT UTILITY AND HETEROGENOUS PREFERENCES

	Model 1 Exponential		**Model 2** Hyperbolic	
	Estimate	Std error	Estimate	Std error
Power function parameter (r)				
Age	-0.008	0.006	-0.009	0.006
White	-0.016	0.020	-0.020	0.022
Male	-0.019	0.017	-0.019	0.018
Commerce faculty	0.008	0.020	0.008	0.022
Financial aid	0.043**	0.019	0.046**	0.021
Risk task first	0.006	0.018	0.009	0.019
Smoker	0.060***	0.022	0.067***	0.024
Constant	0.443***	0.125	0.528***	0.131
PWF parameter (γ)				
Age	-0.003	0.013	-0.003	0.012
White	-0.019	0.062	-0.020	0.059
Male	-0.006	0.056	-0.003	0.053
Commerce faculty	-0.108	0.071	-0.108	0.068
Financial aid	0.014	0.059	0.005	0.056
Risk task first	0.090*	0.052	0.083*	0.050
Smoker	0.022	0.061	0.021	0.058
Constant	0.890***	0.275	0.882***	0.265
PWF parameter (η)				
Age	-0.022	0.019	-0.023	0.020
White	-0.125	0.087	-0.134	0.089
Male	-0.218***	0.081	-0.216***	0.082
Commerce faculty	-0.202**	0.102	-0.209**	0.104
Financial aid	0.061	0.086	0.059	0.087
Risk task first	0.130*	0.078	0.133*	0.080
Smoker	0.158*	0.096	0.168*	0.098
Constant	1.458***	0.389	1.550***	0.398
Discounting parameter (δ)				
Age	-0.003	0.016	-0.001	0.015
White	-0.101	0.074	-0.095	0.064
Male	0.132**	0.064	0.124**	0.052
Commerce faculty	0.032	0.078	0.018	0.068
Financial aid	0.121	0.075	0.095	0.063
Risk task first	0.024	0.066	0.031	0.058
FED: 1 week	0.059	0.071	0.058	0.062
FED: 2 weeks	-0.004	0.072	0.004	0.063
High Principal	-0.208***	0.036	-0.191***	0.026
Smoker	0.260***	0.070	0.220***	0.062
Constant	0.512	0.320	0.490*	0.289
Risk error (μ)				
Constant	0.169***	0.008	0.159***	0.007
Time error (ν)				
Constant	0.196***	0.048	0.338***	0.070
N	17500		17500	
log-likelihood	-9076.945		-9024.286	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Somewhat surprisingly, I now find that smoking is significantly related to the concavity of the utility function: smokers have less concave utility functions than

non-smokers, and this holds across all of the discounting specifications. This result highlights the importance of a joint estimation approach to risk and time preferences because while smoking status does not affect risk preferences in and of themselves, there are clearly feedback effects from discounting to risk preferences that the models capture.

The pattern of results in the H model is the same as the E model. Specifically, smokers discount more than non-smokers, men discount more than women, and the high principal in the time preference task is associated with significantly lower discounting. In addition, in both the E and H models, the estimate of η is significantly higher among smokers than non-smokers, albeit at the 10% level. Consequently, smokers' PWFs depart less from linearity than non-smokers' PWFs. In other words, smokers tend to overweight and underweight probabilities less than non-smokers.

TABLE XV:B: DISCOUNTING FUNCTION ML ESTIMATES
RANK-DEPENDENT UTILITY AND HETEROGENOUS PREFERENCES

	Model 3 Quasi-Hyperbolic		**Model 4** Weibull	
	Estimate	Std error	Estimate	Std error
Power function parameter (r)				
Age	-0.007	0.005	-0.004	0.005
White	-0.019	0.019	-0.013	0.016
Male	-0.021	0.017	-0.025*	0.015
Commerce faculty	0.007	0.017	0.013	0.016
Financial aid	0.035*	0.019	0.025	0.017
Risk task first	-0.002	0.023	-0.001	0.015
Smoker	0.060***	0.022	0.049**	0.022
Constant	0.414***	0.111	0.339***	0.116
PWF parameter (γ)				
Age	-0.003	0.013	-0.004	0.013
White	-0.019	0.063	-0.018	0.064
Male	-0.006	0.057	-0.006	0.058
Commerce faculty	-0.109	0.072	-0.11	0.073
Financial aid	0.017	0.06	0.023	0.061
Risk task first	0.093*	0.053	0.094*	0.054
Smoker	0.022	0.062	0.024	0.063
Constant	0.892***	0.276	0.898***	0.279
PWF parameter (η)				
Age	-0.021	0.019	-0.018	0.02
White	-0.129	0.087	-0.12	0.087
Male	-0.220***	0.08	-0.225***	0.081
Commerce faculty	-0.203**	0.101	-0.194*	0.1
Financial aid	0.055	0.085	0.048	0.084
Risk task first	0.122	0.08	0.123	0.077
Smoker	0.157	0.097	0.144	0.098
Constant	1.429***	0.387	1.349***	0.403

Table continues on next page

TABLE XV:B: DISCOUNTING FUNCTION ML ESTIMATES (CONTINUED)
RANK-DEPENDENT UTILITY AND HETEROGENOUS PREFERENCES

	Model 3 Quasi-Hyperbolic		**Model 4** Weibull	
	Estimate	Std error	Estimate	Std error
Discounting parameter (δ)				
Age	-0.007	0.013	0.007	0.007
White	-0.088	0.071	-0.049	0.031
Male	0.142**	0.06	0.053***	0.02
Commerce faculty	0.018	0.075	-0.003	0.03
Financial aid	0.103	0.067	0.001	0.026
Risk task first	-0.044	0.065	-0.027	0.028
FED: 1 week	0.344***	0.079	0.139**	0.068
FED: 2 weeks	0.287***	0.066	0.204**	0.088
High Principal	-0.156***	0.03	-0.056***	0.015
Smoker	0.224***	0.073	0.082***	0.028
Constant	0.279	0.266	0.003	0.133
Discounting parameter (β)				
Age	-0.004	0.004	-0.007	0.079
White	-0.002	0.012	0.285	0.306
Male	0.011	0.009	-0.492	0.357
Commerce faculty	0.002	0.012	-0.032	0.238
Financial aid	-0.005	0.012	0.520*	0.284
Risk task first	-0.021*	0.011	0.929**	0.4
FED: 1 week	0.348	0.273	2.840*	1.45
FED: 2 weeks	0.167	0.217	4.155*	2.39
High Principal	0.006**	0.002	0.081	0.107
Smoker	-0.002	0.012	-0.382	0.701
Constant	1.040***	0.067	2.161	1.709
Risk error (μ)				
Constant	0.172***	0.008	0.175***	0.01
Time error (ν)				
Constant	0.195***	0.051	0.130***	0.047
N	17500		17500	
log-likelihood	-8826.297		-8522.212	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In the QH model in Table XV:B, with a full set of covariates, the estimate of δ for smokers is significantly larger than non-smokers, but there are no differences in these groups in terms of present-bias (i.e., the estimate of β). As with δ in the E and H models, men tend to discount more than women, and the high principal is negatively related to the extent of discounting. I find that positive FEDs are associated with significantly more discounting in the estimate of δ; a test of the joint hypothesis that positive FEDs have no effect across the estimates of β and δ, is easily rejected too ($p < 0.001$). My prior was that positive FEDs, which hold subjective transaction costs for the SS and LL rewards constant, would lead to less, rather than more, discounting; clearly this assumption is not borne out by these results. Focussing on the estimates of β, I find that task order has a significant effect on present-bias: subjects who

completed the risk task first tend to make more present-oriented choices than people who completed the time preference task first. Finally, the high principal in the time preference task tends to mitigate the extent of present-bias.

The WB model estimates of δ share the same properties as the QH model. Specifically, smokers discount more than non-smokers, men discount more than women, positive FEDs are linked to greater discounting, and a high principal tends to attenuate the extent of discounting. With regard to β, receiving financial aid, completing the risk task first, and being exposed to positive FEDs, all tend to contract time and thereby produce declining discount rates.

In sum, there is a positive relationship between smoking and discounting behaviour which holds across all of the time preference models estimated in this chapter. This result is also robust to the assumption that EU characterises choice under risk (see Appendix H). However, smokers do not differ from non-smokers with regard to present-bias in the QH model nor in terms of time perception in the WB model.

E. Mixture Models of Discounting Behaviour

The analyses conducted thus far have been based on the implicit assumption that the observations are produced by only one discounting DGP (i.e., E, H, QH, or WB) when more may be present in the data. In other words, the E model may explain some discounting choices better than the H model whereas the H model may explain other choices better than the E model. The assumption that only one DGP characterises all of the data, precludes such a possibility.

As discussed in chapters 3 and 4, finite mixture models[51] allow two or more DGPs to account for the data and also provide a measure of the proportion of the data which is explained by each process. In the current context, one can estimate a mixture model of, say, the E and H discounting functions and then ask the data to determine how much support each function has. To do so one specifies a "grand likelihood" function which is just a probability-weighted average of the likelihoods of the two models.

[51] For detailed discussions of mixture models consult McLachlan and Peel (2000), Harrison and Rutström (2009), and Conte, Hey and Moffatt (2011).

Letting π^E represent the probability that the E model is correct, and $\pi^H = (1 - \pi^E)$ the probability that the H model is correct, the grand likelihood is the probability-weighted average of the two conditional likelihoods L^E and L^H for the E and H models, respectively. Thus, the likelihood for the mixture model is given by:

$$\ln L_i(r, \gamma, \eta, \delta_E, \delta_H, \mu, \nu, \kappa; z, X) = \Sigma_i \ln [(\pi^E \times L^E) + (\pi^H \times L^H)], \quad (28)$$

where κ is a parameter which defines the log odds of the probability of the E model: $\pi^E = 1 / (1 + \exp(\kappa))$. Note that this transformation allows the parameter κ to take on any value during the maximisation process but it constrains the probabilities π^E and π^H to lie within the unit interval. The grand likelihood in (28) is maximised to estimate the parameters of each model and the weight accorded to each model in the data, under the assumptions that RDU and the Prelec PWF characterise choice under risk.

Table XVI presents estimates of the mixture model of the E and H discounting functions assuming homogenous preferences.[52] The estimate of $\pi^E = 0.347$ implies that the E model accounts for approximately 35% of the choices in the data; the H model therefore accounts for roughly 65% of the choices. A hypothesis test that $\pi^E = 0.5$ is easily rejected ($p < 0.001$ as noted in the table) but so too is the hypothesis that $\pi^E = 0$ ($p < 0.001$).[53] Thus, the E and H discounting models both find significant support in the data, even though the H model finds more support. This means that it is a mistake to assume that only one DGP characterises the data.

[52] Appendix J contains the results from all of the two process mixture models that can be estimated from the four discounting specifications used in this chapter. I only present the results from the E and H mixture model in this section because these are the most commonly used discounting functions in the addiction literature and they are representative of the results from the other mixture models.

[53] In all of the mixture models in Appendix J, one of the discounting functions explains significantly more of the choices than the other discounting function. However, all discounting functions find significant support in the data which reinforces the point that it is a mistake to assume only one DGP characterises all discounting choices all of the time. The mixture probability estimates for the other mixture models are: E-QH model - $\pi^E = 0.636$; E-WB model - $\pi^E = 0.406$; H-QH model - $\pi^H = 0.634$; H-WB model - $\pi^H = 0.611$; QH-WB model - $\pi^{QH} = 0.609$.

TABLE XVI: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HOMOGENOUS PREFERENCES

	Estimate	Std Error	*p*-value	95% Confidence Interval	
Rank-dependent utility theory					
Power function parameter (r)	0.336***	0.027	0.000	0.283	0.390
PWF parameter (γ)	0.797***	0.025	0.000	0.749	0.846
PWF parameter (η)	0.893***	0.035	0.000	0.825	0.961
Exponential discounting model					
Discounting parameter (δ_E^{mix})	0.137***	0.017	0.000	0.104	0.169
Mixture probability (π^E)	0.347***	0.034	0.000	0.280	0.414
Hyperbolic discounting model					
Discounting parameter (δ_H^{mix})	0.730***	0.069	0.000	0.596	0.865
Mixture probability (π^H)	0.653***	0.034	0.000	0.586	0.720
Error terms					
Risk Error (μ)	0.167***	0.008	0.000	0.152	0.182
Time Error (ν)	0.051***	0.015	0.001	0.021	0.081
N	17500				
log-likelihood	-8808.992				

H_0: $\pi^E = 0.5$, p-value < 0.001

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The mixture model in Table XVI also shows how discounting parameter estimates are distorted when the E or H models have to account for all the data. Figure VIII plots the discount factors from the E and H models in Table XI (i.e., when they are assumed to account for all the data) and the discount factors from the mixture model in Table XVI.

In Model 1 of Table XI, where the E model was assumed to be the sole DGP, the estimate of $\delta_E = 0.493$. In the mixture model, the estimate of δ_E, which I refer to as δ_E^{mix} in Table XVI and Figure VIII, is far lower at 0.137. This implies that when one tries to make all the data fit the E model, one inflates the estimate of the discounting parameter since 65% of the data "wants" to be modelled as H. Similarly, in Model 2

of Table XI, where the H model was assumed to be the sole DGP, the estimate of δ_H = 0.502. In the mixture model in Table XVI, the estimate of δ_H^{mix} is far higher at 0.730. Thus, by assuming one DGP we are averaging the estimates that we derive when allowing multiple DGPs to characterise the data.[54]

Figure VIII: Exponential and hyperbolic discount factors

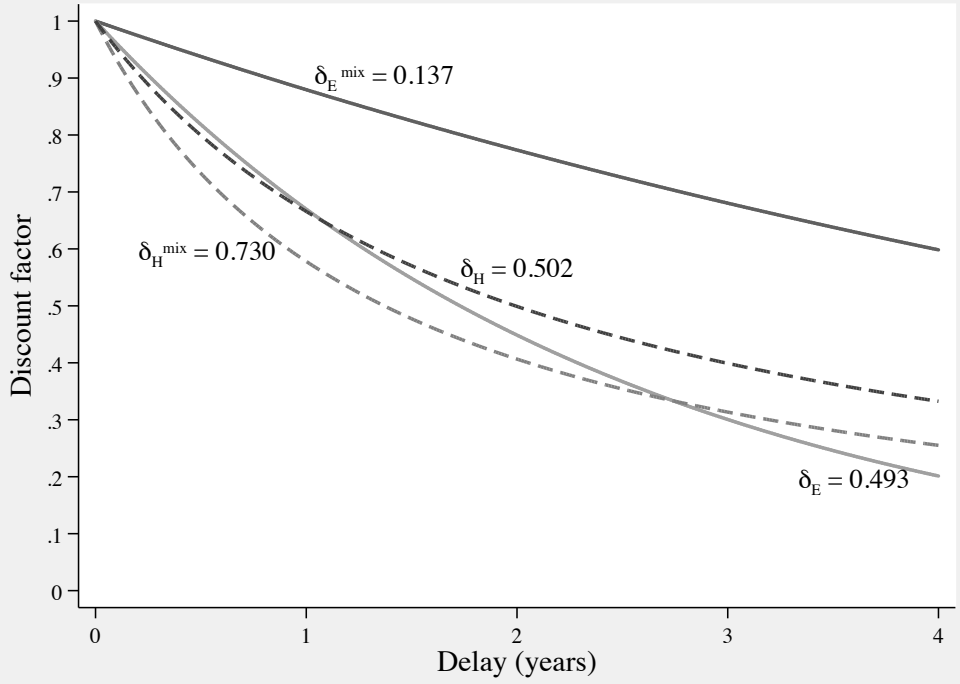

Finally, notice that the estimate of the Fechner error term ν = 0.051 in the mixture model in Table XVI is far lower than the estimates of ν for the E and H models in Table XI. Thus, what was being captured as subject errors in decision making when estimating the E and H models separately is partly the product of forcing the data to fit one DGP.

Mixture models also allow one to explore the hypothesis that smokers are more likely to discount hyperbolically than non-smokers by making the mixture probability a function of smoking status. Recall that with an additively-separable intertemporal

[54] Harrison and Rutström (2009) reach a similar conclusion in the context of choice under risk. Specifically, under the assumption that prospect theory (PT) is the sole DGP, Harrison and Rutström (2009, p. 146) find limited evidence of loss aversion, no significant evidence of probability weighting, and utility function estimates for the gain and loss frames which are not significantly different to one another. In a mixture model of EU and PT, by contrast, the researchers find substantial evidence of loss aversion, significant probability weighting, and utility function estimates which differ significantly across the gain and loss frames.

utility function the E model implies time-consistent preferences whereas the H model may yield time-inconsistent choices. If one finds that current smokers are more likely to discount according to the H model as opposed to the E model, this suggests that they may be more prone to time inconsistency.

TABLE XVII: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HETEROGENOUS PREFERENCES

	Estimate	Std error	*p*-value	95% Confidence Interval	
Power function parameter (r)					
Smoker	0.055**	0.028	0.047	0.001	0.110
Constant	0.302***	0.028	0.000	0.247	0.358
PWF parameter (γ)					
Smoker	0.036	0.052	0.488	-0.066	0.138
Constant	0.778***	0.039	0.000	0.701	0.854
PWF parameter (η)					
Smoker	0.139*	0.082	0.091	-0.022	0.300
Constant	0.812***	0.048	0.000	0.718	0.905
Discounting parameter (δ_E^{mix})					
Smoker	0.051*	0.026	0.052	-0.000	0.102
Constant	0.111***	0.020	0.000	0.073	0.150
Discounting parameter (δ_H^{mix})					
Smoker	0.164**	0.081	0.044	0.004	0.323
Constant	0.622***	0.077	0.000	0.472	0.773
Mixture probability (π^E)					
Smoker	-0.149**	0.074	0.044	-0.293	-0.004
Constant	0.440***	0.064	0.000	0.315	0.565
Error terms					
Risk error (μ)	0.167***	0.008	0.000	0.152	0.182
Time error (ν)	0.048***	0.013	0.000	0.023	0.073
N	17500				
log-likelihood	-8657.459				

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table XVII presents estimates of the mixture model of the E and H discounting functions where the parameters of interest are allowed to vary by smoking status. The results for the risk and time preference parameters mirror those from Table XV:A. Specifically, smokers have significantly less concave utility functions than non-smokers and the estimate of η is significantly higher among smokers than non-smokers. With regard to the time preference parameters, smokers discount significantly more than non-smokers under both the E and H models. Finally, I find that smokers are significantly less likely to discount exponentially, which means they are significantly more likely to discount hyperbolically, than non-smokers. The magnitude of this result is economically significant in that smokers are 15% more likely to discount hyperbolically than non-smokers. Consequently, under the

assumption of an additively-separable intertemporal utility function, smokers may be more likely to make time-inconsistent choices than non-smokers.

TABLE XVIII: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HETEROGENOUS MIXTURE PROBABILITY

	Estimate	Std error	*p*-value	95% Confidence Interval	
Power function parameter (r)					
Smoker	0.053*	0.029	0.069	-0.004	0.111
Constant	0.314***	0.029	0.000	0.257	0.371
PWF parameter (γ)					
Smoker	0.037	0.052	0.478	-0.065	0.138
Constant	0.777***	0.039	0.000	0.701	0.853
PWF parameter (η)					
Smoker	0.138*	0.083	0.096	-0.024	0.300
Constant	0.822***	0.050	0.000	0.725	0.920
Discounting parameter (δ_E^{mix})					
Smoker	0.054**	0.026	0.038	0.003	0.105
Constant	0.118***	0.020	0.000	0.078	0.158
Discounting parameter (δ_H^{mix})					
Smoker	0.154*	0.088	0.080	-0.019	0.327
Constant	0.652***	0.083	0.000	0.489	0.814
Mixture probability (π^E)					
Age	0.003	0.015	0.858	-0.026	0.032
White	0.109	0.085	0.200	-0.058	0.275
Male	-0.105	0.083	0.207	-0.268	0.058
Commerce faculty	0.009	0.083	0.912	-0.153	0.171
Financial aid	-0.093	0.069	0.181	-0.229	0.043
Risk task first	-0.045	0.070	0.519	-0.183	0.092
FED: 1 week	-0.084	0.080	0.291	-0.241	0.072
FED: 2 weeks	-0.014	0.083	0.864	-0.178	0.149
High Principal	0.199***	0.021	0.000	0.158	0.239
Smoker	-0.141	0.089	0.116	-0.316	0.035
Constant	0.385	0.290	0.183	-0.182	0.953
Error terms					
Risk error (μ)	0.166***	0.007	0.000	0.151	0.180
Time error (ν)	0.054***	0.014	0.000	0.027	0.080
N	17500				
log-likelihood	-8552.916				

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The finding that smokers are more likely to discount hyperbolically than non-smokers should be tempered by the results in Table XVIII. This table presents estimates of the mixture model with a heterogenous mixture probability which varies according to the demographic characteristics and task parameters from Table III. With a full set of covariates for the mixture probability, the point estimate of the smoker variable is very similar to the estimate in Table XVII, but it does not quite reach statistical significance ($p = 0.116$). The only factor which significantly affects the mixture probability in this model is a high principal in the time preference task: at the higher

principal of R250, people are approximately 20% more likely to discount exponentially than hyperbolically.

Figure IX shows a kernel density plot, based on the estimates in Table XVIII, of the mixture probability π^E for smokers and non-smokers. The distribution of estimates for smokers, in comparison to non-smokers, is concentrated at low values of π^E, which highlights the finding from Table XVII that smokers are more likely to discount hyperbolically than non-smokers. However, as Table XVIII shows, the evidence for this finding is inconclusive once the full set of demographic characteristics and task parameters have been included in the mixture probability.

Figure IX: Smoking status, exponential and hyperbolic discounting

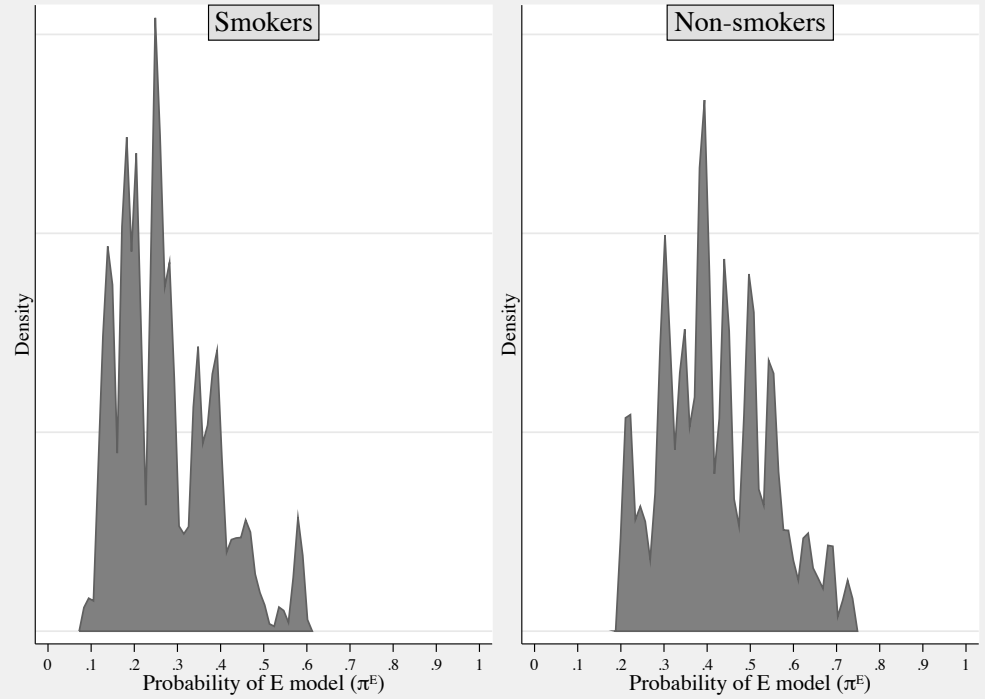

In sum, the results from the mixture model analyses suggest that it is mistaken to assume that only one model of discounting behaviour accurately characterises all discounting choices by all subjects. Recognising that some choices are better explained by one specification while other choices are better explained by another specification, allows one to draw more accurate inferences about the type and extent of discounting behaviour. Finally, the results suggest that smokers may be more likely

to discount hyperbolically than exponentially but conclusive evidence of this assertion is still lacking.

VI. DISCUSSION AND CONCLUSIONS

This chapter provided a detailed review of studies investigating the relationship between risk preferences and smoking behaviour. This review brought to light a number of patterns in the literature: 1) the majority of studies were conducted in the US; 2) the samples, although relatively small, were predominantly drawn from the general population; 3) there was a roughly equal split between choice and titration elicitation mechanisms; 4) every study except HLR adopted a two-step approach to statistical analysis; 5) the majority of studies used real rewards or probabilistic payment schemes; 6) most studies adopted the PD approach[55] to risk preferences where risk attitudes are determined solely by the shape of the PWF; and, unlike the literature on time preferences and smoking, 7) the results linking risk preferences to smoking behaviour are equivocal.

I sought to analyse the relationship between risk and time preferences and smoking behaviour by drawing a relatively large sample of smokers and non-smokers from the student population at UCT. The experimental tasks were carefully designed so as to promote the truthful revelation of preferences and the statistical approach to data analysis allowed me to draw accurate inferences about the preferences of smokers and non-smokers.

I found that both probability weighting and utility function curvature affect attitudes to risk in this sample but I found no statistically significant relationship between risk preferences and smoking status. This result was robust to different theories of choice under risk, different PWFs, and different utility functions which admit IRRA, DRRA, and CRRA.

[55] Chapter 4 showed that the PD model is just Yaari's (1987) dual theory of choice under risk limited to a circumscribed class of lotteries and with a specific PWF: $\pi(p) = p / [p + \gamma(1 - p)]$; if $\gamma > 1$ this represents probability pessimism and risk aversion. Recall that to derive this form for the PWF, Rachlin, Raineri and Cross (1991) severed the purported link between probability and delay which they argued was so central to the way people interpret probabilities. Thus, the specific form of this PWF lacks consistent theoretical foundations.

To analyse the time preferences of smokers and non-smokers I adopted the methodology of HLR which jointly estimates utility function curvature and discounting behaviour so as to characterise time preferences over utility flows, not flows of money. I found that controlling for the concavity of the utility function led to a dramatic decline in estimates of δ, thereby replicating the result of AHLR. I also allowed RDU to characterise choice under risk so as to apportion risk preferences into their utility curvature and probability weighting components.

Despite the presence of probability weighting in this sample, I found that EU estimates of the power function parameter *r* were not significantly different to the estimates I obtained under RDU and the Prelec PWF. Consequently, assuming EU or RDU had no significant impact on the estimates from the discounting models. Nevertheless, I chose to estimate the discounting models assuming RDU and the Prelec PWF so as to incorporate this feature (i.e., probability weighting) of the data.

In the homogenous preferences discounting models, the estimate of the Fechner error term was positive and statistically significant, which highlights the importance of formally incorporating behavioural errors in decision making when analysing time preference data. I also found that the estimate of β in the QH model was significantly less than 1, and the estimate of β in the WB model was significantly greater than 1, which generates declining discount rates over time. Finally, a comparison of the models' log-likelihoods in Table XI suggested that the WB model best characterises the data, followed by the QH model, the H model, and the E model.

To formally adjudicate between the discounting models, Clarke (2007) non-nested model selection tests were conducted. These tests confirmed that the WB model outperformed its competitors. In addition, the Clarke (2007) tests provided a transitive ranking of the discounting specifications: WB > H > E > QH. Thus, despite its relatively high log-likelihood, the QH model was the least successful in explaining all of the discounting choices in these data.

The relationship between time preferences and smoking behaviour was explored in three ways. First, a smoking status covariate was added to the discounting models to

capture the "total effect" of smoking on discounting behaviour without controlling for other factors like age, gender, etc. Across every discounting model, the estimate of δ for smokers was positive and statistically significant, implying smokers discount more heavily than non-smokers. However, smoking status was not related to the extent of present-bias in the QH model nor in the perception of time in the WB model.

Second, to investigate whether smoking intensity was related to discounting behaviour, I estimated the four time preference models and allowed the parameters of interest to vary as a quadratic function of number of cigarettes smoked per day. These analyses revealed a concave relationship between smoking intensity and estimates of the discounting parameter δ. Specifically, every additional cigarette leads to an increase in discounting, but at a decreasing rate until a maximum is reached, after which every additional cigarette leads to a decrease in discounting.

Finally, to explore the "marginal effect" of smoking status on time preferences, I estimated the discounting models and made the parameters of interest a linear function of observable characteristics and task parameters. Across all specifications, the estimate of δ for smokers was positive and statistically significant, which thereby replicated the earlier result while controlling for other variables which may mediate the relationship between smoking and discounting. In Appendix H I also test to see whether these results are robust to the assumption that EU characterises choice under risk: the results are qualitatively identical to those in Section V. In Appendix I I estimate the four discounting specifications under the assumption of linear utility. The appendix shows that one can reach different qualitative conclusions when one does not jointly estimate utility function curvature and discounting behaviour. Appendix I further shows, therefore, that it is imperative to incorporate the shape of the utility function when estimating time preferences.

The final set of statistical analyses was the estimation of mixture models of the different discounting specifications, assuming that RDU and the Prelec PWF capture choice under risk. These mixture models showed that multiple decision processes characterise the discounting of delayed rewards and that it is a mistake to force all discounting choices to fit one particular model. In addition, the mixture models allowed me to explore the hypothesis that smokers are more likely to discount

hyperbolically than non-smokers by making the mixture probability a function of smoking status. I found that smokers are significantly more likely than non-smokers to discount hyperbolically, but this result does not attain statistical significance when the full set of demographic characteristics and task parameters are added to the mixture probability.

This research makes a number of contributions to the literature. When analysing risk preferences and smoking behaviour, I allowed risk attitudes to be determined both by utility function curvature and probability weighting. Prior studies in the literature were always open to the critique that the other source of risk attitudes, i.e., the one not explored in the paper, differed between smokers and non-smokers. Incorporating both utility function curvature and probability weighting in estimates of risk attitudes, allowed me to make strong claims about differences in the risk preferences of smokers and non-smokers which were immune to this critique.

This is only the second study in the smoking-discounting literature to incorporate utility function curvature in the estimation of time preference models, and it is the first which allows RDU to characterise choice under risk. Although the discounting estimates do not differ significantly across the EU and RDU specifications, it is nevertheless theoretically appropriate to apportion risk preferences into their utility curvature and probability weighting components.

This is the first study in this literature to identify a nonlinear effect of smoking intensity on discounting behaviour. Smoking more cigarettes tends to increase discounting but only up to a point, after which each additional cigarette tends to lower discounting. This nonlinear effect may explain why some studies, which only recruited heavy smokers and never-smokers, failed to find a difference in discounting behaviour between these groups.

In addition, this nonlinear effect of smoking intensity may provide an explanation for patterns of cigarette consumption. It has long been assumed that the marked modal clustering around 20 cigarettes per day in mature smokers simply reflects the fact that cigarettes are typically sold in packs of 20. It may be the case though that cigarette companies learned to sell cigarettes in packs of 20 because that is where the

psychofunctional, and not merely the homeostatic, equilibrium lies for the majority of mature smokers.

This research also reiterates the point made in HLR and Chapter 3, that multiple decision processes characterise the discounting of delayed rewards. It is crucial for researchers to be cognisant of this fact when exploring the smoking-discounting relationship. As the analyses in this chapter showed, smokers may be more likely than non-smokers to discount hyperbolically and this may be a factor in addiction. To my knowledge, this is the first study in the literature to have identified this difference between a set of smokers and non-smokers.

This research also suffers from a number of limitations. Clearly a young, university sample of smokers is not representative of smokers in general, so the external validity of these results for the South African population as a whole is questionable. This study is an important first step though towards population-based studies of risk and time preferences and smoking behaviour. The tasks and instructions developed for this study could be used to elicit a representative sample of South African's attitudes toward risk and time, which can then be analysed using the tools adopted in this chapter. Funding permitting, this will be taken up in future research.

A perennial problem with economic experiments is whether the rewards on offer are salient enough to incentivise truthful revelation of preferences over an income domain which is relevant to policy analysis. The rewards in this study were large in comparison to those typically paid out in the literature but still larger incentives would help to alleviate concerns about the extent to which these results scale up as the magnitude of the prizes increases.

Another potential issue with these results is whether the sample in this study is representative of smokers and non-smokers at UCT. As discussed earlier, a large number of people applied to take part in the study so people in the smoking and non-smoking groups were randomly selected to form part of the study pool. It may be the case that those who were selected are not representative of their group. Ideally I would use information on the population of smokers and non-smokers at UCT to

correct for any sample selection issues present in the data.[56] Unfortunately, I do not have any additional information on the population of smokers and non-smokers at UCT to perform these sample selection corrections.

These issues notwithstanding, this chapter provides a rigorous framework within which to analyse risk and time preferences and smoking behaviour. In this sample, there is no statistically significant relationship between smoking and risk preferences, a result which, to the uninitiated, may suggest that it is unnecessary to collect these data in future studies (but see Appendix I). As this chapter hopefully showed though, it is still crucial to collect risk preference data so as to define time preferences over utility flows so that accurate inferences about discounting behaviour can be drawn.

[56] Harrison, Lau and Rutström (2009) and Harrison and Lau (2014) analyse the effect of sample selection bias on estimated risk preference parameters. They used the Danish registry to gather information on people who were invited to participate in their experiments but who did not take part and this allowed them to make sample selection corrections for the sample of people who were invited and who did participate in the experiments. Harrison, Lau and Rutström (2009) find that correcting for sample selection bias leads to attenuated risk aversion estimates, implying that their sample was more risk averse than the population from which it was drawn. Similarly, Harrison and Lau (2014) find that sample selection corrections lead to lower estimates of risk aversion.

REFERENCES

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583-618.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2013): "Discounting Behaviour and the Magnitude Effect: Evidence from a Field Experiment in Denmark," *Economica*, 80, 670-697.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2014): "Discounting Behavior: A Reconsideration," *European Economic Review*, 71, 15-33.

ANDERSON, L. R., AND J. M. MELLOR (2008): "Predicting Health Behaviors with an Experimental Measure of Risk Preference," *Journal of Health Economics*, 27, 1260-1274.

BRADFORD, W. D. (2010): "The Association between Individual Time Preferences and Health Maintenance Habits," *Medical Decision Making*, 30, 99-112.

CLARKE, K. A. (2007): "A Simple Distribution-Free Test for Nonnested Model Selection," *Political Analysis*, 15, 347-363.

COLLER, M., G. W. HARRISON, AND E. E. RUTSTRÖM (2012): "Latent Process Heterogeneity in Discounting Behavior," *Oxford Economic Papers*, 64, 375-391.

COLLER, M., AND M. B. WILLIAMS (1999): "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 107-127.

CONTE, A., J. D. HEY, AND P. G. MOFFAT (2011): "Mixture Models of Choice under Risk," *Journal of Econometrics*, 162, 79-88.

FECHNER, G. (1966): *Elements of Psychophysics*. New York, NY: Holt, Rinehart and Winston.

GONZALEZ, R., AND G. WU (1999): "On the Shape of the Probability Weighting Function," *Cognitive Psychology*, 38, 129-166.

HARRISON, G. W., S. J. HUMPHREY, AND A. VERSCHOOR (2009): "Choice under Uncertainty: Evidence from Ethiopia, India and Uganda," *The Economic Journal*, 120, 80-104.

HARRISON, G. W., AND M. I. LAU (2014): "Risk Attitudes, Sample Selection and Attrition in a Longitudinal Field Experiment," Working Paper WP 2014-04, Center for the Economic Analysis of Risk (CEAR) Georgia State University.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2007): "Estimating Risk Attitudes in Denmark: A Field Experiment," *Scandinavian Journal of Economics*, 109, 341-369.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2009): "Risk Attitudes, Randomization to Treatment, and Self-Selection into Experiments," *Journal of Economic Behavior and Organization*, 70, 498-507.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2010): "Individual Discount Rates and Smoking: Evidence from a Field Experiment in Denmark," *Journal of Health Economics*, 29, 708-717.

HARRISON, G. W., AND E. E. RUTSTRÖM (2008): "Risk Aversion in the Laboratory," in *Research in Experimental Economics: Volume 12. Risk Aversion in Experiments*, ed. by J. C. Cox, and G. W. Harrison. Bingley: Emerald, 41-196.

HARRISON, G. W., AND E. E. RUTSTRÖM (2009): "Expected Utility Theory and Prospect Theory: One Wedding and a Decent Funeral," *Experimental Economics*, 12, 133-158.

HEATHERTON, T. F., L. T. KOZLOWSKI, R. C. FRECKER, AND K. FAGERSTRÖM (1991): "The Fagerström Test for Nicotine Dependence: A Revision of the Fagerström Tolerance Questionnaire," *British Journal of Addiction*, 86, 1119-1127.

HEY, J. D., AND C. ORME (1994): "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62, 1291-1326.

HOLT, C. A., AND S. K. LAURY (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92, 1644-1655.

INGERSOLL, J. (2008): "Non-Monotonicity of the Tversky-Kahneman Probability-Weighting Function: A Cautionary Note," *European Financial Management*, 14, 385-390.

JAMISON, D. T., AND J. JAMISON (2011): "Characterizing the Amount and Speed of Discounting Procedures," *Journal of Benefit-Cost Analysis*, 2, 1-53.

JUSOT, F., AND M. KHLAT (2013): "The Role of Time and Risk Preferences in Smoking Inequalities: A Population-Based Study," *Addictive Behaviors*, 38, 2167-73.

LAWYER, S. R., F. SCHOEPFLIN, R. GREEN, AND C. JENKS (2011): "Discounting of Hypothetical and Potentially Real Outcomes in Nicotine-Dependent and Nondependent Samples," *Experimental and Clinical Psychopharmacology*, 19, 263-74.

LOOMES, G., AND R. SUGDEN (1998): "Testing Different Stochastic Specifications of Risky Choice," *Economica*, 65, 581-598.

MAZUR, J. E. (1984): "Tests of an Equivalence Rule for Fixed and Variable Reinforcer Delays," *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 426-437.

MCLACHLAN, G., AND D. PEEL (2000): *Finite Mixture Models*. New York, NY: Wiley.

MILLER, L., D. E. MEYER, AND J. T. LANZETTA (1969): "Choice among Equal Expected Value Alternatives: Sequential Effects of Winning Probability Level on Risk Preferences," *Journal of Experimental Psychology*, 79, 419-423.

MITCHELL, S. H. (1999): "Measures of Impulsivity in Cigarette Smokers and Non-Smokers," *Psychopharmacology (Berlin)*, 146, 455-64.

MYERSON, J., L. GREEN, AND M. WARUSAWITHARANA (2001): "Area under the Curve as a Measure of Discounting," *Journal of the Experimental Analysis of Behavior*, 76, 235-43.

OEHLERT, G. W. (1992): "A Note on the Delta Method," *The American Statistician*, 46, 27-29.

OHMURA, Y., T. TAKAHASHI, AND N. KITAMURA (2005): "Discounting Delayed and Probabilistic Monetary Gains and Losses by Smokers of Cigarettes," *Psychopharmacology (Berlin)*, 182, 508-15.

POLTAVSKI, D. V., AND J. N. WEATHERLY (2013): "Delay and Probability Discounting of Multiple Commodities in Smokers and Never-Smokers Using Multiple-Choice Tasks," *Behavioural Pharmacology*, 24, 659-67.

PRATT, J. W. (1964): "Risk Aversion in the Small and in the Large," *Econometrica*, 32, 122-136.

PRELEC, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497-527.

QUIGGIN, J. (1982): "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3, 323-343.

RACHLIN, H., A. RAINERI, AND D. CROSS (1991): "Subjective Probability and Delay," *Journal of the Experimental Analysis of Behavior*, 55, 233-44.

REYNOLDS, B. (2006): "The Experiential Discounting Task Is Sensitive to Cigarette-Smoking Status and Correlates with a Measure of Delay Discounting," *Behavioural Pharmacology*, 17, 133-42.

REYNOLDS, B., K. KARRAKER, K. HORN, AND J. B. RICHARDS (2003): "Delay and Probability Discounting as Related to Different Stages of Adolescent Smoking and Non-Smoking," *Behavioural Processes*, 64, 333-344.

REYNOLDS, B., M. PATAK, P. SHROFF, R. B. PENFOLD, S. MELANKO, AND A. M. DUHIG (2007): "Laboratory and Self-Report Assessments of Impulsive Behavior in Adolescent Daily Smokers and Nonsmokers," *Experimental and Clinical Psychopharmacology*, 15, 264-71.

REYNOLDS, B., J. B. RICHARDS, K. HORN, AND K. KARRAKER (2004): "Delay Discounting and Probability Discounting as Related to Cigarette Smoking Status in Adults," *Behavioural Processes*, 65, 35-42.

RICHARDS, J. B., L. ZHANG, S. H. MITCHELL, AND H. DE WIT (1999): "Delay or Probability Discounting in a Model of Impulsive Behavior: Effect of Alcohol," *Journal of the Experimental Analysis of Behavior*, 71, 121-43.

SAHA, A. (1993): "Expo-Power Utility: A 'Flexible' Form for Absolute and Relative Risk Aversion," *American Journal of Agricultural Economics*, 75, 905-913.

STOTT, H. P. (2006): "Cumulative Prospect Theory's Functional Menagerie," *Journal of Risk and Uncertainty*, 32, 101-130.

SZREK, H., L. W. CHAO, S. RAMLAGAN, AND K. PELTZER (2012): "Predicting (Un)Healthy Behavior: A Comparison of Risk-Taking Propensity Measures," *Judgment and Decision Making*, 7, 716-727.

TANAKA, T., C. F. CAMERER, AND Q. NGUYEN (2010): "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam," *American Economic Review*, 100, 557-571.

TVERSKY, A., AND D. KAHNEMAN (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297-323.

VAN DER POL, M., AND M. RUGGERI (2008): "Is Risk Attitude Outcome Specific within the Health Domain?," *Journal of Health Economics*, 27, 706-17.

VUONG, Q. H. (1989): "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 57, 307-333.

WAKKER, P. P. (2008): "Explaining the Characteristics of the Power (CRRA) Utility Family," *Health Economics*, 17, 1329-1344.

WILCOX, N. T. (2011): "'Stochastically More Risk Averse:' A Contextual Theory of Stochastic Discrete Choice under Risk," *Journal of Econometrics*, 162, 89-104.

YAARI, M. E. (1987): "The Dual Theory of Choice under Risk," *Econometrica*, 55, 95-115.

YI, R., W. D. CHASE, AND W. K. BICKEL (2007): "Probability Discounting among Cigarette Smokers and Nonsmokers: Molecular Analysis Discerns Group Differences," *Behavioural Pharmacology*, 18, 633-9.

6. Conclusion

This thesis considered experiments investigating the relationship between instantaneous risk and time preferences and addiction. Addiction is a pressing public health concern and economic models of this phenomenon suggest that instantaneous risk and time preferences play important roles in the causes, course, and consequences of substance dependence. Detailed reviews of the experimental literature linking these preferences to smoking status highlighted a number of methodological and statistical flaws in the way these data have been collected and analysed. Consequently, this thesis, following Harrison, Lau and Rutström (2014), motivated for an approach to the study of addiction which views theory, experimental design, and econometrics as complementary.

Chapter 2 provided a motivation for studying addiction and then focussed on economic models of this phenomenon to isolate the preferences which may be a factor in substance dependence. Emphasis was given to the roles that instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs play in these models. The empirical focus of the thesis shifted to instantaneous risk and time preferences, with research into intertemporal risk preferences and subjective beliefs, and their relationship to addiction, deferred for future study.

Chapter 3 used two relatively large datasets, which were generated by choices on the Kirby, Petry and Bickel (1999) discounting task, to explore the relationship between time preferences and smoking behaviour. Unfortunately these data did not include information which would allow for the joint estimation of utility function curvature and discounting behaviour so the common assumption of linear utility was employed. The chapter introduced a full information maximum likelihood statistical framework, which is consistent with the data generating processes (DGP) proposed by structural theories and accounts for subject errors in decision making, to estimate four distinct discounting models: exponential, hyperbolic, quasi-hyperbolic, and Weibull.

Non-nested model selection tests were conducted and the exponential discounting model better characterised the data than the hyperbolic discounting model, which

challenges a maintained assumption in the discounting literature in psychology and addiction studies. Despite the better fit of the exponential function relative to the hyperbolic function, mixture model analyses conducted later in the chapter showed that multiple decision making processes characterise the discounting of delayed rewards and it is a mistake, therefore, to declare any single model the "winner" in terms of goodness of fit.

The relationship between time preferences and smoking behaviour was explored in two ways. The first simply added a smoking status covariate to the discounting models so as to capture the "total effect" of smoking. The second used a full set of covariates and task parameters in the models so as to estimate the "marginal effect" of smoking, over and above what could be accounted for by other observable characteristics. Across the different discounting models, and the two approaches to investigating the relationship between smoking and time preferences, a robust picture emerged: current smokers tended to discount the future more heavily than never-smokers. These results replicated previous findings in the literature but with a valid statistical approach to data analysis.

Chapter 3 then presented a mixture model of exponential and hyperbolic discounting specifications and found that the exponential model accounted for roughly 54% of the choices in the data whereas the hyperbolic model accounted for the remaining 46%. A test that the two models explained the same proportion of the data could not be rejected, which highlights the point made earlier that multiple decision making processes characterise the discounting of delayed rewards. The mixture model results also showed that by forcing the data to conform to one DGP this yields biased estimates of the underlying parameters. Finally, a mixture model employing a full set of covariates and task parameters showed that current smokers were no more likely than ex-smokers and never-smokers to discount hyperbolically.

Chapter 4 traced the development of the probability discounting (PD) model which has become a popular framework in psychology and addiction studies for exploring people's atemporal attitudes to risk. This review critiqued the theoretical foundations of the model and reinterpreted it in language familiar to economists, statisticians, and decision theorists. Specifically, the chapter showed that the PD model is just Yaari's

(1987) dual theory of choice under risk limited to a certain class of lotteries and employing a specific probability weighting function (PWF). The efficacy of this PWF was then tested against others which are commonly used in the literature on instantaneous risk preferences: the PWF popularised by Tversky and Kahneman (1992) and the Prelec (1998) PWF. Non-nested model selection tests suggested that both the Tversky and Kahneman (1992) and Prelec (1998) PWFs better characterised the data than the PD model's PWF. In addition, in a mixture model of the PD and TK PWFs, a test that the mixture probability assigned to the PD model was zero could not be rejected. Chapter 4 therefore highlighted the theoretical and empirical inadequacy of the PD model.

Chapter 5 embodied the logic on which this thesis is based: that theory, experimental design, and econometrics are complementary. The chapter discussed a relatively large study of instantaneous risk and time preferences conducted at the University of Cape Town (UCT) in 2012. For this study, risk and time preference tasks were carefully designed to promote the truthful revelation of preferences and to facilitate estimation of several different theories of choice under risk and over time. In addition, a theoretically-motivated joint estimation approach to time preferences was adopted.

Across expected utility and rank-dependent utility models of choice under risk, and across different utility and probability weighting functions, no statistically significant relationship between atemporal attitudes to risk and smoking behaviour could be identified. By contrast, across four distinct discounting functions, which were estimated jointly with the concavity of the utility function, smokers robustly discounted more heavily than non-smokers. In addition, a concave relationship between smoking intensity and discounting behaviour was found: every additional cigarette leads to an increase in discounting, but at a decreasing rate until a maximum is reached, after which every additional cigarette leads to a decrease in discounting.

A mixture model of the exponential and hyperbolic discounting specifications showed that the exponential model accounted for approximately 35% of the choices in the data whereas the hyperbolic model accounted for the remaining 65%. A test that the exponential model explained 50% of the data was easily rejected but so too was a test that the exponential model explained none of the choices in the dataset. This mixture

model therefore reiterated the conclusion from Chapter 3 that multiple decision making process characterise the discounting of delayed rewards and that when one tries to force all the data to fit one discounting model, this yields biased estimates of the underlying parameters.

To test whether smokers were more likely than non-smokers to discount hyperbolically as opposed to exponentially, a smoking status covariate was added to the mixture model. In this specification, smokers were significantly less likely than non-smokers to discount exponentially, implying that they were significantly more likely to discount hyperbolically. However, when the mixture probability was allowed to vary as a linear function of the full set of demographic characteristics and task parameters, the smoking status covariate was not statistically significant. Thus, while there was some evidence that smokers were more likely than non-smokers to discount hyperbolically rather than exponentially, this result was not robust to the inclusion of a full set of covariates.

This thesis makes a number of contributions to the literature. First, it advocates an approach to the study of addiction which views theory, experimental design, and econometrics as complementary. Unfortunately, the experimental literature on addiction is plagued by studies that make strong, and unwarranted, assumptions about the nature of instantaneous risk and time preferences, experimental designs which do not satisfy Smith's (1982) precepts, and statistical tools which are not valid. Embracing the theory, experimental design, and econometric trinity allows one to avoid the pitfalls of misconstrued theory, ineffectual experimental design, and sloppy statistical analysis.

Second, this thesis is the first examination of the relationship between time preferences and smoking behaviour which estimated four distinct discounting specifications. These estimates were produced using a statistical framework which recognises, and takes into account, the uncertainty of discounting parameter estimates and the potential for subject errors in decision making. Hence, this framework makes it possible to draw accurate statistical inferences. That the smoking-discounting relationship was robust to different discounting models lends further credence to the hypothesis that smokers discount the future more heavily than non-smokers.

Third, this thesis conducted the first valid statistical tests to adjudicate between the different time preference models. In Chapter 3, non-nested model selection tests showed that the exponential discounting model outperformed the hyperbolic discounting model, but in Chapter 5 these same tests showed that the hyperbolic model outperformed the exponential model. In addition, the Weibull discounting function had the least explanatory power in the data used in Chapter 3 but it had the most explanatory power in the data used in Chapter 5. These results show that the efficacy of different discounting functions varies from population to population. Consequently, it is not appropriate to assume that one discounting function is always superior to others; this hypothesis can and should be tested on a case-by-case basis using valid statistical tests.

Fourth, this is the first study in the smoking-discounting literature to have estimated "profile likelihoods" to determine the optimal shape of the utility function conditional on the subjects' discounting choices. This approach was adopted by Harrison and Rutström (2009) in the context of choice under risk but it has not been applied to time preference data. Recall that the data in Chapter 3 did not include information which would make it possible to estimate discounting parameters jointly with the concavity of the utility function. However, Appendix B shows that by incrementally varying the parameter defining an assumed utility function, one can identify the optimal shape of the utility function conditional on the subjects' discounting choices. The appendix discusses the logic of this method and presents results for the four discounting specifications used in the chapter.

Fifth, this thesis provides the first explanation of the PD model in language familiar to economists. It is also the first study to have conducted valid statistical tests of the PD model's ability to characterise risk preference data and the first study which estimated a mixture model of the PD model's PWF and the PWF popularised by Tversky and Kahneman (1992).

Sixth, Chapter 5's analyses of the relationship between atemporal attitudes to risk and smoking behaviour are the first to have allowed both utility function curvature and probability weighting to characterise choice under risk. It is also the first study in the

addiction literature to use multiple utility and probability weighting functions in the analysis of risk preference data.

Seventh, Chapter 5 is only the second study in the literature on smoking and time preferences to have jointly estimated utility function curvature and discounting behaviour, and the first to employ a rank-dependent utility characterisation of choice under risk when jointly estimating the parameters of time preference theories. By using the rank-dependent utility model, Chapter 5 apportioned risk preferences into their concave utility and probability weighting components, which can have important implications for the inferences one draws about discounting behaviour; the time preference results in Chapter 5 were not very sensitive though to the assumption that rank-dependent utility theory rather than expected utility theory characterised choice under risk.

Eighth, this thesis is the first research to have identified a nonlinear effect of smoking intensity on discounting behaviour. Across all of the discounting specifications, estimates of the long-term discounting parameter δ revealed that every additional cigarette leads to an increase in discounting, but at a decreasing rate until a maximum is reached, after which every additional cigarette leads to a decrease in discounting. These results may provide an explanation for why cigarettes are typically sold in packs of 20 and why some studies that recruited heavy smokers and never-smokers failed to find a significant difference in discounting behaviour.

Finally, this thesis is only the second study in the addiction literature to have estimated mixture models of discounting behaviour. The crucial insight which these models provide is that multiple decision making processes characterise the discounting of delayed rewards and it is a mistake, therefore, to force the data to conform to one model. In addition, the mixture model of exponential and hyperbolic discounting in Chapter 5 provided the first, tentative evidence that smokers may be more likely to discount hyperbolically rather than exponentially. If this finding is replicated it will sharpen our understanding of addiction. Recall that under an additively-separable intertemporal utility function, exponential discounting yields time consistency whereas hyperbolic discounting may yield time inconsistency. Thus, if smokers are more likely to discount hyperbolically they may be more prone to time

inconsistency, and this inconsistency may be an important factor in their susceptibility to addiction.

This thesis also suffers from a number of limitations. In Chapter 3, only a small subset (N = 174) of the experimental subjects were incentivised, and they were only incentivised probabilistically. Thus, to a large extent, the inferences which were drawn were based on the choices people think they would make when faced with those contingencies, or the choices they thought the experimenter wanted them to make, rather than the choices they made and experienced.

In addition, the Kirby, Petry and Bickel (1999) discounting task is a relatively blunt instrument for making precise inferences about discounting behaviour. A task employing more choices, different reward magnitudes, a larger range of implied interest rates, and a front end delay to the smaller, sooner reward, would lend itself more to the precise estimation of discounting parameters.

While this thesis estimated more models of time preference than other studies in the addiction literature, there are a host of alternative discounting models which were not entertained. In future research it may be worthwhile to estimate additional time preference models to determine whether these alter, or refine, our conclusions concerning the relationship between smoking and discounting behaviour.

The research in Chapter 5 was limited solely to university students so this sample is clearly not representative of the South African population as a whole, and the external validity of the results is questionable. But, as discussed in Chapter 5, this research is an important first step towards population-based studies of instantaneous risk and time preferences and smoking behaviour. The tasks and instructions developed for this study could be used to elicit a representative sample of South African's attitudes to risk and time, which can then be analysed using the tools adopted in this thesis.

Another potential issue with the results in Chapter 5 is whether the sample was representative of smokers and non-smokers at UCT. Given that a large number of people applied to take part in the study, people in the smoking and non-smoking groups were randomly selected to form part of the study pool. This process may have

generated a sample which is not representative of the UCT population. Unfortunately without additional information on the population of smokers and non-smokers at UCT, it was not possible to perform any sample selection corrections on the data.

The review of economic models in Chapter 2 highlighted the roles that instantaneous risk and time preferences, intertemporal risk preferences, and subjective beliefs play in the initiation, maintenance, and resolution of addiction. But this thesis only empirically examined the instantaneous risk and time preferences of smokers and non-smokers. I am unaware of any incentivised experimental studies that compare the subjective beliefs and intertemporal risk preferences of addicts and non-addicts. In relation to smoking, it would be valuable to elicit the subjective beliefs of smokers and non-smokers concerning, say, the proportion of lung cancer cases attributable to smoking. Smokers may have better- or worse-informed subjective beliefs, relative to non-smokers, about the potential harms from smoking and incentivised experimental research could shed light on this issue. Andersen, Fountain, Harrison and Rutström (2013) discuss a method for eliciting and jointly estimating instantaneous risk preferences and subjective beliefs, and this approach could be put to good use with a sample of smokers and non-smokers.

Andersen, Harrison, Lau and Rutström (2014) conducted a set of instantaneous risk and time preference experiments, coupled with an intertemporal risk preference experiment, on a representative sample of 413 people in Denmark. The researchers used a full information maximum likelihood statistical framework to jointly estimate instantaneous risk and time preference parameters and an intertemporal risk preference parameter. They found that, on average, the sample displayed intertemporal risk aversion but that about 5% of the sample was intertemporally risk neutral or risk seeking. As discussed in Chapter 2, it would be valuable to use this robust experimental design and statistical framework to explore potential differences in the intertemporal risk preferences of addicts and non-addicts.

In sum, this thesis adopted an approach to scientific research which views theory, experimental design, and econometrics as complementary. Put simply, only with a firm grasp of theory is it possible to design experiments and conduct statistical analyses which lend themselves to accurate and reliable inferences about the latent

constructs of interest. The economic models of addiction in Chapter 2 guided the search for preferences which may be a factor in addiction. The detailed reviews of the extant literature on instantaneous risk preferences, time preferences and smoking then highlighted the ways in which these preferences have been investigated. Given some of the issues which these reviews raised, statistical tools were chosen so as to be consistent with the DGPs proposed by structural theories and to account for subject errors in decision making. Finally, the experimental design in Chapter 5 was formulated to promote the estimation of several different theories of choice under risk and over time, and to allow for the theoretically-motivated joint estimation approach to time preferences which incorporates the curvature of the utility function. Thus, this thesis sought to embody the theory, experimental design, and econometric trinity.

This thesis challenged some of the maintained assumptions in the addiction literature; e.g., that the probability discounting and hyperbolic discounting models best characterise choice under risk and over time, respectively. But it also replicated a previous finding; i.e., that smokers tend to discount the future more heavily than non-smokers. The thesis shows, therefore, that some results withstand careful methodological and statistical scrutiny, whereas other results do not. Even though the experimental literature linking addiction to choice under risk and over time is at least 15 years old, this thesis provides a new set of methodological and statistical tools to better investigate the behavioural correlates of addiction, and thereby hopefully breathe new life into this important area of scientific enquiry.

REFERENCES

ANDERSEN, S., J. FOUNTAIN, G. W. HARRISON, AND E. E. RUTSTRÖM (2013): "Estimating Subjective Probabilities," Working Paper WP 2010-06, Center for the Economic Analysis of Risk (CEAR) Georgia State University.

ANDERSEN, S., G. W. HARRISON, M. I. LAU, AND E. E. RUTSTRÖM (2014): "Multiattribute Utility Theory, Intertemporal Utility and Correlation Aversion," Working Paper WP 2011-03, Center for the Economic Analysis of Risk (CEAR) Georgia State Univerity.

HARRISON, G. W., M. I. LAU, AND E. E. RUTSTRÖM (2014): "Theory, Experimental Design and Econometrics Are Complementary (and So Are Lab and Field Experiments)." in *The Methods of Modern Experimental Economics*, ed. by G. Frechette, and A. Schotter. New York, NY: Oxford University Press.

HARRISON, G. W., AND E. E. RUTSTRÖM (2009): "Expected Utility Theory and Prospect Theory: One Wedding and a Decent Funeral," *Experimental Economics*, 12, 133-158.

KIRBY, K. N., N. M. PETRY, AND W. K. BICKEL (1999): "Heroin Addicts Have Higher Discount Rates for Delayed Rewards Than Non-Drug-Using Controls," *Journal of Experimental Psychology: General*, 128, 78-87.

PRELEC, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497-527.

SMITH, V. L. (1982): "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, 923-955.

TVERSKY, A., AND D. KAHNEMAN (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297-323.

YAARI, M. E. (1987): "The Dual Theory of Choice under Risk," *Econometrica*, 55, 95-115.

7. APPENDICES

APPENDIX A

In this appendix I investigate an inferential issue that arises when using the KPB task to estimate or infer annual, rather than daily, discount rates. This problem was first uncovered when I sought to evaluate annual equivalents of the daily discount rates normally reported in the psychology and addiction literatures. To transform a daily exponential discount rate into an annual discount rate, one commonly[1] uses the following equation:

$$\delta^A = (1 + \delta^D)^{365} - 1, \qquad (1)$$

where δ^A is the annual discount rate and δ^D is the daily discount rate from the exponential discounting model.

Table A:I shows the 27 discounting questions on the KPB task and the implied exponential daily and annual discount rates, assuming linear utility. The table also includes the ratio of the right endpoint and left endpoint of daily and annual discount rate intervals.[2] As one proceeds down the table, the ratio of daily rate intervals varies between 2.17 and 2.67. By contrast, the ratio of annual rate intervals increases from 2.57 to 9894905759808.40, albeit with less extreme values in between. The similarity in daily rate interval ratios and the large increase in annual rate interval ratios in Table A:I has important implications for the inferences one can draw from the KPB task.

[1] In the finance literature, the exponent of 365 in equation (1) is often replaced by 252 to represent the number of trading days in a year. I adopt the banking convention and make use of 365 days to transform a daily exponential discount rate into an annual exponential discount rate.

[2] The KPB task uses six discounting questions to define an exponential discount rate interval: three questions determine the left endpoint of the interval and three questions determine the right endpoint of the interval. For example, the first three rows of Table A:I are consistent with a daily discount rate of 0.00016 and the next three rows of the table are consistent with a daily discount rate of 0.0004; together, these six rows define the discount rate interval (0.00016, 0.0004). To determine the ratio of this interval, I divided the right endpoint (i.e., 0.0004) by the left endpoint (i.e., 0.00016) which provides a daily rate interval ratio of 2.50. In some cases, particularly with annual discount rates, a set of three discounting questions is not consistent with the same discount rate. To calculate the ratio of an interval where the set of three questions is not consistent with a particular rate, I took the largest discount rate from the set as the right endpoint of the interval and the smallest discount rate of the set as the left endpoint of the interval.

Given that the daily rate interval ratios vary between 2.17 and 2.67, the KPB task is sensitive to 3-fold differences in daily discount rates regardless of the pattern of choices on the task. For example, if someone has a daily discount rate of 0.0002 and another person has a daily discount rate of 0.0006 (i.e., a 3-fold difference in daily discount rates), these people will fall into different daily discount rate intervals on the KPB task. Similarly, if someone has a daily discount rate of 0.01 and another person has a daily discount rate of 0.03 (i.e., a 3-fold difference in daily discount rates), these people will fall into different daily discount rate intervals on the KPB task. The daily discount rates in the second example are clearly larger than the daily discount rates in the first example. But the relative difference of the discount rates in the first example is the same as the relative difference in discount rates in the second example, which means the people in the first example fall into different discount rate intervals and so too do the people in the second example. Thus, the uncertainty associated with daily discount rate estimates on the KPB task (i.e., the size of the intervals) is relatively uniform, in a ratio sense, irrespective of choice patterns.

By contrast, the uncertainty associated with annual discount rates increases as one moves down Table A:I, which implies that the sensitivity of the KPB task to differences in annual discount rates is tied to choice patterns. For example, if someone has an annual discount rate of 0.08 and another person has an annual discount rate of 0.24 (i.e., a 3-fold difference in annual discount rates), they will fall into different annual discount rate intervals on the KPB task. However, if someone has an annual discount rate of 1.45 and another person has an annual discount rate of 4.35 (i.e., a 3-fold difference in annual discount rates), they will fall into the *same* annual discount rate interval on the KPB task. In other words, while the KPB task can always a detect a 3-fold difference in daily discount rates, the same is not true of annual discount rates. Note that by row 16 of Table A:I, the KPB task cannot detect a 20-fold difference in annual discount rates, and by row 19 of Table A:I, the KPB task cannot detect a 600-fold difference in annual discount rates.

The fact that the KPB task's sensitivity to differences in annual discount rates is tied to choice patterns has at least two implications. The first is that the precision or reliability of annual discount rate estimates decreases as more SS choices are made on

the task. With only 3 SS choices, the annual discount rate falls into a relatively small interval: (0.06, 0.15). A subject who makes 3 SS choices may have an annual discount rate of 0.08 or an annual discount rate of 0.14, and one cannot determine his "true" annual discount rate based on his pattern of choices. Nevertheless, the relatively small annual discount rate interval of (0.06, 0.15) provides relatively tight bounds on the person's annual discount rate. With 18 SS choices on the KPB task, however, the annual discount rate falls into a relatively large interval: (108, 68776). In this case, one is uncertain whether the person's "true" annual discount rate is 120 or 60000.

If one wanted to use choices on the KPB task to provide estimates of annual discount rates for use in public policy or cost-benefit analyses, this annual rate interval imprecision would likely be problematic. In addition, the large gaps in implied annual discount rates as one moves down the table, makes pooled estimation of a *smooth* annual discounting function difficult. Specifically, the estimation procedure has to contend with some discount rate intervals which are small and others which are large; relatively precise estimates can be obtained for the smaller intervals but relatively imprecise estimates will be obtained for the larger intervals.

The second implication of the KPB task's sensitivity to certain choice patterns with annual discount rates is that one's ability to conduct group comparisons is tied to these choice patterns. For example, if one wants to compare the discounting behaviour of smokers and non-smokers, then it is easier to detect group differences if a large proportion of the people in these groups make a small number of SS choices on the task (i.e., where the task's sensitivity to annual discount rate differences is greatest). For example, suppose the average non-smoker has an annual discount rate of 0.1 and the average smoker has an annual discount rate of 0.2. In this case, the average non-smoker would switch to the LL reward on row 4 of the table and the average smoker would switch to the LL reward on row 7 of the table, thereby making it possible, in principle, to detect a group difference. By contrast, suppose the average non-smoker has an annual discount rate of 6 and the average smoker has an annual discount rate of 100. In this case, the average smoker and the average non-smoker will switch to the LL reward on row 16, thereby making it difficult to detect a group difference.

TABLE A:I

KPB DISCOUNTING TASK WITH IMPLIED ANNUAL RATES AND INTERVAL RATIOS

		Reward values				Daily rate		Annual rate	Fraction choosing LL		
Row	Order	SS	LL	Delay[1]	Daily discount rate	interval ratio	Annual discount rate	interval ratio	Smokers	Non-smokers[2]	Combined
1	13	$34	$35	186	0.00016		0.06		0.01	0.03	0.03
2	1	$54	$55	117	0.00016		0.06		0.08	0.11	0.10
3	9	$78	$80	162	0.00016	2.50	0.06	2.57	0.03	0.05	0.05
4	20	$28	$30	179	0.0004		0.15		0.02	0.08	0.07
5	6	$47	$50	160	0.0004		0.15		0.06	0.07	0.07
6	17	$80	$85	157	0.0004	2.25	0.15	2.70	0.02	0.05	0.05
7	26	$22	$25	136	0.0009		0.41		0.04	0.08	0.08
8	24	$54	$60	111	0.0009		0.41		0.05	0.09	0.09
9	12	$67	$75	119	0.0009	2.56	0.41	3.16	0.03	0.10	0.09
10	22	$25	$30	80	0.0023		1.30		0.06	0.13	0.12
11	16	$49	$60	89	0.0023		1.29		0.15	0.21	0.21
12	15	$69	$85	91	0.0023	2.17	1.31	4.34	0.17	0.27	0.26
13	3	$19	$25	53	0.005		5.62		0.18	0.29	0.28
14	10	$40	$55	62	0.005		5.52		0.13	0.23	0.22
15	2	$55	$75	61	0.005	2.60	5.40	21.93	0.28	0.43	0.41
16	18	$24	$35	29	0.013		114		0.29	0.36	0.35
17	21	$34	$50	30	0.013		108		0.42	0.55	0.53
18	25	$54	$80	30	0.013	2.31	118	636.25	0.59	0.67	0.65
19	5	$14	$25	19	0.03		68776		0.53	0.60	0.59
20	14	$27	$50	21	0.03		44796		0.61	0.73	0.71
21	23	$41	$75	20	0.03	2.33	61172	479086.87	0.75	0.79	0.78
22	7	$15	$35	13	0.07		21461293543		0.71	0.75	0.75
23	8	$25	$60	14	0.06		8178058828		0.86	0.90	0.90
24	19	$33	$80	14	0.07	2.67	10627887353	9894905759808.40	0.84	0.88	0.87
25	11	$11	$30	7	0.15		52496766684372100000000		0.91	0.94	0.94
26	27	$20	$55	7	0.16		80921121402027500000000		0.88	0.88	0.88
27	4	$31	$85	7	0.15		69431358425185700000000		0.90	0.93	0.93

Source: Author's construction but based on Kirby, Petry and Bickel (1999)

Notes: The "Order" column lists the order of each question as presented to subjects. SS refers to the smaller, sooner reward. LL refers to the larger, later reward. "Daily discount rate" and "Annual discount rate" give the parameter values at which the SS and LL amounts are of equal value as determined by the exponential discounting function. "Daily rate interval ratio" and "Annual rate interval ratio" provide the ratio of the right endpoint and the left endpoint of a discount rate interval.

[1]Delays are in days

[2]Non-smokers include both never-smokers and ex-smokers.

Table A:I shows that up to row 15, the KPB task can detect a 5-fold difference in annual discount rates. After row 15 though, the task cannot detect a 20-fold difference in annual discount rates. With reference to the data used in Chapter 3, Table A:I shows that on row 16, 71% of smokers and 64% of non-smokers selected the SS reward. In other words, more than half of the sample selected the SS reward on row 16 of Table A:I, which is where the KPB task's sensitivity to differences in annual discount rates starts to decrease. In addition, in every row of the table, except row 26, a larger proportion of smokers selected the SS reward than non-smokers. This raises the possibility that the precision of the annual discount rate estimates for smokers may be lower than non-smokers.

My aim in Chapter 3 was to compare the discounting behaviour of smokers and non-smokers and not to provide precise estimates of annual discount rates for use in public policy. Consequently, the second implication of the KPB task's sensitivity to certain choice patterns with annual discount rates (i.e., that it may be difficult to detect group differences with a large number of SS choices in certain rows) is more relevant to the inferences which I drew. Note that large annual rate intervals are irrelevant *if* one draws the same qualitative conclusions with daily and annual rates. In the remainder of this appendix, I will show that estimating, or inferring, *annual* discount rates with the KPB task can mask differences in the discounting behaviour of smokers and non-smokers.

To provide a concrete example of this inferential issue using the data reported in Chapter 3, Table A:II presents results from the exponential discounting model which only includes the current smoker covariate.[3] The table includes four sets of results: Model 1 was directly estimated with daily rates; Model 2 was directly estimated with annual rates which were then used to infer daily rates; Model 3 was directly estimated with annual rates; and Model 4 was directly estimated with daily rates which were then used to infer annual rates.

[3] In Chapter 3 I distinguished between current smokers, ex-smokers and never-smokers. To keep things as simple as possible in this appendix, I focus on current smokers and non-smokers, where non-smokers include both ex-smokers and never-smokers.

TABLE A:II: EXPONENTIAL DISCOUNTING ML ESTIMATES
SMOKING TOTAL EFFECT AT DAILY AND ANNUAL RATES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Daily	Inferred Daily	Annual	Inferred Annual
Discounting parameter (δ)				
Smoker	0.00573***	0.00573***	177.27630	177.27650
	(0.00167)	(0.00167)	(119.57760)	(119.57830)
Constant	0.00894***	0.00894***	24.73189***	24.73184***
	(0.00044)	(0.00044)	(4.07623)	(4.07622)
Error (μ)				
Constant	14.91274***	14.91275***	14.91275***	14.91274***
	(0.37902)	(0.37902)	(0.37902)	(0.37902)
N	32535	32535	32535	32535
log-likelihood	-14668.52	-14668.52	-14668.52	-14668.52

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Note that irrespective of whether I estimate with daily rates (Model 1) or estimate with annual rates and then infer daily rates (Model 2), I get identical results. Similarly, regardless of whether I estimate with annual rates (Model 3) or estimate with daily rates and then infer annual rates (Model 4), I get identical results. Moreover, the log-likelihoods of the models with daily and annual rates are the same, and so too are the estimates of the Fechner error term. However, the estimate for smokers with daily rates is positive and statistically significant ($p = 0.001$) whereas the estimate for smokers with annual rates is not statistically significant ($p = 0.138$). Thus, one reaches a different conclusion about the relationship between smoking and discounting when estimating with daily or annual rates. Which set of results reflect the "true" relationship between smoking and discounting?

As a first pass at answering this question note that the estimate of the constant term in Model 1 of Table A:II is 0.009 and the estimate of the smoker coefficient is 0.006. Looking at Table A:I, these daily rate estimates imply that the average non-smoker will switch to the LL reward on row 16 and the average smoker will switch to the LL reward on row 19, thereby potentially revealing a group difference in daily rates. However, with reference to annual rates, we know that the KPB task's sensitivity to differences in annual rates decreases after row 15 and this may be why I do not detect a group difference between smokers and non-smokers when estimating, or inferring, annual discount rates.

To provide a more detailed answer to this question I simulated KPB discounting data so that I control the DGP and know *a priori* whether there is a difference in the discounting behaviour of smokers and non-smokers. If it turns out that the estimates with daily rates reflect the underlying features of the data but the estimates with annual rates do not, then this suggests that one should estimate daily rates with the KPB task so as to draw accurate inferences.

For a simulation to shed light on this issue it should capture the salient features of the KPB data that I used in Chapter 3. Recall that 13.5% of the subjects in the sample were current smokers while the rest were either ex-smokers or never-smokers. As mentioned earlier, I want to keep things as simple as possible in this appendix so I will simply focus on current smokers and non-smokers, where non-smokers include both ex-smokers and never-smokers. Given the relatively low proportion of smokers in the sample, my simulation generated datasets which limited the proportion of smokers to a maximum of 25% of the sample.

Another important feature of the KPB data used in Chapter 3 is the relatively high prevalence of inconsistent choices. Recall that a person's choices on the KPB task are not always deterministically consistent with a particular discounting parameter. For example, a subject may choose the SS reward on the first 3 rows of Table A:I but then fail to consistently choose the LL reward thereafter. In other words, the subject may select the LL reward on rows 4, 5, and 6, before switching back to the SS reward on rows 7, 8, and 9, and then switching back to the LL reward for the remaining rows in the table. Note that for a subject's choices to be deterministically consistent with a particular discounting parameter on the KPB task, he can only switch from the SS reward to the LL reward once on the table, or not at all (i.e., either select the SS reward on every row of the table or the LL reward on every row of the table).

Figure A:I shows the distribution of the total number of switches which subjects made on the KPB task using the data reported in Chapter 3. Note that only 26% of the sample switched 0 or 1 times on the task, implying that only about a quarter of the sample made choices that were deterministically consistent with a particular discounting parameter. By contrast, 34% of the sample switched 3 times on the task and 20% of the sample switched 5 times on the task. As shown on the figure, the

mean number of switches on the task is 3.87, with a standard deviation of 3.09. Thus, my simulation should incorporate this feature of the data for it to mimic the process which generated the data.

Figure A:I: Distribution of the total number of switches on the KPB task

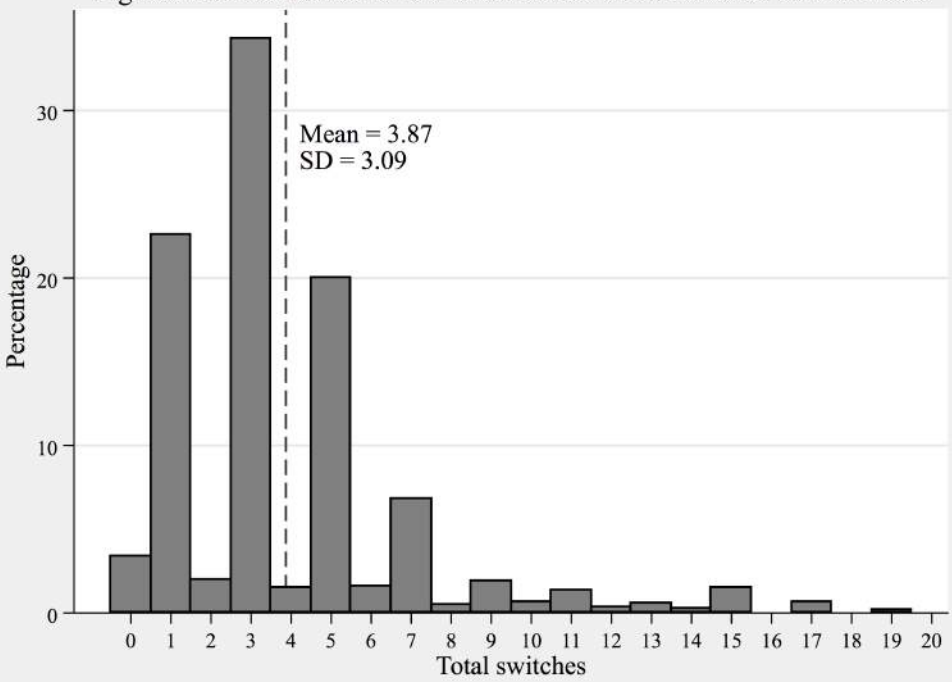

An important DGP design choice was whether daily or annual discount rates would be used to define the discounting behaviour of the sample. Taking draws from a distribution of daily discount rates may yield reliable daily rate estimates but unreliable annual rate estimates. Conversely, taking draws from a distribution of annual discount rates may provide reliable estimates of annual rates but unreliable estimates of daily rates. To determine whether daily rate estimates are more reliable than annual rate estimates with the KPB task, I conducted 75,000 replications where I took draws from distributions of daily discount rates, and 50,000 replications where I took draws from distributions of annual discount rates.

Thus, my simulation was set up to run 125,000 times and on each replication it performed a number of functions which will be detailed below. The first such task was the creation of a dataset consisting of 1,000 subjects who would make 27 choices on the KPB task. This sample was then randomly split into smokers and non-smokers

based on draws from a uniform distribution which ensured that the minimum proportion of smokers in the sample was 0.05 and the maximum was 0.25. Figure A:II is a kernel density plot of the fraction of smokers in the sample across the 125,000 replications in the simulated dataset. The average proportion of smokers in the sample across all replications is 0.150, with a standard deviation of 0.031.

Figure A:II: Fraction of smokers in the sample across replications

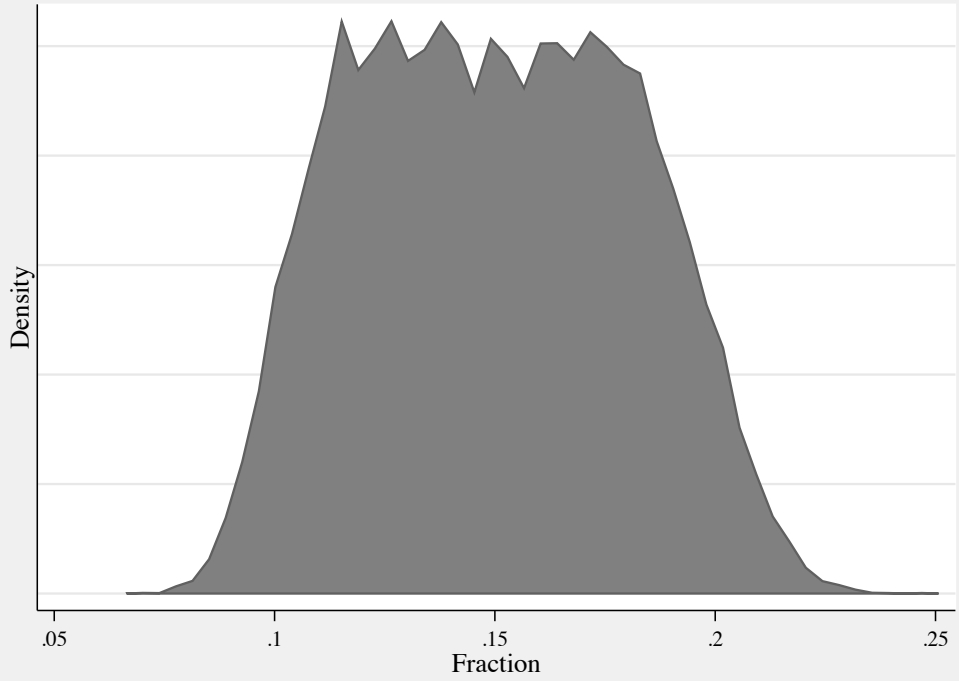

After randomly allocating subjects to smoking status, the simulation used draws from distinct gamma distributions to define each subject's discount rate. The gamma distribution $\Gamma(k, \theta)$ is a two-parameter family of continuous probability distributions which exhibits considerable flexibility. The distribution is defined by two parameters k and θ: k is the shape parameter, and θ is the scale parameter. The mean, mode and variance of the gamma distribution are simple functions of these two parameters. Specifically, the mean of a gamma distribution with parameters k and θ is $k\theta$, the mode is $(k - 1)\theta$ for $k > 1$, and the variance is $k\theta^2$. The support of the gamma distribution is $(0, \infty)$ which is relevant in the present context because I do not have to be concerned with the possibility of generating negative discount rates by taking draws from this distribution.

Figure A:III:A: Gamma PDFs for daily discount rates

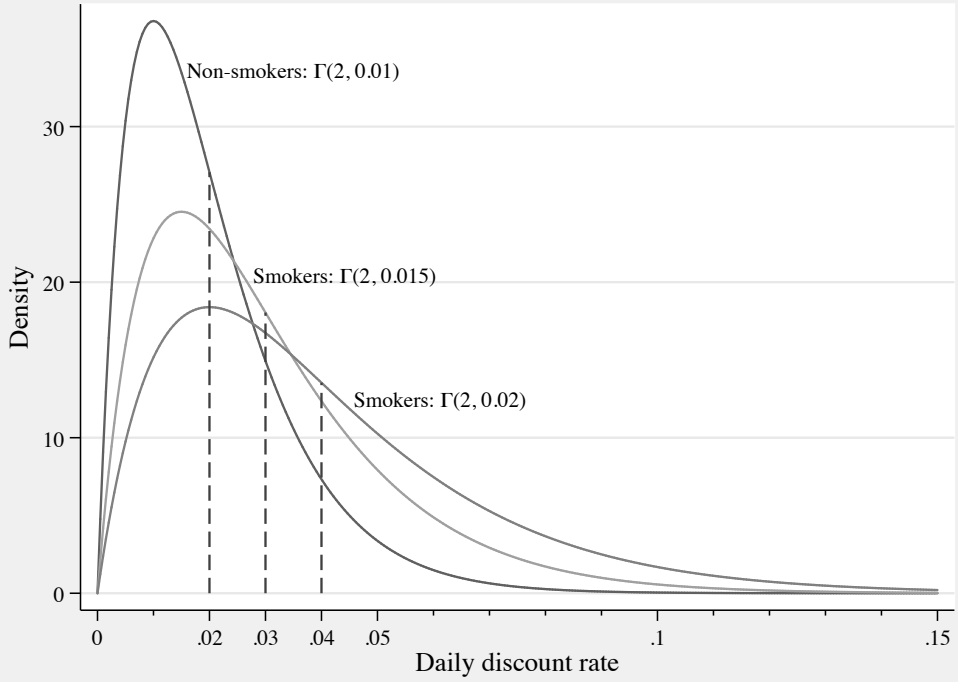

As mentioned previously, I conducted 75,000 replications where I took draws from distributions of daily discount rates, and 50,000 replications where I took draws from distributions of annual discount rates. Non-smokers' *daily* discount rates were determined by taking draws from Γ(2, 0.01) while smokers' *daily* discount rates were determined by taking draws from Γ(2, U[0.01, 0.02)) where U[0.01, 0.02) refers to a continuous uniform distribution on the interval [0.01, 0.02).[4] Figure A:III:A plots the gamma probability density function (PDF) from which non-smokers' *daily* discount rates were drawn (i.e., Γ(2, 0.01)), and two example gamma PDFs (i.e., Γ(2, 0.015), Γ(2, 0.02)) from which smokers' *daily* discount rates were drawn.[5] Focussing on the non-smokers' gamma PDF, we see that the mode of the distribution is $(k - 1)\theta = 0.01$, the mean is $k\theta = 0.02$, and the variance is $k\theta^2 = 0.0002$; the means of the gamma PDFs are represented by the dashed lines on the figure. By contrast, one of the example gamma PDFs for smokers (i.e., Γ(2, 0.015)) has a mode of $(k - 1)\theta = 0.015$,

[4] Note that the draw from the uniform distribution (e.g., 0.015) was taken prior to the draws from the gamma distribution, which means that, during one replication of the simulation, I was taking draws from the same gamma distribution (e.g., Γ(2, 0.015)).

[5] Technically, smokers' daily discount rates were never drawn from Γ(2, 0.02) because draws from Γ(2, U[0.01, 0.02)) preclude a scale parameter θ of 0.02. In other words, draws from the left-closed, right-open interval [0.01, 0.02) do not include 0.02; an interval which is closed and open is referred to as a "clopen" interval. However, I plot the gamma PDF Γ(2, 0.02) because this incorporates the scale parameter's supremum for distributions from which smokers' daily discount rates were drawn.

a mean of $k\theta = 0.03$, and variance of $k\theta^2 = 0.00045$. Finally, the other example gamma PDF for smokers (i.e., $\Gamma(2, 0.02)$) has a mode of $(k - 1)\theta = 0.02$, a mean of $k\theta = 0.04$, and variance of $k\theta^2 = 0.0008$.

The draws from these two gamma distributions ensured that, in almost all cases, smokers had higher average *daily* discount rates than non-smokers. Figure A:IV:A shows the distributions of smokers' and non-smokers' average daily discount rates across the 75,000 replications where I took draws from daily discount rate distributions.

Figure A:IV:A: Distribution of average daily discount rates across replications

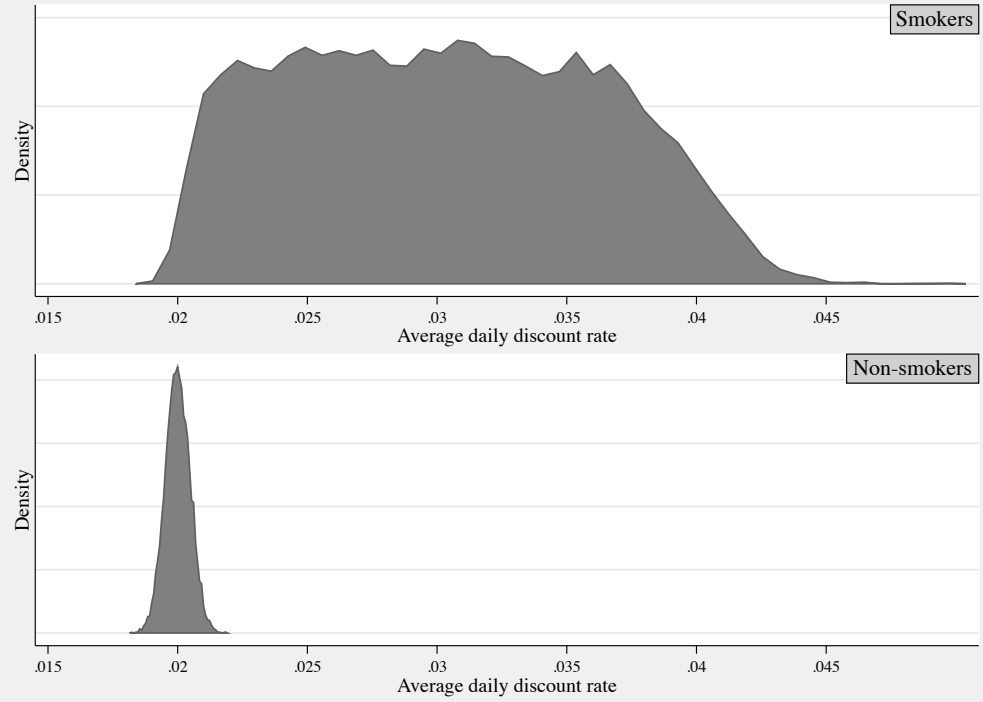

For the 50,000 replications where I took draws from annual discount rate distributions, non-smokers' *annual* discount rates were determined by taking draws from $\Gamma(2, 55)$ while smokers' *annual* discount rates were determined by taking draws from $\Gamma(2, U[55, 110))$ where $U[55, 110)$ refers to a continuous uniform distribution on the interval $[55, 110)$.[6] Note that I chose to use the same shape parameter for the gamma distributions from which daily and annual discount rates were drawn. In

[6] The draw from the uniform distribution (e.g., 80) was taken prior to the draws from the gamma distribution, which means that, during one replication of the simulation, I was taking draws from the same gamma distribution (e.g., $\Gamma(2, 80)$).

addition, across both daily and annual rate simulations, the uniform distribution which determined the scale parameter for smokers used the non-smokers' scale parameter as the left endpoint and twice this value for the right endpoint. Finally, the daily rate scale parameter for non-smokers implies that the modal non-smoker would switch to the LL reward on row 16 of Table A:I. Similarly, the annual rate scale parameter for non-smokers implies that the modal non-smoker would switch to the LL reward on row 16 of Table A:I. Thus, although discount rates were drawn from distinct daily and annual rate distributions, these daily and annual rate distributions yield similar choice patterns on the KPB task.

Figure A:III:B: Gamma PDFs for annual discount rates

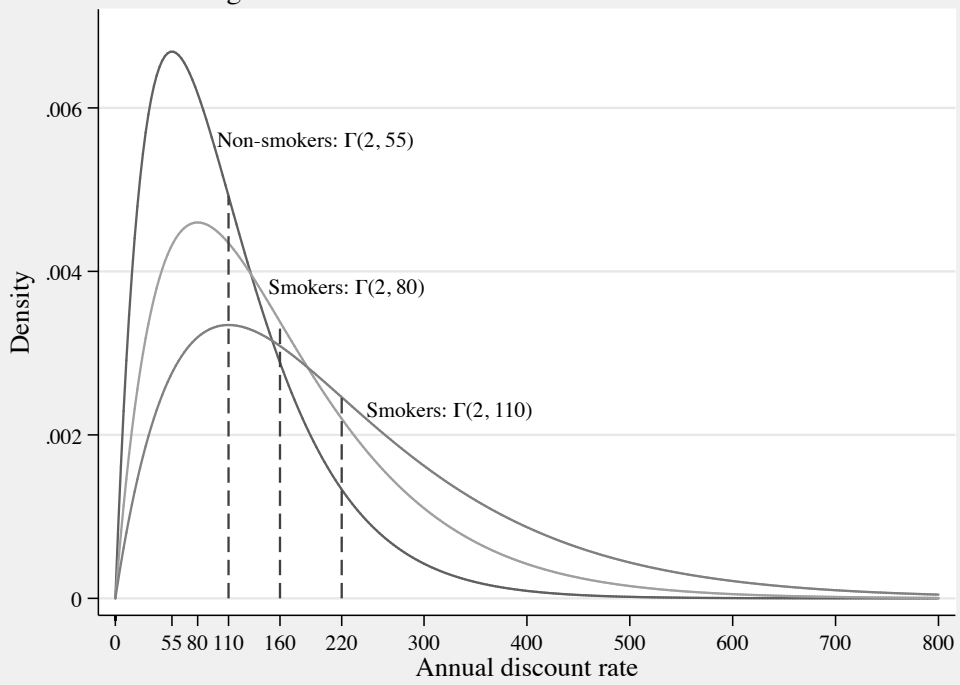

Figure A:III:B plots the gamma PDF from which non-smokers' *annual* discount rates were drawn (i.e., Γ(2, 55)), and two example gamma PDFs (i.e., Γ(2, 80), Γ(2, 110)) from which smokers' *annual* discount rates were drawn.[7] Focussing on the non-smokers' gamma PDF, we see that the mode of the distribution is $(k - 1)\theta = 55$, the mean is $k\theta = 110$, and the variance is $k\theta^2 = 6{,}050$; the means of the gamma PDFs are

[7] Technically, smokers' annual discount rates were never drawn from Γ(2, 110) because draws from Γ(2, U[55, 110)) preclude a scale parameter θ of 110. In other words, draws from the clopen interval [55, 110) do not include 110. However, I plot the gamma PDF Γ(2, 110) because this incorporates the scale parameter's supremum for distributions from which smokers' annual discount rates were drawn.

represented by the dashed lines on the figure. By contrast, one of the example gamma PDFs for smokers (i.e., $\Gamma(2, 80)$) has a mode of $(k - 1)\theta = 80$, a mean of $k\theta = 160$, and variance of $k\theta^2 = 12{,}800$. Finally, the other example gamma PDF for smokers (i.e., $\Gamma(2, 110)$) has a mode of $(k - 1)\theta = 110$, a mean of $k\theta = 220$, and variance of $k\theta^2 = 24{,}200$.

The draws from these two gamma distributions ensured that, in almost all cases, smokers had higher average *annual* discount rates than non-smokers. Figure A:IV:B shows the distributions of smokers' and non-smokers' average annual discount rates across the 50,000 replications where I took draws from annual discount rate distributions.

Figure A:IV:B: Distribution of average annual discount rates across replications

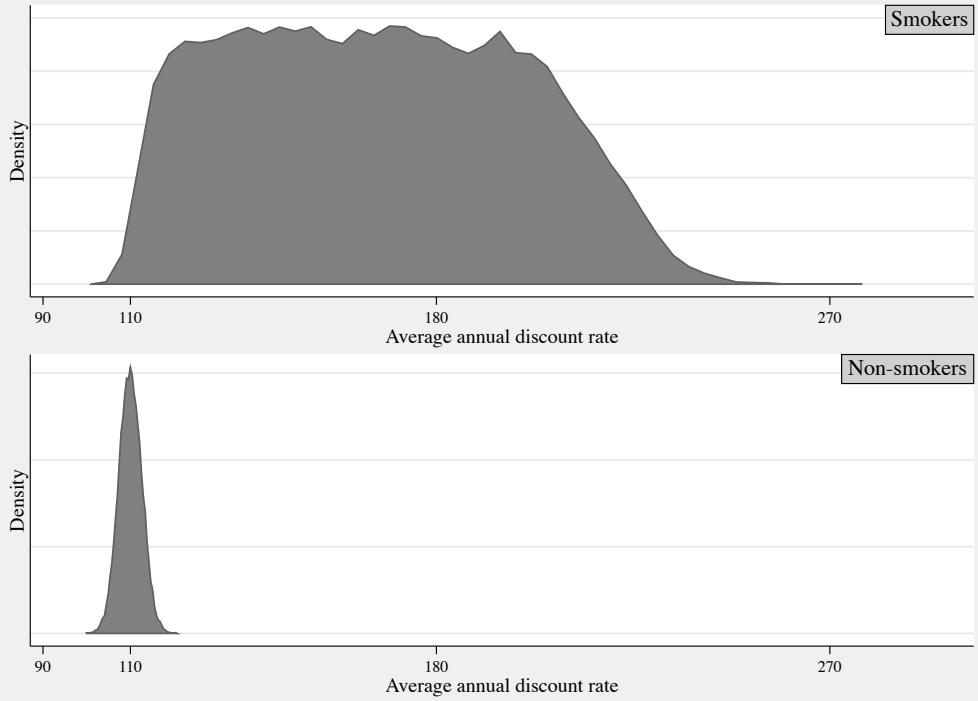

Once each subject had been assigned a daily or annual discount rate, a draw from a gamma distribution $\Gamma(U[0.1, 3), U[0.01, 6))$ determined the Fechner error term for that subject, where $U[\cdot,\cdot)$, once again, refers to a continuous uniform distribution on

the clopen interval $[\cdot,\cdot)$.[8] With a daily or annual discount rate and a Fechner error term, each subject's choice behaviour could then be determined.

Recall that according to the exponential discounting model, under the assumptions of linear utility and a Fechner behavioural error specification, a subject compares the present values (PVs) of SS and LL rewards (where the present values are computed using the subject's discount rate) and then *stochastically* chooses the one with the higher present value. Specifically, using the notation in Chapter 3, a subject computes the PV of the SS reward (which is always available immediately in the KPB task):

$$PV_{SS} = y_0, \tag{2}$$

and the PV of the LL reward[9]:

$$PV_{LL} = [1 / (1 + \delta)^{\tau}]y_{\tau}, \tag{3}$$

and then forms the PV index:

$$\nabla PV = (PV_{SS} - PV_{LL}) / \mu, \tag{4}$$

where μ is the Fechner error term. The value of the PV index, when passed through a cumulative distribution function, then *stochastically* determines the subject's choice. Specifically, in the case of the cumulative logistic distribution function, the probability that the subject chooses the SS reward is given by the following expression:

$$\Pr(\text{Choose SS reward}) = \Lambda(\nabla PV) \tag{5}$$

[8] Note that the draws from the two uniform distributions (U[0.1, 3) and U[0.01, 6)) were taken prior to the draws from the gamma distribution, which means that, during one replication of the simulation, I was taking draws from the same gamma distribution.

[9] Under the assumption of daily discount rates, τ in expression (3) is simply the number of days (e.g., 119 days) between receipt of the SS reward and receipt of the LL reward. Under the assumption of annual discount rates, by contrast, τ is divided by 365 in expression (3) so that it represents the delay between the SS and LL rewards as a fraction of a year (e.g., 119/365 = 0.326 years).

For example, if ∇PV = 2, then the likelihood of the SS reward being chosen is Λ(2) = 0.88 (i.e., there is an 88% chance that the SS reward will be chosen). To operationalise this stochastic choice procedure in the simulation, I took draws from U[0, 1). If the draw from this uniform distribution was below the choice likelihood Λ(∇PV) then the choice was implemented, but if the draw from this distribution was above the choice likelihood, then the opposite choice was made. Expanding on the example where ∇PV = 2, if the draw from the uniform distribution for this particular choice pair was, say, 0.75 then the SS reward was chosen because Λ(2) = 0.88, but if the draw from the uniform distribution was 0.9 then the LL reward was chosen.

Figure A:V: Average number of switches per subject across replications

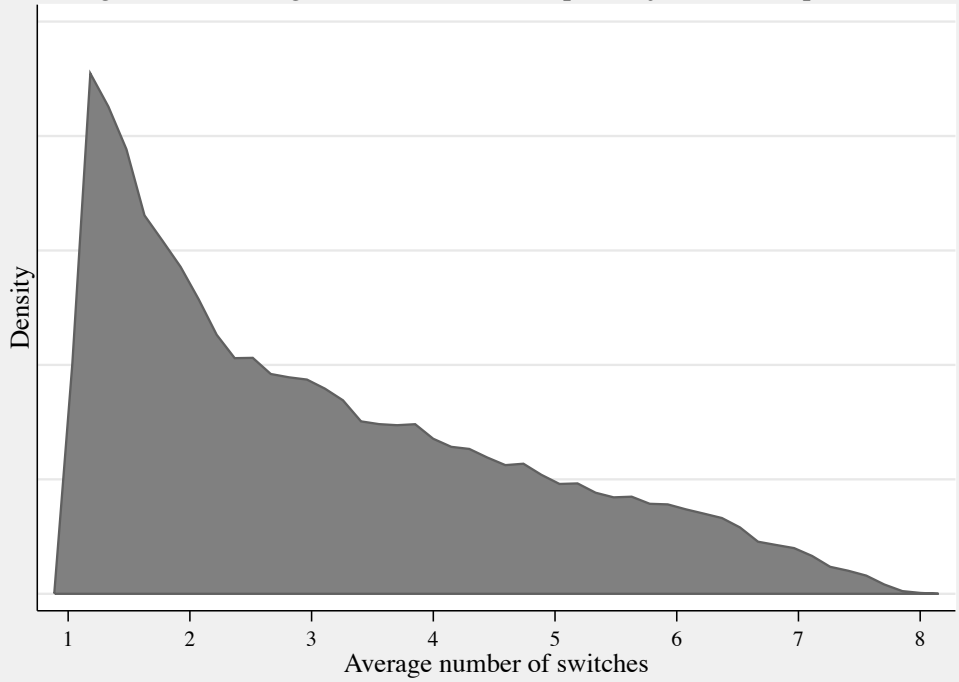

Note that by taking draws for the Fechner error term from different gamma distributions on each replication (e.g., draws from Γ(0.1, 0.001) during one replication and draws from Γ(2, 5) during another replication), I was able to vary the level of inconsistent choice behaviour exhibited by the sample for that replication. From (4) it is clear that as μ increases, ∇PV falls. In the limit, as $\mu \rightarrow \infty$, $\nabla PV \rightarrow 0$, and choice behaviour is essentially random (i.e., Λ(∇PV = 0) = 0.5 so there is a 50% chance that the SS reward is chosen and a 50% chance that the LL reward is chosen). Figure A:V

is a kernel density plot of the average number of switches per subject across replications of the simulation; it resembles the general shape of Figure A:I.

Once the simulation had generated the choice data in the manner described above, I estimated the exponential discounting model with a Fechner error term and smoking status covariate assuming daily and annual discount rates. I also used these estimates to infer annual and daily rates, respectively. The simulation then saved the results from these estimations before starting up again. This process was repeated 125,000 times: 75,000 times under the assumption that daily discount rates generated the data, and 50,000 times under the assumption that annual discount rates generated the data. In what follows, I will discuss the results from the daily rate DGP prior to discussing the results from the annual rate DGP.

Daily Discount Rate DGP

Figure A:VI plots estimates of the Fechner error term as a function of the average number of switches per subject across the 75,000 replications where daily discount rates define the DGP. The figure also includes a kernel-weighted local polynomial regression "line" which clearly shows that as the average number of switches per subject increases, so too does the estimate of the Fechner error term. Thus, estimates of the Fechner error term closely track the level of inconsistency exhibited by the sample.

As I found in Table A:II, the daily estimates and inferred daily estimates were practically identical during each replication and so too were the annual estimates and inferred annual estimates. In addition, estimates of the Fechner error term across daily and annual rates were the same, and so too were the models' log-likelihoods. Pronounced differences emerged though when I focussed on the statistical significance of the smoker covariate across daily and annual rates.

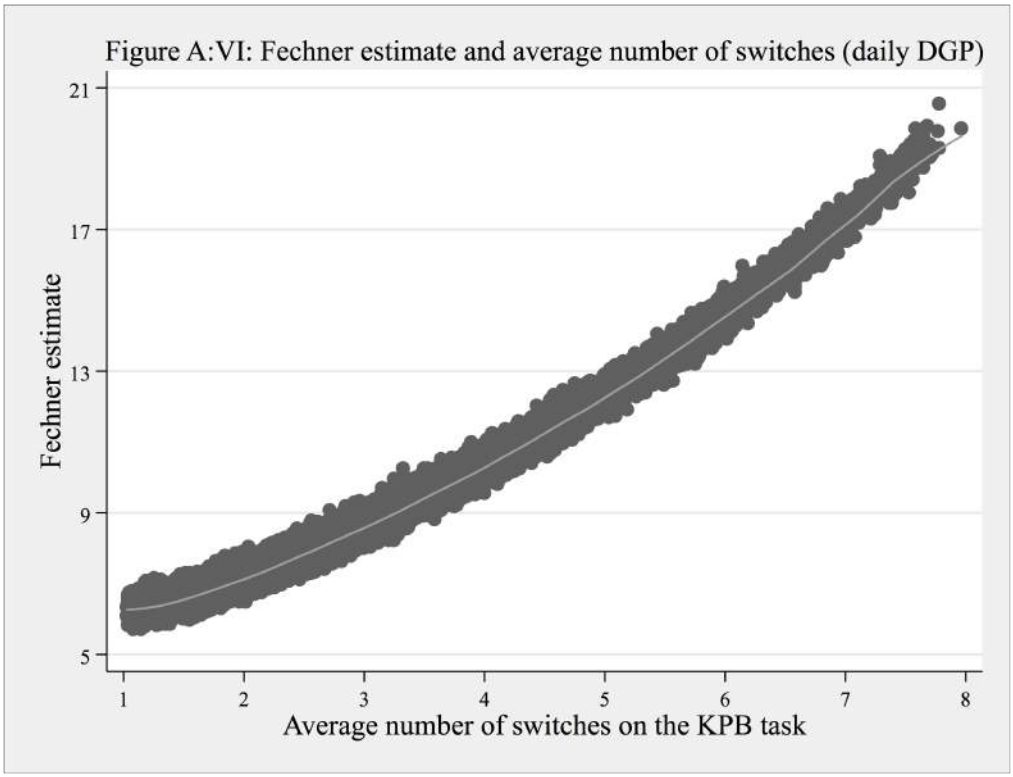

Figure A:VII plots the *p*-values for the smoker covariate when using annual and inferred annual rates, daily and inferred daily rates, and annual and daily rates. Note that the horizontal and vertical dashed lines are plotted at $p = 0.1$, which is the level of statistical significance used in this thesis. Note further that the figure includes a kernel-weighted local polynomial regression "line." In the top panel of Figure A:VII, the *p*-values for the smoker covariate when estimating with annual rates, or estimating with daily rates and inferring annual rates, track each other closely. Similarly, in the middle panel of the figure, the *p*-values for the smoker covariate when estimating with daily rates, or estimating with annual rates and inferring daily rates, track each other almost perfectly. However, in the bottom panel of the figure, there is a marked divergence in the smoker covariate *p*-values when estimating with daily rates or estimating with annual rates. Given the small size of the figure in the bottom panel, this divergence can be quite difficult to see so Figure A:VIII provides a close-up view of the relationship between annual and daily *p*-values.

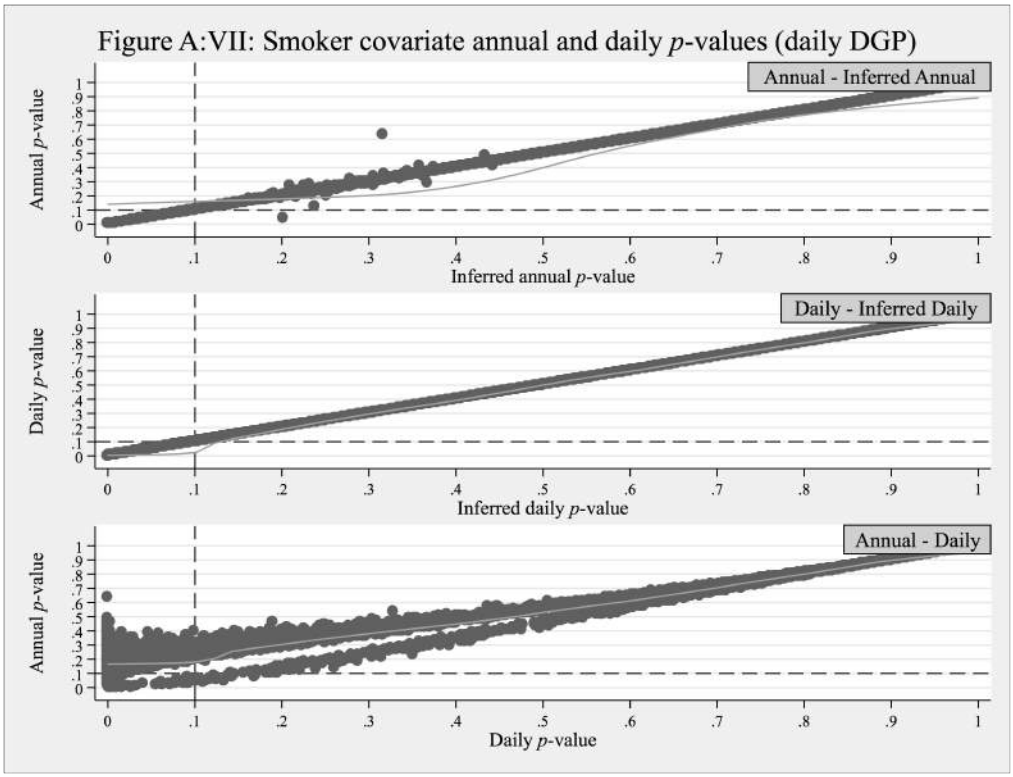

Figure A:VIII plots the smoker covariate annual *p*-values as a function of the smoker covariate daily *p*-values, for daily *p*-values which are statistically significant (i.e., $p < 0.1$). If the smoker covariate was statistically significant with both annual and daily discount rates, then the annual *p*-values would all lie below the horizontal dashed line where the annual $p = 0.1$. However, only 15% of the annual *p*-values are less than 0.1 when the associated daily *p*-values are less than 0.1. Furthermore, the average annual *p*-value = 0.167 (standard deviation = 0.065) for daily *p*-values which are strictly less than 0.1. The kernel-weighted local polynomial regression line in Figure A:VIII highlights this fact and shows that the average annual *p*-value is not statistically significant when the daily *p*-value is statistically significant.

Thus, in only 15% of cases would one reach the same qualitative conclusion about the effect of smoking on discounting behaviour when using annual discount rates as opposed to daily discount rates. However, this does not prove that the estimates with daily rates reveal the correct relationship between smoking and discounting while the estimates with annual rates do not. To determine which set of estimates uncovers the true relationship, it is helpful to focus on the smoker covariate daily and annual *p*-

values as a function of the difference in simulated average *daily* discount rates of smokers and non-smokers.

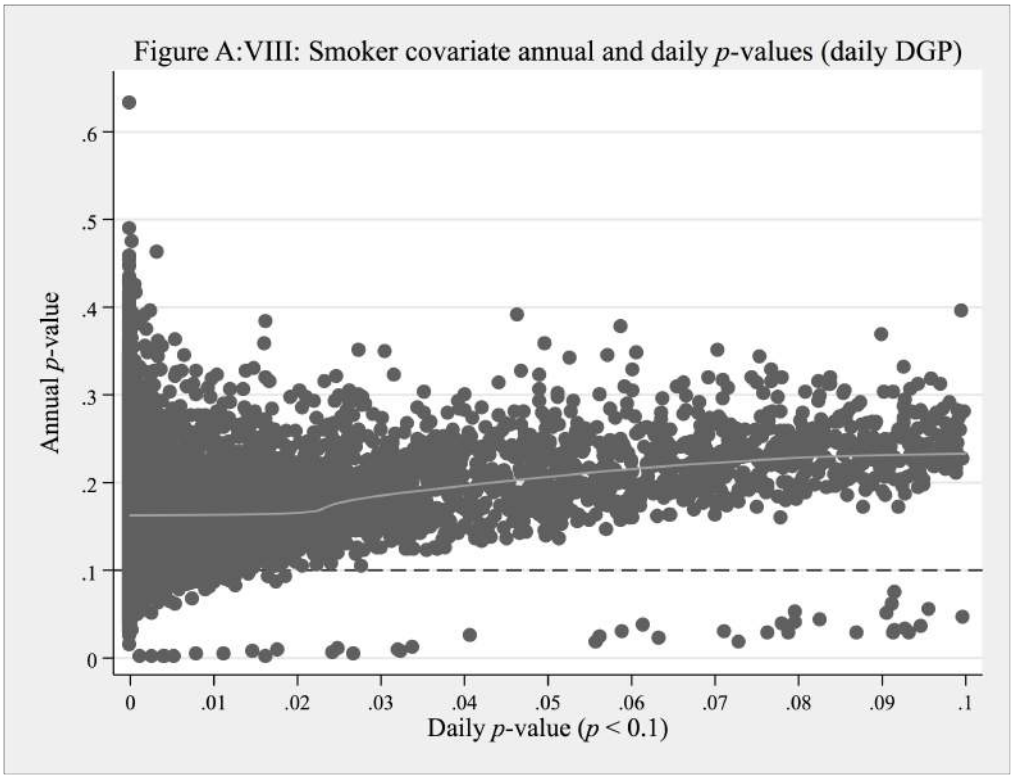

Figure A:IX plots the smoker covariate daily and annual *p*-values as a function of the difference in simulated average *daily* discount rates of smokers and non-smokers. In the top panel of the figure, I find that the smoker covariate daily *p*-value converges to 0 as the difference in smokers' and non-smokers' average discount rates increases; this is precisely what one would expect if the daily rate estimates track the difference in discounting behaviour of smokers and non-smokers. By contrast, the bottom panel of the figure shows a U-shaped relationship between the smoker covariate annual *p*-value and the difference in average *daily* discount rates of smokers and non-smokers. The kernel-weighted local polynomial regression line in the bottom panel never actually crosses the horizontal dashed line where $p = 0.1$, and it gets further from the dashed line as the difference in average discount rates of smokers and non-smokers increases. Thus, Figure A:IX shows that the daily rate estimates reliably track the difference in average *daily* discount rates of smokers and non-smokers whereas the annual rate estimates do not.

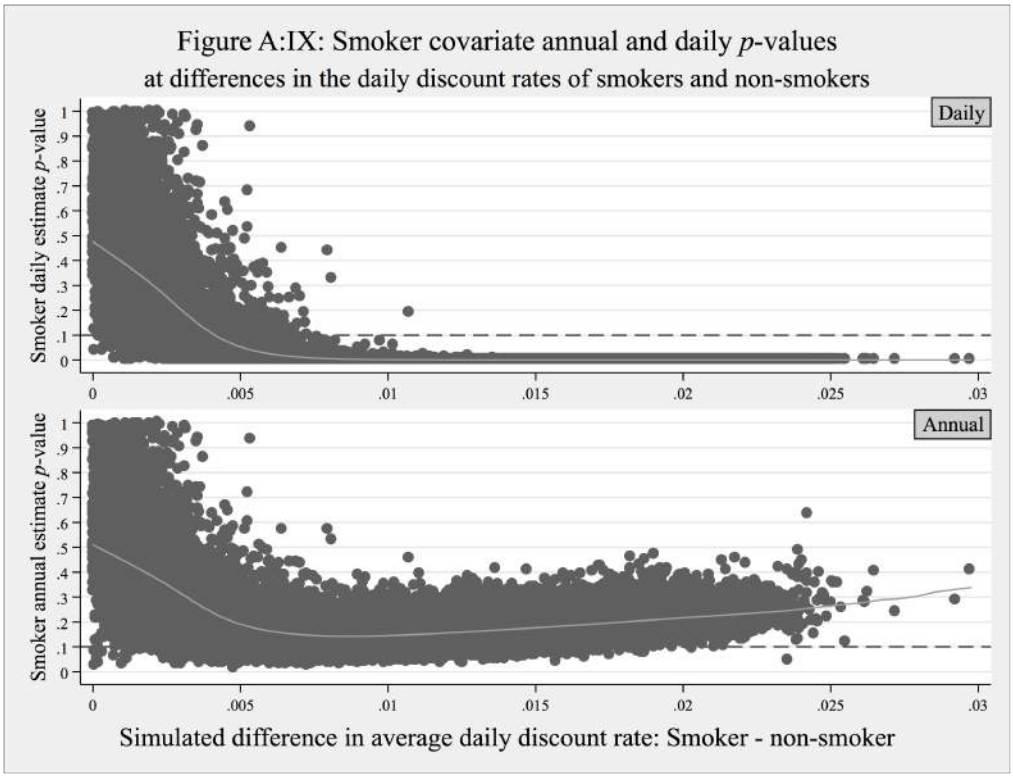

Figure A:IX: Smoker covariate annual and daily *p*-values at differences in the daily discount rates of smokers and non-smokers

To further investigate this issue, I conducted some parametric analyses of the relationship between smoker covariate *p*-values and features of the simulation; these complement the non-parametric analyses. The dependent variables in these analyses are the smoker covariate daily and annual *p*-values. As *p*-values fall within the closed unit interval [0, 1], I make use of Papke and Wooldridge's (1996) model for fractional response variables which has the log-likelihood function:

$$\ln L_i(\beta; y, X) = y_i \ln[\Phi(X_i\beta)] + (1 - y_i)\ln[1 - \Phi(X_i\beta)], \qquad (6)$$

where y is the fractional dependent variable of the model (i.e., the smoker covariate annual or daily *p*-value), X is a vector of covariates, β is the coefficient vector linked to the vector of covariates, and Φ is the standard cumulative normal distribution function.

Table A:III presents results from the fractional dependent variable model where the smoker covariate annual (Model 1) and daily (Model 2) *p*-values are estimated as a function of the difference in average *daily* discount rates of smokers and non-smokers, the average number of switches on the KPB task, and the fraction of

smokers in the sample. In both models, I find that *p*-values tend to increase as the level of inconsistency (i.e., the average number of switches on the KPB task) rises, and *p*-values tend to decrease as the fraction of smokers in the sample increases.

However, the relationship between smoker covariate *p*-values and differences in the average *daily* discount rates of smokers and non-smokers, differs across the two models, which replicates the result in Figure A:IX. Specifically, in the annual model, there is a U-shaped relationship between the smoker covariate *p*-value and differences in the average *daily* discount rates of smokers and non-smokers. In other words, annual *p*-values decrease as the difference in discount rates increase but only up to a point, after which increases in the difference in discount rates, lead to increases in annual *p*-values. In the daily model, by contrast, increases in the difference in discount rates of smokers and non-smokers lead to a monotonic decline in the smoker covariate *p*-value.

TABLE A:III: SMOKER ESTIMATE STATISTICAL SIGNIFICANCE (DAILY DGP)
PAPKE-WOOLDRIDGE FRACTIONAL RESPONSE MODEL

	Model 1	**Model 2**
	Annual	Daily
Difference in daily discount rate: Smoker - non-smoker	-177.217***	-250.353***
	(2.830)	(10.246)
(Difference in daily discount rate: Smoker - non-smoker)^2	7111.990***	-20392.365***
	(133.293)	(1801.042)
Average number of switches on the KPB task	0.041***	0.039***
	(0.003)	(0.006)
Fraction of smokers in the sample	-4.836***	-3.151***
	(0.172)	(0.315)
Constant	0.584***	0.445***
	(0.030)	(0.052)
N	75000	75000
log-likelihood	-37015.382	-10718.75

Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table A:IV maps out the response surface for daily and annual smoker covariate *p*-values at different values of the difference in average *daily* discount rates of smokers and non-smokers. Note that when the difference in discount rates is low, the conditional marginal effect of an increase in the difference is negative in both models. However, as the difference in discount rates reaches 0.015, the conditional marginal effect in the annual model becomes positive, while the conditional marginal effect in the daily model remains negative. Thus, the daily estimates reliably track the

underlying differences in *daily* discounting behaviour of smokers and non-smokers, whereas the annual estimates do not.

TABLE A:IV: CONDITIONAL MARGINAL EFFECTS FOR P-VALUES (DAILY DGP)

	Model 1	**Model 2**
	Annual	Daily
Difference in discount rates		
0.000	-177.217 (2.830)	-250.353 (10.246)
0.005	-106.097 (1.600)	-454.277 (10.280)
0.010	-34.977 (0.818)	-658.201 (27.479)
0.015	36.143 (1.527)	-862.124 (45.313)
0.020	107.262 (2.747)	-1066.048 (63.248)
0.025	178.382 (4.039)	-1269.971 (81.215)

Standard errors in parentheses

In sum, the preceding analyses show that when daily discount rates define the DGP of the KPB task, daily rates provide more reliable estimates than annual rates when drawing inferences about the relationship between smoking and discounting behaviour. However, this does not imply that one should estimate daily discount rates instead of annual discount rates when the DGP of the KPB task is defined by annual rates. I explore the relationship between smoking and discounting behaviour under the assumption of an annual discount rate DGP in the next section.

Annual Discount Rate DGP

Figure A:X plots estimates of the Fechner error term as a function of the average number of switches per subject for both the daily and annual DGP simulations. The figure also includes kernel-weighted local polynomial regression lines which clearly show that as the average number of switches per subject increases, so too does the estimate of the Fechner error term. Thus, estimates of the Fechner error term closely track the level of inconsistency exhibited by the sample. However, estimates of the Fechner error term with the daily DGP lie everywhere above estimates of the Fechner error term with the annual DGP. A non-parametric Wilcoxon rank-sum test confirms that estimates of the Fechner error term with the daily DGP are significantly higher than estimates of the Fechner error term with the annual DGP ($p < 0.001$). It is unclear why this difference in Fechner estimates emerges with daily and annual DGPs but it is certainly a noteworthy difference.

As I found in Table A:II and under the daily discount rate DGP assumption, the daily estimates and inferred daily estimates were practically identical during each replication of the annual rate DGP simulation, and so too were the annual estimates and inferred annual estimates. In addition, estimates of the Fechner error term across daily and annual rates were the same, and so too were the models' log-likelihoods. Of crucial importance now is whether I replicate the result of differences in the statistical significance of the smoker covariate across daily and annual rates.

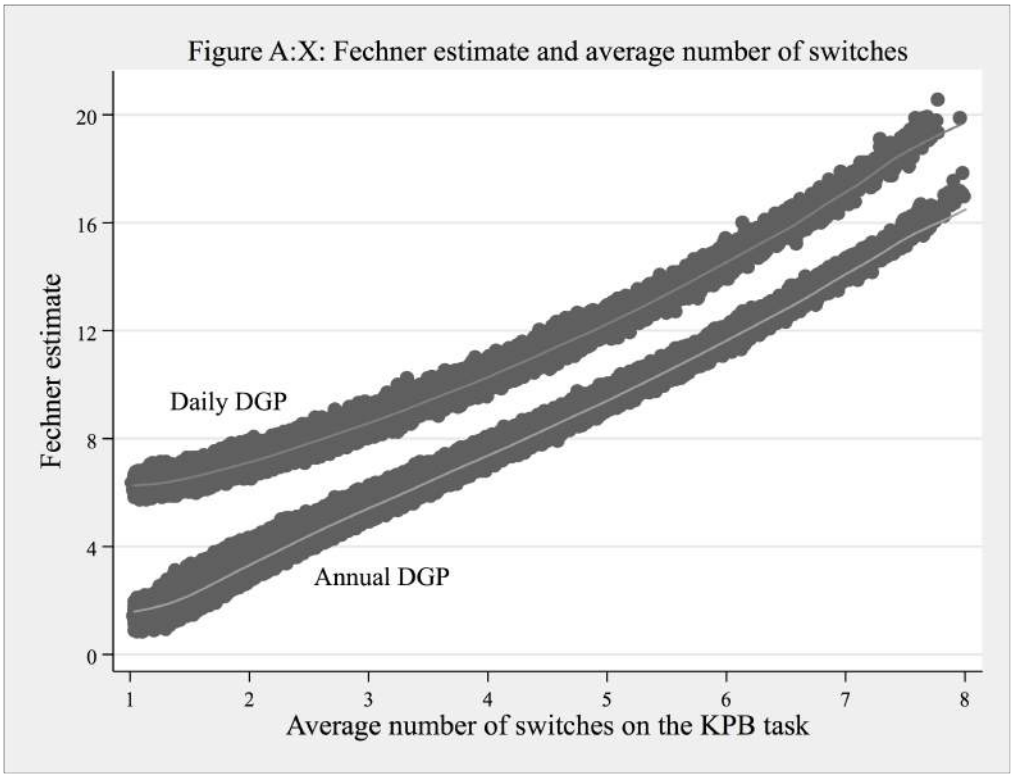

Figure A:XI plots the *p*-values for the smoker covariate when using annual and inferred annual rates, daily and inferred daily rates, and annual and daily rates. The horizontal and vertical dashed lines are plotted at $p = 0.1$ and the figure includes a kernel-weighted local polynomial regression line. In the top panel of Figure A:XI, the *p*-values for the smoker covariate, when estimating with annual rates or estimating with daily rates and inferring annual rates, track each other very closely. Similarly, in the middle panel of the figure, the *p*-values for the smoker covariate when estimating with daily rates or estimating with annual rates and inferring daily rates, track each other almost perfectly. However, in the bottom panel of the figure, there is a divergence, although not nearly as pronounced as with the daily DGP in Figure

A:VIII, in the smoker covariate *p*-values when estimating with daily rates or estimating with annual rates.

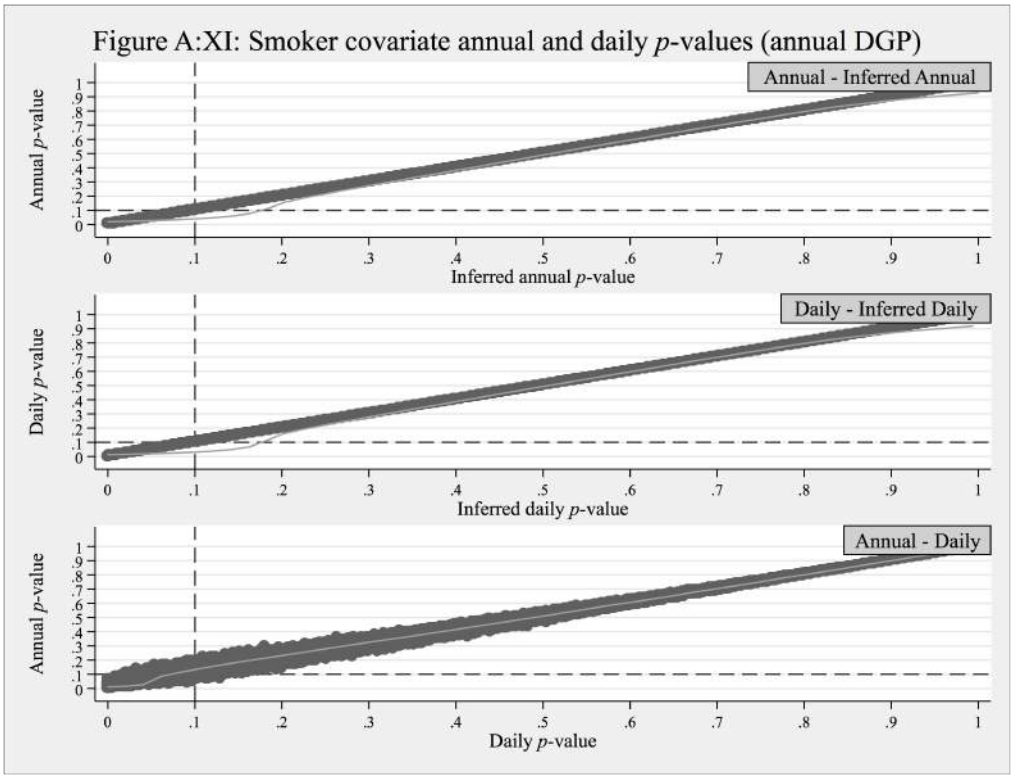

To focus on this divergence in statistical significance, Figure A:XII plots the smoker covariate annual *p*-values as a function of the smoker covariate daily *p*-values where the daily $p < 0.1$. If the smoker covariate was statistically significant with both annual and daily discount rates, then the annual *p*-values would all lie below the horizontal dashed line where the annual $p = 0.1$. Figure A:XII shows that for low daily *p*-values (i.e., $p < 0.05$), annual *p*-values are predominately statistically significant, but at higher daily *p*-values (i.e., $0.1 > p > 0.05$) a lower proportion of annual *p*-values are statistically significant. Specifically, when daily $p < 0.05$, 99% of annual *p*-values are less than 0.1. However, when daily $0.05 < p < 0.1$, only 31% of annual *p*-values are less than 0.1. Clearly these differences in statistical significance across daily and annual rates are not nearly as pronounced under the assumption of an annual DGP as they were under the assumption of a daily DGP. But they raise the question: which set of estimates (daily or annual) more reliably track differences in discounting behaviour?

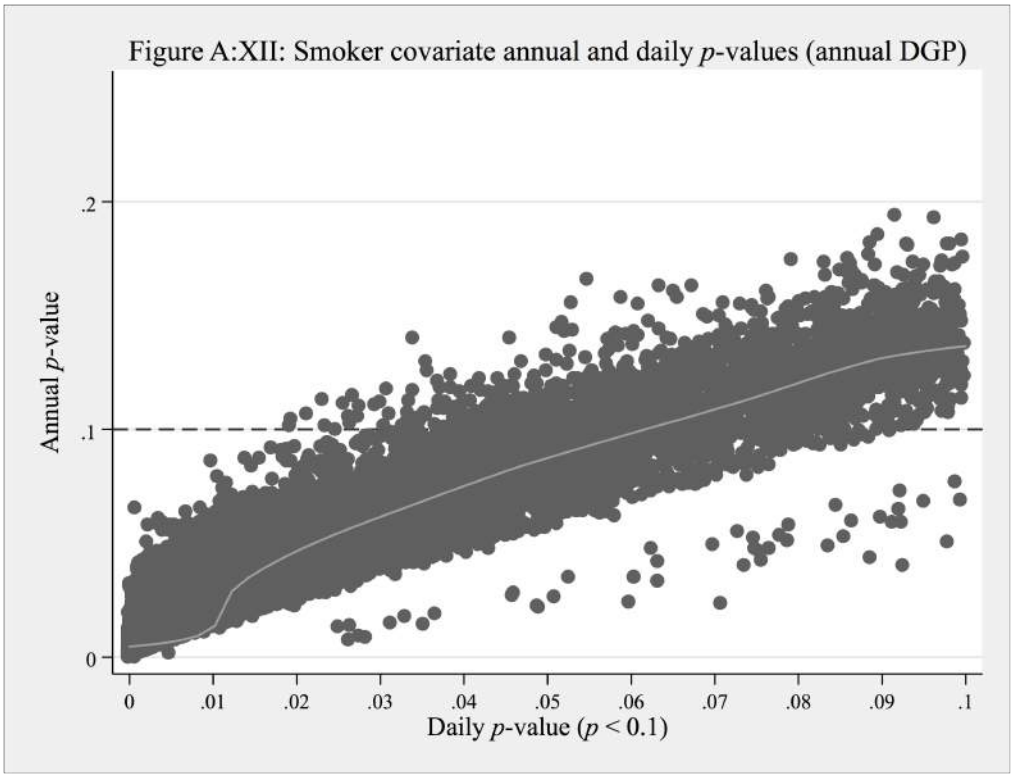

As discussed in the previous section, the fact that the daily rate *p*-values tend to be lower than the annual rate *p*-values does not prove that the estimates with daily rates reveal the correct relationship between smoking and discounting while the estimates with annual rates do not. To determine which set of estimates uncovers the true relationship, it is helpful to focus on the smoker covariate daily and annual *p*-values as a function of the difference in simulated average *annual* discount rates of smokers and non-smokers.

Figure A:XIII plots the smoker covariate daily and annual *p*-values as a function of the difference in simulated average *annual* discount rates of smokers and non-smokers. The kernel-weighted local polynomial regression lines in the two panels show that the smoker covariate daily and annual *p*-values tend to 0 as the difference in simulated average annual discount rates increases; this is what one would expect if the estimates track the underlying relationship between smoking and discounting.

However, unlike smoker covariate daily *p*-values in Figure A:IX (i.e., under the assumption of a daily rate DGP), there is far more variation in daily and annual *p*-values as the difference in simulated discount rates increases. In other words, even at

large simulated differences in *annual* discount rates, a fairly high proportion of smoker covariate *p*-values are not statistically significant. Specifically, 6% of daily smoker covariate *p*-values are not statistically significant when the difference in average annual discount rates is greater than 55, whereas 10% of annual smoker covariate *p*-values are not statistically significant when the difference in average annual discount rates is greater than 55. With the daily DGP simulations, only 0.001% of smoker covariate daily *p*-values were not statistically significant when the difference in simulated *daily* discount rates was greater than 0.01.[10]

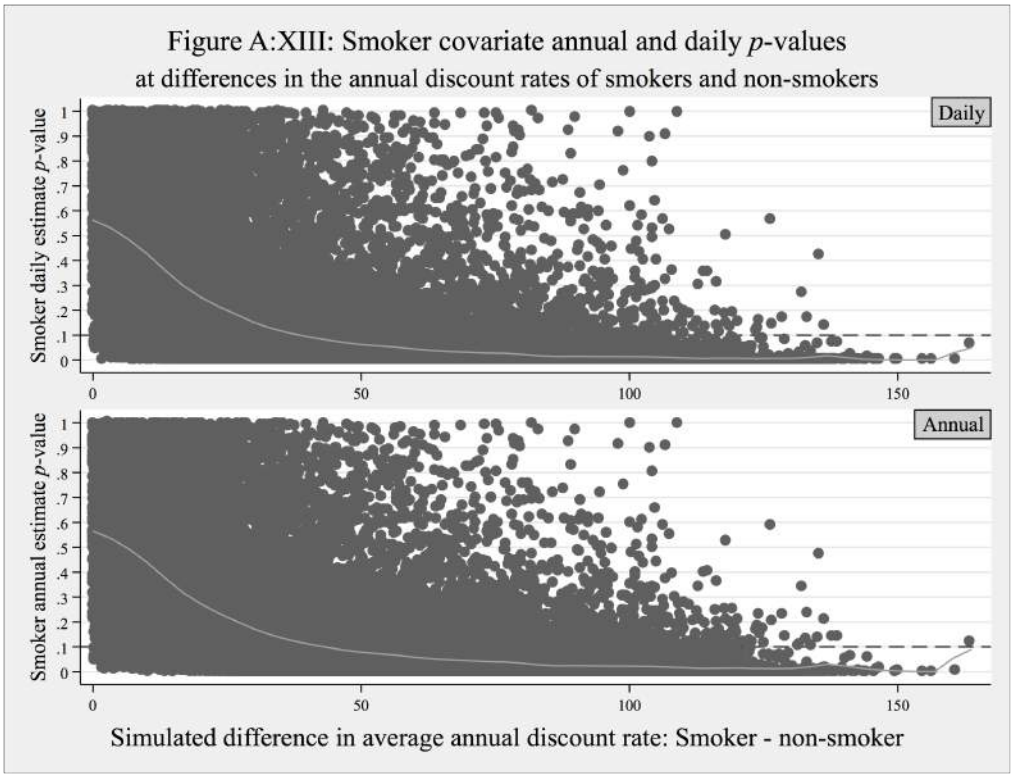

Table A:I provides a possible explanation for this difference in the KPB task's sensitivity to differences in the discount rates of smokers and non-smokers across the daily and annual rate DGP simulations: while the KPB task can always detect a 3-fold difference in daily discount rates, it cannot detect a 20-fold difference in annual

[10] In the daily rate DGP simulations, the average non-smoker daily discount rate was 0.02 so an average *difference* in the daily discount rates of smokers and non-smokers of 0.01 means that, on average, smokers' daily discount rates were 50% greater than non-smokers' daily discount rates. In the annual rate DGP simulations, the average non-smoker annual discount rate was 110 so an average *difference* in annual discount rates of smokers and non-smokers of 55 means that, on average, smokers' annual discount rates were 50% greater than non-smokers' annual discount rates. Thus, a difference of 0.01 in the average daily discount rates of smokers and non-smokers in the daily DGP simulations, is directly comparable to a difference of 55 in the average annual discount rates of smokers and non-smokers in the annual rate DGP simulations.

discount rates after row 15 of the table. This is relevant to the data used in Chapter 3 because more than half of all smokers and non-smokers selected the SS reward on row 16 of the KPB task, and this is where the task's precision, at least with annual discount rates, starts to decrease.

TABLE A:V: SMOKER ESTIMATE STATISTICAL SIGNIFICANCE (ANNUAL DGP)
PAPKE-WOOLDRIDGE FRACTIONAL RESPONSE MODEL

	Model 1	**Model 2**
	Annual	Daily
Difference in annual discount rate: Smoker - non-smoker	-0.0409***	-0.0433***
	(0.0008)	(0.0009)
(Difference in annual discount rate: Smoker - non-smoker)^2	0.0002***	0.0002***
	(0.0000)	(0.0000)
Average number of switches on the KPB task	0.2345***	0.2246***
	(0.0048)	(0.0050)
Fraction of smokers in the sample	-3.1558***	-2.7084***
	(0.2649)	(0.2770)
Constant	-0.1572***	-0.2005***
	(0.0450)	(0.0469)
N	50000	50000
log-likelihood	-14854.404	-13521.594

Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table A:V presents results from the Papke and Wooldridge (1996) fractional response model where smoker covariate annual (Model 1) and daily (Model 2) *p*-values are estimated as a function of the difference in average *annual* discount rates of smokers and non-smokers, the average number of switches on the KPB task, and the fraction of smokers in the sample.[11] As I found in Table A:III under the assumption of a daily rate DGP, *p*-values tend to increase as the level of inconsistency (i.e., the average number of switches on the KPB task) rises, and *p*-values tend to decrease as the fraction of smokers in the sample increases.

However, unlike Table A:III, there is a weak U-shaped relationship between smoker covariate *p*-values and differences in the average annual discount rates of smokers and non-smokers in both the daily and annual models. The qualifier "weak" was used in the previous sentence because the estimate of the quadratic term is far smaller than the estimate of the linear term; in Table A:III, the estimates of the quadratic terms

[11] The huge variation in smoker covariate *p*-values in Figure A:XIII suggests that subjecting these data to statistical analysis may be of limited value. Nevertheless, I conduct these analyses for completeness' sake and to draw comparisons between the results from the daily rate and annual rate DGP simulations.

dwarfed the estimates of the linear terms. In other words, the estimates in Table A:V imply that smoker covariate *p*-values tend to decrease over almost the entire range of differences in the average annual discount rates of smokers and non-smokers and then rise very little. Specifically, the turning point in this relationship occurs at an average annual difference in discount rates of 120.25 in Model 1 (i.e., the annual model), and an average annual difference in discount rates of 127.50 in Model 2 (i.e., the daily model).

Note that the estimate of the linear term in the daily model is significantly smaller ($p < 0.001$) than the estimate of the linear term in the annual model. This implies that an infinitesimal increase in the average annual discount rates of smokers and non-smokers leads to a larger fall in smoker covariate *p*-values in the daily model compared to the annual model, at least for values of the difference in discount rates which produce declines in smoker covariate *p*-values.

In sum, the results in this subsection suggest that it is more difficult to detect differences in the discounting behaviour of smokers and non-smokers when annual discount rates, as opposed to daily discount rates, define the DGP of the KPB task. In addition, the results in this subsection provide some support for the contention that estimating daily discount rates is more reliable than estimating annual discount rates, even when the DGP is defined by annual rates. Recall from the discussion of Figure A:XIII that 6% of daily smoker covariate *p*-values are not statistically significant when the difference in average annual discount rates is greater than 55, whereas 10% of annual smoker covariate *p*-values are not statistically significant when the difference in average annual discount rates is greater than 55. Thus, in about 4% of cases, the daily rate estimates detect the difference in discounting behaviour of smokers and non-smokers, whereas the annual rate estimates do not.

The weight of the simulation evidence in this appendix suggests that to detect group differences in discounting behaviour on the KPB task, it is more reliable to estimate or infer daily, as opposed to annual, discount rates. This result holds strongly in the case of a daily discount rate DGP but only weakly with an annual discount rate DGP. It is still somewhat of a numerical puzzle as to why the daily rate estimates are more reliable at detecting group differences than the annual rate estimates, but to err on the

side of caution when drawing inferences with the KPB task, this appendix suggests that it is preferable to estimate daily discount rates as opposed to annual discount rates.

ADDITIONAL REFERENCES

PAPKE, L. E., AND J. M. WOOLDRIDGE (1996): "Econometric Methods for Fractional Response Variables with an Application to 401 (K) Plan Participation Rates," *Journal of Applied Econometrics*, 11, 619-632.

APPENDIX B

In Section V of Chapter 3 I discussed the estimation of discounting models under the assumption that utility $U(\cdot)$ is linear in income y. Andersen, Harrison, Lau and Rutström (2008) devised a method for jointly estimating the curvature of the utility function and the parameters of discounting models. These researchers used a subject's choices on a risk preference task to pin down the shape of the utility function and then used the subject's choices on a discounting task to estimate the discounting parameters, conditional on the shape of the utility function. As discussed earlier, unfortunately the UCLA and USC studies did not elicit risk preferences so the method of Andersen, Harrison, Lau and Rutström (2008) cannot be used to jointly estimate utility function curvature and discounting behaviour.

However, if one makes some parametric assumptions about the form of the utility function (e.g., a constant absolute risk aversion (CARA) or constant relative risk aversion (CRRA) utility function) then one can estimate the parameters from discounting models subject to particular parameter values of the utility function. By incrementally varying the parameters defining the utility function and re-estimating the discounting parameters, one can then track the evolution of the model's log-likelihood. This process, referred to as the estimation of "profile likelihoods,"[12] allows one to find the set of parameters which yield the highest log-likelihood value for the model. This set of parameters defines, therefore, the optimal shape of the utility function conditional on the subjects' discounting choices.

To illustrate this method, assume that utility of income is defined by a power utility function which exhibits constant relative risk aversion (CRRA):

$$U(y) = y^r, \tag{1}$$

where the parameter r determines the shape of the utility function. With this parametric structure, one can re-write (12) in the main text of Chapter 3 to incorporate the power utility function:

[12] See Harrison and Rutström (2009) for an example of this approach in the context of choice under risk.

$$y_0^r = [1 / (1 + \delta)^\tau] y_\tau^r \qquad (2)$$

If $r = 1$ then (2) collapses to (12) in the main text of Chapter 3 but this is not necessarily the case and one can use the estimation of profile likelihoods to find the optimal value of r, conditional on the subjects' discounting choices. Note that as r varies over the interval (0, 1.5) the utility function changes from concave ($r < 1$), to linear ($r = 1$), to convex ($r > 1$), which implies diminishing, constant, and increasing marginal utility, respectively.

I will vary the value of r in increments of 0.05 over the interval (0, 1.5) and then track the evolution of the model's parameters and the model's log-likelihood. I will perform this exercise for the E, H, QH, and WB models.

Figure B:I plots estimates of the discounting parameter δ, the error term μ, and the log-likelihood, as a function of the power utility function parameter r, for the E model. The log-likelihood value is highest when $r = 0.15$ and a dashed line is included in the figure to indicate this point; there is also a dashed line at $r = 1$ which was the assumption adopted in Section V of Chapter 3. At $r = 0.15$, $\delta = 0.00145$, $\mu = 0.11455$, and the log-likelihood value is -14100.69. The estimate of δ increases steadily as r rises which highlights the point made earlier that if $U(\cdot)$ is concave then the assumption of linear utility will bias estimates of the discounting parameter δ upwards.

Figure B:II plots estimates of the discounting parameter δ, the error term μ, and the log-likelihood, as a function of the power utility function parameter r, for the H model. The log-likelihood value is highest when $r = 0.15$ and a dashed line is included in the figure to indicate this point; there is also a dashed line at $r = 1$ which was the assumption adopted in Section V of Chapter 3. At $r = 0.15$, $\delta = 0.00155$, $\mu = 0.11216$, and the log-likelihood value is -14079.93. Note that as in Figure B:I, the estimates of μ and δ increase as the parameter r rises.

Figure B:I: E model profile likelihood

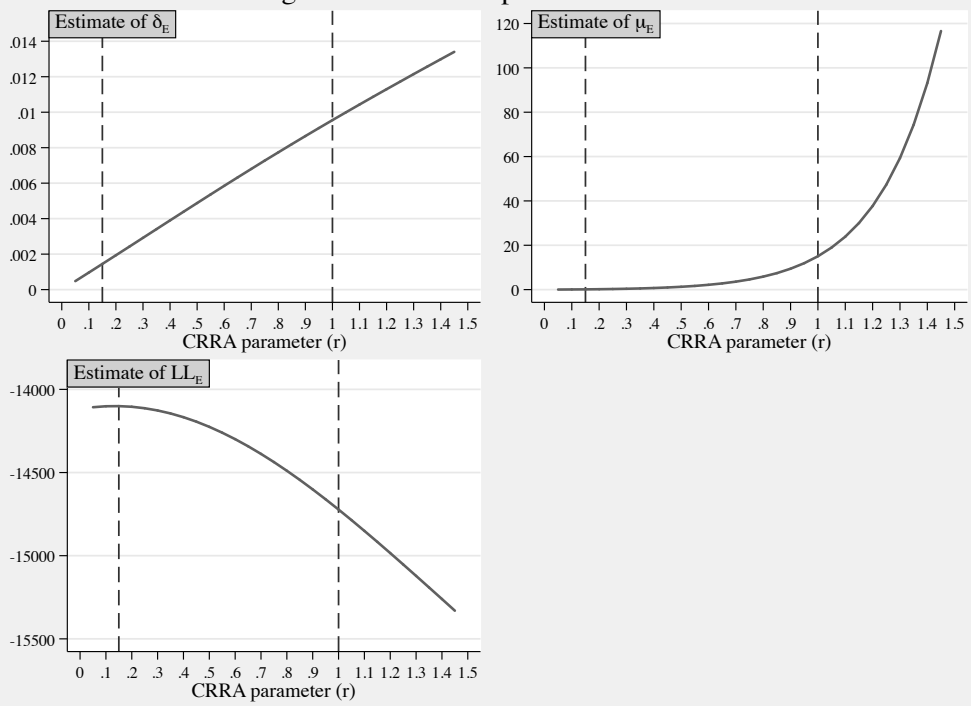

Figure B:II: H model profile likelihood

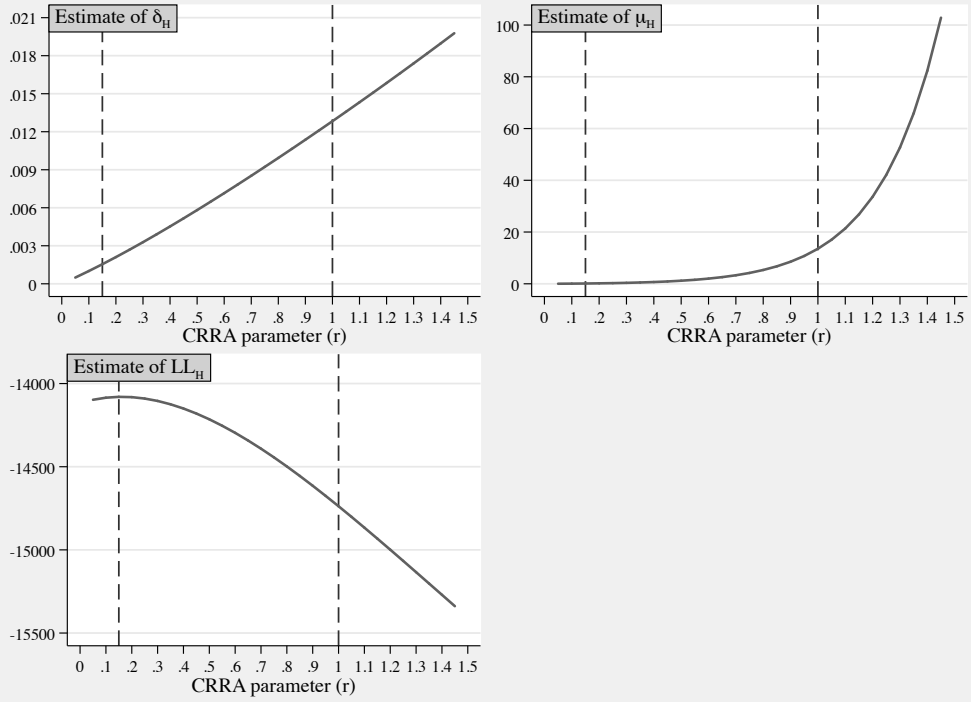

Figure B:III plots estimates of the discounting parameters δ and β, the error term μ, and the log-likelihood, as a function of the power utility function parameter *r*, for the QH model. The log-likelihood value is highest when $r = 0.3$ and a dashed line is

included in the figure to indicate this point; there is also a dashed line at $r = 1$ which was the assumption adopted in Section V of Chapter 3. At $r = 0.3$, $\delta = 0.00129$, $\beta = 0.92765$, $\mu = 0.26045$, and the log-likelihood value is -14084.34. The estimates of δ and β decline sharply over the interval (0, 0.3), but δ rises and β declines slowly thereafter. Note that β is greater than 1 for values of $r < 0.3$, it drops below 1 at $r = 0.3$, and then remains below 1 at higher values of r. These results show that quasi-hyperbolic discounting (i.e., $\beta < 1$) is only present at higher values of r. The evolution of the model's log-likelihood is interesting because it is relatively flat for low values of r, it spikes at $r = 0.3$ and then declines steadily as r continues to rise. As in the previous figures, the error term μ increases as r rises.

Figure B:III: QH model profile likelihood

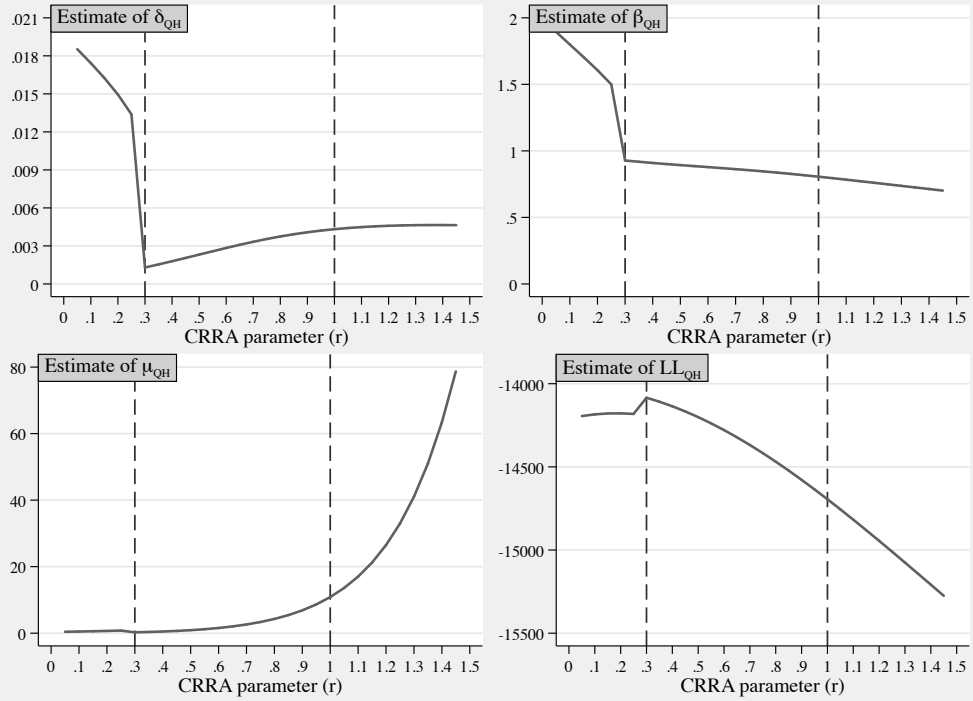

Figure B:IV plots estimates of the discounting parameters δ and β, the error term μ, and the log-likelihood, as a function of the power utility function parameter r, for the WB model. The log-likelihood value is highest when $r = 0.05$ and a dashed line is included in the figure to indicate this point; there is also a dashed line at $r = 1$ which was the assumption adopted in Section V of Chapter 3. At $r = 0.05$, $\delta = 0.00239$, $\beta = 1.63457$, $\mu = 0.02136$, and the log-likelihood value is -14018.49. The estimate of δ follows an interesting path: it increases at a decreasing rate until $r = 0.6$, after which it

increases at an increasing rate. β declines steadily until $r = 0.85$, after which it rises slowly. Note that the estimate of β is always greater than 1, implying that subjects perceive time to be "slowing down" when discounting delayed rewards. The model's log-likelihood declines and the error term μ increases as r rises.

Figure B:IV: WB model profile likelihood

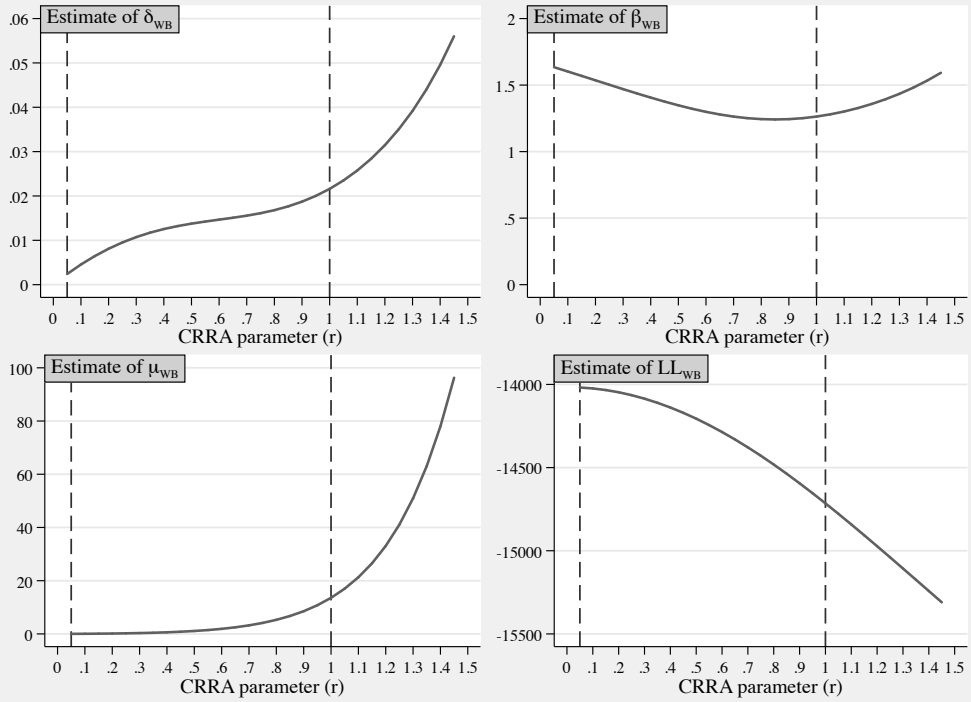

The results in this appendix highlight the link between the shape of a utility function and the inferences that one draws about discounting behaviour. For every model, the estimate of δ is higher under the assumption that $r = 1$, than it is when the profile likelihood attains its maximum (i.e., $r = 0.15$ for the E and H models, $r = 0.3$ for the QH model, and $r = 0.05$ for the WB model). Unfortunately, without data on risk attitudes, one needs to make strong parametric assumptions about the form of a person's utility function, and one needs to assume that the same utility function applies to all individuals, which means it is important to collect risk preference data when estimating discounting parameters so that valid inferences can be drawn.

APPENDIX C

In this appendix I present the results from the five other two process mixture models that can be estimated from the four discounting functions used in Chapter 3. The take home message from this appendix is that all discounting models find significant support in the data, although in certain cases one function finds significantly more support than another.

TABLE C:I: MIXTURE MODEL ML ESTIMATES
HOMOGENOUS PREFERENCES AND HOMOSKEDASTIC ERRORS

	Estimate	Std error	*p*-value	95% Confidence interval	
		Exponential discounting model			
Discounting parameter (δ_E^{mix})	0.00235***	0.00015	0.00000	0.00206	0.00263
Mixture probability (π^E)	0.43412***	0.02073	0.00000	0.39348	0.47476
		Quasi-Hyperbolic discounting model			
Discounting parameter (δ_{QH}^{mix})	0.01268***	0.00048	0.00000	0.01173	0.01362
Discounting parameter (β_{QH}^{mix})	0.67336	0.01644	0.00000	0.64114	0.70558
Mixture probability (π^{QH})	0.56588***	0.02073	0.00000	0.52524	0.60652
		Fechner error term			
Error (μ)	7.14859***	0.15164	0.00000	6.85138	7.44580
N	32535				
log-likelihood	-14317.31				

H_0: $\pi^E = 0.5$, p-value = 0.0015

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE C:II: MIXTURE MODEL ML ESTIMATES
HOMOGENOUS PREFERENCES AND HOMOSKEDASTIC ERRORS

	Estimate	Std error	*p*-value	95% Confidence interval	
		Exponential discounting model			
Discounting parameter (δ_E^{mix})	0.00281***	0.00016	0.00000	0.00250	0.00312
Mixture probability (π^E)	0.49634***	0.0255	0.00000	0.45222	0.54045
		Weibull discounting model			
Discounting parameter (δ_{WB}^{mix})	0.14458***	0.02467	0.00000	0.09623	0.19292
Discounting parameter (β_{WB}^{mix})	1.83431	0.13392	0.00000	1.57184	2.09678
Mixture probability (π^{WB})	0.50366	0.0225	0.00000	0.45954	0.54778
		Fechner error term			
Error (μ)	7.79150***	0.21057	0.00000	7.37879	8.20422
N	32535				
log-likelihood	-14384.83				

H_0: $\pi^E = 0.5$, *p*-value = 0.87080

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE C:III: MIXTURE MODEL ML ESTIMATES
HOMOGENOUS PREFERENCES AND HOMOSKEDASTIC ERRORS

	Estimate	Std error	*p*-value	95% Confidence interval	
		Hyperbolic discounting model			
Discounting parameter (δ_H^{mix})	0.00277***	0.00019	0.00000	0.00240	0.00313
Mixture probability (π^H)	0.44928***	0.02107	0.00000	0.40798	0.49057
		Quasi-Hyperbolic discounting model			
Discounting parameter (δ_{QH}^{mix})	0.01259***	0.00049	0.00000	0.01162	0.01355
Discounting parameter (β_{QH}^{mix})	0.66077	0.01654	0.00000	0.62834	0.69319
Mixture probability (π^{QH})	0.55072***	0.02107	0.00000	0.50943	0.59202
		Fechner error term			
Error (μ)	6.97251***	0.15555	0.00000	6.66763	7.27739
N	32535				
log-likelihood	-14315.18				

H_0: $\pi^H = 0.5$, *p*-value = 0.01610

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE C:IV: MIXTURE MODEL ML ESTIMATES
HOMOGENOUS PREFERENCES AND HOMOSKEDASTIC ERRORS

	Estimate	Std error	*p*-value	95% Confidence interval	
		Hyperbolic discounting model			
Discounting parameter (δ_H^{mix})	0.00336***	0.00021	0.00000	0.00294	0.00377
Mixture probability (π^H)	0.51502***	0.02278	0.00000	0.47037	0.55966
		Weibull discounting model			
Discounting parameter (δ_{WB}^{mix})	0.17057***	0.02522	0.00000	0.12113	0.22001
Discounting parameter (β_{WB}^{mix})	1.98864	0.13548	0.00000	1.72309	2.25417
Mixture probability (π^{WB})	0.48498	0.02278	0.00000	0.44034	0.52963
		Fechner error term			
Error (μ)	7.46597***	0.19772	0.00000	7.07845	7.8535
N	32535				
log-likelihood	-14381.94				

H_0: $\pi^H = 0.5$, *p*-value = 0.50970

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE C:V: MIXTURE MODEL ML ESTIMATES
HOMOGENOUS PREFERENCES AND HOMOSKEDASTIC ERRORS

	Estimate	Std error	*p*-value	95% Confidence interval	
		Quasi-Hyperbolic discounting model			
Discounting parameter (δ_{QH}^{mix})	0.01168***	0.00039	0.00000	0.01092	0.01244
Discounting parameter (β_{QH}^{mix})	0.62259***	0.02328	0.00000	0.57697	0.66821
Mixture probability (π^{QH})	0.50988***	0.02728	0.00000	0.45642	0.56334
		Weibull discounting model			
Discounting parameter (δ_{WB}^{mix})	0.01578**	0.00621	0.01100	0.00361	0.02795
Discounting parameter (β_{WB}^{mix})	1.64602	0.21145	0.00000	1.23158	2.06046
Mixture probability (π^{WB})	0.49012	0.02728	0.00000	0.43666	0.54358
		Fechner error term			
Error (μ)	6.41381***	0.20157	0.00000	6.01873	6.80888
N	32535				
log-likelihood	-14311.45				

H_0: $\pi^{QH} = 0.5$, *p*-value = 0.7172

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

APPENDIX D

In this appendix I present the results from the two other two process mixture models that can be estimated from the three PWFs used in Chapter 4. Estimating a mixture model of the TK and Prelec functions was straightforward and the global maximum that was attained provided sensible estimates of the parameters of these two functions. Unfortunately, estimating a mixture model of the Prelec and PD functions was fraught with difficulty and I will detail these issues in what follows. I first present the results from the mixture model of the TK and Prelec functions in Table D:I below.

TABLE D:I: MIXTURE MODEL ML ESTIMATES
PRELEC AND TK FUNCTIONS

	Estimate	Std error	*p*-value	95% Confidence interval	
Prelec probability weighting function					
PWF parameter (γ^{Prelec})	0.331**	0.150	0.027	0.038	0.625
PWF parameter (η^{Prelec})	1.465***	0.095	0.000	1.279	1.651
Mixture probability (π^{Prelec})	0.164*	0.084	0.052	-0.001	0.329
TK probability weighting function					
PWF parameter (γ^{TK})	0.859***	0.046	0.000	0.770	0.949
Mixture probability (π^{TK})	0.836***	0.084	0.000	0.671	1.001
Sigma					
Constant (σ)	0.105***	0.015	0.000	0.075	0.136
N	480				
log-likelihood	310.772				

H_0: $\pi^{TK} = 1$, *p*-value = 0.052

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table D:I shows that the TK function accounts for approximately 84% of the choices in the data and that the Prelec function, therefore, accounts for the remaining 16%; I can reject the hypothesis that π^{TK} is equal to 1 ($p = 0.052$). The estimates of the Prelec parameters (γ and η) yield an inverse S-shaped function and so too does the estimate of $\gamma = 0.859$ in the TK model.

When I first tried to estimate a mixture model of the PD and Prelec functions, I did not constrain any of the parameter estimates. Estimation of profile likelihoods, to find appropriate starting values which yield a global maximum, showed that two maxima (one local and one global) were obtained. The local maximum (log-likelihood = 299.572) provided sensible estimates for the parameters of the Prelec function but an estimate of γ^{PD} which is not significantly different to 0. This estimate is not plausible because it implies that the PD PWF is a constant which is equal to 1. These estimates are presented in Table D:II below.

TABLE D:II: MIXTURE MODEL ML ESTIMATES
PD AND PRELEC FUNCTIONS (UNCONSTRAINED LOCAL MAXIMUM)

	Estimate	Std error	*p*-value	95% Confidence interval	
		PD probability weighting function			
PWF parameter (γ^{PD})	4.658	3.063	0.128	-1.344	10.661
Mixture probability (π^{PD})	0.138	0.118	0.240	-0.092	0.369
		Prelec probability weighting function			
PWF parameter (γ^{Prelec})	0.829***	0.067	0.000	0.698	0.959
PWF parameter (η^{Prelec})	0.959***	0.110	0.000	0.744	1.175
Mixture probability (π^{Prelec})	0.862***	0.118	0.000	0.631	1.092
		Sigma			
Constant (σ)	0.105***	0.0080	0.0000	0.0890	0.1210
N	480				
log-likelihood	299.572				

H_0: $\pi^{Prelec} = 1$, *p*-value = 0.240

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

The starting values yielding a global maximum (log-likelihood = 313.579) provided plausible estimates for γ^{PD} but completely implausible estimates for γ^{Prelec} and η^{Prelec}; these results are presented in Table D:III below. Recall that $\eta, \gamma > 0$ in the Prelec model and yet γ^{Prelec} in the mixture model is not significantly different to 0, its 95% confidence interval includes values less than 0, and the 95% confidence interval for

η^{Prelec} also includes values less than 0. As γ^{Prelec} is not significantly different to zero, the PWF is a step function.

TABLE D:III: MIXTURE MODEL ML ESTIMATES
PD AND PRELEC FUNCTIONS (UNCONSTRAINED GLOBAL MAXIMUM)

	Estimate	Std error	*p*-value	95% Confidence interval	
		PD probability weighting function			
PWF parameter (γ^{PD})	1.266***	0.133	0.000	1.006	1.527
Mixture probability (π^{PD})	0.883***	0.038	0.000	0.810	0.957
		Prelec probability weighting function			
PWF parameter (γ^{Prelec})	-0.285	0.315	0.366	-0.901	0.332
PWF parameter (η^{Prelec})	0.567*	0.321	0.078	-0.063	1.120
Mixture probability (π^{Prelec})	0.117***	0.038	0.002	0.043	0.190
		Sigma			
Constant (σ)	0.105***	0.008	0.000	0.089	0.121
N	480				
log-likelihood	313.579				
H_0: $\pi^{PD} = 1$, *p*-value = 0.002					

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Given these issues I decided to re-estimate the models and constrain γ^{PD}, γ^{Prelec}, and $\eta^{Prelec} > 0$. The global maximum (log-likelihood = 310.034) of the constrained model provided sensible estimates for γ^{PD} but I could not recover estimates for γ^{Prelec} (i.e., the coefficient estimate as recovered using the delta method was 0 and no standard errors could be estimated). As noted previously, when $\gamma^{Prelec} = 0$, this yields a step function. These results are presented in Table D:IV below.

Another issue I encountered was that a number of different starting values led to the "same" (i.e., down to the sixth decimal point) global maximum but every time I used a different set of starting values I got different estimates for $\exp(\gamma^{Prelec})$[13]; this issue is

[13] Note that $\exp(\gamma^{Prelec})$ is the nonlinear transformation that was used to ensure that γ^{Prelec} was strictly positive.

somewhat incidental because regardless of the estimate of $\exp(\gamma^{Prelec})$ I could not recover γ^{Prelec} using the delta method.

TABLE D:IV: MIXTURE MODEL ML ESTIMATES
PD AND PRELEC FUNCTIONS (CONSTRAINED GLOBAL MAXIMUM)

	Estimate	Std error	*p*-value	95% Confidence interval	
		PD probability weighting function			
PWF parameter (γ^{PD})	1.210***	0.134	0.000	0.947	1.473
Mixture probability (π^{PD})	0.849***	0.060	0.000	0.732	0.966
		Prelec probability weighting function			
PWF parameter (γ^{Prelec})	0	.	.	.	.
PWF parameter (η^{Prelec})	0.853***	0.264	0.001	0.335	1.371
Mixture probability (π^{Prelec})	0.151**	0.060	0.011	0.034	0.268
		Sigma			
Constant (σ)	0.107***	0.010	0.000	0.087	0.127
N	480				
log-likelihood	310.034				

H_0: $\pi^{PD} = 1$, *p*-value = 0.012

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

I then decided to estimate a mixture model of the PD and one-parameter Prelec (i.e., $\pi(p) = \exp[\ -(-\ln p)^{\gamma}\ ]$) functions, with estimates constrained to be positive, in case the issue had something to do with the two-parameter Prelec function. My estimate of $\gamma^{Prelec\text{-}I}$ was not significantly different to zero.

In sum, I was unable to reliably estimate a mixture model of the PD and Prelec PWFs. I document this issue in this appendix for completeness' sake.

APPENDIX E

The introductory presentation, which all experimental subjects were taken through, is included in this appendix. It explains the nature of the risk and time preference tasks and it includes a detailed discussion of the physical randomisation devices used in the experiments. Every effort was taken to ensure that subjects understood how their choices ultimately led to the earnings they received so as to incentivise the truthful revelation of preferences.

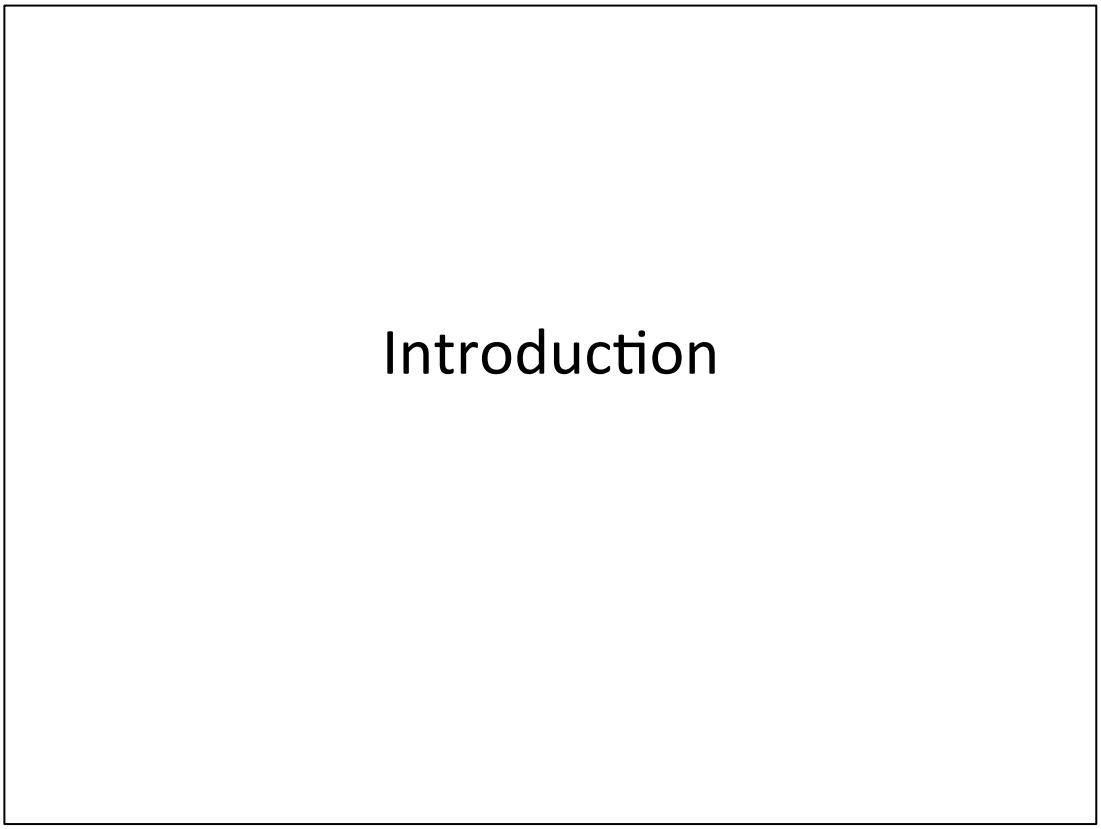

Welcome

- Hello everyone and welcome to today's research session.
- Thank you for agreeing to take part in this study, your views and choices will be very informative and helpful.
- Before we get started I would like to explain how things are going to work.
- When I have finished this short explanation I will ask you to read and sign a consent form.
- Once that is done, we can begin with the tasks.

2 Tasks and a Questionnaire

- You will take part in 2 tasks and you will have the opportunity to earn money in each task.
- One of the tasks requires you to make choices between lotteries with varying prizes and chances of winning. You will make 40 of these choices.
- The other task asks you to choose between amounts of money available at different points in time. You will make 60 of these choices.
- Once you have completed the 2 tasks, you will need to fill out a short questionnaire.
- Once this is done, we will determine your earnings and you will be free to leave.

Earnings

- You will be paid R20 just for participating in today's session.
- At the end of the session, we will determine your earnings for the tasks.
- Some of this money will be paid to you at the end of the session and the rest of it will be paid to you in the future.
- This is why we need your bank details: to pay you via electronic transfer at a future date.
- To determine your earnings for the tasks, we will ask you to roll some dice.
- Let's go through a quick explanation of the dice you will roll.

10-sided dice

- At the end of today's session we will ask you to roll some dice into a plastic bowl which you can see below.
- Two of the dice that you will roll are 10-sided dice and these are used to select a number between 1 and 100.
- Every number between 1 and 100, and including 1 and 100, is equally likely to occur.
- An example of a dice roll is shown below.

10-sided dice

- Let's look at a close-up of the 10-sided dice.
- As you can see, one of the 10-sided dice has sides which increase in multiples of 10: 00, 10, 20, 30, 40, 50, 60, 70, 80 and 90.
- The other 10-sided die has sides which increase in multiples of 1: 0, 1, 2, 3, 4, 5, 6, 7, 8 and 9.
- You will roll the two 10-sided dice together and add the numbers on the two dice to select a number between 1 and 100.
- In the example below, the number that was rolled is 86 (80 + 6).

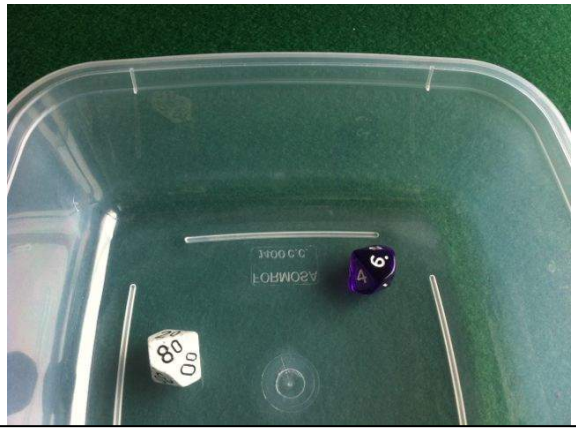

10-sided dice

- To tell the difference between a 6 and a 9 there is a dot at the base of the number.
- This is why the number in the picture on the left below is a 6: there is a dot at the base of the 6.
- 9 looks different because there is a dot at the base of the 9.
- The picture on the right below shows you what a 9 looks like.

10-sided dice

- To roll a number between 1 and 9 you need to roll 00 and a single number between 1 and 9.
- As you can see in the picture on the left below, the number that was rolled is 5 (00 + 5).
- In the case where you roll 00 and 0, this will be treated as 100.
- As you can see in the picture on the right below, the number that was rolled is 100 (00 and 0).

Consent Form

- We have now finished the introductory explanation.
- To continue with today's session I need you to read and sign the consent form.
- This form explains your rights as a research participant and by signing it, you give your consent to participate in the study.
- If you have any questions please raise your hand and someone will come to answer them.
- You may read through the consent form now.

APPENDIX F

This appendix includes the presentation which was used to explain the risk preference task. The presentation goes into a lot of detail about the computer environment within which choices are made, the lotteries on offer and how to interpret them, the payment scheme that is used to determine earnings, and the fact that there are no right or wrong choices in the task. As discussed previously, all of this was done to make the task transparent and easy to understand.

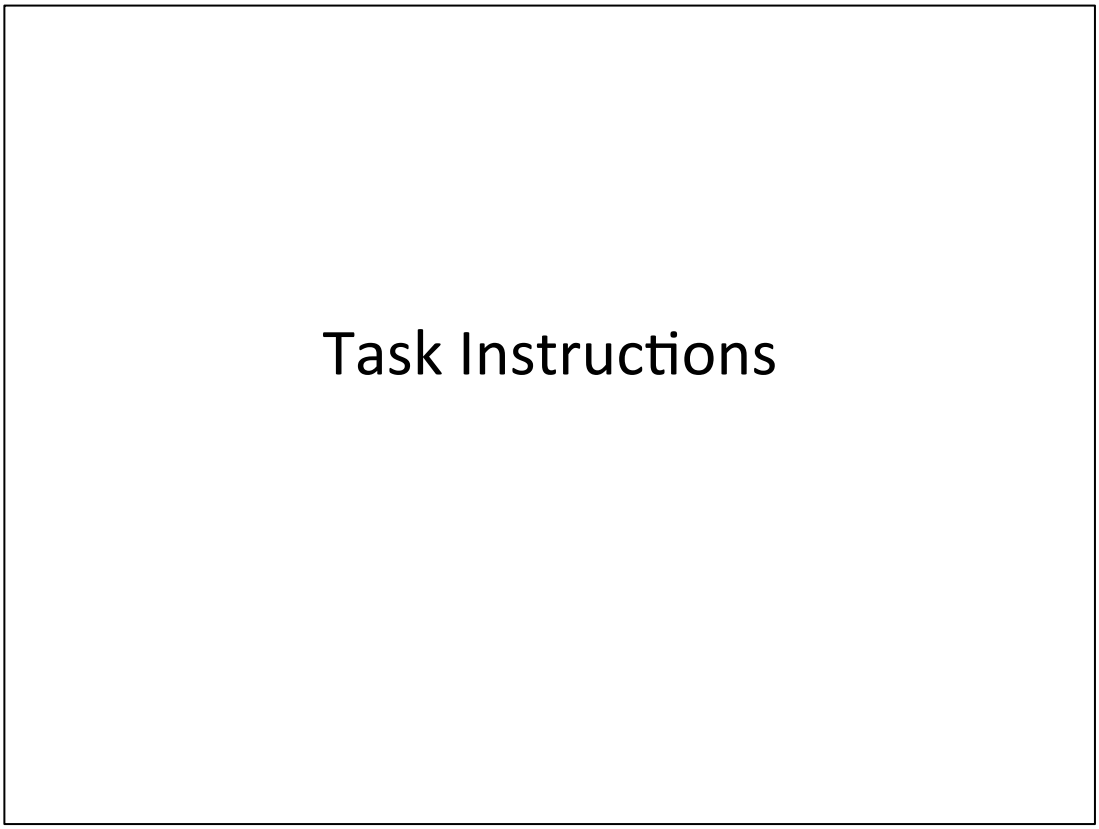

Introduction

- In this task you will choose between lotteries with varying prizes and chances of winning.
- On each computer screen you will be presented with a pair of lotteries and you will need to choose one of them.
- There are 40 pairs of lotteries in this task.
- For each pair of lotteries, you should choose the lottery that you would prefer to play.
- You will actually get the chance to play one of the lotteries you choose, and you will be paid according to the outcome of this lottery.
- So you should think carefully about which lottery you prefer in each pair.

Computer Display

- All of the choices in this task will be made on a computer.
- This is what the computer display will look like:

- The display on your screen will be larger and easier to read.

Computer Display

- On the computer screen there are two lotteries: a “Left” lottery and a “Right” lottery.
- Let’s look at the Left lottery together.
- For the Left lottery there is a 20% chance of winning R0, a 60% chance of winning R100, and a 20% chance of winning R200.
- The coloured areas of the pie chart and the text below the pie chart represent these chances.

The Left Lottery

- As you can see, 20% of the pie chart is blue and this means you have a 20% chance of winning R0.
- The is what the blue text below the pie chart tells you: “Chance of winning R0 is 20%”.
- Similarly, 60% of the pie chart is red which means there is a 60% chance of winning R100.
- The is what the red text below the pie chart tells you: “Chance of winning R100 is 60%”.
- Finally, 20% of the pie chart is green which means there is a 20% chance of winning R200.
- This is what the green text below the pie chart tells you: “Chance of winning R200 is 20%”.

The Right Lottery

- If we look at the Right lottery we see that there is a 10% chance of winning R0 and a 90% chance of winning R100.
- 10% of the pie chart is blue and this means there is a 10% chance of winning R0.
- This is what the blue text below the pie chart tells you: “Chance of winning R0 is 10%”.
- 90% of the pie chart is red which means there is a 90% chance of winning R100.
- This is what the red text below the pie chart tells you: “Chance of winning R100 is 90%”.

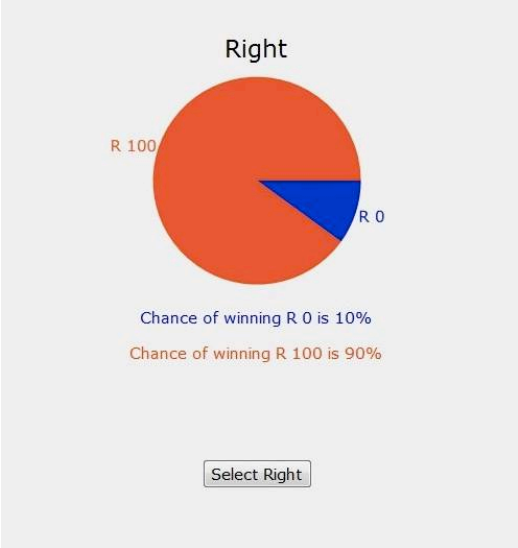

Your Lottery Winnings

- The amount that you win from a lottery will be determined by the draw of a random number between 1 and 100.
- Each number between 1 and 100, and including 1 and 100, is equally likely to occur.
- You will draw this number yourself by rolling two 10-sided dice.
- One of the 10-sided dice has sides which increase in multiples of 10: 00, 10, 20, 30, 40, 50, 60, 70, 80 and 90.
- The other 10-sided die has sides which increase in multiples of 1: 0, 1, 2, 3, 4, 5, 6, 7, 8 and 9.
- You will roll the two 10-sided dice together and add the numbers on the two dice to select a number between 1 and 100.
- For example, suppose the one 10-sided die lands on 70 and the other 10-sided die lands on 5.
- Then we will select number 75.
- We will work through an actual example of this later.

Choices

- Now, suppose that you prefer the Left lottery in the example below.
- To choose the Left lottery just click the button saying “Select Left”.
- This is what the display will then look like if you choose the Left lottery.
- You can then click the button saying “Confirm” to move on to the next screen with a new pair of lotteries.
- If you would like to change your choice then just click “Cancel”.

Choices

- Suppose instead that you prefer the Right lottery in the example below.
- To choose the Right lottery just click the button saying “Select Right”.
- This is what the display will then look like if you choose the Right lottery.
- You can then click the button saying “Confirm” to move on to the next screen with a new pair of lotteries.
- If you would like to change your choice then just click “Cancel”.

Total Number of Choices

- You will need to make 40 choices across 40 screens.
- On each screen there is a different lottery pair and you will need to choose either the Left lottery or the Right lottery.
- The Rand amounts under the lotteries change on each screen.
- In addition, the chances of winning the Rand amounts change for each lottery on each screen.
- So please pay careful attention when making each choice.
- At the end of the session today we will determine your earnings for this task in the following way.

Payment

- First, you will select one of the lottery pairs from this task by rolling a 4-sided die and then a 10-sided die.
- You will roll the 4-sided die to select 10 lottery pairs.
- If the die lands on 1, you will select lottery pairs 1-10; if the die lands on 2, you will select lottery pairs 11-20; if the die lands on 3, you will select lottery pairs 21-30; and if the die lands on 4, you will select lottery pairs 31-40.
- You will then roll the 10-sided die to select one lottery pair from this set of 10 pairs.
- For example, if the 4-sided die lands on 3, you will select lottery pairs 21-30.
- If you then roll a 7 on the 10-sided die, you will select lottery pair 27.
- Once you have selected the lottery pair, we will look at the choice that you made: the Left lottery or the Right lottery.
- We will then determine your winnings from this lottery by rolling two 10-sided dice, as explained earlier.
- Let's see what this means for the example we looked at earlier.

Payment

- Suppose that the lottery pair we looked at earlier gets selected for payment when you roll the 4-sided die and then the 10-sided die.
- And suppose that you chose the Left lottery on this screen.
- For the Left lottery there is a 20% chance of winning R0, a 60% chance of winning R100, and a 20% chance of winning R200.
- You will now roll two 10-sided dice to determine your winnings.
- As you can see on the screen, if you roll a number between 1 and 20, you will win R0. Thus, you have a 20% chance of winning R0.

Payment

- If you roll a number between 21 and 80 you will win R100. Thus, you have a 60% chance of winning R100.
- Finally, if you roll a number between 81 and 100 you will win R200. Thus, you have a 20% chance of winning R200.
- Suppose you roll the two 10-sided dice and one 10-sided die lands on 60 while the other 10-sided die lands on 7.
- Then we will select number 67.
- Because 67 is between 21 and 80, you will win R100.

Payment

- Thus, payment for this task is determined by three things:
  1. The lottery pair that is chosen to be played out using the 4-sided die and the 10-sided die.
  2. Your choice of the Left lottery or the Right lottery in each pair.
  3. The outcome of that lottery when you roll the two 10-sided dice.
- All winnings will be paid in cash at the end of today's session.

Choose the Option you Prefer

- The lottery you prefer in each pair is a matter of personal taste.
- The person next to you may have different tastes so their choices should not matter to you.
- Please work silently and make your choices by thinking carefully about each lottery.
- Since there is a chance that any one of your 40 choices could be selected for payment, you should approach each pair of lotteries as if it is the one that you will be paid for.
- If you have any questions please raise your hand and someone will come to answer them.
- You may begin the task.

APPENDIX G

This appendix includes the presentation which was used to explain the time preference task. The presentation explains the computer environment within which choices are made, the calendar which shows subjects the dates at which SS and LL rewards are available, the payment scheme that is used to determine earnings, and the fact that there are no right or wrong choices in the task. As discussed previously, all of this was done to make the task transparent and easy to understand.

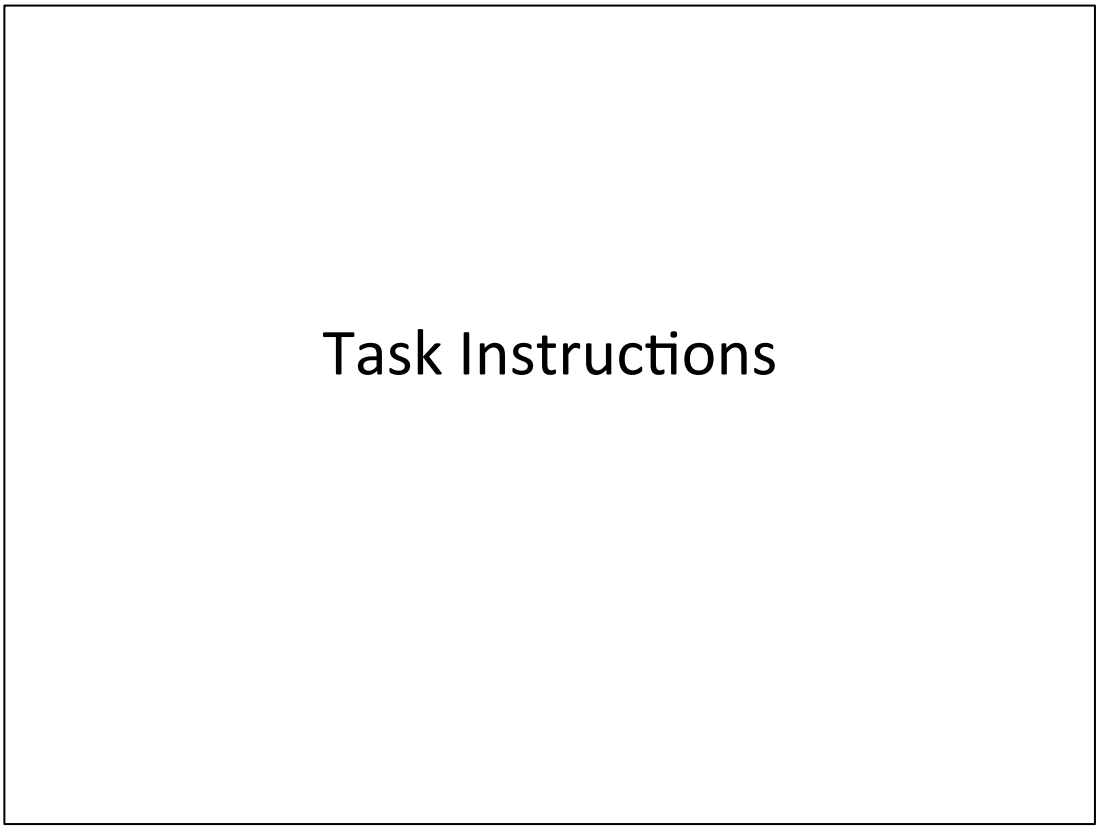

Introduction

- In this task you will choose between different amounts of money available at different times.
- You will need to make 60 choices in total.
- For each choice you will decide between a smaller amount of money which is available sooner and a larger amount of money which is available later.
- One of your 60 choices will be selected at random for payment and you will receive the amount of money that you chose on the appropriate date.

Computer Display

- All of these choices will be made on a computer.
- This is what the computer display will look like:

- The display on your screen will be larger and easier to read.

Computer Display

- At the top of the display is a calendar showing you today's date in purple (11 July 2012) and a future date in green (25 July 2012).
- Below the calendar are two columns: a purple column with amounts of money available today and a green column with amounts of money available in 14 days from today.
- You need to make 4 choices on this screen.
- Each choice appears on a different row.

Choices

- Let's look at the first row together.
- In the first row, you need to choose between receiving R200 today or R201.15 in 14 days from today.
- Note that R200 is the smaller of the two amounts but it is available today.
- R201.15 is the larger of the two amounts but it is only available after 14 days.
- Suppose that you prefer R200 today rather than R201.15 in 14 days.
- To choose R200 today just click the button saying "Select" under "R200 today".
- This is what the display will look like if you choose R200 today.

Choices

- Suppose instead that you prefer R201.15 in 14 days rather than R200 today.
- To choose R201.15 in 14 days just click the button saying “Select” under “R201.15 in 14 days”.
- This is what the display will look like if you choose R201.15 in 14 days.
- Once you have made your choice on the first row you can move on to the other rows on the screen.
- Note that you need to make 4 choices on the screen before you can move on to the next set of 4 choices on a new screen.

Choices

- Once you have made all your choices on the screen, the display will look something like the example below.
- You can then click the button saying “Confirm” to move on to the next screen.
- If you would like to change your choices then click “Cancel”.

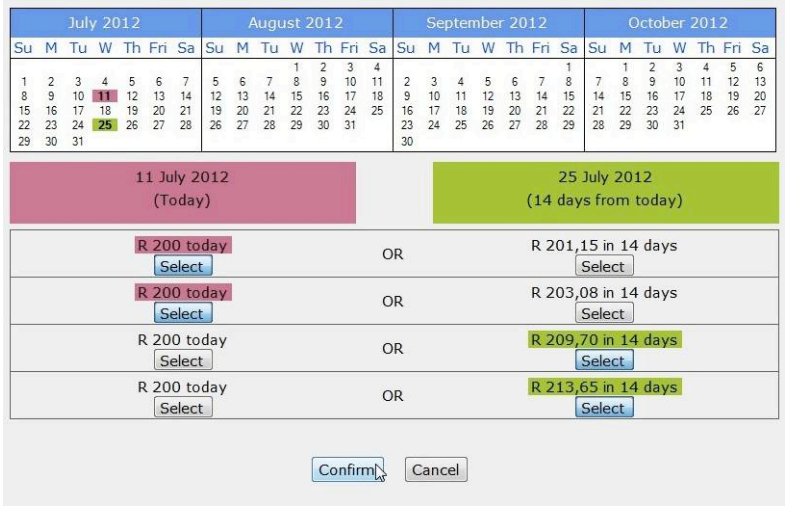

Total Number of Choices

- You will need to make 60 choices in total across 15 screens.
- The Rand amounts change on each row of each screen.
- In addition, the times for delivery of the Rand amounts change across the screens.
- For example, on the screen we just looked at, you had to choose between an amount of money available today and an amount of money available in 14 days.
- On a different screen, you may need to choose between an amount of money available in 14 days and another amount of money available in 21 days.
- So please pay careful attention when making your choices.
- At the end of the session today we will determine your earnings for this task in the following way.

Payment

- First, you will select one of the 15 screens from this task by rolling a 20-sided die.
- If the die lands on 1, you will select screen 1; if the die lands on 7, you will select screen 7; if the die lands on 12, you will select screen 12; and so on.
- If the die lands on 16, 17, 18, 19 or 20, you will roll the die again until it lands on a number between 1 and 15.
- Once you have selected a screen, you will roll a 4-sided die to select 1 of the 4 rows on the screen.
- If the die lands on 1, you will select row 1; if the die lands on 2, you will select row 2; and so on.
- Once you have selected the row, we will look at the choice you made on that row.
- You will then be paid for the choice that you made on that row on the date listed for that choice.
- Let's see what this means for the example we looked at earlier.

Payment

- Suppose that the screen we looked at earlier gets selected for payment when you roll the 20-sided die.
- You will then need to roll the 4-sided die to select a row for payment.
- Suppose that the die lands on 3, then you will select row 3.
- On row 3, you chose R209.70 in 14 days so you will be paid R209.70 in 14 days via electronic transfer.
- You will receive a payment confirmation when the transaction has taken place.

Choose the Option you Prefer

- Note that the option you prefer on each row is a matter of personal taste.
- The people next to you may have different tastes so their choices should not matter to you.
- Please work silently and make your choices by thinking carefully about each option.
- Since there is a chance that any of your 60 choices could be selected for payment, you should approach each choice as if it is the one that you will be paid for.
- If you have any questions please raise your hand and someone will come to answer them?
- You may begin the task.

APPENDIX H

In this appendix I analyse the smoking-discounting relationship with the four discounting models used in Chapter 5 but under the assumption that EU, rather than RDU, characterises choice under risk. Table H:I presents results from the four discounting specifications where the parameters of interest vary according to smoking status. In all models the estimate of δ for smokers is positive and statistically significant which means smokers discount more heavily than non-smokers. The estimates of β in the QH and WB models, by contrast, do not differ according to smoking status. Thus, I have replicated the results from Table XII in Chapter 5 where RDU was assumed to characterise choice under risk.

TABLE H:I: DISCOUNTING FUNCTION ML ESTIMATES
CONCAVE UTILITY AND HETEROGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Power function parameter (r)				
Smoker	0.019	0.020	0.017	0.017
	(0.016)	(0.017)	(0.016)	(0.015)
Constant	0.266***	0.293***	0.258***	0.244***
	(0.033)	(0.028)	(0.034)	(0.038)
Discounting parameter (δ)				
Smoker	0.205***	0.167***	0.185***	0.086**
	(0.062)	(0.047)	(0.063)	(0.035)
Constant	0.369***	0.361***	0.323***	0.170***
	(0.061)	(0.047)	(0.058)	(0.033)
Discounting parameter (β)				
Smoker			0.001	-0.144
			(0.007)	(0.264)
Constant			0.988***	1.664***
			(0.006)	(0.243)
Risk error (μ)				
Constant	0.186***	0.176***	0.189***	0.195***
	(0.011)	(0.009)	(0.012)	(0.014)
Time error (ν)				
Constant	0.143***	0.181***	0.131***	0.116***
	(0.042)	(0.044)	(0.041)	(0.040)
N	17500	17500	17500	17500
log-likelihood	-9389.003	-9358.082	-9314.103	-9169.951

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

To investigate whether time preferences are associated with smoking intensity, I also estimated the discounting models as a quadratic function of number of cigarettes smoked per day. Table H:II shows that for the estimate of δ in all models, the linear

term is positive and significant while the quadratic term is negative and significant. Thus, there is a concave relationship between discounting behaviour and smoking intensity, which is precisely what I found under the assumption that RDU characterises choice under risk. Similarly, I find that smoking intensity is not associated with the estimate of β in the QH model but that the linear term is significant, at the 10% level, in the estimate of β in the WB model.

TABLE H:II: DISCOUNTING FUNCTION ML ESTIMATES
CONCAVE UTILITY AND HETEROGENOUS PREFERENCES

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Power function parameter (r)				
Number of cigarettes	0.001	0.001	0.001	0.002
	(0.003)	(0.004)	(0.003)	(0.003)
(Number of cigarettes)^2	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	0.273***	0.299***	0.266***	0.243***
	(0.034)	(0.029)	(0.036)	(0.041)
Discounting parameter (δ)				
Number of cigarettes	0.043***	0.034***	0.041***	0.028***
	(0.013)	(0.010)	(0.013)	(0.009)
(Number of cigarettes)^2	-0.002***	-0.002***	-0.002***	-0.001***
	(0.001)	(0.000)	(0.001)	(0.000)
Constant	0.389***	0.378***	0.339***	0.151***
	(0.062)	(0.047)	(0.057)	(0.032)
Discounting parameter (β)				
Number of cigarettes			0.001	-0.105*
			(0.002)	(0.056)
(Number of cigarettes)^2			0.000	0.004
			(0.000)	(0.002)
Constant			0.986***	1.937***
			(0.006)	(0.259)
Risk error (μ)				
Constant	0.186***	0.176***	0.189***	0.197***
	(0.012)	(0.009)	(0.013)	(0.015)
Time error (ν)				
Constant	0.141***	0.177***	0.132***	0.109***
	(0.043)	(0.045)	(0.042)	(0.040)
N	17500	17500	17500	17500
log-likelihood	-9394.161	-9365.653	-9320.929	-9164.661

Results account for clustering at the individual level
Standard errors in parentheses
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table H:III maps out the response surface for estimates of δ in the four time preference models evaluated at different values of number of cigarettes smoked per day. At low values of number of cigarettes, the conditional marginal effect of additional cigarettes is positive. By 15 cigarettes though, the conditional marginal

effect of additional cigarettes is negative. Thus, Table H:III replicates the result in Table XIV in Chapter 5 and highlights the nonlinear effect of smoking intensity on discounting behaviour.

TABLE H:III: NUMBER OF CIGARETTES CONDITIONAL MARGINAL EFFECTS FOR δ

	Model 1	**Model 2**	**Model 3**	**Model 4**
	Exponential	Hyperbolic	Quasi-Hyperbolic	Weibull
Number of cigarettes				
0	0.043 (0.013)	0.034 (0.010)	0.041 (0.013)	0.028 (0.010)
5	0.024 (0.008)	0.019 (0.006)	0.023 (0.008)	0.017 (0.006)
10	0.005 (0.004)	0.003 (0.003)	0.006 (0.004)	0.006 (0.003)
15	-0.014 (0.006)	-0.012 (0.006)	-0.012 (0.005)	-0.005 (0.002)
20	-0.033 (0.012)	-0.028 (0.011)	-0.030 (0.010)	-0.016 (0.005)
25	-0.052 (0.018)	-0.043 (0.015)	-0.048 (0.016)	-0.027 (0.008)

Standard errors in parentheses

Finally, tables H:IV:A and H:IV:B analyse the smoking-discounting relationship by making the parameters of interest a linear function of observable characteristics and task parameters. In the E and H models in Table H:IV:A, there is a positive and significant relationship between smoking and discounting behaviour: smokers discount the future more heavily than non-smokers.

Similarly, in the QH and WB models in Table H:IV:B, the estimate of δ for smokers is positive and statistically significant. However, smoking status is not significantly related to the extent of present-bias in the QH model nor in the way people perceive time in the WB model (i.e., in the estimates of β). Thus, the results in tables H:IV:A and H:IV:B replicate those in tables XV:A and XV:B in Chapter 5.

In sum, the preceding results show that the smoking-discounting relationship is robust to the assumption that EU, rather than RDU, characterises choice under risk.

TABLE H:IV:A: DISCOUNTING FUNCTION ML ESTIMATES
CONCAVE UTILITY AND HETEROGENOUS PREFERENCES

	Model 1 Exponential		**Model 2** Hyperbolic	
	Estimate	Std error	Estimate	Std error
Power function parameter (r)				
Age	-0.006	0.005	-0.007	0.005
White	-0.012	0.018	-0.014	0.019
Male	-0.012	0.016	-0.010	0.017
Commerce faculty	0.011	0.018	0.011	0.020
Financial aid	0.039**	0.018	0.040**	0.019
Risk task first	0.004	0.016	0.006	0.017
Smoker	0.052***	0.020	0.056***	0.021
Constant	0.382***	0.110	0.419***	0.111
Discounting parameter (δ)				
Age	-0.004	0.014	-0.003	0.011
White	-0.091	0.066	-0.081	0.052
Male	0.114**	0.055	0.098**	0.042
Commerce faculty	0.027	0.067	0.014	0.055
Financial aid	0.111*	0.067	0.085	0.052
Risk task first	0.022	0.058	0.025	0.047
FED: 1 week	0.053	0.063	0.048	0.051
FED: 2 weeks	-0.004	0.064	0.002	0.052
High Principal	-0.178***	0.033	-0.149***	0.021
Smoker	0.232***	0.062	0.187***	0.050
Constant	0.485*	0.276	0.444**	0.224
Risk error (μ)				
Constant	0.181***	0.011	0.170***	0.008
Time error (ν)				
Constant	0.155***	0.042	0.201***	0.044
N	17500		17500	
log-likelihood	-9163.252		-9117.061	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE H:IV:B: DISCOUNTING FUNCTION ML ESTIMATES
CONCAVE UTILITY AND HETEROGENOUS PREFERENCES

	Model 3 Quasi-Hyperbolic		**Model 4** Weibull	
	Estimate	Std error	Estimate	Std error
Power function parameter (r)				
Age	-0.006	0.005	-0.003	0.005
White	-0.015	0.018	-0.010	0.015
Male	-0.013	0.016	-0.019	0.014
Commerce faculty	0.010	0.016	0.016	0.016
Financial aid	0.031*	0.018	0.023	0.016
Risk task first	-0.004	0.021	-0.002	0.015
Smoker	0.053***	0.020	0.044**	0.020
Constant	0.368***	0.101	0.307***	0.101
Discounting parameter (δ)				
Age	-0.007	0.011	0.007	0.007
White	-0.082	0.065	-0.047	0.029
Male	0.132**	0.056	0.051***	0.020
Commerce faculty	0.018	0.069	-0.002	0.029
Financial aid	0.094	0.063	0.000	0.025
Risk task first	-0.041	0.060	-0.026	0.027
FED: 1 week	0.313***	0.075	0.133**	0.067
FED: 2 weeks	0.265***	0.064	0.203**	0.092
High Principal	-0.139***	0.029	-0.054***	0.013
Smoker	0.206***	0.067	0.080***	0.026
Constant	0.260	0.242	0.005	0.129
Discount parameter (β)				
Age	-0.003	0.003	-0.009	0.079
White	-0.002	0.011	0.288	0.306
Male	0.011	0.009	-0.499	0.360
Commerce faculty	0.002	0.011	-0.026	0.239
Financial aid	-0.005	0.011	0.526*	0.281
Risk task first	-0.020*	0.010	0.922**	0.397
FED: 1 week	0.360	0.279	2.839*	1.452
FED: 2 weeks	0.140	0.213	4.379*	2.455
High Principal	0.005**	0.002	0.076	0.106
Smoker	-0.002	0.011	-0.382	0.706
Constant	1.035***	0.062	2.200	1.712
Risk error (μ)				
Constant	0.184***	0.011	0.187***	0.012
Time error (ν)				
Constant	0.163***	0.046	0.119***	0.039
N	17500		17500	
log-likelihood	-8912.317		-8606.675	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

APPENDIX I

In this appendix I estimate the marginal effect of smoking status across the four discounting specifications using the demographic characteristics and task parameters in Table III of Chapter 5 but under the assumption of linear utility. On a qualitative level (i.e., in terms of the sign and significance of the estimates), the results across linear and concave utility specifications are very similar, except in the WB and, to a lesser extent, the QH models. The different estimates one obtains with the WB model can have important implications for the inferences one draws about the relationship between smoking and discounting behaviour.

Table I:I:A presents the results from the E and H discounting models and Table I:I:B presents the results from the QH and WB models. Table I:I:A shows that the results for the E and H models are qualitatively identical under the assumption of linear utility and under the assumption that RDU and the Prelec weighting function characterise choice under risk (i.e., in Table XV:A in Chapter 5).

TABLE I:I:A: DISCOUNTING FUNCTION ML ESTIMATES
LINEAR UTILITY AND HETEROGENOUS PREFERENCES

	Model 1 Exponential		**Model 2** Hyperbolic	
	Estimate	Std error	Estimate	Std error
Discounting parameter (δ)				
Age	-0.018	0.142	-0.007	0.047
White	-0.319	0.515	-0.189	0.202
Male	0.864*	0.518	0.336*	0.175
Commerce faculty	0.335	0.514	0.080	0.213
Financial aid	0.759	0.615	0.295	0.207
Risk task first	0.001	0.530	0.092	0.190
FED: 1 week	0.201	0.500	0.205	0.195
FED: 2 weeks	0.191	0.526	0.165	0.213
High Principal	-2.515***	0.525	-0.882***	0.113
Smoker	1.512***	0.521	0.615***	0.193
Constant	3.747	2.717	1.722*	0.901
Time error (ν)				
Constant	22.124***	1.661	22.061***	1.700
N	10500		10500	
log-likelihood	-5083.655		-4945.307	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table I:I:B, by contrast, shows that one reaches different qualitative conclusions under the assumption of linear utility; I will first report the differences before

discussing their implications. The only difference in the QH model, across linear and concave utility specifications, is that estimate of the "FED: 1 week" coefficient in the β equation is statistically significant in Table I:I:B, whereas the same coefficient is not statistically significant in Table XV:B in Chapter 5.

TABLE I:I:B: DISCOUNTING FUNCTION ML ESTIMATES
LINEAR UTILITY AND HETEROGENOUS PREFERENCES

	Model 3 Quasi-Hyperbolic		**Model 4** Weibull	
	Estimate	Std error	Estimate	Std error
Discounting parameter (δ)				
Age	-0.021	0.099	0.048**	0.02
White	-0.266	0.46	-0.134	0.125
Male	0.977**	0.408	0.244*	0.129
Commerce faculty	0.045	0.451	-0.039	0.107
Financial aid	0.633	0.613	-0.093	0.139
Risk task first	-0.26	0.404	-0.177	0.183
FED: 1 week	1.281**	0.522	0.389***	0.129
FED: 2 weeks	1.421***	0.468	0.377***	0.132
High Principal	-1.633***	0.39	-0.259***	0.053
Smoker	1.160**	0.559	0.094	0.115
Constant	2.15	1.874	-0.21	0.524
Discount parameter (β)				
Age	-0.012	0.011	-0.134	0.094
White	-0.006	0.034	0.14	0.846
Male	0.037	0.027	-0.724	1.302
Commerce faculty	-0.006	0.038	0.068	0.592
Financial aid	-0.007	0.036	1.017	0.715
Risk task first	-0.059*	0.031	1.352**	0.58
FED: 1 week	0.551**	0.267	-0.042	0.632
FED: 2 weeks	0.168	0.177	0.383	0.496
High Principal	0.032**	0.016	-0.018	0.268
Smoker	-0.004	0.036	0.667	0.661
Constant	1.153***	0.203	3.847**	1.559
Time error (ν)				
Constant	25.314***	2.79	19.540***	1.314
N	10500		10500	
log-likelihood	-4876.24		-4537.238	

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In the WB model, the estimate of the smoker variable in the δ equation is not statistically significant in Table I:I:B, but this coefficient estimate is statistically significant in Table XV:B in Chapter 5 where I incorporate the curvature of the utility function. By contrast, age in the δ equation is statistically significant in Table I:I:B but it is not significant in Table XV:B in Chapter 5. With regard to the β equation, I find that the coefficient estimates of the "Financial aid," "FED: 1 week," and "FED: 2

weeks" variables are not statistically significant in Table I:I:B whereas the same coefficients are statistically significant in Table XV:B in Chapter 5. Finally, the coefficient estimate of the "FED: 1 week" variable actually switches sign across the estimates in Table I:I:B and Table XV:B in Chapter 5.

As Chapter 5 focussed on the relationship between smoking and discounting behaviour, the fact that the smoker variable is not statistically significant in the β and δ equations of the WB model can have important implications for the inferences one draws. Recall that the WB model better characterises all of the data in Chapter 5 than the other discounting specifications. As discussed in Chapter 3, it is common in addiction studies for researchers to compare the efficacy of different discounting models using invalid statistical tests, pick the one which "best" explains the data, and then use it to conduct multivariate analyses. If this approach was adopted with the data in Chapter 5, a researcher would select the WB model for multivariate analysis. If this researcher also failed to adopt the joint estimation approach to time preferences, which incorporates the concavity of the utility function, he may conclude that there is no relationship between smoking and discounting behaviour.

This example highlights some important issues. It is useful to estimate a range of discounting models so as to check the sensitivity and reliability of the results. If the hypothetical researcher above estimated E, H, QH, and WB models assuming linear utility he would be hard-pressed to claim that smoking is not related to discounting. In addition, it is crucial to let theory guide experimental design and analysis. In this context, the hypothetical researcher should be cognisant of the fact that time preferences are defined over time-dated utility flows, not flows of money. If the researcher was determined to use only the "best" discounting function for multivariate analysis, but he estimated it jointly with the curvature of the utility function, he would conclude that smoking and discounting behaviour are related. The preceding discussion may have made it seem like having a good grasp of theory and estimating several models are substitutes; this was not my intention. The ideal hypothetical researcher would combine his knowledge of theory with a desire to investigate relationships of interest across multiple structural theories.

In sum, this appendix showed that one can reach different conclusions about the relationship between smoking and discounting when one abandons the theoretically-motivated joint estimation approach to time preferences which incorporates the concavity of the utility function.

APPENDIX J

In this appendix I present the results from the five other two process mixture models that can be estimated from the four discounting functions used in Chapter 5. Note that RDU and the Prelec PWF are assumed to characterise choice under risk in the joint estimation of these models. The take home message from this appendix is that all discounting models find significant support in the data, even though in each mixture model one function finds significantly more support than the other.

TABLE J:I: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HOMOGENOUS PREFERENCES

	Estimate	Std Error	*p*-value	95% Confidence Interval	
Rank-dependent utility theory					
Power function parameter (r)	0.293***	0.027	0.000	0.240	0.346
PWF parameter (γ)	0.797***	0.025	0.000	0.748	0.846
PWF parameter (η)	0.853***	0.032	0.000	0.790	0.916
Exponential discounting model					
Discounting parameter (δ_E^{mix})	0.815***	0.107	0.000	0.605	1.026
Mixture probability (π^E)	0.636***	0.039	0.000	0.561	0.712
Quasi-Hyperbolic discounting model					
Discounting parameter (δ_{QH}^{mix})	0.105***	0.014	0.000	0.077	0.133
Discounting parameter (β_{QH}^{mix})	0.994***	0.002	0.000	0.991	0.998
Mixture probability (π^{QH})	0.364***	0.039	0.000	0.288	0.439
Error terms					
Risk Error (μ)	0.175***	0.008	0.000	0.158	0.191
Time Error (ν)	0.034***	0.010	0.001	0.015	0.053
N	17500				
log-likelihood	-8775.861				

H_0: $\pi^E = 0.5$, p-value < 0.001

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE J:II: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HOMOGENOUS PREFERENCES

	Estimate	Std Error	*p*-value	95% Confidence Interval	
Rank-dependent utility theory					
Power function parameter (r)	0.301***	0.027	0.000	0.249	0.353
PWF parameter (γ)	0.797***	0.025	0.000	0.748	0.846
PWF parameter (η)	0.860***	0.032	0.000	0.798	0.923
Exponential discounting model					
Discounting parameter (δ_E^{mix})	0.153***	0.017	0.000	0.119	0.186
Mixture probability (π^E)	0.406***	0.034	0.000	0.339	0.474
Weibull discounting model					
Discounting parameter (δ_{WB}^{mix})	0.534***	0.131	0.000	0.279	0.790
Discounting parameter (β_{WB}^{mix})	5.940**	2.722	0.029	0.605	11.275
Mixture probability (π^{WB})	0.594***	0.034	0.000	0.526	0.661
Error terms					
Risk Error (μ)	0.173***	0.008	0.000	0.157	0.189
Time Error (ν)	0.055***	0.014	0.000	0.027	0.084
N	17500				
log-likelihood	-8720.160				

H_0: $\pi^E = 0.5$, p-value = 0.006

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE J:III: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HOMOGENOUS PREFERENCES

	Estimate	Std Error	*p*-value	95% Confidence Interval	
Rank-dependent utility theory					
Power function parameter (r)	0.329***	0.027	0.000	0.277	0.381
PWF parameter (γ)	0.797***	0.025	0.000	0.749	0.846
PWF parameter (η)	0.886***	0.035	0.000	0.819	0.954
Hyperbolic discounting model					
Discounting parameter (δ_H^{mix})	0.717***	0.067	0.000	0.585	0.849
Mixture probability (π^H)	0.634***	0.036	0.000	0.564	0.704
Quasi-Hyperbolic discounting model					
Discounting parameter (δ_{QH}^{mix})	0.119***	0.015	0.000	0.090	0.148
Discounting parameter (β_{QH}^{mix})	0.994***	0.002	0.000	0.990	0.998
Mixture probability (π^{QH})	0.366***	0.036	0.000	0.296	0.436
Error terms					
Risk Error (μ)	0.168***	0.008	0.000	0.153	0.184
Time Error (ν)	0.046***	0.012	0.000	0.021	0.070
N	17500				
log-likelihood	-8752.007				

H_0: $\pi^H = 0.5$, p-value < 0.001

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE J:IV: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HOMOGENOUS PREFERENCES

	Estimate	Std Error	*p*-value	95% Confidence Interval	
Rank-dependent utility theory					
Power function parameter (r)	0.328***	0.026	0.000	0.276	0.379
PWF parameter (γ)	0.797***	0.025	0.000	0.749	0.846
PWF parameter (η)	0.884***	0.034	0.000	0.817	0.952
Hyperbolic discounting model					
Discounting parameter (δ_H^{mix})	0.720***	0.069	0.000	0.585	0.855
Mixture probability (π^H)	0.611***	0.039	0.000	0.535	0.688
Weibull discounting model					
Discounting parameter (δ_{WB}^{mix})	0.072***	0.009	0.000	0.053	0.090
Discounting parameter (β_{WB}^{mix})	1.759***	0.206	0.000	1.355	2.164
Mixture probability (π^{WB})	0.389***	0.039	0.000	0.312	0.465
Error terms					
Risk Error (μ)	0.169***	0.008	0.000	0.153	0.184
Time Error (ν)	0.044***	0.012	0.000	0.021	0.068
N	17500				
log-likelihood	-8703.874				

H_0: $\pi^H = 0.5$, p-value = 0.004

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

TABLE J:V: MIXTURE MODEL ML ESTIMATES
RANK-DEPENDENT UTILITY AND HOMOGENOUS PREFERENCES

	Estimate	Std Error	*p*-value	95% Confidence Interval	
Rank-dependent utility theory					
Power function parameter (r)	0.291***	0.027	0.000	0.238	0.343
PWF parameter (γ)	0.797***	0.025	0.000	0.748	0.846
PWF parameter (η)	0.851***	0.032	0.000	0.788	0.913
Quasi-Hyperbolic discounting model					
Discounting parameter (δ_{QH}^{mix})	0.797***	0.111	0.000	0.579	1.015
Discounting parameter (β_{QH}^{mix})	0.996***	0.003	0.000	0.990	1.002
Mixture probability (π^{QH})	0.609***	0.044	0.000	0.524	0.695
Weibull discounting model					
Discounting parameter (δ_{WB}^{mix})	0.066***	0.009	0.000	0.047	0.084
Discounting parameter (β_{WB}^{mix})	1.720***	0.188	0.000	1.352	2.087
Mixture probability (π^{WB})	0.391***	0.044	0.000	0.305	0.476
Error terms					
Risk Error (μ)	0.175***	0.008	0.000	0.159	0.192
Time Error (ν)	0.034***	0.010	0.000	0.015	0.052
N	17500				
log-likelihood	-8715.090				

H_0: $\pi^{QH} = 0.5$, p-value = 0.012

Results account for clustering at the individual level
* $p<0.10$, ** $p<0.05$, *** $p<0.01$