Addiction and Intertemporal Risk Attitudes

by

Glenn W. Harrison, Andre Hofmeyr, Harold Kincaid, Don Ross and J. Todd Swarthout⌘

June 2019

ABSTRACT

Addictions are typically characterised by cycles of abstinence and relapse over many years, with a variety of resolution states. Economic models of addiction assume intertemporal dependencies in the consumption of addictive goods, thereby incorporating attitudes to risk over time in explanations of addictive behaviour. We are the first to study the intertemporal risk attitudes of addicts. Focussing on smoking behaviour, we compare experimentally elicited risk preferences of addicts, former addicts, and controls. Contrary to an assumption taken up in standard economic models of addiction, smokers do not exhibit intertemporal risk seeking behaviour. Instead, our sample is characterised by high levels of intertemporal risk aversion which varies by smoking intensity and smoking severity in men, but not in women. Our results are the first to demonstrate the role that intertemporal risk attitudes, together with atemporal risk attitudes and discounting behaviour, play in the onset and persistence of addiction.

Keywords: Addiction, intertemporal risk attitudes, atemporal risk preferences, time preferences, tobacco.

JEL codes: C91, D01, D81, D91, I12.

⌘ Department of Risk Management & Insurance and Center for the Economic Analysis of Risk (CEAR), Robinson College of Business, Georgia State University, USA (Harrison); School of Economics, University of Cape Town, South Africa (Hofmeyr, Kincaid); School of Sociology, Philosophy, Criminology, Government, and Politics, University College Cork, Ireland, School of Economics, University of Cape Town, South Africa, and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA (Ross); and Department of Economics, Andrew Young School of Policy Studies, Georgia State University, USA (Swarthout). Harrison is also affiliated with the School of Economics, University of Cape Town. E-mail contacts: gharrison@gsu.edu, andre.hofmeyr@uct.ac.za, harold.kincaid@uct.ac.za, don.ross931@gmail.com and swarthout@gsu.edu. We are grateful to the National Research Foundation (NRF) of South Africa for funding this research; in all other respects the funder had no involvement in the research project.

1. INTRODUCTION

Economic models of addiction highlight the importance of atemporal risk preferences, time preferences, and intertemporal risk preferences for the onset, persistence, and resolution of addiction. The relevance of these preferences in the context of addiction is clear: the consumption of addictive goods occurs under conditions of risk and uncertainty (*atemporal risk preferences*); it involves an intertemporal trade-off between current benefits and future costs (*time preferences*); and there is serial correlation in addiction outcomes given that present consumption tends to influence future consumption while simultaneously increasing the risk of a decline in long-term welfare (*intertemporal risk preferences*).

Intertemporal risk aversion[1] refers to any aversion to variability of outcomes *over* time, just as atemporal risk aversion refers to any aversion to variability of outcomes *at* a point in time. Smoking addictions are often associated with multiple attempts to quit, usually successfully for short periods of time, accompanied by eventual relapse. It is precisely this standard profile of a smoking addict that interacts with attitudes to risk over time. Measures of risk attitudes at a point in time logically need have no relation at all with risk attitudes over time. Hence it is critical to differentiate these two measures of risk attitudes to see the complete way in which risk attitudes in general interact with smoking behaviour.

Despite the theoretical importance of these preferences, they have received little attention in the empirical economic literature on addiction. Harrison, Hofmeyr, Ross and Swarthout (HHRS) [2018] review the experimental literature on atemporal risk preferences, time preferences, and smoking behaviour, and find that most studies use inappropriate statistical methods or preference elicitation mechanisms that lack incentive compatibility. Furthermore, we know of no experimental or other empirical studies that analyse the relationship between *intertemporal* risk preferences and addiction.

[1] The literature on intertemporal risk preferences emerged from the literature on multi-attribute utility theory (see Keeney and Raiffa [1976] for a review). A multi-attribute utility function captures the idea that agents may take into account the 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, the times at which different goods or amounts of money are received can be regarded as 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. Epstein and Zin [1989] also introduced a concept of intertemporal risk aversion that does not rely on a multi-attribute utility function.

We evaluate an incentive-compatible experiment designed to elicit the atemporal risk preferences, time preferences, and intertemporal risk preferences of a sample of student (n = 145) and staff (n = 111) smokers, ex-smokers, and non-smokers at the University of Cape Town (UCT) in 2016-2017. We adopt a full information maximum likelihood statistical approach to estimate a structural model of intertemporal risk preferences jointly with a rank-dependent utility (RDU) model of atemporal risk preferences, due to Quiggin [1982], and a quasi-hyperbolic (QH) model of time preferences, due to Phelps and Pollak [1968] and Laibson [1997]. We also estimate a range of alternative specifications to test the robustness of our results.

We make a number of contributions. First, we identify significant heterogeneity in intertemporal risk preferences but find, contrary to the assumption employed by standard economic models of addiction, that smokers do not exhibit intertemporal risk seeking behaviour. Instead, our sample is characterised by intertemporal risk aversion, which does not differ significantly according to smoking status, but does differ according to the smoking intensity and smoking severity of men.

Second, we replicate the finding of HHRS that atemporal risk aversion does not differ according to smoking status and smoking intensity, measured by the number of cigarettes smoked per day, while extending this null result to a measure of smoking severity, the Fagerström [2012] Test for Cigarette Dependence. Third, we replicate the finding of economically and statistically significant differences in the time preferences of smokers and non-smokers, and add nuance to this result by incorporating ex-smokers in the sample: ex-smokers discount at a level between smokers and non-smokers. Finally, we identify a positive relationship between smoking intensity and discounting behaviour that makes it harder for heavier smokers to quit because the long-term costs of continuing to smoke and the long-term benefits that result from successful abstention are discounted heavily.

Section 2 discusses the theory of intertemporal risk preferences and their importance in economic models of addiction. Section 3 describes our experimental design and presents summary statistics for the sample. Section 4 outlines our statistical approach for jointly estimating atemporal risk preferences, time preferences, and intertemporal risk preferences, along with measures of smoking behaviour. Section 5 presents the results and Section 6 concludes.

2. THEORY

HHRS conduct a detailed review of the literature on atemporal risk preferences, time preferences, and smoking behaviour so we limit the discussion to the relationship between intertemporal risk preferences and addiction.

A. Intertemporal Risk Attitudes

Intertemporal risk preferences are determined by properties of the intertemporal utility function. Consider the following intertemporal choice model:

$$U(x_0, x_1, x_2, \ldots) = E\left[\theta\left(\sum_{t=0}^{n} D_t\, u(x_t)\right)\right], \tag{1}$$

where $u(x_t)$ is the atemporal utility function over money at time t, $D_t > 0$ is a discount factor for time horizon t, and θ is the identity function when $U(\cdot)$ is additively separable, but we allow for departures from this assumption below.[2]

Richard [1975] is credited with introducing intertemporal risk preferences to the economic literature, although the concept apparently first appeared in de Finetti [1952]. Richard [1975] basically extended the notion of risk preferences over one variable to risk preferences over multiple variables and referred to the latter as multivariate risk aversion.[3]

To flesh out this idea, Table 1 includes two intertemporal lotteries (A and B) that yield outcomes in two time periods. Under intertemporal lottery A, you flip a coin and if it lands on heads you receive $50 today *and* $5 in 14 days, but if it lands on tails you receive $5 today *and* $50 in 14 days. By contrast, under intertemporal lottery B, you flip a coin and if it lands on heads you receive $50 today *and* $50 in 14 days, but if it lands on tails you receive $5 today *and* $5 in 14 days. The outcomes in intertemporal lottery A are negatively serially correlated whereas the outcomes in intertemporal lottery B are positively serially correlated.

---

[2] The expectation $E(\cdot)$ in (1) typically denotes a probability-weighted average, but to incorporate the possibility that decision makers subjectively distort objective probabilities, we use $E(\cdot)$ to denote a decision-weighted average, in the terminology of RDU theory.

[3] Researchers in the field of intertemporal risk preferences typically employ the risk averse component of these preferences in their terminology. Keeney [1973] uses the term "conditional risk aversion," Richard [1975] refers to "multivariate risk aversion," Epstein and Tanny [1980] define "correlation aversion," Strzalecki [2013] employs "long-run risk aversion," and Andersen, Harrison, Lau and Rutström [2018] use "intertemporal risk aversion" or "intertemporal correlation aversion." We prefer the term "intertemporal risk preferences" because it does not presuppose an aversion to lotteries with positive serial correlation.

Table 1
Intertemporal Lotteries and Intertemporal Risk Preferences

State of nature	Probability	Intertemporal Lottery A	Intertemporal lottery B
Heads	0.5	\$50 today *and* \$5 in 14 days	\$50 today *and* \$50 in 14 days
Tails	0.5	\$5 today *and* \$50 in 14 days	\$5 today *and* \$5 in 14 days

If a decision maker chooses intertemporal lottery A over intertemporal lottery B this is evidence of intertemporal risk aversion because, as Richard [1975, p. 12] remarked, "… the decision maker 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.'"[4] If the decision maker is indifferent between the lotteries, she is intertemporally risk neutral, and if the decision maker prefers intertemporal lottery B to intertemporal lottery A this is indicative of intertemporal risk seeking behaviour.

Richard [1975] shows that the sign of the cross partial derivatives of the intertemporal utility function determines preferences toward serially correlated lotteries. Specifically, if $\partial^2 U(\mathbf{x})/(\partial x_{t-1}\partial x_t) \leq 0$ in (1) the decision maker is intertemporally risk averse. In words, if the intertemporal utility function's cross partial derivative is non-positive then the decision maker prefers lotteries where the outcomes are negatively serially correlated because the marginal utility of current consumption is decreasing in past consumption. By contrast, if $\partial^2 U(\mathbf{x})/(\partial x_{t-1}\partial x_t) = 0$ the decision maker is intertemporally risk neutral, whereas if $\partial^2 U(\mathbf{x})/(\partial x_{t-1}\partial x_t) \geq 0$ the decision maker is intertemporally risk seeking. Thus, the form of a decision maker's intertemporal utility function determines her intertemporal risk preferences, just as the form of a decision maker's atemporal utility function determines her atemporal risk preferences under expected utility theory (EUT).[5] Critically, the sign of $\partial^2 U(\mathbf{x})/(\partial x_{t-1}\partial x_t)$ has no formal or economic connection to the sign of $\partial^2 U(\mathbf{x})/(\partial x_{t-1}^2)$ or $\partial^2 U(\mathbf{x})/(\partial x_t^2)$.

[4] Andersen, Harrison, Lau and Rutström [2018, p. 538] provide another intuitive definition for intertemporal risk aversion when drawing an analogy between atemporal risk aversion and intertemporal risk aversion: "The [intertemporally risk] averse individual prefers to have non-extreme payoffs *across* periods, just as the [atemporally] risk averse individual prefers to have non-extreme payoffs *within* periods."

[5] Much as the literature on choice under atemporal risk has evolved to incorporate rank and sign dependence, as has the literature on intertemporal risk preferences (see Fishburn [1984] and Miyamoto and Wakker [1996]).

The standard model of intertemporal choice in economics employs an additively-separable intertemporal utility function, so that $\theta(\cdot)$ is the identity function in (1). Additive separability implies intertemporal risk neutrality because consumption at different points in time is independent, so the cross partial derivatives of the intertemporal utility function are necessarily zero. Thus, even though a decision maker may be risk averse, risk neutral, or risk seeking over *atemporal* lotteries, an additively-separable intertemporal utility function[6] yields *intertemporal* risk neutrality.[7]

The example in Table 1 highlights the restrictive nature of the additively-separable assumption. In this example, it may appear that a preference for intertemporal lottery A over intertemporal lottery B can be rationalised by the curvature of the atemporal utility function over outcomes $u(\cdot)$, which would imply that there is no need to abandon the assumption of additive separability and invoke intertemporal risk aversion. To see that curvature of the atemporal utility function cannot account for a preference of intertemporal lottery A over intertemporal lottery B, assume $S < L$ and $t < t + \tau$. Consider the "safe" intertemporal lottery A where a decision maker receives $(L_t, S_{t+\tau})$ with probability p and $(S_t, L_{t+\tau})$ with probability $1 - p$; in Table 1, $p = 0.5$, $L = \$50$, $S = \$5$, $t = 0$, and $\tau = 14$ days. Under the assumption of additive separability, the stochastic discounted utility (SDU) of intertemporal lottery A is

$$\mathrm{SDU_A} = \omega(p) \times [D_t\, u(L_t) + D_{t+\tau}\, u(S_{t+\tau})] + [1 - \omega(p)] \times [D_t\, u(S_t) + D_{t+\tau}\, u(L_{t+\tau})],$$

where $\omega : p \rightarrow [0, 1]$ with $\omega'(p) > 0$, $D_t > 0$ is a discount factor for time horizon t, and $u : x \in \{S, L\} \rightarrow \mathbb{R}$ with $u'(x) > 0$.

Now consider the "risky" intertemporal lottery B where a decision maker receives $(L_t, L_{t+\tau})$ with probability p and $(S_t, S_{t+\tau})$ with probability $1 - p$. Under the assumption of additive separability, the SDU of inertemporal lottery B is

$$\mathrm{SDU_B} = \omega(p) \times [D_t\, u(L_t) + D_{t+\tau}\, u(L_{t+\tau})] + [1 - \omega(p)] \times [D_t\, u(S_t) + D_{t+\tau}\, u(S_{t+\tau})].$$

[6] The term "intertemporal utility function" encompasses both deterministic and stochastic choice contexts. To emphasise the stochastic nature of our intertemporal risk preference task, we instead use the term "stochastic discounted utility" below.

[7] The fact that an additively-separable intertemporal utility function yields intertemporal risk neutrality has the unfortunate implication that the intertemporal elasticity of substitution equals the inverse of atemporal risk attitudes. In economic models of addiction where the intertemporal utility function is not additively separable, this link between instantaneous risk attitudes and the intertemporal elasticity of substitution is broken (see Bommier [2007], Bommier, Kochov and Le Grand [2017] and Andersen, Harrison, Lau and Rutström [2018]).

Taking the difference in the SDU of intertemporal lottery A and intertemporal lottery B we find that

$$\text{SDU}_\text{A} - \text{SDU}_\text{B} = \text{D}_{t+\tau}\,[2\omega(p) - 1][u(S_{t+\tau}) - u(L_{t+\tau})].$$

Evaluating this expression, $\text{D}_{t+\tau}$ is strictly positive and $[u(S_{t+\tau}) - u(L_{t+\tau})]$ is strictly negative, so the sign of $\text{SDU}_\text{A} - \text{SDU}_\text{B}$ is determined by $[2\omega(p) - 1]$. Thus, under the assumption of additive separability, curvature of the *atemporal* utility function $u(\cdot)$ has *no effect on the ranking of intertemporal lotteries*. The only factor that can affect this ranking is the function $\omega(\cdot)$. Specifically, $\text{A} \succ \text{B} \Leftrightarrow \omega(p) < 0.5$. This condition must hold for any value of p to rationalise a preference for intertemporal lottery A over intertemporal lottery B. In the case of EUT, $\omega(p) = p$, so if $p = 0.5$, as per the example in Table 1, this implies that $\omega(p) = 0.5$, $2\omega(p) - 1 = 0$, and $\text{SDU}_\text{A} = \text{SDU}_\text{B}$. Thus, EUT, coupled with additive separability, implies indifference between intertemporal lottery A and intertemporal lottery B. Consequently, to account for a preference of A over B when $p = 0.5$ one must adopt a non-EUT model of choice under atemporal risk, such that $\omega(0.5) < 0.5$, or abandon the assumption of additive separability. In Section 5 we show that the *atemporal* risk preferences of our sample do indeed depart from EUT, on average, but that this cannot account for our subjects' choices over intertemporal lotteries without also incorporating *intertemporal* risk aversion.

B. Implications for Models of Addiction

Economic models of addiction abandon the additive-separability assumption to incorporate intertemporal dependencies in the consumption of addictive goods so it is important to understand the implications of this departure from the standard model for the intertemporal risk preferences of decision makers.

The first well-known economic model of addiction was developed by Becker and Murphy (BM) [1988]. This model assumes that agents consume addictive and non-addictive goods, where consumption of the former increases the decision maker's stock of addictive capital. The model also assumes that the agent's intertemporal utility function is not additively separable over time in the addictive and non-addictive goods, because their marginal utilities are influenced by the decision maker's 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. Specifically, the BM model assumes that higher consumption of the addictive good in the past raises the marginal utility of present consumption, implying that the more the addict has consumed, the greater the benefit from consumption now. 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.[8] This assumption was central to economic models of addiction in the tradition that descended from BM, because it provides a 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 BM model and its extensions are special cases of more general models of consumption *habit formation*, which, as Bommier and Rochet [2006, p. 725-726] recognise, place strong restrictions on the intertemporal utility function through the assumption of adjacent complementarity. Specifically, 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, $\partial^2 U(\mathbf{x})/(\partial x_{t-1}\partial x_t) > 0$, implying that agents in these models are *typically* intertemporally risk seeking.[9]

Psychologists have not generally regarded the BM model as providing an accurate specification of addiction, precisely because it mispredicts the dynamics of the typical life-course of an addiction. Almost all addicts eventually achieve abstinence or controlled, moderate consumption, and most do so without clinical intervention or therapy (see Heyman [2009]). Orphanides and Zervos [1995] generalise the deterministic BM model by incorporating a stochastic element in addictive consumption and thereby allow the hypothesis that the addict makes an initial forecasting error when she chooses to experiment with the

[8] Adjacent complementarity implies that anything that 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 that 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 see Chaloupka and Warner [2001]. For critiques of this literature see Ferguson [2000] and Baltagi [2007].

[9] Although stochastic economic models of addiction typically assume, sometimes implicitly, that agents are intertemporally risk seeking, Bommier and Rochet [2006] and Lichtendahl, Chao and Bodily [2012] show that intertemporal risk aversion is in fact compatible with habit formation.

addictive target. Later self-control could then be explained as resulting from correction of this initial error. But this directly conflicts with the fact that most addicts also achieve recovery only after first experiencing multiple periods of attempts at control, during which anhedonic costs of withdrawal are fully paid, followed by relapse. It is plausible that people make initial forecasting errors preceding addiction. However, the psychological literature has never endorsed the suggestion that people experience the full suite of addictive onset, effects, increased tolerance, welfare losses, and withdrawal symptoms, then repeat their initial forecasting errors, and indeed do so multiple times.[10]

A more recent wave of economic models of addiction aligns with the general view of psychologists and psychiatrists that addiction is not "rational" in the sense of BM, that is, that it necessarily implies intertemporal and perhaps even synchronous, preference ambivalence. There are two general strategies for representing ambivalence found in this literature. One approach, taken up by Bernheim and Rangel [2004] and others[11], is to model choice as resulting from competition between dual systems with differing preferences. Though this general idea has been widely promoted, particularly by Kahneman [2011], it lacks independent neuropsychological evidence or an accepted canonical model (see Grayot [2019]).

The limited influence of dual system models on the applied economic literature on addiction policy and interventions might thus be explained by perception that it is not clear how such models might best be integrated into the more general economic theory of the consumer. This is reflected in their unclear implications with respect to intertemporal risk attitudes. Bernheim and Rangel [2004], for example, need not appeal to adjacent complementarity if it is supposed that an atemporally risk seeking system becomes more likely to control behaviour as addiction takes hold. The model is thus compatible with representing observed choices of the whole person as reflecting intertemporal risk aversion. But this again leads to ambiguity in modelling the risk preferences, both atemporal and intertemporal, of an agent who is trying to escape addiction. Psychologists and clinicians

[10] Chaiton et al. [2016] use the longitudinal Ontario Tobacco Survey to estimate the average number of quit attempts it takes to quit smoking successfully. They evaluate four different statistical methods and find that the average number of quit attempts is 6.1 using the standard cross-sectional "recalled lifetime quit attempts" metric. This average rises to 30 quit attempts using their preferred method to estimate the probability of a successful quit attempt on the basis of observed quit rates, which allows for smoking relapse over time.
[11] Examples are Laibson [2001], Loewenstein, O'Donoghue and Rabin [2003], Bénabou and Tirole [2004], Benhabib and Bisin [2004], Fudenberg and Levine [2006][2011][2012], and Gul and Pesendorfer [2007].

generally suppose that unsuccessful attempts at quitting generate learning that typically eventually leads to cessation or controlled consumption (see West and Brown [2013]). Dual system models shed no light on whether such learning might involve adjustment in atemporal risk preferences, intertemporal risk preferences, or their interaction.

A second strategy for representing ambivalent choice behaviour is to suppose that a unified agent that unambiguously prefers not to be addicted faces special costs imposed by exogenous temptations that can exceed its budget for self-control, and that addiction dynamically increases these costs over time (see Loewenstein [1999] and Laibson [2001]).[12] Since these models assume that addicts and non-addicts need not systematically differ in their preferences at all, they are silent on possible influences of risk preferences except insofar as these effect atemporal costs and benefits.

In general, psychologists do not view addiction as a variety of habit formation, in either the everyday sense or in the economist's technical sense. Certainly addicts often form habits, for example, in being willing to pay costs to stick to their usual brand of cigarettes. And furthermore, familiar experiences associated with such habits can cue addictive cravings through mechanisms of associative learning. But addiction is not regarded by psychologists or clinicians as a kind of habit.[13] This reflects their emphasis, not substantially taken up in the economic literature, on the dynamics by which addicts learn about *both* the costs and benefits of addiction through becoming addicted, and by which they typically learn to manage or cease addictive consumption. Learning, as a temporal process, plausibly interacts with intertemporal risk preferences, but empirical investigation of such interaction has not been undertaken.

The phenomenology of addiction is likewise ambiguous where intertemporal risk attitudes are concerned. Ainslie [2001] distinguishes between two common patterns in addicts' quitting strategies. Many addicts seek to avoid temporally concentrated utility

[12] Loewenstein, O'Donoghue and Rabin [2003] combine elements of both strategies.

[13] Addiction arises when the simple conditioned learning system implemented in the primitive ventral striatum of the brain, that humans share with other mammals, learns that a stereotyped action sequence (e.g., taking out the cigarette pack, taking out the cigarette, lighting it) reliably produces a strong reward that the brain perceives as varying stochastically across the whole estimation interval that the system scans. This sets a prediction problem for it that the system cannot solve. But the brain cannot stop *trying* to solve this prediction problem given a behaviourally or perceptually associated cue. The stereotyped behaviour sequences might be regarded as a kind of habit. But this is not habitual *consumption* of the kind associated in economic models with intertemporal risk seeking behaviour.

discontinuities by tapering consumption. For examples, smokers often formulate and implement consumption schedules that increase the intervals between cigarettes. Another pattern is for a major "cold turkey" cessation event to be prepared in advance, often made salient by a staged public ceremony and accompanied by "hard" or "soft" commitment devices.[14] Since the tapering strategy prolongs addiction, relative to the cold turkey strategy, but avoids a transitional consumption shock, one might hypothesise that addicts would favour the former to the extent that they are intertemporally risk averse. Ainslie [2001] reviews different experiential obstacles and learning dynamics associated with the two strategies, but does not frame these by reference to intertemporal risk preferences.

In this context, a natural first step in empirical investigation is to compare elicited intertemporal risk preferences among samples of people who have not been addicted, people who are current addicts, and people who are former addicts.

3. EXPERIMENTAL DESIGN AND SUMMARY STATISTICS

A. General Procedures

After receiving ethics approval and permission to access students and staff at UCT, we sent out emails describing the study to all students and approximately 20% of staff members.[15] Given our interest in smoking behaviour the emails included a web link to an online, sign-up survey that contained the following three questions about smoking: 1) "Have you ever smoked cigarettes?" (Yes/No); 2) "If you answered Yes to question 1), have you smoked at least 100 cigarettes in your life?" (Yes/No); 3) "If you answered Yes to question 1), do you currently smoke cigarettes, occasionally or regularly?" (Yes/No). A pool of over 2,000 students and 220 staff members completed the sign-up survey to take part in the study.

---

[14] The literature has come to use the phrase "hard commitment" to designate establishment of an inescapable material cost to commitment failure, e.g., payment of a cash fine. "Soft commitment" then refers to public pledges, which might lead to loss of face or social status if the pledge is breached. Both forms of commitment sometimes lead to subsequent consumption being carried on evasively. Evasion itself typically implies higher costs of consumption.

[15] UCT prevents researchers from emailing all staff members (approximately 6,700 people) because they do not want staff to be inundated with requests to participate in research studies. Consequently, researchers are given a spreadsheet containing basic information on all staff members, e.g., faculty, pay class, gender, etc., and are instructed to select approximately 20% of the people on the spreadsheet so that emails can be sent out to them. We were advised not to select staff members in the lowest pay classes (pay classes 1 – 4) because they do not have regular access to email, and we chose not to include any staff members from UCT's satellite campuses. Random selection of the remaining staff members produced the 20% sample that was used for recruitment.

We sampled from the student and staff groups separately, and the two groups were never mixed within a given laboratory session.

For students, we defined two groups from which to randomly select study participants: a smoker group (defined by answering yes to questions 1, 2, and 3 above) of approximately 500 people, and a group of approximately 1,000 people comprising ex-smokers (defined by answering yes to questions 1 and 2 but no to question 3) and non-smokers (defined by answering no to question 1).[16] Those people who were randomly selected to take part in the study (260 smokers and 160 ex-smokers and non-smokers) were added to a dedicated, restricted-access site on the university's virtual learning environment which allowed them to sign up for an experimental session that did not conflict with their academic timetable. A total of 8 sessions were conducted with students between November 2016 and March 2017. Given the limited number of staff members who applied to take part in the study, all of the 220 people who filled in the sign-up survey were added to a dedicated, restricted-access site on the university's virtual learning environment so that they could sign up for an experimental session that suited their work schedule. These 5 sessions were conducted in August 2017. In total, 256 people participated in the lab sessions: 145 students and 111 staff members.

The experiment took place in a computer lab at UCT that had been set up to run the experimental software developed by us, which is discussed in more detail below. Subjects were separated by partitions and were asked not to talk to each other. We employed a team of three research assistants (RAs) to help run the sessions, administer payments, and answer questions.

Upon arrival at the lab, subjects were randomly allocated to a computer terminal and were asked to read and sign a consent form. When everyone had signed the consent form, an RA went through a short presentation[17] which provided a description of what would take place in the session. At the end of the introductory presentation, subjects were asked to read atemporal risk preference task instructions and to raise their hands when they were finished.

[16] There were approximately 600 students who answered yes to question 1 but no to question 2. They were excluded from the sampling frame so that we could focus on smokers, ex-smokers who had smoked more than 100 cigarettes, and non-smokers who had never smoked cigarettes.

[17] Appendix A includes the introductory presentation, the atemporal risk preference task instructions, the time preference task instructions, and the intertemporal risk preference task instructions.

When a subject raised her hand, an RA asked the subject to put on a set of headphones and watch a video[18] that we developed to further explain the task and familiarise the participant with the screen-based, decision-making environment. This approach was adopted for all of the tasks: subjects received written and audio-visual instructions and were required to go through both of them before completing the task. After finishing the atemporal risk preference task video, the subject raised her hand and was then allowed to complete the task. After completing the choice task itself, the subject rolled two 10-sided dice to randomly select one of the choices that was made, and then rolled the two 10-sided dice again to resolve the chosen lottery. An RA recorded the subject's earnings for the task on a payment receipt that would be used to determine the subject's final earnings at the end of the experimental session.

The subject was then asked to read the written instructions for the time preference task before proceeding to the audio-visual instructions.[19] When the video was finished, the subject completed the task and then rolled a 20-sided die and a 4-sided die to randomly select one of the choices that was made. An RA recorded this amount and the payment date on the subject's payment receipt.

Following the completion of the time preference task, subjects read through instructions for the intertemporal risk preference task before watching the audio-visual instructions.[20] After completing the task, the subject rolled a 4-sided die and a 10-sided die to randomly select one of the choices that was made before rolling a 10-sided die to resolve the chosen intertemporal lottery. An RA then recorded the amounts and payment dates for this choice on the subject's payment receipt.

Subjects then completed a task that elicited their subjective beliefs about the mortality risks of smoking. After reading through and watching audio-visual instructions explaining the task, subjects responded to 10 questions, e.g., "For adults 35 years of age and older, what percentage of deaths from lung cancer are associated with smoking in the United States between 2005 and 2009?"[21]

---

[18] See https://cear.gsu.edu/files/2019/03/Risk_Demo.mp4.
[19] See https://cear.gsu.edu/files/2019/03/Time_Demo.mp4.
[20] See https://cear.gsu.edu/files/2019/03/Risk_and_Time_Demo.mp4.
[21] The subjective beliefs task is not the focus of our analysis.

When a subject had finished all four tasks, she then completed a questionnaire on the computer which included 10 questions on demographic and socio-economic characteristics as well as a number of modules designed to gather information on smoking behaviour and other potentially co-occurring mental disorders, e.g., anxiety, depression, and alcohol use disorder. With regard to smoking behaviour, we included the tobacco and nicotine use module from the National Epidemiological Survey on Alcohol and Related Conditions (NESARC) described by Grant and Dawson [2006]; the Fagerström Test for Cigarette Dependence (FTCD) described by Heatherton, Kozlowski, Frecker and Fagerström [1991] and Fagerström [2012]; and the diagnostic criteria for tobacco use disorder and tobacco withdrawal in the Diagnostic and Statistical Manual of Mental Disorders, Fifth Edition (DSM-5) documented in American Psychiatric Association [2013].

While subjects completed the questionnaire, an RA determined their total earnings for all of the tasks. All subjects received a show-up fee of R40. The show-up fee together with earnings for the atemporal risk preference task and subjective beliefs task were paid out immediately in cash, and earnings for the time preference task and intertemporal risk preference task were paid out on the dates corresponding to the subject's choices on the randomly selected questions. Delayed payments were effected via electronic transfer and subjects received a payment notification on their cell phones as soon as the transfers took place. Such transfers are a common means of payment in South Africa and were used to reduce the transaction costs which subjects would have had to incur by coming to collect their delayed payments from us. Experimental sessions lasted approximately 1.5 hours and subjects earned R920 (roughly $150 at purchasing power parity (PPP) at the time) on average.

B. Atemporal Risk Preference Task

The atemporal risk preference task interface was based on Hey and Orme [1994]. It presented subjects with a choice between two lotteries on a screen, displayed in Figure 1 as pie charts with accompanying text that listed the probabilities and monetary amounts of the prizes. Subjects made 90 choices in the task and then rolled dice to randomly select one choice for payment.

The 90 lottery pairs were drawn from the designs of Wakker, Erev and Weber [1994], Loomes and Sugden [1998], Cox and Sadiraj [2008, p. 33], and Harrison, Martínez-Correa and Swarthout [2015]. These lottery pairs were chosen to provide good coverage of the probability space, to facilitate the estimation of non-EUT models of choice under atemporal risk, to investigate the calibration puzzle of Hansson [1988] and Rabin [2000], and to determine whether subjects satisfy the reduction of compound lotteries axiom.[22] Each lottery pair was drawn randomly, without replacement, from this battery and presented to subjects sequentially. The task used prize magnitudes between R0 and R700 ($0 - $112 at PPP) and probabilities which varied in increments of 0.05 between 0 and 1.

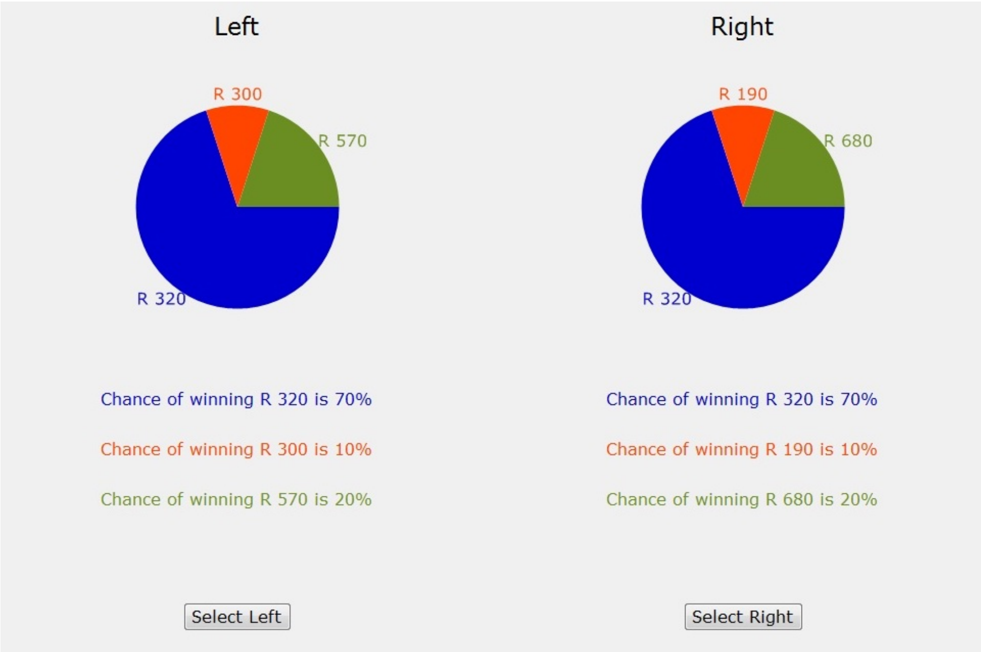

Figure 1: Atemporal Risk Preference Task Interface

C. Time Preference Task

The time preference task presented subjects with choices between smaller, sooner (SS) and larger, later (LL) rewards, illustrated in Figure 2. On each screen subjects made 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 to show subjects when they would receive the amounts of money they chose.

Following Coller and Williams [1999], two front end delays (FEDs) to the SS rewards were used: zero days and 7 days. This design allows one to hold subjective transaction costs

[22] Appendix B includes a detailed discussion of the lottery pairs that were used in the atemporal risk preference task.

constant for the SS and LL rewards at the positive 7-day FED. It also facilitates estimation of the parameters of a QH or β-δ 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 7-day FED allows one to then identify the long-term discounting parameter δ.[23] Subjects in an experimental session were exposed to both of these FED treatments.

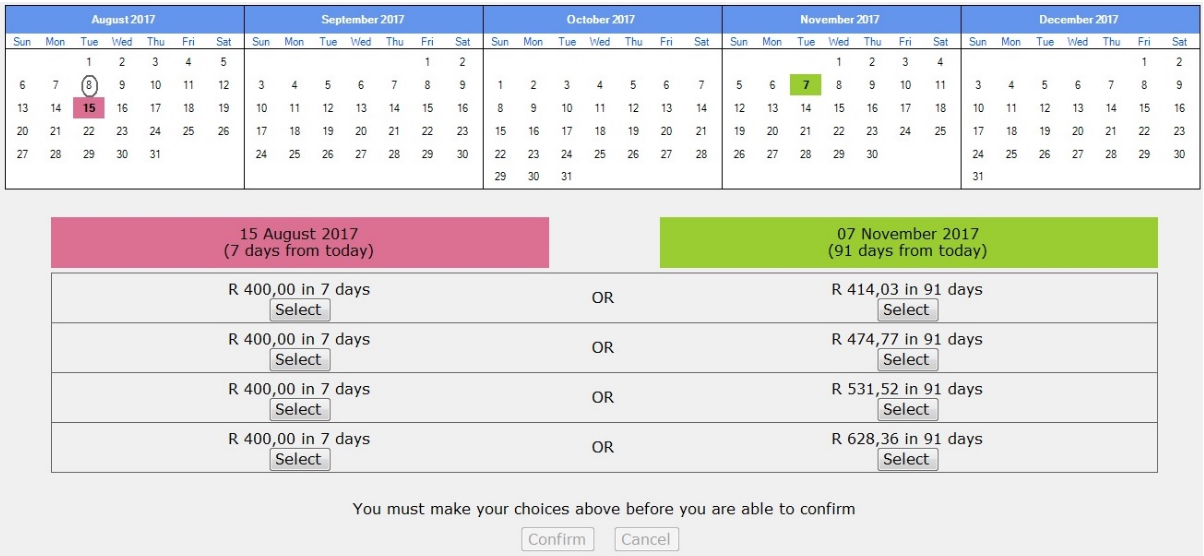

Figure 2: Time Preference Task Interface

Two principals (R250 and R400; $40 and $64 at PPP), four time horizons (7, 14, 42, and 84 days), and nominal annual interest rates between 5% and 250% were used in the time preference task. These parameters, together with the FEDs, 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 time preference task, the subject rolled dice to randomly select one of these choices for payment.

D. Intertemporal Risk Preference Task

The intertemporal risk preference task interface was based on Andersen, Harrison, Lau and Rutström (AHLR) [2018]. On each screen, illustrated by Figure 3, it presented subjects with a choice between two risky profiles of outcomes that were paid out at different points in time (viz., intertemporal lotteries). Probabilities were communicated by text and pie

[23] To easily distinguish between the two parameters of the QH discounting model, we refer to the "present-bias" discounting parameter β and the "long-term" discounting parameter δ.

charts, prizes were listed numerically, and the dates on which the prizes would be paid out were displayed in text and on a calendar.

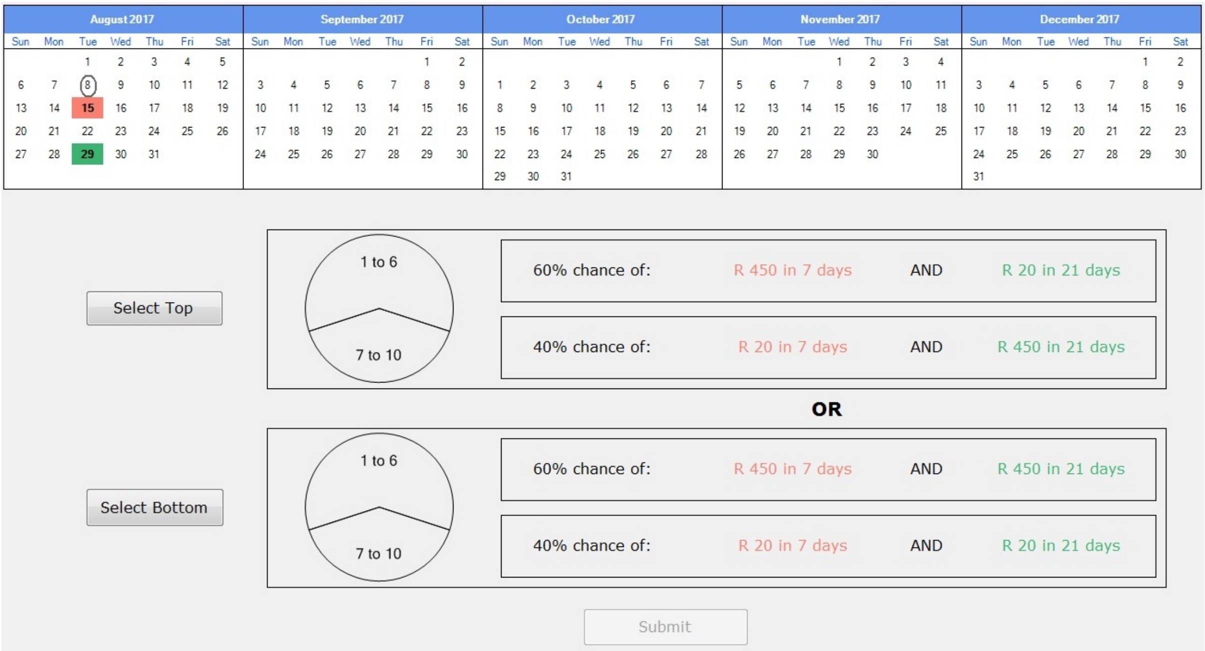

Figure 3: Intertemporal Risk Preference Task Interface

The pairs of intertemporal lotteries were structured in the following way. For a particular pair, lottery A assigned a probability of, say, 0.6 to receiving a larger amount L_t at time t *and* a smaller amount $S_{t+\tau}$ at time $t+\tau$ $(L_t, S_{t+\tau})$ and a probability of 0.4 to receiving the smaller amount S_t at time t *and* the larger amount $L_{t+\tau}$ at time $t+\tau$ $(S_t, L_{t+\tau})$. Lottery B, by contrast, assigned a probability of 0.6 to receiving L_t and $L_{t+\tau}$ and a probability of 0.4 to receiving S_t and $S_{t+\tau}$. In this example, lottery A is the "safe" intertemporal lottery because the subject always earns $L + S$, whereas lottery B is the "risky" intertemporal lottery because the subject either earns $2L$ or $2S$. We constructed 40 of these intertemporal lotteries, broken down into 4 sets of 10, with prizes $S_t = S_{t+\tau}$ and $L_t = L_{t+\tau}$ in each set. Each set of 10 intertemporal lotteries included prizes with probability $p(L_t, S_{t+\tau}) = p(L_t, L_{t+\tau})$ starting at 0.1, and increasing by 0.1 until the last choice was between two degenerate intertemporal lotteries. Using the example above, the last choice in this set was between lottery A that pays $(L_t, S_{t+\tau})$ with certainty and lottery B that pays $(L_t, L_{t+\tau})$ with certainty; lottery B clearly dominates lottery A in this pair and is a test of subject comprehension or monotonicity of preferences.

To construct our battery of intertemporal lotteries, we used a 7-day FED to the sooner reward, two time horizons of 14 days and 42 days between the rewards, and the following two sets of larger (*L*) amounts and smaller (*S*) amounts: (R450, R20; $72, $3 at PPP) and (R260, R10; $42, $1.50 at PPP). Each intertemporal lottery pair was drawn at random, without replacement, from this battery and presented to subjects sequentially. The order in which the intertemporal lotteries appeared (i.e., whether the "safe" lottery appeared as the "Top" choice or the "Bottom" choice in Figure 3) varied randomly across screens. At the end of the task, the subject rolled dice to randomly select one choice for payment.

E. Summary Statistics

Table 2 presents summary statistics for the sample of 256 subjects. The average age in the sample is approximately 30 years old, 27% of the sample is White[24], 43% is male, and, coincidentally, 43% of the sample is made up of staff. Subjects were asked to rate their current financial situation on a scale of 1 to 5, where 1 represented "very broke" and 5 represented "in very good shape." The mean response of 2.80 implies that, on average, subjects were neither broke nor in good shape at the time of their experimental session. Non-smokers make up 52% of the sample, ex-smokers constitute 12% of the sample, and smokers comprise the remaining 36% of the sample.[25]

An estimate from the South African National Health and Nutrition Examination Survey of the mean number of cigarettes smoked per day by current smokers in South Africa is 7.4. For people residing in the Western Cape, where the present study was conducted, the mean is 8.5 (see Shisana et al. [2013, p. 114-115]). The smokers in our study reported the average number of cigarettes they typically smoked in a day, and the mean across all responses was 8.129 with a standard deviation of 5.317.

[24] Designation of population groups or "races" follows the traditional categorisation in South Africa that is still employed in affirmative action and related policies, notwithstanding recognition that it involves cultural and historical discriminations that are without biological significance. Approximately 30% of the sample is Black and 31% is Coloured, a culturally salient population group in South Africa composed of individuals of mainly Malaysian and Indonesian descent who speak Afrikaans as a first language. Of the remaining sample, 9% is Indian and 3% preferred not to classify their race.

[25] According to The Tobacco Atlas (see www.tobaccoatlas.org and Drope et al. [2018]) 26.5% of men and 5.5% of women smoke tobacco daily in South Africa. The prevalence rate for men is lower than in other medium Human Development Index (HDI) countries but the prevalence rate for women is higher than in other medium-HDI countries. Prevalence rates for selected high-income countries are: US – men: 14.4%, women: 11.7%; UK – men: 19.9%, women: 18.1%; Australia – men: 15.6%, women: 13.3%; Germany – men: 25.1%, women: 17.1%.

Smokers also completed the FTCD, which is a measure of smoking severity that scores people on a scale of 0 to 10, with higher numbers indicating greater severity. The average FTCD score among smokers is 2.495 with a standard deviation of 2.119.[26] In the experimental literature on atemporal risk preferences, time preferences, and smoking behaviour, reviewed in detail by HHRS, researchers often try to maximise the difference between smokers and non-smokers by selecting heavy smokers to take part in the study, e.g., at least 20 cigarettes smoked per day for the last 5 years and a FTCD score of at least 6 in Bickel, Odum and Madden [1999]. We recruited smokers across the entire spectrum of severity to determine whether being a smoker, irrespective of intensity, is associated with atemporal risk preferences, time preferences, and intertemporal risk preferences. This also allows us to explore the relationship between atemporal risk preferences, time preferences, intertemporal risk preferences, and smoking intensity.

Table 2

Summary Statistics

Variable	Mean	Standard Deviation
Demographics		
Age	29.948	11.887
White	0.266	0.443
Male	0.434	0.497
Financial situation today	2.840	1.041
Staff	0.434	0.497
Non-smoker	0.520	0.501
Ex-smoker	0.117	0.322
Current smoker	0.363	0.482
FTCD score	2.495	2.119
Average cigarettes per day	8.129	5.317
Treatments - Time Preferences		
FED: 0 days	0.502	0.500
FED: 7 days	0.498	0.500
High Principal	0.501	0.500

Table 2 shows that randomisation in the time preference task ensured that FED treatments were split evenly across the sample, and 50% of choices in the time preference task involved the high principal of R400.

[26] Fagerström and Furberg [2008] compare smokers' FTCD scores using nationally-representative studies in 13 countries and find that these scores range from 2.8 to 4.6. FTCD scores are highest in Sweden and the United States and lowest in Germany and Norway.

4. ECONOMETRICS

We adopt the statistical approach of AHLR to estimate the parameters of an intertemporal utility function jointly with the parameters defining atemporal risk preferences and time preferences.

Our intertemporal risk preference experiment used intertemporal lotteries that paid out amounts of money at two different points in time. Taking this into account, (1) can be simplified as follows

$$U(x_t, x_{t+\tau}) = E\big[\theta\big(D_t u(x_t) + D_{t+\tau} u(x_{t+\tau})\big)\big]. \tag{2}$$

To admit the possibility that the intertemporal utility function is not additively separable, we use a power function for $\theta(\cdot)$:

$$\theta\big(D_t u(x_t) + D_{t+\tau} u(x_{t+\tau})\big) = \big(D_t u(x_t) + D_{t+\tau} u(x_{t+\tau})\big)^{\rho} \tag{3}$$

where $\theta(z) = \ln z$ if $\rho = 0$, and $\theta(z) = -z^{\rho}$ if $\rho < 0$, following Wakker [2008]. With this power function specification, $\rho = 1$ yields the standard additively-separable model and intertemporal risk neutrality, $\rho < 1$ denotes intertemporal risk aversion, and $\rho > 1$ represents intertemporal risk seeking behaviour.

The intertemporal lotteries in our experiment only had two possible states of nature. Consider the "safe" intertemporal lottery A where the decision maker receives $(L_t, S_{t+\tau})$ with probability p and $(S_t, L_{t+\tau})$ with probability $1 - p$. Given the assumption that $\theta(z) = z^{\rho}$, the SDU of intertemporal lottery A is

$$\mathrm{SDU_A} = \omega(p) \times [\mathrm{D}_t\, u(L_t) + \mathrm{D}_{t+\tau}\, u(S_{t+\tau})]^{\rho} + [1 - \omega(p)] \times [\mathrm{D}_t\, u(S_t) + \mathrm{D}_{t+\tau}\, u(L_{t+\tau})]^{\rho}. \tag{4}$$

Apart from the specific functional form for $\theta(\cdot)$, equation (4) is completely general because we have not made any parametric assumptions about $u(\cdot)$, D_t, and $\omega(\cdot)$.

Now consider the "risky" intertemporal lottery B where the decision maker receives $(L_t, L_{t+\tau})$ with probability p and $(S_t, S_{t+\tau})$ with probability $1 - p$. The SDU of intertemporal lottery B is

$$\mathrm{SDU_B} = \omega(p) \times [\mathrm{D}_t\, u(L_t) + \mathrm{D}_{t+\tau}\, u(L_{t+\tau})]^{\rho} + [1 - \omega(p)] \times [\mathrm{D}_t\, u(S_t) + \mathrm{D}_{t+\tau}\, u(S_{t+\tau})]^{\rho}. \tag{5}$$

To write out the likelihood function for the choices the subjects made and estimate the parameters of the SDU model, we need to parameterise the functions $u(\cdot)$, D_t, and $\omega(\cdot)$. We consider the simplest case of EUT and exponential discounting first, and then discuss extensions to non-EUT and non-exponential specifications.

Under EUT, $\omega(p) = p$, and under exponential discounting $D^{E}_t = 1 / (1 + \delta)^t$. We let atemporal utility be defined by a power utility function that displays constant relative risk aversion

$$u(x) = x^r, \tag{6}$$

where $u(x) = \ln x$ if $r = 0$, and $u(x) = -x^r$ if $r < 0$.

With these assumptions, we can jointly estimate the atemporal risk preference parameter r, the time preference parameter δ, and the intertemporal risk preference parameter ρ by forming a latent $\triangledown$SDU index that captures the difference in the stochastic discounted utility of intertemporal lotteries A and B. We adopt the contextual utility behavioural error specification of Wilcox [2011] and define the latent index as

$$\triangledown\text{SDU} = [(\text{SDU}_\text{B} - \text{SDU}_\text{A}) / \lambda] / \psi, \tag{7}$$

where ψ is a behavioural error term for the intertemporal risk preference task and the term λ normalises the difference in SDU of intertemporal lotteries A and B to lie within the unit interval.

The likelihood of the intertemporal risk preference choices, conditional on the SDU specification being true, depends on the estimates of r, μ, δ, υ, ρ, and ψ, where μ is a behavioural error term for the atemporal risk preference task and υ is a behavioural error term for the time preference task, just as ψ is the behavioural error term for the intertemporal risk preference task. The conditional log-likelihood is

$$\ln L(r, \mu, \delta, \upsilon, \rho, \psi; c, \boldsymbol{X}) = \textstyle\sum_i[(\ln \Lambda(\triangledown\text{SDU} \times I(c_i = 1)) + (\ln \Lambda(\triangledown\text{SDU} \times I(c_i = 0))], \tag{8}$$

where $c_i = 1(0)$ denotes the choice of intertemporal lottery B(A) in intertemporal risk preference task i, Λ is the logistic cumulative distribution function, and $\boldsymbol{X}$ is a vector of individual characteristics capturing smoking status, gender, age, etc.

The joint likelihood of the atemporal risk preference, time preference, and intertemporal risk preference responses can then be written as

$$\ln L(r, \mu, \delta, \upsilon, \rho, \psi; c, \boldsymbol{X}) = \ln L^{\text{ARP}} + \ln L^{\text{TP}} + \ln L^{\text{SDU}}, \tag{9}$$

where $\ln L^{\text{ARP}}$ is the conditional log-likelihood of the atemporal risk preference choices, $\ln L^{\text{TP}}$ is the conditional log-likelihood of the time preference choices, and $\ln L^{\text{SDU}}$ is defined by (8).

It is straightforward to extend (9) to incorporate non-EUT models of choice under atemporal risk and non-exponential discounting specifications. For example, in the case of QH discounting, we replace $D^{\text{E}}{}_t = 1 / (1 + \delta)^t$ with $D^{\text{QH}}{}_t = \beta / (1 + \delta)^t$. We then form the latent $\bigtriangledown$SDU index in (7) and proceed as before with one additional parameter (β) to estimate in (9).

5. RESULTS

We estimate the SDU model (9) jointly with the parameters defining atemporal risk preferences and time preferences. Based on analyses of the atemporal risk preference data and time preference data, we extend the econometric specification in (9) to incorporate a RDU model of choice under atemporal risk and a QH model of time preferences. This specification is then used to analyse the relationship between atemporal risk preferences, time preferences, intertemporal risk preferences, and three measures of smoking behaviour: smoking status; smoking intensity, measured by the number of cigarettes smoked per day; and smoking severity, measured by smokers' scores on the FTCD.

A. Baseline Estimates

We estimate the homogenous preference SDU model (9) under the assumptions that EUT characterises choice under atemporal risk and that discounting is exponential; Table C1 in the appendix presents the results. The estimate of the atemporal risk preference parameter $r = 0.409$, which is significantly less than 1 ($p < 0.001$), implies that the sample is moderately risk averse, whereas the estimate of the exponential discounting parameter $\delta = 0.782$ indicates that future rewards are discounted at the relatively high rate of 78.2% per annum. The estimate of the intertemporal risk preference parameter ρ is -1.043, which is significantly less than 1 ($p < 0.001$), and shows that the sample is characterised by a high level of intertemporal risk aversion. Recall that when $\rho = 1$, the SDU model is additively separable, which implies intertemporal risk neutrality. Thus, our results show that the most common model of intertemporal choice in economics, viz., the additively-separable model, is not an accurate description of the intertemporal risk preferences of our sample. This echoes the result in AHLR but with a much higher level of intertemporal risk aversion in our sample.[27]

HHRS emphasise the importance of appropriately characterising a sample's atemporal risk attitudes when drawing inferences about its discounting behaviour because, as Andersen, Harrison, Lau and Rutström [2008] showed, estimates of utility function curvature significantly affect estimates of discounting parameters. Under EUT, atemporal risk preferences are determined solely by the curvature of the utility function over outcomes, whereas under RDU atemporal risk preferences are determined jointly by the curvature of the utility function and the probability weighting function (PWF). This implies that if there is evidence of probability weighting in a sample then this needs to be taken into account when estimating time preference models or else this probability-weighting source of atemporal risk preferences will show up in the curvature of the utility function under EUT and bias discounting parameter estimates.

This logic extends naturally to the SDU model because it is estimated jointly with the parameters defining atemporal risk preferences and time preferences. Consequently, it is important to accurately identify atemporal risk preferences *and* time preferences when estimating a SDU model because these atemporal risk preference and time preference estimates propagate into inferences drawn from the SDU model. We therefore investigate whether there is evidence of non-EUT and non-exponential discounting in our data so that the SDU model can be extended to incorporate these features.

[27] AHLR use the following functional form for $u(\cdot)$ and $\theta(\cdot)$ in their structural econometric model: $u(x) = x^{1-r} / (1-r)$ and $\theta(z) = z^{1-\rho} / (1-\rho)$. Hence their estimates are not directly comparable to ours because we use power functions for $u(\cdot)$ and $\theta(\cdot)$ in (9). Estimating model (9) on the data of AHLR, the estimate of r is 0.449, the estimate of δ is 0.077, and the estimate of ρ is 0.563, which shows that we find much higher levels of intertemporal risk aversion in our sample, where the estimate of ρ is -1.043. Alternatively, if we use the AHLR functional forms in (9), the estimates on our sample are: $r = 0.591$, $\delta = 0.782$, and $\rho = 2.043$. The estimates of AHLR [Table 1, p. 544], using this same specification for $u(\cdot)$ and $\theta(\cdot)$, are: $r = 0.35$, $\delta = 0.114$, and $\rho = 0.32$. This shows again that we find much higher levels of intertemporal risk aversion in our sample.

We begin by estimating EUT and RDU models of choice under atemporal risk; see Table C2 in the appendix for the results. A crucial ingredient of a RDU model is the specification of the PWF. Owing to its flexibility, we use the Prelec [1998] function

$$\pi(p) = \exp[-\eta(-\ln p)^{\varphi}], \tag{10}$$

which is defined for $1 > p > 0$, $\eta > 0$ and $\varphi > 0$. This function nests a power PWF when $\eta = 1$, and it nests a one-parameter function when $\varphi = 1$ that admits linear, inverse S-shaped, and S-shaped forms.

Figure 4 graphs the estimates of the Prelec PWF from the RDU model in Table C2 along with the implied decision weights for 2, 3, and 4 outcome equi-probable reference lotteries.[28] We cannot reject the hypothesis that $\eta = 1$, but the estimate of $\varphi = 0.629$ is significantly less than 1 ($p < 0.001$), which yields an inverse S-shaped PWF with overweighting of low probabilities and underweighting of moderate to high probabilities.[29] For 3-outcome and 4-outcome reference lotteries, this form of probability weighting implies decision weights for the highest and lowest ranked lottery prizes that exceed the corresponding probabilities, and decision weights for intermediate prizes that are less than the corresponding probabilities. This subjective distortion of objective probabilities leads to a statistically significant increase ($p < 0.001$) in the estimate of $r = 0.553$ under the RDU model compared to the estimate of $r = 0.408$ under the EUT model. In turn, this increase in the power function parameter r under RDU leads to a statistically significant increase ($p < 0.001$) in the estimate of the exponential discount rate $\delta = 1.192$ compared to the estimate of $\delta = 0.785$ under EUT; Table C3 in the appendix presents the results from exponential discounting models under the assumption that either EUT or RDU characterises choice under atemporal risk. Thus the statistically significant evidence of probability weighting has an economically significant impact on estimates of the exponential discount rate. This again demonstrates the

[28] An equi-probable reference lottery is one where the probabilities assigned to prizes are equal. Thus, in the case of a 2-outcome equi-probable reference lottery, each prize has a probability of 0.5. For a 3-outcome equi-probable reference lottery, each prize has a probability of 1/3; and for a 4-outcome equi-probable reference lottery, each prize has a probability of 0.25. The dashed lines in the right panel of Figure 4 represent these reference probabilities: 0.25, 0.33, and 0.5.

[29] Table C2 shows that the sample as a whole is better characterised by RDU than EUT, but this does not imply that every person in our sample probability weights and, therefore, departs from EUT. Figure C1 in the appendix shows the results from an individual-level analysis where we estimate EUT and RDU specifications for each subject and then test whether $\omega(p) = p$ in the RDU model. Using a 5% level of statistical significance, we cannot reject the hypothesis that $\omega(p) = p$ for 57% of the sample, implying that at least half of the people in our study are better characterised by EUT than RDU. However, the remaining 43% of the sample exhibits statistically significant evidence of nonlinear probability weighting. As we show, it is necessary to take this probability weighting into account when drawing inferences about time preferences and intertemporal risk preferences at the level of the *sample of subjects*.

importance of correctly characterising atemporal risk preferences when estimating time preferences.

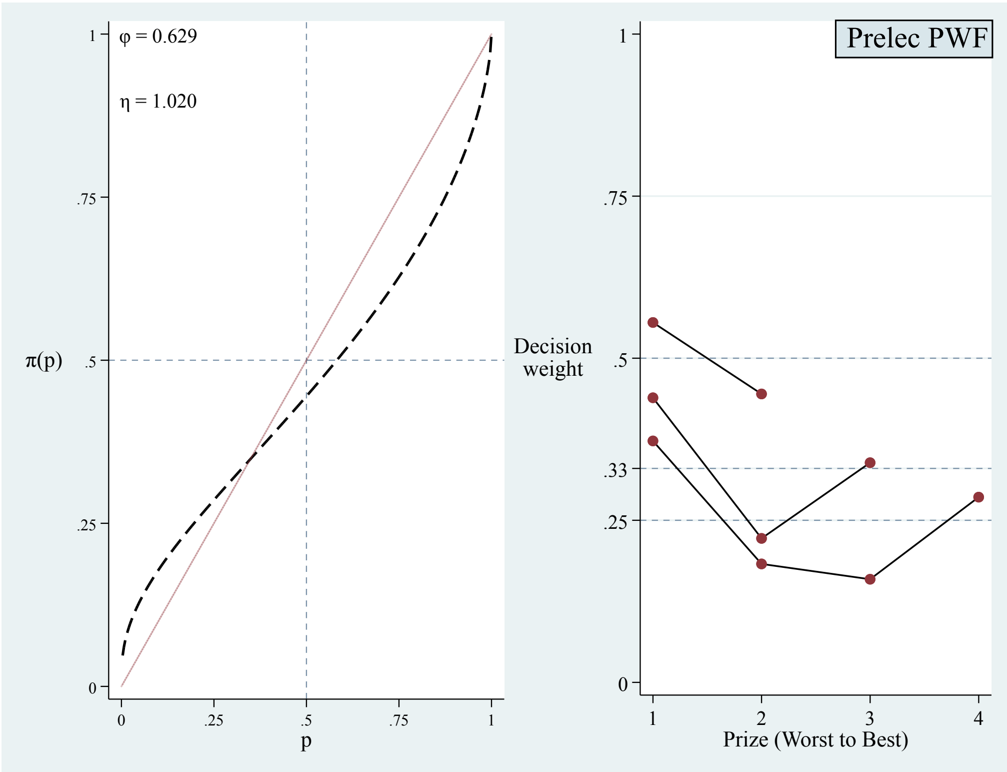

Figure 4: Estimated Probability Weighting Function and Implied Decision Weights

Similarly, we find statistically significant evidence of non-exponential discounting when estimating a QH discounting function jointly with a RDU model; Table C4 in the appendix presents the results. Our estimate of β is 0.960, and is significantly less than 1 ($p < 0.001$), which generates declining discount rates over time and a lower estimate of the long-term discount rate δ = 0.885 compared to the estimate of δ = 1.192 under the assumption of exponential discounting.

In sum, analyses of the atemporal risk preference data and time preference data suggest that we should estimate our SDU model jointly with a RDU model to incorporate nonlinear probability weighting in choice under atemporal risk, and a QH discounting model to account for a present-bias in intertemporal decision making. Table 3 presents the results from this model.

Table 3
Intertemporal Risk Preference ML Estimates
RDU, Quasi-Hyperbolic Discounting
Homogenous Preferences

	Model
	Estimate
Atemporal Risk Preferences	
Power function parameter (*r*)	0.522***
	(0.019)
PWF parameter (φ)	0.716***
	(0.019)
PWF parameter (η)	1.027***
	(0.028)
Error (μ)	0.140***
	(0.005)
Time Preferences	
Discounting parameter (β)	0.961***
	(0.003)
Discounting parameter (δ)	0.840***
	(0.066)
Error (υ)	0.874***
	(0.140)
Intertemporal Risk Preferences	
Power function parameter (ρ)	-0.644***
	(0.218)
Error (ψ)	0.286***
	(0.022)
N	48640
log-likelihood	-28351.110

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

The atemporal risk preference estimates indicate a moderate level of utility function curvature and statistically significant evidence of inverse S-shaped probability weighting. The time preference results show that there is a discontinuous $\beta = 0.961$ drop in the value of a reward if it is not available immediately but this drop asymptotes toward the long-term discount rate $\delta = 0.840$ over time, which is significantly lower ($p < 0.001$) than the estimate of $\delta = 1.122$ under the assumption of exponential discounting. With regard to intertemporal risk preferences, there is a marked and statistically significant ($p < 0.001$) increase in the estimate of $\rho = -0.644$ in Table 3 relative to the estimate of $\rho = -1.043$ under the assumptions of EUT and exponential discounting; see Table C1 in the appendix. These results show that in a joint estimation framework, atemporal risk preference, time preference, and intertemporal risk preference estimates are inextricably linked. Hence, correct specification of the constituent parts of a SDU model is necessary for valid statistical inference. The effects of allowing for intertemporal risk aversion on risk premia can be subtle. We can show that allowing for intertemporal risk aversion, compared to imposing intertemporal risk neutrality, can make intertemporally “safe” (“risky”) lotteries more (less) attractive, and typically have more dramatic effects on the risk premium for intertemporally “risky” lotteries.[30] We analyse the relationship between atemporal risk preferences, time preferences, intertemporal risk preferences, and smoking behaviour using the statistical specification in Table 3. We also estimate alternative specifications to test the robustness of our results.

B. Smoking Status

HHRS, using a sample of 175 UCT students in 2012, find that atemporal risk preferences do not differ as a function of smoking status but do find that smokers discount the future significantly more heavily than non-smokers. We evaluate these findings with a larger sample of UCT students and staff that has more variation in demographic and socio-economic characteristics. In addition, we specifically recruited ex-smokers so as to draw comparisons between smokers, ex-smokers, and non-smokers. Finally, our experiment elicited intertemporal risk preferences so we can analyse the relationship between smoking status and the curvature of the intertemporal power function parameter ρ in our SDU model.

Table D1 in the appendix presents results from the SDU model under the assumptions that a RDU model with a power utility function and the Prelec PWF characterise choice under atemporal risk, and that discounting is QH. We allow the parameters of the model to vary as a linear function of smoking status, demographics, and socio-economic

[30] To understand the role that intertemporal risk aversion plays in the characterisation of attitudes to risk over time, we calculate certainty equivalents and then evaluate risk premia for intertemporal lotteries A and B, under the standard assumption of intertemporal risk neutrality and using the intertemporal risk preference estimates in Table 3. Figure C2 in the appendix shows that risk premia for intertemporal lottery A are *smaller* when using the estimates in Table 3 than when we re-estimate the model and impose intertemporal risk neutrality. By contrast, Figure C3 in the appendix shows that risk premia for intertemporal lottery B are *larger* when calculated using the estimates in Table 3 in comparison to when we re-estimate the model and impose intertemporal risk neutrality. Thus, intertemporal risk aversion, relative to intertemporal risk neutrality, makes the “safe” intertemporal lottery A more attractive to a decision maker, while making the “risky” intertemporal lottery B less attractive to a decision maker. Figure C4 combines Figure C2 and Figure C3, with a common *y*-axis in all of the panels, to emphasise the economic significance of incorporating intertemporal risk aversion as opposed to assuming intertemporal risk neutrality. Figure C4 shows that intertemporal risk aversion generates a larger difference in risk premia, relative to intertemporal risk neutrality, for the “risky” intertemporal lottery B compared to the “safe” intertemporal lottery A across most of the probability space. This difference in risk premia shows the economic significance of allowing intertemporal risk aversion as opposed to assuming intertemporal risk neutrality in the characterisation of attitudes to risk over time, particularly in relation to “risky” intertemporal lotteries.

characteristics. There are no statistically significant differences in the atemporal risk preferences of smokers, ex-smokers, and non-smokers, which accords with the findings of HHRS. This result is robust to the assumptions that EUT characterises choice under atemporal risk and that discounting is exponential.[31]

With regard to time preferences, the estimate of δ_{Smokers} is 0.356, and implies that smokers discount at a significantly higher rate than non-smokers ($p < 0.001$). This difference in discounting behaviour is economically significant: the long-term discount rate of smokers is 36 percentage points higher than non-smokers. The comparable results in Table E1 show that under the assumptions of EUT and exponential discounting, smokers discount at a 32 percentage point higher rate than non-smokers. By contrast, there are no statistically significant differences in the long-term δ discounting behaviour of ex-smokers and non-smokers ($p = 0.440$) and of smokers and ex-smokers ($p = 0.161$). Moreover, there are no statistically significant differences between smokers, ex-smokers, and non-smokers in terms of present-bias β.

Figure 5 shows a kernel-weighted local polynomial regression, with a 95% confidence interval, of the fraction of LL choices by smokers, ex-smokers, and non-smokers at the nominal annual interest rates in the time preference task. At each interest rate, the estimate of smokers' LL choice fraction is far below that of non-smokers, and the 95% confidence intervals do not overlap, implying that smokers discount the future at a significantly higher rate than non-smokers. By contrast, the estimates and 95% confidence intervals for ex-smokers overlap with those of smokers and non-smokers, suggesting that ex-smokers discount at a level between smokers and non-smokers. Figure 5 provides visual confirmation of the results in Table D1.

[31] Appendix E presents the results of a comparable set of SDU models assuming EUT and exponential discounting as opposed to RDU and QH discounting. The results in appendices D and E are similar and differences are noted where necessary. Table E1 shows that the atemporal risk preference parameter estimate for smokers is 0.033, and is significantly higher ($p = 0.067$) than the estimate for non-smokers. This result is not economically significant and is a product of the covariance of estimates in a joint estimation framework. Analyses of the atemporal risk preference data alone show that there are no statistically significant differences in the atemporal risk preferences of smokers, ex-smokers, and non-smokers under EUT and RDU specifications.

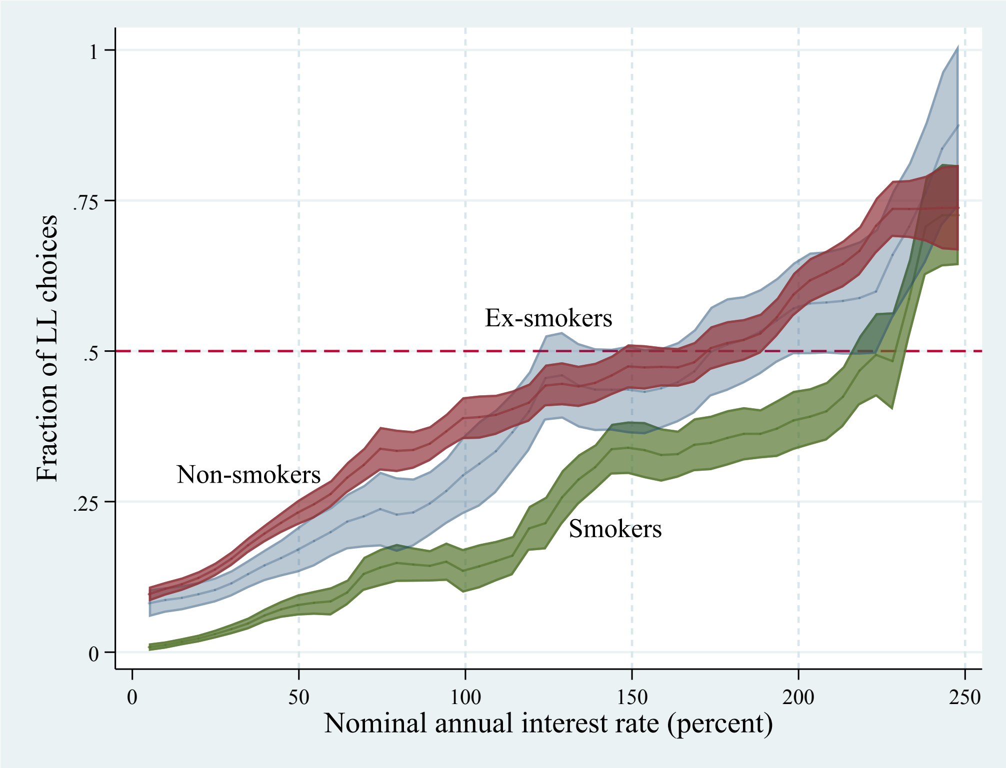

Figure 5: Fraction of LL Choices by Smoking Status

Figure 6 shows a kernel density plot of the intertemporal risk preference parameter ρ, based on predictions of ρ for each subject using the covariate estimates in Table D1. The distribution is skewed towards high levels of intertemporal risk aversion and exhibits significant heterogeneity according to demographic and socio-economic characteristics. The coefficient estimate of ρ for men is 0.841 ($p < 0.05$) and implies that they are significantly less intertemporally risk averse than women. There is also a strong association between a subject's financial situation on the day of the experiment and estimates of intertemporal risk aversion. Specifically, every one category improvement on the financial situation scale is associated with a 0.611 ($p < 0.05$) *increase* in intertemporal risk aversion, implying that subjects in better financial situations are more intertemporally risk averse than subjects in worse financial situations. Of course, this is correlation: we are agnostic about causation. There are no statistically significant differences in intertemporal risk preferences between smokers and ex-smokers or between smokers and non-smokers, but ex-smokers are significantly more intertemporally risk averse than non-smokers at the 10% level. This latter result is not robust to the assumption that EUT and exponential discounting characterise atemporal risk preferences and time preferences, respectively; see Table E1 in the appendix.

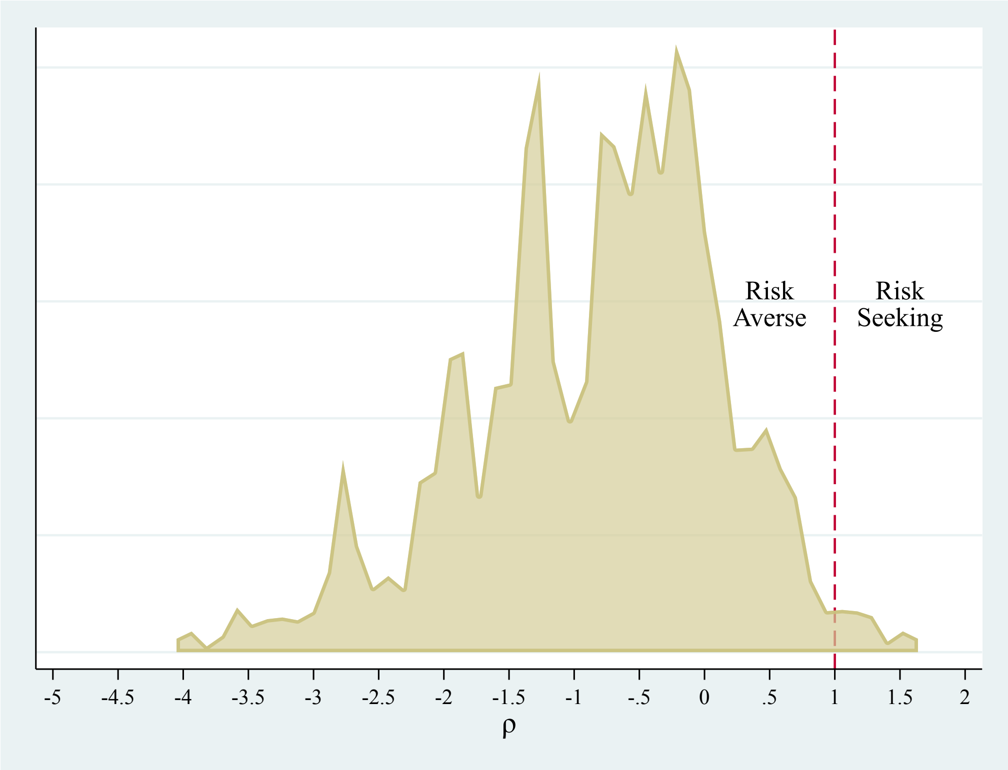

Figure 6: Distribution of the Intertemporal Risk Preference Parameter (ρ)

In sum, the sample is characterised by a high degree of heterogeneity in intertemporal risk preferences that varies as a function of gender and financial situation but not by smoking status. The analyses in this section suggest that the only robust behavioural difference between smokers, ex-smokers, and non-smokers appears to be in their long-term discounting behaviour, with smokers discounting the most, non-smokers discounting the least, and ex-smokers discounting at a level between the two other groups.

C. Smoking Intensity and Smoking Severity

Given the historical differences in smoking prevalence for men and women[32], together with the statistically and economically significant difference in their intertemporal risk preferences, we split the sample by gender to analyse the relationship between atemporal

[32] Thun et al. [2013] review historical differences in male and female smoking prevalence and smoking behaviour in the United States since the early 20th century. They also examine male and female death rates and relative risks attributed to cigarette smoking during three time periods: 1959-1965, 1982-1988, and 2000-2010. They find marked disparities in relative risks between male and female smokers in the earlier cohorts, but convergence in relative risks in the most recent cohort, leading them to quote former US Secretary of Health, Education, and Welfare, Joseph A. Califano, Jr. [1979, i], who wrote, “Women who smoke like men die like men who smoke.”

risk preferences, time preferences, intertemporal risk preferences, and measures of smoking intensity and smoking severity.

Table 4 presents results from the SDU model estimated jointly with a RDU model, power utility function, and Prelec PWF for choice under atemporal risk, and a QH discounting function. Following HHRS, we investigate whether there is a relationship between smoking intensity, measured by the number of cigarettes smoked per day, and atemporal risk preferences, time preferences, and intertemporal risk preferences. Unlike HHRS, who find a concave relationship between smoking intensity and discounting behaviour, estimates of the quadratic term across all parameters in our model are not statistically significant so we only include a linear term for smoking intensity.

Table 4 shows that for both men and women there is a large and statistically significant relationship between the number of cigarettes smoked per day and the long-term discounting parameter δ. Specifically, every additional cigarette smoked per day is associated with a 5 percentage point increase in the long-term discounting of men, whereas every additional cigarette smoked per day is associated with a 3 percentage point increase in the long-term discounting of women.[33] These economically significant estimates explain why heavier smokers find it harder to quit: the long-term benefits that result from successful abstention are discounted heavily and do not exceed the short-term costs of quitting. By contrast, there is no statistically significant relationship between smoking intensity and the present-bias parameter β.

Table 4 also suggests that there is a relationship between smoking intensity and the atemporal risk attitudes of men and women. For women, the number of cigarettes smoked per day is statistically significant in the PWF parameter φ. For men, the number of cigarettes smoked per day is statistically significant in the atemporal risk preference parameter r and the PWF parameter φ. However, the statistically significant estimate for r is 0.005 ($p < 0.05$), and is not economically significant: a 10 cigarette increase in the smoking intensity of men is only associated with a 0.05 increase in the atemporal risk preference parameter r, implying only a modest decrease in atemporal risk aversion. Furthermore, this statistically significant result is a product of our joint estimation statistical framework, because Table D2 in the

[33] Table E2 in the appendix shows that these results are robust to the assumptions of EUT and exponential discounting.

appendix shows that when analysing the atemporal risk preference data alone, there is no relationship between smoking intensity and curvature of the atemporal utility function of men. Similarly, the statistically significant estimate of the number of cigarettes smoked per day by women in the PWF parameter φ is also not present in the atemporal risk preference data alone.

By contrast, the statistically significant estimate of the number of cigarettes smoked per day by men in the PWF parameter φ is present in the atemporal risk preference data alone; see Table D2. But this estimate of -0.015 is not *economically significant*: a 10 cigarette increase in the smoking intensity of men leads to a 0.15 decrease in the PWF parameter φ, implying only a small change in probability weighting. Thus, while there is a robust *statistical* relationship between the number of cigarettes smoked per day and the probability weighting of men this does not lead to substantive *economic* changes in atemporal risk attitudes.

Table 4
Intertemporal Risk Preference ML Estimates
RDU, Quasi-Hyperbolic Discounting
Smoking Intensity: Number of Cigarettes Smoked per Day

	Model 1		**Model 2**	
	Male		Female	
	Estimate	Std error	Estimate	Std error
Atemporal risk preference parameter (*r*)				
Age	-0.001	0.002	-0.002	0.001
White	-0.025	0.031	0.022	0.027
Financial situation	-0.002	0.012	0.024	0.015
Staff member	0.027	0.043	0.011	0.034
Number of cigarettes	0.005**	0.002	0.003	0.003
Constant	0.589***	0.065	0.442***	0.053
PWF parameter (φ)				
Age	0.004	0.005	-0.001	0.003
White	0.175**	0.076	0.039	0.066
Financial situation	-0.024	0.032	0.022	0.028
Staff member	0.030	0.111	-0.013	0.090
Number of cigarettes	-0.015***	0.005	0.012***	0.004
Constant	0.733***	0.130	0.617***	0.107
PWF parameter (η)				
Age	0.016*	0.009	-0.001	0.005
White	-0.029	0.086	0.024	0.098
Financial situation	-0.008	0.047	0.108***	0.041
Staff member	-0.184	0.152	-0.130	0.111
Number of cigarettes	0.004	0.008	0.009	0.008
Constant	0.648***	0.245	0.811***	0.151

Table 4 (Continued)

	Model 1		Model 2	
	Male		Female	
	Estimate	Std error	Estimate	Std error
Discounting parameter (β)				
Age	<0.001	0.001	<0.001	<0.001
White	0.009	0.011	0.016**	0.007
Financial situation	0.001	0.005	0.003	0.004
Staff member	0.003	0.009	<0.001	0.008
Number of cigarettes	-0.001	0.001	<0.001	<0.001
Constant	0.966***	0.018	0.964***	0.016
Discounting parameter (δ)				
Age	0.005	0.009	-0.002	0.005
White	-0.399**	0.173	-0.182	0.121
Financial situation	-0.274***	0.091	-0.185*	0.097
Staff member	-0.226	0.258	-0.125	0.166
Number of cigarettes	0.052**	0.023	0.034**	0.014
Constant	1.755***	0.353	1.404***	0.368
Intertemporal risk preference parameter (ρ)				
Age	-0.050	0.034	0.025	0.048
White	-0.096	0.351	-3.995	8.671
Financial situation	-0.153	0.194	-1.038*	0.534
Staff member	0.972	0.664	1.353	1.290
Number of cigarettes	-0.047**	0.024	0.059	0.075
Constant	1.664*	0.945	-0.128	1.238
Error terms				
μ	0.127***	0.007	0.149***	0.008
υ	0.989***	0.276	0.559***	0.140
ψ	0.208***	0.024	0.314***	0.028
N	20900		26410	
log-likelihood	-11635.506		-15082.686	

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

Finally, Table 4 shows that smoking intensity is related to the intertemporal risk attitudes of men ($p < 0.05$) but not women ($p = 0.434$). For men, every additional cigarette smoked per day is associated with a 0.05 *increase* in intertemporal risk aversion, implying that heavier male smokers in our sample tend to be more intertemporally risk averse. This estimate is economically significant because an increase of 10 cigarettes smoked per day is associated with a 0.5 increase in intertemporal risk aversion. The point estimate for women of the number of cigarettes smoked per day is of the opposite sign and has a large standard error.[34] This shows the importance of splitting the sample by gender when analysing

[34] Table E2 in the appendix shows that these intertemporal risk preference results are robust to the assumptions of EUT and exponential discounting.

intertemporal risk attitudes, because the statistically significant estimate for men is washed out by the large standard error of the estimate for women when the sample is pooled; see Table D3 in the appendix for the pooled estimates.

Table D4 in the appendix presents results from the SDU model where the parameters are allowed to vary as a linear function of demographics, socio-economic characteristics, and smoking severity, measured by smokers' scores on the FTCD. For men and women, there are no statistically significant relationships between present-bias β, long-term discounting δ, and smoking severity. Similarly, there are no substantive relationships between smoking severity and the atemporal risk preferences of men and women.[35] However, there is a large and statistically significant ($p < 0.05$) relationship between smoking severity and the intertemporal risk attitudes of men. A 1-unit increase in FTCD score is associated with a 0.44 *increase* in intertemporal risk aversion, suggesting that as smoking severity increases, male smokers become much more intertemporally risk averse. Echoing the results for smoking intensity, there is no statistically significant relationship between smoking severity and the intertemporal risk preferences of women.[36]

6. DISCUSSION AND CONCLUSIONS

We investigate the relationship between atemporal risk preferences, time preferences, intertemporal risk preferences, and three measures of smoking behaviour using an incentive-compatible experimental design and structural econometric framework. We find statistically and economically significant evidence of non-linear probability weighting in choice under atemporal risk but no substantive differences in atemporal risk preferences by smoking status, smoking intensity, and smoking severity.

By contrast, time preferences are related both to smoking status and smoking intensity. Smokers discount significantly more heavily than non-smokers, and ex-smokers discount at a level between these two groups. In addition, there is a large and statistically

[35] Table D5 in the appendix shows the results for men and women from the RDU model estimated on the atemporal risk preference data alone. There are no statistically significant relationships between smoking severity and the atemporal risk preferences of women. For men, the FTCD score is statistically significant in the power function parameter r and the PWF parameter φ but only at the 10% level in both cases, and the coefficient estimates are not economically significant.

[36] Table E3 shows that these results are robust to the assumptions of EUT and exponential discounting.

significant relationship between smoking intensity and the discounting behaviour of men and women: every additional cigarette is associated with a 3-5 percentage point increase in the long-term discounting parameter δ. However, smoking severity is not related to long-term discounting, and there are no statistically significant differences in present-bias according to smoking status, smoking intensity, and smoking severity.

The intertemporal risk preferences of our sample are characterised by a large degree of heterogeneity and high levels of intertemporal risk aversion that varies according to gender and financial situation but not by smoking status. Analyses conducted on subsamples of men and women reveal that smoking intensity and smoking severity are associated with the intertemporal risk attitudes of men but not the intertemporal risk attitudes of woman. These results are robust to different models of choice under atemporal risk and alternative discounting specifications.

Our research makes a number of contributions to the experimental economic literature on addiction. First, we replicate the finding of HHRS, using a larger sample with more demographic and socio-economic variation, that atemporal risk aversion does not differ substantively according to smoking status and smoking intensity, while extending this null result to smoking severity as measured by the FTCD. These results suggest that despite the clear risks involved in tobacco smoking, atemporal risk preferences are not a robust behavioural marker of addiction.

Second, we replicate the finding of economically and statistically significant differences in the time preferences of smokers and non-smokers, and add nuance to this result by including ex-smokers in the sample: ex-smokers discount at a level between smokers and non-smokers. This suggests a causal relationship between discounting behaviour and smoking status, about which we can only speculate. One hypothesis is that time preference is a persistent trait that remains constant through the onset and course of addiction, and that the mid-range discounting of ex-smokers makes them more likely than non-smokers to smoke in the first place, but more likely than current smokers to try to quit and to succeed conditional on trying. A problem with this hypothesis is that it ignores the fact that most smokers in the sample will eventually *become* ex-smokers, notwithstanding their higher discounting. An alternative hypothesis is that both smoking and quitting affect discounting in opposite directions. Expected observations, on either hypothesis, are sensitive to the age distribution in a sample, since younger ex-smokers will be expected to have quit after shorter smoking careers and fewer failed attempts. This also leads to ambiguous causal hypotheses. Intuitively, one might expect that a sample of younger ex-smokers should resemble a sample of non-smokers more closely than a sample of older ex-smokers resembles non-smokers. However, suppose that smokers' experience with trying to quit teaches them to manage their discounting, as suggested by Ainslie [2001]. In that case, veterans of longer and more self-conscious quitting campaigns might undergo more learning of reduced discounting than smokers who quit with less effort sooner after starting. Evidently, a panel study is needed to disambiguate these possible relationships.

Third, we identify a large, positive relationship between smoking intensity and the discounting behaviour of men and women that has important implications for treatment of tobacco use disorder. Heavier smokers tend to have higher discount rates, which will make it harder for them to quit because the long-term costs of continuing to smoke and the long-term benefits that result from successful abstention are discounted heavily. These differences in smoking intensity and discounting behaviour could be leveraged in the design of smoking cessation programmes. For example, the reinforcement schedules of a contingency management smoking cessation intervention, which provides monetary incentives for biochemically-verified abstinence, could be tailored to the smoking intensity and discounting behaviour of smokers. Heavier smokers could be assigned to a front-loaded reinforcement schedule where they are given a large first payment for successful abstention to get them over the initial hump. Lighter smokers, on the other hand, could be given a uniform-incentive reinforcement schedule where the rewards for abstinence are held constant across visits. Acknowledging differences between smokers and adjusting cessation interventions accordingly may make treatment of tobacco use disorder more efficacious in general, and particularly fruitful in the case of hard-to-treat smokers.

Fourth, this is the first study to have investigated the intertemporal risk preferences of smokers, specifically, and addicts, generally. Building on the work of AHLR, we provide a template for conducting this investigation that uses incentive-compatible economic experiments and a structural econometric framework to estimate a SDU model jointly with atemporal risk preference specifications and discounting functions.

Fifth, we show the importance of accurate identification of atemporal risk preferences and discounting behaviour when drawing inferences about intertemporal risk attitudes. As errors and uncertainty at all levels of a joint estimation framework propagate, as they should theoretically, it is theoretically appropriate and empirically necessary to apportion atemporal risk preferences into their utility function curvature and probability weighting components, and incorporate non-constant discounting behaviour, if it is present, when estimating SDU models.

Finally, we identify significant heterogeneity in intertemporal risk preferences but find, contrary to the assumption employed by standard economic models of addiction, that smokers do not exhibit intertemporal risk seeking behaviour. Our sample is characterised by a high level of intertemporal risk aversion, which does not differ significantly according to smoking status. However, measures of smoking intensity and smoking severity are related to the intertemporal risk attitudes of men: increases in smoking intensity and smoking severity are associated with statistically and economically significant increases in intertemporal risk aversion. By contrast, the intertemporal risk preferences of women do not differ as a function of smoking intensity or smoking severity.

As discussed earlier, initial economic models of addiction, in taking it to be a form of habitual consumption, implicitly conjectured that such intertemporal risk seeking preferences would be fundamental to choice-based accounts. As demonstrated by Bommier and Rochet [2006] and Lichtendal, Chao, and Bodily [2012], however, there is no strict implication of intertemporal risk seeking preferences from modelling addiction as habitual consumption. However, a theorist might venture the following hypothesis. Stereotyped behavioural sequences are cues for addictive cravings (see West and Brown [2013]). Such sequences are habits, though not habits of consumption. Perhaps, then, people with stronger dispositions to adopt habits in general should be more vulnerable to pathological development of the form of neural associative learning that underlies addiction. This reasoning would generate the opposite prediction from the BM model of addiction: we would expect to find statistically *higher* intertemporal risk *aversion* in addicts. Our findings relating smoking severity and the consumption of cigarettes to the intertemporal risk preferences of men, support this idea. But the lack of these relationships amongst women, at least in our sample, point to the need for more research into the aetiology of smoking in women (and men).

Our experimental methodology points the way to further research. We noted the potential value of a panel design to explore possible causal relationships between discounting, the *ex ante* probability of taking up smoking, and the *ex post* probability of a smoker having quit after a given length of smoking career and number of quit attempts. We earlier cited the observation by Ainslie [2001] of two distinct common strategies by which smokers try to quit, tapering and cold turkey, and pointed out that the former pattern is suggestive of higher intertemporal risk aversion than the latter pattern. If there are two types of addicts who could be reliably distinguished by experimental measurements of intertemporal risk aversion, as our results for men suggest, this could be useful for clinicians choosing from menus of therapeutic interventions for patients with varying characteristics. In general, experimental operationalisation of structural models of heterogeneous behavioural response within populations is an increasingly emphasised aim and achievement of laboratory economics.

Taken together, these results have two implications for the behavioural analysis of smoking. First, heterogeneity of smoking behaviour, rather than the binary classifications of "ever" or "never" smokers, clearly interacts with risk and time preferences. This is particularly evident in the smoking intensity results where increases in the number of cigarettes smoked per day are associated with large increases in long-term discounting behaviour on the one hand, and increases in the intertemporal risk aversion of men, but not women, on the other. Second, evidence for atemporal *and* intertemporal risk aversion, coupled with moderate levels of discounting, point to the potential role that poorly calibrated subjective beliefs about smoking might play in the onset and persistence of addiction.

REFERENCES

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

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

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 (2018): "Multiattribute Utility Theory, Intertemporal Utility, and Correlation Aversion," *International Economic Review*, 59, 537-555.

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 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.

BICKEL, W. K., A. L. ODUM, AND G. J. MADDEN (1999): "Impulsivity and Cigarette Smoking: Delay Discounting in Current, Never, and Ex-Smokers," *Psychopharmacology*, 146, 447-54, 91460447.213 [pii].

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

BOMMIER, A., A. KOCHOV, AND F. LE GRAND (2017): "On Monotone Recursive Preferences," *Econometrica*, 85, 1433-1466.

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

CALIFANO JR, J. A. (1979): "The Secretary's Foreward," in *Smoking and Health: A Report of the Surgeon General*. Washington, DC: Department of Health, Education, and Welfare (DHEW publication no. (PHS) 79-50066).

CHAITON, M., L. DIEMERT, J. E. COHEN, S. J. BONDY, P. SELBY, A. PHILIPNERI, AND R. SCHWARTZ (2016): "Estimating the Number of Quit Attempts It Takes to Quit Smoking Successfully in a Longitudinal Cohort of Smokers," *BMJ Open*, 6, e011045.

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.

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

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.

DROPE, J., N. SCHLUGER, Z. CAHN, J. DROPE, S. HAMILL, F. ISLAMI, A. LIBER, N. NARGIS, AND M. STOKLOSA (2018): *The Tobacco Atlas*. Atlanta, GA: American Cancer Society and Vital Strategies.

EPSTEIN, L. G., AND S. M. TANNY (1980): "Increasing Generalized Correlation: A Definition and Some Economic Consequences," *Canadian Journal of Economics*, 13, 16-34.
EPSTEIN, L. G., AND S. E. ZIN (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937-969.
FAGERSTRÖM, K. (2012): "Determinants of Tobacco Use and Renaming the FTND to the Fagerström Test for Cigarette Dependence: A Revision of the Fagerström Tolerance Questionnaire," *Nicotine & Tobacco Research*, 14, 75-78.
FAGERSTRÖM, K., AND H. FURBERG (2008): "A Comparison of the Fagerström Test for Nicotine Dependence and Smoking Prevalence across Countries," *Addiction*, 103, 841-845, 10.1111/j.1360-0443.2008.02190.x.
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.
GRANT, B. F., AND D. A. DAWSON (2006): "Introduction to the National Epidemiologic Survey on Alcohol and Related Conditions," *Alcohol Research and Health*, 29, 74-78.
GRAYOT, J. (2019): *Agency and Choice*. Doctoral dissertation, Erasmus University, Rotterdam.
GUL, F., AND W. PESENDORFER (2007): "Harmful Addiction," *The Review of Economic Studies*, 74, 147-172.
HANSSON, B. (1988): "Risk Aversion as a Problem of Conjoint Measurement," in *Decision, Probability, and Utility*, ed. by P. Gardenfors, and N.-E. Sahlin. New York, NY: Cambridge University Press.
HARRISON, G. W., A. HOFMEYR, D. ROSS, AND J. T. SWARTHOUT (2018): "Risk Preferences, Time Preferences, and Smoking Behavior," *Southern Economic Journal*, 85, 313-348.
HARRISON, G. W., J. MARTÍNEZ-CORREA, AND J. T. SWARTHOUT (2015): "Reduction of Compound Lotteries with Objective Probabilities: Theory and Evidence," *Journal of Economic Behavior & Organization*, 119, 32-55.
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.
HEYMAN, G. M. (2009): *Addiction: A Disorder of Choice*. Cambridge, MA: Harvard University Press.
KAHNEMAN, D. (2011): *Thinking Fast and Slow*. New York: Farrar, Straus and Giroux.
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.
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.
LICHTENDAHL, K. C., R. O. CHAO, AND S. E. BODILY (2012): “Habit Formation from Correlation Aversion,” *Operations Research*, 60, 625-637.
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.
LOOMES, G., AND R. SUGDEN (1998): “Testing Different Stochastic Specifications of Risky Choice,” *Economica*, 65, 581-598.
MIYAMOTO, J. M., AND P. P. WAKKER (1996): “Multiattribute Utility Theory without Expected Utility Foundations,” *Operations Research*, 44, 313-326.
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.
PRELEC, D. (1998): “The Probability Weighting Function,” *Econometrica*, 66, 497-527.
QUIGGIN, J. (1982): “A Theory of Anticipated Utility,” *Journal of Economic Behavior & Organization*, 3, 323-343.
RABIN, M. (2000): “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica*, 68, 1281-1292.
RICHARD, S. F. (1975): “Multivariate Risk Aversion, Utility Independence and Separable Utility Functions,” *Management Science*, 22, 12-21.
SHISANA, O., D. LABADARIOS, T. REHLE, L. SIMBAYI, K. ZUMA, A. DHANSAY, P. REDDY, W. PARKER, E. HOOSAIN, P. NAIDOO, C. HONGORO, Z. MCHIZA, N. P. STEYN, N. DWANE, M. MAKOAE, T. MALULEKE, S. RAMLAGAN, S. ZUNGU, M. G. EVANS, L. JACOBS, M. FABER, AND SANHANES-1 TEAM (2013): *South African National Health and Nutrition Examination Survey (SANHANES-1)*. Cape Town: HSRC Press.
STRZALECKI, T. (2013): “Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion,” *Econometrica*, 81, 1039-1074.
THUN, M. J., B. D. CARTER, D. FESKANICH, N. D. FREEDMAN, R. PRENTICE, A. D. LOPEZ, P. HARTGE, AND S. M. GAPSTUR (2013): “50-Year Trends in Smoking-Related Mortality in the United States,” *The New England Journal of Medicine*, 368, 351-64.
WAKKER, P. P. (2008): “Explaining the Characteristics of the Power (CRRA) Utility Family,” *Health Economics*, 17, 1329-1344.
WAKKER, P. P., I. EREV, AND E. U. WEBER (1994): “Comonotonic Independence: The Critical Test between Classical and Rank-Dependent Utility Theories,” *Journal of Risk and Uncertainty*, 9, 195-230.
WEST, R., AND J. BROWN (2013): *Theory of Addiction*. Oxford: Wiley-Blackwell (Second Edition).
WILCOX, N. T. (2011): “‘Stochastically More Risk Averse:’ a Contextual Theory of Stochastic Discrete Choice under Risk,” *Journal of Econometrics*, 162, 89-104.

APPENDIX A
[ONLINE WORKING PAPER]

The introductory presentation, atemporal risk preference task instructions, time preference task instructions, and intertemporal risk preference task instructions are included in this appendix. The introductory presentation provides an overview of the session and includes a detailed discussion of the physical randomisation devices used in the experiment. The atemporal risk preference task instructions, time preference task instructions, and intertemporal risk preference task instructions discuss the computer environment within which choices are made, the options between which the subjects must choose and how to interpret them, and the payment scheme that is used to determine earnings. The presentation and instructions were designed to promote comprehension and ensure that subjects understood how their choices ultimately led to the earnings they received so as to incentivise the truthful revelation of preferences.

A. Introductory Presentation

Introduction

Consent Form

- Before we can begin today's session you need to read and sign a consent form which you will find in the folder in front of you
- You will notice that there are 2 consent forms in the folder and one of them is for you to take home so please place it in your bag now
- The consent form explains your rights as a research participant and, by signing it, you give your consent to participate in the study
- You need to sign the consent form on the last page and when you have done so please raise your hand
- Once everyone has signed their consent forms, we can continue
- If you have any questions please raise your hand and someone will come to answer them
- You may read through the consent form now

Welcome

- 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
- Once that is done, we can begin with the tasks
- If you have any questions, please do not ask them out loud – raise your hand and someone will come over to you

4 Tasks and a Questionnaire

- You will take part in 4 tasks and you will have the opportunity to earn money in each task
- We will determine your payment for each task once you have finished that task and write it down on a payment sheet you will have beside you
- Once you have completed all 4 tasks, you will need to fill out a short questionnaire
- We will then total up your payments privately, as discussed in a moment
- Once this is done, you will be free to leave

Earnings

- You will be paid R40 just for participating in today's session
- At the end of each task, we will determine your earnings for that task
- Some of this money will be paid to you at the end of the session today, in private, 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 each task 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 dice 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 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 new picture below shows you what a 9 looks like

Note: This slide contained animations

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 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 new picture below, the number that was rolled is 100 (00 and 0)

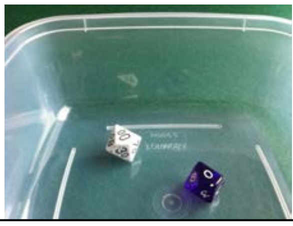

Note: This slide contained animations

The Tasks

- We have now finished the introductory explanation
- You will find instructions for the first task that you need to complete in the folder in front of you
- Please read through this and when you are finished raise your hand so that an experimenter can play a video for you which provides further details on the task
- When this is finished you will begin the first task

B. Atemporal Risk Preference Task Instructions

Task Instructions

This is a task where 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 90 pairs of lotteries in this task. For each pair of lotteries, you should choose the lottery you 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 that lottery, so you should think carefully about which lottery you prefer.

Here is an example of what the computer display of such a pair of lotteries might look like.

The outcome of the lotteries will be determined by the draw of a random number between 1 and 100. Each number between, and including, 1 and 100 is equally likely to occur. In fact, you will be able to draw the number yourself using two 10-sided dice.

In the above example, the Left lottery pays R20 with a 55% chance, R160 with a 25% chance and R190 with a 20% chance. So when you roll the two 10-sided dice if the number drawn is between 1 and 55 you will be paid R20, if the number is between 56 and 80 you will be paid R160, and if the number is between 81 and 100 you will be paid R190. The blue colour in the pie chart corresponds to 55% of the area and illustrates the chances that the number drawn will be between 1 and 55 and your prize will be R20. The orange area in the pie chart corresponds to 25% of the area and illustrates the chances that the number drawn will be between 56 and 80 and your prize will be R160. The green area in the pie chart corresponds to 20% of the area and illustrates the chances that the number drawn will be between 81 and 100 and your prize will be R190.

Now look at the Right lottery in the example. It pays R20 with a 75% chance, and R250 with a 25% chance. So when you roll the two 10-sided dice if the number drawn is between 1 and 75 you will be paid R20, and if the number is between 76 and 100 you will be paid R250. The blue colour in the pie chart corresponds to 75% of the area and illustrates the chances that the number drawn will be between 1 and 75 and your prize will be R20. The green area in the pie chart corresponds to 25% of the area and illustrates the chances that the number drawn will be between 76 and 100 and your prize will be R250.

Each pair of lotteries is shown on a separate screen on the computer. On each screen, you should indicate which lottery you prefer to play by clicking on one of the buttons beneath the lotteries.

You could also get a pair of lotteries in which one of the lotteries will give you the chance to play "Double or Nothing." For instance, the Right lottery in the following screen image pays "Double or Nothing" if the Green area is selected. The right pie chart indicates that there is a 50% chance that you get R0. So if you roll the two 10-sided dice and the number drawn is between 1 and 50 you will be paid R0. However, if the number is between 51 and 100 you will toss a coin to determine if you get double the amount listed in green (R210). If the coin comes up Heads you get R420, otherwise you get nothing. The prizes listed underneath each pie refer to the amounts <u>before</u> any "Double or Nothing" coin toss.

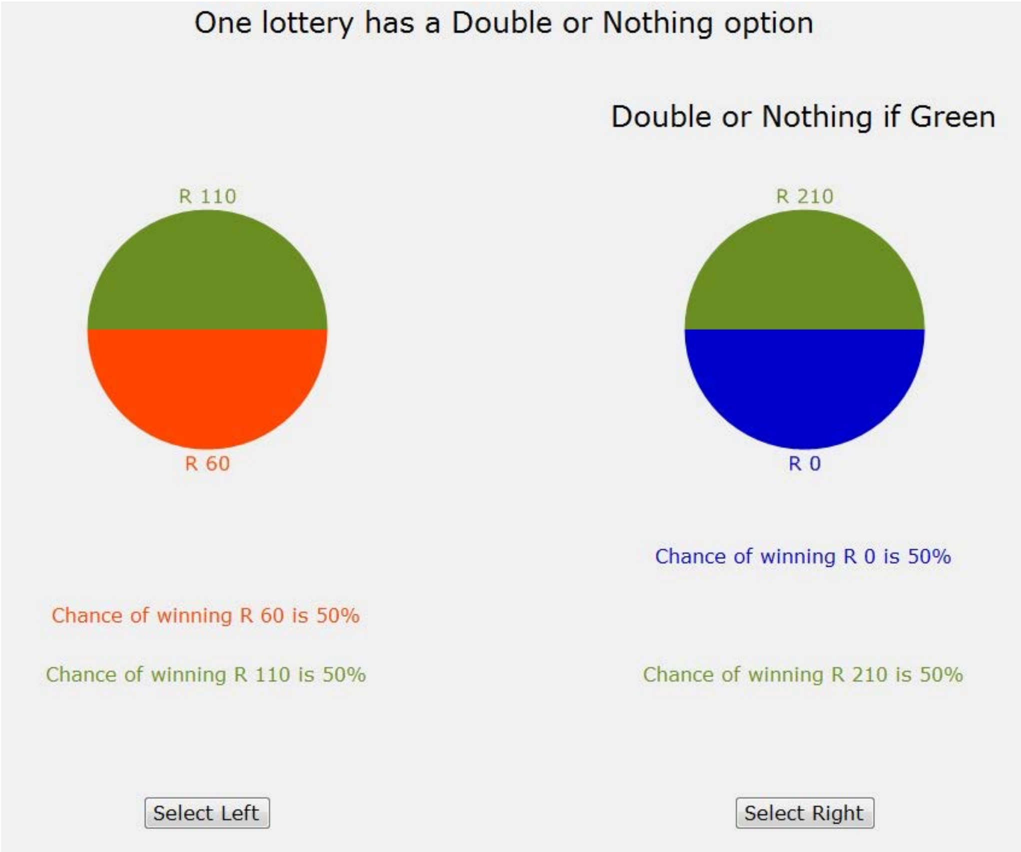

For instance, suppose you picked the lottery on the left in the last example. If the random number drawn was 37, you would win R60; if it was 93, you would get R110.

If you picked the lottery on the right and drew the number 37, you would get R0; if instead you drew 93, you would have to toss a coin to determine if you get "Double or Nothing." If the coin comes up Heads then you get R420. However, if it comes up Tails you get nothing from your chosen lottery.

After you have worked through all of the 90 pairs of lotteries, raise your hand and an experimenter will come to you to determine your payment for this task. You will roll two 10-sided dice until a number between 1 and 90 comes up to determine which pair of lotteries will be played out. Since there is a chance that any of your 90 choices could be played out for real, you should approach each pair of lotteries as if it is the one that you will play out. Finally, you will roll the two ten-sided dice again to determine the outcome of the lottery you chose, and if necessary you will then toss a coin to determine if you get "Double or Nothing."

It is also possible that you will be given a lottery in which there is a "Double or Nothing" option no matter what number you roll with the two 10-sided dice. The screen image below illustrates this possibility. The Right lottery in the example pays "Double or Nothing" for any number that is drawn with the two 10-sided dice. So if you select the Right lottery and roll a number between 1 and 50 you will toss a coin to see whether you get R0 or R120 (double R60). If you roll a number between 51 and 100 you will toss a coin to see whether you get R0 or R420 (double R210).

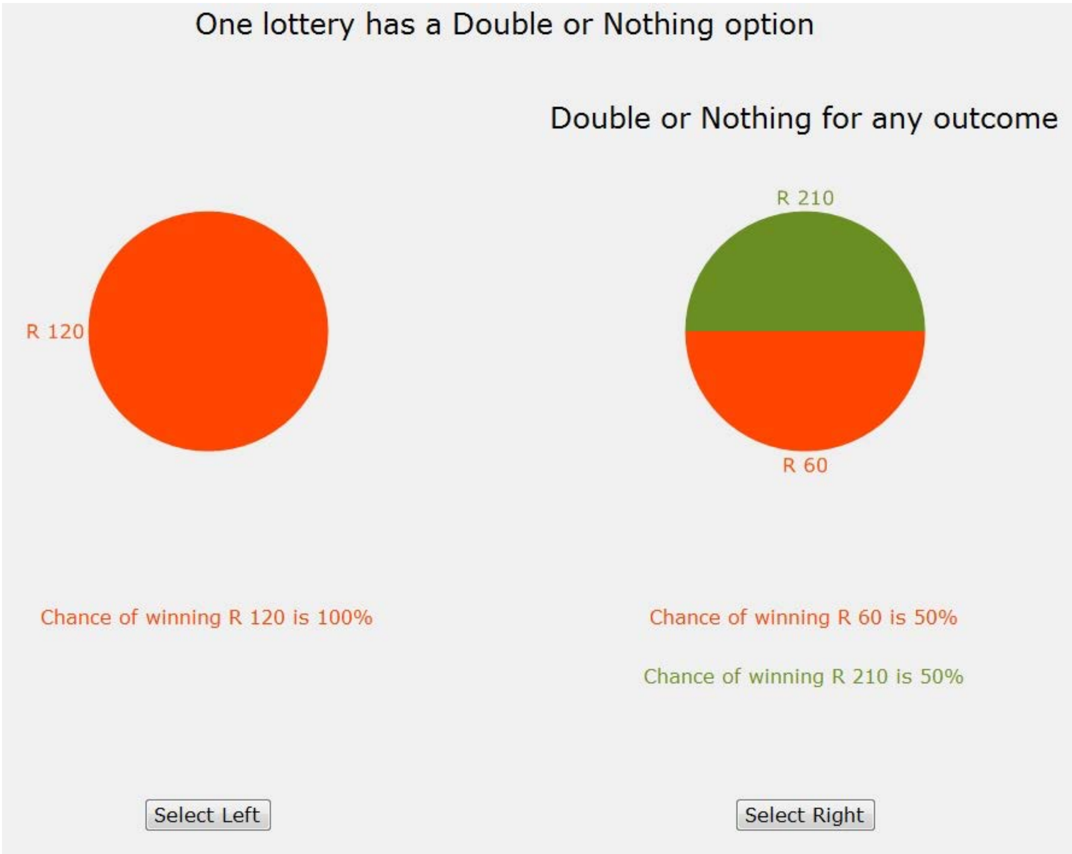

Therefore, your earnings for this task are determined by four things:

- by which lottery you selected, the Left or the Right, for each of these 90 pairs;
- by which lottery pair is chosen to be played out in the set of 90 such pairs using the two 10-sided dice;
- by the outcome of that lottery when you roll the two 10-sided dice; and
- by the outcome of a coin toss if the chosen lottery outcome is of the "Double or Nothing" type.

Which lotteries you prefer is a matter of personal taste. The people next to you may be presented with different lotteries, and may have different preferences, so their responses should not matter to you. Please work silently, and make your choices by thinking carefully about each lottery.

Payment for this task is in cash, and is in addition to the R40 show-up fee that you receive just for being here. When you have finished the task, please raise your hand and an experimenter will come to you to determine your payment for this task.

C. Time Preference Task Instructions

Task Instructions

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 you chose at the appropriate date.

All of these choices will be made on a computer and here is an example of what the computer display might look like:

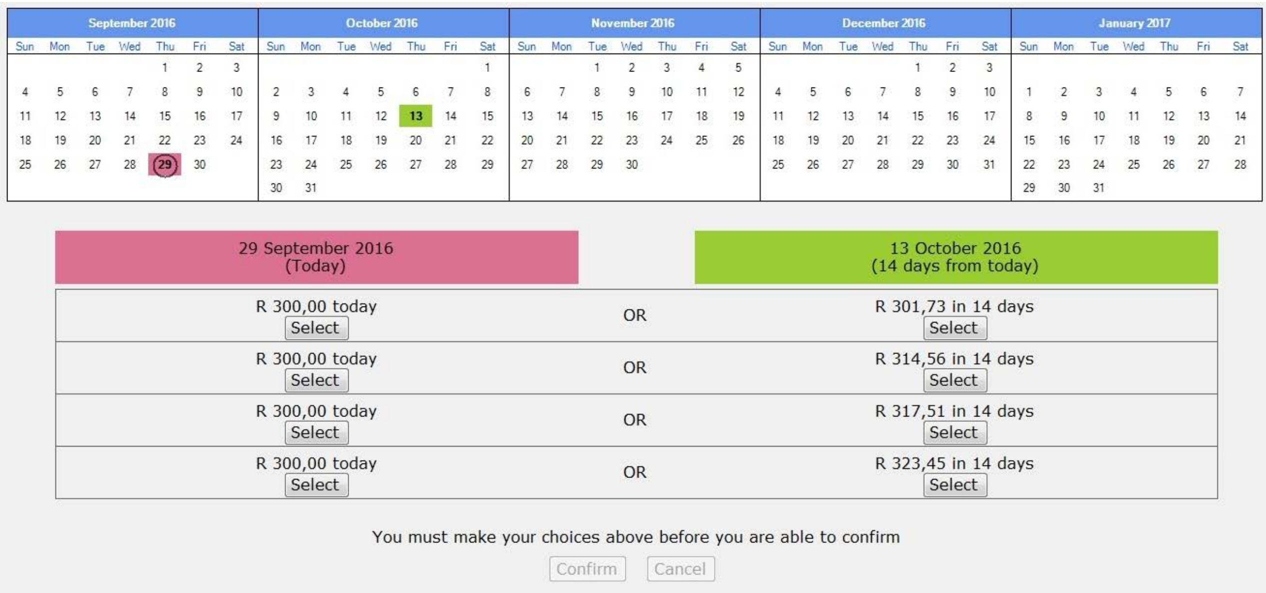

For the purpose of explaining this task, assume for the moment that today is 29 September, 2016. At the top of the display is a calendar showing you today's date in a circle (29 September 2016). This date is also highlighted in purple and a future date is highlighted in green (13 October 2016). Below the calendar are two columns: a purple column with amounts of money available at an earlier date (today) and a green column with amounts of money available at a later date (in 14 days from today). You need to make 4 choices on this screen. Each choice appears on a different row.

In the first row, you need to choose between receiving R300 today or R301.73 in 14 days from today. Note that R300 is the smaller of the two amounts but it is available today. R301.73 is the larger of the two amounts but it is only available after 14 days. Suppose that you prefer R300 today over R301.73 in 14 days from today. To choose R300 today just click the button saying "Select" under "R300 today".

Suppose instead that you prefer R301.73 in 14 days rather than R300 today. To choose R301.73 in 14 days just click the button saying "Select" under "R301.73 in 14 days".

Once you have made your choice on the first row you can move on to the other rows on the screen. 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. Once you have made all of your choices on the screen you can click the button saying "Confirm" to move on to the next screen. If you would like to change your choices then click "Cancel".

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 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 7 days and another amount of money available in 21 days. So please pay careful attention when making your choices.

When you are finished the task, please raise your hand and an experimenter will come to you to determine your payment for this task. You will select one of the 15 screens from this task by rolling a 20-sided dice. If the dice lands on 1, you will select screen 1; if the dice lands on 7, you will select screen 7; if the dice lands on 12, you will select screen 12; and so on. If the dice lands on 16, 17, 18, 19 or 20, you will roll the dice again until it lands on a number between 1 and 15.

Once you have selected a screen, you will roll a 4-sided dice to select 1 of the 4 rows on the screen. If the dice lands on 1, you will select row 1; if the dice lands on 2, you will select row 2; and so on. Once you have selected the row, we will look at the choice that 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. For instance, in the last example, suppose that row 3 is selected for payment. If you chose R300 today, you will be paid R300 at the end of today's session. If you chose R317.51 in 14 days then you will be paid R317.51 in 14 days via electronic transfer into your bank account and you will receive a payment notification on your cellphone when the transaction has taken place. That is why we need your bank account details: to pay you via electronic transfer, if necessary.

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 for 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.

Task Instructions

In this task you will make a number of choices between two options that we can think of as the TOP and BOTTOM options. An example of a choice that you will need to make is shown below.

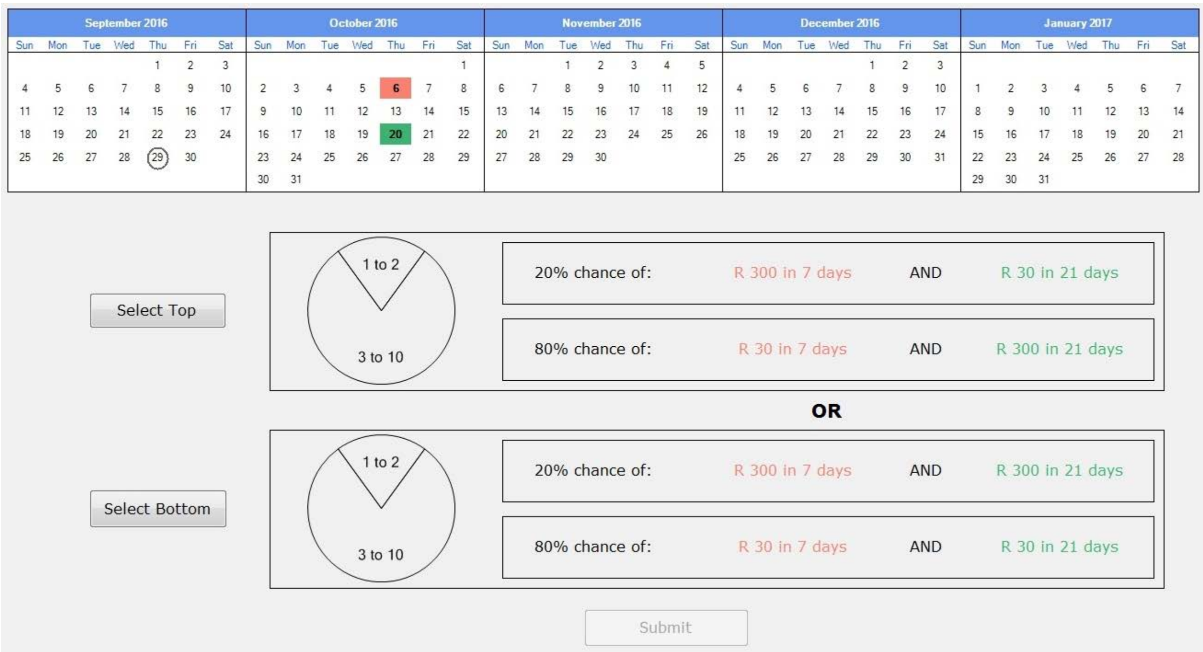

You will need to make 40 choices in total across 40 screens. On each screen, you should choose the option you prefer.

The outcome of each option will be determined by the draw of a random number between 1 and 10. Each number is equally likely to occur, and you will draw the number yourself using a 10-sided dice.

In the example, the TOP option pays R300 in 7 days AND R30 in 21 days if the number is 1 or 2. It pays R30 in 7 days AND R300 in 21 days if the number is between 3 and 10.

The BOTTOM option pays R300 in 7 days AND R300 in 21 days if the number is 1 or 2. It pays R30 in 7 days AND R30 in 21 days if the number is between 3 and 10.

When you are finished the task, please raise your hand and an experimenter will come to you to determine your payment for this task. You will be paid for one of your choices in this task. You will select one of the 40 choices you made by rolling a 4-sided dice and a 10-sided dice. If you roll 1 on the 4-sided dice, you will select choices 1-10; if you roll 2 on the 4-sided dice, you will select choices 11-20; if you roll 3 on the 10-sided dice, you will select choices 21-30; and if you roll 4 on the 4-sided dice, you will select choices 31-40. You will then roll the 10-sided dice to select a number between one of these ranges. For example, suppose you roll 3 on the 4-sided dice. Then you will select choices 21-30. If you then roll 7 on the 10-sided dice you will select choice 27. Once the choice has been selected, you will then roll the 10-sided dice again to determine the payment for the decision that you made. Any future payments will be made via electronic transfer into your bank account and you will receive a payment notification on your cellphone when the transaction has taken place. That is why we need your bank account details: to pay you via electronic transfer.

If the example above is selected for payment and you chose the TOP option, you will roll the 10-sided dice to determine your earnings for this task. If you roll a 8 then you will be paid R30 in 7 days AND R300 in 21 days.

By contrast, if the example above is selected for payment and you chose the BOTTOM option, you will roll the 10-sided dice to determine your earnings for this task. If you roll a 5 then you will be paid R30 in 7 days AND R30 in 21 days.

Note that the option you prefer 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 40 choices could be selected for payment, you should approach each choice as if it is the one you will be paid for.

APPENDIX B
[ONLINE WORKING PAPER]

Risk Preference Task Lotteries

The 90 lottery pairs used in the risk preference task were drawn from the designs of Wakker, Erev and Weber (WEW) [1994], Loomes and Sugden (LS) [1998], Cox and Sadiraj (CS) [2008, p. 33], and Harrison, Martínez-Correa and Swarthout (HMS) [2015].

WEW constructed a battery of lotteries to test the "comonotonic independence" axiom of rank-dependent utility (RDU) theory, due to Quiggin [1982]. Their main lottery pairs consist of 6 sets of 4 pairs. The logic of their design can be seen by considering the first set [WEW, p. 204, Figure 3.1]. The second and third prizes in each pair stay the same within the set of 4 lottery pairs. The only thing that varies from pair to pair is the monetary value of the first prize, and that is common to the two lotteries within each pair. Since the first prize is a common consequence in both lotteries within a pair, the independence axiom of expected utility theory (EUT) implies that it should not affect choices. In the 1^{st} pair the first prize is only \$0.50, and it is the lowest ranked prize for both lotteries. The first prize increases to \$3.50 in the 2^{nd} pair, but it is again the lowest ranked prize for both lotteries. Consequently, rank-dependence should have no effect on choice patterns as the subject moves from the 1^{st} to the 2^{nd} pair. By contrast, the first prize in the 3^{rd} pair is \$6.50, which makes it the second highest ranked prize for both lotteries; this is where RDU *could* generate a different prediction to EUT, depending on the nature and extent of probability weighting. Finally, in the 4th pair the common consequence of \$9.50 is the highest ranked prize for both lotteries, again allowing RDU to predict something different to EUT, and to the choices in the 3^{rd} pair. This design does not formally require a RDU decision-maker to choose differently to an EUT decision-maker, but simply allows it for *a priori* reasonable levels of probability weighting. We used all 24 of the main WEW lottery pairs and scaled the prizes considerably.

LS designed 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 EUT. We used 30 lottery pairs from the LS design which provided a thorough and well-balanced coverage of the Marschak-Machina (MM) triangle and captured the full range of risk preferences, under the null hypothesis of EUT: risk seeking - gradients less than 1; risk neutral - gradients equal to 1; and risk averse - gradients greater than 1.

CS designed a simple test of the calibration puzzle posed by Hansson [1988] and Rabin [2000]: that the risk aversion which is observed with small stakes in the lab yields implausible levels of risk aversion with larger stakes. The logic of the CS design is as follows: give people choices between safe and risky lotteries, where the safe lotteries are certain amounts of money, and the risky lotteries are a 50:50 chance of -y/+x either side of the certain amount of money in the safe lottery. For each lottery pair, x > y so that the expected value of the risky lottery is slightly larger than the value of the safe lottery. Across a set of lottery pairs, the value of the safe prize varies, but x and y are held constant. The idea behind this test of the calibration puzzle is to regard the safe lottery as "lab wealth," and then see if subjects are risk averse as one varies lab wealth. For example, suppose -y/+x = -$10/+$15, then consider two binary choices: one where the safe lottery is $20 and another where the safe lottery is $100. The subject then makes two choices: take $20 for certain, or take a 50:50 chance of $10 or $35; and take $100 for certain, or take a 50:50 chance of $90 or $115. The Hansson-Rabin premiss is that one gets risk aversion in both cases, with a majority of people picking the safe lottery. We used 6 lottery pairs from the implementation of the CS design in Harrison, Lau, Ross and Swarthout [2017]: 3 pairs where -y/+x = -R60/+R70 and the safe options were R120, R320, and R520; and 3 pairs where -y/+x = -R30/+R40 and the safe options were R60, R340, and R540.

HMS designed lotteries to test the reduction of compound lotteries (ROCL) axiom, which states that a decision-maker is indifferent between a multi-stage compound lottery and the actuarially-equivalent simple lottery where the probabilities of the stages of the compound lottery have been multiplied out. Given a simple (S) lottery and compound (C) lottery, HMS create an actuarially-equivalent (AE) lottery from a two-stage C lottery by multiplying out the probabilities of the two-stages, and then construct three pairs of lotteries: a S-C pair, a S-AE pair, and an AE-C pair. They used probabilities drawn from {0, ¼, ½, ¾, 1} and final prizes of {$0, $10, $20, $35, $70}. The compound lotteries were created using a "double or nothing" (DON) procedure so the first-stage prizes in a compound lottery were drawn from {$5, $10, $17.5, $35}. The second-stage DON procedure then provides the set of final prizes above, which is either $0 or double the stakes of the first stage.

Most of the HMS compound lotteries used a conditional version of DON, where the initial lottery triggered the DON procedure only if a particular outcome was realised in the initial lottery. For example, consider the compound lottery formed by an initial lottery that pays $10 and $20 with equal probability. The DON stage is reached if the outcome of the initial lottery is $10. Then, in the subsequent DON lottery, the subject has an equal chance of winning $20 (i.e., double $10) or $0 (i.e., nothing). Alternatively, if the realised outcome of the initial lottery is $20, the DON stage is not triggered, and the subject earns $20. Figure 2 in HMS [p. 35] shows a tree representation of this compound lottery and the corresponding actuarially-equivalent simple lottery. The benefit of using a conditional DON lottery is that it allows one to obtain better coverage of the MM triangle relative to unconditional DON (see p. 35-36 of HMS for more details). This allows for variation in both prizes and probability distributions so that one can identify source-dependent preferences that take into account attitudes toward variability in prizes and variability in probabilities. We used 30 lottery pairs from the HMS design: 15 S-C pairs and 15 S-AE pairs. Hence we have a data-based metric, between 0 and 15, for each subject's consistency with the ROCL axiom.

ADDITIONAL REFERENCES

HARRISON, G. W., M. I. LAU, D. ROSS, AND J. T. SWARTHOUT (2017): "Small Stakes Risk Aversion in the Laboratory: A Reconsideration," *Economics Letters*, 160, 24-28.

APPENDIX C

In this appendix we estimate the SDU model (9) in the main text under the assumptions that EUT characterises choice under atemporal risk and that discounting is exponential. We then analyse the atemporal risk preference data and the time preference data to determine whether RDU, as opposed to EUT, better characterises choice under atemporal risk and whether there is evidence of QH discounting. We find statistically significant evidence of nonlinear probability weighting which has an economically significant impact on estimates of the power utility function parameter r and the exponential discount rate. While the sample as a whole is better characterised by RDU than EUT, we conduct supplementary analyses at the level of the individual where we estimate EUT and RDU specifications for each subject and then test whether $\omega(p) = p$ in the RDU model. We find that at least half of the sample is better characterised by EUT than RDU, which adds nuance to our atemporal risk preference results. In analyses of the atemporal risk preference and time preference data, we find evidence of QH discounting which has an economically significant effect on estimates of the long-term discounting parameter δ. Finally, we calculate and graph risk premia for intertemporal lotteries A and B to show the effect of intertemporal risk aversion on attitudes to risk over time.

Table C1 presents estimates of the SDU model (9) under the assumptions that EUT characterises choice under atemporal risk and that discounting is exponential. As discussed in the main text, the estimate of the intertemporal risk preference parameter ρ = -1.043 implies a high level of intertemporal risk aversion.

Table C1
Intertemporal Risk Preference ML Estimates
EUT, Exponential Discounting
Homogenous Preferences

	Model
	Estimate
Atemporal Risk Preferences	
Power function parameter (*r*)	0.409***
	(0.019)
Error (μ)	0.167***
	(0.007)
Time Preferences	
Discounting parameter (δ)	0.782***
	(0.062)
Error (υ)	0.360***
	(0.062)
Intertemporal Risk Preferences	
Power function parameter (ρ)	-1.043***
	(0.290)
Error (ψ)	0.312***
	(0.021)
N	48640
log-likelihood	-28973.300

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

Table C2 presents estimates of an RDU model with a power utility function and the Prelec PWF. The estimate of φ = 0.629 is significantly less than 1 ($p < 0.001$) which gives the PWF an inverse S-shape form. The estimate of *r* under the RDU model is statistically significantly higher than the estimate of *r* under the EUT model ($p < 0.001$), implying that it is necessary to estimate time preference models jointly with an RDU model of atemporal risk preferences.

Table C2
Atemporal Risk Preference ML Estimates
Homogenous Preferences

	Model 1	**Model 2**
	EUT	RDU
Power function parameter (*r*)	0.408***	0.553***
	(0.019)	(0.023)
PWF parameter (φ)		0.629***
		(0.020)
PWF parameter (η)		1.020***
		(0.031)
Error (μ)	0.167***	0.145***
	(0.007)	(0.005)
N	23040	23040
log-likelihood	-15030.136	-14696.023

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

While the sample as a whole is better characterised by RDU than by EUT, this does not imply that every person in the sample probability weights and, therefore, departs from EUT. Figure C1 shows the results from an individual-level analysis where we estimate EUT and RDU specifications for each subject and then test whether $\omega(p) = p$ in the RDU model. Using a 5% level of statistical significance, we cannot reject the hypothesis that $\omega(p) = p$ for 57% of the sample, implying that at least half of the people in our study are better characterised by EUT than RDU. However, the remaining 43% of the sample exhibits statistically significant evidence of nonlinear probability weighting. Hence it is necessary to take this probability weighting into account when drawing inferences about time preferences and intertemporal risk preferences at the level of the *sample of subjects*.

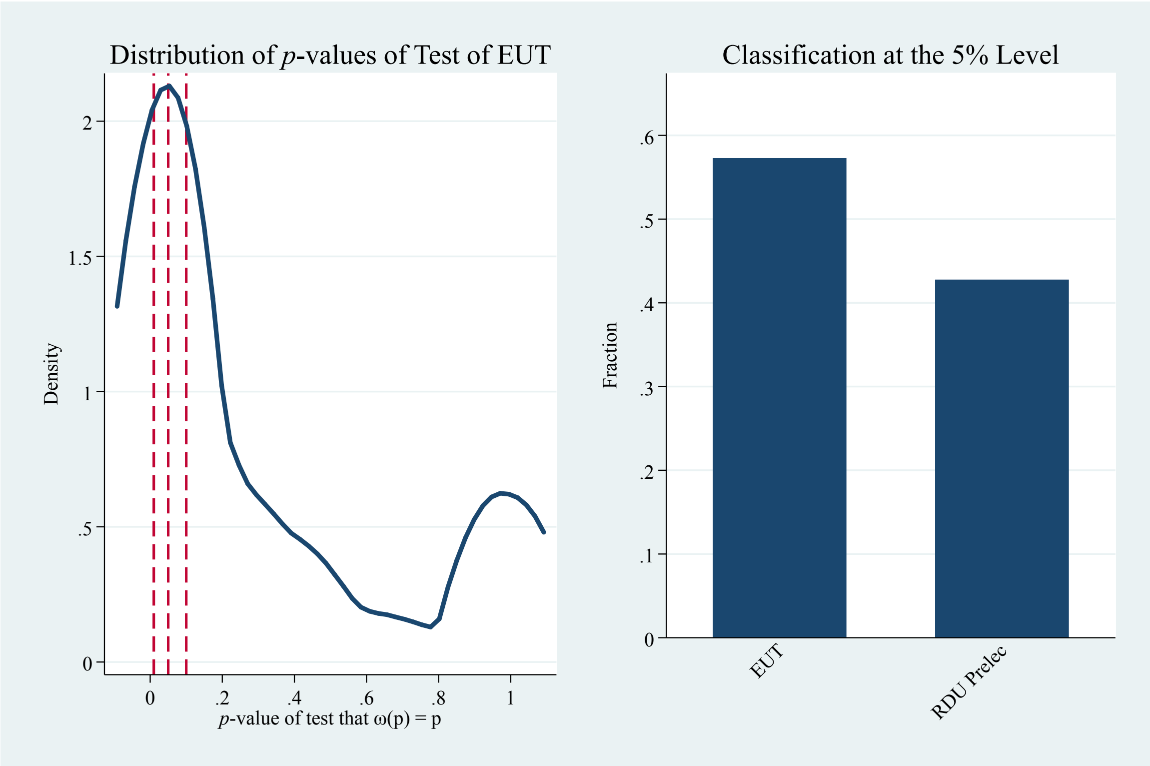

Figure C1: Classifying Subjects as EUT or RDU

Table C3 presents estimates of two exponential discounting models. Model 1 assumes that EUT and a power utility function characterise choice under atemporal risk whereas Model 2 assumes RDU with a power utility function and the Prelec PWF. The estimate of the long-term discounting parameter δ is statistically significantly higher under the RDU model relative to the EUT model ($p < 0.001$), which highlights the way in which atemporal risk preference estimates propagate into estimates of discounting parameters.

Table C3
Exponential Discounting Function ML Estimates
EUT and RDU, Homogenous Preferences

	Model 1 EUT	**Model 2** RDU
Power function parameter (*r*)	0.410***	0.553***
	(0.019)	(0.022)
PWF parameter (φ)		0.629***
		(0.020)
PWF parameter (η)		1.021***
		(0.031)
Discounting parameter (δ)	0.785***	1.192***
	(0.063)	(0.101)
Risk error (μ)	0.167***	0.145***
	(0.007)	(0.005)
Time error (ν)	0.362***	1.106***
	(0.062)	(0.199)
N	38400	38400
log-likelihood	-23681.626	-23345.72

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

Table C4 shows the statistically and economically significant influence of QH discounting on estimates of the long-term discounting parameter δ. Under the QH model, there is a sharp drop in the value of a reward if it is not available immediately but this drop asymptotes toward the long-term discounting parameter δ over time. In the exponential model, by contrast, the discount rate δ does not vary over time and remains at the far higher level of 1.192.

Table C4
Discounting Function ML Estimates
Rank-Dependent Utility, Homogenous Preferences

	Model 1	**Model 3**
	Exponential	Quasi-Hyperbolic
Power function parameter (*r*)	0.553***	0.541***
	(0.022)	(0.021)
PWF parameter (φ)	0.629***	0.632***
	(0.020)	(0.020)
PWF parameter (η)	1.021***	1.012***
	(0.031)	(0.030)
Discounting parameter (δ)	1.192***	0.885***
	(0.101)	(0.073)
Discounting parameter (β)		0.960***
		(0.003)
Risk error (μ)	0.145***	0.145***
	(0.005)	(0.005)
Time error (ν)	1.106***	1.013***
	(0.199)	(0.176)
N	38400	38400
log-likelihood	-23345.720	-22939.946

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

To understand the role that intertemporal risk aversion plays in the characterisation of attitudes to risk over time, we calculate certainty equivalents and then evaluate risk premia for intertemporal lotteries A and B, under the standard assumption of intertemporal risk neutrality and using the intertemporal risk preference estimates in Table 3 of the main text that incorporate intertemporal risk aversion. Following the approach in AHLR [p. 544-546], the certainty equivalent of intertemporal lottery A can be calculated as

$$\mathrm{CE_A} = [(\mathrm{SDU_A}^{1/\rho}) / (D_t + D_{t+\tau})]^{1/r}, \tag{1}$$

where $\mathrm{CE_A}$ is received in both the sooner and the later time periods, and $\mathrm{SDU_A}$ is the stochastic discounted utility of intertemporal lottery A given by (4) in the main text. The risk premium is then defined as

$$\mathrm{RP_A} = \mathrm{EV_A} - (D_t + D_{t+\tau}) \times \mathrm{CE_A}, \tag{2}$$

where $\mathrm{EV_A} = p \times [\mathrm{D}_t L_t + \mathrm{D}_{t+\tau} S_{t+\tau}] + (1 - p) \times [\mathrm{D}_t S_t + \mathrm{D}_{t+\tau} L_{t+\tau}]$.

Similarly, the certainty equivalent of intertemporal lottery B can be calculated as

$$\mathrm{CE_B} = [\omega(p) \times L^{r\rho} + (1 - \omega(p)) \times S^{r\rho}]^{1/(r\rho)}, \tag{3}$$

where $\mathrm{CE_B}$, like $\mathrm{CE_A}$ in (1), is received in both the sooner and the later time periods. Unlike intertemporal lottery A, the certainty equivalent of intertemporal lottery B is independent of the discount factor.[1] The risk premium is then

$$\mathrm{RP_B} = \mathrm{EV_B} - \mathrm{CE_B}, \tag{4}$$

where $\mathrm{EV_B} = p \times L_t + (1 - p) \times S_t$.

To show the effect of intertemporal risk aversion on attitudes to risk over time, we use the estimates in Table 3 of the main text to calculate risk premia for intertemporal lottery A and intertemporal lottery B. We then re-estimate the model in Table 3 but impose intertemporal risk neutrality ($\rho = 1$). As emphasised in the main text, in a joint estimation framework, atemporal risk preference, time preference, and intertemporal risk preference estimates are inextricably linked, so when we impose the intertemporal risk neutrality constraint this leads to different estimates of the other parameters in the system. Specifically, imposing intertemporal risk neutrality leads to a much lower estimate of the atemporal risk preference parameter r. In effect, the atemporal risk preference parameter r has to account for the risk aversion that the intertemporal risk preference parameter ρ identifies when it is not

[1] The certainty equivalent of intertemporal lottery B is independent of the discount factor because in our experimental design $S_t = S_{t+\tau}$ and $L_t = L_{t+\tau}$, implying that the subject either receives the smaller or larger reward in the sooner *and* later time periods when choosing intertemporal lottery B. With an experimental design where intertemporal lottery B paid out different amounts in the two time periods, the certainty equivalent would not be independent of the discount factor.

constrained to 1. This change in the estimate of *r* has a marked impact on the calculation of risk premia for intertemporal lottery A but a far smaller impact on the calculation of risk premia for intertemporal lottery B.

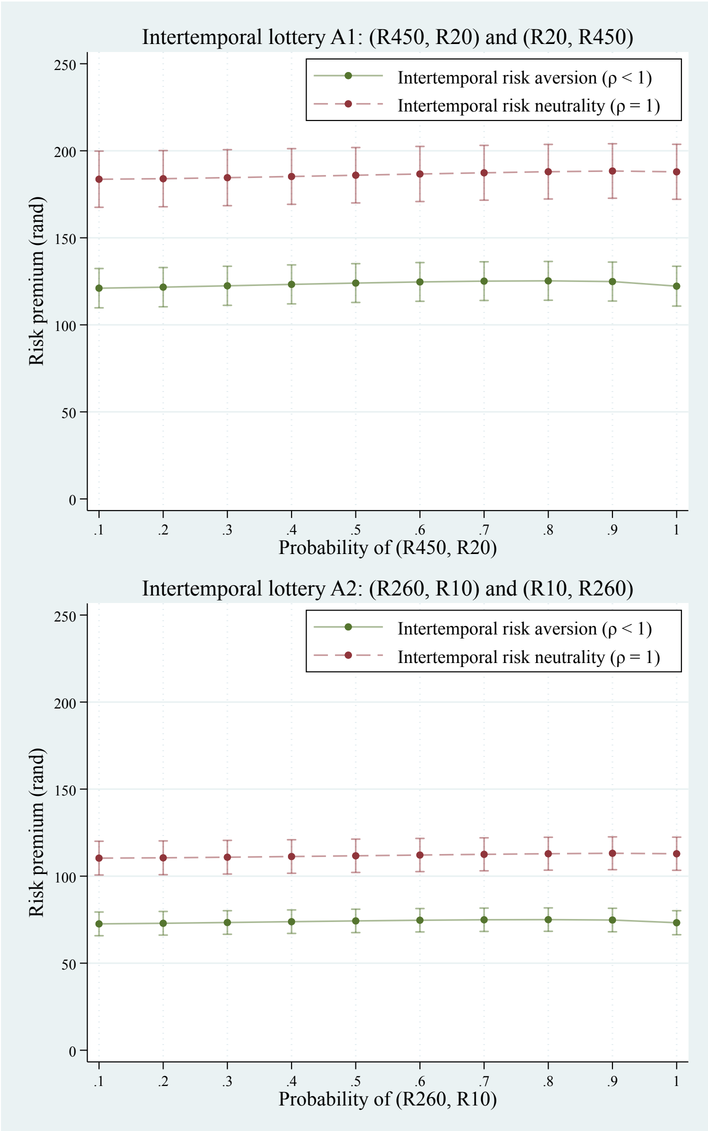

Figure C2: Risk Premia for Intertemporal Lottery A

Figure C2 shows risk premia, with 95% confidence intervals, for intertemporal lottery A using the estimates in Table 3 of the main text, which incorporate intertemporal risk aversion, and when we impose intertemporal risk neutrality.[2] The risk premia are significantly larger for intertemporal lottery A under the assumption of intertemporal risk

[2] Equations (1) and (2) in this appendix show that the risk premia for intertemporal lottery A are a function of the estimated discount factor, which in a QH discounting framework differs according to the time horizon. Figure C2 plots the risk premia for the 42-day horizon between the rewards in the intertemporal risk preference task but the results for the 14-day horizon are qualitatively identical.

neutrality in comparison to intertemporal risk aversion. This difference is driven by the far lower estimate of the atemporal risk preference parameter r when we impose intertemporal risk neutrality, which produces a lower certainty equivalent and higher risk premium, than when we allow for intertemporal risk aversion. Thus, intertemporal risk aversion makes the "safe" intertemporal lottery A more attractive to a decision maker compared to the situation where we impose intertemporal risk neutrality.

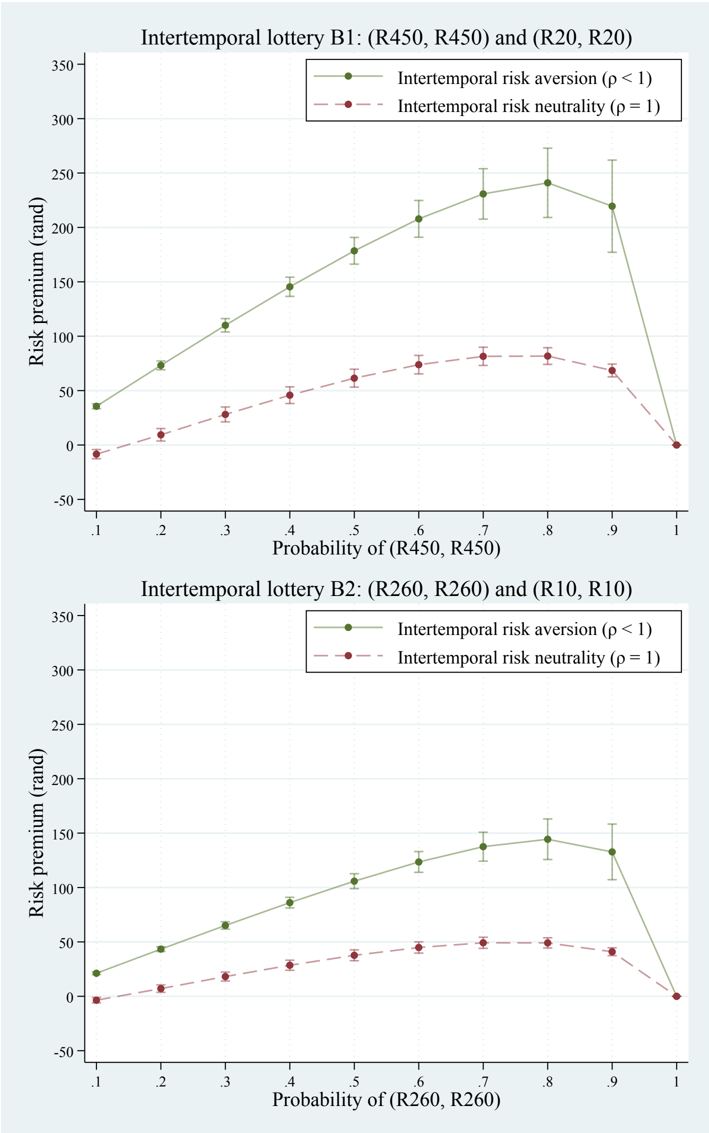

Figure C3: Risk Premia for Intertemporal Lottery B

Figure C3 shows risk premia, with 95% confidence intervals, for intertemporal lottery B using the estimates in Table 3 of the main text, which incorporate intertemporal risk aversion, and when we impose intertemporal risk neutrality. Risk premia for intertemporal lottery B are much higher under intertemporal risk aversion compared to intertemporal risk neutrality. Intertemporal risk aversion clearly accounts for a large amount of the estimated risk premium of intertemporal lottery B, thereby making it far less attractive to the decision maker relative to the situation where we impose intertemporal risk neutrality.

Finally, Figure C4 combines Figure C2 and Figure C3, with a common *y*-axis in all of the panels, to emphasise the economic significance of incorporating intertemporal risk aversion as opposed to assuming intertemporal risk neutrality. Comparing the top two panels of Figure C4, which represent the risk premia associated with intertemporal lottery A1 and intertemporal lottery B1, the *difference* in risk premia for A1, under intertemporal risk aversion relative to intertemporal risk neutrality, is relatively constant as a function of the probability assigned to the rewards (R450, R20) and varies between R61.95 and R65.70. By contrast, the *difference* in risk premia for B1, under intertemporal risk aversion relative to intertemporal risk neutrality, varies markedly across the probability space, starting at R43.99 when $p = 0.1$, increasing steadily until a maximum of R159.24 is reached when $p = 0.8$, and then declining to zero when $p = 1$. The differences in risk premia for intertemporal lottery B1 are larger than the differences in risk premia for intertemporal lottery A1 for $p \in [0.3, 0.9]$, which shows the large *economic* impact that intertemporal risk aversion has on estimated risk premia for the "risky" intertemporal lottery B1.

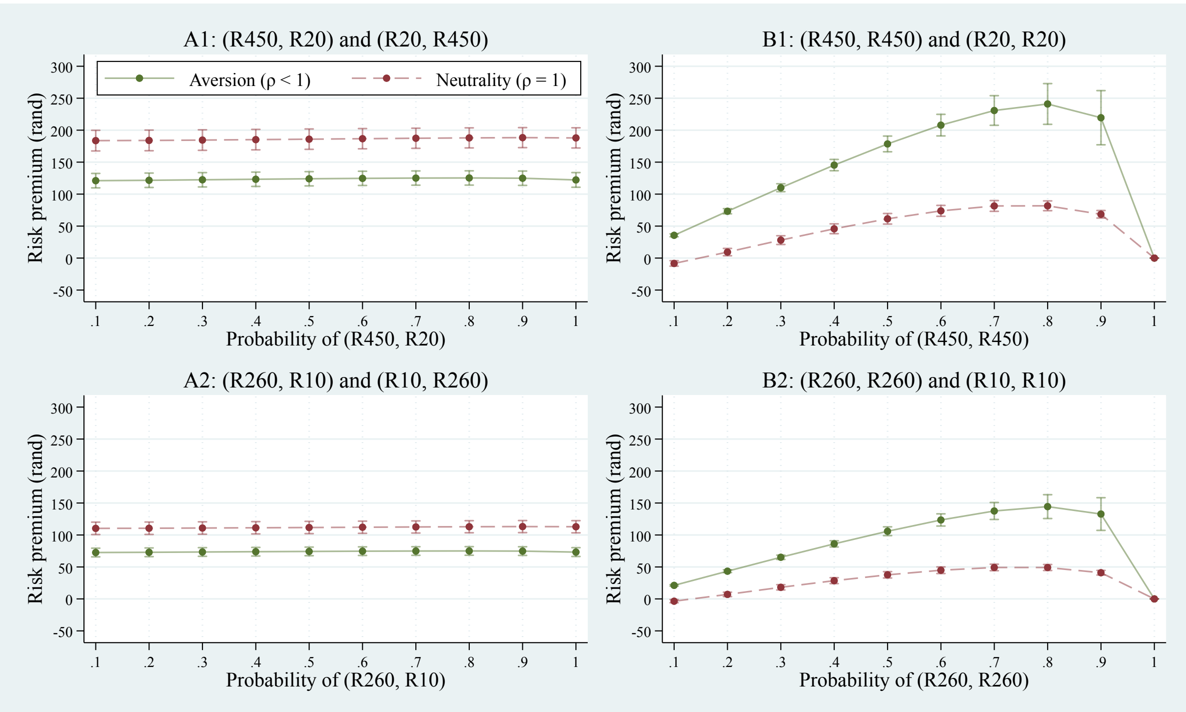

Figure C4: Risk Premia for Intertemporal Lottery A and Intertemporal Lottery B

The same qualitative pattern holds for the bottom two panels of Figure C4, which represent the risk premia associated with intertemporal lottery A2 and intertemporal lottery B2. The *difference* in risk premia for A2, under intertemporal risk aversion relative to intertemporal risk neutrality, is relatively constant as a function of the probability assigned to the rewards (R260, R10) and varies between R37.41 and R39.63. By contrast, the *difference* in risk premia for B2, under intertemporal risk aversion relative to intertemporal risk neutrality, varies significantly across the unit interval, starting at R24.73 when $p = 0.1$, increasing steadily until a maximum of R95.25 is reached when $p = 0.8$, and then declining to zero when $p = 1$. The differences in risk premia for intertemporal lottery B2 are larger than the differences in risk premia for intertemporal lottery A2 for $p \in [0.3, 0.9]$, which shows again the large *economic* impact that intertemporal risk aversion has on estimated risk premia for the "risky" intertemporal lottery B2.

APPENDIX D
[ONLINE WORKING PAPER]

In this appendix we present results from the SDU model (9) in the main text estimated jointly with a RDU model, power utility function, and Prelec PWF for choice under atemporal risk, and a QH discounting function. We allow the parameters of the SDU model to vary as a linear function of demographics, socio-economic characteristics, and three measures of smoking behaviour: smoking status (Table D1); smoking intensity, measured by the number of cigarettes smoked per day (Table D3); and smoking severity, measured by smokers' scores on the FTCD (Table D4). These latter two tables are split according to gender given the historical differences in smoking prevalence between men and women, and the economically and statistically significant differences in their intertemporal risk preferences. We also present atemporal risk preference results in Table D2 and Table D5 to corroborate the discussion in the main text that there are no substantive differences in the atemporal risk preferences of men and women as a function of smoking intensity and smoking severity, respectively.

Table D1
Intertemporal Risk Preference ML Estimates
RDU, Quasi-Hyperbolic Discounting
Heterogenous Preferences

	Model	
	Estimate	Std error
Atemporal risk preference parameter (*r*)		
Age	-0.002	0.001
White	0.010	0.020
Male	0.008	0.016
Financial situation	0.010	0.009
Staff member	0.020	0.026
Ex-smoker	0.003	0.032
Smoker	0.028	0.019
Constant	0.526***	0.041
PWF parameter (φ)		
Age	<0.001	0.003
White	0.087*	0.051
Male	0.137***	0.044
Financial situation	0.009	0.020
Staff member	0.036	0.075
Ex-smoker	0.026	0.072
Smoker	0.013	0.043
Constant	0.596***	0.088

Table D1 (Continued)

	Model	
	Estimate	Std error
PWF parameter (η)		
Age	0.003	0.004
White	-0.018	0.066
Male	-0.080	0.058
Financial situation	0.069**	0.032
Staff member	-0.118	0.092
Ex-smoker	0.172	0.107
Smoker	0.039	0.061
Constant	0.823***	0.124
Discounting parameter (β)		
Age	<0.001	<0.001
White	0.014**	0.006
Male	-0.002	0.005
Financial situation	0.001	0.003
Staff member	0.003	0.006
Ex-smoker	-0.005	0.009
Smoker	-0.007	0.005
Constant	0.970***	0.012
Discounting parameter (δ)		
Age	<0.001	0.005
White	-0.254**	0.100
Male	0.020	0.096
Financial situation	-0.211***	0.061
Staff member	-0.186	0.134
Ex-smoker	0.111	0.143
Smoker	0.356***	0.130
Constant	1.521***	0.230
Intertemporal risk preference parameter (ρ)		
Age	<0.001	0.027
White	-0.569	0.500
Male	0.841**	0.382
Financial situation	-0.611**	0.261
Staff member	1.064	0.688
Ex-smoker	-1.958*	1.180
Smoker	-0.288	0.411
Constant	0.537	0.832
Error terms		
μ	0.140***	0.005
υ	0.761***	0.135
ψ	0.275***	0.023
N	47310	
log-likelihood	-26859.893	

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

Table D2

Atemporal Risk Preference ML Estimates

Rank-Dependent Utility Theory

Smoking Intensity: Number of Cigarettes Smoked per Day

	Model 1		Model 2	
	Male		Female	
	Estimate	Std Error	Estimate	Std Error
Power function parameter (*r*)				
Age	-0.003	0.005	-0.010***	0.004
White	-0.070	0.075	-0.009	0.069
Financial situation	0.036	0.041	0.020	0.030
Staff member	0.104	0.123	0.211*	0.122
Number of cigarettes	-0.003	0.007	0.010	0.007
Constant	0.605***	0.138	0.653***	0.115
PWF parameter (φ)				
Age	0.008*	0.004	0.003	0.003
White	0.099	0.069	0.061	0.060
Financial situation	-0.040	0.029	0.014	0.026
Staff member	-0.021	0.096	-0.035	0.083
Number of cigarettes	-0.014***	0.005	0.005	0.004
Constant	0.623***	0.105	0.442***	0.112
PWF parameter (η)				
Age	0.016*	0.010	-0.008	0.007
White	-0.095	0.079	0.029	0.141
Financial situation	0.013	0.050	0.109**	0.048
Staff member	-0.145	0.156	0.065	0.140
Number of cigarettes	-0.002	0.006	0.008	0.012
Constant	0.608**	0.256	0.964***	0.167
Error (μ)				
Constant	0.128***	0.007	0.157***	0.009
N	9900		12510	
log-likelihood	-6269.601		-7901.332	

Results account for clustering at the individual level

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table D3

Intertemporal Risk Preference ML Estimates

RDU, Quasi-Hyperbolic Discounting

Smoking Intensity: Number of Cigarettes Smoked per Day

	Model	
	Estimate	Std error
Atemporal risk preference parameter (*r*)		
Age	-0.002	0.001
White	0.005	0.019
Male	0.007	0.016
Financial situation	0.012	0.008
Staff member	0.018	0.025
Number of cigarettes	0.004**	0.002
Constant	0.522***	0.040
PWF parameter (φ)		
Age	<0.001	0.003
White	0.094*	0.050
Male	0.137***	0.044
Financial situation	0.006	0.020
Staff member	0.021	0.073
Number of cigarettes	-0.003	0.004
Constant	0.611***	0.086
PWF parameter (η)		
Age	0.003	0.005
White	-0.004	0.067
Male	-0.086	0.059
Financial situation	0.066**	0.033
Staff member	-0.119	0.093
Number of cigarettes	0.004	0.005
Constant	0.830***	0.126
Discounting parameter (β)		
Age	-0.000	<0.001
White	0.014**	0.006
Male	-0.002	0.004
Financial situation	0.002	0.003
Staff member	0.003	0.006
Number of cigarettes	-0.000	<0.001
Constant	0.967***	0.011
Discounting parameter (δ)		
Age	-0.001	0.005
White	-0.271***	0.101
Male	0.027	0.097
Financial situation	-0.205***	0.059
Staff member	-0.168	0.135
Number of cigarettes	0.040***	0.012
Constant	1.535***	0.221

Table D3 (Continued)

	Model	
	Estimate	Std error
Intertemporal risk preference parameter (ρ)		
Age	-0.004	0.029
White	-0.740	0.550
Male	0.914**	0.403
Financial situation	-0.572**	0.273
Staff member	1.013	0.678
Number of cigarettes	-0.013	0.041
Constant	0.326	0.782
Error terms		
μ	0.141***	0.005
υ	0.751***	0.128
ψ	0.280***	0.024
N	47310	
log-likelihood	-26883.870	

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

Table D4

Intertemporal Risk Preference ML Estimates

RDU, Quasi-Hyperbolic Discounting

Smoking Severity: Fagerström Test for Cigarette Dependence

	Model 1		**Model 2**	
	Male		Female	
	Estimate	Std error	Estimate	Std error
Atemporal risk preference parameter (*r*)				
Age	0.002	0.004	-0.003**	0.002
White	-0.030	0.039	0.068	0.059
Financial situation	-0.013	0.016	-0.006	0.014
Staff member	0.018	0.051	0.054	0.042
FTCD score	0.005	0.007	<0.001	0.008
Constant	0.521***	0.106	0.570***	0.076
PWF parameter (φ)				
Age	0.014	0.015	-0.003	0.005
White	0.198	0.147	0.057	0.083
Financial situation	-0.134**	0.057	0.025	0.043
Staff member	-0.153	0.191	-0.032	0.116
FTCD score	-0.049**	0.022	-0.007	0.014
Constant	0.871***	0.298	0.806***	0.145
PWF parameter (η)				
Age	-0.017	0.028	0.001	0.010
White	-0.143	0.133	-0.077	0.192
Financial situation	0.118*	0.069	0.135**	0.062
Staff member	-0.087	0.222	-0.096	0.197
FTCD score	0.064**	0.027	0.003	0.039
Constant	0.918	0.574	0.742***	0.258
Discounting parameter (β)				
Age	<0.001	0.001	<0.001	<0.001
White	0.006	0.012	0.013	0.011
Financial situation	-0.004	0.007	-0.004	0.004
Staff member	-0.001	0.014	<0.001	0.015
FTCD score	<0.001	0.002	0.001	0.001
Constant	0.970***	0.018	0.979***	0.014
Discounting parameter (δ)				
Age	0.035	0.039	-0.002	0.011
White	-0.334	0.283	0.130	0.374
Financial situation	-0.192	0.167	-0.098	0.144
Staff member	-0.109	0.449	-0.515*	0.289
FTCD score	0.078	0.065	-0.006	0.044
Constant	0.832	0.949	1.535***	0.456

Table D4 (Continued)

	Model 1		Model 2	
	Male		Female	
	Estimate	Std error	Estimate	Std error
Intertemporal risk preference parameter (ρ)				
Age	0.008	0.077	0.022	0.083
White	0.259	0.930	1.232	1.538
Financial situation	-0.994**	0.497	-0.656	0.555
Staff member	0.809	0.907	0.463	1.876
FTCD score	-0.435**	0.213	0.409	0.403
Constant	3.040**	1.493	-1.177	2.876
Error terms				
μ	0.131***	0.011	0.146***	0.014
υ	0.606**	0.252	0.439**	0.178
ψ	0.224***	0.047	0.221***	0.063
N	9120		8170	
log-likelihood	-4797.759		-4401.497	

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

Table D5

Atemporal Risk Preference ML Estimates

Rank-Dependent Utility Theory

Smoking Severity: Fagerström Test for Cigarette Dependence

	Model 1		**Model 2**	
	Male		Female	
	Estimate	Std Error	Estimate	Std Error
Power function parameter (*r*)				
Age	0.046	0.092	-0.017	0.014
White	0.069	0.196	0.130	0.120
Financial situation	-0.067	0.099	-0.024	0.052
Staff member	-0.237	0.580	0.456	0.468
FTCD score	-0.069*	0.040	0.033	0.030
Constant	-0.056	1.743	0.824**	0.348
PWF parameter (φ)				
Age	0.028	0.035	0.010*	0.005
White	0.127	0.122	0.115	0.087
Financial situation	-0.098*	0.052	0.019	0.044
Staff member	-0.267	0.223	-0.202	0.153
FTCD score	-0.038*	0.021	-0.016	0.017
Constant	0.343	0.760	0.390**	0.174
PWF parameter (η)				
Age	0.050	0.208	-0.001	0.014
White	-0.031	0.256	0.053	0.191
Financial situation	0.063	0.108	0.116*	0.062
Staff member	-0.475	1.229	0.107	0.306
FTCD score	0.003	0.028	0.032	0.057
Constant	-0.262	4.244	0.678*	0.369
Error (μ)				
Constant	0.132***	0.020	0.149***	0.013
N	4320		3870	
log-likelihood	-2687.46		-2440.704	

Results account for clustering at the individual level

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

APPENDIX E
[ONLINE WORKING PAPER]

The tables in this appendix complement those presented in Appendix D because we estimate the SDU model (9) in the main text jointly with an EUT model and an exponential discounting function. We allow the parameters of the SDU model to vary as a linear function of demographics, socio-economic characteristics, and three measures of smoking behaviour: smoking status (Table E1); smoking intensity, measured by the number of cigarettes smoked per day (Table E2); and smoking severity, measured by smokers' scores on the FTCD (Table E3). These latter two tables are split according to gender given the historical differences in smoking prevalence between men and women, and the economically and statistically significant differences in their intertemporal risk preferences.

Table E1
Intertemporal Risk Preference ML Estimates
EUT, Exponential Discounting
Heterogenous Preferences

	Model	
	Estimate	Std error
Atemporal risk preference parameter (*r*)		
Age	-0.002	0.001
White	0.010	0.018
Male	0.019	0.015
Financial situation	0.002	0.008
Staff member	0.024	0.025
Ex-smoker	-0.005	0.029
Smoker	0.033*	0.018
Constant	0.429***	0.039
Discounting parameter (δ)		
Age	<0.001	0.004
White	-0.250***	0.090
Male	0.033	0.083
Financial situation	-0.160***	0.050
Staff member	-0.140	0.118
Ex-smoker	0.097	0.130
Smoker	0.318***	0.107
Constant	1.268***	0.197

Table E1 (Continued)

	Model	
	Estimate	Std error
Intertemporal risk preference parameter (ρ)		
Age	0.006	0.033
White	-0.915	0.670
Male	0.918*	0.487
Financial situation	-0.712**	0.332
Staff member	1.207	0.853
Ex-smoker	-2.274	1.528
Smoker	-0.275	0.525
Constant	0.213	1.106
Error terms		
μ	0.165***	0.007
υ	0.311***	0.056
ψ	0.306***	0.023
N	47310	
log-likelihood	-27526.615	

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

Table E2

Intertemporal Risk Preference ML Estimates

EUT, Exponential Discounting

Smoking Intensity: Number of Cigarettes Smoked per Day

	Model 1		**Model 2**	
	Male		Female	
	Estimate	Std error	Estimate	Std error
Atemporal risk preference parameter (*r*)				
Age	-0.002	0.002	-0.001	0.001
White	-0.020	0.027	0.015	0.023
Financial situation	-0.005	0.011	0.009	0.010
Staff member	0.038	0.046	0.017	0.028
Number of cigarettes	0.005**	0.002	0.005*	0.002
Constant	0.540***	0.071	0.345***	0.043
Discounting parameter (δ)				
Age	0.003	0.010	-0.001	0.004
White	-0.417**	0.179	-0.170*	0.094
Financial situation	-0.242***	0.091	-0.114**	0.058
Staff member	-0.154	0.254	-0.116	0.122
Number of cigarettes	0.049**	0.019	0.026***	0.010
Constant	1.702***	0.368	1.020***	0.222
Intertemporal risk preference parameter (ρ)				
Age	-0.034	0.033	0.037	0.059
White	-0.235	0.378	-23.331***	3.991
Financial situation	-0.228	0.203	-1.207*	0.706
Staff member	0.911	0.681	1.443	1.652
Number of cigarettes	-0.053*	0.029	0.115	0.078
Constant	1.344	0.921	-1.014	1.605
Error terms				
μ	0.155***	0.009	0.176***	0.011
υ	0.523***	0.159	0.191***	0.039
ψ	0.233***	0.023	0.351***	0.030
N	20900		26410	
log-likelihood	-11933.148		-15477.688	

Results account for clustering at the individual level

* p<0.10, ** p<0.05, *** p<0.01

Table E3

Intertemporal Risk Preference ML Estimates

EUT, Exponential Discounting

Smoking Severity: Fagerström Test for Cigarette Dependence

	Model 1		**Model 2**	
	Male		Female	
	Estimate	Std error	Estimate	Std error
Atemporal risk preference parameter (*r*)				
Age	0.005	0.005	-0.004**	0.002
White	-0.015	0.037	0.067	0.059
Financial situation	-0.029**	0.012	-0.012	0.014
Staff member	0.011	0.051	0.045	0.042
FTCD score	0.001	0.007	0.001	0.008
Constant	0.449***	0.118	0.483***	0.078
Discounting parameter (δ)				
Age	0.031	0.040	-0.004	0.008
White	-0.340	0.255	0.083	0.273
Financial situation	-0.174	0.140	-0.066	0.090
Staff member	-0.022	0.455	-0.344*	0.204
FTCD score	0.054	0.056	-0.007	0.034
Constant	0.893	0.931	1.287***	0.379
Intertemporal risk preference parameter (ρ)				
Age	-0.039	0.097	0.048	0.133
White	-0.411	1.274	1.407	1.910
Financial situation	-0.901**	0.391	-0.623	0.772
Staff member	0.942	1.217	0.120	2.750
FTCD score	-0.376*	0.194	0.608	0.598
Constant	3.490*	1.948	-2.856	4.622
Error terms				
μ	0.169***	0.014	0.165***	0.014
υ	0.328**	0.147	0.168***	0.051
ψ	0.254***	0.050	0.250***	0.068
N	9120		8170	
log-likelihood	-5006.841		-4513.919	

Results account for clustering at the individual level

* p<0.10, ** p<0.05, *** p<0.01