Intrapersonal Externalities: When Decisions Help or Hinder your Future Self

David J. Stillwell, MSc.

Thesis submitted to The University of Nottingham

for the degree of Doctor of Philosophy

July 2013

"Underlying the choice to smoke…"

Not a good word to describe why people smoke.

You are writing for a journal called Addiction.

Correction from Senior Editor

Abstract

INTRAPERSONAL EXTERNALITIES are outcomes from a decision which affect the payoffs from future decisions. Positive intrapersonal externalities represent self-investment in one's future self, such as learning a musical instrument or exercising. Negative intrapersonal externalities represent a disinvestment in one's future self, such as eating sugary food or smoking a cigarette. An individual who succumbs to the lure of an immediate payoff in lieu of future potential earnings is considered to be choosing myopically, and it is hypothesised that addictions are one such situation. This failure of the decision-making system to optimise its choices contrasts with delay discounting, in which an individual is assumed to choose rationally but to discount the future, and impulsive disinhibition, in which an individual is assumed to be unable to control their actions even while they state a preference for an action they do not take. This thesis starts by measuring the relationship between delay discounting and impulsive disinhibition with a range of addictive behaviours in a large online study, and finds that while they are consistent predictors there is further variance to explain. It then examines the validity of the intrapersonal externalities model, and finds that it adds independent variance beyond delay discounting and impulsivity in predicting smoking behaviour; there is no evidence that performance in intrapersonal externalities tasks is related to trait impulsivity. It also uses intrapersonal externalities in the laboratory to study advice seeking and taking in an impulsive decision-making context, which would be incompatible with delay discounting or impulsive disinhibition theory. This thesis concludes that the intrapersonal externalities model has been shown to be a viable third model to understand addictive decisions, and suggests that it could be extended to study social addictions.

Acknowledgements

With thanks to my supervisor, Richard Tunney, for his intellectual discussions, for helping me to navigate the arcane world of academia, and for his trust and patience when my research interests strayed to the unconventional.

Thanks to Miles Cox and Peter Chapman, my viva examiners, for their close reading of this thesis. After 3 years of work, I appreciated the time you took to give me a rigorous and engaging examination.

Thanks to Facebook, for releasing the apps platform and suddenly making hundreds of millions of people around the world potential participants in my experiments.

Thanks to my parents for organizing and supporting me, and for being proud of what I do.

Thanks to my online friends for their conversation and jokes when I worked late at night, and to my offline friends for not letting me work too hard during the day.

This work was supported by the Economic and Social Research Council [grant number ES/F021801/1].

Contents

CHAPTER 1

General Introduction

A young daughter has to be alternately encouraged and threatened by her parents before she will practice her piano. Reading the notes and moving her hands fast enough to play a melody is difficult, and her clumsy attempts at playing Chopsticks are unrewarding. She is often lured by the temptation of hanging out with her friends, or watching a television show. Fast forward ten years, and the daughter is on the cusp of adulthood; she plays the piano for her local church – a new song nearly every Sunday, and because of her sightreading skill she only has to practice the song once to get the hang of it – and she studies Music at school, where she finds it easy to dissect songs in order to explain how they are so emotionally powerful.

Some actions have consequences that outlast the initial decision and change the potential outcomes of future decisions. For the daughter, practicing a musical instrument is difficult and unrewarding initially but it becomes more enjoyable as her skill improves and she can sightread new songs. Crucially, it benefits her even when she is not playing the instrument: when she listens to music she not only hears the sounds, she also appreciates the performer's finesse and the skill gone into shaping the music, which makes the experience all the more enjoyable. Differences in the adult brain after musical training in childhood are detectable neurologically (Skoe & Kraus, 2012).

The negative situation is more well studied – where actions have initially positive consequences but in the longer term they lead to a suboptimal result. For example, eating sugary food may feel good immediately afterwards but over the long term eating too much will lead to obesity, which will negatively affect the individual's quality of life in all areas – even when lying in bed. The same situation occurs in all addictions, whether 'hard' addictions like substance abuse, or 'soft' (behavioural) addictions like gambling (Potenza, 2006).

I identify two major classes of models that attempt to explain why individuals choose to consume addictive substances. Psychopharmacological models emphasise individual differences in the cognitive system that make one more or less susceptible to impulsivity, such that an individual cannot avoid actions that they know are bad for them. Normative economic models emphasise individual differences in the degree to which one discounts the future, such that an individual subjectively undervalues future negative consequences because of the delay before they occur. The implications of the psychopharmacological model are that individuals do not have self-control and so while they realise that consuming a substance will negatively affect their happiness, they cannot stop themselves. In contrast, the normative economic model states that individuals are rational when they make a decision - they just do not care enough about the future. Both of these models have drawbacks. If individuals are slaves to their impulses, then why do some individuals manage to stop their addictive consumption? And why do impulsive people not become addicted to everything? If individuals are rationally addicted, then why do many addicts claim to want to stop?

The central claim that this thesis examines is whether there is evidence for a third class of model originating in the cognitive decision-making and behavioural economic literatures. Essentially it posits that the decision-making system has difficulty learning the mix of immediate and future consequences that stem from an action, and this leads people to not taking into account the long-term consequences of some decisions. These long-term consequences are called *intrapersonal externalities*. In economics, externalities refer to consequences of an individual's decisions that affect others' wellbeing (e.g. pollution). A selfish economic actor will not take externalities into account when making a decision. In the same vein, interpersonal externalities refer to consequences that change the value of options available to one's future self, such that an individual ignores their future self's potential choice payoffs when making a decision.

According to the model, my happiness today is based on my history of previous decisions, including the time 3 months ago when I decided to go for a run rather than stay in bed. That choice made me slightly more healthy today than I would have been otherwise, which ever so slightly increases my potential subjective utility from both staying in bed and going for a run next time I wake up. This set of consequences is hard to learn, especially as one of the decisions (staying in bed 3 months ago) was not made, and so individuals concentrate on the immediate consequences of actions rather than trying to decipher the long term ones.

No one can wake up one day and decide to be a concert pianist, and no one decides to become an alcoholic. That is not to say that people do not wake up and decide to be a concert pianist, or even to succumb to the lure of alcoholism, but ultimately these decisions have to be followed up by repeated actions over a long period. Such grand choices are commitments that can be broken, as so many New Years resolutions are every year, rather than decisions. The intrapersonal externalities model recognises this, and so the long-term eventualities are the result of a set of individual choices. Both eventualities occur after repeated choices to practice the piano, or have another drink, over many years. The unit of analysis is therefore the individual decision. Positive intrapersonal externalities represent a self-investment in one's future, whereas negative intrapersonal externalities represent a disinvestment.

The DSM-IV-TR (APA, 2000) includes three criteria in its definition of substance dependence which mirror some of the decision-making processes that occur in situations with conflicting long and short-term incentives. First, an individual has increased tolerance of the substance, so that they need a markedly increased amount of the substance to achieve the same effect as previously. In decision-making terms, this indicates that the history of drug-taking behaviour has an effect on the utility of current drug-taking behaviour. Second, an individual desires to cut down on use of the substance. In decision-making terms, this could indicate that the individual is aware of the long-term consequences of consuming the substance, even while they succumb to its short-term lure. Third, the individual continues using the substance despite knowledge that it causes or exacerbates physical or psychological problems. In decision-making terms, this again represents an understanding of the long-term negative consequences of consuming the substance over a long period. The DSM-V is expected to contain a new category of "behavioural addictions", which includes behaviours such as pathological gambling which do not involve consumptino of exogenous pharmacological substances. Nevertheless, although parts of this thesis concentrate on addictive substances, mostly nicotine smokers, both these and behavioural addictions are simply an important example of a situation in which a choice has both long and short-term consequences which conflict. Actually these conflicting choices describe many situations which are understood to lead to sub-optimal outcomes; individuals over-indulge in everything from sweet foods to Internet use, and under-invest in everything from exercise to studying. Addiction is, in the context of this thesis, an extreme example of a situation where long and short-term incentives are not in alignment.

This general introduction starts by (1.1) relating the origin of the normative economic model, which in the research field of choice over time is reflected in delay reward discounting. This is followed by (1.2) some extensions to the decision-making model from behavioural economists, which use psychological insights to take account of situations where behaviour does not match the normative accounts. In the area of choice over time this leads to a description of intrapersonal externalities as another situation where individuals do not conform to normative expectations. Then it moves away from economic decision-making and (1.3) describes some of the psychopharmacological accounts, which have a separate origin. Finally, (1.4) the ways that intrapersonal externalities could theoretically overlap or integrate with the other models are explored.

1.1 The Normative Decision-Making Model: Delay Discounting as a Rational Choice

1.1.1 Rational Choice and Expected Utility Theory

Expected utility theory (EUT) was first proposed by Bernoulli (1738, cited in Starmer, 2000) but was not taken up in the literature until von Neumann and Morgenstern (1947, cited in Starmer, 2000) showed that it could be derived from axioms that seemed logical, and could be applied to strategic situations in which individuals had to make decisions; this began the field of game theory. EUT logically describes what an individual should do if they are rational. Each option in a decision has outcomes and respective probabilities that each outcome will occur, thus in EUT the expected utility of any option is the probability-weighted sum of utilities of consequences. EUT takes account of an agent's attitude to risk in the shape of the utility function, which can be linear if the agent is risk neutral, concave if the agent is risk averse or convex if the agent is risk loving. The rational agent should act to maximise expected utility (Schoemaker, 1982).

1.1.2 Delay Discounting

In the area of choice over time, delay discounting has been used to examine individual decisions under the assumption of rational choice (Becker & Murphy, 1988). In order to account for individuals' preferences for immediate rewards over future ones, a discounting function is applied based on the time until the reward is received. Delay discounting has been studied both from a psychological and an economic perspective (Rachlin, Raineri, & Cross, 1991). This is because it has the advantage that it presents a logical framework that allows it to fit into mathematical models of decision making, and so has been used to make policy recommendations (Monterosso & Ainslie, 2007).

Delay discounting is typically researched by giving people multiple forced choices between an immediate reward and a larger delayed reward. Afterwards, an indifference point is calculated where the individual subjectively values both the immediate and delayed options the same, and after calculating several of these for a range of delay periods it is possible to calculate a discount curve. The shape of the discounting curve was debated for some time, with the two candidate curves being the exponential curve and the hyperbolic curve. These imply different choice dynamics in intertemporal choice situations (Ainslie, 1975). The exponential curve assumes that as the delay increases by equal increments the proportion of the reward that is discounted is constant (e.g. 10% per day).

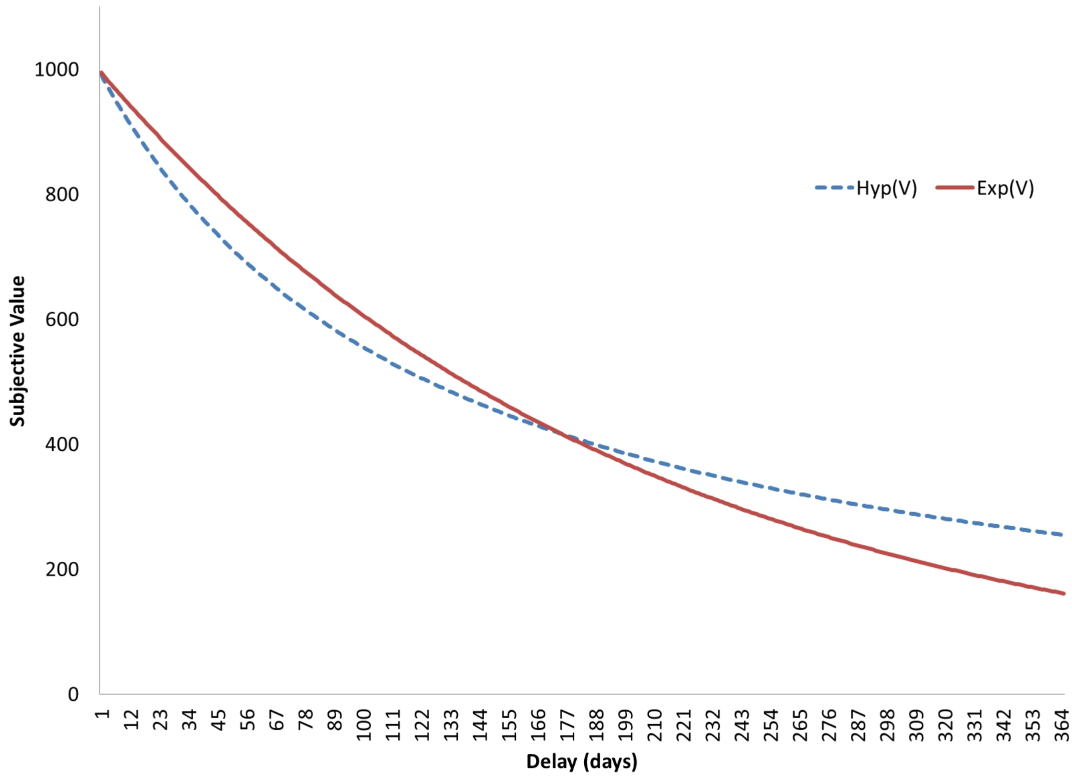

Figure 1.1: Exponential and hyperbolic delay reward discounting curves

Therefore, if two options are both discounted according to the exponential function then preference between them will be constant as a function of time; this is time consistency. In contrast, the hyperbolic curve assumes that the proportion of a reward that is discounted changes as a function of time, so that an equal increment in delay will produce a larger decrement in reward value at short delays than long ones (Vuchinich & Simpson, 1998). The two curves are plotted in Figure 1.1, where the lines crossing indicates that hyperbolic discounting leads to steeper immediate discounting than exponential discounting, but that the rate of discounting decreases over the longer term. This leads to the situation where an individual might say that they prefer a course of action in the future that would be advantageous in the long term (such as refraining from drinking), however as the time draws nearer they switch to the short-term preference

(Kirby & Herrnstein, 1995). Empirical research generally favours the hyperbolic function (Bickel, Odum, & Madden, 1999; Rachlin et al., 1991; Rachlin, Siegel, & Cross, 1994), implying that individuals have time inconsistent preferences (Rachlin, 2000). This distinction is explained in full detail in Chapter 2, where the two models of discounting are empirically compared.

The validity of delay discounting, as an individual difference that is associated with addiction, is evidenced by the finding that individuals with a range of different addictions discount future rewards more than control subjects do (Bickel et al., 1999; Kirby & Petry, 2004; Madden, Petry, Badger, & Bickel, 1997; Vuchinich & Simpson, 1998). Additionally, there is evidence that people change their addictive substance consumption habits in advance of an announced price rise (Gruber & Koszegi, 2001). This behaviour is surprisingly forward thinking, because individuals must recognise that in the future they will choose myopically, and therefore it is inconsistent with the view that addicts choose compulsively.

1.2 Intrapersonal Externalities: A Psychologically-Informed Model of Choice Over Time

1.2.1 Systematic Deviations from Rational Choice

Historically, economics has argued that its assumptions about how individuals make decisions are inconsequential, as long as the theories explain the aggregate data (Lewin, 1996). However, experimental research has noted many systematic deviations from EUT that are troubling. Two such situations were created by Allais (1953). In the common consequence effect, two decisions are made between pairs of gambles. The second pair of gambles is derived from the first pair, except that a common consequence of both pairs is removed. If individuals conform to EUT, then if they choose the first choice from the first pair, they should also choose the first choice from the second pair, since the second pair is the same as the first except that both choices in it have a common consequence removed. However, it has been found that individuals instead systematically switch from one choice to the other between pairs (Rubinstein, 1988). A similar situation is the common ratio effect, in which again two decisions are made between pairs of gambles. The second pair of gambles is again derived from the first pair, except that a common ratio of their probabilities is removed from both gambles. Given that both pairs of gambles are the same except that the second pair of gambles has lower probabilities, rational agents should choose the same gamble in each pair in both decisions – but experimental research finds that humans switch systematically. Recent evidence however seems to show that these effects are not stable over time. de Kuilen and Wakker (2006) set up a repeated common consequence effect and found that performance converged towards the expected utility level when participants had experience of the payoffs and even when they repeatedly played the game as a thought experiment. This suggests that the irrational decisions noted by Allais (1953) can be overcome and that the cognitive decision-making mechanism is not inherently irrational.

Preference reversals are another systematic deviation from rational choice that have been extensively studied (Lichtenstein & Slovic, 1971; Lindman, 1971). In a typical task, participants are presented with a choice between two gambles, one with a high probability of winning a small reward and the other with a low probability of winning a large reward. Participants are then given a distraction task, before being asked to put a dollar value on each of the two previous gambles. It has been found that participants typically put a higher dollar value on the gamble with a large reward, but choose the gamble with the small reward in the first instance (Tversky, Slovic, & Kahneman, 1990). Given that these two methods are just different ways of asking the same question, its findings are puzzling for EUT. However, preference reversals can be reduced when gambles are presented as relative frequencies (Tunney, 2006). This suggests that the difficulty that human participants have is in utilising the Bayesian probability formats that are usually used in such research (Gigerenzer & Hoffrage, 1995).

Another deviation from EUT was reported by Tversky and Kahneman (1981), who described the phenomenon of framing, whereby decisions are affected by the context in which they are presented. Specifically, participants tend to be loss aversive. When presented with two pairs of stochastically identical gambles, in which the outcomes of the first gambles are presented as gains and the outcomes of the second gambles are presented as losses, participants systematically switch choices between the two pairs. This suggests that people treat gain probabilities and loss probabilities differently from one another.

Another problem was reported by Rabin (2000), who noted that given someone who rejects a gamble with a 50% likelihood of gaining $11 and a 50% likelihood of losing $10, EUT predicts an implausibly high level of risk aversion if the gamble is scaled up to a potential loss of $1000, where participants should not accept the gamble even if the potential reward is $100,000. Similarly, Samuelson (1963, cited in Rabin, 2000) showed that EUT predicts that if people will turn down a gamble where they have a 50% chance of losing $100 and a 50% chance of gaining $200, they will also turn down 100 of these gambles strung together. This suggests two things; first, it is further evidence that people treat potential gains and losses differently from one another, and second, that there must be some mechanism whereby individuals group gambles together.

EUT was not developed based on psychological plausibility. There was no direct evidence that humans used probabilities and valuations at all when making decisions, so the utility function may have just been a theoretical concept without basis in the brain. However, research has found neurons in the monkey cortex that seem to fire at the rate at which probabilities and valuations are presented, suggesting a decision-making system very similar to that presented by EUT (Platt & Glimcher, 1999). This does not of course mean that EUT is wholly psychologically plausible, but it does provide evidence that at least the brain uses the general concepts of valuation and probability are used in making decisions.

1.2.2 Prospect Theory and Social Preferences: Extensions to EUT

Prospect theory (Kahneman & Tversky, 1979) and later Cumulative Prospect theory (Tversky & Kahneman, 1992) aimed to overcome the weaknesses of EUT by introducing two important alterations. First, outcomes are weighted by way of an inverse 'S' function which centres around a reference point (see Figure 1.2). The reference point can be considered as the position that the agent was in before making the decision, or could also include some aspect of the agent's expectations. That is, in the terms of the laboratory experiments where most research is done, if the agent expects to be paid a certain amount for their participation, then they might be disappointed to be paid less and so consider that to be a loss (Heath, Huddart, & Lang, 1999).

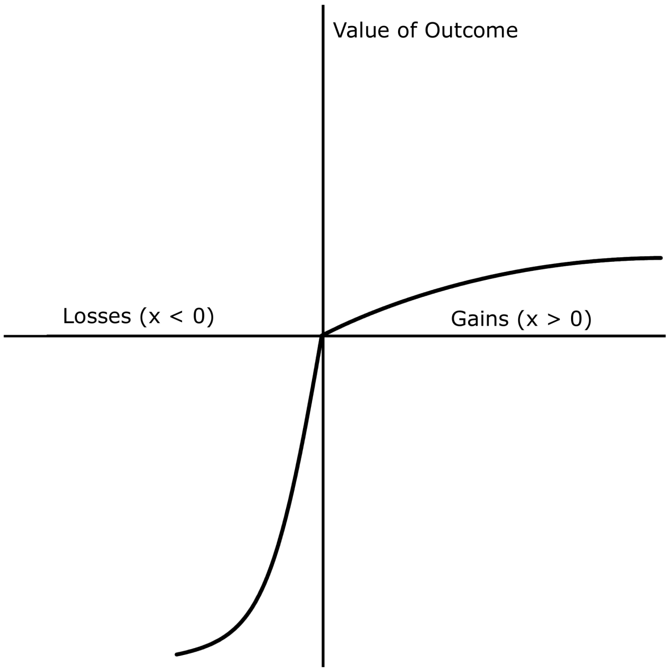

Figure 1.2: How gains and losses are weighted in Cumulative Prospect theory. From Starmer (2000).

The reference point addresses the finding from framing research that gains are not treated the same as losses. Since people are loss aversive, the losses portion of the function is steeper than the gains portion, which has the effect that losses are given a higher weighting than equivalent gains. The curve in both portions of the function addresses the finding that people do not treat twice the amount of a gain as giving twice as much utility. The power law function was chosen because it fits the data, but it was also noted that this is psychologically plausible, since it reflects the principle of diminishing sensitivity, a general principle of perception (Camerer & Ho, 1994).

The second alteration to EUT is similar, but is instead applied to probabilities. Probabilities are weighted according to an inverse 'S' function shown in Figure 1.3 (Lattimore, Baker, & Witte, 1992). Therefore, probabilities that have a low chance of occurring are over-weighted, and those that are almost certain to happen are under-weighted. This addresses Rabin's (2000) criticism of EUT whereby risk aversion at relatively low stakes will imply an implausible degree of risk aversion at high stakes, although Rabin's criticism came much later, and probability weighting was included for other reasons detailed below.

These two alterations together also address the phenomena of the common consequence and common ratio effects. In the common consequence effect, removing a common outcome from a pair of gambles will make a difference to how the utility of the gambles are conceived since they will be weighted differently. Small outcomes will be over-weighted, and large ones will be under-weighted. Similarly, in the common ratio effect, removing the common probability of an outcome occurring in a pair of gambles will make a difference to how the utility

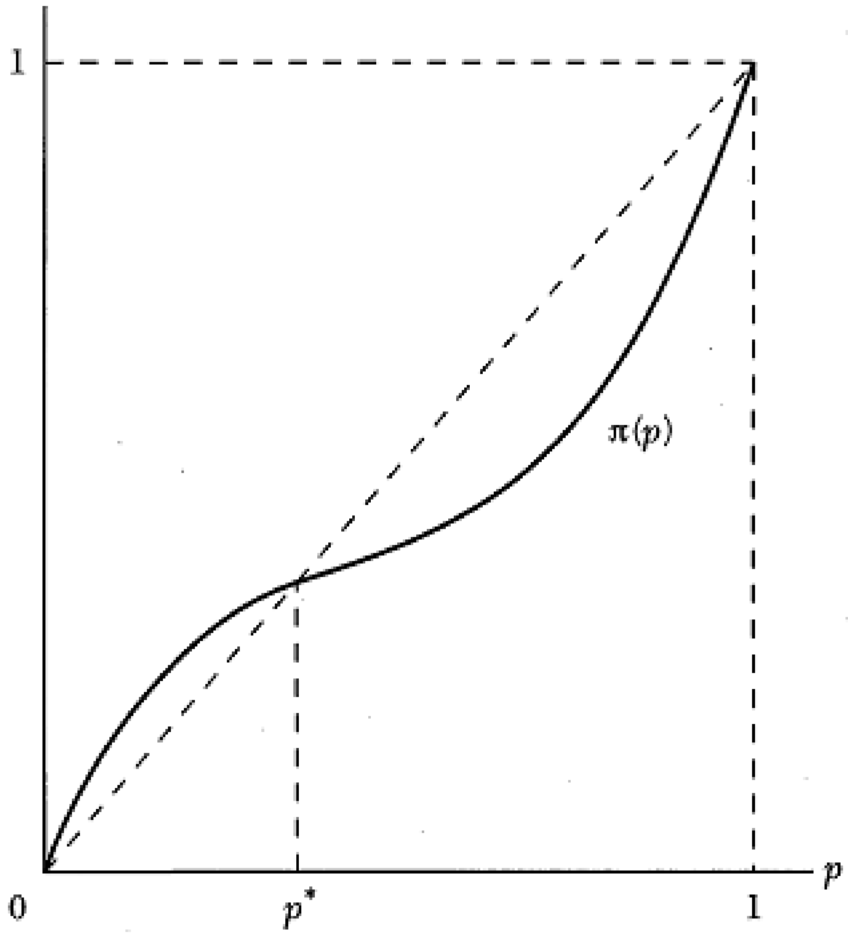

Figure 1.3: Probability weighting in Cumulative Prospect theory. From Starmer (2000).

of the gambles are conceived, since they will be weighted differently. Gambles with a low probability will be over-weighted, and gambles with a high probability will be over-weighted.

It is significant that these alterations to EUT were made primarily in order to address findings such as the common consequence and common ratio effects. However, even though the evidence for these has been weakened and they may no longer need addressing, the two alterations are still on solid ground. This may be due to their firm psychological foundation. The probability weighting function and gains-losses function were modeled by Stewart, Chater, and

Brown (2006) in their Decision by Sampling theory. The authors started from a psychological perspective by assuming that people can only make decisions by comparing a current option's attributes (probability and payoff) with other attributes stored in memory. By ranking the attribute involved in the current decision with others experienced in the past, it is possible to weight the attribute accordingly. This requires only that humans can sample the environment and store it in memory, make binary comparisons, and have a concept of frequency. The authors analysed the frequency of credit and debit transactions to UK bank accounts and found that both roughly followed a power law, whereby there were many more small than large transactions. This produces a curve whereby the subjective value of a gain/loss will not necessarily double when the amount of the gain/loss doubles. They also found that the curve for losses was steeper than that for gains, representing the fact that there were more small losses than small gains. This produces two different curves around the reference point, as in prospect theory, with people acting loss-aversive. In practice this means that a person's utility function is relative to their usual transactions, so millionaires would consider some amounts to be small that others would not. This seems intuitively plausible.

A similar analysis was carried out but with probabilities, to argue that small probabilities are over-weighted and larger ones are under-weighted. The authors analysed the frequency with which various probability-related words (e.g. "unlikely" or "almost impossible") are used in language, since probabilities and probability-related words are both used well by humans (Wallsten & Budescu, 1995). They asked participants to relate the words to probabilities, and then compared the words' frequencies in the British National Corpus, a collection of the frequencies of words in the English language designed to be representative of the language as a whole, including both spoken and written English. They found that words which denote small probabilities are over-represented compared to words which denote large probabilities.

Beyond individual decision making, a wealth of evidence from social strategic experiments has found that people do not conform to selfish "rational" expectations in social situations. In the well-known ultimatum game (Guth, Schmittberger, & Schwarze, 1982), player one is presented with a sum of money and then asked to offer a share of it to player two. The second player then decides whether to accept or reject the offer. If it is accepted, both players take away the sum of money that was agreed, but if the offer is rejected then both players get nothing. The rational theory predicts that player one will offer a very small amount of money to player two, and that player two will accept it since even a very small amount of money is better than nothing. However, it is found that on average about 40-50% of the total amount is offered. Furthermore, in the dictator game, the second player is forced to accept the offer, so player one does not have to worry about the offer being rejected, and so has an even bigger incentive to offer very little money. However, again, it is found that players routinely offer around 50% of the total money.

One model was developed to explain this finding, inequality aversion (Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999). In the model, unequal outcomes (especially unequal against the individual) alter the payoff schedules such that the ratio of money for each participant in an experiment becomes important, rather than simply the amount of money received by the individual.

An elegant finding contradicted the theory that it was inequality that particularly galled the participants in ultimatum games. Falk, Fehr, and Fischbacher (2003) carried out a series of ultimatum games that showed that intentions are perhaps the most important aspect. Four different games were used, in all of which player one had the option of two different splits of the money to offer. In all four games, player one had the option of an 8/2 split (in favour of player one), but the other option was either 5/5, 2/8, 8/2 (so there was effectively no option), or 10/0. When player two found that they had been offered 2, when they could have been offered 5, they were much more likely to reject than when they were offered 2 and they knew that the alternative was that they would be offered nothing.

A theory of reciprocity (Falk & Fischbacher, 2006) is intended to take this finding into account, which as in the inequality aversion model modifies an individual's payoff schedules based on their desire to respond to kindness with kindness and to revenge behaviours that they deem intentionally unfair. One potential theoretical stumbling block for this theory is that when strategic games are played with multiple people, the theory assumes that a kindness value is generated for each other individual in the group. It may be that in games with a lot of people, individuals are not considered individually but rather the whole group is evaluated as if it was one individual. It would probably be difficult to ascertain in which situations individuals are considered alone and when they are considered as part of a group.

Like Prospect theory, these social experiments which demonstrate diversions from EUT seem to be due to limitations in the cognitive architecture that makes decisions, and so they have been solved by augmenting EUT to take these psychological factors into account. This is a recurring theme in decision making research and one which this thesis adds to.

1.2.3 Addiction as a Cognitive Bias: Intrapersonal Externalities and the Harvard Game

Addiction presents a relatively widespread situation where choices do not maximize long-term utility (Gruber & Koszegi, 2001). Smokers often tell others about their quitting intention in order to make it socially embarrassing to smoke (Prochaska, Crimi, Lapsanski, Martel, & Reid, 1982), they sometimes set up aversive consequences for smoking such as tearing up dollar bills for each cigarette smoked above a daily maximum, and youth smokers overestimate their chances of quitting in five years time (Gruber & Koszegi, 2001). Thirty-eight million of the 46 million smokers in America in 1993 had tried to quit, most several times (Gruber & Koszegi, 2001).

Choice over time can be studied from a cognitive decision making perspective using intrapersonal externalities. Models of intrapersonal externalities postulate that when an individual makes a decision, they do not fully take into account its effects on the utility of options available to the future self, and so in effect these are a kind of externality applied to the future self (Read, 2001). By ignoring the consequences of a choice on the payoffs available from future choices, an individual is led down the Primrose Path to addiction (A. Goldstein & Kalant, 1990). At each choice point they choose the option with the higher immediate utility, yet they arrive at a suboptimal situation.

Intrapersonal externalities have been operationalized in the Harvard Game (R. J. Herrnstein, Loewenstein, Prelec, & Vaughan, 1993). In this cognitive task, participants are asked to make repeated choices between two options. One option (Meliorate) gives an immediately higher payoff than the other option, but reduces the payoffs from both options over a set number of future trials. Choosing the other option (Maximize) leads to a consistently average payoff which ultimately leads to a higher overall payoff. The Harvard Game therefore preserves the tension in addictive decisions between an immediately higher payoff versus the long-term negative consequences that that payoff leads to. Unlike delay discounting however, the consequences of making the suboptimal decision in the Harvard Game do not have to be delayed for any significant length of time. The Harvard Game is described operationally in full detail in Chapter 3, where it is used.

In this respect intrapersonal externalities have already been used to make policy recommendations (Gruber & Koszegi, 2001; Laux, 2000). Intrapersonal externalities have a similar advantage to delay discounting, whereby they extend current decision making mathematical models in order to account for apparently irrational long-term choices.

The Harvard Game is unique in the area of choice over time, as it requires participants to learn the payoff schedules through experience rather than just being explicitly told their options. This is psychologically plausible (Koritzky & Yechiam, 2010), as one cannot ever be told the precise subjective payoffs from drinking or not drinking, but must learn them from experience. It also means that researchers can study in the laboratory how individuals learn to make self-controlled or impulsive decisions. No one chooses to become an addict (or to become a concert pianist). Each choice (to have another drink, or to practice the piano) is made in isolation, but they add up to a payoff that is not the sum of its parts.

Studying learning through experience has become a niche research area within the decision-making literature, ever since it was found that individuals do not make the same decisions once they have experienced the payoffs, compared to being told their options (Rakow & Newell, 2010). As mentioned, individuals seem to be risk-seeking when presented with small probabilities of large gains but risk-averse when presented with small probabilities of large losses; however, when individuals experience repeated small probabilities of large gains they are risk-averse, and when they experience repeated small probabilities of large losses they are risk-seeking (Barron & Erev, 2003).

1.3 Psychopharmacological Models of Addiction

There is a cluster of theories of myopic choice which are predominantly discussed in the psychopharmacological literature. They are sometimes described collectively as impulsivity, with each theory a different dimension of the same construct (de Wit, 2009; Meda et al., 2009).

1.3.1 Impulsive Disinhibition

Impulsivity encompasses a variety of definitions, including a tendency to act without forethought, insensitivity to consequences, and an inability to inhibit inappropriate behaviours (Reynolds, Ortengren, Richards, & de Wit, 2006). Each definition lends themselves to different self-report questionnaires or behavioural tasks which are designed to measure them. Numerous studies find that drug users score higher than non-users on self report measures of impulsivity, sensation seeking and inattention (Sher & Trull, 1994; Slater, 2003; Tercyak & Audrain-McGovern, 2003; Zuckerman, Ball, & Black, 1990), and also on behavioural measures of impulsivity (Fillmore & Rush, 2002; Lejuez et al., 2003; S. H. Mitchell, 1999). Although behavioural measures and self-report measures tend to be unrelated (Reynolds et al., 2006), indicating that this class of measures is broad and inconsistent. Intriguingly, impulsivity seems to increase when an individual starts using an addictive substance, which may cause a negative feedback loop (de Wit, 2009; Dick et al., 2010).

1.3.2 Sensitivity to Reward

After alcohol intoxication the heart rate increases, and it has been suggested that this may reflect an epigenetic (where the social environment cues genetic triggers) predisposition to alcohol dependence (Conrod, Pihl, & Vassileva, 1998). As the blood alcohol content increases (but not when it decreases), some alcoholics and men at risk of alcoholism have an exaggerated heart rate response (Conrod, Peterson, Pihl, & Mankowski, 1997; J. R. Wilson & Nagoshi, 1988). Individuals experience the stimulating effects of alcohol when their blood alcohol content is increasing (Davidson, Hutchison, Dagon, & Swift, 2002; Friedman, Carpenter, Lester, & Randall, 1980), and so the higher heart rate at this point may reflect the subjective experience of more stimulating effects (Holdstock, King, & de Wit, 2000) and fewer sedative effects (King, Carroll, Newton, & Dornan, 2002). Furthermore, questionnaire data show that the heart rate response to alcohol is related to reward seeking (Brunelle et al., 2004). Essentially, in the language of economic decision-making, individuals who consume addictive substances experience higher subjective utility from them (through reward sensitivity) than individuals who do not. The evidence is, however, inconsistent. Gilman, Ramchandani, Crouss, and Hommer (2012) found that heavy drinkers, compared to social drinkers, reported lower subjective alcohol effects and showed a reduced neural response to alcohol in the brain's reward system. This result suggests that heavy drinkers consume more alcohol because they need more before they get the same subjective effect. One way to consolidate these disparate findings is to split the subjective utility of alcohol intoxication (which is high in heavy drinkers), from the sensitivity to alcohol intoxication (which is low in heavy drinkers, although this may be the result of heavy drinking rather than an antecedent. It could be that heavy drinkers have a short time to habituation).

1.3.3 Risk-Taking

When clinical psychologists and economists characterize behaviour as 'risky', they mean different things (Schonberg, Fox, & Poldrack, 2011). For economists, risk is a variance in results but with the same expected value (e.g. 10% chance of £1000 and 90% chance of nothing is more risky than 100% chance of £100, but they have the same expected value). For clinical psychologists, risk is a behaviour that has potential negative consequences – it can harm oneself or others. Here the clinical psychologists' meaning is used.

Addictions are risky behaviours; almost a third of all deaths from cancer in the United States each year are directly caused by tobacco use (Lejuez et al., 2003). Both adults and adolescents are aware of this, yet they continue to smoke, suggesting that they are willfully taking a risk. Smoking is related to other real-world risky behaviours; daily smokers, as opposed to nonsmokers, are more likely to be in traffic accidents (Difranza, Winters, Goldberg, Cirillo, & Biliouris, 1986), are less likely to wear seat belts (Eiser, Sutton, & Wober, 1979), and engage in more high-risk sexual behaviours (Valois, Oeltmann, Waller, & Hussey, 1999). In the laboratory, the Balloon Analogue Risk Task (BART; Lejuez et al., 2002) measures risk taking. Participants press a button to pump up an imaginary balloon; each time they pump the balloon they earn more money, but at some point the balloon pops and they lose the money they earned. Participants can stop pumping at any time before it pops and keep the money they have earned up until that point. The number of pumps indicates how risk seeking the participant is. Results have demonstrated test-retest reliability (T. L. White, Lejuez, & Wit, 2008), and validity, in that daily smokers take more risks than nonsmokers (N=60; Lejuez et al., 2003), although this result was not replicated when smokers with co-morbidities were excluded (N=64; Dean, Sugar, Hellemann, & London, 2011). MDMA drug use was also found to be related to risk taking in the BART (Hopko et al., 2006).

1.3.4 Time Perspective

The Zimbardo Time Perspective Inventory measures the degree to which temporal information about the present and future influences behaviour (Zimbardo & Boyd, 1999). It includes three scales: two present-mindedness scales and a future-mindedness scale. The Consideration of Future Consequences Scale measures time perspective on a single scale between present- and future-mindedness (Strathman, Gleicher, Boninger, & Edwards, 1994).

Daugherty and Brase (2010) found that both time perspective scales were associated with alcohol, tobacco and drug use, as well as a range of other future oriented behaviours, and that the measures were still predictive even after demographics, Big Five personality, and delay discounting were controlled for, although the sample size (N=467) was not large enough to consistently distinguish between subfactors of time perspective. This indicates not only that an individual's time perspective influences their decisions, but that time perspective does not overlap with delay discounting. Gellert, Ziegelmann, Lippke, and Schwarzer (2012) found that the relationship between time perspective and health behaviours was more subtle; time perspective moderated the relationship between planning on health behaviours, such that for those who had a limited time perspective it was important to plan in order to behave healthily.

1.3.5 Mindfulness

Mindfulness reflects living in the here and now. Like impulsivity, mindfulness emphasises the present, but unlike impulsivity, mindfulness stems from a recognition that all things are impermanent (Marlatt, 2002). It means that those who are mindful would actually be less likely to act on impulse, because they are mindful of the consequences of their decisions and so do not act rashly. They act under conscious control rather than from automatic thoughts or impulses (C. Murphy & MacKillop, 2012). C. Murphy and MacKillop (2012) found zero-order correlations between mindfulness and alcohol use and adverse consequences of drinking, but these effects did not survive when including trait impulsivity as a measure.

1.4 Conclusion - Three Separate Models?

This thesis uses a cognitive perspective to approach the question of how decisions are biased by being made over time, and emphasises the advantages of logical models over descriptive ones. Although initial economic models were implausible relative to the cognitive apparatus present in the human brain, later models have utilised the insights from psychological research whilst still retaining their logical basis. This is important because logical models make specific predictions that are testable, and they extend other logical theories rather than creating another theory for the specific domain they were created for. Decisions are made in many domains and it is parsimonious to assume that there is only one decision making system in the brain for all of them.

With this in mind, how do intrapersonal externalities integrate with delay discounting and the impulsivity models? With regards to delay discounting, the delay until a reward is received is another mechanism that makes it difficult for an individual to decipher the long-term outcomes of each decision. While it is difficult to learn that a previous decision is affecting the results of a current decision, it is even harder to learn this if the previous decision was made a month or more ago. Additionally, a future consequence can still be discounted according to the individual's personal delay discounting function.

Some impulsivity models can be integrated. Sensitivity to reward could be considered simply as part of an individual's method of assigning utility to various options, leading to a higher subjective utility being given to a drug reward.

Similarly, risk taking is already part of the expected utility model (it is within the boundaries of the model to discount a risky reward), and Prospect theory even accounts for the psychological irrationalities of over or under-weighting certain probabilities. For time perspective and mindfulness, the evidence for their association with addictive behaviour is weak, especially once trait impulsivity is taken into account.

This leaves impulsive disinhibition, where the way forward to integration is less clear, since it suggests that actions are taken on a whim, due to reasons that are independent of deliberation. A strict reading of expected utility theory would state that impulses reflect utility, but since regret can set in very swiftly after an impulsive decision is made, this would imply an implausibly steep rate of discounting. In the experiments that follow, impulsive disinhibition is measured as part of a broad trait impulsivity questionnaire in order to determine its relationship with intrapersonal externalities.

There are other factors that are involved in the formation or perseverance of substance use beyond the three literatures described above. They are not measured in this thesis due to the practical constraint of measuring the breadth of factors which are potentially involved, because they were not broad enough to account for a range of addictive behaviours, or because it was not clear whether and how far they overlap with the literatures above. Some of them are nevertheless included in the Discussion chapter.

1.5 Thesis Framework

This thesis examines the concept of intrapersonal externalities as measured through the Harvard Game. It is useful to conceive of the following chapters as contributing to the psychometric steps used when creating a new questionnaire (Rust & Golombok, 2008). Before creating a new questionnaire, it should be determined that a psychological concept is not comprehensively measured elsewhere, and so there is a need for the new measure. In Chapter 2 a large online study demonstrates that despite using several methods of predicting impulsive behaviour, there is further variance to be explained. A new questionnaire should also demonstrate that it is a valid measure, which means that it should be related to other theoretically similar concepts; in Chapter 3 and 5 the relationship between intrapersonal externalities, and smoking behaviour, delay reward discounting and impulsivity are examined. Additionally, the criterion-related validity (Rust & Golombok, 2008), whether the Harvard Game tests the way it should given the theory of intrapersonal externalities presented in this chapter, is studied in Chapter 3 by examining the relationship between intrapersonal externalities and participants' understanding of the task. Finally, Chapter 4 extends the Harvard Game to use it as a platform to research other questions. If the Harvard Game represents a situation where participants struggle to learn to choose between long and short-term rewards, then it should be possible to use it to study the effects of advice-seeking and advice-taking on impulsive decisions.

CHAPTER 2

Impulsivity, Delay Discounting and Addictive Behaviours: An Online Study

2.1 Chapter Introduction

2.1.1 Motivation for experiments: Overlap between measures and unique variance accounted for

As described in Chapter 1, I have identified two major schools of thought regarding the origin of short-term decision making. Either individuals discount the future according to a personal discounting function but make decisions using this function in an economically rational manner, or they are victim to their psychopharmacology, which does not sufficiently control their impulses. This chapter examines these competing, but not necessarily mutually exclusive, approaches to test how far they explain addictive substance use and dependency, as examples of commonly understood short-term behaviours. Also included is a broad Big Five measure of personality (Digman, 1990), since this taxonomy has been shown to encompass a wide range of behaviours (Paunonen, 2003), including relationships with addictive substance use.

In order to address this subtle question, a large sample of individuals is necessary to achieve reliable results. This is because after including the demographics of age and gender, there are nine predictors, three of which might be expected to be correlate with one another (delay discounting, trait impulsivity and conscientiousness). For this reason data were collected through a large online social network application.

This chapter also aims to establish how well trait impulsivity, delay discounting, personality and demographics together are associated with use and dependency on addictive substances. This will determine how much space there is for an alternative model of choice over time, or whether the current theories simply need to be understood in more detail.

2.1.2 The myPersonality project: Example of successful utilisation of online social networks for large-scale social research

The Internet has made it possible for social scientists to inexpensively reach large samples of research participants, sometimes hundreds of thousands of people or more (Rentfrow, Gosling, & Potter, 2008). These samples are more diverse and representative than the WEIRD (Western, Educated, Industrialized, Rich and Democratic; Henrich, Heine, & Norenzayan, 2010) samples traditionally used by researchers (Gosling, Sandy, John, & Potter, 2010; Gosling, Vazire, Srivastava, & John, 2004), and the quality of data collected is as good or better than in traditional pencil and paper methods (Pettit, 2002) or even national surveys acquired via telephone polling (Chang & Krosnick, 2009).

In recent years there has been a remarkable shift towards more social and less anonymous Internet use. Interactions between people using anonymous nicknames, email addresses, or avatars are increasingly being replaced by interactions within Online Social Networks (OSNs) that are based on real identities and connections that largely mirror offline social links (Clouston, Amin, Verdery, & Gauthier, 2009; Ellison, Steinfield, & Lampe, 2007; Mayer & Puller, 2008). This presents a great opportunity in terms of research design, solves problems related to anonymity that were characteristic to previous web studies, and makes the data available on OSNs both far richer and on a larger scale than has been available to social scientists in the past (Kleinberg, 2008).

The most popular OSN today, Facebook (http://www.facebook.com), boasts over 800 million active users (as of December 2011). Facebook allows its users to create a personal profile, connect and communicate with their friends, and interact in various applications and games. Users can share and comment on content such as photos, videos, links or micro-blog entries. Importantly, Facebook conveniently keeps track of friends' activities and profile updates by automatically notifying its users about their friends' actions.

Despite its popularity in society, Facebook is far from being fully understood or even popular among social scientists, although a growing number of studies involve Facebook (410 in 2011; R. E. Wilson, Gosling, & Graham, 2012). Due to their role in the lives of millions of individuals and whole societies, Facebook and other OSNs constitute an interesting research area in their own right, but OSNs also provide excellent research grounds to collect data and carry out experiments related to virtually any of the social sciences.

In May 2007 Facebook released an Application Programming Interface (API) that allows developers unaffiliated with Facebook to create applications, such as games or quizzes, and integrate them with the Facebook platform. Basic OSN applications can be developed relatively easily as they are simply websites with some additional fragments of code that facilitate interaction with the OSN platform. Consequently, existing web questionnaires or other web-based applications used for research may be converted into an OSN application without much additional work. Studies using OSNs such as Facebook have multiple advantages:

1. Size and Global Reach

   As of December 2011, Facebook has over 845 million unique active users per month, and 483 million users per day. It is available in more than 70 languages, and 80% of users are from outside the U.S. and Canada (http://newsroom.fb.com/content/default.aspx?NewsAreaId=22). At the end of 2011, more than 7 million applications and websites integrated with Facebook (http://newsroom.fb.com/content/default.aspx?NewsAreaId=137).

2. User related information

   With the user's permission, applications can conveniently access and record an unparalleled set of data about the user through the Facebook API. Facebook stores most of its data in a network graph format, creating friendship links between users, "like" links between users and brands, statements or activities, "working at" links between users and employers, and "studying at" links between users and schools or universities. This format is extremely convenient compared to the plain text format of most of the Internet, which requires Natural Language Processing (NLP) algorithms to interpret.

   At the time of writing, the Facebook API contains data that falls into the following broad categories:

   (a) **Demographic Profile** containing unique id, name, profile picture, age, gender, relationship status, romantic interests, geographical location, place of origin, work and education history, political views, biography, link to personal website, and time zone.

(b) **Likes**, which represent endorsement or preference connection with virtually anything. Users express their likes (and thus create a connection) by pressing a like button that is pervasive on Facebook and, increasingly, in other parts of the web. For instance, one may like a friend's picture, a comment on a friend's picture, the "Jesus" fan page, a local pub, or an article on the BBC website.

(c) **Social network**, composed of friendship connections, connections between user's friends (thus, egocentric network), user's significant connections (e.g. spouse, parents, kids), and basic information about user's friends (depending on friends' privacy settings)

(d) **Photos and videos** in which user has been tagged or that were published by the user. This provides a direct link to users' social activity away from the web.

(e) **Textual writings** of users, including: status updates, comments, links, and notes published by the user or their friends

3. Facebook of real people

Users' unique identities benefit research methodologically by allowing a researcher to note when a user retakes a questionnaire. Instead of multiple submissions from the same user being a drawback of online research (Birnbaum, 2004), they present an opportunity to do longitudinal studies. A unique identity also makes it trivial to compare different questionnaire responses together, and so as more studies and questionnaires are made available the data becomes exponentially richer. It is also unnecessary to ask people to complete the same questionnaire or enter the same demographic information twice for two different studies. Finally, it is easy to set up an electronic system for users to automatically delete their research information. This means that users have more control over their data than in offline settings.

4. Convenient and Ecological

   People already use OSNs extensively so it is convenient for them to participate when they have the time and inclination, and in a physical location that they choose (Reips, 2000). This makes it unnecessary for a researcher to bring them into a laboratory. If a research project provides a service that allows users to interact with their friends, then these interactions are ecologically valid as they become part of the social behaviour that users do of their own volition.

5. Low Cost

   Recruitment of participants is a difficulty with web studies (Birnbaum, 2004). But if an OSN application is interesting for users, then they will spread it themselves through the social network.

In June 2007 I developed a Facebook application called myPersonality, soon after the launch of the Facebook API that for the first time allowed independent developers to integrate their applications with the Facebook platform. myPersonality offers Facebook users a set of genuine personality and ability measures, and then gives them personalized feedback on their results.

The number of measures provided by myPersonality is constantly expanding. As of January 2012 there were more than 25 scales available that ranged from the 336 item long IPIP NEO personality facets questionnaire to our own adaptive IQ test based on a 120 item long item bank. To date, over 6 million users participated in the most popular questionnaire, the IPIP version of the NEO Personality Inventory (Goldberg et al., 2006), and nearly half of them have allowed us to anonymously record their scores and information stored in their Facebook profiles. The great majority of this data, including records of online behaviour, rich demographic data, friendship networks, and scores on the other psychometric measures were made freely available to other academic researchers on the myPersonality project website (Stillwell & Kosinski, 2011).

Users do not receive a payment or course credit for using myPersonality. This provides better quality data, since unmotivated or careless participants may answer randomly, skip, or misread the test items (Kurtz & Parrish, 2001). Chang and Krosnick (2009) found that nonprobability Internet samples yielded the most accurate self-reports of political views, although they were a more biased sample than telephone interviewing or probability Internet samples. In myPersonality's 336 item IPIP proxy for Costa and McCrae's NEO-PI-R facets, users actually pay a small sum (about $4) to participate, which helps to ensure that participants are highly motivated to complete this long questionnaire accurately. The myPersonality sample is also more representative than those obtained in traditional studies:

1. Gender

   According to Gosling et al. (2004), 71% of participants in 510 traditional samples from 156 articles published in JPSP were female. In comparison, in their sample of web studies only 57% of participants were female. Among our participants 63% are female (Figure 2.1), a gender bias that is stronger than in on-line studies, but weaker than in the published samples.

2. Age

   The Facebook population is still predominantly young, and so are myPersonality users. The average participant is 23.5 years old and nearly half of them (47%) are between 18 and 24 years old (Figure 2.1). However, the age composition of the myPersonality sample seems to be very similar to that of traditional studies. As estimated by Gosling et al. (2004), the average age in the JPSP samples was 23 years, roughly equal to the web studies.

3. Geographic region

   There are few international barriers in web-surfing and thus online based studies are inexpensively available to geographically detached populations. myPersonality attracts more than 42% of respondents from outside the US. Additionally, 44 countries are represented by more than 1000 respondents. Consequently, the geographical diversity of our respondents (note that myPersonality is currently available exclusively in English) is

2.5 times higher than in the published samples reviewed by Gosling et al. (2004), which had 17% of non-US participants.

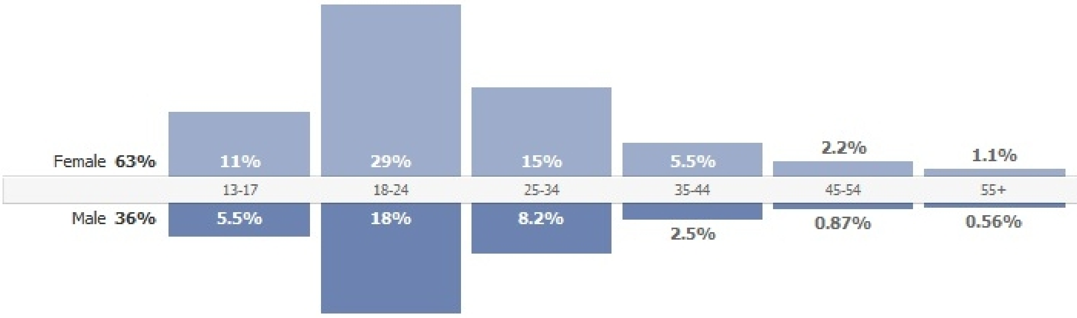

Figure 2.1: Gender and age distribution of myPersonality users

At least in case of the basic demographic properties of gender, age and geographic region, our samples collected on Facebook are not less representative than those obtained in traditional studies. Obviously, populations are defined by a number of other demographic traits (e.g. race, education, income) but it is not obvious why samples collected on Facebook should be particularly biased for any of those traits. On the contrary, Facebook reaches groups often excluded from samples, such as shy, socially or physically handicapped, or extremely busy individuals. For example, as it is shown on Figure 2.2, racial composition of the Facebook user base in the United States looks increasingly similar to the one in the general population – a situation still rare in many other contexts (e.g. Universities).

The following two experiments present data from three myPersonality questionnaires, two of which were added specifically for the purpose of this thesis. 'Today or Tomorrow?' measures delay reward discounting, the degree to which an individual discounts the value of a future reward due to the delay

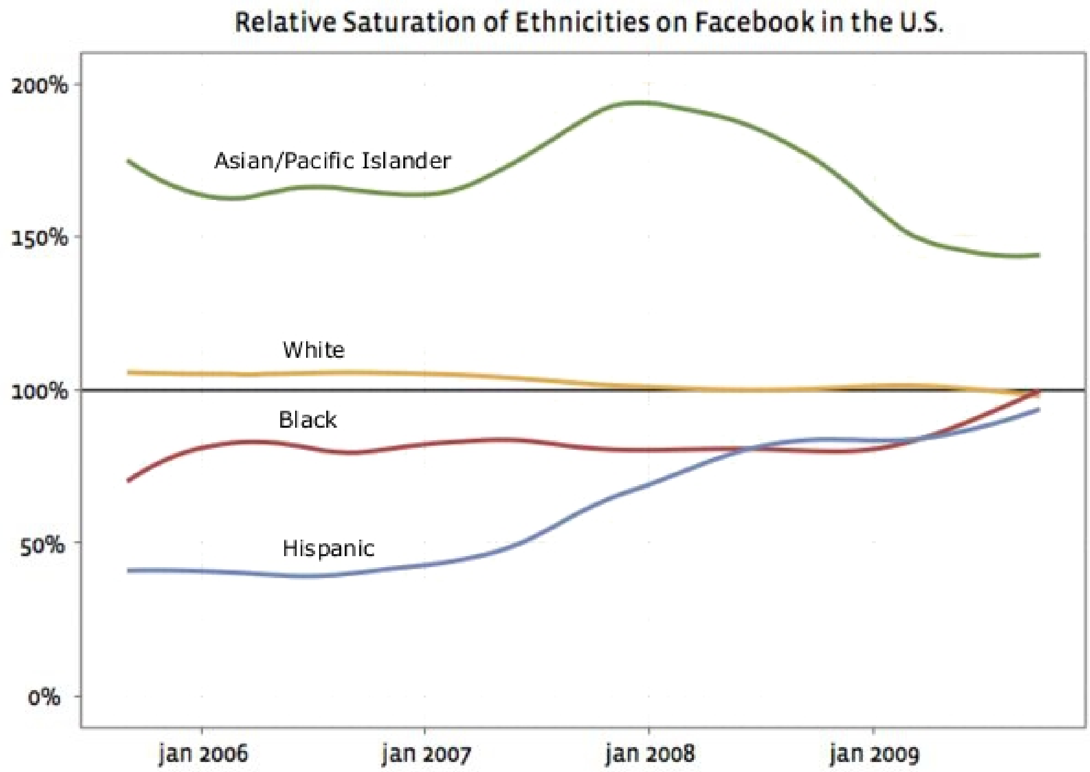

Figure 2.2: Relative saturation of the ethnicities on Facebook in the US (Marlow, 2009)

to receiving it, and it also includes questions on addictive substance use and dependency. It has been completed almost 15,000 times since it was introduced on 30th June 2010. The Barratt Impulsivity Scale (BIS-11; Stanford et al., 2009) was added in October 2010 and has been completed more than 18,250 times. The Big Five personality test was the first test added to myPersonality in June 2007, and is the default test that the vast majority of users take when they use the application for the first time. It has been answered by more than 6 million individuals, as well as being retaken almost 1 million times. This provides an unprecedented dataset to study individual differences in cognition.

2.2 Experiment 1: Individual Differences Related to Addictive Substances: Use and Dependency

The following results have been submitted for publication.

2.2.1 Introduction

About one third of tobacco smokers develop tobacco dependence, 15% of drinkers become alcohol dependent and 15% of users of other drugs become dependent (Anthony, Warner, & Kessler, 1994). This poses the question of why some people try an addictive substance and become lifelong users whereas others stop. Even in controlled animal studies which eliminate social influences, 17% of animals persist in self-administering a drug despite negative reinforcement (Deroche-Gamonet, Belin, & Piazza, 2004). This suggests there is an individual difference in vulnerability to addiction and that addiction is not just due to opportunity or other social factors.

Many studies and meta-analyses of humans have established links between personality measures and tobacco, alcohol and drug-use, however across all addictive substances there are more studies that relate substance-use to personality than substance dependency to personality. Since the majority of users of an addictive substance start when they are young (Grant & Dawson, 1998), this confounds use and dependency, since older users of an addictive substance are more likely to have high dependency. Although age can be statistically controlled for, it would be useful to compare whether there are different personality markers of substance use and substance dependency.

Leading theories in the psychopharmacological literatures assume that there is a similar process underlying all addictions , whether it is an automatic stimulus-response habit (Everitt, Dickinson, & Robbins, 2001; Tiffany, 1990), a drug-induced loss of impulse control in the frontostriatal system (Jentsch & Taylor, 1999), sensitization of the system that responds to incentives (Robinson & Berridge, 1993), or disruption of hedonic homeostasis (Koob & LeMoal, 1997). If there is an analogous process, then to the extent that personality factors are related to cognitive and neurological substrates of impulsivity (Adelstein et al., 2011; Caseras et al., 2006; DeYoung et al., 2010), similar personality factors should be predictive of use and dependency across different types of addictive substance, yet co-occurrence of addictions is rare (Sussman et al., 2011). In this study we compare tobacco, alcohol, and other drugs.

Big five personality (Digman, 1990; Goldberg, 1990) measures openness (liberal vs. traditional), conscientiousness (organized vs. flexible approach to life), extraversion (outgoing vs. less talkative), agreeableness (concerned by social cohesion vs. assertive), and neuroticism (easily stressed vs. emotionally stable). Personality traits are associated with the outcome of therapeutic interventions for substance dependency (Hooten et al., 2005), and understanding the individual differences in drug use and misuse could help the development of prevention programs and interventions to reduce addictive substance problems (Mezquita, Stewart, & Ruiperez, 2010).

Many researchers have reported links between big five personality and whether someone currently smokes. Malouff, Thorsteinsson, and Schutte (2006) meta-analysed 9 studies with 4,730 participants and found that lower conscientiousness, lower agreeableness, higher neuroticism and higher extraversion (outside the USA and Canada) correlated with being a smoker. Terracciano and Costa (2004) studied an elderly American population and found that higher neuroticism, lower agreeableness and lower conscientiousness correlated with being a smoker. Munafo and Black (2007) found that higher extraversion and higher neuroticism, as measured when participants were 16 years old, predicted future smoking between the age of 20 and 53. Hong and Paunonen (2009) used peer reports of personality and found that lower conscientiousness, lower agreeableness, higher neuroticism and lower extraversion were related to smoking status. Additionally, a meta-analysis of another model of personality, Eysenck's tripartite taxonomy (measuring extraversion, psychoticism and neuroticism), found that higher extraversion and higher neuroticism correlated with being a smoker Munafo and Black (2007). Two studies also reported analyses of interactions between personality factors; Hong and Paunonen (2009) found that being both lower in conscientiousness and lower in agreeableness were related to smoking, and Terracciano and Costa (2004) found that being both higher in neuroticism and lower in conscientiousness were related to smoking. Littlefield and Sher (2012) found that decreasing neuroticism over a 17 year time span in young adulthood is associated with desisting smoking during that time. Overall, there is broad agreement that individuals with higher neuroticism, lower conscientiousness, and higher extraversion are more likely to be smokers. There have not been enough studies looking at interactions between factors to show a clear weight of evidence.

Three studies reported the relationship between smoking dependency and big five personality traits. Munafo and Black (2007) did not find a relationship between personality measured at age 16 and number of cigarettes smoked between 20 and 53 for any of the big five traits using a sample size of 217. Littlefield and Sher (2012) found no relationship between changing personality in young adulthood and changes in the number of cigarettes smoked. Shadel and colleagues (2004) found that higher openness was related to motivation to quit and the number of quit attempts in the past year.

Alcohol consumption is more common than smoking across most cultures; one study at a Spanish university reported that just 3.5% of university undergraduates did not consume alcohol (Mezquita et al., 2010). Consequently, researchers rarely examine alcohol use except among adolescents (Nees et al., 2012) and instead the emphasis is on the level of alcohol consumption. A meta-analysis examining links between alcohol involvement and personality (Malouff, Thorsteinsson, Rooke, & Schutte, 2007) found that alcohol involvement was associated with lower conscientiousness, lower agreeableness and higher neuroticism. Clark and colleagues (2012) also found a relationship between lower conscientiousness and weekly alcohol use, but not for the other big five traits. Mezquita et al. (2010) found that higher extraversion and lower conscientiousness were related to alcohol consumed, whereas higher neuroticism was only related to alcohol problems rather than alcohol consumed. Nees et al. (2012) studied the correlates of early alcohol use in 14 year olds from a personality, behavioural (risk taking, risk adjustment and delay aversion) and neuroimaging (activation in reward-related brain regions) perspective. They found that personality dispositions accounted for by far the largest proportion of variance in alcohol use (16%) compared with behaviour (0.6%) and brain responses (0.4%). Of the big five traits, higher extraversion was a significant marker. Bottlender and Soyka (2005) found that outpatients with higher neuroticism and lower conscientiousness scores were more likely to relapse 12 months after treatment for alcohol dependence. Hopwood et al. (2007) found that there were few if any traits that differentiated individuals with current alcohol abuse problems from those with a history of alcohol problems, suggesting that the traits represent stable predispositions to substance use.

In contrast to alcohol, consumption of drugs is rarely separated into use and dependency, probably because any drug use is not socially acceptable, even though epidemiological studies find that half the US population have tried a drug for a non-medical reason Anthony et al. (1994). A meta-analysis of substance abuse disorder, excluding alcohol, found that higher neuroticism, lower conscientiousness and lower agreeableness were related to diagnosis (Kotov, Gamez, Schmidt, & Watson, 2010). Higher neuroticism was found to be related to not completing a substance abuse treatment program (Scalise, Berkel, & Van Whitlock, 2010). Bogg and Roberts (2004) studied conscientiousness in depth and found that lower scores were related to drug use. Among specific drugs; ter Bogt, Engels, and Dubas (2006) found that MDMA users are more extroverted than non-users but found no difference between heavy (three or more tablets in a session) or light users. Terracciano, Lockenhoff, Crum, Bienvenu, and Costa (2008) found that marijuana use was linked with higher openness, lower agreeableness and lower conscientiousness, and cocaine/heroin use was linked with higher neuroticism and lower conscientiousness.

Impulsivity and drug use are co-morbid (Stanford et al., 2009). Impulsivity is a behavioural endophenotype that precedes drug exposure (Ersche, Turton, Pradhan, Bullmore, & Robbins, 2010; Tarter, Kirisci, Habeych, Reynolds, & Vanyukov, 2004) and increases with drug exposure (Dallery & Locey, 2005). Furthermore, impulsivity is a multifaceted construct that includes inhibitory control and heightened preference for immediate reward (delay discounting), both of which have been consistently related to tobacco (Bickel et al., 1999; Daugherty & Brase, 2010; Heyman & Gibb, 2006; Reimers, Maylor, Stewart, & Chater, 2009; Stillwell & Tunney, 2012; Sweitzer, Donny, Dierker, Flory, & Manuck, 2008), alcohol (Daugherty & Brase, 2010; Reimers et al., 2009), and drug use (Belin, Mar, Dalley, Robbins, & Everitt, 2008; Reimers et al., 2009).

Measures of impulsivity are rarely used in tandem with big five personality, and when they have inconsistent results have been found. Daugherty and Brase (2010) found that steep delay discounting is related to frequency of tobacco and alcohol use (but not frequency of drug use), but this effect was no longer significant after accounting for time perspective and big five personality. Hair and Hampson (2006) found that conscientiousness and trait impulsivity both predicted alcohol consumption in female college students but both accounted for the same variance. Conversely, conscientiousness has a small correlation to various behavioral measures of impulsivity (Edmonds, Bogg, & Roberts, 2009). This raises the question of whether the broad personality measures are explaining additional variance beyond the narrower impulsivity construct.

We make three unique contributions. First, it is known that impulsivity is related to addictive behaviors and that big five personality is also related to addictive behaviors, but not whether they both are together. We measure big five personality, a trait measure of impulsivity, and the behavioural measure of delay discounting in order to test whether and how broad personality traits extend our understanding of addictive behaviours beyond impulsivity. Second, we study tobacco use, alcohol use and drug use in all participants, and also include a measure of dependency for each. This allows comparison of how individual differences relate to substance use and dependency separately, and to see whether the same measures are related to the three addictive substance categories. Finally, we look at interactions between the big five factors as markers of substance use, since this is an understudied topic in the literature.

2.2.2 Method

Procedure

In order to maximize the sample size we used an online social network application to collect a large dataset. All users of the Facebook online social network are able to add the myPersonality application to their account (Stillwell & Kosinski, 2011), and more than 6m users have completed a questionnaire so far. Users are initially presented with the Big Five questionnaire, and are then given the option of completing more than 20 other personality measures. Participants are only incentivized by the prospect of receiving individualized feedback on the results of each questionnaire. The myPersonality application appears within the Facebook website, surrounded by Facebook navigation and advertising. Participants can stop using the application at any time by closing the window or navigating away from the page.

The 'Today or Tomorrow' questionnaire received ethical approval from the University of Nottingham School of Psychology Ethics Committee. Respondents gave consent for their results to be used in an anonymised format for research, were able to stop at any time without pressure by simply closing the browser window, and were able to delete their information after completing the study through an automatic form mechanism within the 'My Personality' application.

Measures

Big five personality was recorded on myPersonality using the IPIP version of the NEO personality inventory (Goldberg et al., 2006). Users are initially presented with a consent form, explaining how their information will be used for research and that they can delete their information at any time by following the automatic mechanism. Users then completed the Big Five questionnaire. At the bottom of the questionnaire, they were asked to opt in to sharing their

Facebook profile information, and from this their age and gender was recorded. Users then received a graphical and textual description of their personality results.

Delay discounting was recorded on myPersonality using a questionnaire called 'Today or Tomorrow?' Users were told that they would receive feedback on their delay discounting function, which measures how much a user prefers the present over the future. Participants were asked to make repeated choices between two monetary amounts; varying amounts now compared to a larger amount at some future point. A delay discounting parameter was calculated for each participant using a hyperbolic discounting curve as explained in Stillwell and Tunney (2012), such that greater values reflect a sharper decline in the subjective value of money as the delay to obtain the money increases.

Substance use was assessed at the beginning of the Today or Tomorrow questionnaire under a section labelled as optional. Users answered the question 'Do you smoke?' by choosing between 'daily or more', 'less than daily', or 'never'. Users were then asked 'If you smoke at all, please answer the following questions', and then presented with the questions from the Cigarette Dependence Scale (CDS-5; Etter, Le Houezec, & Perneger, 2003). Users then answered 'Do you drink alcohol?' by choosing between 'weekly or more often', 'less than once a week', or 'never'. Users who drink at all were then asked to complete the Alcohol Use Questionnaire (AUQ; Townshend & Duka, 2005). Finally users were asked 'Do you use drugs?' and given the option between 'weekly or more often', 'less than once a week' or 'never'. Users who used drugs at all were asked to complete the Assessment of Substance Misuse in Adolescents (ASMA; Willner, 2000).

Trait impulsivity was assessed on myPersonality using the Barratt Impulsivity Scale (BIS-11; Stanford et al., 2009) in its own section of the site. Users were told that they would receive feedback on their overall impulsivity score and on the three second order factors: attentional impulsiveness, non-planning impulsiveness and motor impulsiveness.

Participants

A total of 18,278 users entered their smoking status, alcohol status or drug-use status, although the vast majority of participants did not complete all of the demographic, personality and impulsivity measures and so most analyses use a subset. Raw correlations among all of the variables are shown in Table 2.1, including the number of participants, which indicates how many completed the various measures. 64% of participants were female, and the average age was 23.02 years (SD=9.2 years).

Table 2.1: Spearman 2-tailed correlations among the variables analysed. Correlations in bold are significant at p=.001 or below. Sex is coded 0 = male, 1 = female. Nicotine, alcohol and drug-use are coded 0 = no, 1 = yes. Numbers in parentheses are the N for that comparison.

	Sex	Ope	Con	Ext	Agr	Neu	BIS-11	Delay Discounting	Nicotine	Alcohol	Drug	CDS-5	AUQ	ASMA
Age	-.006 (10526)	.023 (10297)	**.185** **(10297)**	**.033** **(10297)**	**.056** **(10297)**	**-.075** **(10095)**	-.028 (1812)	-.027 (6069)	**.306** **(10175)**	**.422** **(10090)**	**.135** **(6794)**	**.332** **(2333)**	-.019 (4372)	.042 (1160)
Sex		.007 (10936)	-.017 (10936)	**.036** **(10936)**	**.049** **(10936)**	**.179** **(10717)**	**.091** **(1813)**	.027 (6455)	<.001 (10717)	-.022 (10613)	-.029 (7167)	**.073** **(2385)**	**.058** **(4498)**	.089 (1181)
Ope			**-.082** **(17768)**	**.205** **(17768)**	**.120** **(17768)**	**-.073** **(17402)**	.009 (1810)	**-.058** **(9003)**	**.059** **(16899)**	**.084** **(16653)**	**.100** **(11154)**	**-.062** **(3680)**	.006 (6784)	-.064 (1622)
Con				**.231** **(17769)**	**.223** **(17769)**	**-.319** **(17403)**	**-.409** **(1810)**	**-.093** **(9003)**	-.006 (16900)	.018 (16654)	**-.107** **(11154)**	-.017 (3680)	**-.064** **(6784)**	**-.124** **(1622)**
Ext					**.238** **(17769)**	**-.357** **(17403)**	**.117** **(1810)**	**.051** **(9003)**	**.086** **(16900)**	**.089** **(16654)**	**.050** **(11154)**	-.024 (3680)	**.075** **(6784)**	**-.108** **(1622)**
Agr						**-.342** **(17403)**	**-.134** **(1810)**	**-.042** **(9003)**	**-.067** **(16900)**	**-.046** **(16654)**	**-.056** **(11154)**	-.016 (3680)	-.023 (6784)	**-.130** **(1622)**
Neu							**.179** **(10717)**	**.092** **(8807)**	**.038** **(16552)**	<.001 (16314)	.023 (10923)	**.137** **(3593)**	.036 (6588)	**.213** **(1572)**
BIS-11								**.200** **(1163)**	**.147** **(1743)**	**.127** **(1729)**	**.189** **(1386)**	.093 (470)	**.156** **(949)**	**.268** **(284)**
Delay Discounting									**.113** **(8920)**	.025 (8844)	**.072** **(6143)**	**.115** **(2004)**	**.085** **(3982)**	**.098** **(1025)**
Nicotine										**.376** **(17889)**	**.402** **(12003)**		**.217** **(7230)**	**.182** **(1690)**
Alcohol											**.303** **(11936)**	**-.070** **(3950)**		-.061 (1691)
Drug												-.006 (2524)	**.208** **(4802)**	
CDS-5													**.089** **(2908)**	**.228** **(1009)**
AUQ														.075 (1287)

2.2.3 Results

Raw correlations among all measures are presented in Table 2.1. In order to assess the ability to predict whether a given user would use a particular addictive substance, three hierarchical binary logistic regressions were used. They examined the unique associations between the demographic, personality and impulsivity measures, and smoking status, alcohol status, and drug-use status as the criterion variables (coded as 1 = yes, 0 = no). In the first step, age and gender were included as predictors, and in the second step the psychological predictors were added. For all models there was a significant difference in variance between step 1 and step 2. The models are presented in Table 2.2 and show that age, openness, lower agreeableness, BIS-11 impulsivity and steeper delay discounting predicted smoking. Age, extraversion, lower agreeableness and BIS-11 impulsivity predicted alcohol use. Openness, lower conscientiousness, lower agreeableness, BIS-11 impulsivity, and steeper delay discounting predicted drug-use.

In order to assess the ability to predict whether a substance user would show greater dependency, only users of a particular substance completed the Cigarette Dependence Scale (CDS-5; Cronbach's α=.88; Etter et al., 2003) which measured smoking behaviour, the Alcohol Use Questionnaire (AUQ; This scale counts the number of units consumed and so a reliability measure is inappropriate; Townshend & Duka, 2005) which measured their alcohol intake, and/or the Assessment of Substance Misuse in Adolescents (ASMA; Cronbach's α=.81; Willner, 2000) which measured their drug dependence. Three further hierarchical linear regressions examined the unique associations between the demographic, personality and impulsivity measures, and smoking dependency, alcohol dependency and drug dependency. In the first step, age and gender were included as predictors, and in the second step the psychological predictors were added. For all models the difference in variance accounted for between step 1 and step 2 was significant. The models are presented in Table 2.2 and show that age, neuroticism, and delay discounting were associated with smoking dependence. Only BIS-11 trait impulsivity was associated with alcohol dependence, and neuroticism and BIS-11 impulsivity were associated with drug dependence.

Table 2.2: Regressions showing the relationship between demographic, personality and impulsivity measures, and addictive behaviours. Bold relationships are significant at p<.05

Behaviour	Smoking					Alcohol					Drug-Use				
	Yes/No Logistic Regression N = 1082			CDS-5 Linear Regression N = 282		Yes/No Logistic Regression N = 1081			AUQ Linear Regression N = 607		Yes/No Logistic Regression N = 869			ASMA Linear Regression N = 185	
	Wald χ^2	p	Odds	t	p	Wald χ^2	p	Odds	t	p	Wald χ^2	p	Odds	t	p
Step 1															
Demographics															
Age	**32.73**	**<.001**	**1.04**	**6.12**	**<.001**	**43.233**	**<.001**	**1.06**	**-2.44**	**0.02**	0.05	0.827	1	0.66	0.51
Sex	0.04	0.85	1.03	1.39	0.17	0.002	0.97	1	0.35	0.73	0.51	0.47	0.89	1.85	0.07
R^2	Cox & Snell R^2 =.03			R^2 = .134		Cox & Snell R^2 =.05			R^2 = .01		Cox & Snell R^2 =.001			R^2 = .02	
Step 2															
Personality															
Openness	**14.99**	**<.001**	**1.72**	-0.04	0.97	3.28	0.07	1.26	-0.11	0.91	**10.76**	**0.001**	**1.77**	-0.85	0.4
Conscientiousness	0.74	0.39	0.9	-0.15	0.88	0.14	0.7	0.96	-0.43	0.67	**6.18**	**0.01**	**0.71**	-0.04	0.97
Extraversion	3.11	0.08	1.2	0.33	0.74	**11.7**	**0.001**	**1.39**	0.4	0.69	0.52	0.47	1.09	-0.94	0.35
Agreeableness	**12.45**	**<.001**	**0.65**	1.08	0.28	**8.89**	**0.003**	**0.71**	0.02	0.98	**4.74**	**0.03**	**0.73**	0.09	0.92
Neuroticism	1.54	0.22	1.14	**2.58**	**0.01**	0.09	0.77	1.03	0.91	0.36	1.45	0.23	0.87	**2.58**	**0.01**
Impulsivity															
BIS-11	**13.44**	**<.001**	**2.57**	0.22	0.83	**12.3**	**<.001**	**2.4**	**2.79**	**0.01**	**10.61**	**0.001**	**2.71**	**4.14**	**<.001**
Delay Discounting	**24.94**	**<.001**	**1.96**	**3.32**	**0.001**	1.84	0.18	1.19	1.33	0.18	**6.82**	**0.01**	**1.52**	1.57	0.12
Sig.	χ^2 = 143.21, p<.001			F=7.51, p<.001		χ^2 = 114.48, p<.001			F=2.72, p=.004		χ^2 = 65.33, p<.001			F=5.18, p<.001	
R^2	Cox & Snell R^2 = .124			R^2 = .199		Cox & Snell R^2 = .100			R^2 = .039		Cox & Snell R^2 = .072			R^2 = .209	

In order to distinguish which sub factor of BIS-11 impulsivity was predictive of both substance use and substance dependency, all models were rerun using the second order factors of the BIS-11 (attentional impulsiveness, motor impulsiveness, and non-planning impulsiveness) instead of the overall score. It was found that in all models except one, the only significant second order predictor was non-planning impulsiveness. The exception was when predicting ASMA scores, the only significant second order predictor was motor impulsiveness. In all models, using the second order BIS-11 factors rather than the overall score did not change the pattern of other significant predictors, except when predicting whether a given user smokes, extroversion changed from p=0.08 to p=0.036.

It was surprising that we did not find a relationship between smoking and the personality traits of higher neuroticism and lower conscientiousness given the broad agreement in the literature of a link (Bogg & Roberts, 2004; Hong & Paunonen, 2009; Malouff et al., 2006; Munafo & Black, 2007; Munafo, Zetteler, & Clark, 2007; Raynor & Levine, 2009; Terracciano & Costa, 2004; Terracciano et al., 2008). We therefore repeated the hierarchical binary logistic regression but removed delay discounting and trait impulsivity as predictors, leaving age, gender, and the big five personality traits (N=9715). This analysis found that neuroticism (Wald χ^2 = 25.99, p<.001) and lower conscientiousness (Wald χ^2= 23.44, p<.001) were significant predictors. We also repeated the analysis for alcohol use (N=9637) and again found that both neuroticism (Wald χ^2=7.59, p=.006) and lower conscientiousness (Wald χ^2=27.78, p<.001) were significant predictors. This indicates that these broad personality traits are capturing variance that can be more narrowly understood as trait impulsivity or delay reward discounting.

We also tested for interactions between personality traits associated with use of the three addictive substances. We used the hierarchical method followed by Terracciano and Costa (2004) where age, gender, and the five personality traits are added, followed by the interactions between all pairs of traits to see if they explain additional variance. For tobacco use (N=9715) there was a significant interaction between openness and neuroticism (Wald χ^2=8.28, p=.004) such that being both open to experience and being emotionally unstable is related to smoking. For alcohol use (N=9637) there was a significant interaction between openness and agreeableness (Wald χ^2=4.75, p=.029), such that being both open to experience and high in agreeableness is related to drinking. For drug use (N=6434) there was again a significant interaction between openness and agreeableness (Wald χ^2=4.65, p=.031), such that being both open and agreeable is related to drug use.

2.2.4 Discussion

We used a social network application to collect data on demographics, personality, impulsivity, and addictive substance use and dependency from a large sample of international participants. The percentage of variance accounted for by the predictors was between 20.9% for drug dependency and just 3.9% for alcohol consumption, indicating large differences in the explanatory power of these same predictors between different drugs. Age and gender were poor predictors overall; for smoking and alcohol, age was a significant marker, where being older made one more likely to be a smoker, to be more dependent on cigarettes, to use alcohol, and to drink more per month. None of the analyses showed an effect of gender.

Among the big five personality factors, those which were associated with substance use were not associated with substance dependency and vice versa. The most consistent marker of substance use was low agreeableness, which was related to smoking, alcohol and drug use but not dependency. Loukas, Krull, Chassin, and Carle (2000) reported that individuals lower in agreeableness reported stronger coping motives to use alcohol and weaker upbringing reasons to limit drinking, which may generalise to other addictive substances in that individuals lower in agreeableness are more resistant to parental pressure to avoid addictive substances. Higher openness was also related to substance use but not dependency; it predicted tobacco use and drug use and was marginal for alcohol use (p=.07). Being open to new experiences includes being open to trying addictive substances, but it does not relate to becoming dependent. Surprisingly, low conscientiousness was only a factor in drug use (but not dependency), and high extraversion was only associated with alcohol use (but not dependency).

The only big five personality trait that was associated with dependency was high neuroticism, which was related to both nicotine dependency and drug dependency, whereas it was not related to use of any of the addictive substances. In this study we used the Alcohol Use Questionnaire as our measure of alcohol dependency, which simply counts the number of units consumed, but large amounts of alcohol consumed is not necessarily the same as having an alcohol problem and the point at which consumption becomes problematic may be different between individuals. This may explain why we did not find a relationship between higher neuroticism and our measure of alcohol dependency – Mezquita et al. (2010) found that neuroticism was related to alcohol problems but not alcohol consumed. It may be that neuroticism determines the point at which an individual considers their addictive substance use to be problematic.

We also looked for interactions between the big five traits relating to addictive substance use. For drinking and drug use there was a significant interaction between higher openness and higher agreeableness. This was surprising since in our data low agreeableness is consistently associated with substance use, and it could suggest that whereas low agreeableness allows one to resist parental pressure, a combination of high openness and high agreeableness could make one more susceptible to peer pressure. For smoking there was a significant interaction between higher openness and higher neuroticism, indicating that if one is open to trying new experiences and also prone to anxiety, then one is more likely to try smoking. This effect does not match those found by previous researchers; Hong and Paunonen (2009) found an interaction between lower conscientiousness and lower agreeableness, and Terracciano and Costa (2004) found one between higher neuroticism and lower conscientiousness. This difference may be related to our inclusion of measures of impulsivity.

The link between substance use and the personality traits of higher neuroticism and lower conscientiousness disappeared when trait impulsivity and delay discounting were included as predictors. This indicates that these broad personality traits are capturing variance in substance use that can be more narrowly understood as trait impulsivity or delay discounting, which replicates the conclusion of Hair and Hampson (2006) who studied alcohol consumption. This finding demonstrates the value of including multiple individual difference measures that would be expected to be related to addictive behaviour, so that they control for one another and the potential mechanisms underlying the broad traits can be recognised.

Trait impulsivity as measured by the BIS-11 was predictive of both use and dependency for alcohol and drugs, and of tobacco use but not dependency. Post-hoc analyses of the three second order factors of the BIS-11 found that in all cases the effect was driven by non-planning impulsiveness, except for drug dependency which was linked to motor impulsiveness. Steep delay discounting was still associated to addictive substances in half of the analyses even after taking into account conscientiousness and BIS-11 impulsivity. It predicted smoking and smoking dependency, and drug-use, and for all analyses the effect was in the expected direction even if it did not reach significance.

Given that the percentage of variance accounted for by this range of predictors was as low as 3.9% (for alcohol consumption), and that even when a range of predictors reached significance as in drug use which was associated with three personality traits and two measures of impulsivity and yet just 7.2% of variance was accounted for, we accept that there are many other potentially important factors which are not included in this study. Other factors identified in the literature include risk taking (Lejuez et al., 2003), impulsive disinhibition (Reynolds et al., 2006), sensitivity to reward (Brunelle et al., 2004), time perspective (Daugherty & Brase, 2010), mindfulness (C. Murphy & MacKillop, 2012), and cognitive decision biases (Heyman & Dunn, 2002).

Two limitations with the study were first that the sample was a convenience one. On the other hand, users were incentivized to answer honestly in all questionnaires because their only reward was to receive feedback on their results. Additionally, the substance use questions were marked as being optional, so users had no reason to lie. A second limitation is that we asked generic questions about drug use, whereas given the results of this research we would expect different individual differences to be associated with different drugs.

Using responses from a large online sample, we conclude that big five personality traits explain unique variance in addictive substance use and dependency even after accounting for trait impulsivity and delay discounting. Low agreeableness and high openness are consistently associated with addictive substance use, whereas high neuroticism is associated with substance dependency. There are also other personality traits which are related to specific substances; high extraversion is related to alcohol use and low conscientiousness is related to drug use. We also found interactions between personality traits which hint at a more complex relationship; although further studies should be undertaken as results are inconsistent. Researchers should separate studies on the use of addictive substances from those on dependency, as different individual differences underlie the two processes.

2.3 Experiment 2: Effects of measurement methods on the relationship between smoking and delay reward discounting

The following results were published in Addiction (Stillwell & Tunney, 2012)

2.3.1 Introduction

Delay reward discounting (DRD) measures how a reward's subjective utility decreases as the interval before it is obtained increases. Steeper DRD has been found to be related to a range of addictive behaviours, including smoking (Baker, Johnson, & Bickel, 2003; Bickel et al., 1999; Heyman & Gibb, 2006; Johnson, Bickel, & Baker, 2007; Reynolds et al., 2007; Sweitzer et al., 2008), the success of smoking cessation (Krishnan-Sarin et al., 2007; MacKillop & Kahler, 2009), the initiation of regular smoking in adolescents (Audrain-McGovern et al., 2009), drinking (Vuchinich & Simpson, 1998), heroin and cocaine (Heil, Johnson, Higgins, & Bickel, 2006; Kirby & Petry, 2004), opioids (Madden et al., 1997) and marijuana (Johnson et al., 2010). However, the measurement of DRD varies and it is possible that this may affect estimates of effect size or cross-experiment comparisons (Mackillop et al., 2011). We therefore compared multiple methods of measuring DRD to estimate whether the relationship between DRD and smoking changed depending upon the measurement method.

In humans, DRD is typically measured by giving individuals repeated choices between various immediate amounts and a larger amount delayed for various lengths of time. Normally the rewards are hypothetical, but equivalent results have been found whether DRD is measured using real or hypothetical rewards (Johnson & Bickel, 2002). For example, an individual is asked to choose either \$500 today or \$1000 in a month. The point at which an individual switches from preferring a delayed reward to an immediate reward allows estimation of how much they subjectively value the delayed reward in today's money.

In normative economics, it has been assumed that individuals have an exponential discounting function. This has the implication that individuals are time consistent. Using an exponential curve, a parameter k can be calculated which represents the individual's degree of DRD. This is shown in Equation 2.3.1, where D is the delay in months, A is the amount of the reward and V is the subjective value of the reward as determined by an individual's preferences.

$$V = Ae^{-kD} \tag{2.3.1}$$

A large k indicates that the individual steeply discounts the future, whereas a low k indicates that the individual is more willing to wait. However, in humans DRD is typically found to follow a hyperbolic curve, such that small delays have a proportionately larger impact than longer delays (Rachlin et al., 1991). This is shown in Equation 2.3.2.

$$V = \frac{A}{1 + kD} \tag{2.3.2}$$

This accounts for time inconsistent preferences, where an individual switches their preference from a delayed reward to an immediate reward as the time before the reward is available decreases. The exact form of the DRD function has important economic implications, as the regrets inherent in hyperbolic discounting could be corrected through taxation (Gruber & Koszegi, 2001).

Although DRD has been conceptualised as a single stable trait that underlies decisions about delay in all domains (Odum, 2011), it has been found in the past that different methods or parameters used when measuring delay discounting leads to different DRD parameter estimates. It is possible that these variations could cause illusory or differing effects to be found between addictive behaviours and DRD. We set out to confirm the effect of these methodological differences and to see how they might interact with different groups of smokers.

The first is that there is evidence that the order that rewards are presented can affect the derived DRD parameter. Randomizing the order of the immediate rewards leads to the highest rate of discounting, followed by putting the rewards in an ascending order, and then in a descending order (Robles & Vargas, 2006, 2008; Robles, Vargas, & Bejarano, 2009). People who smoke tend to be higher in impulsivity and impulsivity itself may lead to spurious findings that smoking is associated with delay discounting. For example, assessing delay discounting requires people to repeatedly rate a higher distant reward that remains constant against a smaller reward that changes. One measure of impulsivity, the Barratt Impulsivity Scale (BIS-11; Koff & Lucas, 2011; Stanford et al., 2009), measures a dimension called attentional impulsiveness (difficulty maintaining attention), which has been found by some researchers to be most closely related to DRD (de Wit, Flory, Acheson, McCloskey, & Manuck, 2007; Kirby, Petry, & Bickel, 1999). Smokers might switch to the immediate reward earlier simply because they want to stop choosing the same delayed reward. To mitigate order effects, researchers have used a titrating procedure, where the immediate amount decreases after the immediate amount is chosen, and increases after the delayed amount is chosen, leading to an increasingly accurate assessment of the switch point. Rodzon (2011) found no difference in the derived DRD parameter between a fixed procedure and a titrating procedure (although with a sample size of just 24), but a titrating procedure still requires an initial amount, against which future amounts could be anchored by the participant. This could lead to unexpected effects due to trait impulsivity differentially interacting with the anchoring effect rather than differences in DRD. In the current experiment, participants were randomly allocated to three groups where the immediate rewards were either presented in an ascending, descending, or randomized order.

The second parameter difference is that typically in DRD studies multiple delays are used with anything from 6 hours to 25 years. Within a study, the DRD parameters for each delay are calculated and then averaged. This averaging is understandable if it is assumed that an individual has a single delay DRD parameter accurately measured by their discounting function and that variations around this are noise, however it may be that some individuals discount some delays more than others. For example, smokers could plausibly be expected to discount the distant future comparatively more than they discount tomorrow, since smoking involves negative returns over a decadal timeframe. In the current experiment, participants' DRD parameters were calculated separately for each delay and then compared to see whether smokers discounted certain timeframes more than others.

The third parameter difference is the size of the delayed amount. Often $1000 is used, however when other amounts are used studies find that that smaller amounts are discounted more steeply than large amounts (Grace & McLean, 2005; Loewenstein & Prelec, 1992; S. H. Mitchell & Wilson, 2010); this is commonly known as the magnitude effect. It is possible that larger amounts may be discounted more by smokers than smaller amounts, since smoking involves repeatedly choosing a smaller reward over a long-term health decline. In the current study, $1000 and $100 were compared within-subjects at a time delay of 1 month.

We aimed to test whether these differences in DRD methodology would systematically bias the observed relationship between smoking and DRD, and to see whether the relationship is only found under certain conditions. Since these effects are subtle, a large sample size was necessary to achieve reliable results. We therefore used an application which runs on the Facebook social network. Nearly 9500 international users completed a multiple-item DRD task in return for feedback on their results, and agreed to share their data with the researchers.

2.3.2 Method

Delay Reward Discounting Measure

Participants were asked to make repeated choices between two monetary amounts; various amounts now compared to larger amounts at some future point. The delays and amounts are a subset of those used in previous DRD research (Bickel et al., 1999; Rachlin et al., 1991). The 15 immediate monetary rewards were $1000, $950, $900, $850, $750, $600, $500, $400, $250, $150, $100, $60, $20, $10, and $1, and the six delays were 1 week, 2 weeks, 1 month, 6 months, 1 year, and 5 years. These were all compared to $1000 at the future time point. Participants were also asked an extra set of questions which asked for their preferences of rewards at a delay of 1 month that were a tenth of the size of those above ($100, $95, etc.) and compared to $100 at the future time point. This totalled seven sets of questions, which were presented in a randomized order for each participant.

Participants were randomly allocated to one of three groups. Within each set of questions, the amounts were either presented in an ascending order ($1, $10, [. . .], $1000), descending order ($1000, $950, [. . .], $1), or randomized order.

To calculate the participant's hyperbolic DRD parameter (k) for each delay in both their feedback and for the research, an indifference point was established by calculating an average between the maximum immediate monetary amount chosen and the minimum delayed monetary amount chosen (Bickel et al., 1999).

The parameter was then calculated according to the hyperbolic DRD formula mentioned above. Since the distribution of k is often found to be non-normal (Bickel et al., 1999; Rachlin et al., 1991), the data were approximately normalised using the natural-log transformation. The k parameter reflects the steepness of the discount curve, whereby greater k values reflect a sharper decline in the subjective value of money as the delay to obtain that money increases.

Procedure

Users of the 'My Personality' application on the Facebook social network (Stillwell & Kosinski, 2011) were invited to participate in a new questionnaire called 'Today'or Tomorrow?' and told that they would receive feedback on their results. Users who chose to start the 'Today or Tomorrow' questionnaire were given further information about the study (Appendix A), including being told that their DRD function would be estimated but that the rewards were hypothetical. As 'My Personality' is used by an international audience, users were given the option to choose a currency that they either used or were most familiar with from 9 of the most widely used world currencies (British Pound, Canadian Dollar, Euro, Filipino Peso, Indian Rupee, Indonesian Rupiah, Singapore Dollar, South African Rand, United States Dollar). The delayed amounts were based on a published study that used US Dollars, and so the amounts were converted to the other currencies using the exchange rate from Google's exchange rate function on 22nd June 2010 (see Table 2.3).

Table 2.3: Currencies used, including exchange rates from Google's exchange rate function on 22nd June 2010

Name	Symbol used on 'My Personality'	Conversion per $1 USD	N	Mean Age (SD)	% Male	% daily/ non-daily/ non-smokers	Mean(log(k)) (SD)
British Pound	£	0.68	915	25.0 (10.2)	38%	24/7/70	-.93 (.58)
Canadian Dollar	$	1.02	436	22.4 (9.2)	33%	17/6/76	-1.09 (.58)
Euro	€	0.81	644	24.7 (8.1)	42%	29/10/62	-.98 (.57)
Filipino Peso	P	45.45	159	22.5 (6.1)	42%	11/14/75	-.73 (.60)
Indian Rupee	Rs	45.65	156	23.1 (6.1)	56%	15/6/78	-.75 (.57)
Indonesian Rupiah	Rp	9009	51	22.9 (6.1)	59%	16/6/78	-.59 (.67)
Singapore Dollar	$	1.38	184	20.3 (6.4)	41%	10/4/86	-.91 (.58)
South African Rand	R	7.51	63	26.7 (9.1)	36%	32/5/63	-1.03 (.52)
United States Dollar	$	1	6430	22.8 (9.2)	36%	17/7/75	-.98 (.56)

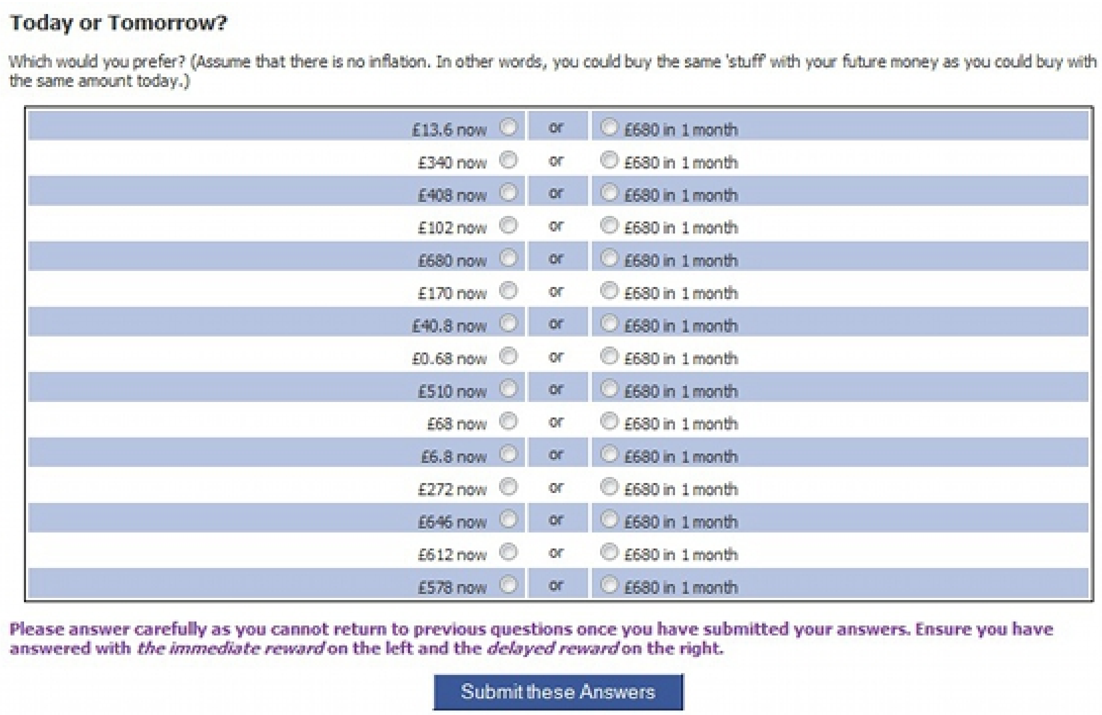

Figure 2.3: Screenshot of 'Today or Tomorrow' questions on 'My Personality'. Immediate amounts are presented in descending order and have been converted into British Pounds.

Users were also asked to answer the question "Do you smoke?" with answers "daily or more", "less than daily", or "never", and complete the Cigarette Dependence Scale (CDS-5; Etter et al., 2003). Users also completed the Alcohol Use Questionnaire (AUQ; Townshend & Duka, 2005), and the Assessment of Substance Misuse in Adolescence (ASMA; Willner, 2000) although these results are not presented here. Users were told that completion of these questionnaires was optional and would not affect their feedback. Users then completed the DRD measure (Figure 2.3), including short instructions telling them to assume no inflation (Kawashima, 2006). Finally, for feedback users were told which quartile their calculated DRD parameter was in, and shown a graph with their personal DRD curve compared to the mean DRD curve (Figure 2.4).

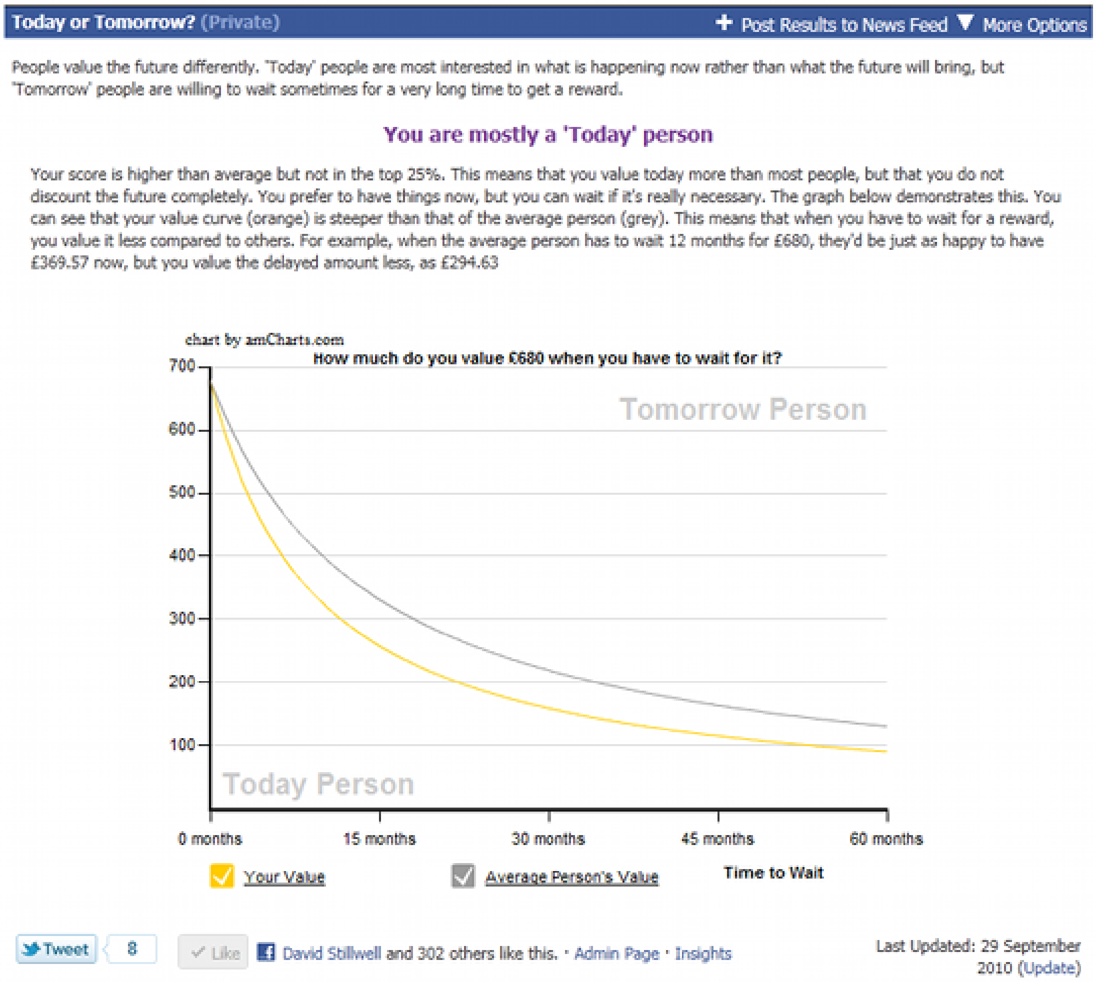

Figure 2.4: Screenshot of 'Today or Tomorrow' feedback on 'My Personality'. The lighter curve is the participant's personal DRD curve calculated from their DRD parameter. The darker curve is that of the average person.

The 'Today or Tomorrow' questionnaire received ethical approval from the University of Nottingham School of Psychology Ethics Committee. Respondents gave consent for their results to be used in an anonymised format for research, were able to stop at any time without pressure by simply closing the browser window, and were able to delete their information after completing the study through an automatic form mechanism within the 'My Personality' application.

Participants

From September 2010 to June 2011, 9454 unique users completed the full 'Today or Tomorrow' questionnaire, allowing their DRD parameter to be calculated. Respondents were tracked using their Facebook user ID, which allowed repeat respondents to be removed from the analysis. Of the 6549 who entered their gender, there were 2504 males (38%) and 4045 females (62%). Of the 6154 who entered their date of birth, the mean was 23.1 years old (SD = 9.1 years). Due to the random condition allocation process, there were 2964 participants in the ascending condition, 3195 in the descending condition and 3295 in the randomized condition. Table 2.3 shows the currencies chosen by participants and the demographic and smoking characteristics of users of each currency.

2.3.3 Results

In order to examine whether the exponential or hyperbolic discounting function fit individuals' DRD best, curves were fit for each participant using their 6 delays (between 1 week and 5 years). Sums of squared deviations were calculated for each participant's 6 indifference points from the curves generated by the discounting functions (Lagorio & Madden, 2005). It was found that the mean deviation from that predicted by the discounting function was smaller for the hyperbolic function (mean = 24150, SD = 32418, absolute difference = $155.40) than the exponential function (mean = 31558, SD = 41236, absolute difference = $177.65). For 86.8% of participants the hyperbolic function was a better fit to their DRD than the exponential function, which agrees with previous research (Bickel et al., 1999; Lagorio & Madden, 2005; Madden, Bickel, & Jacobs, 1999). There was no significant difference in the percentage of daily, non-daily, and non-smokers who were best described by the hyperbolic function.

Since the hyperbolic discount function best describes most individuals' DRD, for each participant estimated DRD log(k) parameters were averaged over the six delays which used $1000 as the delayed amount. To test whether self-reported smoking behaviour affected any effect of DRD order condition, participants were split into three groups: daily smokers, non-daily smokers, and non-smokers. The validity of separating the daily and non-daily smoking groups was confirmed by the CDS-5 scores of the two groups; daily smokers scored 16.03 (SD = 4.60, N = 1520), and non-daily smokers scored 7.06 (SD = 2.67, N = 502). In all analyses, the group sizes are unequal and we therefore used a Type III sum of squares ANOVA method in which each effect is controlled for every other effect, which ameliorates the problem.

Using the 9038 respondents who reported their smoking behaviour, a 3 x 3 between groups ANOVA examined whether the order that the items were presented in is related to delay discounting. Condition (ascending, descending, and randomized) and smoking status (daily smoker, N = 1592; non-daily smoker, N = 669; non-smoker, N = 6777) were the factors and log(k) was the dependent variable (Figure 2.5). A main effect of smoking status was found ($F(2,9038) = 58.98$, $MSE = .32$, $p < .001$), and planned t-tests showed that daily smokers had a steeper DRD curve than non-daily smokers ($t(5887) = 3.02$, $p = .003$), who in

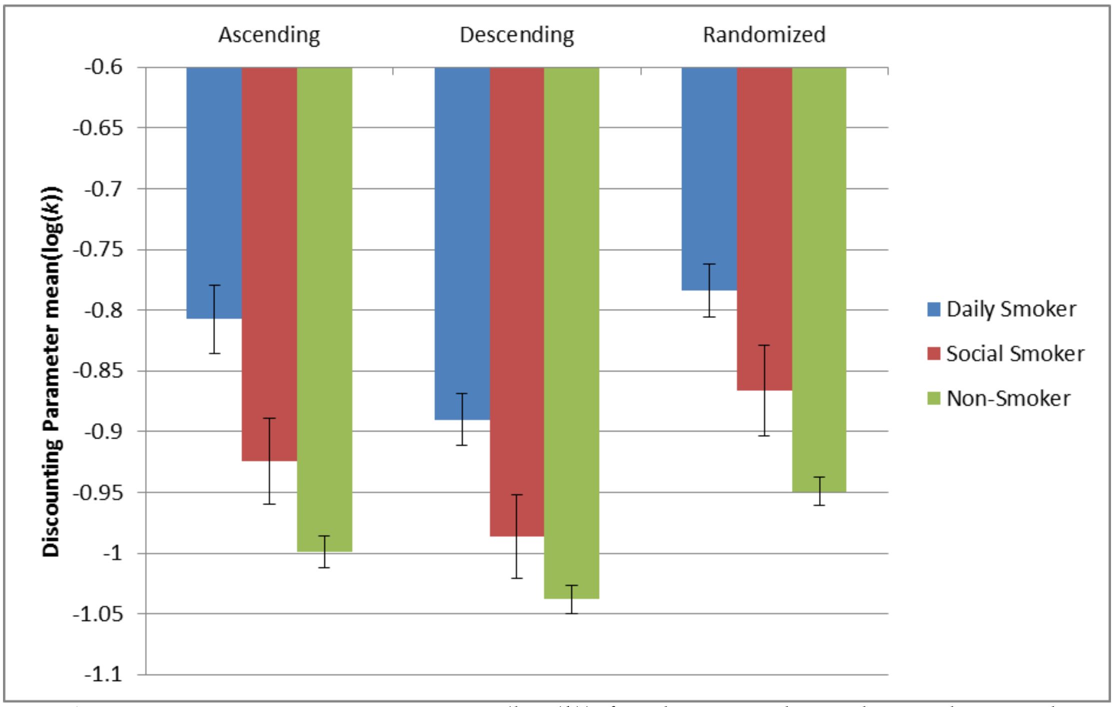

Figure 2.5: Mean DRD parameter (log(k)) for the ascending, descending and randomized conditions, separated by smoking status. The identical differences for each condition illustrate that the three smoking factors are affected equally by the three order conditions. Error bars are standard errors.

turn had a steeper DRD curve than non-smokers (t(7444) = 3.24, p = .001). A main effect of condition was found (F(2, 9038) = 11.20, MSE = .32, p < .001) and planned t-tests found that all three conditions had different parameters from one another (ascending was steeper than descending t(6157) = 2.48, p = .013; randomized was steeper than ascending t(6257) = 3.37, p = .001; randomized was steeper than descending t(6488) = 6.24, p < .001). However, no interaction was found between smoking status and condition (F(4, 9038) = .38, MSE = .32, p = .82);.the differences between smoking groups within each condition are the same, so the different smoking groups are affected equally by the DRD order condition.

To find out whether there were differences in the estimated DRD parameters

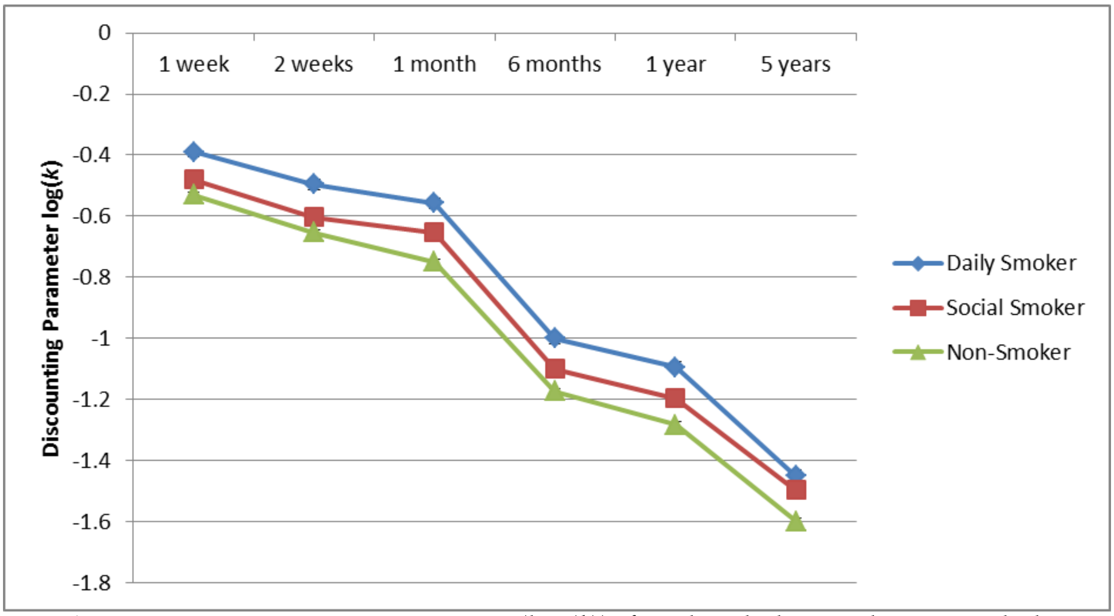

Figure 2.6: Mean DRD parameter (log(k)) for the daily smoker, non-daily smoker, and non-smoker conditions, at various delays. The similar curves indicate that the same differences between smoking statuses would be obtained no matter what delay length was used. Standard error bars are too small to be seen.

at different delay lengths, and whether any differences were affected by smoking status, a 3x3x6 mixed ANOVA was conducted with condition (ascending, descending, and randomized) and smoking status (daily smoker, non-daily smoker, non-smoker) as the between groups factors and delay length (1 week, 2 weeks, 1 month, 6 months, 1 year, 5 years) as the within groups factor. There was an effect of delay length ($F(5, 45145) = 3195.15$, $MSE = .21$, $p < .001$), such that shorter delays led to a steeper discounting parameter being estimated. The three-way interaction was not significant ($F(20, 45145) = 1.37$, $MSE = .21$, $p = .13$), but there was an interaction between delay length and smoking status ($F(10, 45145) = 2.16$, $MSE = .21$, $p = .017$). Figure 2.6 illustrates that for each of the smoking groups, the estimated DRD parameters become shallower as the delay length increased, but there was no clear pattern of changing differences between the smoking groups, indicating that similar results would be obtained no matter which delay period was used. Table 2.4 presents correlations between individuals' estimated log(k) parameters for various delays. It can be seen that as the time difference between two delays increases, the correlation between individuals' estimated DRD curves decreases (the smallest correlation is between 1 week and 5 years). This indicates that the delay length effect found in the ANOVA above is not simply due to a missing parameter in the hyperbolic discounting function which would affect the whole cohort of participants but not groups, and so we might still have expected to find differences between smoking groups. This nevertheless calls into question the estimation of a single DRD parameter across various delay lengths, as an individual's DRD parameter at one delay may not be strongly predictive of their parameter at another delay.

Table 2.4: Pearson correlations of individuals' estimated log(k) delay reward discounting parameters at various delays. All correlations are significant at p<.001

Delay	1 week	2 weeks	1 month	6 months	1 year
2 weeks	0.736				
1 month	0.647	0.736			
6 months	0.531	0.619	0.7		
1 year	0.439	0.529	0.65	0.767	
5 years	0.302	0.394	0.48	0.651	0.715

In order to investigate the effects of differing delayed rewards, the DRD parameters for $1000 in 1 month were compared to the parameters for $100 in 1 month. A 3 x 2 mixed ANOVA with smoking status (daily smoker, non-daily smoker, and non-smoker) as the between-groups factor, and delayed amount ($100, $1000) as the within-subjects factor found a main effect of delayed amount ($F(1, 8940) = 259.51$, MSE = .17, $p < .001$), but no interaction ($F(2, 8940) = 1.73$, MSE

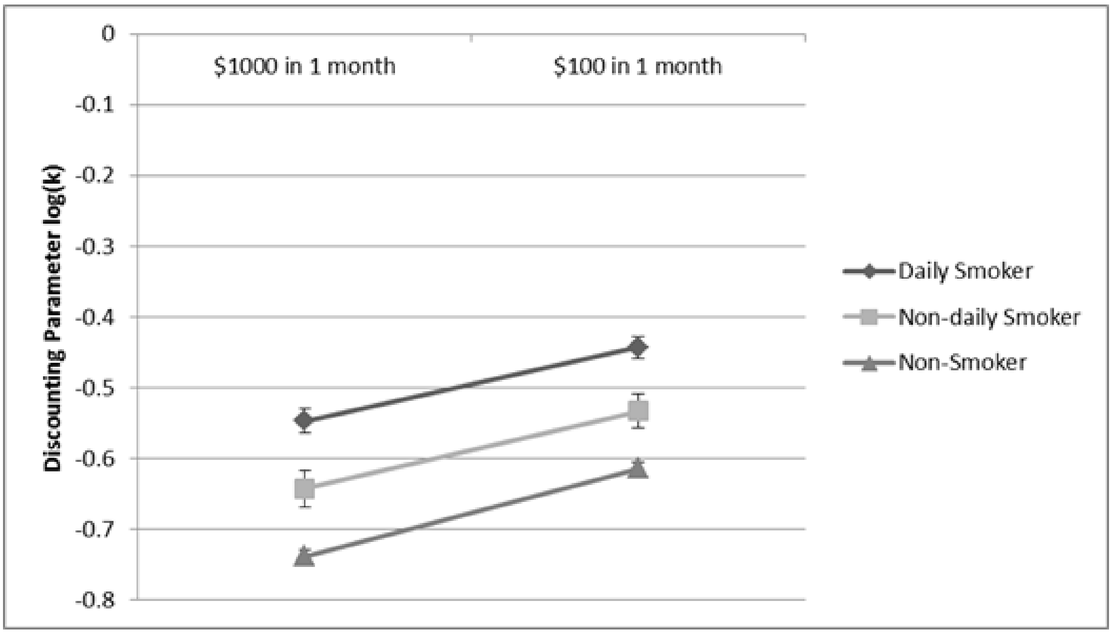

Figure 2.7: Mean DRD parameter (log(k)) for the $1000 and $100 delayed amounts in 1 month. The matching curves indicate that the three smoking statuses are identically affected by the two delayed amounts. Error bars are standard errors.

= .17, p = .18). These effects are shown in Figure 2.7 where the $100 delayed amount is discounted more steeply than the $1000 delayed amount and daily smokers discount more steeply than non-daily smokers. However, since the shape of the curves are identical the same differences between smoking groups would have been found no matter what delayed amount was used in the DRD procedure. This analysis relies on the subjective difference between $100 and $1000 being the same for all currencies. Despite our attempt to equalise them using the exchange rate, it may still be that our monetary values are perceived as larger or smaller for some currencies depending upon local prices. In order to control for this effect, we repeated the analysis using only respondents who picked the US Dollar as their currency. We again found a main effect of delayed amount ($F(1, 6362) = 159.44$, MSE = .17, $p < .001$) but no interaction ($F(2, 6362) = .23$, MSE = .17, $p = .79$), which corroborates our findings.

The daily, non-daily and non-smoker groups differed by demographic characteristics, particularly age; daily smokers were older than non-daily smokers, who were older than non-smokers. To control for any effects of age, the main analyses were repeated including age and gender as covariates (these are presented in Appendix B). The pattern of results replicated the current results, indicating that irrespective of which method of measuring discounting is used, the same effect of smoking group would be found.

2.3.4 Discussion

We found three methodological differences that may challenge the assumption of a single DRD parameter across all situations. The order that immediate rewards were presented in led to differing derived discounting parameters, with the randomized order demonstrating the steepest discounting, followed by the ascending order, followed by the descending order. Also, smaller delay lengths led to a steeper discounting function, and small rewards were discounted more than large rewards. Finally, daily smokers demonstrated the steepest discounting, followed by non-daily smokers, followed by non-smokers. Nevertheless, despite the large sample size, there were no interactions between smoking status and any of the above effects. This indicates that the difference in DRD between different smoking groups is remarkably robust to the method of measurement: it does not matter what order the items are presented in, what delayed time is used, or what delayed amount is used. The results of this study do however underline the importance of not directly comparing discounting parameters between studies, as the method of eliciting them is important.

Consistent with previous research (Bickel et al., 1999; Lagorio & Madden, 2005; Madden et al., 1999) it was found that a hyperbolic DRD curve fit better than an exponential curve. This implies that individuals have inconsistent time preferences and so may change their minds as the time before two rewards are available decreases. This has important implications for the external consequences that economists use when calculating the negative impacts of behaviours such as smoking (Gruber & Koszegi, 2001), such that an individual may not take their future self's preferences fully into account.

There are other variables that could affect DRD that were not studied here. It is known that individuals differentially discount different types of rewards (Tsukayama & Duckworth, 2010), so it might be that daily smokers discount cigarettes more than non-daily smokers. Additionally, this study does not shed light on whether the difference in DRD is due to an acute nicotine effect or whether it is a trait that may explain why people smoke, although there is evidence that steep DRD leads to smoking rather than the reverse (Audrain-McGovern et al., 2009).

Running an online study was successful in obtaining a large sample, and from countries that are underrepresented in traditional research (Gosling et al., 2004), but it did present a methodological problem specific to DRD in that the currencies used by participants differ. In order to standardize amounts the exchange rate was used, but it could be more appropriate to use a measure of purchasing power which takes account of unequal prices for the same goods between different countries. Additionally, the exchange rates on one date were used throughout the 10 month study, however, exchange rates fluctuated during this time, for example the US Dollar to British Pound exchange rate peaked at 0.6513 and floored at 0.5984. Nevertheless, whichever exchange rate was used, the calculated discounting rate would remain unaffected, and so discounting rates in our experiment are comparable between currencies. However, since we found that larger amounts are discounted less steeply, the purchasing power of a currency could affect how long people are willing to wait to receive it.

We found differences in DRD rate between users who picked different currencies. Indonesian Rupiah users had the steepest DRD functions whereas Canadian Dollar users were most self-controlled. Nevertheless, we view these differences extremely cautiously because as well as the differences in purchasing power between currencies we also cannot be sure that the selection biases for 'My Personality' users are the same for each country. For example, users of 'My Personality' in the United States are more likely to be sociologically representative than users in countries such as India where internet use is less common and where only a certain sociographic would use our English-language Facebook application. In conclusion, irrespective of how we measured DRD – varying the order of the items, the length of the delay and the magnitude of the delay – daily smokers had a shorter temporal horizon than non-daily smokers, who in turn had a shorter horizon than non-smokers. This is strong evidence that smoking is reliably related to DRD as a generalised behavioural preference function, rather than only a particular method of measuring discounting.

It also indicates to researchers that the method they use to measure discounting should depend upon their convenience. However, researchers should only compare discounting functions within studies, or effect sizes between studies, as small changes in DRD methodology can significantly affect the derived DRD parameter.

2.4 Chapter Discussion

This chapter examined how trait impulsivity, delay discounting and personality is related to addictive substance use and dependency. The general findings from Experiment 1 were that the conscientiousness personality trait did not explain unique variance beyond the impulsivity measures, however delay discounting and trait impulsivity were independent predictors. This conclusion was further supported by Experiment 2, which found that the methods used to measure delay discounting did not have an effect on its association with smoking. Since the methods were hypothesized to be related to BIS-11 impulsivity, specifically the motor impulsiveness subfactor, this indicates that trait impulsivity does explain why how an individual responds to the delay discounting task is related to smoking.

Importantly, Experiment 1 found that orthogonal personality traits were related to substance use and dependency. This indicates that in order to answer the question of what individual differences confer vulnerability to addiction, it is necessary to study individuals who have at least tried the addictive substance.

Otherwise, a researcher might be measuring traits which determine whether an individual tries a substance, rather than vulnerability to its addictive properties.

Between 4% (for alcohol dependency) and 21% (for drug dependency) of variance in substance use and dependency was accounted for by the psychological and demographic measures included. This indicates that there is substantial variance in addictive behaviour still to be accounted for, even though we used a broad set of predictors. The following chapter concentrates on smoking, as it is a well-defined addictive behaviour (as opposed to drug use) that is common enough to find in a student sample but rare enough that it does not represent a large proportion of the population (26% of myPersonality users smoked). In our models, 12.4% of variance in smoking use was accounted for, which was average among the addictive substances.

CHAPTER 3

Validity of the Simplified Harvard Game

3.1 Chapter Introduction

Impulsivity is a risk factor for drug dependence, but the mechanisms underpinning this association are unclear. One possibility is that impulsive individuals have an impaired capacity to represent abstract global consequences, allowing immediate rewards to dominate behaviour despite incurring a long term disadvantage. This could occur in either the the domain of losses or gains, or in both. There is also a question of whether participants learn a conscious understanding of the payoffs involved in their choices.

In this chapter, the validity of a task called the Simplified Harvard Game (SHG) is tested by examining whether an impairment in representing global consequences is related to smoking behaviour. The payoffs in the SHG are then reversed into the loss domain in the Credit Card Game, the results of which suggest that impulsive disinhibition may not be the reason that smoking behaviour is related to performance on the SHG. In the second half of the Chapter, two experiments probe participants' understanding of the SHG.

3.2 Experiment 3: Impulsivity and acute nicotine exposure are associated with discounting global consequences in the Harvard game

The following results are accepted for publication in Human Psychopharmacology: Clinical and Experimental (Hogarth, Stillwell & Tunney, in press).

3.2.1 Introduction

Although impulsivity is a risk factor for drug dependence, it is not clear how impulsivity influences decision making to drive addiction. The existing literature demonstrates that impulsivity and drug use are co-morbid (Stanford et al., 2009), and impulsivity is found both prior to drug exposure (Ersche et al., 2010; Tarter et al., 2004; Verdejo-Garcia, Perales, & Perez-Garcia, 2007) and to be augmented by drug exposure (Dallery & Locey, 2005; Heil et al., 2006; Setlow, Mendez, Mitchell, & Simon, 2009; Winstanley, 2007). The implication is that impulsivity and drug exposure are reciprocal, and individual susceptibility to this vicious circle arguably determines the longitudinal transition to clinical dependence (de Wit, 2009; Everitt et al., 2008; Perry & Carrol, 2008). What is unclear, is precisely what impairment in decision making allows drug use to dominate behaviour.

One proposal is that impulsivity reflects a dominance of sub-cortical stimulus-response reward learning systems over prefrontal cortical goal-directed learning systems (Everitt et al., 2008; Koob & Volkow, 2010; Matsuo et al., 2009; Nelson & Killcross, 2006; Wallis, 2007). On this view, reduced prefrontal cortical volume in impulsive individuals impairs their capacity to represent the abstract, long term consequences of behaviour, and thereby they fail to acquire complex goal-directed strategies that are ultimately more optimal. Rather, impulsive individuals acquire automatic stimulus-response (S-R) reflexes, such that reward-seeking behaviour is elicited directly by reward paired cues or contexts, without retrieving a representation of the consequences. This decoupling of reward seeking from intentional determination renders the behaviour less tractable to regulation by knowledge of global consequences.

The idea that retrieval of consequences is impaired in impulsive individuals is supported by several studies. First, impulsivity assessed by the Barratt Impulsivity Scale (BIS-11 Patton, Stanford, & Barratt, 1995) has been associated with an impairment in goal-directed control of action selection in the outcome devaluation paradigm (Hogarth, 2011). In this procedure, participants first acquired two instrumental responses for different rewarding outcomes (food and water), before one outcome was devalued by specific satiety. Finally, choice between the two responses was tested in extinction, such that any reduction in choice of the devaluated outcome must be mediated by integration of knowledge about the response-outcome (R-O) contingencies acquired in training with knowledge about the current incentive value of the outcome to determine the propensity to perform that response. The results showed that individuals high in BIS impulsivity showed a reduced devaluation effect in this extinction test despite reporting equivalent decline in the hedonic evaluation of the devalued outcome. The study thus demonstrated that impulsivity was associated with impairment in the retrieval of a representation of the current value of the consequences of behaviour to make adaptive choices.

In a related study, Hogarth (in press) found that the relationship between subjective craving to smoke and the number of puffs consumed in an ad libitum smoking session became progressively decoupled with increasing BIS impulsivity. In accordance with Tiffany (1990), therefore, this finding suggested that high BIS impulsivity conferred a propensity for drug taking behaviour to become automatic, in the sense of no longer being governed by subjective desire for the outcome. Again these results implicate a weakening of intentional determination of behaviour in impulsivity.

Several other studies provide converging evidence for this view. For instance, psychostimulant exposure increases impulsivity (Setlow et al., 2009) and impairs goal-directed control of reward seeking, suggesting a predominance of S-R learning (Jedynak, Uslaner, Esteban, & Robinson, 2007; Nelson & Killcross, 2006). Similarly, impulsivity increases the perseveration of drug self-administration that incurs shock punishment (Belin, Jonkman, Dickinson, Robbins, & Everitt, 2009; Economidou, Pelloux, Robbins, Dalley, & Everitt, 2009) non-reward (Diergaarde et al., 2008) and financial costs (Dallery & Raiff, 2007; Mueller et al., 2009), suggesting an autonomy of the behaviour from its consequences. Impulsivity also reduces sensitivity to error feedback signals (Groen et al., 2008; Potts, George, Martin, & Barratt, 2006) (see also: Franken, van Strien, Franzek, & de Wetering, 2007; Garavan & Stout, 2005; Hester, Simoes-Franklin, & Garavan, 2007), an impaired ability to modify established stimulus-response control of performance by instructions (Enticott, Ogloff, & Bradshaw, 2006) and a narrowing of the perceptual dimensions considered when making free categorisation decisions (Milton, Longmore, & Wills, 2008, ; Experiment 4). These data all converge on the view that impulsivity reflects a decoupling of action selection from control by the abstract representations of consequences, in favour of automatic S-R based control.

The Harvard Game

A rational decision-maker should dispassionately weigh up all future consequences of a decision before they decide whether to pursue a course of action or not. Nevertheless, experiments have long suggested that human decision-making does not take into account every eventuality and that there are situations where consequences are systematically ignored. One such situation is that of intrapersonal externalities (R. J. Herrnstein & Prelec, 1991), where changes to the utility of options available to the future self are not taken into account when making a decision, in the same way that externalities in economics refer to situations in which consequences to others are not taken into account when individuals make decisions. Intrapersonal externalities lead to an under-investment in activities that exhibit increasing average returns to the rate of consumption (for example, exercise becomes increasingly rewarding with increased practice, and also has positive effects on how rewarding other life activities are), and an over-investment in activities that exhibit decreasing average returns to the rate of consumption (for example, use of addictive substances becomes decreasingly rewarding with increased use, and also crucially has negative effects on how rewarding other life activities are).

In an experiment reported by R. J. Herrnstein et al. (1993), participants were presented with a repeated binary choice task using payoff schedules similar to those shown in Figure 3.1. The payoff from each choice was determined by the proportion of the previous ten choices that was allocated to each option, where choosing the optimal long-term option (maximize) would lead to a lower immediate payoff but would slightly increase the payoff from both options over the next ten trials, leading to a higher overall payoff. Conversely, choosing the optimal short-term option (meliorate) would lead to a higher immediate payoff but would decrease the payoff over the next ten trials. Consistently choosing the long-term option ultimately leads to the highest payoff, although on any single trial the short-term option would give the greatest number of points. The authors found that participants did not learn to optimise their behaviour, choosing instead the option with the greatest short-term payoff. R. Herrnstein and Vaughan (1980) suggested that choices are made according to the principle of melioration, in which the option with the highest immediate payoff is selected irrespective of the consequences for future payoffs.

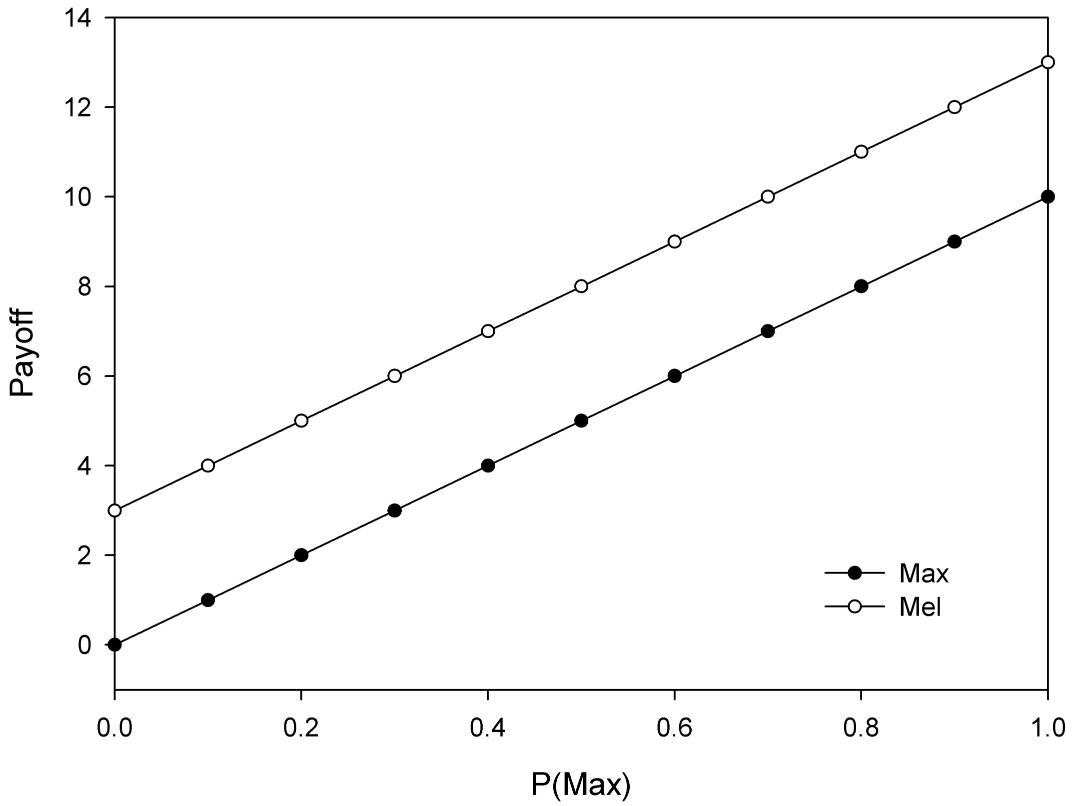

Figure 3.1: Example payoff schedules used in the task. P(max) represents the proportion of Max choices over the previous 10 trials.

It is not new to postulate that addicts do not fully take into account internalities, although past models have generally used hyperbolic time discounting as the theoretical reason for inconsistent preferences (Gruber & Koszegi, 2001). Nevertheless, evidence for the ecological validity of intrapersonal externalities was found by Heyman and Dunn (2002), who found that patients recovering in drug-clinics were more likely to choose sub-optimally than control patients, suggesting that addicts may be worse than others in taking into account the full consequences of their decisions. Further evidence for the link between intrapersonal externalities and addictions has come from neuropsychological research which found that the level of prefrontal brain activity is associated with performance in the task (Yarkoni, Braver, Gray, & Green, 2005; Yarkoni, Gray, et al., 2005). This is the same area that is implicated in studies using the Iowa Gambling Task (Bechara, Damasio, Damasio, & Anderson, 1994; Ernst et al., 2002), which is related to abuse of various substances, including alcohol and stimulants (Bechara et al., 2001). This suggests that intervention at this behavioural level could be effective in reducing addictive behaviour if a method were found to improve participants' decisions on tasks involving intrapersonal externalities.

The use of the Harvard game to study impulsivity is hampered by the fact that even in normal populations the majority of participants fail to optimise their performance and select the global option. In an attempt to correct this limitation, several experiments have sought to help participants to learn to optimise their performance without success (R. J. Herrnstein et al., 1993; Kudadjie-Gyamfi & Rachlin, 1996; Warry, Remington, & Sonuga-Barke, 1999). Stillwell and Tunney (2009) argued that the difficulty in learning the original Harvard game stems from the fact that on each trial participants received a single payoff based on both their current choice and history of choices, making it difficult to discriminate the differential outcomes produced by each choice. In the absence of knowledge that the local response was responsible for long term costs, participants would have no basis for reducing their choice of that response. To examine this possibility, Stillwell and Tunney (2009) designed the simplified Harvard game, where choosing the local option immediately reported both the monetary payoff and the costs to global utility signalled by a loss of ”game units”, that is, the number of trials that remain in which a choice could be made. With the local option now clearly associated with both a higher immediate payoff and higher overall costs (loss of game units), participants were more successful in learning to choose the global option across exposure to the task. In the present study, this simplified Harvard Game was employed to determine whether distribution of responses between the local and global option would show lawful variation with individual differences.

One question about the Harvard game concerns its association with delay discounting. In delay discounting, participants choose between a lower immediate reward versus a larger delayed reward, which is similar to the Harvard game where participants also choose between a lower immediate reward versus a larger delayed reward, although the length of the time delay is measured in seconds. The question, therefore, is what psychological processes are unique to each task. To address this question, participants in the current experiment also completed a delay discounting task, to determine whether there would be any dissociation with the Harvard game. On the basis of previous studies using this procedure, it was anticipated that delay discounting should be associated with BIS nonplanning impulsivity (de Wit et al., 2007) and higher levels of smoking dependency (Bickel et al., 1999; Heyman & Gibb, 2006; Johnson et al., 2007; Reynolds et al., 2007; Sweitzer et al., 2008).

Study of the Harvard Game has left a number of open questions as to whether the higher proportion of local responses in the drug user group was mediated by trait impulsivity, chronic drug exposure, or acute drug exposure. The objective of the current study was to try and untangle the impact of these variables on performance in the Harvard game. Participants in the current study were all smokers. The rationale for recruiting this group is that they show higher BIS impulsivity scores compared to the general population of non-smokers (Bickel et al., 1999; S. H. Mitchell, 1999, 2004; Reynolds et al., 2007; Skinner, Aubin, & Berlin, 2004; Spillane, Smith, & Kahler, 2010) and it was thought that individual differences in Harvard game and delay discounting performance might be more apparent at the high end of the impulsivity spectrum. In addition, this sample has the advantage of allowing us to untangle smoking related variables that might potentially impact on task performance (Dallery & Locey, 2005; Heyman & Gibb, 2006; Sweitzer et al., 2008). Specifically, by contrasting non-daily and daily smokers who had not smoked prior to the experiment, the effect of dependence level/chronic nicotine exposure could be tested. Further, by contrasting these two groups with daily smokers who had smoked prior to the experiment, the effect of acute nicotine exposure could be tested. The question at stake was whether performance in the Harvard game and delay discounting task would be differentially related to BIS impulsivity, prior smoking or daily smoking status. The primary hypothesis was that BIS impulsivity would confer a preference for the local choice in the Harvard game.

3.2.2 Method

Participants

Fifty-one healthy students from Nottingham University volunteered to take part in the experiment after being recruited by e-mails, posters and leaflets (54.9% female, mean age = 21.6 years, SD = 4.2 years). Only participants who self-identified as smokers for greater than 2 years were included in the study. Participants reported their age, gender, and then answered the questions (1) 'How often do you usually smoke', by ticking the phrase 'At least once per day' or 'Less than once per day' (which defined smoking status: Daily, non-daily), and (2) 'Did you smoke at any point prior to the experiment today', by ticking the answer 'yes' or 'no' (which defined prior smoking).

Participants then completed the BIS-11 impulsivity questionnaire (Patton et al., 1995). This questionnaire contains three subscales: (1) Motor impulsivity, e.g. "I do things without thinking", assesses propensity for action without thought, (2) nonplanning impulsivity, e.g. "I plan tasks carefully", assesses capacity for purposive future action, and (3) attentional impulsivity, e.g. "I don't pay attention", assesses capacity for sustained attention. These subscales were examined following analysis of the total BIS score to determine if there was any selectivity. The sample's mean BIS score was 67.2 (SD=10.5) which is slightly higher than a published norm of 62.3 (Stanford et al., 2009).

Procedure

Participants provided written informed consent to participate followed by the Harvard game, delay discounting task, and BIS-11 questionnaire in counterbalanced order. Finally, participants were debriefed and paid in accordance with their performance in the Harvard game (see below).

Harvard Game

Participants first read the following instructions on the screen:

> Your task is simple. You will have to repeatedly choose between two buttons, marked ”Left” and ”Right”. Simply click on a button with the mouse to register your choice. As a result of your choices you will win Points. After every choice you will be shown your Points from each choice as well as your cumulative Points. As you gain more Points, Pacman will eat more dots and get larger! However, choices will also use up Game Units. After every choice you will be shown the Game Units used up from each choice as well as your Game Units remaining. Once these have run out then the game is over. You will play the game 8 times. Try and beat your previous score in every game! Your payment from this experiment will be based on the number of points that you gain during the games. This will be calculated on the basis of 0.12p/point. This means you can earn between £5 and £7.20. After each game, you will be shown

> your current earnings so far. That's all there is to it - just try to win as many Points from the computer as you can before you run out of Game Units. Take as much time as you wish and please do not write anything down during the experiment.

Participants were then told that payment for the whole experiment would be based on the number of points that they gained and multiplied by 0.12p/point to calculate their earnings.

The Harvard task (Stillwell & Tunney, 2009) consisted of 8 games, each with 150 game units. This equated to between 53 and 150 trials per game, depending upon participants' choices. Figure 3.2 shows the main screen of the game. At the top centre of the screen a horizontal bar labelled 'Game Units Remaining' provided clear feedback concerning the payoff associated with each key. Above this, another horizontal bar labelled 'Points' depicted the total number of points gained during that game. This bar was based around an animated Pac-Man figure which moved from left to right and grew larger as the total number of Points increased. These were designed in order to increase the salience of feedback concerning the payoff (rewards and costs) associated with each key. Participants made their choices by selecting one button or another using the mouse. On each trial the points gained and game units lost information was updated, overwriting the outcome of the previous trial, so that participants would learn the differential payoff structure associated with left and right responses (counterbalanced between participants).

After each choice, both buttons were disabled for between 0.5 and 1.5 seconds in proportion to the number of game units lost on that trial. This variable delay was imposed to ensure that the total duration of each game was the same irrespective of the proportion of local or global response, given that choosing the local key used up more game units and so would otherwise result in faster termination of the experiment which may act as an incentive.

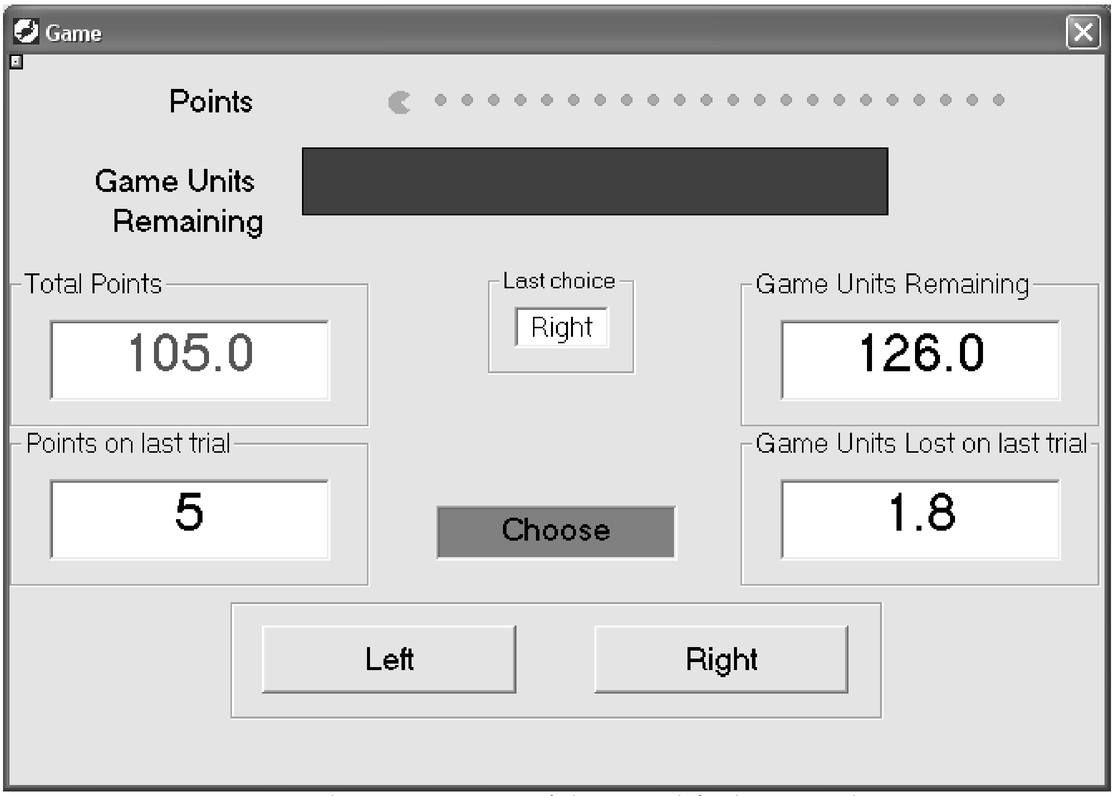

Figure 3.2: The main screen of the Simplified Harvard Game.

At the end of each game, a new screen summarised the total points gained during that game and the previous games (Figure 3.3). The top-centre of the screen displayed textually the total points gained during the game, how much those points were worth monetarily, and the maximum number of points that it was possible to gain during a game. Below this, a cartoon face was presented, contingent upon whether the participant gained more points during the recently completed game than the previous game. If the participant gained more points, the face smiled; if an equal number of points were earned, it was neutral; and if a lesser number of points were collected, it frowned. Beneath these, a bar chart graphically detailed the total points gained on that game and on all previous games.

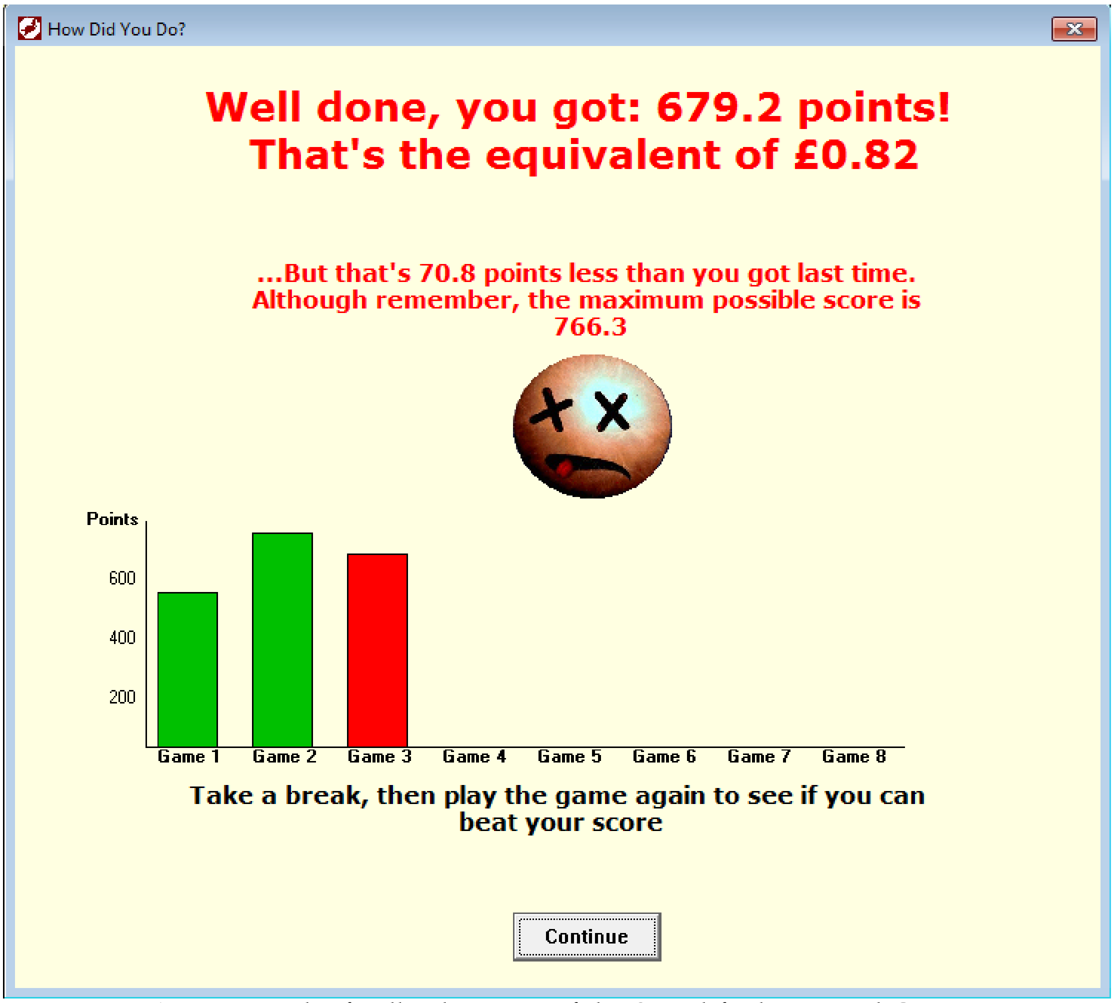

Figure 3.3: The feedback screen of the Simplified Harvard Game.

Payoff Schedules

Participants received points for every choice that they made, but lost game units. Choosing the local button returned 10 points per trial, but increased the rate at which game units were lost over the next 10 trials for both buttons. In contrast, choosing the global key returned 5 points per trial but cost fewer game units in the long term. The number of game units lost after each choice was determined by Equation 3.2.1.

$$GU = 1 + 2(P(b_{1-10})) \tag{3.2.1}$$

Where GU is the number of game units lost, and $P(b_{1-10})$ is the proportion of choices allocated to the local button over the preceding 10 trials. In essence, choosing the local button over the global button earned twice the number of points but used up three times as many game units over the next 10 trials, and so the lost game units represent a lost opportunity to earn money in the future.

By choosing the local key the participant is committed to losing an extra 0.2 game units for the next 10 trials, with each additional lost unit relinquishing the potential to consistently earn 5 points per game unit. This is why although for any single trial the participant would earn more by choosing the local option, in the long term it is optimal to choose the global option. Over the 150 game units of each game, consistently choosing the global button would return a cumulative payoff of 750 points whereas consistently choosing the local button would return a cumulative payoff of 530 points.

Each Simplified Harvard Game can be split into two blocks (Stillwell & Tunney, 2009); in one block participants should choose the option which maximises their long-term earnings, and in the other block near the end of each Game participants should choose the short-term option because there is no long-term remaining. In this experiment and the following one, only the analysis of the first block of trials is presented because the experimental hypothesis concerns self-control of melioration behaviour which is tested in the first block.

The average number of points gained by participants was 5188 (SD=420), leading to an average payment of £6.23 (maximum obtained: £7.12; minimum obtained: £5.26). The proportion of responses allocated to the global button was recorded across eight games. A low proportion of responses allocated to the global button indicates increased discounting of the overall payoff in favour of the immediately higher payoff.

Delay Discounting Task

Participants first read instructions on the screen: ″You are going to be asked to make choices between an immediate monetary reward and another monetary reward delayed by a certain length of time. Please use the mouse to select the option that you would prefer. Assume that the delayed amount will be adjusted for inflation. The rewards are hypothetical, however please make the choices as if you were given the choice for real. An example of the kind of choice you will be asked to make is between £50 now and £100 tomorrow. Take a second to consider which you would prefer, and then press the OK button to continue to the experiment.″ The experimenter then verbally explained that adjusting for inflation would mean that the delayed amount would still have the same purchasing power as having the same amount today. Participants were then repetitively presented with various monetary rewards on the left-hand side of the computer screen to be earned immediately versus £1000 delayed by various lengths of time on the right-hand side. Participants pressed the button below the option that they preferred with the mouse, which highlighted the button, and then chose 'New Choice' in the centre, which confirmed their choice and updated the screen with the next available choice. This meant that for each decision the participant's mouse started in the centre of the screen, and so it did not bias a particular choice. Participants could change their mind by choosing the other option before pressing 'New Choice'.

The delays and amounts used have been commonly used in previous delay discounting research (Bickel et al., 1999; Rachlin et al., 1991), although the amounts were denominated in British Pounds rather than US Dollars. The 27 immediate monetary rewards were £1000, £990, £960, £920, £850, £800, £750, £700, £650, £600, £550, £500, £450, £400, £350, £300, £250, £200, £150, £100, £80, £60, £40, £20, £10, £5, and ££1. The seven delays associated with the delayed £1000 alternative were 1 week, 2 weeks, 1 month, 6 months, 1 year, 5 years, and 25 years. Each delay was presented in a block, with the monetary amount randomised across trials. Moreover, the order of the delays was randomised as sequential ordering of delays can alter participants' preferences (Robles & Vargas, 2008). Between each delay a popup message alerted the participant to the change in delay.

To calculate a participant's delay discounting parameter (k), an indifference point was first established by calculating an average between the maximum immediate monetary amount chosen and the minimum delayed monetary amount chosen (Bickel et al., 1999). This value reflects the point at which the participant is indifferent between the immediate reward and the delayed reward. Next, non-linear regression was used to fit the seven indifference points from each participant to a hyperbolic function, according to the methodology established by Bickel et al. (1999). Since the distribution of k is often found to be non-normal (Bickel et al., 1999; Rachlin et al., 1991), the data were approximately normalised using the natural-log transformation. The k parameter reflects the steepness of the discount curve, whereby greater k values reflect a sharper decline in the subjective value of money as the delay to obtain that money increases.

3.2.3 Results

Participants

Of the 51 participants, there were 16 non-daily smokers and 11 daily smokers who had not smoked prior to the experiment, and 24 daily smokers who had smoked prior to the experiment. This smoker group variable (3) was analysed to determine the impact of dependence level (daily vs. non-daily) and acute nicotine exposure (prior vs. no prior smoking). Total BIS scores were entered into the same ANCOVA as a continuous variable. Subsequent analysis was undertaken to examine BIS sub-scales and potential confounding variables.

Harvard Game

The proportion of global responses obtained in the eight games were examined with smoker group (3) as a categorical variable and BIS impulsivity as a covariate in ANCOVA. This analysis yielded a main effect of game, $F(7,329) = 2.10$, $p < .05$, indicating that participants acquired an increased choice of the global response across training on the Harvard game. There was no interaction between game and smoker group or BIS impulsivity, $Fs < 1$, indicating that these groups acquired a preference for the global response at a comparable rate.

There was a main effect of BIS impulsivity, $F(1,47) = 6.70$, $p < .02$, shown in Figure 2A/B. This effect could not be explained by age, because multiple regression with the proportion of global responses as the dependent variable, and BIS impulsivity and age as predictors indicated that BIS impulsivity served as an independent predictor in this model, $t = -2.18$, $p < .05$, whereas age did not, $t = -.56$, $p = .58$. Moreover, there was no significant difference in BIS impulsivity between the smokers groups, $F < 1$, suggesting the effect shown in Figure 3.4A/B could not be attributed to either differential tobacco dependence or prior smoking across levels of BIS. Finally, separate assessment of the three BIS scales (as continuous variables) indicated that the greater proportion of global responses was predicted more strongly by the BIS motor scale, $F(1,49) = 5.30$, $p < .03$, compared to the nonplanning, $F(1,49) = 1.84$, $p = .18$, or attentional scales, $F(1,49) = 2.51$, $p = .12$. However, none of the three scales emerged as an independent predictor of global responses in multiple regression, $ts < -1.06$, $ps > .28$, suggesting that these scales are too highly interrelated to partial out their independent effect. Overall, therefore, the data suggest that BIS impulsivity is associated with discounting global consequences in the Harvard game.

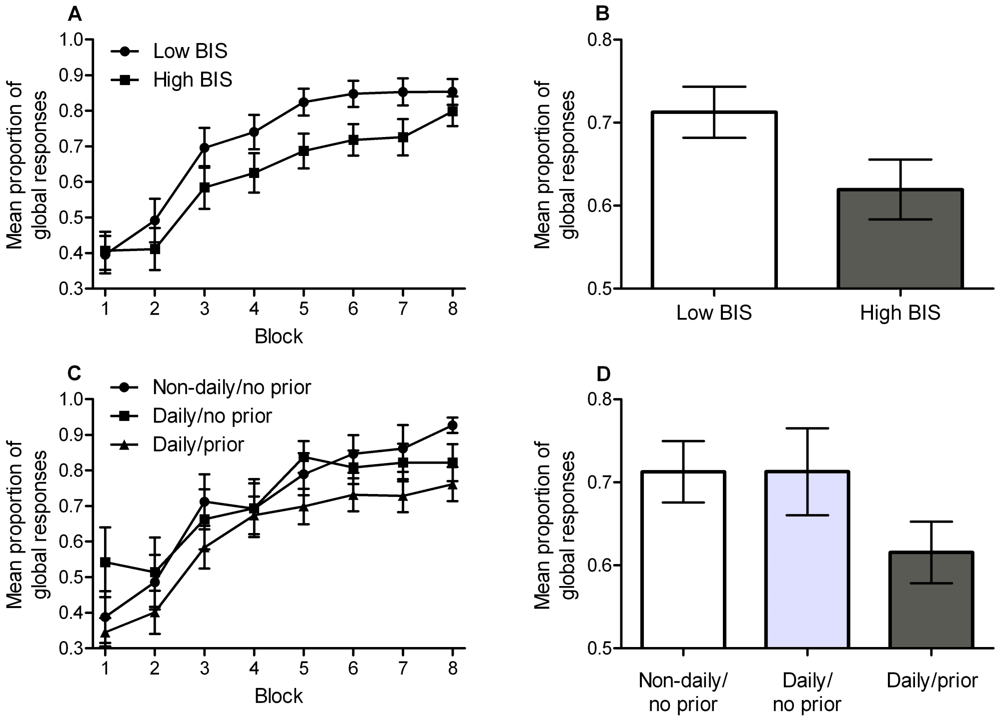

Figure 3.4: Proportion of global responses in the Harvard game. Data are shown for each of the eight games in the impulsive (A) and smoker groups (C), as well as the average across the eight games in the impulsive (B) and smoker groups (D).

The initial ANCOVA also yielded a main effect of smoker group, $F(2,47) = 3.22$, $p < .05$, shown in Figure 3.4C/D. Exploration of this effect indicated that there was no reliable difference between daily and non-daily smokers who had not smoked prior to the experiment, $F < 1$, whereas these two group collapsed chose the global response significantly more frequently than daily smokers who had smoked prior to the experiment, $F(1,49) = 4.24$, $p < .05$. This effect of prior smoking on global responses remained significant when age and BIS impulsivity were included as covariates, $F(1,47) = 6.12$, $p < .02$, indicating that these potentially confounding variables could not explain the effect. Therefore, these analyses suggest that acute smoking prior to the experiment increased discounting of global consequences in the Harvard game.

Delay Discounting Task

The delay discounting parameters (k) were natural-log transformation entered into ANCOVA as the dependent variable, with smoker group (3) as a categorical independent variable and BIS impulsivity as a covariate. This analysis yielded no significant main effects of smoker group, $F < 1$, or BIS impulsivity, $F(1,47) = 1.39$, $p = .24$.

3.2.4 Discussion

The current experiment found that preference for the local choice in the Harvard game was associated with BIS impulsivity and acute prior smoking in a sample of 51 young adult smokers. The results showed that the proportion of global responses in the Harvard game increased with training, suggesting participants acquired knowledge about the higher overall payoff of the global choice with training, and there were no group differences in the acquisition of this knowledge. However, BIS impulsivity and acute prior smoking were associated with a greater overall proportion of local responses. These two group effects were independent in that they were not confounded, and could not be attributed to the other potentially confounding variables, age or daily smoking status. Overall, therefore, these data support the view that BIS impulsivity and acute nicotine exposure reduced the impact of global consequences on choice in favour of the immediate higher payoff.

It is difficult to pinpoint the psychological process underpinning the higher proportion of local responses in the high BIS impulsive group and prior smokers. Arguably, these effects were not due to a failure to acquire knowledge of the contingency between the global response and the higher overall payoff, because the groups showed equivalent acquisition of the global choice over training. Rather, these effects were presumably driven by differential weighting of the consequences produced by each choice. Figure 3.2 shows that there were several sources of feedback following a choice: Points won on that trial, total points won in that game, game units lost on that trial, and total game units remaining in that game. Accordingly, the high BIS group and prior smokers may have been hypersensitive to the greater increment in points won following a local choice, hyposensitive to greater loss of game units following the local choice, hyposensitive to the greater total points won following global choices, and hyposensitive to the slower rate of game units lost following global choices. Finally, the high BIS impulsive group and prior smokers may have been more intolerant of delay. This argument can be made because the local option used up game points more quickly, thus participants may have believed that the local option brought the experiment to a close more quickly and thus acted as an incentive to choose the local option (despite the lost overall earnings). Although this possibility was mitigated by adding a proportionate delay at the end of each trial, such that the experimental duration was matched irrespective of the portion of local/global responses, nevertheless, participants may have been differentially sensitive to this feedback.

Although the specific impairment in impulsivity cannot be isolated with the Harvard game, the finding of reduced global choice nevertheless converges on the view that the core impairment in impulsivity is a weaker capacity to retrieve an abstract global representation of consequences to control goal-directed action selection. As noted in the introduction, converging support for this claim comes from the finding that impulsivity is associated with a decoupling of intention and action (Hogarth, 2011; Hogarth, Chase, & Baess, 2012), perseveration of drug self-administration under shock punishment (Belin et al., 2009; Economidou et al., 2009) extinction (Diergaarde et al., 2008) and financial costs (Dallery & Raiff, 2007; Mueller et al., 2009), reduced sensitivity to error feedback (Groen et al., 2008; Potts et al., 2006), impaired top down regulation of prepotent responding (Enticott et al., 2006), and restricted use of perceptual dimensions in free categorisation (Milton et al., 2008). Moreover, Impulsivity has been associated with reduced volume in the prefrontal cortex (Matsuo et al., 2009), and this region is known to synthesise multiple dimensions of outcome expectancies (Wallis, 2007), and play a role in goal-directed action (de Wit, 2009; Valentin, Dickinson, & O'Doherty, 2007). Finally, impulsivity is believed to facilitate the transition to automatic drug use behaviour by weakening the control of self-administration by a representation of the consequences (Everitt et al., 2008, 2001; R. Z. Goldstein et al., 2009; Koob & Volkow, 2010). The current study supports this view by demonstrating that impulsivity is associated with impaired use of knowledge about the global consequences of behaviour to determine choice.

The other key finding was that prior smoking was associated with increased local choice. This finding corresponds with a number of published studies. Specifically, greater delay discounting, which is arguably related to Harvard game performance (R. J. Herrnstein et al., 1993; R. J. Herrnstein & Prelec, 1991), has been found to be greater in smokers who report a shorter latency to smoke in the morning (Sweitzer et al., 2008), in daily versus non-daily smokers (Heyman & Gibb, 2006) and in smokers versus non-smokers (Bickel et al., 1999; Johnson et al., 2007; Reynolds et al., 2007). Although these groups tend to conflate dependence status, chronic and acute nicotine exposure, our data go some way to untangling these variables. Specifically, daily and non-daily smokers did not differ in global choice, suggesting that dependence status and chronic nicotine exposure are not important for this decision bias. Rather, only smoking prior to the study was related to reduced proportion of global choices. Of course, individuals who smoked prior to the study are arguably more dependent and have received more chronic nicotine exposure, but if these two factors were important, a difference would have been expected in the contrast between the non-daily and daily/no prior smoking group, which was not found. Most decisively, however, animal studies have shown that acute nicotine administration does increase delay discounting (Dallery & Locey, 2005; Locey & Dallery, 2009) and a similar effect has been found with cocaine administration (Mendez et al., 2010) see also (Heil et al., 2006; Setlow et al., 2009; Tsutsui-Kimura et al., 2010). Thus, acute drug administration does induce a preference for short term gains in the face of global costs. Arguably, this acute drug induced preference for short term rewards may underpin the drug induced priming of drug self-administration behaviour (Fuchs, Evans, Parker, & See, 2004; Leri & Stewart, 2001; Schmidt, Anderson, Famous, Kumaresan, & Pierce, 2005). On this view, the loss of a representation of global consequences following acute drug administration disinhibits automatic control of drug self-administration by drug paired stimuli. Our effect of acute smoking on local choice is at least consistent with this account.

The current study sheds light on the only other published study of individual differences in the Harvard game (Heyman & Dunn, 2002). These researchers found that illicit drug users (n=22) showed a greater proportion of local responses in the Harvard game compared to non-drug users (n=21), supporting a role for this decision impairment in drug dependence. However, they found no difference in BIS impulsivity between these two groups, and BIS did not correlate with the proportion of global responses, in contrast to the current study. It is difficult to account for these discrepant findings, given the multiple differences between the studies, but key sources of variance worth considering are the demographic status of the participants and the instantiation of the Harvard game. More decisively, the null group effect in impulsivity found by Heyman and Dunn (2002) suggests that acute drug use may have instead driven the increased proportion of local choices in the drug-users. The upshot of this analysis is that both impulsive trait and acute drug state must be taken into account when considering individual differences in Harvard game performance.

One final question is why the delay discounting task was insensitive to group variables compared to the Harvard game. These null effects contradict studies which have found differential delay discounting as a function of BIS impulsivity (de Wit, 2009) and smoker group (Bickel et al., 1999; Heyman & Gibb, 2006; Reynolds et al., 2007; Sweitzer et al., 2008). One explanation for the differential sensitivity of the two tasks is suggested by Heyman and Gibb (2006). In this study, daily and non-daily smokers (n=71) completed a delay discounting task in which the monetary rewards were hypothetical (as in the current delay discounting task), or actual, in that participants kept the money they earned (as in the current Harvard game). Only in the actual reward condition did daily smokers show greater delay discounting than non-daily smokers, whereas no group effects were seen in the hypothetical reward condition. Impulsive decision making for actual rewards seems to be more sensitive to the effects of acute nicotine exposure than hypothetical rewards; however, Sweitzer et al. (2008) found an effect of morning smoking on hypothetical delay discounting with a much larger sample (n=710) and another study de Wit et al. (2007) found an effect of BIS impulsivity on delay discounting in a much larger sample (n=606). Thus, although impulsive choice for actual rewards may be more sensitive to group variables than choice for hypothetical rewards, groups effects can be detected with hypothetical rewards given sufficient power.

To conclude, the study found that BIS impulsivity and acute nicotine exposure were associated with discounting long term consequences in favour of short term gains in the Harvard game. These data accord with the proposed role of impulsivity in the aetiology of drug dependence, wherein impaired retrieval of a mental representation of global consequences in the service of intentional action favours predominance of automatic control over behaviour. Similarly, the acute effect of nicotine exposure on increasing local choice accords with the view that acute drug priming of self-administration is mediated by loss of intentional regulation by knowledge of global consequences disinhibiting automatic control by drug paired stimuli. Overall, these data favour the view that individual differences in the capacity to represent abstract consequences is a key process underpinning the aetiology of drug dependence.

3.3 Experiment 4: Relationship Between the Credit Card Game, Smoking and Impulsivity

3.3.1 Introduction

The subjective value of losses and gains are asymmetrical in humans. When individuals are presented with two pairs of probabilistically identical gamples, one of which is presented as a gain and the other as a loss, participants systematically switch choices between the two pairs, from preferring a high probability of a small gain to preferring a low probability of a big loss. This inconsistency demonstrates that the framing of the gambles as either gains or losses affects subjective values (Tversky & Kahneman, 1981).

Delayed losses have also been compared with delayed gains in delay discounting experiments. Whereas in reward discounting, participants consistently discount the rewards as they are delayed, in loss discounting S. H. Mitchell and Wilson (2010) found that a "substantial number" of participants reverse-discounted losses, such that they would rather incur the cost now rather than delay it. However, S. H. Mitchell and Wilson (2010) and J. G. Murphy, Vuchinich, and Simpson (2001) both found that those who steeply discounted gains also steeply discounted losses, indicating that the individual difference is inconsistent across both frames.

The Simplified Harvard Game was modified from one of gaining rewards to one of avoiding costs. Practically, this meant that participants were given a high number of points to begin with but they lost points each time they made a choice. Choosing one button meant that the participant lost a small number of points after each choice, but the number of choices remaining until the end of the experiment reduced slowly. Choosing the other button meant that the participant lost a high number of points after each choice, but the number of choices remaining until the end of the experiment reduced at a faster rate, such that the participant had to make fewer choices overall and so ultimately kept more points.

The choice between a big charge and smaller payments over a long time span mirrors the decision that borrowers have to make when their credit card statement arrives in the post. Borrowers can either pay back the principle quickly, incurring a large immediate expense, or they can pay the minimum amount each month, but their debt lasts longer and ultimately costs more in interest charges. The loss framed Simplified Harvard Game is therefore named the Credit Card Game.

Having such a clear analogy for the task is useful, since paying off a credit card each month vs. carrying a balance forward has been associated with both self-control and impulsivity in college students (N=165; Mansfield, Pinto, & Parente, 2003), although overall level of credit card debt was not related to impulsivity (Norvilitis, Szablicki, & Wilson, 2003) in college students (N=227). This suggests that how one uses a credit card, rather than overall economic circumstance, is related to impulsivity.

The independent variables assessed in this experiment were the same as in Experiment 3: smoking behaviour and impulsivity. Two competing hypotheses are to be examined. If impulsivity determines performance on the task, in the sense that participants choose the immediately most rewarding option despite knowing that the other option is better in the longer term, then since losses subjectively loom larger than gains (Tversky & Kahneman, 1981) all participants should perform poorly and daily smokers should perform particularly poorly. On the other hand, the choice to smoke overweights short-term gains over long-term losses, so daily smokers' poor performance on the task relative to non-daily smokers may reflect difficulty learning about small gains relative to large losses. As the Credit Card Game gives users a choice between small losses relative to large gains, if daily smokers are not worse than non-daily smokers then this suggests that the task is revealing a deficit learning about gains.

3.3.2 Method

The method was the same as in Experiment 3, except that the Credit Card Game used reversed payoff schedules where after each choice participants lost points rather than gained them (explained below).

Participants

Forty-four healthy students from Nottingham University volunteered to take part in the experiment after being recruited by e-mails, posters and leaflets (61.4% female, mean age = 20.9 years (SD=2.4 years)). Only participants who self-identified as smokers for greater than 2 years were included in the study. Participants reported their age, gender, and then answered the questions (1) 'How often do you usually smoke', by ticking the phrase 'At least once per day' or 'Less than once per day' (which defined smoking status: Daily, non-daily), and (2) 'Did you smoke at any point prior to the experiment today', by ticking the answer 'yes' or 'no' (which defined prior smoking). As in Experiment 3, participants also answered the Barratt Impulsivity Scale (BIS-11; Patton et al., 1995). The sample's mean BIS score was 67.6 (SD=10.0) which is slightly higher than a published norm of 62.3 (Stanford et al., 2009).

Credit Card Game

Participants read the following instructions on the screen (text in bold has been added later to emphasise the main text changed from Experiment 3):

> "Your task is simple. You will have to repeatedly choose between two buttons, marked "Left" and "Right". Simply click on a button with the mouse to register your choice.
>
> **At the beginning of the experiment you will be given 1250 Points, but as a result of your choices you will lose points. After every choice you will be shown the points you lost as well as the number of points remaining.**
>
> Choices will also use up Game Units. After every choice you will be shown the Game Units used up from that choice as well as your Game Units remaining. Once these have run out then that game is over and you keep whatever number of points you have still remaining.
>
> You will play the game 8 times. Try and keep more and more points after every game!
>
> Your payment from this experiment will be based on the number of points that you manage to keep during the games. This will be calculated on the basis of 0.12p/point. This means you can earn between £5 and £7.35. After each game, you will be shown your current earnings so far.

> That's all there is to it – just try to keep as many points as you can from the computer by the end of each game. Take as much time as you wish but please do not write anything down during the experiment."

The Credit Card Game (Stillwell & Tunney, 2009) consisted of 8 games, each with 150 game units. This equated to between 53 and 150 trials per game, depending upon participants' choices. Figure 3.2 shows the main screen of the game. As can be seen, the game is created to mirror as closely as possible the Harvard Game used in Experiment 3. The differences were that participants lost points on each trial, and the animated PacMan figure moved from right to left, eating points as it went.

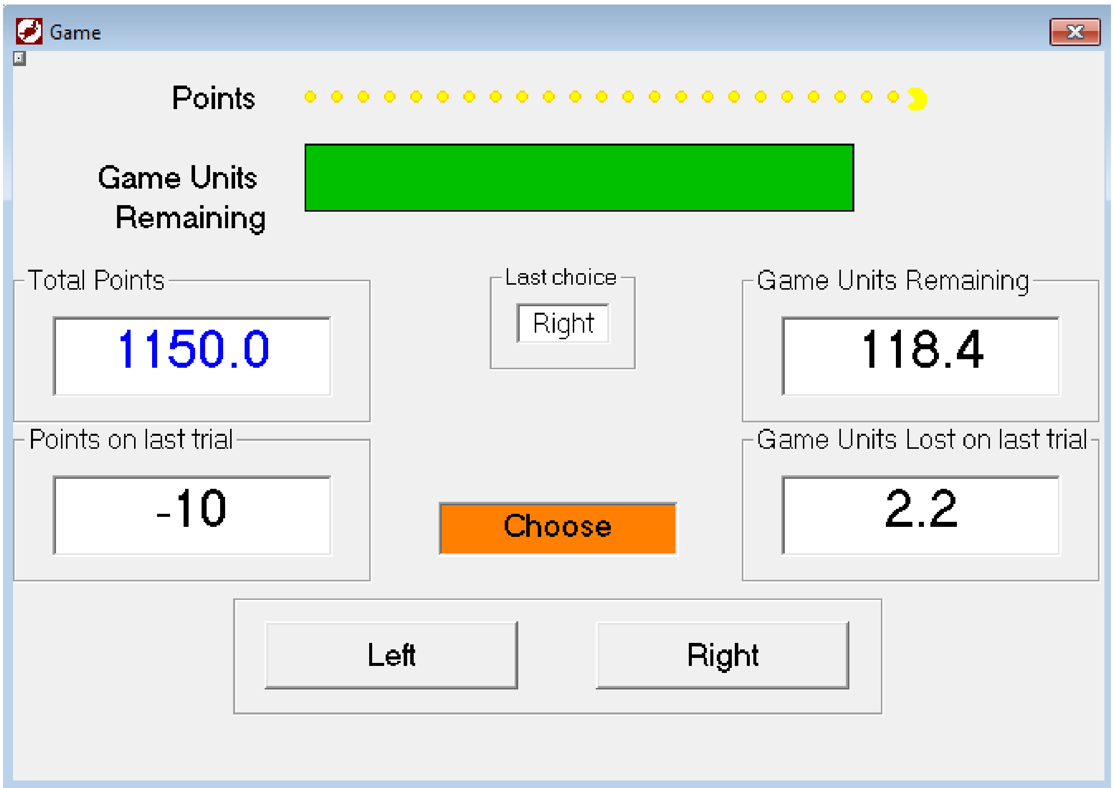

Figure 3.5: The main screen of the Credit Card Game.

Payoff Schedules

Participants lost both game units and points for every choice that they made. Choosing the local button lost 5 points per trial, but decreased the rate at which game units were lost over the next 10 trials for both buttons. In contrast, choosing the global key lost 10 points per trial but more game units were lost in the long term. The number of game units lost after each choice was determined by Equation 3.3.1.

$$GU = 3 - 2(P(b_{1-10})) \tag{3.3.1}$$

Where GU is the number of game units lost, and $P(b_{1-10})$ is the proportion of choices allocated to the local button over the preceding 10 trials. In essence, choosing the local button over the global button lost half the number of points but used up a third as many game units over the next 10 trials, and so the remaining game units represent a commitment to losing further points in the future.

By choosing the local key the participant is committed to losing 0.2 fewer game units for the next 10 trials, with each additional unit remaining representing a commitment to losing a further 5 points per game unit. This is why although for any single trial the participant would lose less by choosing the local option, in the long term it is optimal to choose the global option. Over the 150 game units of each game, consistently choosing the global button would return a cumulative payoff of 750 points whereas consistently choosing the local button would return a cumulative payoff of 545 points.

The average number of points gained by participants was 5157 (SD=308), leading to an average payment of £6.19 (maximum obtained: £6.93; minimum obtained: £5.43). These averages are remarkably close to those in Experiment 3, indicating that the two tasks are of similar difficulty. The proportion of responses allocated to the global button was recorded across eight games. A low proportion of responses allocated to the global button indicates increased discounting of the overall payoff in favour of the immediately smaller loss.

3.3.3 Results

Participants

Of the 44 participants, there were 17 non-daily and 7 daily smokers who had not smoked prior to the experiment, and 20 daily smokers who had smoked prior to the experiment. As in Experiment 3, the smoker group variable was analysed to determine if dependence level and acute nicotine exposure were related to performance in the task.

Total BIS scores were entered into the ANCOVA as a continuous variable. The addition of a covariate should not change the main effect of a within-subjects factor because the two are independent. If a student scores 20% for task 1 and 30% for task two, this should be unaffected by the student's age, which is the same. Similarly, the means for the whole group across the two tasks should not depend on each participant's age. However, the covariate may interact with the repeated measure. The students' ages may determine whether they are better at task 1 or task 2. This leads to an adjustment in the sums of square which makes the main effect of the repeated measure weaker (Delaney & Maxwell, 1981). One solution is to mean-centre the covariate prior to running the ANCOVA (Delaney & Maxwell, 1981), which is employed in the following results.

Credit Card Game

The proportion of global responses obtained in the eight games were examined with smoker group (3) as a categorical variable and BIS impulsivity as a covariate in the ANCOVA.

This analysis yielded a main effect of game, $F(7,280) = 41.02$, $p < .001$, indicating that participants acquired an increased choice of the global response across training on the Harvard game. There was no interaction between game and smoker group or BIS impulsivity, $Fs < 1$, indicating that these groups acquired a preference for the global response at a comparable rate.

Unlike the previous experiment, there was no main effect of BIS impulsivity, $F(1,40) = 1.21$, $p = .28$. The initial ANCOVA also found no main effect of smoker group, $F(2,40) = .18$, $p = .83$.

In conclusion, there is no evidence that BIS impulsivity or acute smoking prior to the experiment affected performance in the Credit Card Game.

Delay Discounting

The delay discounting parameters (k) were natural-log transformation entered into ANCOVA as the dependent variable, with smoker group (3) as a categorical independent variable and BIS impulsivity as a continuous independent variable. This analysis yielded no significant main effects of smoker group, $F(2,40) = 1.6$, $p=.21$, or BIS impulsivity, $F(1,40) = 2.13$, $p = .15$.

3.3.4 Discussion

Preference for the local choice in the Credit Card Game was not associated with the smoking variables of dependency or acute nicotine influence, or with trait impulsivity as measured by the BIS. This is unlikely to be due to the task difficulty, as participants learned to improve their performance on the task as the experiment progressed, and also made a similar number of global choices as in the Simplified Harvard Game in Experiment 3. These results contrast with those of the previous experiment, where it was found that acute nicotine influence was associated with the number of global choices in the Simplified Harvard Game.

Of the two hypotheses, the results support the contention that the poor performance by daily smokers on the Simplified Harvard Game is due to a deficit learning about small gains relative to large losses. In the Credit Card Game, participants choose between small losses and large gains, which does not expose this deficit. The results do not support the hypothesis that the poor performance by daily smokers in Experiment 3 was due to impulsive disinhibition, since in the Credit Card Game the subjective difference between the two choices would be larger than in the Simplified Harvard Game because losses loom larger than gains Tversky and Kahneman (1981). If individuals choose based on impulsivity then this should have caused poorer performance by all individuals on the task, and this should particularly have affected smokers' performance. Nevertheless, this is far from strong evidence, since the statistical power is not high, so it would be ideal to test it on a larger sample.

In conclusion, in the Credit Card Game the difficulty was similar to the Simplified Harvard Game, but there was no difference between daily smokers and non-daily smokers. This suggests that smokers perform poorly in tasks with intrapersonal externalities due to diffculty learning about small gains rather than disinhibition of impulsivity.

3.4 Experiment 5: Participants' Understanding of the Simplified Harvard Game

The following results and those presented in the next experiment have been published in Judgment and Decision Making (Stillwell & Tunney, 2012).

3.4.1 Introduction

A series of laboratory experiments has tried without success to guide participants to choose optimally in the Harvard Game. Warry et al. (1999) attempted to reduce the motivation for participants to choose impulsively by reducing the immediate differential between the two options. They found that this decreased the propensity for participants to choose sub-optimally, however by the end of their experiment participants were still choosing around chance levels, and the authors noted that extrapolation of their data suggested that participants would reach asymptote at a level that was non-optimal. Two experiments have also attempted to guide participants' explicit understanding of the payoff schedules by providing a fairly explicit hint on how participants could maximize their payoffs. R. J. Herrnstein et al. (1993, ; Experiment B) found that choices were only briefly improved by the hint but soon returned to sub-optimal levels. KudadjieGyamfi and Rachlin (1996) provided a similar hint, but found no corresponding improvement at all. Nevertheless, Tunney and Shanks (2002) showed that participants could overcome sub-optimal behaviour, as long as they were given regular feedback about how their behaviour compared to the optimal outcome, and they were given around 1,000 trials to learn the schedules. This suggests that suboptimal choices in the Harvard Game may not be a stable decision-making bias or a failure of impulse control, but rather due to a failure to fully learn the payoff schedules.

Normally in experiments studying intrapersonal externalities participants' choices either affect the number of points gained (e.g. Yarkoni, Braver, et al., 2005) or the number of choices remaining until the end of the experiment (e.g. R. J. Herrnstein et al., 1993). However Stillwell and Tunney (2009) modified the schedules so that both the number of points gained and the number of choices remaining until the end of the experiment were affected by participants' choices. This allowed the two outcomes from each decision to be separated so that the immediate effects were visible through the number of points gained on each trial, and the number of choices remaining until the end of the experiment decreased at differing rates depending upon the participant's history of choices. In other words, choosing myopically led to earning high payoffs through the experiment, but ultimately the experiment ended prematurely and the participant lost the opportunity to earn further payoffs. Separating the consequences from each decision made the outcomes of each decision more easily discernable, and resulted in participants learning to choose optimally much earlier than had previously been demonstrated. This also suggests that suboptimal behaviour in the task may not be a failure of impulse control, but rather a failure to fully understand the payoff schedules.

In nature the outcomes from choices that are made may not be so easily divided into separate simple categories. So, if the results from laboratory intrapersonal externality experiments are to be useful in understanding the suboptimal decision-making that occurs in addictions, the process whereby participants learn to choose optimally in the simpler version needs to be understood. One process could be the result of conscious insight into the payoff schedules that participants are able to report. This mirrors research into the Iowa Gambling task which found that participants were able to explicitly report their understanding of the task (Maia & McClelland, 2004). The present experiments attempted to test explicit knowledge of the payoff schedules, to find out what participants who behave optimally are able to report about the payoff schedules. The experiments tested participants' knowledge by asking a series of questions designed to cover every scenario in the task. They used a quantitative test of participants' understanding, as these have been shown to be more sensitive than qualitative tests (Maia & McClelland, 2004).

The experiments also tested whether making participants' payment contingent on the number of points that they gained during the task had an effect on their choices. Particularly in the economics literature, it is seen as crucial to incentivize participants in this way (Hertwig & Ortmann, 2001). Participants in the Contingent condition were paid based on the number of points they earned, whereas those in the Certain condition were paid a fixed amount. It is possible that giving points-contingent payments could cause participants to have more motivation and thus gain more points, or that participants would not explore as fully as they would otherwise and so would settle on a suboptimal strategy, leading to fewer points (Beeler & Hunton, 1997). If, however, it is not a motivation failure that leads to poor performance during experiments using intrapersonal externalities, but rather the cognitive failure to understand the payoff schedules, then a difference might not be expected. This would suggest a cognitive component in decisions that have both long and short-term consequences that has not been fully explored.

3.4.2 Method

Participants

Forty-nine students or staff from Nottingham University volunteered to take part in the experiment; 33 (67%) women (mean age=27.3 years, SD=7.3 years). Participants were randomly assigned to one of two conditions: 26 participants in the Certain condition and 23 participants in Contingent condition.

Design and procedure

Participants were given standardized instructions explaining that the experiment was a decision-making experiment (See appendix A). Also, those in the Contingent condition were told that their payment from the experiment would be based on the number of points that they gained, which would be multiplied by 0.08p/point, and those in the Certain condition were told that they would earn a guaranteed payment of £4. Pilot data had shown that the mean number of points gained over eight sessions was 5220, and so both conditions would on average earn a similar payment (minimum 4000 points = £3.20; maximum 6120 points = £4.90).

The experiment consisted of 8 sessions, each with 150 game units. This equated to between 53 and 150 choices per session, depending upon what the participant's choices were. On each trial, the points gained and game units lost outcome boxes were updated with feedback from the previous trial, and then two buttons were enabled marked '#' and ''. Participants were then prompted to make a choice of one of these buttons. The symbol presented on each button was counterbalanced between participants, as was the payoff schedules attached to each button.

After each choice, both buttons were disabled for between 0.5 and 1.5 seconds, and the points gained and game units lost outcome boxes were cleared to ensure that participants were aware that the outcome boxes indicated feedback from the preceding trial rather than the expected payoff. Choosing the short-term button used up more game units and so the experiment ended prematurely. Consequently, in order to reduce the motivation to finish the experiment quickly, choosing the short-term button increased this delay so that the total delay over the experiment was similar whichever button was chosen. The formula used is shown in Equation 3.4.1.

$$Delay = 0.5 + (P(short - term_{1-10})) \tag{3.4.1}$$

Where $Delay$ is in seconds, and $P(short - term_{1-10})$ is the proportion of choices allocated to the short-term button over the preceding 10 trials. Between each session, participants were shown the points that they gained on the previous sessions. They then completed questions from four scenarios designed to probe their awareness and understanding of the payoff schedules used. The scenarios were the same each time, although they were presented in a random order after each session.

To ensure that participants understood the task, the experimenter sat with the participant for the first session, its feedback, and the first set of scenarios. Participants were then allowed to finish the other sessions in private, although the experimenter was available to answer questions.

Payoff Schedules

Participants received points for every choice that they made, but lost game units. Choosing the short-term key returned 10 points per single trial, however it increased the rate at which game units were lost over the next 10 trials. In contrast, choosing the long-term key returned 5 points per single trial but used up fewer game units over the next 10 trials, so that as long as there were more than 10 game units remaining choosing the long-term key would optimise participants' points payoff. The number of game units lost after each choice was determined by Equation 3.2.1.

To calculate the payoff at the beginning of each session, participants started with a history of ten successive long-term button choices. Over the 150 game units of each session, consistently choosing the long-term button would return a cumulative payoff of 750 points. Consistently choosing the short-term button would return a cumulative payoff of 530 points. However, the optimal solution was to switch from the long-term to the short-term key towards the end of the session, for which a maximum payoff of 765 points was possible.

Stimuli

Figure 3.6 shows the main game screen. Two buttons marked '#' and '' were displayed horizontally next to one another on the computer screen. Above these two buttons, on the left side of the screen were two outcome boxes marked 'Points gained on previous trial' and 'Total points'. On the right side of the screen were another two outcome boxes marked 'Game Units lost on previous trial' and 'Game Units remaining'. At the top centre of the screen a horizontal bar labelled 'Game Units' depicted graphically how many game units remaining there were. The colour of the bar was dependent upon the number of game units remaining; between 51 and 150 it was green, between 11 and 50 it was yellow, and between 0 and 10 it was red. Above this, another horizontal bar labelled 'Points' depicted the total number of points gained during that game. This bar was based around an animated Pac-Man figure which moved from left to right and grew larger as the total number of points increased. These were designed in order to increase the saliency of the feedback. Participants made their choices by selecting one button or another using the mouse.

At the end of each session, a new screen summarised the total points gained during that session and the previous sessions. The top-centre of the screen displayed textually the total points gained during the session, and the maximum number of points that it was possible to gain during a session. Participants in the Contingent condition were also informed how much their session's points were worth monetarily. Below this, a cartoon face was presented, depending upon whether the participant gained more points during the recently com-

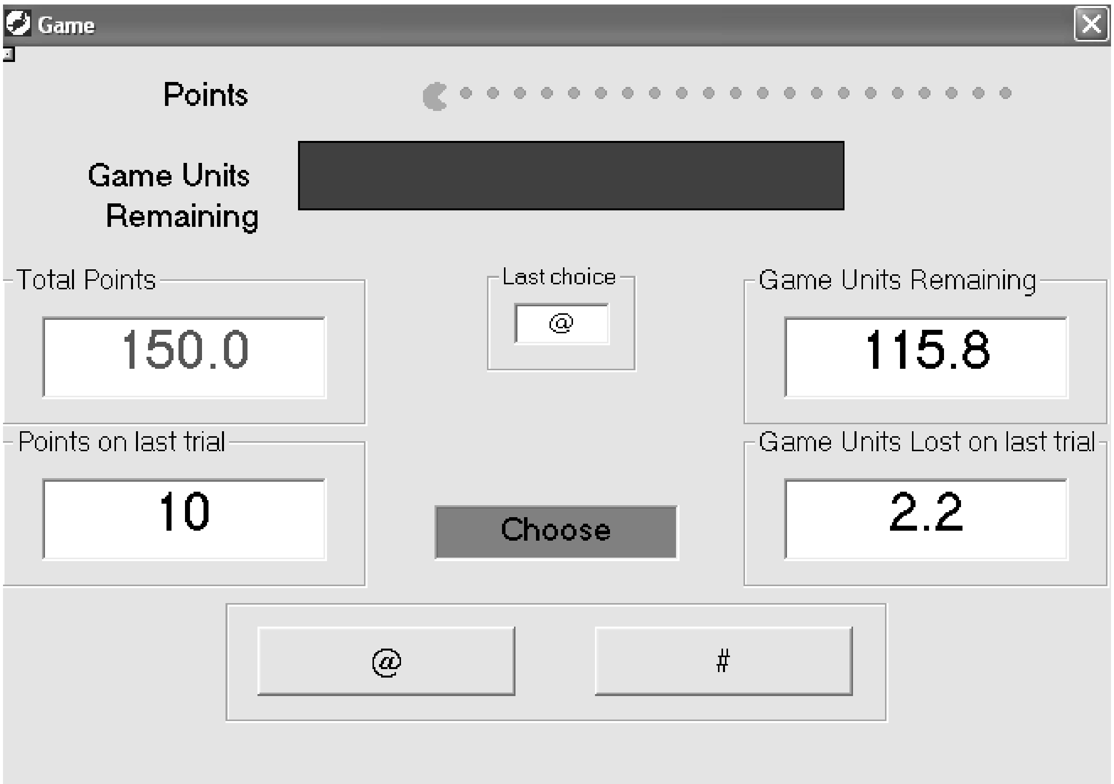

Figure 3.6: Screenshot of main game screen.

pleted session than the previous session. If the participant gained more points, the face smiled; if an equal number of points were earned, it was neutral; and if a lesser number of points were collected, it frowned. Beneath these, a bar chart graphically detailed the total points gained on that session and on all previous sessions.

Contingency Knowledge Probe

Four scenarios were consecutively presented to each participant between each session. For each scenario the participants were asked to answer how many game units would be lost and points gained if the person in it chose to continue pressing the same button (a), or what the outcomes would be if they switched to the other button (b). Participants were given a free response, and did so by typing their answer. The scenarios and their correct answers are shown in Table 3.1. Participants were not given any feedback on their contingency knowledge.

Table 3.1: Awareness probe scenarios and their correct answers. In this example the "left-hand button" refers to the long-term button, and the "right-hand button" refers to the short-term button, but the positions of these were randomized.

Scenario	Points Gained	Game Units Lost
1. John has been choosing the left-hand button repeatedly for the last 20 turns. The last time he chose the left-hand button, he lost 1 game unit and gained 5 points. a) What would happen if he chose the left-hand button again? b) What would happen if he chose the right-hand button next time?		
1 (a)	5	1
1 (b)	10	1.2
2. Jane has been choosing the left-hand button repeatedly for the last 20 turns. However, 3 turns ago she switched to the right-hand button. The last time she chose the right-hand button, she lost 1.6 game units and gained 10 points. a) What would happen if she chose the right-hand button again? b) What would happen if she chose the left-hand button next time?		
2 (a)	10	1.8
2 (b)	5	1.6
3. Bob has been choosing the right-hand button repeatedly for the last 20 turns. However, 3 turns ago he switched to the left-hand button .The last time he chose the left-hand button, he lost 2.4 game units and gained 5 points. a) What would happen if he chose the left-hand button again? b) What would happen if he chose the right-hand button next time?		
3 (a)	5	2.2
3 (b)	10	2.4
4. Sarah has been choosing the right-hand button repeatedly for the last 20 turns. The last time she chose the right-hand button; she lost 3 game units and gained 10 points. a) What would happen if she chose the right-hand button again? b) What would happen if she chose the left-hand button next time?		
4 (a)	10	3
4 (b)	5	2.8

The points gained class of questions reflects whether participants knew that the short-term button always gave 5 points whereas the long-term button always gave 10 points. The game units lost questions reflect different levels of knowledge. Participants that answer question 1a correctly know that the lowest number of game units lost is always 1, and conversely question 4a reflects that the maximum number of game units lost is 3. Questions 1b and 2a reflect that game units lost usually increases compared to the previous trial when the short-term button is pressed, and questions 3a and 4b reflect that game units lost usually decreases compared to the previous trial when the long-term button is pressed. For these, participants do not necessarily have to understand that it is the history of choices that determines the number of game units lost, only that one button usually increases them and the other button usually decreases them. However, questions 2b and 3b are both examples of a situation where the number of game units lost does not always decrease or increase compared to the previous trial when the long-term or short-term button is pressed respectively. For a participant to correctly answer this question, it must be understood that it is the history of choices that determines the number of game units lost.

3.4.3 Results

The average number of points gained by participants in the Certain condition was 5026 (SD=451). Participants in the Contingent condition were paid based on the number of points gained during the experiment. The average number of points gained was 5021 (SD=454) leading to an average payment of £4.02 (maximum obtained: £4.75; minimum obtained: £3.45).

Learning across conditions in the Simplified Harvard Game

The proportion of responses allocated to two buttons was recorded across eight sessions. Each session started with 150 game units, and was split into two blocks for purposes of analysis, based on the number of game units remaining. Any choices made while there were more than 10 game units remaining were allocated to Block 1, whereas any choices made while there were 10 or fewer game units remaining were allocated to Block 2. These two blocks represent the strategies that should be followed; for most of the game participants should choose the long-term button, but at the end of the experiment it becomes optimal to choose the short-term button. The precise optimal switching point for each session depends upon the participant's choices in Block 1, but choosing the short-term button when there were fewer than 10 game units remaining was better than choosing the long-term button. Thus, it was optimal to switch at the beginning of Block 2.

The mean proportions of responses allocated to the long-term button in block 1 for both conditions and in each session are shown in Figure 3.7, and the frequencies of the proportions in each session are shown in Figure 3.8. The data were entered into a repeated-measures ANOVA with Session as the within-subjects factor and Condition as the between-subjects factor. In this and in all further analyses, degrees of freedom were adjusted using the Greenhouse-Geisser method in cases where the assumption of sphericity is violated. The ANOVA revealed a reliable effect of Session ($F(5.34, 250.9) = 24.31$, $MSE = .06$, $p < .001$, $\eta_p^2 = .34$) and a reliable linear contrast indicative of an increasing trend towards maximization as the experiment progressed ($F(1, 47) = 85.86$, MSE = 0.08, $p < .001$, $\eta_p^2 = .65$). However, there was no evidence of an effect of Condition (($F1, 47) = .01$, MSE = .40, $p > .05$), suggesting that rewarding participants for gaining points did not affect their performance.

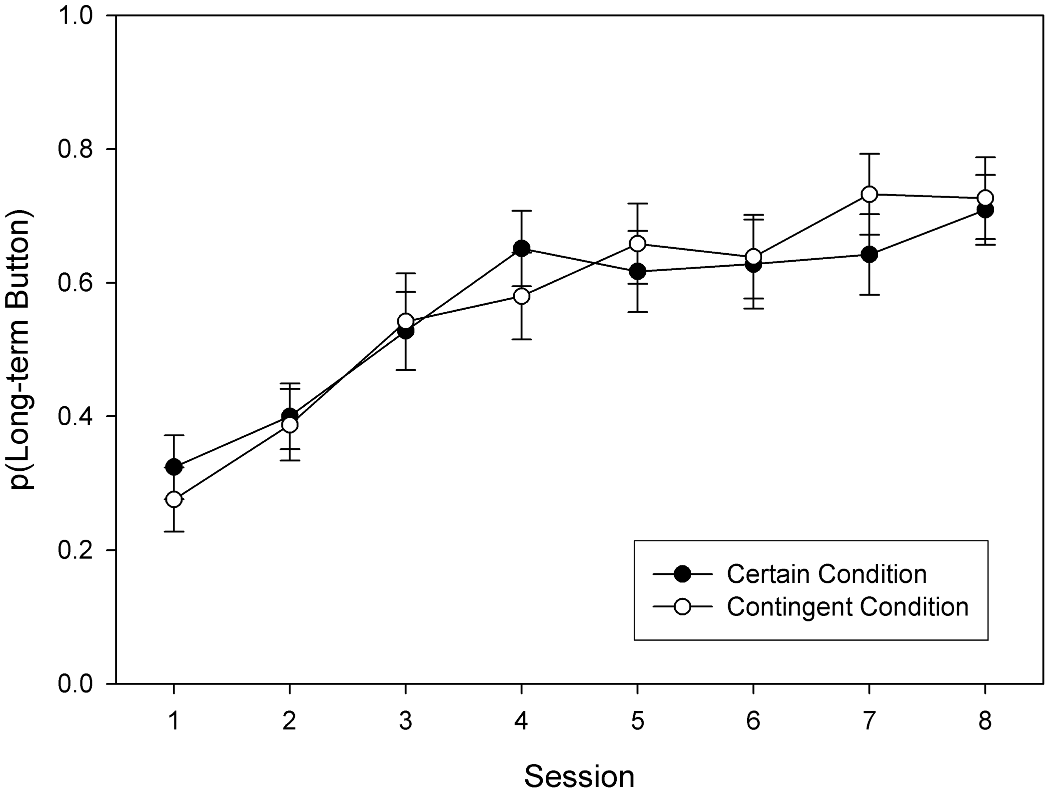

Figure 3.7: Proportion of responses in block 1 allocated to the long-term button as a function of session. Error bars are standard errors of the mean.

If participants' behaviour is aimed at maximizing expected utility, then they should switch from the long-term to the short-term button at the beginning of Block 2. The proportions of long-term button responses for the two blocks of each session are shown in Figure 3.9 and show that toward the end of each session participants increasingly exhibit switching behaviour. To test this we compared the proportions of long-term button responses in Block 1 and Block

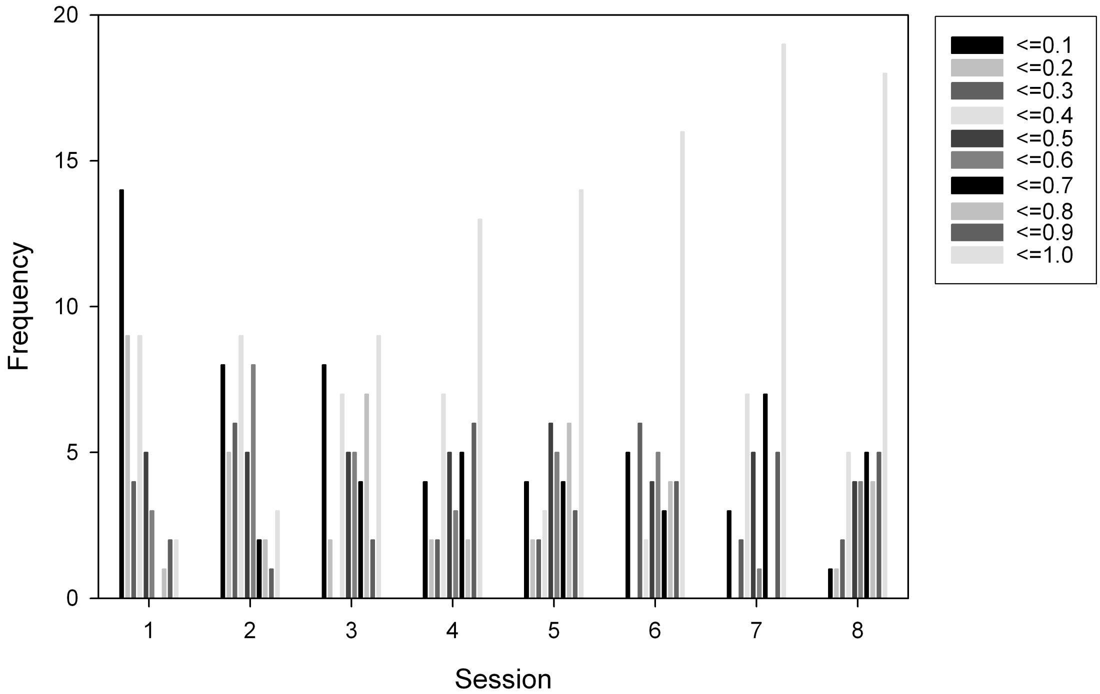

Figure 3.8: Frequencies of responses allocated to the long-term button in each of the 8 sessions of Experiment 1. X axis labels indicate the frequencies of respondents who scored lower than or equal to the label.

2 of each session. These data were entered into a 2x8x2 repeated-measures ANOVA with Session and Block as within-subjects factors, and Condition as the between-subjects factor. The ANOVA showed an effect of Block, signifying that participants were switching responses at the end of each session (F1, 47 = 30.17, MSE = .16, p < 0.001, η_p^2 = .39). A reliable interaction was found between Session and Block (F5.4, 252.3 = 3.05, MSE = .04, p < 0.01, η_p^2 = .06). The linear trend across Sessions differed between Blocks (F1,47 = 8.91, MSE = .06, P < .01, η_p^2 = .16) revealing that participants' choices between Block 1 and Block 2 increasingly diverged as the experiment progressed, and pair-wise comparisons revealed that across the 8 sessions participants increasingly switched from the long-term button in Block 1 to the short-term button in Block 2. This switch-

ing behaviour became consistently apparent after the fourth session. There was no reliable between-subjects main effect of Condition, nor any interactions between Condition and Block or Session, suggesting that rewarding participants for gaining points did not mediate their switching behaviour.

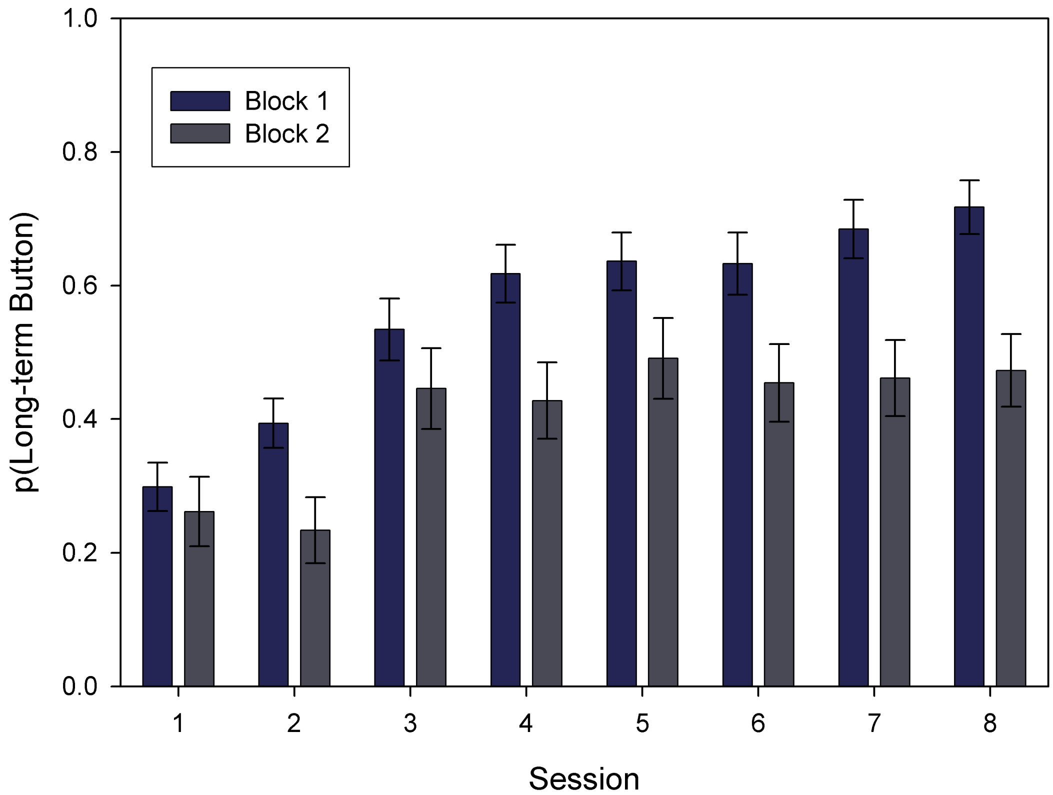

Figure 3.9: Proportion of long-term button choices in each session, in block 1 compared to block 2. Error bars are standard errors of the mean. Probe responses indicative of optimal behaviour.

Between each session, participants answered questions from four scenarios designed to test their knowledge of the payoff schedules used in the experiment. One type of question asked how many points would be gained after the next choice was made. However, ceiling effects were found as most people answered the questions correctly even in the early sessions. Therefore, as this type of question was not able to discriminate between individuals, they were not analysed further.

The second type of question asked how many game units would be lost after the next choice was made. The percentage of participants who answered each question correctly is presented in Table 3.2, as well as an analysis of whether participants learned to answer the question correctly as the sessions progressed. In order to test how effective these were in relating the knowledge that answering the question correctly represented to behaviour in the experiment, participants' correct or incorrect answers to the probe questions were used to predict performance in the next session. Mixed effects models using the nested data, with response to the probe question (correct or incorrect) as a first level predictor and participant ID as the grouping variable, found significant positive effects for Q1A, Q1B and Q2A. To our surprise, it also found a significant negative correlation for Q3B and a sizable (but non-significant) correlation for Q4A, and so these were explored further. It was found that participants who learned to choose the long-term button during Block 1 made systematic errors on Q3B and Q4A; incorrectly assuming that the game units lost would increase from the previous trial (in Q3B, from 2.4 to 2.6, and in Q4A from 3 to 3.2) when it would not. In fact, further mixed effects models found that the misconception in Q3B was associated with optimal behaviour in Block 1.

The systematically incorrect answers by participants who performed optimally in Block 1 of their next game go some way to explaining the pattern of correlations. In order to answer questions 1B and 2A correctly, which participants who chose optimally were more likely to do, it is necessary to know that usually pressing the short-term key leads to an increase in the number of game units lost compared to the previous trial. It is perhaps not surprising that there is such a close relationship between answering question 1A correctly and choosing optimally. Participants who chose optimally would have extensive experience of scenario one's position on of the payoff schedules. It is noteworthy that 3A and 4B were not associated with optimal decision making. These reflect a mistaken understanding that pressing the short-term button always leads to an increase in the number of game units lost on the next trial. Due to the 10 trial history used when calculating the payoff schedules this was not always the case, but as a heuristic it is correct more often than not.

Table 3.2: Point-biserial correlations in Experiment 1 between correct answers on probe questions and the proportion of long-term button choices in block 1 of the next session.

	Q1A	*Q1B*	*Q2A*	*Q2B*	*Q3A*	*Q3B*	*Q4A*	*Q4B*	*Q3B*	*Q4A*
	Corr.	*Corr.*	*Corr.*	*Corr.*	*Corr.*	*Corr.*	*Corr.*	*Corr.*	*Err.b*	*Err.b*
% Chosen	73%	51%	57%	40%	39%	30%	59%	38%	35%	28%
Δ % Chosen between session 2 and session 8[a]	37%**	41%**	39%**	27%**	10%	4%	4%	10%	24%*	4%
Fixed effect (t)	3.56**	3.79**	2.73**	-0.12	-0.14	-2.11*	-1.3	0.57	3.72**	0.1

[a] Evidence of learning across all sessions was tested using a within-subjects ANOVA with the 7 payoff contingency probes as the within subjects factor and the proportion of correct responses as the dependent variable.
[b] The incorrect answer for Q3B was 2.6, and for Q4A it was 3.2. Both of these errors are participants incorrectly assuming that the game units lost would increase from the previous trial.
* = p<.05; ** = p<.01

The overall understanding of the payoff schedules, using the proportion of correct answers across all eight questions as the dependent variable, was compared between the Certain condition and Contingent condition using a 7x8x2 mixed ANOVA with Session and Question as within-subjects factors, and Condition as a between groups factor. The analysis did not find a reliable main effect of

Condition ($F_{1,\,47}$ = 2.05, MSE = 3.99, p = .16), nor did it find an interaction between Condition and Question ($F_{4.5,\,213}$ = .47, MSE = .45, p = .78) suggesting that both groups performed equally well on each question, however it did find a reliable interaction between Session and Condition, ($F_{3.54,\,166.5}$ = 2.76, MSE = .78, p = .04, η_p^2 = .06). The linear contrast for the Session and Condition interaction ($F_{1,47}$ = 5.20, MSE = .36, P = .03, η_p^2 = .10) suggests that despite the lack of a main effect, participants in the Certain condition learned the correct answers to questions at a faster rate compared to participants in the Contingent condition. It can be seen from Figure 3.10 that participants who were paid contingent on their performance stopped improving their understanding of the payoff schedules after the second game, whereas those who were given a guaranteed payment continued to improve beyond this.

In conclusion, participants in both conditions learned to optimize their behaviour and even to switch towards the immediately beneficial option towards the end of the experiment. There was however no evidence that paying participants contingently on their choices changed their choice behaviour, although despite this, those who were given a certain payment for participation benefitted from extended learning of the payoff schedules and were ultimately able to better predict the outcome of choices on the task. There is some indication that participants who performed optimally did not form a full understanding of the historical interaction between the two options affecting the number of game units lost. Instead they generalised that the short-term key increased the number of game units lost compared to the previous trial, leading to systematically

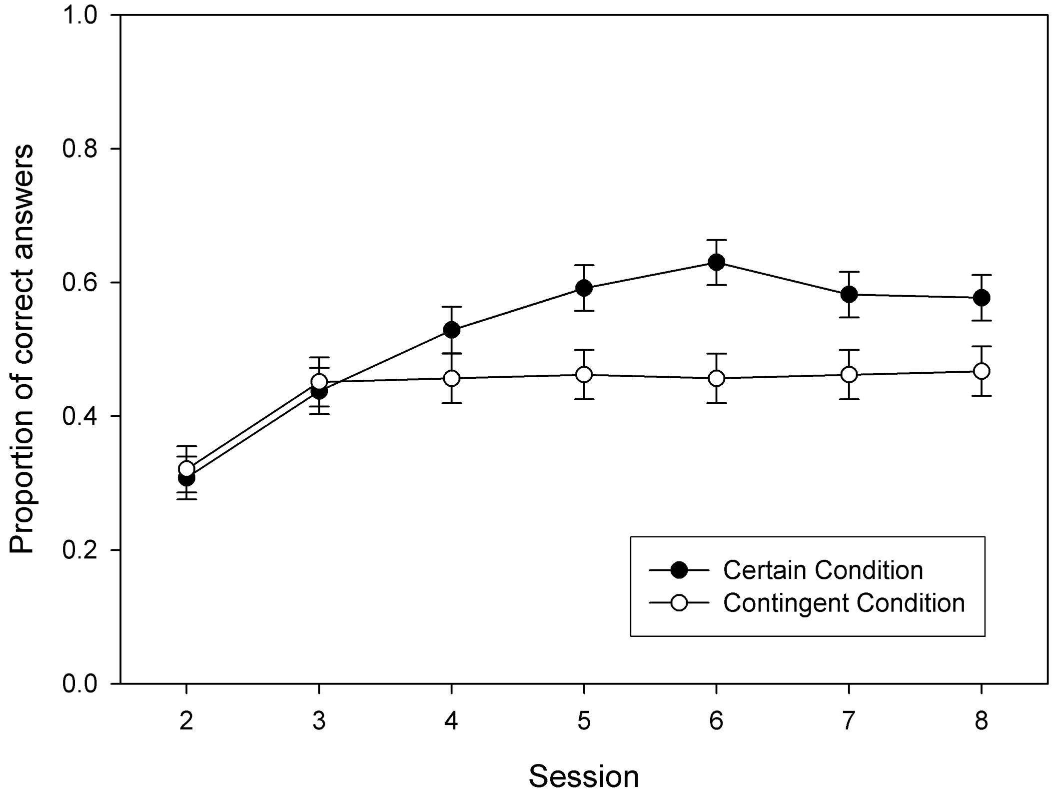

Figure 3.10: Proportion of correct answers to contingency probe questions in each session, separated by condition. Error bars are standard errors of the mean.

incorrect answers on some questions. This may explain why the Contingent condition had a poorer overall understanding of the task but still performed equally well as the Contingent condition.

3.5 Experiment 6: Participants' Understanding of the Simplified Harvard Game 2

3.5.1 Introduction

The previous experiment modified the Simplified Harvard Game by providing knowledge probe questions. It is possible that by asking these questions they changed participants' learning about the task payoffs. In order to test this, the experiment was repeated with two new groups. One group received the knowledge probes throughout, and the other group only received knowledge probes for the final three sessions.

The previous experiment confounds the wait between trials and performance on the task, such that acting impulsively leads to a longer wait between trials. This had the advantage that the overall length of time on the task would be equal for both groups, as those who behave suboptimally end the experiment after fewer trials but have to wait longer between trials to make up for it. Unfortunately this also means that participants could use the waiting time between trials to gauge how well they were doing on the task. Therefore, in Experiment 2 the inter-trial interval was set to a fixed time rather than based on the history of choices.

3.5.2 Method

Participants

Forty-eight students from Nottingham University volunteered to take part in the experiment; 33 (69%) women and 15 men (mean age=25.0 years, SD=6.3 years). Participants were assigned to one of four conditions based on the order in which they presented.

Design and procedure

The design was largely identical to the previous experiment, with half of the participants paid contingent on their performance and half paid a guaranteed amount. However, an additional condition was added which was the point at which contingency knowledge probe questions were asked. Half the participants were asked after each game, and the other half were asked only after the final three games. Additionally, the inter-trial interval was set to 1 second, rather than the variable interval as in the previous experiment.

Contingency Knowledge Probe

In order to simplify the probe questions, participants' answers were restricted to whether the number of game units lost ”would stay the same”, ”would increase” or ”would decrease”. Otherwise, the questions and were identical to the game unit questions from the previous experiment, shown in Table 1.

3.5.3 Results

The average number of points gained by participants in the Certain condition was 5297 (SD=306). Participants in the Contingent condition were paid based on the number of points gained during the experiment. The average number of points gained was 5383 (SD=342) leading to an average payment of £4.31 (maximum obtained: £4.72; minimum obtained: £3.67).

Learning across conditions in the Simplified Harvard Game

The mean proportions of responses allocated to the long-term button in Block 1 for the four conditions and in each session are shown in Figure 3.11, and the frequencies of the proportions in each session are shown in Figure 3.12. The data were entered into an 8x2x2 repeated-measures ANOVA with Session (1 to 8) as the within-subjects factor (coded numerically) and the two Conditions (contingent vs. certain payment, and early vs. late knowledge probe) as between-subjects factors. The ANOVA revealed a reliable effect of Session (F3.49, 153.42 = 16.08, MSE = .09, p < .001, η_p^2 = .27) and a reliable linear contrast indicative of an increasing trend towards maximization as the experiment progressed (F1, 44 = 42.9, MSE = .10, p < .001, η_p^2 = .49). The effect of Payment Condition (F1, 44 = 1.01, MSE = .21, p > .05) was not significant, suggesting that rewarding participants for gaining points did not affect their task performance. For Probe Condition (F1, 44 = 3.85, MSE = .21) the effect was almost significant. To properly analyse this, we compared the two probe conditions for Sessions

2-6, as neither group had received the probe during the first session, and both groups had received the probe during Sessions 7 and 8. The 5x2 within-subjects ANOVA found an effect of Probe Condition ($F1, 46 = 4.97$, $MSE = .14$, $p = .031$) indicating that the group who received the probes chose the long-term button more often than the group who did not. One explanation for this finding is that asking knowledge questions encouraged participants to gain explicit knowledge of the payoff schedules, which in turn improved their performance on the task.

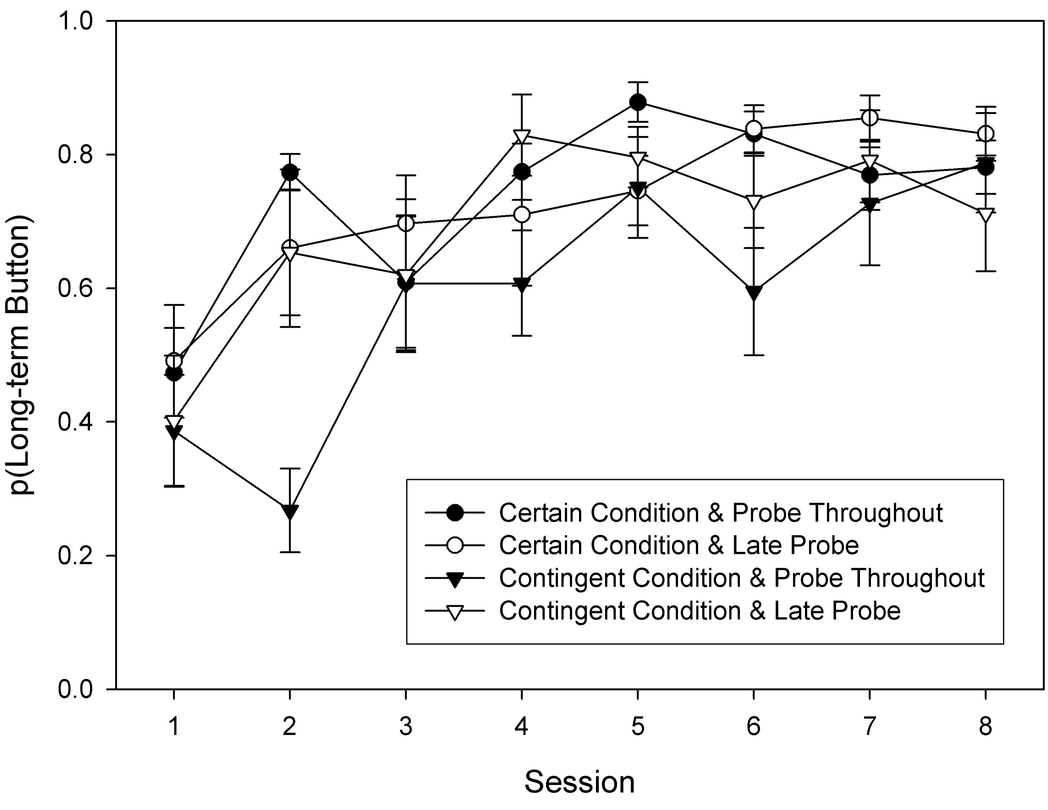

Figure 3.11: Proportion of responses in block 1 allocated to the long-term button as a function of session. Error bars are standard errors of the mean.

If participants' behaviour is aimed at maximizing expected utility, then they should switch from the long-term to the short-term button during the final

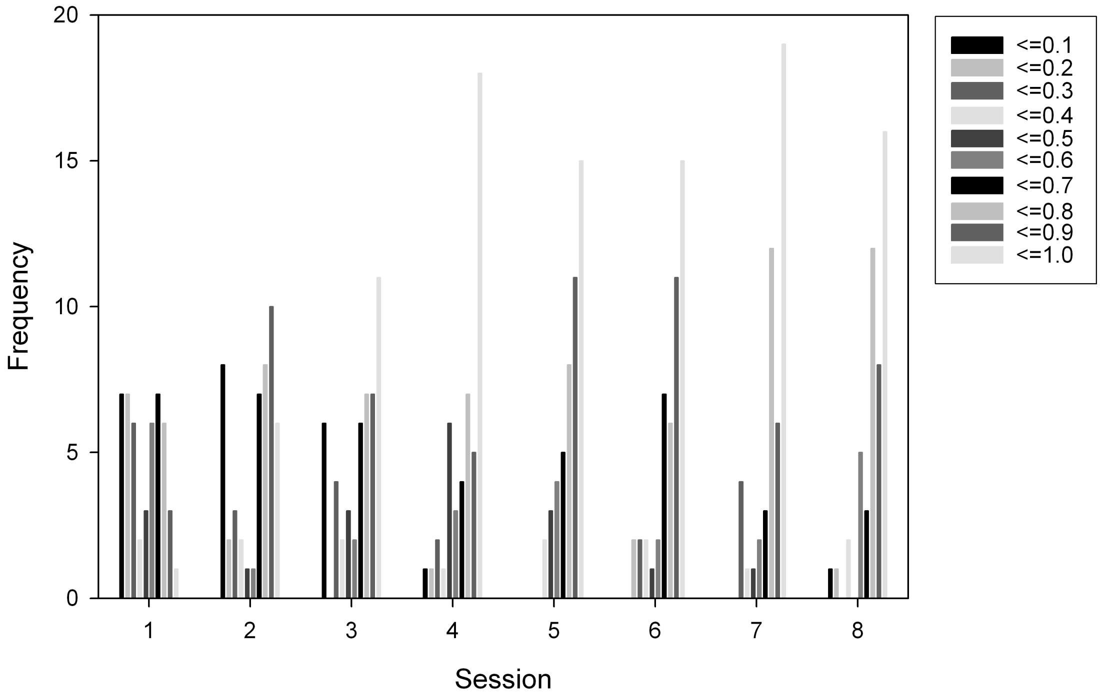

Figure 3.12: Frequencies of responses allocated to the long-term button in each of the 8 sessions of Experiment 2. X axis labels indicate the frequencies of respondents who scored lower than or equal to the label.

block. To test this we compared the proportions of long-term button responses in Block 1 and Block 2 of each session. These data were entered into an 8x2x2x2 repeated-measures ANOVA with Session (1 to 8) and Block (1 to 2) as within-subjects factors, and the two Conditions (contingent vs. certain payment, and early vs. late knowledge probe) as the between-subjects factors. The ANOVA showed an effect of Block, signifying that participants were switching responses at the end of each session ($F1, 44 = 58.34$, $MSE = .21$, $p < 0.001$, $\eta_p^2 = .57$). A reliable interaction was found between Session and Block, signifying that participants switched more as the experiment progressed ($F4.4, 192 = 7.90$, $MSE = .05$, $p < 0.01$, $\eta_p^2 = .15$). The linear contrast for the Session and Block interaction ($F1,44 = 29.74$, $MSE = .04$, $P < .001$, $\eta_p^2 = .40$) revealed that participants switched more as the experiment progressed, and pair-wise comparisons revealed that this switching behaviour became consistent from the second session onwards. There was no reliable interaction between the two Conditions and Block or Session, suggesting that rewarding participants for gaining points did not mediate their switching behaviour.

Probe responses indicative of optimal behaviour

The percentage of participants who answered each probe question correctly is presented in Table 3.3. As in the previous experiment, mixed effects models were used to test for a relationship between responses to the probe questions and proportion of long-term choices in Block 1 of the next Session, while holding the participant ID constant in order to account for the different skill levels amongst participants. A significant positive relationship was found between both Q1A and Q2A with Block 1 of the next session; participants who answered those questions correctly were more likely to choose the long-term button more often in Block 1 of the next Session. As in the previous experiment, correct answers for Q3B and Q4A were negatively related to performance in the next ask (although not significantly in this experiment). In order to be consistent with the previous experiment, the equivalent incorrect answers for these questions were analysed. The direction of the effect was the same as in the previous experiment; participants who behaved optimally incorrectly assumed that the rate that game units were lost would increase from the previous trial, although this was not a significant effect.

The smaller number of significant effects in this experiment may be due to the format of the probe questions. In the previous experiment, participants were given a free response, so guesses were unlikely to be correct by chance. But in the current experiment, participants were given a multiple response option, so guesses were more likely to be correct.

Table 3.3: Point-biserial correlations in Experiment 6 between correct answers on probe questions and the proportion of long-term button choices in block 1 of the next session.

	Q1A	***Q1B***	***Q2A***	***Q2B***	***Q3A***	***Q3B***	***Q4A***	***Q4B***	***Q3B***	***Q4A***
	Corr.	***Corr.***	***Corr.***	***Corr.***	***Corr.***	***Corr.***	***Corr.***	***Corr.***	***Err.***[a]	***Err.***[a]
% Chosen	74%	58%	68%	46%	39%	34%	57%	46%	40%	33%
Fixed effect (t)	3.02**	2.60*	1.44	0.32	0.77	-0.78	-0.2	0.86	1.08	0.52

[a] The incorrect answer for both Q3B and Q4A was that the game units "would increase". * = p<.05; ** = p<.01

The overall understanding of the payoff schedules, using the proportion of correct answers across all eight questions as the dependent variable, was compared between the Certain condition and the Contingent condition using a mean of the final three sessions where all participants answered contingency knowledge probe questions. A 3x8x2x2 mixed ANOVA with Session (6-8) and Question (q1a to q4b) as the within-subjects factors, and Probe Condition and Payment Condition as between groups factors. The analysis found a reliable main effect of Payment Condition (F1, 44 = 4.40, MSE = 1.17, p = .04, η_p^2 =.09), indicating that those who were given a certain payment answered more questions correctly in the final three sessions than those given a contingent payment. There was no main effect of Probe Condition (F1, 44 = .06, MSE = 1.17, p = .81) indicating that the timing of the probes did not affect participants' understanding of the task. It can be seen from Figure 3.13 that participants who were paid contingent on their performance generally stopped improving their understanding of the payoff schedules after the second session, whereas those who were given a certain payment continued to improve beyond this, although they took longer before they reached their optimal understanding of the task.

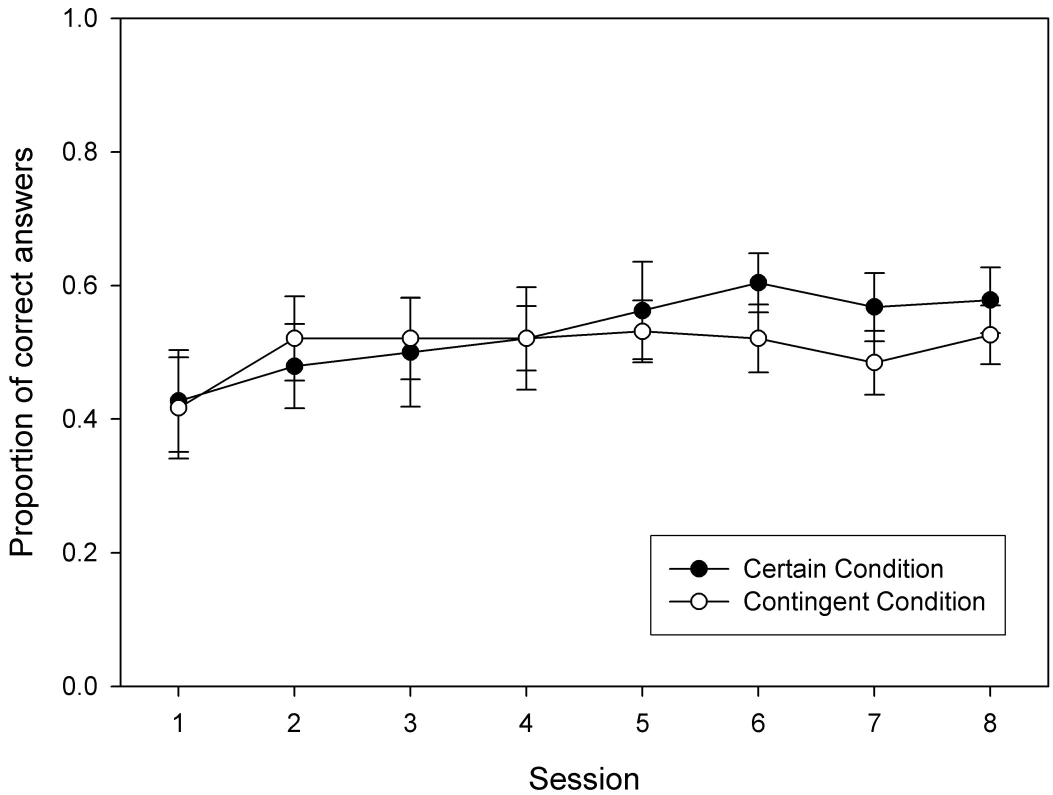

Figure 3.13: Proportion of correct answers to contingency probe questions in each session, separated by condition. Error bars are standard errors of the mean.

In conclusion, this experiment replicated the behavioural and contingency knowledge results of the previous experiment and there was no evidence that the probe questions changed participants' behaviour in the task or their understanding. Participants optimized their behaviour whether they were paid contingent upon their performance or not, and those who performed the best seemed to follow a heuristic that did not fully characterise the complexity of the payoff schedules. There was also evidence that those paid contingently on their responses did not learn as much explicit knowledge about the payoff schedules. Finally, there was no evidence that fixing the inter-trial interval altered participants' performance in the task.

Discussion

In the Harvard Game experiments reported here, participants learned to take account of the intrapersonal externalities inherent in the task, maximizing their expected utility. By the final session, most participants chose the long-term option for the majority of the session, and switched responses to the short-term option towards the end. In real-world terms, participants learned to choose activities that would increase their long-term welfare rather than those that gave immediate gratification, but once they realised that the end was close participants learned to prioritise their short-term needs.

By asking participants to complete quantitative questions about their predictions in the experiment we could distinguish between the different conceptualisations of the task that participants held. Based on the questions that participants who made more optimal choices answered correctly, and the errors that participants made, we found evidence that participants appeared to use a generalised heuristic that one option would usually increase the rate at which choice opportunities remaining in the experiment decreased. There was no evidence that participants who made more optimal choices realised that the other option would usually decrease the rate at which the number of choice opportunities remaining in the experiment decreased.

In addition, the experiments found that paying participants based on their choices had no observable effect on their choice behaviour, despite the claims of Hertwig and Ortmann (2001). The finding that payment type does not affect choice behaviour is consistent with similar research using the Iowa Gambling Task (Bowman & Turnbull, 2003) which found no difference between real or facsimile money.

It is possible that contingent payments increased participants' motivation but this was offset by a decrease in exploratory behaviour. Supporting evidence for this was found from the analysis of how participants' understanding of the payoff schedules changed across the eight sessions. In the group who were paid contingent on their performance, their ability to successfully predict the outcome of choices in the simplified Harvard Game plateaued after the second game, whereas the group who were given a guaranteed payment continued to improve beyond this. In both experiments, the participants who were paid a fixed amount learned more about the payoff schedules by the end of the experiment than participants who were paid contingent upon their performance. This is consistent with previous research by Schwartz (1982), who found that participants' learning of complex sequences was impaired by giving them contingent reinforcement. Schwartz (1982) concluded that participants repeated what worked in the past rather than trying to understand the task. Since explicit knowledge of the payoff schedules is overall related to performance on the task, it is an open question whether, given enough trials, participants not paid contingent on their reinforcement would ultimately understand the task better and so also learn to perform better than those paid contingent on their performance, or whether participants paid contingent on their performance would extend their initial exploratory period so that they gained more understanding before settling into a pattern of responses that seems to work.

The experiments provide insight into what people who take into account intrapersonal externalities understand. As we found evidence that explicit understanding is related to performance in the task, but found no difference in performance based on how participants were paid, this suggests that the task is predominantly a cognitive one. It is possible that an intervention could be piloted to increase the occurrence of behaviours with positive intrapersonal externalities, or decrease the occurrence of behaviours with negative intrapersonal externalities. An intervention could use pervasive digital devices to provide immediate and personalized feedback each time an individual engages in behaviour with intrapersonal externalities.

As well as intrapersonal externalities, other factors are associated with apparently impulsive behaviour. For example high rates of time-discounting are also related to addictive behaviour (Kirby et al., 1999; J. M. Mitchell, Fields, D'Esposito, & Boettiger, 2005; Vuchinich & Simpson, 1998). The difficulty that the human decision-making system has in taking account of intrapersonal externalities should be considered as an additional factor leading to addictive behaviour.

To conclude, if the results of experiments on intrapersonal externalities are to be useful in understanding suboptimal behaviour in the real world, the differences in understanding between participants who learn to behave optimally and those who do not needs to be understood. Our results suggest that explicit awareness is useful for making optimal decisions in the simplified Harvard Game, but that participants who make more optimal choices do not have, and do not need, a full knowledge of the historical interaction that leads to each payoff. Instead participants learn a simpler conception which emphasises that one option generally seems to make an aspect of their situation worse than it was previously, on a myopic choice by choice basis. When the two aspects (history of choices and current choice) are combined into a single outcome in the full Harvard Game, and presumably in real life intrapersonal decision-making, it is much more difficult to learn this relationship and this may explain why the simplified Harvard Game is simpler to learn.

3.6 Chapter Discussion

In this Chapter, Experiment 3 found that performance in the Simplified Harvard Game is related to smoking behaviour; high nicotine dependency was related to myopic choices on the task. Experiment 4 reversed the payoff schedules from the Simplified Harvard Game from a gain frame to a loss frame, and was unable to distinguish smoking behaviour. This suggests that the reason that smokers chose myopically is because they have a deficit in learning about gains, rather than due to impulsive disinhibition. Experiments 5 and 6 found that conscious understanding of the task was associated with high performance, which further backs up the assertion that the Simplified Harvard Game measures learning rather than simply impulsive disinhibition. Additionally, it was found that rewarding participants for performance rather than participation did not affect the number of myopic choices made in the Simplified Harvard Game, which provides further evidence that the Simplified Harvard Game is a learning challenge rather than a challenge of withholding impulses.

CHAPTER 4

Effects of Advice in the Factory Manager Game

Advice is often given to smokers and drinkers, and yet advice-seeking and advice-taking have not been studied experimentally in self-control situations. This is because if poor long-term decisions are due to impulsivity then the individual understands that a decision would be suboptimal but is unable to avoid making it anyway. If poor long-term decisions are due to discounting future rewards then the individual already has the ability to make an optimal decision and so does not need advice. But if the individual is struggling to make an optimal decision given the difficulties of intrapersonal externalities, such as not knowing the payoffs from all choices, then advice might improve their choices.

4.1 Chapter Introduction

Self-control decisions are not made in a social vacuum. Applied research already implicitly accepts this and has studied the prevalence and effectiveness of general practitioners' advice given to smokers (Russell, Wilson, Taylor, & Baker, 1979), problem drinkers (Fleming, Barry, Manwell, Johnson, & London, 1997; Wallace, Cutler, & Haines, 1988), and obese patients (Galuska, Will, Serdula, & Ford, 1999). Nevertheless, although there has been research in the past two decades into the method by which people integrate advice into their decision making, little is known about how individual differences may alter this process. Even less is known about situations where decisions are made that involve self-control, defined as a situation where one has to choose between an immediate smaller reward and a future larger reward.

4.1.1 The Simplified Harvard Game

One reason that self-control situations have not been studied in the context of advice is that many of the laboratory methods that are used as proxies to measure self-control implicitly assume that self-control is a trait and so not influenced by advice. Individuals answering trait questionnaires such as the Barratt Impulsivity Scale (BIS-11; Patton et al., 1995) cannot be given advice as to which response option to choose. Similarly in delay discounting, which asks questions to determine how much immediate reward an individual would give up in order to obtain a delayed reward, individuals cannot be advised on their own preferences. Both of these methods are associated with behaviours such as smoking and problem drinking that are typically seen as demonstrating a lack of self-control (Flory & Manuck, 2009; J. M. Mitchell et al., 2005; S. H. Mitchell, 1999, 2004).

The Simplified Harvard Game is appropriately placed to fulfil this role, as it is a behavioural measure of self-control that has already been demonstrated to be related to both trait impulsivity as measured by the BIS-11 and smoking behaviour (See Experiment 5 and 6). In the SHG, participants choose repeatedly between two options. One option (the local choice) leads to a high immediate reward but decreases the reward for future choices of either option. The other option (the global choice) leads to a lower immediate reward, but has no long-term effects. Choosing the global choice repeatedly leads to a better outcome, even though on any one trial it is advantageous to choose the local option. Participants are never told the reward contingencies for each button, but must learn them through experience. An advisor can encourage the participant to choose the global option, and so effectively encourage them to act in a self-controlled manner.

4.1.2 The Judge-Advisor System (JAS)

In JAS experiments, a judge makes a final decision on a task but receives advice, or at least has advice available, from one or more advisors (Sniezek & Buckley, 1995). JAS experiments can be split into two types, ”choice” tasks and ”judgement” tasks (Bonaccio & Dalal, 2006). Choice tasks are those involving discreet categories, and so the advisor recommends a particular choice, whereas judgement tasks involve quantitative judgements, and the resulting advice is another participant's estimate.

A separate feature that differs across experiments is whether advice is imposed or whether it must be sought by the judge (Gino, Shang, & Croson, 2009). In the following experiments 7 to 10, we investigate advice-seeking, and in experiment 11 we investigate advice-taking. For this reason, we split up our analysis of previous JAS findings into the two aforementioned categories.

4.1.3 Advice Taking

The majority of JAS research has concentrated on how judges change their initial opinions when they are given advice. Generally, judges do not change their initial opinions enough, and so exhibit egocentrism of advice discounting (Harvey & Harries, 2004). Advice commonly weighs around just 20-30% in the final judgement (Harvey & Fischer, 1997). Although this may only be the case on tasks where the judge is confident in their ability to decide, as Gino and Moore (2007) found that in difficult tasks advice was actually over-weighted.

Judges do not weight all advice equally. The opinions of advisors are given more weight when they are believed to be experts (Harvey & Fischer, 1997; Sniezek, Schrah, & Dalal, 2004), when they are believed to be well-intentioned (Sniezek & Van Swol, 2001), when their opinion is not considered to be an outlier (Harries, Yaniv, & Harvey, 2004), and when they have to pay for the advice (whether this is obligatory or not) (Gino et al., 2009). Judges are also more confident in their decisions when they are given advice from lots of advisors (Budescu & Rantilla, 2000), and they solicit advice more often from advisors who have unshared information rather than advice that they already possess (Van Swol & Ludutsky, 2007) although the authors did not find evidence that the judges actually changed their decisions.

Neuropsychological patient research has found that individuals with right dorsolateral and orbito-frontal lesions may be especially likely to discount advice due to interference with assessment and weighting of advice (Gomez-Beldarrain, Harries, Garcia-Monco, Ballus, & Grafman, 2004)). Interestingly, these areas are also related to performance in the Harvard game (Yarkoni, Braver, et al., 2005; Yarkoni, Gray, et al., 2005).

Advice is not homogenous. Dalal and Bonaccio (2010) studied five types of advice. "Receommended for" represented the advice traditionally used in JAS experiments where the advisor recommended one alternative. "Recommended against" counselled against one alternative, "Information" provided details about one or more alternatives, "decision-support" gave assistance on how to make the decision, and "social support" provided interpersonal assistance. They found that in certain situations, judges preferred receiving more information rather than a specific recommendation.

4.1.4 Advice Seeking

Traditionally in the JAS advice is automatically provided to decision-makers. But there has been a paucity of research into when judges ask for advice and whom they ask for it from. It has been found that advisors are subject to reputation formation, and so judges seek advice from advisors who have been shown to be more accurate (Yaniv & Kleinberger, 2000), and judges tend to acquire their own information before they seek advice (Schrah, Dalal, & Sniezek, 2006). Some studies find that judges do not always solicit free advice (Gardner & Berry, 1995, Exp 3), but other studies found that free advice from experts was nearly always solicited (Schrah et al., 2006). Similarly contradictory findings have been found for unsolicited versus explicitly solicited advice, which is either discounted to a greater extent (Gibbons, Sniezek & Dalal, 2003, cited in Bonaccio & Dalal, 2006) because it is intrusive (Goldsmith & Fitch, 1997), or there is no difference (Gino & Moore, 2007).

4.1.5 Individual Differences in Advice Utilisation

Of particular relevance to self-control situations is how individual differences affect the use of advice. Dalal and Bonaccio (2010) found a range of differences in the personalities of those who appreciated different types of advice, including that those high in agreeableness and females are more likely to welcome social support, and those with a dependent decision making style were thankful for decision support. Koestner and colleagues (1999) found individuals high in reflective autonomy followed the recommendations of credible advisors whereas individuals high in reactive autonomy rejected their recommendations. A more subtle finding is that when asked to make decisions about their own likely behaviour, judges were more swayed by advisors who they were told were demographically similar to them, but when asked to make decisions about others they were more swayed by advisors who were dissimilar to them (Gino et al., 2009). Judges stated that this was due to the perceived informativeness of the advice.

4.2 Experiment 7: Advice Seeking with Predominantly Impulsive Advice

4.2.1 Introduction

Given the amount of advice available in the real world for self-control situations, such as choosing whether to smoke or drink, we might expect that individual differences in cognition which manifest themselves in high trait impulsivity might also be related to differences in the acquisition of advice. This would help to go some way to explain why those who score high in trait impulsivity also tend to engage in behaviours that demonstrate low self-control, perhaps not seeking out the advice that would help them.

In order to test this, we created the Factory Manager Game, which was an extension of the Simplified Harvard Game (Stillwell & Tunney, 2009). Instead of choosing between two extremes, to be self-controlled or to meliorate, participants were able to choose on a gradient between them, in order to mostly meliorate for example. This allows us to give the participant advice to be more or less self-controlled, without hinting to the participant that their payoff is based on their history of choices. This concrete setup has face validity for participants also, as they take control of a machine and decide on its production level. Although participants are naïve to begin with, they come to learn that the machine can be run quickly, which produces lots of goods immediately, but that doing this causes the machine to wear out faster than simply running it at a slower speed.

We therefore measured individuals' propensity to seek advice in a task that mirrored the Harvard Game. We expected that individuals high in trait impulsivity would solicit advice differently from individuals low in trait impulsivity, although made no prediction as to whether it would be more or less frequently. On the one hand, impulsive individuals may be more susceptible to peer pressure and thus solicit advice more often, but on the other hand they might be more independent and so solicit advice less often.

4.2.2 Method

Participants

Fifty students from Nottingham University volunteered to take part in the experiment after being recruited by posters and e-mails. Their mean age was 23.5

years (SD=5.2 years) and 32 (64%) were female. Participants completed the BIS-11 impulsivity questionnaire, and a total score was derived in accordance with established criteria (Patton et al., 1995). Participants were categorised as high or low impulsive in this and in all subsequent experiments by a median split of the BIS-11 total score.

Procedure

Following arrival, participants were seated in a testing room where they completed a consent form, followed in counterbalanced order by the BIS-11 (Patton et al., 1995) and TAS-20 (Bagby, Parker, & Taylor, 1994) questionnaires, and the Factory Manager Game. Finally, participants were debriefed and paid in accordance with their performance in the Factory Manager Game (see below).

Factory Manager Game

The experimenter read the following instructions to the participant:

> You have been promoted to the position of factory manager. Your factory uses a machine to produce "Goods". Fortunately the factory manager's job is a simple one. Simply set the speed of the machine in order to produce as many Goods as possible.
>
> As a result of your choices your machine will produce Goods. After every run you will be shown the number of Goods produced that run as well as the total Goods produced. However, your machine

will also wear out. After every run you will be shown the machine integrity lost that time as well as the total machine integrity remaining. Your machines start at 100% but will eventually wear out to 0%. Once this happens then the machine is completely worn out and the game is over.

You have also been given access to some advisors, who were previously factory managers like you. After playing the Factory Manager game, they left a recommendation for the speed that you should run the machine. However, you will only be told whether the speed that you last ran the machine is slower or faster than they recommended. Ultimately it will be your decision to set the machine speed to whatever you want, but there is no cost to receiving advice.

You will play the Factory Manager game 8 times. Try and beat your previous score in every game!

Your payment from this experiment will be based on the number of Goods that you produce. This will be calculated on the basis of 0.12p/good. This means you can earn between £3.20 and £4.90. After each game, you will be shown how much you earned.

That's all there is to it – just try to produce as many Goods as you can before your machine wears out, and use the available advice whenever you want it. Take as much time as you wish and please do not write anything down during the experiment.

The experimenter then took the participant through an on-screen tutorial using the following text:

> Before you play the game for real, there will be a short practice to ensure that you are used to the controls. Beware that the practice console will not reflect how production in the real Factory Manger game works. Please follow the instructions below.
>
> 1) To make a choice, drag the production level slider to the machine speed that you want. Do not let go of the slider until it is at the speed that you want because as soon as you let go then that will be your choice. Go ahead now and set this to 150%.
>
> 2) You do not have to change the production level each time if you do not want to. If you just click once on the slider without dragging it then it will choose the speed you chose last time. Go ahead now and choose 150% again.
>
> 3) You can click on the ask advisor button to get advice. Go ahead now and ask Advisor 1 what (s)he thinks. That's it. Press the continue button to start the Factory Manager game.

In the tutorial, the feedback that appeared on the screen did not reflect the real payoff schedules in the Factory Manager Game. Participants were only able to go through the steps shown; the game prevented them from further exploration.

The simplified Harvard task consisted of 8 games, each with 150 game units. This equated to between 53 and 150 trials per game, depending upon participants' choices. Figure 4.1 shows the main screen of the game. At the top centre of the screen a horizontal bar labelled 'Machine Integrity' provided clear feedback concerning the payoff associated with each key. Above this, another horizontal bar labelled 'Goods' depicted the total number of Goods gained during that Game. These were designed in order to increase the salience of feedback concerning the payoff (rewards and costs) associated with a given production level. Participants made their choices by dragging the production level slider using the mouse. On each trial the points gained and game units lost information was updated, overwriting the outcome of the previous trial, so that participants could learn the payoff structure associated with the production level slider.

After each choice, the production level slider was disabled for between 0.5 and 1.5 seconds, and this variable delay was used to ensure that the total duration of each game was the same irrespective of the participant's choices, given that increasing the production level used up more machine integrity and so the experiment ended more quickly. To reduce the motivation to finish the experiment quickly, increasing the production level increased this delay according to Equation 4.2.1, where D is the delay in seconds, and $\mu(lvl_{1-10})$ is the mean production level over the preceding 10 trials.

$$D = 0.5 + \frac{\mu(lvl_{1-10}) - 100}{100} \tag{4.2.1}$$

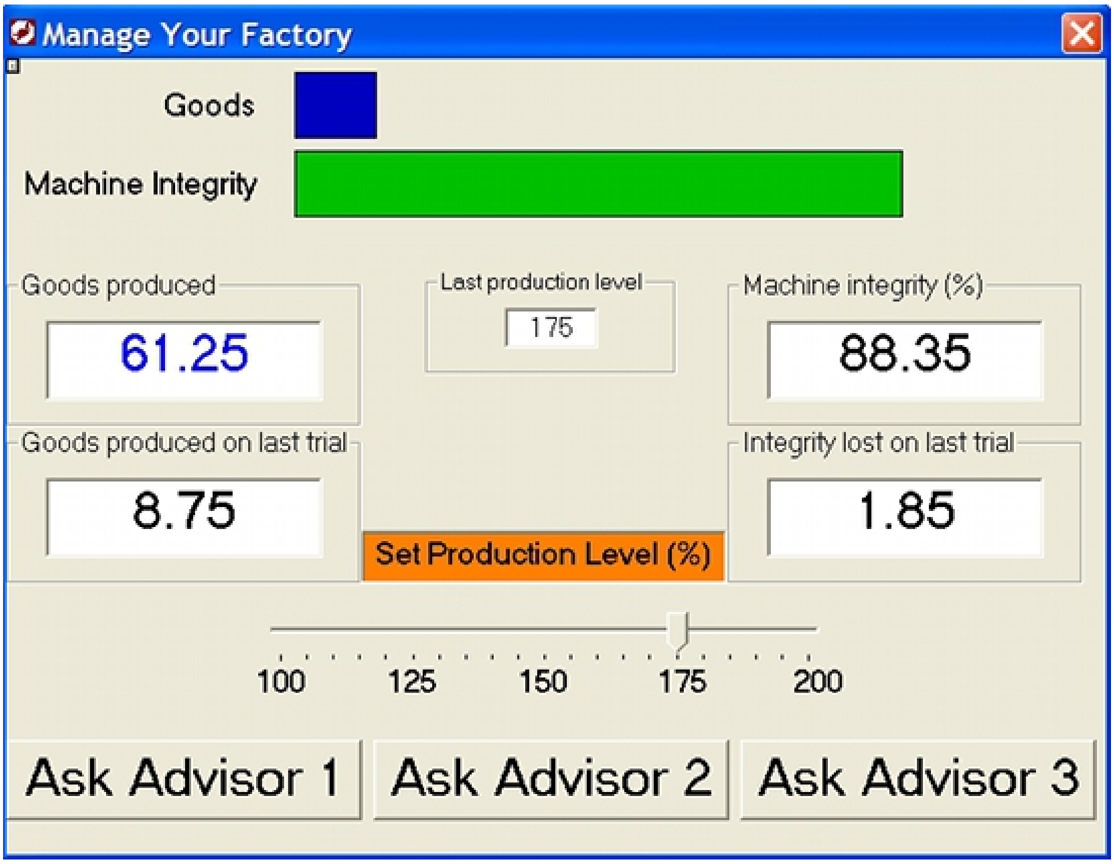

Figure 4.1: Factory Manager main screen

Payoff Schedules

On each production run, participants produced goods but their machine's integrity reduced. Choosing the maximum production level (200%) returned 10 goods per run, but increased the rate at which the machine's integrity reduced over the next 10 trials. In contrast, running the machine at its minimum production level (100%) returned 5 goods per run, but did not have long-term ramifications. The percentage of machine integrity lost after each run was determined by Equation 4.2.2.

$$Itg = 1 + 2^{\frac{\mu(lvl_{1-10}) - 100}{100}} \tag{4.2.2}$$

Where Itg is the amount of integrity lost, and $\mu(lvl_{1-10})$ is the mean production level over the preceding 10 trials. In essence, running the machine at full speed produced twice as many goods but used up three times the machine integrity over the next 10 trials, and so the lost machine integrity acts as a punisher as it represents a lost opportunity to produce more goods.

In each game, consistently running the machine at 100% produced 500 goods, and consistently running the machine at 200% produced 363 goods. However, the optimal solution was to switch from running the machine at 100% to running it at 200% when there was seven percent machine integrity remaining. Adopting this strategy would return a cumulative payoff of 510 goods. Consequently, in the analysis, we examined choice of the two buttons separately in trials either side of seven percent machine integrity remaining. This was the point at which it was optimal to switch from running the machine slowly to quickly. Any choices made while there was more than seven percent machine integrity remaining were allocated to Block 1, whereas any choices after this were allocated to Block 2. Groups are therefore contrasted in Block 1 to examine cost-discounting, and in Block 2 to examine their basic understanding of the payoff structure of the task, as only those who understand the schedules would learn to switch at the end of each session.

At the end of each game, a new screen summarised the total goods produced during that game and the previous games. The top-centre of the screen displayed textually the total goods produced during the game, how much those goods were worth monetarily, and the maximum number of goods that it was possible to gain during a game. Below this, a cartoon face was presented, contingent upon whether the participant produced more goods than in the previous game. If the participant produced more, the face smiled; if an equal number of goods were produced, it was neutral; and if a lesser number of goods were produced, it frowned. Beneath these, a bar chart graphically detailed the total goods produced on that game and on all previous games.

Procedure and Advice

Advice was available from three advisors. Participants could ask for an advisor's recommendation at any point in the game by clicking on one of the "Ask Advisor" buttons. Each of the advisors had an ideal production level in mind, and so every time their advice was asked the advisor would give advice relative to the previous production level set by the participant. They would tell the participant to either set the production level higher, lower, or to leave it about the same. An example of the advice given was, "Advisor 1: Last time you set the production level at 163%. I recommend that you set it higher than that." Participants were told that the advisors were participants who had previously done the experiment, and who had left a recommendation for the speed to run the machine. The three advisors' ideal production levels were 180%, 175% and 120%. It is not uncommon in JAS experiments for the advice to be created by the experimenter but that the participant is told it is real advice (e.g. Brehmer & Hagafors, 1986; Budescu & Rantilla, 2000).

4.2.3 Results

Simplified Harvard Game

Figure 4.2 shows the average production level in block 1 of the 8 Factory Manager games, for the high and low impulsive groups. These data were entered into an 8x2 mixed ANOVA with game as the within-subjects factor and impulsivity as the between-subjects factor. Degrees of freedom in this and all subsequent analyseswere adjusted using the Greenhouse-Geisser method in cases where the assumption of sphericity is violated. This revealed a reliable effect of game (F(4.29, 206.1) = 20.46, MSE = .06, p < .001, η_p^2 = .30), which had a significant linear contrast (F(1, 48) = 51.76, MSE = .09, p < .001, η_p^2 = .52), demonstrating that participants learned to optimise the number of goods produced by running the machine slower as the games progressed. There was no reliable between-subjects effect of impulsivity (F(1, 48) = 2.16, MSE = .36, p = .15, η_p^2 = .04), indicating no difference in performance in the task between those high and low in impulsivity, nor was there an interaction (F(4.29, 206.1) = .32, MSE = .06, p = .87, η_p^2 = .01), indicating that the two groups improved on the task at the same rate.

To maximize their payoff, participants should switch from the lowest production level to the fastest production level between block 1 and block 2 at the end of each game. To test this and any effect of impulsivity on switching we compared mean production level in block 1 and block 2 across each game. These data were entered into an 8x2x2 mixed ANOVA with game and block as the

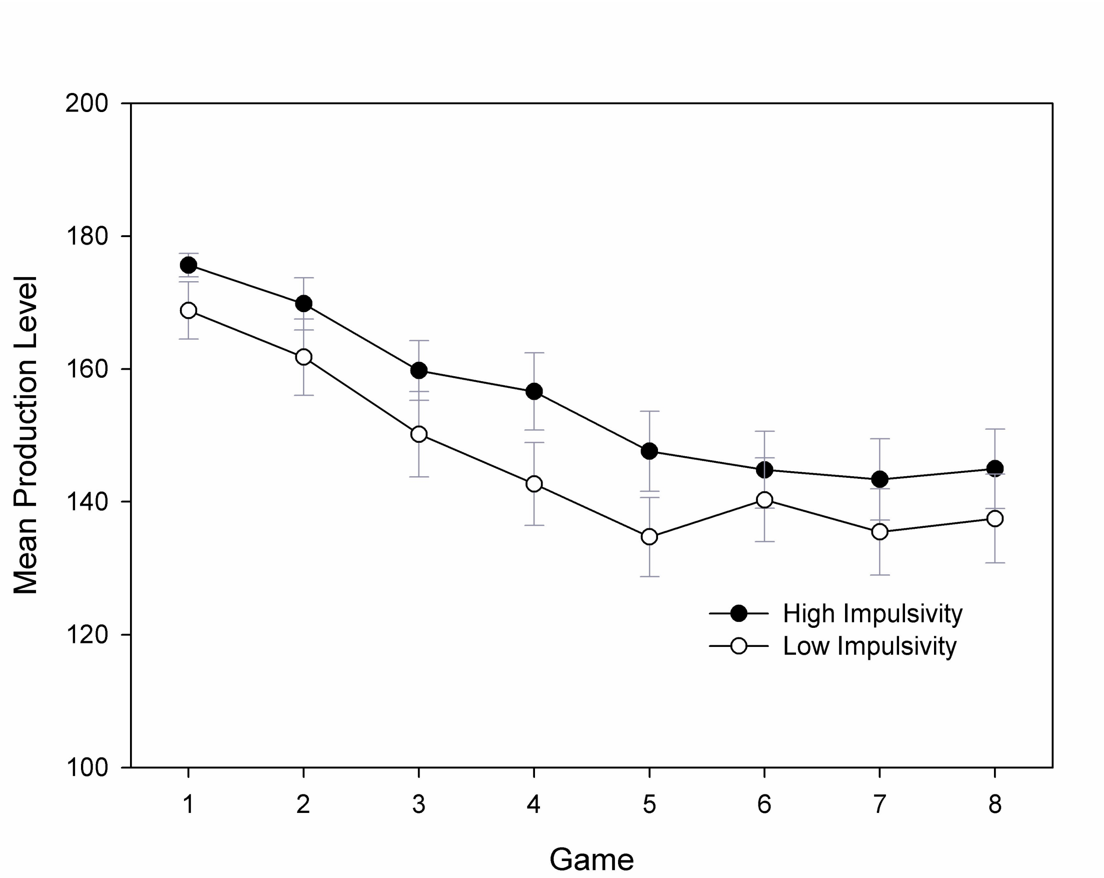

Figure 4.2: Learning in the Factory Manager game

within-subjects factors, and impulsivity as the between-subjects factor. There was an effect of block (F(1, 48) = 7.35, MSE = .05, p = .009, η_p^2 = .13), which shows that participants increased the production level at the end of each game. However, there was no interaction between block and game (F(5.23, 250.85) = 1.81, MSE = .02, p = .11, η_p^2 = .04), or between block and impulsivity (F(1, 48) = .15, MSE = .01, p = .7, η_p^2 < .01), and so there was no evidence that participants increasingly switched as the sessions progressed or that the two impulsivity groups reliably differed in their switching behaviour.

Advice Seeking

Participants asked for advice between 0 and 334 times (mean = 45, SD=60.7). The number of times participants in the two groups asked for advice was compared using Kendall's τ. This is a non-parametric correlation coefficient, which is applicable since the number of times that participants ask for advice is an ordinal variable (Howell, 2007). The test found a significant relationship ($\tau(50) = .34$, $p = .004$), providing some evidence that those high in impulsivity asked for more advice. This was not simply an artefact of the groupings we used, as a correlation between participants' BIS-11 scores and the number of times that they sought advice produced a similar, albeit marginally non-significant, result ($\tau(50) = .19$, $p = .059$).

4.2.4 Discussion

Our results demonstrate that those high in impulsivity sought more advice than those low in impulsivity. A very simple hypothesis (1) from this could be that participants who were high in impulsivity just clicked the advice button more and would have done so whatever the button did, and the BIS-11 does include a subtrait called "motor impulsiveness" (Patton et al., 1995). There are alternative hypotheses however. One could be individual differences in the perceived difficulty of the task (2). People have little reason to pay attention to others' advice when they find the task easy and believe that they are better than others, but when they perform worse than expected they are more likely to believe that others have useful information (Gino & Moore, 2007). It could be that those high in impulsivity found the task difficult, and so solicited more advice. Another explanation could be to do with the type of advice offered (3). Individuals may not see advice that is far from their initial expectations as credible (Yaniv, 2004). As we previously found that those high in impulsivity were more likely to choose myopically in the SHG, it may be that those low in impulsivity saw the generally impulsive advice as poor, and so did not continue asking for it.

Concerning performance in the Factory Manager version of the SHG, we found that switching did not increase as the experiments progressed – this was inconsistent with previous SHG data. This is probably due to the slider allowing participants to hedge rather than have to choose purely myopically or self-controlled. Similarly, there was no main effect of impulsivity, but again this may have been due to allowing participants to hedge and so reducing differences. Finally, participants did not improve on the task as much as in the SHG. Nevertheless, they do learn the task above chance level and they do switch, which demonstrates that they understood the contingencies, so as their performance on the task was not the main focus of our scrutiny, but rather their advice-seeking, we continued with the Factory Manager paradigm. Chapter 5 presents regressions across all of the studies in this thesis which include the BIS-11 questionnaire and a measure of performance on the SHG.

4.3 Experiment 8: Advice Seeking with Predominantly Self-Controlled Advice

4.3.1 Introduction

We tested whether those high in impulsivity still sought more advice when the advice was generally encouraging them to be self-controlled and so was far from their initial expectations. The Discussion section of Experiment 7 forwarded three hypotheses. If hypothesis 1 (impulsivity leads to clicking on advice button) or 2 (participants who find the task hard ask for more advice) was correct, participants high in impulsivity should still seek more advice no matter what the advice was. If hypothesis 3 (participants ask for advice that meets their expectations) was correct, participants high in impulsivity should not seek more advice.

4.3.2 Method

Participants

Fifty-seven students from Nottingham University volunteered to take part in the experiment after being recruited by posters and e-mails. Their mean age was 20.5 years (SD=2.4 years) and 74% were female.

Advice

The experiment was identical to the previous one, except that the advice predominantly suggested that participants should be self-controlled by setting the production level low. The three advisors' ideal production levels were 120%, 125% and 180%.

4.3.3 Results

Simplified Harvard Game

The mean production levels in block 1 of the 8 Factory Manager Games, split by the two impulsivity groups, were entered into an 8x2 mixed ANOVA with game as the within-subjects factor and impulsivity as the between-subjects factor. This revealed a reliable effect of game ($F(4.14, 227.9) = 18.24$, $MSE = .06$, $p < .001$, $\eta_p^2 = .25$), which had a significant linear contrast ($F(1, 55) = 52.03$, $MSE = .09$, $p < .001$, $\eta_p^2 = .49$), demonstrating that participants learned to optimise the number of goods produced by running the machine slower as the games progressed. There was no reliable between-subjects effect of impulsivity ($F(1, 55) = 1.18$, $MSE = .30$, $p = .28$, $\eta_p^2 = .02$), indicating no difference in performance in the task between those high and low in impulsivity, nor was there an interaction ($F(4.14, 227.9) = .86$, $MSE = .06$, $p = .49$, $\eta_p^2 = .02$), indicating that the two groups improved on the task at the same rate.

To maximize their payoff, participants should switch from the lowest production level to the highest production level between block 1 and block 2 at the end of each game. To test this and any effect of impulsivity on switching we compared mean production level in block 1 and block 2 across each game. These data were entered into an 8x2x2 mixed ANOVA with game and block as the within-subjects factors, and impulsivity as the between-subjects factor. There was an effect of block (F(1, 54) = 7.29, MSE = .13, p = .009, η_p^2 = .12), which shows that participants increased the production level at the end of each game. Unlike in the previous experiment, there was a reliable interaction between block and game (F(3.75, 202.43) = 4.59, MSE = .04, p = .002, η_p^2 = .08), demonstrating that participants increasingly switched as the sessions progressed. However, there was no interaction between block and impulsivity (F(1, 54) = .27, MSE = .13, p = .61, η_p^2 = .01), and so there was no evidence that the two impulsivity groups reliably differed in their switching behaviour.

Advice Seeking

Participants asked for advice between 3 and 188 times (mean = 37.5, SD=33.9). A non-parametric correlation between impulsivity group and the number of times that advice was sought did not find a reliable relationship (τ(57) = .16, p = .16). A correlation between participants' BIS-11 scores and the number of times that they sought advice produced a similar non-significant result (τ(57) = .14, p = .14).

4.3.4 Discussion

This experiment provided equivocal evidence that the type of advice given may have affected participants' advice-seeking behaviour, as those high in impulsivity did not seek more advice when the advice generally disagreed with their natural inclinations. Nevertheless, the significance value was not far from borderline, and so we continued this line of enquiry by conducting another experiment using a within-subjects design.

Of note is that we did find increasing switching as the games progressed this time, consistent with previous research using the SHG.

4.4 Experiment 9: Within Subjects Advice Seeking

4.4.1 Introduction

We tested whether the type of advice (either self-controlled or impulsive) that participants were given influenced their advice-seeking behaviour. In order to reduce the inclination for participants to click on the advice button indiscriminately, we added a cost to the advice. This was expected to have the additional benefit of reducing the strongly positively skewed data on the number of times that participants ask for advice. This was useful because we were testing an interaction between advice type and individuals' impulsivity, which requires an ANOVA with normally distributed data.

4.4.2 Method

Participants

Forty-nine students from Nottingham University volunteered to take part in the experiment after being recruited by an online experiment system. Their mean age was 20.4 years (SD=1.6 years) and 59% were female.

Procedure and Advice

The experimental procedure was identical to the previous one, except that the advice system changed. There were four advisors, two of whom predominantly recommended that the participant should set the production level high, so maximizing their short-term gain, and two of whom recommended setting the production level low, so as to be self-controlled. The independent variable was how often participants sought advice from the two impulsive advisors and from the two self-controlled advisors.

The self-controlled advisors' ideal production levels were 119.5, 124.5, and the impulsive advisors' ideal production levels were 175.5 and 180.5. This averages out to setting the production level at 150, the mid-point of the scale. The reason that non-integers were chosen was because the machine's production level was always an integer, so participants were never advised that they had set the production level "about right", thus they were always told to either run the machine faster or slower. This prevented the situation where a participant could quickly learn what an advisor's ideal speed was just by chance, and not have to ask again. Unlike in previous experiments, the position of the advisors was randomised. This was because participants tend to begin by asking advisor 1, and so since in this experiment we were measuring which advisors they asked, as well as how often, it was crucial that there was no positioning bias.

A cost was also added to the advice, so that each time a participant asked for advice it would cost them 0.5p. This was set as low enough that participants would not be too discouraged from asking for advice.

The instructions to participants were changed to reflect the new advice system. The advice paragraph was changed to:

> "You have also been given access to 4 advisors, who previously played the factory manager game like you are about to. After finishing, they left a recommendation for the speed that you should run the machine in order to produce as many Goods as possible. Each time you ask for their advice, they will tell you whether the speed that you last ran the machine is slower or faster than their ideal speed. Each time you ask for advice it will also cost you 0.5p. Ultimately it will be your decision to set the machine speed to whatever you want."

4.4.3 Results

Simplified Harvard Game

The mean production levels in block 1 of the 8 Factory Manager Games, split by the two impulsivity groups, were entered into an 8x2 mixed ANOVA with game as the within-subjects factor and impulsivity as the between-subjects factor. This revealed a reliable effect of game (F(3.8, 180.6) = 23.31, MSE = .07, p < .001, η_p^2 = .33), which had a significant linear contrast (F(1, 47) = 53.86, MSE = .10, p < .001, η_p^2 = .53), demonstrating that participants learned to optimise the number of goods produced by running the machine slower as the games progressed. There was no reliable between-subjects effect of impulsivity (F(1, 47) = 2.43, MSE = .31, p = .13, η_p^2 = .05), indicating no difference in performance in the task between those high and low in impulsivity, nor was there an interaction (F(3.84, 180.6) = .90, MSE = .07, p = .46, η_p^2 = .02), indicating that the two groups improved on the task at the same rate.

To maximize their payoff, participants should switch from the lowest production level to the highest production level between block 1 and block 2 at the end of each game. To test this and any effect of impulsivity on switching we compared mean production level in block 1 and block 2 across each game. These data were entered into an 8x2x2 mixed ANOVA with game and block as the within-subjects factors, and impulsivity as the between-subjects factor. There was an effect of block (F(1, 47) = 8.27, MSE = .09, p = .006, η_p^2 = .15), which shows that participants increased the production level at the end of each game. There was a reliable interaction between block and game (F(3.47, 245.86) = 3.47, MSE = .03, p = .004, η_p^2 = .07), demonstrating that participants increasingly switched as the sessions progressed. However, there was no interaction between block and impulsivity (F(1, 47) = .15, MSE = .09, p = .70, η_p^2 < .01), and so there was no evidence that the two impulsivity groups reliably differed in their switching behaviour.

Advice Seeking

Overall, the number of times that participants asked for advice was positively skewed with a mean comparable to its standard deviation (mean = 8.65, SD = 7.8, skewness = 1.1, range 0 – 31). Since our hypothesis tested an interaction, this is not possible with a non-parametric ANOVA, and so it was necessary to have normally distributed data. Howell (2007) recommends that for data that are counts, where the mean is proportional to the variance, and where values of the counts are less than 10, then the square root transformation in Equation 4.4.1 is appropriate. This was applied and reduced the skewness to within acceptable levels (mean = 2.7, SD = 1.3, skewness = .3, range = .71 – 5.61).

$$Y = \sqrt{X + 0.5} \tag{4.4.1}$$

Participants could solicit advice from four advisors, two of whose advice was to run the machine at a fast pace, which would maximize the short-term outcome, and two of whose advice was to run the machine slower, which would equate

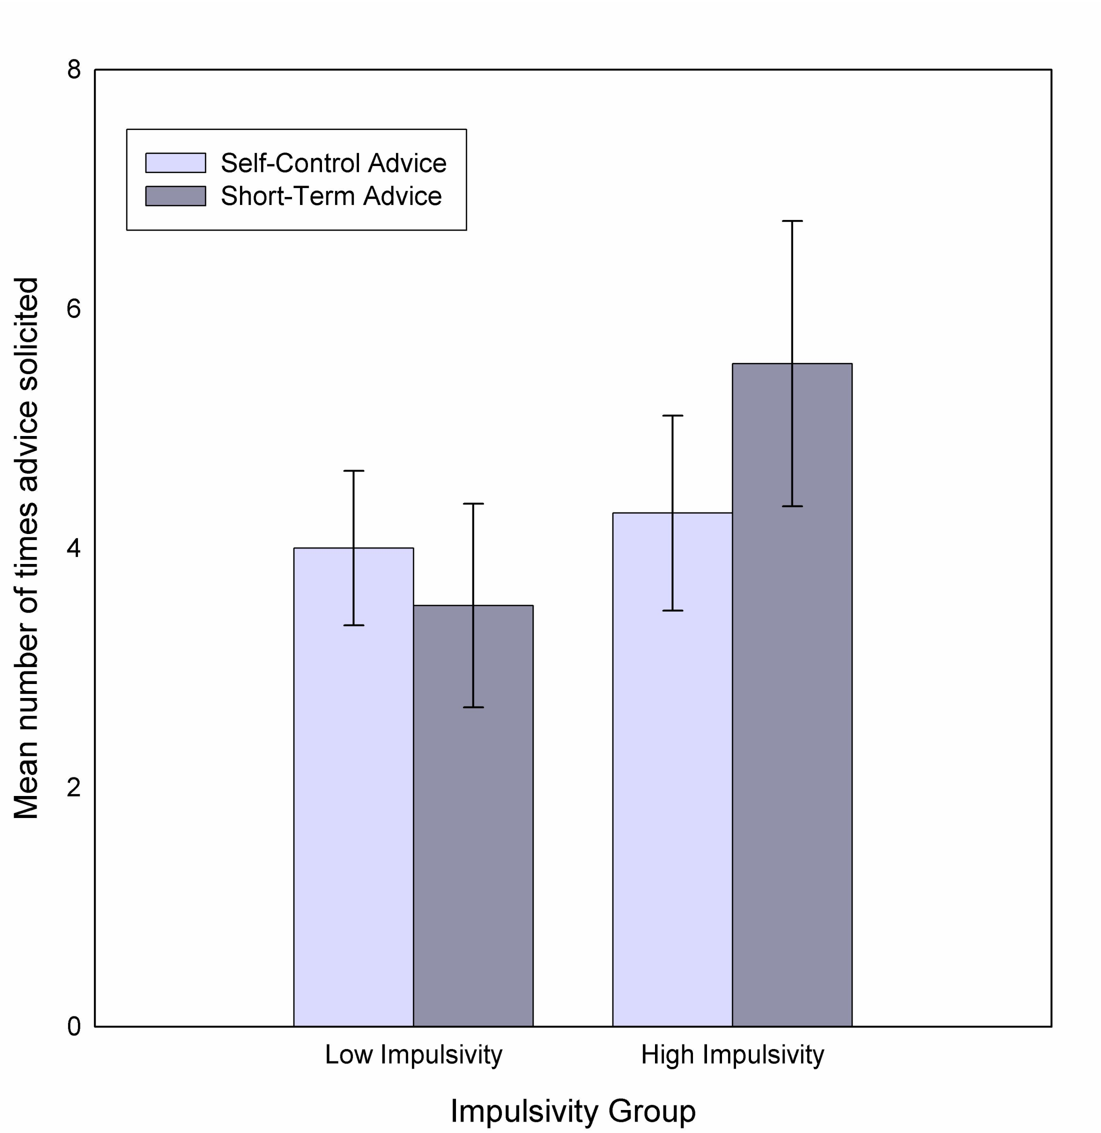

Figure 4.3: Advice seeking in the Factory Manager game

to being self-controlled. We compared the transformed number of times that the two impulsivity groups (high/low) asked for advice from the two types of advisor (short-term / self-controlled). These data were entered into a 2x2 repeated-measures ANOVA with advice type as the within-subjects factor and impulsivity group as the between-subjects factor. There was no main effect of advice type ($F(1, 47) = .01$, $MSE = .33$, $p = .918$, $\eta_p^2 < .01$), indicating that both types of advisor were asked a similar number of times overall. Nor was there a main effect of impulsivity group (F(1, 47) = .85, MSE = 1.5, p = .36, η_p^2 = .02), indicating that the two impulsivity groups asked for advice a similar number of times. However, the interaction was significant (F(1, 47) = 4.11, MSE = .33, p = .048, η_p^2 = .08), indicating that the two groups asked for different types of advice more often. The mean number of untransformed times that the two impulsivity groups solicited advice from the two types of advisor is shown in Figure 4.3. It shows that the two groups asked for advice from the self-controlled advisor a similar number of times, but that the number of times they solicited advice from the short-term advisor differed markedly, with those high in impulsivity asking more often.

4.5 Experiment 10: Advice Seeking by Experience

These data were collected by Claire Watson as part of her third year undergraduate research project. The study was conceived by me, although refined by both of us, and the analysis presented here is my own.

4.5.1 Introduction

It has previously been found in JAS experiments that advisors who judges believe to be experts have their advice weighted more heavily (Harvey & Fischer, 1997; Sniezek et al., 2004). We test whether the propensity for judges to seek advice from experts or novices is related to their trait impulsivity. A naïve hypothesis is that those high in impulsivity may not weight the advice from experts as heavily as those low in impulsivity.

4.5.2 Method

Participants

Forty-eight students from Nottingham University were opportunity sampled. Their mean age was 21.9 years (SD=.84 years) and 46% were female.

Procedure and Advice

The experimental procedure was largely identical to the previous one, except that the recompense and advice systems changed. In this experiment participants were not paid based on their performance in the task. Nevertheless, the task has previously been found to be intrinsically motivating, we found similar results whether participants are paid contingent upon their performance or not (see Experiments 5 & 6).

There were three advisors, each with a different level of experience of playing the SHG, for 2 minutes, 20 minutes or 180 minutes. The independent variable was how often participants sought advice from the three advisors with different levels of experience.

The advisors' ideal production levels were randomized, as was their order in the game as advisor 1, advisor 2 and advisor 3. The ideal production levels were 119.5, 150.5 and 180.5. A cost was also added to the advice, so that each time a participant asked for advice it would cost them 5 goods. This was set as low enough that participants would not be too discouraged from asking for advice.

The instructions to participants were changed to reflect the new advice system. The advice paragraph was changed to:

> You have also been given access to advice from three people who were previously factory managers like you, who played the factory manager game for 2 minutes, 20 minutes or 180 minutes. Afterward, they left a recommendation for the speed that you should run the machine. However you will only be told whether the speed that you last set the machine at is slower or faster than they recommend. Ultimately it is your decision to set the machine speed to whatever you want. It is also your decision as to how many times you ask for advice, but there will be a 5 Good ”cost” each time you do.

4.5.3 Results

Simplified Harvard Game

The mean production levels in block 1 of the 8 Factory Manager Games, split by the two impulsivity groups, were entered into an 8x2 mixed ANOVA with game as the within-subjects factor and impulsivity as the between-subjects factor. This revealed a reliable effect of game (F(4.4, 200.7) = 7.8, MSE = .07, $p < .001$, $\eta_p^2 = .15$), which had a significant linear contrast (F(1, 46) = 17.96, MSE

= .10, p < .001, = .28), demonstrating that participants learned to optimise the number of goods produced by running the machine slower as the games progressed. There was no reliable between-subjects effect of impulsivity (F(1, 46) = 1.44, MSE = .37, p = .24, η_p^2 = .03), indicating no difference in performance in the task between those high and low in impulsivity, nor was there an interaction (F(4.36, 200.7) = .60, MSE = .07, p = .67, η_p^2 = .01), indicating that the two groups improved on the task at the same rate.

To maximize their payoff, participants should switch from the lowest production level to the highest production level between block 1 and block 2 at the end of each game. To test this and any effect of impulsivity on switching we compared mean production level in block 1 and block 2 across each game. These data were entered into an 8x2x2 mixed ANOVA with game and block as the within-subjects factors, and impulsivity as the between-subjects factor. There was an effect of block (F(1, 46) = 10.5, MSE = .12, p = .002, η_p^2 = .19), which shows that participants increased the production level at the end of each game. There was a reliable interaction between block and game (F(5.03, 231.47) = 5.10, MSE = .03, p < .001, η_p^2 = .10), demonstrating that participants increasingly switched as the sessions progressed. However, there was no interaction between block and impulsivity (F(1, 46) = .08, MSE = .12, p = .78, η_p^2 < .01), and so there was no evidence that the two impulsivity groups reliably differed in their switching behaviour.

Advice Taking

Overall, the number of times that participants asked for advice was again positively skewed with a mean comparable to its standard deviation (mean = 10.3, SD = 8.4, skewness = 1.3, range 0 – 33). We again transformed the data using the same square root transform as in the previous experiment. This reduced the skewness to within acceptable levels (mean = 3.0, SD = 1.2, skewness = .4, range .71 – 5.79).

Participants could solicit advice from three advisors, each of whom was said to have differing experience playing the Factory Manager Game. We compared the transformed number of times that the two impulsivity groups (high/low) asked for advice from the three advisors (2 mins, 20 mins, 180 mins). These data were entered into a 2x3 repeated-measures ANOVA with advice type as the within-subjects factor and impulsivity group as the between-subjects factor. There was a main effect of advice type (F(1.78, 82.1) = 23.16, MSE = .623, p < .001, η_p^2 = .34), with a reliable linear contrast (F(1, 46) = 37.45, MSE = .67, p < .001, η_p^2 = .45) indicating that participants asked for more advice from those with more experience. There was not a main effect of impulsivity group (F(1, 46) = .92, MSE = 1.1, p = .34, η_p^2 = .02), indicating that the two impulsivity groups asked for advice a similar number of times. Nor was there an interaction (F(1.78, 82.1) = .83, MSE = .62, p = .43, η_p^2 = .02), indicating that the two groups did not differ in the type of advice that they sought. The mean number of untransformed times that the two impulsivity groups solicited advice from the two types of advisor is shown in Figure 4.4. It shows that the two groups both asked for advice from the more experienced advisors more.

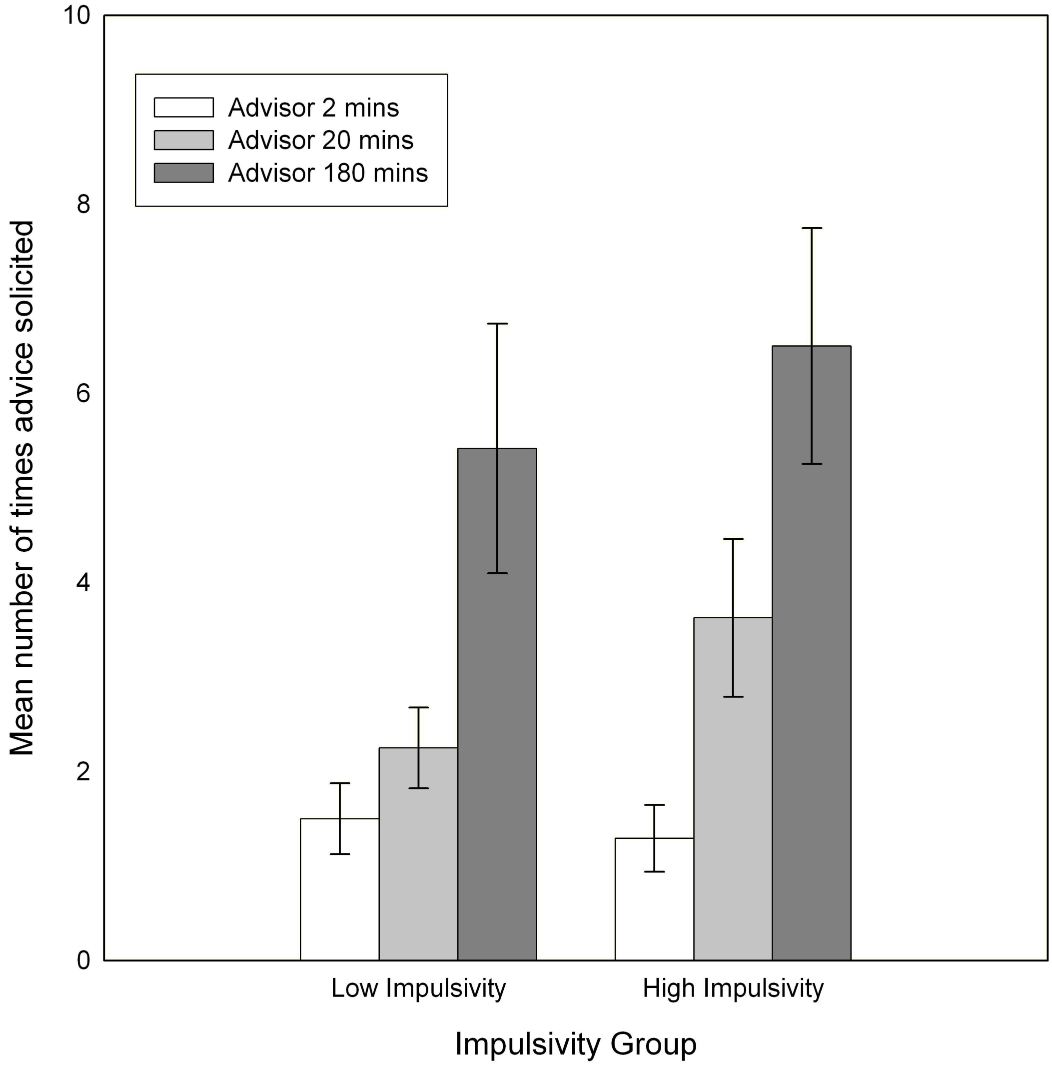

Figure 4.4: Advice type sought in the Factory Manager game

4.5.4 Discussion

We find no evidence that impulsivity affects the relative importance given to seeking advice from experts or novices, although it could be that the effect is small and the statistical power of the experiment was not enough to find it.

4.6 Experiment 11: Advice Taking

4.6.1 Introduction

Egocentric advice discounting has been found to increase as the distance between judges' initial opinions and advisors' recommendations increases (Yaniv, 2004). This, combined with evidence from Experiment 10 that judges continue to seek advice from those whose opinions agrees with their initial views, might suggest that providing participants with the same advice will differentially affect their performance in the Factory Manager Game. We hypothesise that those high in impulsivity will have their performance in the task particularly badly hampered by advice telling them to run the machine quickly and so act myopically. Based on Experiment 10, we expected no difference when participants were given advice to be self-controlled.

4.6.2 Method

Participants

Forty-five students from Nottingham University participated after being recruited by an online system. Their mean age was 21.5 years (SD=3.3 years) and 80% were female.

Procedure and Advice

Participants played exactly the same Factory Manager Game as in previous experiments, and received recompense contingent on their performance. However, participants did not ask for advice but were instead periodically given it during the game.

Participants were allocated to one of two conditions, based on the order that they participated with consecutive participants allocated to alternate conditions. In the self-controlled advice type condition, the advice given told the participant to run the machine relatively slowly, with advice randomly selected between 100% and 130%. In the impulsive advice type condition, the advice given was randomly selected between 170% and 200%, so running the machine relatively quickly and maximizing short-term gain at the expense of the long-term. In both conditions, advice was displayed three times in each game, when the participant went below 90, 55 and then 20 game units. This equated to receiving unsolicited advice 24 times. The advice appeared in front of the Factory Manager console, so participants had to acknowledge it by pressing OK before they could continue.

The instructions to participants were changed to reflect the new advice system. The advice paragraph was changed to:

> You will also be given advice from people who previously played the Factory Manager Game. After playing, each left a recommendation for the speed that you should run the machine in order to

produce as many Goods as possible. Advice from these people will pop up throughout, but you do not have to follow it and ultimately it is your decision to set the machine speed to whatever you want.

4.6.3 Results

Simplified Harvard Game

The mean production levels in block 1 of the 8 Factory Manager Games, split by the two impulsivity and advisor type groups, were entered into an 8x2x2 mixed ANOVA with game as the within-subjects factor and impulsivity and advisor type as the between-subjects factors. This revealed a reliable effect of game (F(4.6, 190.2) = 16.8, MSE = .05, p < .001, η_p^2 = .29), which had a significant linear contrast (F(1, 41) = 51.5, MSE = .07, p < .001, η_p^2 = .56), demonstrating that participants learned to optimise the number of goods produced by running the machine slower as the games progressed. There was no reliable between-subjects effect of impulsivity (F(1, 41) = .93, MSE = .41, p = .34, η_p^2 = .02), indicating no difference in performance in the task between those high and low in impulsivity, nor was there an interaction (F(4.64, 190.2) = 1.34, MSE = .05, p = .25, η_p^2 = .03), indicating that the two groups improved on the task at the same rate. There was however a main effect of advice type (F(1, 41) = 11.05, MSE = .41, p = .002, η_p^2 = .21), indicating that those who received self-control advice subsequently ran the machine slower and produced more goods in the long run. But there was no interaction between advice type and game (F(4.64, 190.2)

= .11, MSE = .01, p = .99, $\eta_p^2 < .01$), indicating that no matter what advice participants were given they improved on the task at the same rate. Finally, there was no three-way interaction between advice type, game, and impulsivity (F(4.64, 190.2) = .40, MSE = .02, p = .84, $\eta_p^2 = .01$), indicating that no matter what impulsivity group participants were in, they used the advice in the same way to change their performance on the task. Figure 4.5 shows the mean production level in block 1 across all 8 Factory Manager Games for each of the 4 groups.

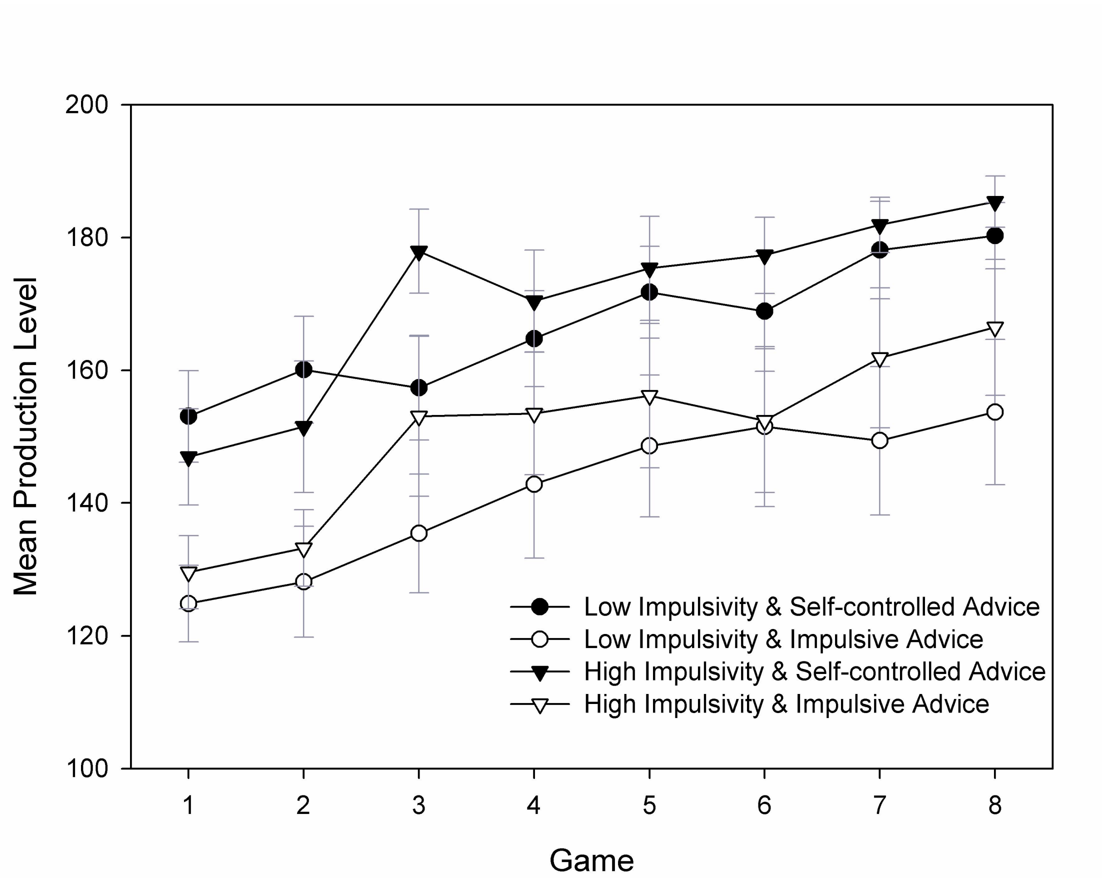

Figure 4.5: Performance in the Factory Manager game by impulsivity and advice-type groups

To maximize their payoff, participants should switch from the lowest production level to the highest production level between block 1 and block 2 at the end of each game. To test this and any effect of impulsivity or advice type on switching we compared mean production level in block 1 and block 2 across each game. These data were entered into an 8x2x2x2 mixed ANOVA with game and block as the within-subjects factors, and impulsivity and advice type as the between-subjects factors. There was an effect of block ($F(1, 41) = 22.2$, $MSE = .13$, $p < .001$, $\eta_p^2 = .35$), which shows that participants increased the production level at the end of each game. There was a reliable interaction between block and game ($F(4.93, 201.9) = 3.83$, $MSE = .04$, $p = .003$, $\eta_p^2 = .09$), demonstrating that participants increasingly switched as the sessions progressed. However, there was no interaction between block and impulsivity ($F(1, 41) = .01$, $MSE = .13$, $p = .93$, $\eta_p^2 < .01$), or between block and advice type ($F(1, 41) = 1.04$, $MSE = .13$, $p = .31$, $\eta_p^2 = .03$) and so there was no evidence that the two levels of trait impulsivity or the advice type they were given reliably changed their switching behaviour.

4.6.4 Discussion

We found no reliable relationship between trait impulsivity and utilisation of different types of advice. This may suggest that impulsivity does not affect the reputation formation system that is used to assess advisors' advice (Yaniv & Kleinberger, 2000), such that both impulsivity groups learned to discount the impulsive advisors' advice to the same degree.

A limitation is that it is possible that as advice was provided for free and was unsolicited, that this may have caused participants to underweight its value (Gibbons, Shiezek & Dalal, 2003, cited in Bonaccio & Dalal, 2006; Gino et al., 2009). However, the consistent and long-lasting difference in performance in the Factory Manager Game between groups given self-control advice and groups given impulsive advice disagrees with this conclusion.

4.7 Chapter Discussion

Perhaps surprisingly, impulsivity mediated who advice was sought from, such that advice was sought from advisors who agreed with the judge's initial inclinations, but when it came to actual performance on the self-control task, judges did not underweight advice from advisors who disagreed with their inclinations. This supports the importance in the literature of separating advice-seeking from advice-utilisation. It also supports the relatively new interest in individual differences in JAS experiments. The evidence suggests that those who are high in impulsivity, and thus more likely to engage in myopic behaviours such as smoking and drinking, will tend to seek advice from those whose advice agrees with what they already think. However, they will solicit the opinions of experts, and when they are given unsolicited advice, this will sway their behaviour as much as those low in impulsivity. These would suggest that policy makers should concentrate on making as much advice available as possible from credible sources, in order that it counteracts the natural tendency of those high in impulsivity to seek the opinions of people who they know are likely to tell them to act impulsively.

There is also contradictory findings in the literature concerning whether judges solicit free advice (Gardner & Berry, 1995; Schrah et al., 2006). In most of my advice-seeking experiments, there were participants who, despite being told about the advice and even being given a tutorial where they had to ask for it, did not ask for advice at all during the task. This is puzzling, given that the cost to soliciting advice was so low, and it is worthy of follow-up study to understand who these independent judges are.

CHAPTER 5

Individual Differences Related to Intrapersonal Externalities

5.1 Chapter Introduction

Through the course of this thesis, there have been seven experiments using intrapersonal externalities. Participants in all of these experiments completed the Barratt Impulsivity Scale. The following chapter analyses the relationship between trait impulsivity and performance on tasks with intrapersonal externalities.

5.2 Experiment 12: BIS-11 Impulsivity and Intrapersonal Externalities

5.2.1 Introduction

Experiment 3 found that performance on the Simplified Harvard Game was related to smoking status, independently from trait impulsivity. This agrees with other research that has compared psychometric measures of impulsivity with behavioural measures, such as delay discounting or the Balloon Analog Risk Task (Crean, de Wit, & Richards, 2000; S. H. Mitchell, 1999; Reynolds et al., 2006; Reynolds, Richards, Horn, & Karraker, 2004; J. L. White et al., 1994), although some other experiments do show links (Kirby et al., 1999; Richards, Zhang, Mitchell, & de Wit, 1999; Swann, Bjork, Moeller, & Dougherty, 2002).

In the individual experiments it was not advisable to compare performance on the tasks with intrapersonal externalities with trait impulsivity scores, because both are noisy variables with measurement error, so fifty participants would not provide enough experimental power; however, by combining all of the intrapersonal externality tasks the pattern should be clear.

5.2.2 Method

Participants were pooled from Experiments 3, 4, 7, 8, 9, 10 and 11. This totalled 351 participants who had completed a task with intrapersonal externalities. Participants gave their e-mail addresses when they participated in an experiment, which were used to ensure that there were no repeat participants across experiments. The characteristics of participants are shown in Table 5.1. There are more participants in Experiment 4 than were included in the report of the experiment because some smokers participated but later inspection revealed that they had not smoked long enough (more than 2 years) to meet the criterion for the experiment, and so they were excluded.

Table 5.1: Summary of individual characteristics of participants in intrapersonal externalities experiments

Exp. Number	N	% Male	Mean Age/yrs (SD)	% Smoker	Mean BIS-11 (SD)
3	51	45	21.6 (4.2)	100	67.3 (10.5)
4	51	39	20.6 (2.4)	100	66.9 (9.8)
7	50	36	23.5 (5.2)	12	63.6 (8.0)
8	57	26	20.5 (2.4)	4	65.0 (9.3)
9	49	41	20.4 (1.6)	16	63.2 (9.6)
10	48	54	21.0 (0.8)	29	67.3 (11.1)
11	45	20	21.3 (3.3)	13	65.4 (9.8)
Total	351	37	21.3 (2.9)	39	65.5 (9.7)

5.2.3 Results

To test the link between performance in experiments with intrapersonal externalities and impulsivity, a regression was run with the proportion of maximising responses in the task (pmax) as the dependent variable. The predictors were experiment number as dummy categorical variables, age, gender, smoker type (1 = daily, 0 = non-daily or non-smoker), and BIS-11 impulsivity. Table 5.2 summarises the results of the personality and demographic predictors; impulsivity was not a significant predictor.

The regression was repeated twice; with the six first order factors (Table 5.3),

Table 5.2: Regression coefficients showing the relationship between demographics, smoker type, and impulsivity, with performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.07	-1.37	0.17
Age	-0.07	-1.18	0.24
Smoker Type	0.00	0.06	0.95
Impulsivity (BIS-11)	0.01	0.25	0.81

and with the three second order factors (Table 5.4). Both analyses found no subfactors related to performance on the tasks.

Table 5.3: Regression coefficients showing the relationship between demographics, smoker type, and BIS-11 first order factors, with performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.08	-1.38	0.17
Age	-0.06	-1.12	0.26
Smoker Type	0.03	0.33	0.74
BIS-11 Attention	0.02	0.25	0.80
BIS-11 Motor	0.03	0.57	0.57
BIS-11 Self-Control	0.07	1.06	0.29
BIS-11 Cognitive Complexity	-0.06	-0.98	0.33
BIS-11 Perseverence	-0.05	-0.95	0.34
BIS-11 Cognitive Instability	-0.04	-0.64	0.52

Three further analyses were conducted to exhaust all possibilities, but all supported the above findings. In order to exclude the possibility that extra noise was added by mixing Simplified Harvard Games with Factory Manager Games, the above analyses were repeated but only for the Factory Manager Games (Experiments 7-11; reported in Appendix D). In order to exclude the possibility that a significant effect was not found because smoking group accounted for variance that would have been accounted for by impulsivity, the above analyses were repeated without smoking group (reported in Appendix E). In an abundance of caution, the above analyses were repeated using as the depen-

Table 5.4: Regression coefficients showing the relationship between demographics, smoker type, and BIS-11 second order factors, with performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.08	-1.40	0.16
Age	-0.07	-1.21	0.23
Smoker Type	0.03	0.32	0.75
BIS-11 Attentional Impulsiveness	-0.01	-0.15	0.88
BIS-11 Motor Impulsiveness	0.00	0.05	0.96
BIS-11 Nonplanning Impulsiveness	0.02	0.31	0.75

dent variable the z-scores of performance on the task standardized within an individual experiment. This has the advantage that using dummy variables to take account of between-experiment variance is unnecessary (reported in Appendix F). All analyses found no link between BIS-11 trait impulsivity and performance on tasks with intrapersonal externalities.

5.2.4 Discussion

Given the number of statistical analyses reported here, it is remarkable that there were none where trait impulsivity or a subfactor of the BIS-11 scale were predictive of performance on tasks with intrapersonal externalities. This provides consistent evidence that there is no link between performance on tasks with intrapersonal externalities and impulsivity. Although there was no relationship between smoking status and performance on the tasks, this was not unexpected because Experiment 3 found that the distinction between morning smokers and morning abstaining smokers was most indicative of performance, whereas these analyses compared smokers with non-smokers.

These results add to the growing body of research which finds no relationship between psychometric measures of impulsivity and behavioural measures (Crean et al., 2000; S. H. Mitchell, 1999; Reynolds et al., 2006, 2004; J. L. White et al., 1994). These results also support those from Experiment 3, which found that performance on the Simplified Harvard Game and trait impulsivity were independent predictors of smoking status.

Reynolds et al. (2006) analysed four behavioural tasks measuring aspects of impulsivity. Principle components analysis found two components, the first included the Stop Task and the Go/No-Go task. The second was composed of a delay discounting task and the Balloon Analog Risk Task. The authors explain the components by suggesting that the first is related to impulsive disinhibition, because the challenge in the two tasks is to withhold a response, and the second is related to decision-making, because the individual makes a conscious decision. It would be useful to find out where the intrapersonal externalities tasks fit into this model. The second component seems most clearly related; on the one hand, a conscious decision is made on the task, but on the other hand Experiment 3 showed a weak, statistically non-significant, link between delay discounting and choices in the Simplified Harvard Game. Intrapersonal externalities could reflect a separate differential learning component.

CHAPTER 6

General Discussion: Do intrapersonal externalities improve our understanding of addicted behaviour?

6.1 Why do only some people become addicts?

Approximately 15-17% of humans who try a drug become addicted, as defined by pathological use in the face of adverse consequences (Anthony et al., 1994). Well-controlled animal research found a similar percentage (Deroche-Gamonet et al., 2004), indicating that even when individuals experience the same stimuli, there is an internal process which determines whether the individual will go on to becoming addicted.

Two literatures have a bearing on the question of why some people go on to become addicts. From economics, delay discounting assumes that individuals choose in their best interest; the reason that some people smoke is because they steeply discount the long-term future consequences of a decision. From psychopharmacology, models of impulsivity assume that some individuals are compelled to choose an option that they realise is not in their best interest.

A third theory, of intrapersonal externalities, takes a middle road. It draws on research showing that individuals attempt to make the best choices possible, but sometimes they can make systematically irrational decisions. The theory of intrapersonal externalities hypothesises that an individual will not take into account externalities imposed on its future self. Payoffs from decisions with intrapersonal externalities include both long-term and short-term consequences. The short-term consequences are easy to associate with the decision, but the long-term consequences are diffuse and therefore difficult to link with the decision that caused them; therefore, such consequences are difficult to learn because an individual's cognitive architecture is not up to the task (Stillwell & Tunney, 2009).

As intrapersonal externalities are abstract, they explain a range of situations. Typically, behaviours with negative intrapersonal externalities (such as smoking) are over-performed, and those with positive intrapersonal externalities (such as exercising) are under-performed. There are two questions which limit the use of intrapersonal externalities in research. First, are our current explanations of addiction sufficient for describing addicts and for suggesting interventions to solve addictions? Second, little work has been done to validate the theory of intrapersonal externalities, so do they relate to addictions in the real world?

6.2 Intrapersonal externalities are a valid and flexible tool to study addictions

The two experiments in Chapter 2 used individual difference measures which measured aspects of the two broad theoretical areas of why addictive behaviour can be difficult to overcome. In the behavioural economics literature, delay discounting is assumed to measure how willing individuals are to wait for a reward. Delay discounting assumes that individuals make choices in their best interest but future rewards are discounted according to a personal discounting function. Some individuals steeply discount the future and so choose myopically. On the other hand, a cluster of theories from the psychopharmacological literature, which can collectively be described as impulsivity, posit that individuals do not choose in their best interest. This could be because they are unable to inhibit their impulses, or because their cognitions are present-oriented.

In Experiment 1, which investigated the relationship between individual differences and addictive behaviours, demographics were controlled for and big five personality was measured because they are fundamental and broad individual differences which are important across a range of behaviours (Barrick, Mount, & Judge, 2001; Ozer & Benet-Martinez, 2006). Then, trait impulsivity was measured using a broad questionnaire inventory (the Barratt Impulsivity Scale, BIS-11; Stanford et al., 2009), and delay discounting was measured using a behavioural task. The results showed that although the impulsivity and delay discounting measures were reliably related to a range of impulsive behaviours, the percentage of variance accounted for by all of the demographic, personality, impulsivity and discounting measures was in some cases as low as 3.9% (for alcohol dependency), and never higher than 20.9% (for drug dependency). Although there is measurement error from the questionnaires and behavioural tasks, which if accounted for would boost the variance explained, there is space for additional explanation.

In experiment 2, the method of measuring delay discounting was investigated in detail. It was possible that delay discounting was confounded with impulsivity by its method of measurement, and therefore the relationship between delay discounting and addictive behaviours could have been illusory. When individuals complete a delay discounting questionnaire, it is possible that impulsive individuals would switch earlier than self-controlled individuals simply to finish the questionnaire earlier. The results showed that delay discounting was related to smoking, no matter how it was measured, and with an equal effect size. This provides strong support that the amount that a reward's value is discounted the longer the individual has to wait for it is an independent predictor of addictive behaviour.

Chapter 3 took the first step to validating whether a task which measured intrapersonal externalities is related to real-world addictive behaviour. In the

Simplified Harvard Game (SHG), participants chose between two options. One option gave a higher immediate reward, but reduced the payoffs from both options over the next ten choices. Ultimately, the best strategy was to choose the button with the lowest immediate reward. In Experiment 3, participants who did not take intrapersonal externalities into account and so performed poorly on the SHG, were more likely to be daily smokers than non-daily smokers. In Experiment 4, the Credit Card Game, participants were given points at the beginning of each game, but lost points at a differential rate depending upon their choices in the task. The best strategy was to choose the button with the highest immediate loss, because the lowest immediate loss would prolong the task and ultimately lead to a greater loss. This mirrors the choices made each month when a credit card bill arrives. An individual can either pay off the full balance, taking a large hit immediately, or pay off the minimum amount but accepting the interest charge. The results from the experiment were that there was no relationship between smoking and intrapersonal externalities when the SHG was presented as one of avoiding costs. Taken together, these experiments indicate that the reason that smokers find it difficult to learn about intrapersonal externalities is due to a difficulty in learning about small gains.

Experiments 5 and 6 explored participants' understanding of the Simplified Harvard Game by asking questions that probed their knowledge of the payoff schedules, and found that participants who performed well on the task explicitly learned that one option made their situation worse, but did not understand the full schedule of payoffs. This result suggests a way to change behaviour, by telling an individual the full consequences of their behaviour immediately after (or before) making a decision. Current anti-smoking campaigns emphasise the risk to the *average person*. Individuals may respond better to personalised feedback based on their individual risk of developing a disease. This could be calculated and presented through a smartphone application.

Having established the validity of intrapersonal externalities as a marker for addictive behaviour, Chapter 4 used the Factory Manager Game to study advice-seeking and advice-taking. The Factory Manager Game differs from the Simplified Harvard Game in two important ways. First, by allowing the participant to choose a production level on a sliding scale, it was possible to advise the participant on the speed to set the machines without having to give away the fact that the ultimate payoff was determined by the history of choices. Second, setting the speed of a machine is a concrete task and therefore more amenable to providing advice.

Theoretically, using intrapersonal externalities to study advice-seeking and advice-taking is a new development because delay discounting assumes that the individual already has an internal discounting function which is presumably largely immune to outside influence, and impulsive disinhibition states that an individual makes a choice outside of their control. The effect of advice is consequently difficult to study under these paradigms in the laboratory. In contrast, the challenge of the Factory Manager Game is to learn to maximize income with limited knowledge and experience, and so advice can help or hinder.

The experimental results of advice seeking were that impulsivity mediated who advice was sought from; judges preferred to ask advisors who agreed with their initial inclinations. By contrast, when advice was mandatory, judges did not underweight advice from advisors who disagreed with their inclinations. Practically, this supports the UK government's current policy, to advertise advice from credible sources in order to counter individuals' natural tendency to seek advice from those who support their impulsive inclinations.

Chapter 5 analysed the results of all Simplified Harvard Games and Factory Manager Games so that trait impulsivity could be compared with performance in the games with a large sample size. The results showed no relationship between trait impulsivity as measured by the Barratt Impulsivity Scale (Stanford et al., 2009) and performance in the task. Additionally, in Experiment 3, there was no relationship between delay discounting and trait impulsivity. Exploration of the impulsivity literature found that it is common to find no relationship between questionnaire measures and behavioural measures of impulsivity with modest sample sizes ($N = 70$; Reynolds et al., 2006, cites 5 other instances as well as their own), despite different questionnaire measures correlating with one another. One explanation for this discrepancy is that the behavioural measures are narrow but objective, whereas the self-report measures are subjective but broad and so are able to account for a range of motives behind impulsive behaviours (Reynolds et al., 2006).

6.3 Further theories of addictive behaviour

This thesis concentrated on impulsivity, delay discounting and intrapersonal externalities, because these are broad constructs which account for a range of addictive behaviours, because they do not theoretically overlap, and because they are important in the literature and it is not practically possible to measure all constructs. Nevertheless, researchers have identified other factors that are involved in the formation or perseverance of substance use, some of which are evaluated here as potential missing links.

6.3.1 Emotion Regulation

The process by which individuals modulate their emotions in order to appropriately respond to life's demands is known as emotion regulation (Aldao, Nolen-Hoeksema, & Schweizer, 2010). Poor emotion regulation has been linked with a range of disorders, including pathological gambling (Williams, Grisham, Erskine, & Cassedy, 2012), eating disorders (Davis, Strachan, & Berkson, 2004) and substance abuse (Fox, Axelrod, Paliwal, Sleeper, & Sinha, 2007). The most important regulatory strategy in determining psychopathologies, including substance abuse, is rumination (Aldao et al., 2010). Ruminating on a problem is where an individual repetitively focuses on their experience of the emotion and its causes and consequences.

Tice, Bratslavsky, and Baumeister (2001) linked emotion regulation with impulse control and hypothesised that impulsive tendencies were indulged when an individual was feeling in a short-term bad mood, which they believed could be solved by a short-term hedonic solution. On the other hand, individuals who believed their bad mood was unchangable were less likely to indulge their impulses. In decision-making terms, this indicates a remarkably optimal self-awareness; an individual evaluates the expected utility of an addictive behaviour, taking into account their current mood and whether they think it can be changed. But evaluating the expected utility of options is central to conventional utility theory, so it is unclear how far emotion regulation overlaps

Emotion regulation is unique because it takes into account that some individuals may be more likely to experience (and notice) short-term negative moods than others, and thus would be more likely to benefit from short-term hedonic boosts. Additionally, it emphasises how an individual deals cognitively with the negative emotion.

6.3.2 Executive Function and the Iowa Gambling Task

Executive function broadly describes a cortical system which controls other functions. In order for the brain to estimate economic utility, it must predict and then signal the value of future events, taking into account the current state for comparison. Neuroscientists have located this system within the pre-frontal cortex (Schoenbaum, Roesch, & Stalnaker, 2006). A problem in this system, which predicts the expected value of decisions, could theoretically underlie the maladaptive choices that addicts make, and there is consistent evidence that addicts have abnormalities in blood flow in the orbito-frontal cortex (Dom, Sabbe, Hulstijn, & Van Den Brink, 2005).

Although it is interesting to find neurological substrates which underlie behaviour, the evidence does not help to explain the cognitions that occur or how an individual could be helped to overcome them. For that, the Iowa Gambling Task (IGT; Bechara et al., 1994) measures executive function. In the IGT, individuals are given a starting amount of money and are told to maximize profit by choose among four decks of cards. Two of the decks give consistently large profits, but these are more than offset by rare but massive losses, such that these decks ultimately incur a loss. The other two decks give consistently small profits, which are slightly offset by small losses, such that these decks are the most advantageous. The IGT has recently become available as a clinical instrument (Bechara, 2007).

The IGT has been used as a behavioural measure of risky decision making, and in this capacity significant impairments on the IGT have been found to be related to dependency on various substances, including alcohol (Bechara et al., 2001; Fein, Klein, & Finn, 2004; Goudriaan, Grekin, & Sher, 2011), cocaine (Bartzokis et al., 2000; Bolla et al., 2003), opioids (Mintzer & Stitzer, 2002) and marijuana (Bolla, Eldreth, Matochik, & Cadet, 2005). Similarly, pathological gamblers do not learn to switch to the advantageous decks (Glicksohn & Zilberman, 2010; Goudriaan, Oosterlaan, de Beurs, & Van den Brink, 2004). Finally, individuals with damage to the ventromedial prefrontal cortex (Bechara et al., 2001), as well as those with more diffuse frontal lobe lesions (Manes et al., 2002), choose poorly on the IGT.

The limitations of the IGT in understanding addiction are that it is unclear whether premorbid personality affects both addiction and performance on the IGT, and the IGT is a complex behavioural measure that correlates with brain structure but does not measure it (Buelow & Suhr, 2009).

6.3.3 Attentional Bias

Individuals who abuse alcohol, tobacco, or illicit drugs have an attentional bias towards cues that reflect their addiction (Field & Cox, 2008). The bias towards drug cues even predicts future drinking behaviour in excessive drinkers (Cox, Pothos, & Hosier, 2007). Field et al. (2011) found that heavy drinkers looked for longer at alcohol pictures than control pictures, regardless of whether the pictures were likely to lead to a reinforcement, whereas light drinkers only showed an attentional bias when the alcohol pictures were paired with reinforcement.

Attentional bias would increase the automaticity of drug-seeking behaviour in addicts, supporting the impulsivity view of addiction. On the other hand, it is presumably driven by a (perhaps pre-conscious) valuation of expected utility from that reward, which attests to expected utility theory. This presumes that expected utility is computed and so affects attentional bias, but another hypothesis could be that the amount of attention given to a cue determines its expected utility.

6.4 Global addictions: the Global Warming Game

This thesis has predominantly limited itself to discussion of individual decision-making, but intrapersonal externalities can theoretically apply to group decision-making. Essentially, the Harvard Game is a series of repeated Prisoner's Dilemmas (Rapoport & Chammah, 1965) but played with one's future self rather than another player (Rachlin, 2000). Choosing to defect (meliorate) means that one's future self only has choices between two suboptimal outcomes. Choosing to co-operate (self-invest) means that one's future self has the best possible options available. Unlike some versions of the Prisoner's Dilemma, one's future self has no ability to affect one's current choice.

The Tragedy of the Commons (Hardin, 1968) generalised the Prisoner's Dilemma to social dilemmas. Each member of a group has the option of investing in the group or keeping their income. At the end of each round, the group's investment is multiplied and distributed evenly to all members of the group. Investing in the group is essentially the same as co-operating, with the risk that other members of the group might decide to keep their money and so free-ride.

The Tragedy of the Commons can further be generalised to take account of the fact that consequences take some time before they become apparent. I propose a new game called the Global Warming Game. At the end of each round, participants can choose to pollute (defect) or recycle (co-operate). Polluting leads to a higher immediate payoff than recycling, but unlike the Tragedy of the Commons, negative effects of polluting, or positive effects of recycling, do not occur immediately but accrue over time. The challenge is for the group to learn to maximise their income, given the individual incentive to be selfish and the group delay in experiencing the full consequences of each decision.

Several research questions immediately present themselves from such a game. First, how can participants be helped to overcome the challenge of the task? This is a topical question, given concerns over climate change, and it is also relevant in other areas of intergenerational conflict, such as financing pensions (each generation's selfish incentive is to pay as little as possible yet take as much as possible from future generations). It is common to call the world's use of oil an "addiction" even in academic papers (Mason, Prior, Mudd, & Giurco, 2011; Trevors, 2011; Vilenchik, Peled, & Andelman, 2010). Mason et al. (2011) defines society's addiction to a resource as "its centrality and criticality to economic, social and environmental systems" (p.958), but it might be, however, that society is also more literally addicted to the resource, in that it has trouble taking into account the embedded future consequences of a decision, in the manner of an individual's addiction.

In experimental economics, there has been much fanfare over altruistic punishment as a method of helping groups to maximise their income in Tragedy of the Commons experiments (Fehr & Gachter, 2002). This is where participants have the option of sacrificing some of their own income in order to punish free-riders. It is an open question whether altruistic punishment would also overcome the Global Warming Game. Another method of limiting the socially negative consequences of a Tragedy of the Commons dilemma is to allow participants to communicate (Cooper, Dejong, Forsythe, & Ross, 1992). It would be interesting to see in the Global Warming Game whether communication overcomes the dilemma (it clearly does not in the real world), and also what participants discuss in the game, such as whether they articulate selfish motives towards the future recipients of their payoffs.

A second research area using the Global Warming Game would be in situations where a group passes the long-term consequences of their decisions to another group – the next generation. In game terms, one group would start the Global Warming Game and receive their payoffs based on their performance, and then another group would take over the game from the situation that the first group left behind. This presents an extremely strong challenge to altruism, because both the individual and the group have an incentive to be selfish.

6.5 A new way to do cognitive research?

It is a common in undergraduate psychology reports that a student will lament that the sample size was too small and that if only it was multiplied then statistically significant effects would be found. I will do the same, but hopefully with more persuasion.

In terms of sample size, typically individual differences or personality research has a few hundred participants, whereas cognitive research has around 30 per condition. Bertamini and Munafo (2012) checked studies published in Cognition in 2007 and found a median number of participants (*N*) of 80 for regular articles and 52 for brief articles. Partly this is due to effect sizes: cognitive research expects big effect sizes (due to well controlled studies) whereas individual difference research has to use lots of people in order to find effects. But addiction is a fusion of individual differences and cognition; many people try addictive substances but few become addicted. It therefore presents unique challenges. It can be frustrating to read studies where the N was small, such as the two validations of the Balloon Analog Risk Task mentioned in the Introduction, which had 60 and 64 participants (Dean et al., 2011; Lejuez et al., 2003) and found conflicting results.

More generally in Psychology, there is growing concern that samples are homogeneous and tend to be taken from among university students. Henrich et al. (2010) called the typical participant WEIRD (Western, educated, industrialized, rich, democratic), but more importantly for cognitive psychology they reported studies showing that such participants are frequently outliers compared to the rest of the species. This included studies of cooperation and fairness from behavioural economics, and so closely relates to decision-making.

More generally across all the sciences, there has been concern that apparently large effect sizes shrink over time and studies cannot be replicated. Begley and Ellis (2012) attempted to replicate 53 landmark studies from the field of preclinical medical research, but were only successful for 11% of them. Ioannidis even goes as far as to say that "simulations show that for most study designs and settings, it is more likely for a research claim to be false than true" (p.0696; Ioannidis, 2005). Bertamini and Munafo (2012) identified the phenomenon of "bite-size science" as one that has undesired side effects such as decreasing the number of participants.

The challenge is to increase the *N* without increasing research costs to the same power. The Internet is a solution to this. Once a study is online it can be participated in by 10 million people almost as easily as 50 people, as proved by myPersonality. Motivated online participants ("citizen scientists"; Wiederhold, 2011) who want to take part also provide good quality data as well as increasing the inclusiveness of the scientific method.

There is no reason that Internet studies should be limited to those measuring individual differences - who are just the researchers who got there first because they need large *N*s. Today or Tomorrow, the online version of a delay discounting task on myPersonality, showed that cognitive tasks will also be answered by thousands of people, providing they are created to provide interesting feedback. Harvard's Project Implicit (www.projectimplicit.net) has already collected data online using Implicit Association Tests (Greenwald, McGhee, & Schwartz, 1998) from thousands of individuals on the thoughts and feelings that are, according to the authors, outside of their conscious control.

Facebook also, with users' consent, can provide a key to bring studies together by providing a unique identity for each user. Whereas many people may use a single computer, only one person tends to use a Facebook account. Additionally, Facebook is already where users' online social interactions take place, so it comes with access to a wealth of online behavioural data. For example: knowing a user's Facebook friends could allow a researcher to measure whether impulsive individuals tend to be friends with other impulsive individuals, knowing what people write about in their status updates can allow a researcher to measure what kinds of things impulsive people talk about and how they say it, and knowing that an individual who retakes a study 12 months later is the same individual would allow longitudinal analyses which makes the data exponentially more valuable over time. There are legitimate privacy concerns, and individuals should have control over their data. But the potential benefits call for greater attention by the academic community.

6.6 Conclusions

Intrapersonal externalities provide an additional theory and method to understand addictions. Whereas rational economic actors in the behavioural economics literature discount future rewards, and irrational individuals in the psychopharmacological literature cannot suppress their impulses, decision-makers subject to intrapersonal externalities struggle to learn the full consequences of their decisions, and so choose suboptimal outcomes even while they attempt to maximise their subjective utility. As an extension of the economic decision-making literature, the theory also has the advantage that it can be described mathematically.

This thesis established the necessity to build upon the two currently dominant theories of impulsive behaviour, in that current measures do not account for a large proportion of variance in addictive behaviour, and in that intrapersonal externalities do not highly correlate with either delay discounting or trait impulsivity, but do correlate with smoking behaviour. It then shows that using the theory of intrapersonal externalities has advantages in both the laboratory and potentially in practice. In the laboratory, it allows a researcher to explore areas that are normally incompatible with the research area of choice over time, such as advice seeking and taking. In practice, it suggests a new way to change behaviour, by providing individual advice and feedback on choices.

The theory of intrapersonal externalities is still relatively unexplored; there are potential areas of unexplored timely research such as extending them to group situations. They still need to be validated in a wider set of self-investment circumstances (including in positive self-investment such as in exercise), and the limits of their use in the laboratory need to be quantified. Nevertheless, the theory has demonstrated potential, and so researchers should pay attention.

Appendices

APPENDIX A

Today or tomorrow: Instructions to participants

'Today or Tomorrow' information

1. Important information about the questionnaire

"Today or Tomorrow?" measures what economists call your delay discounting function. This describes how much you prefer the present over the future, or how willing you are to wait for a reward. You will be asked to answer 7 blocks of 15 questions asking whether you would prefer money now or in the future. Once you have done so, we will calculate your delay discounting function which estimates how much any reward drops in value as the delay to receive it increases. We will also compare this to others' delay discounting functions and will tell you whether you value the present or the future more than most others do.

The survey will take about 12 minutes to complete. Allow yourself enough time so that you do not have to rush. If you are interrupted you can return to the questionnaire and the questions that you have already submitted will be saved. To receive the most accurate results, you should complete the questionnaire in a quiet environment. You should specifically avoid taking the questionnaire while others are watching your responses. All answers you provide will be treated as confidential. We will only do research with your answers after the data has been anonymised so that it cannot be linked back to you.

Sadly we cannot actually give you the rewards that you prefer, but we would still ask you to answer as honestly as you can as if the rewards were for real. This will also give you the most accurate results.

The questionnaire has been made available as part of research by David Stillwell and Dr. Richard Tunney from the School of Psychology at the University of Nottingham. We are also grateful to Dr. Lee Hogarth for his invaluable assistance.

I have read and understood all of the above, and I will follow its recommendations. I know that I can withdraw at any time by closing my browser window and not answering further questions.

APPENDIX B

Individual differences in discounting: Controlling for demographics

The following section repeats the analyses from Experiment 2, but controls for age and gender.

B.1 Results

The daily, non-daily, and non-smoker groups differed by demographic characteristics, particularly age (Table B.1); daily smokers were older than non-daily smokers, who were older than non-smokers. This would be unlikely to cause spurious delay reward discounting (DRD) effects, because previous results have found that age negatively correlates with discounting rate (Green, Myerson, & Ostaszewski, 1999), and so this group age difference would likely decrease the relationship between smoking and DRD. In order to control for the effects of age and gender, the major analyses of the three methodological DRD differences were repeated with age and gender as covariates.

Table B.1: Demographic characteristics of smoking groups

Smoking Group	% Male	Mean Age (SD)
Daily Smokers	36%	28.3 (9.5)
Non-Daily Smokers	41%	23.6 (8.2)
Non-Smokers	38%	21.8 (8.6)

Using the 5953 respondents who reported their smoking behaviour, gender and age, a 3 x 3 between groups ANOVA was conducted to test whether the order that the items are presented in is related to delay discounting. Condition (ascending, descending, and randomized) and smoking status (daily smoker, N = 1073; non-daily smoker, N = 453; non-smoker, N = 4427) were the factors, age and gender were covariates, and log(k) was the dependent variable. A main effect of smoking status was found ($F(2,5942) = 58.72$, $MSE = .32$, $p <.001$). A main effect of condition was found ($F(2, 5942) = 9.27$, $MSE = .32$, $p < .001$). However, no interaction was found between smoking status and condition ($F(4, 5942) = .67$, $MSE = .32$, $p = .61$). Both gender ($F(1, 5942) = 6.19$, $MSE = .32$, $p = .013$) and age ($F(1, 5942) = 36.27$, $MSE = .32$, $p < .001$) were significant predictors.

To find out whether there were differences in the estimated DRD parameters at different delay lengths, and whether any differences were different for the three smoking statuses, a 3x3x6 mixed ANOVA was conducted with condition (ascending, descending, and randomized) and smoking status (daily smoker, non-daily smoker, non-smoker) as the between groups factors, delay length (1 week, 2 weeks, 1 month, 6 months, 1 year, 5 years) as the within groups factor, and the covariates of age and gender. There was a strong effect of delay length (F(5, 29710) = 179.48, MSE = .20, p < .001), such that shorter delays led to a steeper discounting parameter being estimated. The three-way interaction was not significant (F(20, 29710) = .76, MSE = .20, p = .77), but there was an interaction between delay length and smoking status (F(10, 29710) = 2.80, MSE = .20, p = .002).

In order to investigate the effects of differing delayed rewards, the DRD parameters for $1000 in 1 month were compared to the parameters for $100 in 1 month. A 3 x 2 mixed ANOVA with smoking status (daily smoker, non-daily smoker, and non-smoker) as the between-groups factor, delayed amount ($100, $1000) as the within-subjects factor, and age and gender as covariates found a main effect of delayed amount (F(1, 5847) = 61.11, MSE = .17, p < .001), but no interaction (F(2, 5847) = 1.28, MSE = .17, p = .28).

Overall, when controlling for age and gender the pattern of results replicated the findings in the main analysis.

APPENDIX C

Individuals' insight into intrapersonal externalities: Instructions to participants

Thank-you for agreeing to take part in this experiment.

Your task is simple. You will have to repeatedly choose between two buttons, marked # and . Simply click on a button with the mouse to register your choice.

As a result of your choices you will win Points. After every choice you will be shown your Points from each choice as well as your cumulative Points. As you gain more Points, Pacman will eat more dots and get larger!

However, choices will also use up Game Units. After every choice you will be shown the Game Units used up from each choice as well as your Game Units remaining. Once these have run out then *the game is over*.

You will play the game 8 times. Try and beat your previous score in every game!

At the end of every game, you will be asked a series of questions relating to potential scenarios within the game. Your answers to these questions will have no effect on the games that you play.

Your payment from this experiment will be £4.

[Participants in condition contingent were shown the following text in replace of the previous paragraph]

Your payment from this experiment will be based on the number of points that you gain during the games. This will be calculated on the basis of 0.08p/point. After each game, you will be shown your current earnings so far.

That's all there is to it – just try to win as many Points from the computer as you can before you run out of Game Units. Take as much time as you wish and please do not write anything down during the experiment.

APPENDIX D

Relationship between impulsivity and performance in Factory Manager Games

The following analyses show the relationship between trait impulsivity, as measured by the Barratt Impulsivity Scale, and performance on the Factory Manager Games. All analyses use linear regressions, and include Experiment number as a dummy variable. Table D.1 shows the overall BIS-11 score relationship, Table D.2 shows the relationship for the BIS-11 first order factors, and Table D.3 shows the relationship for BIS-11 second order factors. There are no signicant relationships for any of the impulsivity scores.

Table D.1: Regression coefficients showing the relationship between demographics, smoker type, and impulsivity, with performance in the Factory Manager Game.

Measure	Beta	t	p
Gender	-0.09	-1.40	0.16
Age	-0.07	-0.95	0.34
Smoker Type	-0.01	-0.14	0.89
Impulsivity (BIS-11)	0.05	0.73	0.47

Table D.2: Regression coefficients showing the relationship between demographics, smoker type, and BIS-11 first order factors, with performance in the Factory Manager Game.

Measure	Beta	t	p
Gender	-0.09	-1.38	0.17
Age	-0.06	-0.90	0.37
Smoker Type	-0.01	-0.07	0.94
BIS-11 Attention	0.01	0.13	0.89
BIS-11 Motor	0.06	0.83	0.41
BIS-11 Self-Control	0.09	1.04	0.30
BIS-11 Cognitive Complexity	-0.07	-0.87	0.39
BIS-11 Perseverence	-0.04	-0.53	0.60
BIS-11 Cognitive Instability	-0.04	-0.58	0.56

Table D.3: Regression coefficients showing the relationship between demographics, smoker type, and BIS-11 second order factors, with performance in the Factory Manager Game.

Measure	Beta	t	p
Gender	-0.09	-1.38	0.17
Age	-0.07	-0.97	0.33
Smoker Type	-0.01	-0.13	0.90
BIS-11 Attentional Impulsiveness	-0.01	-0.16	0.88
BIS-11 Motor Impulsiveness	0.04	0.50	0.61
BIS-11 Nonplanning Impulsiveness	0.03	0.37	0.71

APPENDIX E

Relationship between impulsivity and performance in intrapersonal externality tasks: Excluding smoking status

The following analyses show the relationship between trait impulsivity, as measured by the Barratt Impulsivity Scale, and performance on the intrapersonal externality tasks. All analyses use linear regressions, and include Experiment number as a dummy variable. Smoking group was excluded, in case it accounted for variance that would have been accounted for by trait impulsivity. Table E.1 shows the overall BIS-11 score relationship, Table E.2 shows the relationship for the BIS-11 first order factors, and Table E.3 shows the relationship for BIS-11 second order factors. There are no signicant relationships for any of the impulsivity scores.

Table E.1: Regression coefficients showing the relationship between demographics and impulsivity, with performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.07	-1.38	0.17
Age	-0.07	-1.18	0.24
Impulsivity (BIS-11)	0.01	0.24	0.81

Table E.2: Regression coefficients showing the relationship between demographics and BIS-11 first order factors, with performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.08	-1.36	0.17
Age	-0.06	-1.12	0.26
BIS-11 Attention	0.02	0.24	0.81
BIS-11 Motor	0.03	0.54	0.59
BIS-11 Self-Control	0.07	1.10	0.27
BIS-11 Cognitive Complexity	-0.06	-1.00	0.32
BIS-11 Perseverence	-0.05	-0.93	0.35
BIS-11 Cognitive Instability	-0.04	-0.64	0.53

Table E.3: Regression coefficients showing the relationship between demographics and BIS-11 second order factors, with performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.08	-1.40	0.16
Age	-0.07	-1.21	0.23
BIS-11 Attentional Impulsiveness	-0.01	-0.14	0.88
BIS-11 Motor Impulsiveness	0.00	0.04	0.97
BIS-11 Nonplanning Impulsiveness	0.02	0.33	0.74

APPENDIX F

Relationship between impulsivity and performance in intrapersonal externality tasks: Using z-scores

The following analyses show the relationship between trait impulsivity, as measured by the Barratt Impulsivity Scale, and performance on the intrapersonal externality tasks. Performance is measured using z-scores, standardized within experiment. This means that there is no need to use Experiment number as a dummy variable. Table F.1 shows the overall BIS-11 score relationship, Table F.2 shows the relationship for the BIS-11 first order factors, and Table F.3 shows the relationship for BIS-11 second order factors. There are no signicant relationships for any of the impulsivity scores.

Appendix F: Relationship between impulsivity and performance in intrapersonal externality tasks: Using z-scores

Table F.1: Regression coefficients showing the relationship between demographics, smoker type, and impulsivity, with within-experiment standardized performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.07	-1.35	0.18
Age	-0.06	-1.18	0.24
Smoker Type	0.03	0.48	0.63
Impulsivity (BIS-11)	0.00	0.06	0.95

Table F.2: Regression coefficients showing the relationship between demographics, smoker type, and BIS-11 first order factors, with within-experiment standardized performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.08	-1.34	0.18
Age	-0.06	-1.14	0.26
Smoker Type	0.04	0.66	0.51
BIS-11 Attention	0.02	0.28	0.78
BIS-11 Motor	0.03	0.50	0.62
BIS-11 Self-Control	0.07	1.04	0.30
BIS-11 Cognitive Complexity	-0.07	-1.09	0.28
BIS-11 Perseverence	-0.06	-1.05	0.29
BIS-11 Cognitive Instability	-0.04	-0.61	0.54

Table F.3: Regression coefficients showing the relationship between demographics, smoker type, and BIS-11 second order factors, with within-experiment standardized performance in games with intrapersonal externalities.

Measure	Beta	t	p
Gender	-0.08	-1.36	0.18
Age	-0.07	-1.19	0.23
Smoker Type	0.03	0.48	0.63
BIS-11 Attentional Impulsiveness	-0.01	-0.10	0.92
BIS-11 Motor Impulsiveness	0.00	-0.07	0.95
BIS-11 Nonplanning Impulsiveness	0.01	0.20	0.84

References

Adelstein, J. S., Shehzad, Z., Mennes, M., DeYoung, C. G., Zuo, X. N., Kelly, C., ... Milham, M. P. (2011). Personality is reflected in the brain's intrinsic functional architecture. *Plos One*, *6*(11), 12.

Ainslie, G. (1975). Specious reward - behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, *82*(4), 463-496.

Aldao, A., Nolen-Hoeksema, S., & Schweizer, S. (2010). Emotion-regulation strategies across psychopathology: A meta-analytic review. *Clinical Psychology Review*, *30*(2), 217-237.

Allais, M. (1953). Behavior of the rational man before risk - criticism of american school postulates and axioms. *Econometrica*, *21*(4), 503-546.

Anthony, J. C., Warner, L. A., & Kessler, R. C. (1994). Comparative epidemiology of dependence on tobacco, alcohol, controlled substances, and inhalants: Basic findings from the national comorbidity survey. *Experimental and Clinical Psychopharmacology*, *2*(3), 244-268.

APA. (2000). *Diagnostic and statistical manual of mental disorders* (4th ed., text rev. ed.). Washington, DC.

Audrain-McGovern, J., Rodriguez, D., Epstein, L. H., Cuevas, J., Rodgers, K., &

Wileyto, E. P. (2009). Does delay discounting play an etiological role in smoking or is it a consequence of smoking? *Drug and Alcohol Dependence*, *103*(3), 99-106.

Bagby, R. M., Parker, J. D. A., & Taylor, G. J. (1994). The 20-item toronto-alexithymia-scale .1. item selection and cross-validation of the factor structure. *Journal of Psychosomatic Research*, *38*(1), 23-32.

Baker, F., Johnson, M. W., & Bickel, W. K. (2003). Delay discounting in current and never-before cigarette smokers: Similarities and differences across commodity, sign, and magnitude. *Journal of Abnormal Psychology*, *112*(3), 382-392.

Barrick, M. R., Mount, M. K., & Judge, T. A. (2001). Personality and performance at the beginning of the new millennium: What do we know and where do we go next? *International Journal of Selection and Assessment*, *9*(1-2), 9-30.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*(3), 215-233.

Bartzokis, G., Lu, P. H., Beckson, M., Rapoport, R., Grant, S., Wiseman, E. J., & London, E. D. (2000). Abstinence from cocaine reduces high-risk responses on a gambling task. *Neuropsychopharmacology*, *22*(1).

Bechara, A. (2007). *Iowa gambling task professional manual*. Lutz, FL: Psychological Assessment Resources.

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, *50*(1-3), 7-15.

Bechara, A., Dolan, S., Denburg, N., Hindes, A., Anderson, S. W., & Nathan, P. E. (2001). Decision-malting deficits, linked to a dysfunctional ventromedial prefrontal cortex, revealed in alcohol and stimulant abusers. *Neuropsychologia*, *39*(4), 376-389.

Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy*, *96*(4), 675-700.

Beeler, J. D., & Hunton, J. E. (1997). The influence of compensation method and disclosure level on information search strategy and escalation of commitment. *Journal of Behavioral Decision Making*, *10*(2), 77-91.

Begley, C. G., & Ellis, L. M. (2012). Raise standards for preclinical cancer research. *Nature*, *483*(7391), 531-533.

Belin, D., Jonkman, S., Dickinson, A., Robbins, T. W., & Everitt, B. J. (2009). Parallel and interactive learning processes within the basal ganglia: Relevance for the understanding of addiction. *Behavioural Brain Research*, *199*(1), 89-102.

Belin, D., Mar, A. C., Dalley, J. W., Robbins, T. W., & Everitt, B. J. (2008). High impulsivity predicts the switch to compulsive cocaine-taking. *Science*, *320*(5881), 1352-1355.

Bertamini, M., & Munafo, M. R. (2012). Bite-size science and its undesired side effects. *Perspectives on Psychological Science*, *7*(1).

Bickel, W. K., Odum, A. L., & Madden, G. J. (1999). Impulsivity and cigarette smoking: delay discounting in current, never, and ex-smokers. *Psychopharmacology*, *146*(4), 447-454.

Birnbaum, M. H. (2004). Human research and data collection via the internet. *Annual Review of Psychology*, *55*, 803-832.

Bogg, T., & Roberts, B. W. (2004). Conscientiousness and health-related behaviors: A meta-analysis of the leading behavioral contributors to mortality. *Psychological Bulletin*, *130*(6), 887-919.

Bolla, K. I., Eldreth, D. A., London, E. D., Kiehl, K. A., Mouratidis, M., Contoreggi, C., … Ernst, M. (2003). Orbitofrontal cortex dysfunction in abstinent cocaine abusers performing a decision-making task. *Neuroimage*, *19*(3), 1085-1094.

Bolla, K. I., Eldreth, D. A., Matochik, J. A., & Cadet, J. L. (2005). Neural substrates of faulty decision-making in abstinent marijuana users. *Neuroimage*, *26*(2).

Bolton, G. E., & Ockenfels, A. (2000). Erc: A theory of equity, reciprocity, and competition. *American Economic Review*, *90*(1), 166-193.

Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, *101*(2), 127-151.

Bottlender, M., & Soyka, M. (2005). Impact of different personality dimensions (neo five-factor inventory) on the outcome of alcohol-dependent patients 6 and 12 months after treatment. *Psychiatry Research*, *136*(1), 61-67.

Bowman, C. H., & Turnbull, O. H. (2003). Real versus facsimile reinforcers on the iowa gambling task. *Brain and Cognition*, *53*(2), 207-210.

Brehmer, B., & Hagafors, R. (1986). Use of experts in complex decision-making - a paradigm for the study of staff work. *Organizational Behavior and Human Decision Processes*, *38*(2), 181-195.

Brunelle, C., Assaad, J. M., Barrett, S. P., Avila, C., Conrod, P. J., Tremblay, R. E., & Pihl, R. O. (2004). Heightened heart rate response to alcohol intoxication is associated with a reward-seeking personality profile. *Alcoholism-Clinical and Experimental Research*, *28*(3), 394-401.

Budescu, D. V., & Rantilla, A. K. (2000). Confidence in aggregation of expert opinions. *Acta Psychologica*, *104*(3), 371-398.

Buelow, M. T., & Suhr, J. A. (2009). Construct validity of the iowa gambling task. *Neuropsychology Review*, *19*(1), 102-114.

Camerer, C. F., & Ho, T. H. (1994). Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty*, *8*(2), 167-196.

Caseras, F. X., Fullana, M. A., Riba, J., Barbanoj, M. J., Aluja, A., & Torrubia, R. (2006). Influence of individual differences in the behavioral inhibition system and stimulus content (fear versus blood-disgust) on affective startle reflex modulation. *Biological Psychology*, *72*(3), 251-256.

Chang, L., & Krosnick, J. A. (2009). National surveys via rdd telephone interviewing versus the internet. *Public Opinion Quarterly*, *73*(4), 641-678.

Clark, A., Tran, C., Weiss, A., Caselli, G., Nikcevic, A. V., & Spada, M. M. (2012). Personality and alcohol metacognitions as predictors of weekly levels of alcohol use in binge drinking university students. *Addictive be-*

haviors, *37*(4), 537-40.

Clouston, S., Amin, S., Verdery, A., & Gauthier, G. R. (2009). The structure of undergraduate association networks: A quantitative ethnography. *Connections*, *29*(2), 18-31.

Conrod, P. J., Peterson, J. B., Pihl, R. O., & Mankowski, S. (1997). Biphasic effects of alcohol on heart rate are influenced by alcoholic family history and rate of alcohol ingestion. *Alcoholism-Clinical and Experimental Research*, *21*(1), 140-149.

Conrod, P. J., Pihl, R. O., & Vassileva, J. (1998). Differential sensitivity to alcohol reinforcement in groups of men at risk for distinct alcoholism subtypes. *Alcoholism-Clinical and Experimental Research*, *22*(3), 585-597.

Cooper, R., Dejong, D. V., Forsythe, R., & Ross, T. W. (1992). Communication in coordination games. *Quarterly Journal of Economics*, *107*(2).

Cox, W. M., Pothos, E. M., & Hosier, S. G. (2007). Cognitive-motivational predictors of excessive drinkers' success in changing. *Psychopharmacology*, *192*(4), 499-510. (Times Cited: 41)

Crean, J. P., de Wit, H., & Richards, J. B. (2000). Reward discounting as a measure of impulsive behavior in a psychiatric outpatient population. *Experimental and Clinical Psychopharmacology*, *8*(2).

Dalal, R. S., & Bonaccio, S. (2010). What types of advice do decision-makers prefer? *Organizational Behavior and Human Decision Processes*, *112*(1), 11-23.

Dallery, J., & Locey, M. L. (2005). Effects of acute and chronic nicotine on impulsive choice in rats. *Behavioural Pharmacology*, *16*(1), 15-23.

Dallery, J., & Raiff, B. R. (2007). Delay discounting predicts cigarette smoking in a laboratory model of abstinence reinforcement. *Psychopharmacology*, *190*(4), 485-496.

Daugherty, J. R., & Brase, G. L. (2010). Taking time to be healthy: Predicting health behaviors with delay discounting and time perspective. *Personality and Individual Differences*, *48*(2), 202-207.

Davidson, D., Hutchison, K., Dagon, C., & Swift, R. (2002). Assessing the stimulant effects of alcohol in humans. *Pharmacology Biochemistry and Behavior*, *72*(1-2), 151-156.

Davis, C., Strachan, S., & Berkson, M. (2004). Sensitivity to reward: implications for overeating and overweight. *Appetite*, *42*(2).

Dean, A. C., Sugar, C. A., Hellemann, G., & London, E. D. (2011). Is all risk bad? young adult cigarette smokers fail to take adaptive risk in a laboratory decision-making test. *Psychopharmacology*, *215*(4), 801-811.

de Kuilen, G. V., & Wakker, P. P. (2006). Learning in the allais paradox. *Journal of Risk and Uncertainty*, *33*(3), 155-164.

Delaney, H. D., & Maxwell, S. E. (1981). On using analysis of covariance in repeated measures designs. *Multivariate Behavioral Research*, *16*(1), 105-123.

Deroche-Gamonet, V., Belin, D., & Piazza, P. V. (2004). Evidence for addiction-like behavior in the rat. *Science*, *305*(5686), 1014-1017.

de Wit, H. (2009). Impulsivity as a determinant and consequence of drug use:

a review of underlying processes. *Addiction Biology*, *14*(1), 22-31.

de Wit, H., Flory, J. D., Acheson, A., McCloskey, M., & Manuck, S. B. (2007). Iq and nonplanning impulsivity are independently associated with delay discounting in middle-aged adults. *Personality and Individual Differences*, *42*(1), 111-121.

DeYoung, C. G., Hirsh, J. B., Shane, M. S., Papademetris, X., Rajeevan, N., & Gray, J. R. (2010). Testing predictions from personality neuroscience: Brain structure and the big five. *Psychological Science*, *21*(6), 820-828.

Dick, D. M., Smith, G., Olausson, P., Mitchell, S. H., Leeman, R. F., O'Malley, S. S., & Sher, K. (2010). Understanding the construct of impulsivity and its relationship to alcohol use disorders. *Addiction Biology*, *15*(2), 217-226.

Diergaarde, L., Pattij, T., Poortvliet, I., Hogenboom, F., de Vries, W., Schoffelmeer, A. N. M., & De Vries, T. J. (2008). Impulsive choice and impulsive action predict vulnerability to distinct stages of nicotine seeking in rats. *Biological Psychiatry*, *63*(3), 301-308.

Difranza, J. R., Winters, T. H., Goldberg, R. J., Cirillo, L., & Biliouris, T. (1986). The relationship of smoking to motor-vehicle accidents and traffic violations. *New York State Journal of Medicine*, *86*(9), 464-467.

Digman, J. M. (1990). Personality structure emergence of the five-factor model. In *Rosenzweig, m. r. and l. w. porter* (p. 417-440).

Dom, G., Sabbe, B., Hulstijn, W., & Van Den Brink, W. (2005). Substance use disorders and the orbitofrontal cortex - systematic review of behavioural decision-making and neuroimaging studies. *British Journal of Psychiatry*, *187*, 209-220.

Economidou, D., Pelloux, Y., Robbins, T. W., Dalley, J. W., & Everitt, B. J. (2009). High impulsivity predicts relapse to cocaine-seeking after punishment-induced abstinence. *Biological Psychiatry*, *65*(10), 851-856.

Edmonds, G. W., Bogg, T., & Roberts, B. W. (2009). Are personality and behavioral measures of impulse control convergent or distinct predictors of health behaviors? *Journal of Research in Personality*, *43*(5), 806-814.

Eiser, J. R., Sutton, S. R., & Wober, M. (1979). Smoking, seat-belts, and beliefs about health. *Addictive Behaviors*, *4*(4), 331-338.

Ellison, N. B., Steinfield, C., & Lampe, C. (2007). The benefits of facebook "friends": Social capital and college students' use of online social network sites. *Journal of Computer-Mediated Communication*, *12*(4), 28.

Enticott, P. G., Ogloff, J. R. P., & Bradshaw, J. L. (2006). Associations between laboratory measures of executive inhibitory control and self-reported impulsivity. *Personality and Individual Differences*, *41*(2), 285-294.

Ernst, M., Bolla, K., Mouratidis, M., Contoreggi, C., Matochik, J. A., Kurian, V., … London, E. D. (2002). Decision-making in a risk-taking task: A pet study. *Neuropsychopharmacology*, *26*(5), 682-691.

Ersche, K. D., Turton, A. J., Pradhan, S., Bullmore, E. T., & Robbins, T. W. (2010). Drug addiction endophenotypes: Impulsive versus sensation-seeking personality traits. *Biological Psychiatry*, *68*(8), 770-773.

Etter, J. F., Le Houezec, J., & Perneger, T. V. (2003). A self-administered questionnaire to measure dependence on cigarettes: The cigarette dependence scale. *Neuropsychopharmacology*, *28*(2), 359-370.

Everitt, B. J., Belin, D., Economidou, D., Pelloux, Y., Dalley, J. W., & Robbins, T. W. (2008). Neural mechanisms underlying the vulnerability to develop compulsive drug-seeking habits and addiction. *Philosophical Transactions of the Royal Society B-Biological Sciences*, *363*(1507), 3125-3135.

Everitt, B. J., Dickinson, A., & Robbins, T. W. (2001). The neuropsychological basis of addictive behaviour. *Brain Research Reviews*, *36*(2-3), 129-138.

Falk, A., Fehr, E., & Fischbacher, U. (2003). On the nature of fair behavior. *Economic Inquiry*, *41*(1), 20-26.

Falk, A., & Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, *54*(2), 293-315.

Fehr, E., & Gachter, S. (2002). Altruistic punishment in humans. *Nature*, *415*(6868), 137-140.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, *114*(3), 817-868.

Fein, G., Klein, L., & Finn, P. (2004). Impairment on a simulated gambling task in long-term abstinent alcoholics. *Alcoholism-Clinical and Experimental Research*, *28*(10).

Field, M., & Cox, W. M. (2008). Attentional bias in addictive behaviors: A review of its development, causes, and consequences. *Drug and Alcohol Dependence*, *97*(1-2).

Field, M., Hogarth, L., Bleasdale, D., Wright, P., Fernie, G., & Christiansen, P. (2011). Alcohol expectancy moderates attentional bias for alcohol cues in light drinkers. *Addiction*, *106*(6), 1097-1103.

Fillmore, M. T., & Rush, C. R. (2002). Impaired inhibitory control of behavior in chronic cocaine users. *Drug and Alcohol Dependence*, *66*(3), 265-273.

Fleming, M. F., Barry, K. L., Manwell, L. B., Johnson, K., & London, R. (1997). Brief physician advice for problem alcohol drinkers - a randomized controlled trial in community-based primary care practices. *Jama-Journal of the American Medical Association*, *277*(13), 1039-1045.

Flory, J. D., & Manuck, S. B. (2009). Impulsiveness and cigarette smoking. *Psychosomatic Medicine*, *71*(4), 431-437.

Fox, H. C., Axelrod, S. R., Paliwal, P., Sleeper, J., & Sinha, R. (2007). Difficulties in emotion regulation and impulse control during cocaine abstinence. *Drug and Alcohol Dependence*, *89*(2-3), 298-301.

Franken, I. H. A., van Strien, J. W., Franzek, E. J., & de Wetering, B. J. V. (2007). Error-processing deficits in patients with cocaine dependence. *Biological Psychology*, *75*(1), 45-51.

Friedman, H. J., Carpenter, J. A., Lester, D., & Randall, C. L. (1980). Effect of alpha-methyl-p-tyrosine on dose-dependent mouse strain differences in locomotor-activity after ethanol. *Journal of Studies on Alcohol*, *41*(1), 1-7.

Fuchs, R. A., Evans, K. A., Parker, M. P., & See, R. E. (2004). Differential involvement of orbitofrontal cortex subregions in conditioned cue-induced and cocaine-primed reinstatement of cocaine seeking in rats. *Journal of Neuroscience*, *24*(29), 6600-6610.

Galuska, D. A., Will, J. C., Serdula, M. K., & Ford, E. S. (1999). Are health care professionals advising obese patients to lose weight? *Jama-Journal of the American Medical Association*, *282*(16), 1576-1578.

Garavan, H., & Stout, J. C. (2005). Neurocognitive insights into substance abuse. *Trends in Cognitive Sciences*, *9*(4), 195-201.

Gardner, P. H., & Berry, D. C. (1995). The effect of different forms of advice on the control of a simulated complex system. *Applied Cognitive Psychology*, *9*, S55-S79.

Gellert, P., Ziegelmann, J. P., Lippke, S., & Schwarzer, R. (2012). Future time perspective and health behaviors: Temporal framing of self-regulatory processes in physical exercise and dietary behaviors. *Annals of Behavioral Medicine*, *43*(2), 208-218.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve bayesian reasoning without instruction - frequency formats. *Psychological Review*, *102*(4), 684-704.

Gilman, J. M., Ramchandani, V. A., Crouss, T., & Hommer, D. W. (2012). Subjective and neural responses to intravenous alcohol in young adults with light and heavy drinking patterns. *Neuropsychopharmacology*, *37*(2), 467-477.

Gino, F., & Moore, D. A. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making*, *20*(1), 21-35.

Gino, F., Shang, J., & Croson, R. (2009). The impact of information from similar or different advisors on judgment. *Organizational Behavior and Human Decision Processes*, *108*(2), 287-302.

Glicksohn, J., & Zilberman, N. (2010). Gambling on individual differences in decision making. *Personality and Individual Differences*, *48*(5), 557-562.

Goldberg, L. R. (1990). An alternative description of personality - the big-5 factor structure. *Journal of Personality and Social Psychology*, *59*(6), 1216-1229.

Goldberg, L. R., Johnson, J. A., Eber, H. W., Hogan, R., Ashton, M. C., Cloninger, C. R., & Gough, H. G. (2006). The international personality item pool and the future of public-domain personality measures. *Journal of Research in Personality*, *40*(1), 84-96.

Goldsmith, D. J., & Fitch, K. (1997). The normative context of advice as social support. *Human Communication Research*, *23*(4), 454-476.

Goldstein, A., & Kalant, H. (1990). Drug policy - striking the right balance. *Science*, *249*(4976), 1513-1521.

Goldstein, R. Z., Craig, A. D., Bechara, A., Garavan, H., Childress, A. R., Paulus, M. P., & Volkow, N. D. (2009). The neurocircuitry of impaired insight in drug addiction. *Trends in Cognitive Sciences*, *13*(9), 372-380.

Gomez-Beldarrain, M., Harries, C., Garcia-Monco, J. C., Ballus, E., & Grafman, J. (2004). Patients with right frontal lesions are unable to assess and use advice to make predictive judgments. *Journal of Cognitive Neuroscience*, *16*(1), 74-89.

Gosling, S. D., Sandy, C. J., John, O. P., & Potter, J. (2010). Wired but not weird: The promise of the internet in reaching more diverse samples. *Behavioral and Brain Sciences*, *33*(2-3), 94.

Gosling, S. D., Vazire, S., Srivastava, S., & John, O. P. (2004). Should we trust web-based studies? a comparative analysis of six preconceptions about internet questionnaires. *American Psychologist*, *59*(2), 93-104.

Goudriaan, A. E., Grekin, E. R., & Sher, K. J. (2011). Decision making and response inhibition as predictors of heavy alcohol use: A prospective study. *Alcoholism-Clinical and Experimental Research*, *35*(6), 1050-1057.

Goudriaan, A. E., Oosterlaan, J., de Beurs, E., & Van den Brink, W. (2004). Pathological gambling: a comprehensive review of biobehavioral findings. *Neuroscience and Biobehavioral Reviews*, *28*(2), 123-141.

Grace, R. C., & McLean, A. R. (2005). Integrated versus segregated accounting and the magnitude effect in temporal discounting. *Psychonomic Bulletin & Review*, *12*(4), 732-739.

Grant, B. F., & Dawson, D. A. (1998). Age of onset of drug use and its association with dsm-iv drug abuse and dependence: Results from the national longitudinal alcohol epidemiologic survey. *Journal of Substance Abuse*, *10*(2), 163-173.

Green, L., Myerson, J., & Ostaszewski, P. (1999). Discounting of delayed rewards across the life span: age differences in individual discounting functions. *Behavioural Processes*, *46*(1), 89-96.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, *74*(6).

Groen, Y., Wijers, A. A., Mulder, L. J. M., Waggeveld, B., Minderaa, R. B., & Althaus, M. (2008). Error and feedback processing in children with adhd and children with autistic spectrum disorder: An eeg event-related potential study. *Clinical Neurophysiology*, *119*(11), 2476-2493.

Gruber, J., & Koszegi, B. (2001). Is addiction "rational"? theory and evidence. *Quarterly Journal of Economics*, *116*(4), 1261-1303.

Guth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental-analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, *3*(4), 367-388.

Hair, P., & Hampson, S. E. (2006). The role of impulsivity in predicting maladaptive behaviour among female students. *Personality and Individual Differences*, *40*(5), 943-952.

Hardin, G. (1968). Tragedy of commons. *Science*, *162*(3859), 1243-&.

Harries, C., Yaniv, I., & Harvey, N. (2004). Combining advice: The weight of a dissenting opinion in the consensus. *Journal of Behavioral Decision Making*, *17*(5), 333-348.

Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment, and sharing responsibility. *Organizational Behavior and Human Decision Processes*, *70*(2), 117-133.

Harvey, N., & Harries, C. (2004). Effects of judges' forecasting on their later combination of forecasts for the same outcomes. *International Journal of Forecasting*, *20*(3), 391-409.

Heath, C., Huddart, S., & Lang, M. (1999). Psychological factors and stock option exercise. *Quarterly Journal of Economics*, *114*(2), 601-627.

Heil, S. H., Johnson, M. W., Higgins, S. T., & Bickel, W. K. (2006). Delay discounting in currently using and currently abstinent cocaine-dependent outpatients and non-drug-using matched controls. *Addictive Behaviors*, *31*(7), 1290-1294.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*(2-3), 61.

Herrnstein, R., & Vaughan, W. (1980). Melioration and behavioral allocation. In *Limits of action: The allocation of individual behavior* (p. 143-176). New York: Academic Press.

Herrnstein, R. J., Loewenstein, G. F., Prelec, D., & Vaughan, W. (1993). Utility maximization and melioration - internalities in individual choice. *Journal of Behavioral Decision Making*, *6*(3), 149-185.

Herrnstein, R. J., & Prelec, D. (1991). Melioration - a theory of distributed choice. *Journal of Economic Perspectives*, *5*(3), 137-156.

Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences*, *24*(3), 383-+.

Hester, R., Simoes-Franklin, C., & Garavan, H. (2007). Post-error behavior in active cocaine users: Poor awareness of errors in the presence of intact performance adjustments. *Neuropsychopharmacology*, *32*(9), 1974-1984.

Heyman, G. M., & Dunn, B. (2002). Decision biases and persistent illicit drug use: an experimental study of distributed choice and addiction. *Drug and Alcohol Dependence*, *67*(2), 193-203.

Heyman, G. M., & Gibb, S. P. (2006). Delay discounting in college cigarette chippers. *Behavioural Pharmacology*, *17*(8), 669-679.

Hogarth, L. (2011). The role of impulsivity in the aetiology of drug dependence: reward sensitivity versus automaticity. *Psychopharmacology*, *215*(3), 567-580.

Hogarth, L., Chase, H. W., & Baess, K. (2012). Impaired goal-directed behavioural control in human impulsivity. *Quarterly Journal of Experimental Psychology*, *65*(2), 305-316.

Holdstock, L., King, A. C., & de Wit, H. (2000). Subjective and objective responses to ethanol in moderate/heavy and light social drinkers. *Alcoholism-Clinical and Experimental Research*, *24*(6), 789-794.

Hong, R. Y., & Paunonen, S. V. (2009). Personality traits and health-risk behaviours in university students. *European Journal of Personality*, *23*(8), 675-696.

Hooten, W. M., Wolter, T. D., Ames, S. C., Hurt, R. D., Vickers, K. S., Offord, K. P., & Hays, J. T. (2005). Personality correlates related to tobacco abstinence following treatment. *International Journal of Psychiatry in Medicine*, *35*(1), 59-74.

Hopko, D. R., Lejuez, C. W., Daughters, S. B., Aklin, W. M., Osborne, A., Simmons, B. L., & Strong, D. R. (2006). Construct validity of the balloon analogue risk task (bart): Relationship with mdma use by inner-city drug users in residential treatment. *Journal of Psychopathology and Behavioral Assessment*, *28*(2), 95-101.

Hopwood, C. J., Morey, L. C., Skodol, A. E., Stout, R. L., Yen, S., Ansell, E. B., ... McGlashan, T. H. (2007). Five-factor model personality traits associated with alcohol-related diagnoses in a clinical sample. *Journal of Studies on Alcohol and Drugs*, *68*(3), 455-460.

Howell, D. (2007). *Statistical methods for psychology* (6th ed.). Thomson Wadsworth.

Ioannidis, J. P. A. (2005). Why most published research findings are false. *Plos Medicine*, *2*(8).

Jedynak, J. P., Uslaner, J. M., Esteban, J. A., & Robinson, T. E. (2007). Methamphetamine-induced structural plasticity in the dorsal striatum. *European Journal of Neuroscience*, *25*(3), 847-853.

Jentsch, J. D., & Taylor, J. R. (1999). Impulsivity resulting from frontostriatal dysfunction in drug abuse: implications for the control of behavior by reward-related stimuli. *Psychopharmacology*, *146*(4), 373-390.

Johnson, M. W., & Bickel, W. K. (2002). Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the Experimental Analysis of Behavior*, *77*(2), 129-146.

Johnson, M. W., Bickel, W. K., & Baker, F. (2007). Moderate drug use and delay discounting: A comparison of heavy, light, and never smokers. *Experimental and Clinical Psychopharmacology*, *15*(2), 187-194.

Johnson, M. W., Bickel, W. K., Baker, F., Moore, B. A., Badger, Q. J., & Budney, A. J. (2010). Delay discounting in current and former marijuana-dependent individuals. *Experimental and Clinical Psychopharmacology*, *18*(1), 99-107.

Kahneman, D., & Tversky, A. (1979). Prospect theory - analysis of decision under risk. *Econometrica*, *47*(2), 263-291.

Kawashima, K. (2006). The effects of inflation and interest rates on delay discounting in human behavior. *Psychological Record*, *56*(4), 551-568.

King, N., Carroll, C., Newton, P., & Dornan, T. (2002). "you can't cure it so you have to endure it": The experience of adaptation to diabetic renal disease. *Qualitative Health Research*, *12*(3), 329-346.

Kirby, K. N., & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science*, *6*(2), 83-89.

Kirby, K. N., & Petry, N. M. (2004). Heroin and cocaine abusers have higher discount rates for delayed rewards than alcoholics or non-drug-using controls. *Addiction*, *99*(4), 461-471.

Kirby, K. N., Petry, N. M., & Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology-General*, *128*(1), 78-87.

Kleinberg, J. (2008). The convergence of social and technological networks. *Communications of the Acm*, *51*(11), 66-72.

Koestner, R., Gingras, I., Abutaa, R., Losier, G. F., DiDio, L., & Gagne, M. (1999). To follow expert advice when making a decision: An examination of reactive versus reflective autonomy. *Journal of Personality*, *67*(5), 851-872.

Koff, E., & Lucas, M. (2011). Mood moderates the relationship between impulsiveness and delay discounting. *Personality and Individual Differences*, *50*(7), 1018-1022.

Koob, G. F., & LeMoal, M. (1997). Drug abuse: Hedonic homeostatic dysregulation. *Science*, *278*(5335), 52-58.

Koob, G. F., & Volkow, N. D. (2010). Neurocircuitry of addiction. *Neuropsychopharmacology*, *35*(1), 217-238.

Koritzky, G., & Yechiam, E. (2010). On the robustness of description and experience based decision tasks to social desirability. *Journal of Behavioral Decision Making*, *23*(1), 83-99.

Kotov, R., Gamez, W., Schmidt, F., & Watson, D. (2010). Linking "big" personality traits to anxiety, depressive, and substance use disorders: A meta-analysis. *Psychological Bulletin*, *136*(5), 768-821.

Krishnan-Sarin, S., Reynolds, B., Duhig, A. M., Smith, A., Liss, T., McFetridge, A., ... Potenza, M. N. (2007). Behavioral impulsivity predicts treatment outcome in a smoking cessation program for adolescent smokers. *Drug and Alcohol Dependence*, *88*(1), 79-82.

KudadjieGyamfi, E., & Rachlin, H. (1996). Temporal patterning in choice among delayed outcomes. *Organizational Behavior and Human Decision Processes*, *65*(1), 61-67.

Kurtz, J. E., & Parrish, C. L. (2001). Semantic response consistency and protocol validity in structured personality assessment: The case of the neo-pi-r. *Journal of Personality Assessment*, *76*(2), 315-332.

Lagorio, C. H., & Madden, G. J. (2005). Delay discounting of real and hypothetical rewards iii: Steady-state assessments, forced-choice trials, and all real rewards. *Behavioural Processes*, *69*(2), 173-187.

Lattimore, P. K., Baker, J. R., & Witte, A. D. (1992). The influence of probability on risky choice - a parametric examination. *Journal of Economic Behavior & Organization*, *17*(3), 377-400.

Laux, F. L. (2000). Addiction as a market failure: using rational addiction results to justify tobacco regulation. *Journal of Health Economics*, *19*(4), 421-437.

Lejuez, C. W., Aklin, W. M., Jones, H. A., Richards, J. B., Strong, D. R., Kahler, C. W., & Read, J. P. (2003). The balloon analogue risk task (bart) differentiates smokers and nonsmokers. *Experimental and Clinical Psychopharmacology*, *11*(1), 26-33.

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., ... Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: The balloon analogue risk task (bart). *Journal of Experimental Psychology-Applied*, *8*(2), 75-84.

Leri, F., & Stewart, J. (2001). Drug-induced reinstatement to heroin and cocaine seeking: A rodent model of relapse in polydrug use. *Experimental and Clinical Psychopharmacology*, *9*(3), 297-306.

Lewin, S. B. (1996). Economics and psychology: Lessons for our own day from the early twentieth century. *Journal of Economic Literature*, *34*(3), 1293-1323.

Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology*, *89*(1), 46-.

Lindman, H. R. (1971). Inconsistent preferences among gambles. *Journal of*

Experimental Psychology, *89*(2), 390-&.

Littlefield, A. K., & Sher, K. J. (2012). Smoking desistance and personality change in emerging and young adulthood. *Nicotine & Tobacco Research*, *14*(3), 338-342.

Locey, M. L., & Dallery, J. (2009). Isolating behavioral mechanisms of intertemporal choice: Nicotine effects on delay discounting and amount sensitivity. *Journal of the Experimental Analysis of Behavior*, *91*(2), 213-223.

Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice - evidence and an interpretation. *Quarterly Journal of Economics*, *107*(2), 573-597.

Loukas, A., Krull, J. L., Chassin, L., & Carle, A. C. (2000). The relation of personality to alcohol abuse/dependence in a high-risk sample. *Journal of Personality*, *68*(6), 1153-1175.

Mackillop, J., Amlung, M. T., Few, L. R., Ray, L. A., Sweet, L. H., & Munafo, M. R. (2011). Delayed reward discounting and addictive behavior: a meta-analysis. *Psychopharmacology*, *216*(3), 305-21.

MacKillop, J., & Kahler, C. W. (2009). Delayed reward discounting predicts treatment response for heavy drinkers receiving smoking cessation treatment. *Drug and Alcohol Dependence*, *104*(3), 197-203.

Madden, G. J., Bickel, W. K., & Jacobs, E. A. (1999). Discounting of delayed rewards in opioid-dependent outpatients: Exponential or hyperbolic discounting functions? *Experimental and Clinical Psychopharmacology*, *7*(3), 284-293.

Madden, G. J., Petry, N. M., Badger, G. J., & Bickel, W. K. (1997). Impulsive and self-control choices in opioid-dependent patients and non-drug-using control participants: Drug and monetary rewards. *Experimental and Clinical Psychopharmacology*, *5*(3), 256-262.

Maia, T. V., & McClelland, J. L. (2004). A reexamination of the evidence for the somatic marker hypothesis: What participants really know in the iowa gambling task. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(45), 16075-16080.

Malouff, J. M., Thorsteinsson, E. B., Rooke, S. E., & Schutte, N. S. (2007). Alcohol involvement and the five-factor model of personality: A meta-analysis. *Journal of Drug Education*, *37*(3), 277-294.

Malouff, J. M., Thorsteinsson, E. B., & Schutte, N. S. (2006). The five-factor model of personality and smoking: A meta-analysis. *Journal of Drug Education*, *36*(1), 47-58.

Manes, F., Sahakian, B., Clark, L., Rogers, R., Antoun, N., Aitken, M., & Robbins, T. (2002). Decision-making processes following damage to the prefrontal cortex. *Brain*, *125*.

Mansfield, P. M., Pinto, M. B., & Parente, D. H. (2003). Self-control and credit-card use among college students. *Psychological Reports*, *92*(3), 1067-1078.

Marlatt, G., Alan. (2002). Buddhist philosophy and the treatment of addictive behavior. *Cognitive and Behavioral Practice*, *9*(1), 44-50.

Marlow, C. (2009). *How diverse is facebook?* (Vol. 2012).

Mason, L., Prior, T., Mudd, G., & Giurco, D. (2011). Availability, addiction and alternatives: three criteria for assessing the impact of peak minerals on society. *Journal of Cleaner Production*, *19*(9-10), 958-966.

Matsuo, K., Nicoletti, M., Nemoto, K., Hatch, J. P., Peluso, M. A. M., Nery, F. G., & Soares, J. C. (2009). A voxel-based morphometry study of frontal gray matter correlates of impulsivity. *Human Brain Mapping*, *30*(4), 1188-1195.

Mayer, A., & Puller, S. L. (2008). The old boy (and girl) network: Social network formation on university campuses. *Journal of Public Economics*, *92*(1-2), 329-347.

Meda, S. A., Stevens, M. C., Potenza, M. N., Pittman, B., Gueorguieva, R., Andrews, M. M., … Pearlson, G. D. (2009). Investigating the behavioral and self-report constructs of impulsivity domains using principal component analysis. *Behavioural Pharmacology*, *20*(5-6), 390-399.

Mendez, I. A., Simon, N. W., Hart, N., Mitchell, M. R., Nation, J. R., Wellman, P. J., & Setlow, B. (2010). Self-administered cocaine causes long-lasting increases in impulsive choice in a delay discounting task. *Behavioral Neuroscience*, *124*(4), 470-477.

Mezquita, L., Stewart, S. H., & Ruiperez, M. A. (2010). Big-five personality domains predict internal drinking motives in young adults. *Personality and Individual Differences*, *49*(3), 240-245.

Milton, F., Longmore, C. A., & Wills, A. J. (2008). Processes of overall similarity sorting in free classification. *Journal of Experimental Psychology-Human Perception and Performance*, *34*(3), 676-692.

Mintzer, M. Z., & Stitzer, M. L. (2002). Cognitive impairment in methadone maintenance patients. *Drug and Alcohol Dependence*, *67*(1).

Mitchell, J. M., Fields, H. L., D'Esposito, M., & Boettiger, C. A. (2005). Impulsive responding in alcoholics. *Alcoholism-Clinical and Experimental Research*, *29*(12), 2158-2169.

Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology*, *146*(4), 455-464.

Mitchell, S. H. (2004). Measuring impulsivity and modeling its association with cigarette smoking. *Behav Cogn Neurosci Rev*, *3*(4), 261-75.

Mitchell, S. H., & Wilson, V. B. (2010). The subjective value of delayed and probabilistic outcomes: Outcome size matters for gains but not for losses. *Behavioural Processes*, *83*(1), 36-40.

Monterosso, J., & Ainslie, G. (2007). The behavioral economics of will in recovery from addiction. *Drug and Alcohol Dependence*, *90*, S100-S111.

Mueller, E. T., Landes, R. D., Kowal, B. P., Yi, R., Stitzer, M. L., Burnett, C. A., & Bickel, W. K. (2009). Delay of smoking gratification as a laboratory model of relapse: effects of incentives for not smoking, and relationship with measures of executive function. *Behavioural Pharmacology*, *20*(5-6), 461-473.

Munafo, M. R., & Black, S. (2007). Personality and smoking status: A longitudinal analysis. *Nicotine & Tobacco Research*, *9*(3), 397-404.

Munafo, M. R., Zetteler, J. I., & Clark, T. G. (2007). Personality and smoking status: A meta-analysis. *Nicotine & Tobacco Research*, *9*(3), 405-413.

Murphy, C., & MacKillop, J. (2012). Living in the here and now: interrelationships between impulsivity, mindfulness, and alcohol misuse. *Psychopharmacology*, *219*(2), 527-536.

Murphy, J. G., Vuchinich, R. E., & Simpson, C. A. (2001). Delayed reward and cost discounting. *Psychological Record*, *51*(4), 571-588.

Nees, F., Tzschoppe, J., Patrick, C. J., Vollstadt-Klein, S., Steiner, S., Poustka, L., ... Consortium, I. (2012). Determinants of early alcohol use in healthy adolescents: The differential contribution of neuroimaging and psychological factors. *Neuropsychopharmacology*, *37*(4), 986-995.

Nelson, A., & Killcross, S. (2006). Amphetamine exposure enhances habit formation. *Journal of Neuroscience*, *26*(14), 3805-3812.

Norvilitis, J. M., Szablicki, P. B., & Wilson, S. D. (2003). Factors influencing levels of credit-card debt in college students. *Journal of Applied Social Psychology*, *33*(5), 935-947.

Odum, A. L. (2011). Delay discounting: Trait variable? *Behavioural Processes*, *87*(1), 1-9.

Ozer, D. J., & Benet-Martinez, V. (2006). Personality and the prediction of consequential outcomes. *Annual Review of Psychology*, *57*, 401-421.

Patton, J. H., Stanford, M. S., & Barratt, E. S. (1995). Factor structure of the barratt impulsiveness scale. *Journal of Clinical Psychology*, *51*(6), 768-774.

Paunonen, S. V. (2003). Big five factors of personality and replicated predictions of nehavior. *Journal of Personality and Social Psychology*, *84*(2), 411-424.

Perry, J. L., & Carrol, M. E. (2008). The role of impulsive behavior in drug abuse. *Psychopharmacology*, *200*(1), 1-26.

Pettit, F. A. (2002). A comparison of world-wide web and paper-and-pencil personality questionnaires. *Behavior Research Methods Instruments & Computers*, *34*(1), 50-54.

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233-238.

Potenza, M. N. (2006). Should addictive disorders include non-substance-related conditions? *Addiction*, *101*, 142-151.

Potts, G. F., George, M. R. M., Martin, L. E., & Barratt, E. S. (2006). Reduced punishment sensitivity in neural systems of behavior monitoring in impulsive individuals. *Neuroscience Letters*, *397*(1-2), 130-134.

Prochaska, J. O., Crimi, P., Lapsanski, D., Martel, L., & Reid, P. (1982). Self-change processes, self-efficacy and self-concept in relapse and maintenance of cessation of smoking. *Psychological Reports*, *51*(3), 983-990.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, *68*(5), 1281-1292.

Rachlin, H. (2000). *The science of self-control*. Cambridge, Massachusetts: Harvard University Press.

Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective-probability and delay. *Journal of the Experimental Analysis of Behavior*, *55*(2), 233-244.

Rachlin, H., Siegel, E., & Cross, D. (1994). Lotteries and the time horizon. *Psychological Science*, *5*(6), 390-393.

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Be-*

havioral Decision Making, *23*(1), 1-14.

Rapoport, A., & Chammah, A. (1965). *Prisoners dilemma*. Ann Arbor, MI: Univ of Michigan Press.

Raynor, D. A., & Levine, H. (2009). Associations between the five-factor model of personality and health behaviors among college students. *Journal of American College Health*, *58*(1), 73-81.

Read, D. (2001). Intrapersonal dilemmas. *Human Relations*, *54*(8), 1093-1117.

Reimers, S., Maylor, E. A., Stewart, N., & Chater, N. (2009). Associations between a one-shot delay discounting measure and age, income, education and real-world impulsive behavior. *Personality and Individual Differences*, *47*(8), 973-978.

Rentfrow, P. J., Gosling, S. D., & Potter, J. (2008). A theory of the emergence, persistence, and expression of geographic variation in psychological characteristics. *Perspectives on Psychological Science*, *3*(5), 339-369.

Reynolds, B., Ortengren, A., Richards, J. B., & de Wit, H. (2006). Dimensions of impulsive behavior: Personality and behavioral measures. *Personality and Individual Differences*, *40*(2), 305-315.

Reynolds, B., Patak, M., Shroff, P., Penfold, R. B., Melanko, S., & Duhig, A. M. (2007). Laboratory and self-report assessments of impulsive behavior in adolescent daily smokers and nonsmokers. *Experimental and Clinical Psychopharmacology*, *15*(3), 264-271.

Reynolds, B., Richards, J. B., Horn, K., & Karraker, K. (2004). Delay discounting and probability discounting as related to cigarette smoking status in adults. *Behavioural Processes*, *65*(1), 35-42.

Richards, J. B., Zhang, L., Mitchell, S. H., & de Wit, H. (1999). Delay or probability discounting in a model of impulsive behavior: Effect of alcohol. *Journal of the Experimental Analysis of Behavior*, *71*(2), 121-143.

Robinson, T. E., & Berridge, K. C. (1993). The neural basis of drug craving - an incentive-sensitization theory of addiction. *Brain Research Reviews*, *18*(3), 247-291.

Robles, E., & Vargas, P. A. (2006). Functional parameters of delay discounting assessment tasks: Order of presentation. In *Meeting of the society-for-the-quantitative-analyses-of-behavior (sqab 2006)* (p. 237-241). Elsevier Science Bv.

Robles, E., & Vargas, P. A. (2008). Parameters of delay discounting assessment: Number of trials, effort, and sequential effects. *Behavioural Processes*, *78*(2), 285-290.

Robles, E., Vargas, P. A., & Bejarano, R. (2009). Within-subject differences in degree of delay discounting as a function of order of presentation of hypothetical cash rewards. *Behavioural Processes*, *81*(2), 260-263.

Rodzon, K., Berry, M. S., & Odum, A. L. (2011). Within-subject comparison of degree of delay discounting using titrating and fixed sequence procedures. *Behavioural Processes*, *86*(1), 164-167.

Rubinstein, A. (1988). Similarity and decision-making under risk (is there a utility-theory resolution to the allais paradox). *Journal of Economic Theory*, *46*(1), 145-153.

Russell, M. A. H., Wilson, C., Taylor, C., & Baker, C. D. (1979). Effect of general-practitioners advice against smoking. *British Medical Journal*, *2*(6184), 231-235.

Rust, J., & Golombok, S. (2008). *Modern psychometrics: The science of psychological assessment* (Third Edition ed.). Routledge.

Scalise, D. A., Berkel, L., & Van Whitlock, R. (2010). Client factors associated with treatment completion in a substance abuse treatment facility. *Addiction Research & Theory*, *18*(6), 667-680.

Schmidt, H. D., Anderson, S. M., Famous, K. R., Kumaresan, V., & Pierce, R. C. (2005). Anatomy and pharmacology of cocaine priming-induced reinstatement of drug seeking. *European Journal of Pharmacology*, *526*(1-3), 65-76.

Schoemaker, P. J. H. (1982). The expected utility model - its variants, purposes, evidence and limitations. *Journal of Economic Literature*, *20*(2), 529-563.

Schoenbaum, G., Roesch, M. R., & Stalnaker, T. A. (2006). Orbitofrontal cortex, decision-making and drug addiction. *Trends in Neurosciences*, *29*(2).

Schonberg, T., Fox, C. R., & Poldrack, R. A. (2011). Mind the gap: bridging economic and naturalistic risk-taking with cognitive neuroscience. *Trends in Cognitive Sciences*, *15*(1), 11-19.

Schrah, G. E., Dalal, R. S., & Sniezek, J. A. (2006). No decision-maker is an island: Integrating expert advice with information acquisition. *Journal of Behavioral Decision Making*, *19*(1), 43-60.

Schwartz, B. (1982). Reinforcement-induced behavioral stereotypy - how not to teach people to discover rules. *Journal of Experimental Psychology-General*, *111*(1), 23-59.

Setlow, B., Mendez, I. A., Mitchell, M. R., & Simon, N. W. (2009). Effects of chronic administration of drugs of abuse on impulsive choice (delay discounting) in animal models. *Behavioural Pharmacology*, *20*(5-6), 380-389.

Shadel, W. G., Cervone, D., Niaura, R., & Abrams, D. B. (2004). Investigating the big five personality factors and smoking: Implications for assessment. *Journal of Psychopathology and Behavioral Assessment*, *26*(3), 185-191.

Sher, K. J., & Trull, T. J. (1994). Personality and disinhibitory psychopathology - alcoholism and antisocial personality-disorder. *Journal of Abnormal Psychology*, *103*(1), 92-102.

Skinner, M. D., Aubin, H. J., & Berlin, I. (2004). Impulsivity in smoking, nonsmoking, and ex-smoking alcoholics. *Addictive Behaviors*, *29*(5), 973-978.

Skoe, E., & Kraus, N. (2012). A little goes a long way: how the adult brain is shaped by musical training in childhood. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *32*(34).

Slater, M. D. (2003). Sensation-seeking as a moderator of the effects of peer influences, consistency with personal aspirations, and perceived harm on marijuana and cigarette use among younger adolescents. *Substance Use & Misuse*, *38*(7), 865-880.

Sniezek, J. A., Schrah, G. E., & Dalal, R. S. (2004). Improving judgement with prepaid expert advice. *Journal of Behavioral Decision Making*, *17*(3), 173-190.

Sniezek, J. A., & Van Swol, L. M. (2001). Trust, confidence, and expertise in a judge-advisor system. *Organizational Behavior and Human Decision Processes*, *84*(2), 288-307.

Spillane, N. S., Smith, G. T., & Kahler, C. W. (2010). Impulsivity-like traits and smoking behavior in college students. *Addictive Behaviors*, *35*(7), 700-705.

Stanford, M. S., Mathias, C. W., Dougherty, D. M., Lake, S. L., Anderson, N. E., & Patton, J. H. (2009). Fifty years of the barratt impulsiveness scale: An update and review. *Personality and Individual Differences*, *47*(5), 385-395.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, *38*(2), 332-382.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1-26.

Stillwell, D. J., & Kosinski, M. (2011). *mypersonality research wiki.*

Stillwell, D. J., & Tunney, R. J. (2009). Melioration behaviour in the harvard game is reduced by simplifying decision outcomes. *Quarterly Journal of Experimental Psychology*, *62*(11), 2252-2261.

Stillwell, D. J., & Tunney, R. J. (2012). Effects of measurement methods on the relationship between smoking and delay reward discounting. *Addiction*, *107*(5), 1003-12.

Strathman, A., Gleicher, F., Boninger, D. S., & Edwards, C. S. (1994). The consideration of future consequences - weighing immediate and distant outcomes of behavior. *Journal of Personality and Social Psychology*, *66*(4), 742-752.

Sussman, S., Leventhal, A., Bluthenthal, R. N., Freimuth, M., Forster, M., & Ames, S. L. (2011). A framework for the specificity of addictions. *International Journal of Environmental Research and Public Health*, *8*(8), 3399-3415.

Swann, A. C., Bjork, J. M., Moeller, F. G., & Dougherty, D. M. (2002). Two models of impulsivity: Relationship to personality traits and psychopathology. *Biological Psychiatry*, *51*(12), 988-994.

Sweitzer, M. M., Donny, E. C., Dierker, L. C., Flory, J. D., & Manuck, S. B. (2008). Delay discounting and smoking: Association with the fagerstrom test for nicotine dependence but not cigarettes smoked per day. *Nicotine & Tobacco Research*, *10*(10), 1571-1575.

Tarter, R. E., Kirisci, L., Habeych, M., Reynolds, M., & Vanyukov, M. (2004). Neurobehavior disinhibition in childhood predisposes boys to substance use disorder by young adulthood: direct and mediated etiologic pathways. *Drug and Alcohol Dependence*, *73*(2), 121-132.

ter Bogt, T. F. M., Engels, R., & Dubas, J. S. (2006). Party people: Personality and mdma use of house party visitors. *Addictive Behaviors*, *31*(7), 1240-1244.

Tercyak, K. P., & Audrain-McGovern, J. (2003). Personality differences associated with smoking experimentation among adolescents with and without comorbid symptoms of adhd. *Substance Use & Misuse*, *38*(14), 1953-1970.

Terracciano, A., & Costa, P. T. (2004). Smoking and the five-factor model of personality. *Addiction*, *99*(4), 472-481.

Terracciano, A., Lockenhoff, C. E., Crum, R. M., Bienvenu, O. J., & Costa, P. T.

(2008). Five-factor model personality profiles of drug users. *Bmc Psychiatry*, *8*, 10.

Tice, D. M., Bratslavsky, E., & Baumeister, R. F. (2001). Emotional distress regulation takes precedence over impulse control: If you feel bad, do it! *Journal of Personality and Social Psychology*, *80*(1), 53-67.

Tiffany, S. T. (1990). A cognitive model of drug urges and drug-use behavior - role of automatic and nonautomatic processes. *Psychological Review*, *97*(2), 147-168.

Townshend, J. M., & Duka, T. (2005). Binge drinking, cognitive performance and mood in a population of young social drinkers. *Alcoholism-Clinical and Experimental Research*, *29*(3), 317-325.

Trevors, J. T. (2011). Human addictions: Oil, conflicts and super consumerism. *Water Air and Soil Pollution*, *217*(1-4), 1-2.

Tsukayama, E., & Duckworth, A. L. (2010). Domain-specific temporal discounting and temptation. *Judgment and Decision Making*, *5*(2), 72-82.

Tsutsui-Kimura, I., Ohmura, Y., Izumi, T., Yamaguchi, T., Yoshida, T., & Yoshioka, M. (2010). Endogenous acetylcholine modulates impulsive action via alpha 4 beta 2 nicotinic acetylcholine receptors in rats. *European Journal of Pharmacology*, *641*(2-3), 148-153.

Tunney, R. J. (2006). Preference reversals are diminished when gambles are presented as relative frequencies. *Quarterly Journal of Experimental Psychology*, *59*(9), 1516-1523.

Tunney, R. J., & Shanks, D. R. (2002). A re-examination of melioration and rational choice. *Journal of Behavioral Decision Making*, *15*(4), 291-311.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453-458.

Tversky, A., & Kahneman, D. (1992). Advances in prospect-theory - cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297-323.

Tversky, A., Slovic, P., & Kahneman, D. (1990). The causes of preference reversal. *American Economic Review*, *80*(1), 204-217.

Valentin, V. V., Dickinson, A., & O'Doherty, J. P. (2007). Determining the neural substrates of goal-directed learning in the human brain. *Journal of Neuroscience*, *27*(15), 4019-4026.

Valois, R. F., Oeltmann, J. E., Waller, J., & Hussey, J. R. (1999). Relationship between number of sexual intercourse partners and selected health risk behaviors among public high school adolescents. *Journal of Adolescent Health*, *25*(5), 328-335.

Van Swol, L. M., & Ludutsky, C. L. (2007). Tell me something i don't know - decision makers' preference for advisors with unshared information. *Communication Research*, *34*(3), 297-312.

Verdejo-Garcia, A. J., Perales, J. C., & Perez-Garcia, M. (2007). Cognitive impulsivity in cocaine and heroin polysubstance abusers. *Addictive Behaviors*, *32*(5), 950-966.

Vilenchik, Y., Peled, E., & Andelman, D. (2010). Kicking the oil addiction. *Physics World*, *23*(1), 22-26.

Vuchinich, R. E., & Simpson, C. A. (1998). Hyperbolic temporal discounting in social drinkers and problem drinkers. *Experimental and Clinical Psychopharmacology*, *6*(3), 292-305.

Wallace, P., Cutler, S., & Haines, A. (1988). Randomized controlled trial of general-practitioner intervention in patients with excessive alcohol-consumption. *British Medical Journal*, *297*(6649), 663-668.

Wallis, J. D. (2007). Orbitofrontal cortex and its contribution to decision-making. *Annual Review of Neuroscience*, *30*, 31-56.

Wallsten, T. S., & Budescu, D. V. (1995). A review of human linguistic probability processing - general-principles and empirical-evidence. *Knowledge Engineering Review*, *10*(1), 43-62.

Warry, C. J., Remington, B., & Sonuga-Barke, E. J. S. (1999). When more means less: Factors affecting human self-control in a local versus global choice paradigm. *Learning and Motivation*, *30*(1), 53-73.

White, J. L., Moffitt, T. E., Caspi, A., Bartusch, D. J., Needles, D. J., & Stouthamerloeber, M. (1994). Measuring impulsivity and examining its relationship to delinquency. *Journal of Abnormal Psychology*, *103*(2).

White, T. L., Lejuez, C. W., & Wit, H. (2008). Test-retest characteristics of the balloon analogue risk task (bart). *Experimental and Clinical Psychopharmacology*, *16*(6), 565-570.

Wiederhold, B. K. (2011). Citizen scientists generate benefits for researchers, educators, society, and themselves. *Cyberpsychology Behavior and Social Networking*, *14*(12), 703-704.

Williams, A. D., Grisham, J. R., Erskine, A., & Cassedy, E. (2012). Deficits in emotion regulation associated with pathological gambling. *British Journal of Clinical Psychology*, *51*, 223-238.

Willner, P. (2000). Further validation and development of a screening instrument for the assessment of substance misuse in adolescents. *Addiction*, *95*(11), 1691-1698.

Wilson, J. R., & Nagoshi, C. T. (1988). Adult children of alcoholics - cognitive and psychomotor characteristics. *British Journal of Addiction*, *83*(7), 809-820.

Wilson, R. E., Gosling, S. D., & Graham, L. T. (2012). A review of facebook research in the social sciences. *Perspectives on Psychological Science*.

Winstanley, C. A. (2007). The orbitofrontal cortex, impulsivity, and addiction - probing orbitofrontal dysfunction at the neural, neurochemical, and molecular level. *Linking Affect to Action: Critical Contributions of the Orbitofrontal Cortex*, *1121*, 639-655.

Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, *93*(1), 1-13.

Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, *83*(2), 260-281.

Yarkoni, T., Braver, T. S., Gray, J. R., & Green, L. (2005). Prefrontal brain activity predicts temporally extended decision-making behavior. *Journal of the Experimental Analysis of Behavior*, *84*(3), 537-554.

Yarkoni, T., Gray, J. R., Chrastil, E. R., Barch, D. M., Green, L., & Braver, T. S.

(2005). Sustained neural activity associated with cognitive control during temporally extended decision making. *Cognitive Brain Research*, *23*(1), 71-84.

Zimbardo, P. G., & Boyd, J. N. (1999). Putting time in perspective: A valid, reliable individual-differences metric. *Journal of Personality and Social Psychology*, *77*(6), 1271-1288.

Zuckerman, M., Ball, S., & Black, J. (1990). Influences of sensation seeking, gender, risk appraisal, and situational motivation on smoking. *Addictive Behaviors*, *15*(3), 209-220.

List of Figures

List of Tables