The role of unconscious influences on decision-making under uncertainty: Behavioural and computational approaches

Emmanouil Konstantinidis

Department of Experimental Psychology

University College London

Thesis submitted for the degree of
Doctor of Philosophy (PhD)
September 2014

Abstract

How do people make decisions in uncertain environments and what types of knowledge control their choices? Can our decisions be guided by unconscious influences or intuitive "gut" feelings? According to the Somatic Marker Hypothesis, a popular account of the role of affect in decision-making, emotion-based signals can guide our decisions in uncertain environments outside awareness. However, evidence for this claim can be questioned on the grounds of inadequate and insensitive assessments of conscious knowledge.

In this work, variations of a classic experience-based decision-making paradigm, the Iowa Gambling Task (IGT), are employed in combination with subjective measures of awareness in order to investigate the role played by unconscious influences. Specifically, the validity of post-decision wagering as a sensitive and bias-free measure of conscious content is examined and contrasted to confidence ratings and quantitative reports. The results demonstrate the inadequacy of post-decision wagering as a direct measure of conscious knowledge and also question the claim that implicit processes influence decision-making.

In order to measure and understand the cognitive and psychological processes underlying performance on the IGT, computational modeling analyses are undertaken to provide deeper insights into the dynamics of decision-making. Reinforcement-learning models are evaluated using different model comparison techniques and a computational model of confidence ratings in decision-making under uncertainty is developed.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text. This work has not been submitted for any other degree or professional qualification except as specified.

Signature: London, September 10, 2014

(Emmanouil Konstantinidis)

To my parents, Αλέξανδρος and Πένυ,
my sister, Εύη, and my newly born nephew, Φίλιππος.

Acknowledgements

It has been an absolute delight to have David R. Shanks as my supervisor and mentor throughout my doctoral work. I would like to thank David for his invaluable guidance, thoughtful insights, and the innumerable discussions about my research. This thesis would not have been possible without David's critical comments and suggestions.

I have been incredibly lucky to have worked with two great mathematical psychologists, Maarten Speekenbrink and Jerome R. Busemeyer, who introduced me to the world of computational modelling. I would like to thank Maarten for providing guidance and answers to all my statistics/mathematics questions throughout my PhD. Special thanks go to Jerome for hosting me in his lab at Indiana University Bloomington during my US visit in the summer of 2013.

I would also like to thank the former and current members of the ShanksLab, Miguel Vadillo, Chris Berry, Tom Hardwicke, Rosalind Potts, and Emma Ward for making working in our office an enjoyable experience and for discussions about philosophical, psychological, methodological, and mathematical matters. Special thanks go to Tom Hardwicke for offering to proofread my thesis (any spelling or grammatical mistakes should be attributed to him! :). I am also thankful to Helen Steingroever for many discussions and debates about reinforcement-learning models.

I am grateful to the UCL Graduate School for funding my research, and to the Experimental Psychology Society (EPS) for providing financial assistance for my US visit.

Finally, I would like to thank my parents for their love, encouragement, and support throughout my studies. Without them I would never have enjoyed so many opportunities in my life.

Published articles

Experiments 1, 3, 4, and 6 appear in Konstantinidis, E, & Shanks, D. R. (in press). Don't bet on it! Wagering as a measure of awareness in decision making under uncertainty. *Journal of Experimental Psychology: General.*

Experiments 4A and 5 appear in Konstantinidis, E., & Shanks, D. R. (2013). Awareness in decision-making: Implicit influences or measurement biases? In M. Knauff, M. Pauen, N. Sebanz, and I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 811-816). Austin, TX: Cognitive Science Society.

The section *To simulate or not?* in Chapter 5 is based on Konstantinidis, E., Speekenbrink, M., Stout, J. C., Ahn, W.-Y., & Shanks, D. R. (2014). To simulate or not? Comment on Steingroever, Wetzels, and Wagenmakers (2014). *Decision*, *1*, 184-191.

In Preparation:

The cognitive modelling analysis in Chapter 5 appear in Konstantinidis, E., Busemeyer, J. R., Speekenbrink, M., & Shanks, D. R. (2014). *Cognitive modeling of experience-based decision-making: The differential effects of past payoffs and past choices.* Manuscript in preparation.

Contents

List of Figures

List of Tables

1 Introduction

How do people make decisions and adjust their behaviour in uncertain environments and what types of knowledge control their choices? The study of human decision-making has relied on two broad categories of experimental tasks: *Description-based* tasks in which information about the possible courses of actions is available (probabilities and monetary outcomes of an action) and *experience-based* tasks in which probabilities and outcomes have to be learned by observing or exploring the different alternatives (e.g., Hertwig, Barron, Weber, & Erev, 2004; Koritzky & Yechiam, 2010).

In description-based tasks, participants make decisions by choosing between two (or more) verbal descriptions of available options. Each description/option consists of a monetary outcome (e.g., \$5) and a probability (e.g., $p = 0.4$) of this outcome to occur. A typical instantiation of this paradigm includes two options (Figure 1.1–d): a *safe* option which provides a medium magnitude outcome with certainty ($p = 1$), and a *risky* option which produces a high magnitude outcome with a probability p and a low magnitude outcome with probability $1 - p$ (Erev et al., 2010).

In decisions from experience, participants are presented with a number of options (usually two) and are asked to choose from any option they want in order to maximise their overall profit. This paradigm has usually been employed in three different variations: in the *partial-information paradigm* (Figure 1.1–a) participants only receive feedback on the selected option (obtained payoff) whereas in the *full-information paradigm* (Figure 1.1–b) they receive feedback on both the selected and the unselected options (foregone payoff). In both paradigms only the obtained payoffs contribute to the overall sum of points/payoffs. A third variation (the *sampling paradigm*) allows participants to sample from the available options as long as they wish and then make a single decision between the options (Hertwig & Erev, 2009).

The repeated nature of the experience-based paradigm allows a number of important questions to be examined, relating to how much people explore the environment before they exploit the most profitable options and what factors affect this trade-off, the role of cognitive processes such as memory and learning in updating the value of each option, and economic factors which shape decision strategies (see Rakow & Newell, 2010). Also, experience-based tasks permit a more insightful examination of individual differences in cognitive processes and the application of computational models in an attempt to decompose and explain these underlying cognitive and psychological processes (e.g., Busemeyer & Stout, 2002; Erev & Barron, 2005; Erev et al., 2010).

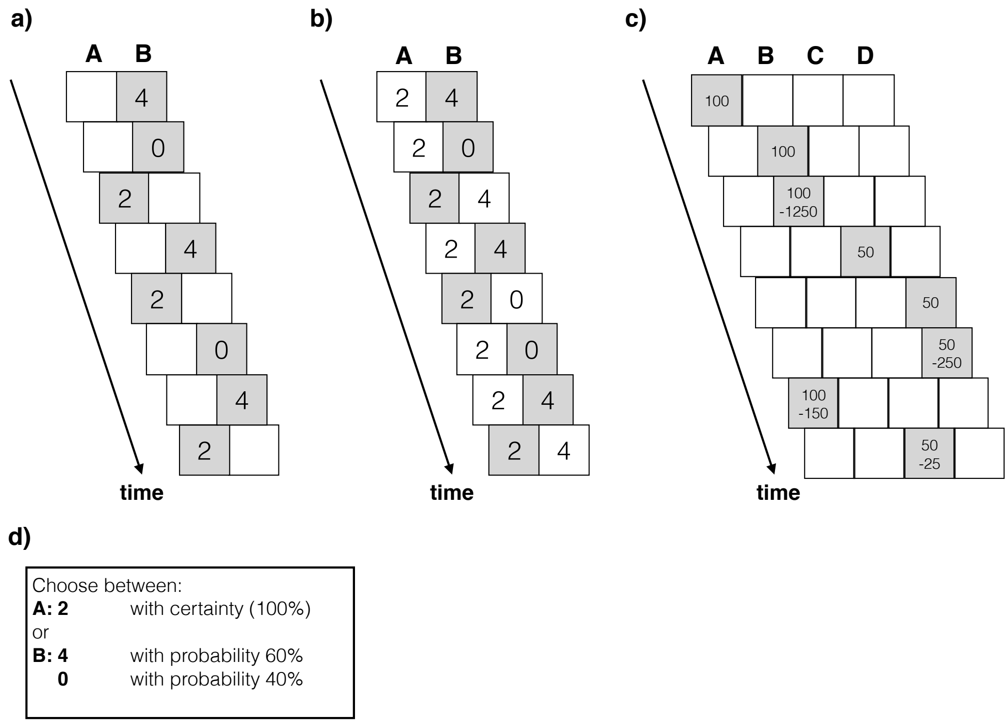

Figure 1.1: Schematic representation of decisions from experience (a,b, and c) and description (d). a) Partial-information, b) Full-information, c) The Iowa Gambling Task.

Traditionally, decision research has focused on how people make decisions based on description and experience and how they process and combine values and their associated probabilities. Different formal and descriptive theories of decision-making assume different integration of values and probabilities. For example, expected value (EV) theory dictates a linear combination of values and probabilities for each option whereas expected utility (EU) theory assumes linear probabilities but non-linear value functions for each option. Finally, prospect theory (Kahneman & Tversky, 1979) assumes non-linear functions and transformations for both values and probabilities associated with each option (see Lejarraga & Gonzalez, 2011).

In recent years several theories have emerged which deviate from normative models and formal theories of decision-making, focusing more on affective and emotional processes that support decision-making under uncertainty (e.g., Dolan, 2002; Finucane, Alhakami, Slovic, & Johnson, 2000; Loewenstein, Weber, Hsee, & Welch, 2001; Mellers & McGraw, 2001; N. Schwarz, 2000). This work has drawn attention to the impact of affect by pointing out that automatic and rapid emotional reactions may serve as input to the decision-making process.

1.1 The Somatic Marker Hypothesis

In the popular TV series “House of Cards”, the resourceful soon-to-be Vice President of the USA, Frank Underwood, is taking a walk in the woods near St. Louis, Missouri. His companion is Raymond Tusk, a friend of the President and a billionaire industrialist who specialises in nuclear power and owns quite a few nuclear power plants. Given the President’s hesitation to provide strong support towards a full-scale reliance on nuclear power, Raymond advocates that nuclear power is “the only option we have right now that does not completely trash the planet. The argument against nuclear power is an *emotional* one”. Frank blandly replies: “And you don’t make decisions based on emotion”. Raymond takes a deep breath and with a condescending and disdainful tone (as if he was talking to a little child) says to Frank: “Decisions based on emotion aren’t decisions at all. [They are] Instincts...Which can be of value. The rational and the irrational complement each other. Individually, they’re far less powerful.”

Raymond’s reply, that decisions based on emotion are of no practical value, reflects the tendency which has dominated decision science, that is a focus on mathematical functional models, based on principles from economics and statistics (Weber & Johnson, 2009). Decision research neglected the influence of *hot*, emotional processes and relied more on *cold*, reason-based explanations to study and investigate decision-making (Peters, Västfjäll, Gärling, & Slovic, 2006). Emotion and affect were considered as "obstacles" which bias people's reasoning and cognitive processes. Despite the fact that affect has played a major role in shaping a variety of social and behavioural theories, its role has rarely been recognised as an important component of judgment and decision making (Finucane et al., 2000). The *emotions revolution* of the past two decades has tried to correct the focus on analytic and computational processes by showing the prevalence of emotional processes as automatic and effort-free inputs that guide and motivate behaviour (Weber & Johnson, 2009).

One popular account of the role of affect in reasoning and decision-making is the Somatic Marker Hypothesis (SMH) proposed by Damasio (1994, 1996). Initially, the SMH was developed to explain deficits in patients with certain kinds of prefrontal brain damage (ventromedial prefrontal cortex, VMPFC) who exhibit severe decision-making impairments in social and personal domains while their cognitive and problem solving abilities remain largely unimpaired (Bechara, Damasio, Tranel, & Damasio, 2005; Saver & Damasio, 1991). Such patients also have difficulties in expressing emotional and affective information. The SMH proposes that these deficits are connected and that decision-making is regulated by neural biasing signals arising from emotion processing (for reviews see Bechara, 2004; Bechara & Damasio, 2005; Dunn, Dalgleish, & Lawrence, 2006). These signals can be marked as either positive or negative and are linked directly to bodily states. When a negative somatic marker is associated with a possible response action, it produces an avoidance reaction; on the other hand, a positive somatic marker indicates that the response option is beneficial. In situations of uncertainty, these somatic markers help guide behaviour by marking response alternatives with an emotional signal, thus providing information useful for guiding the decision process (Damasio, 1994; but see Davis, Love, & Maddox, 2009). Hence, the inability of VMPFC patients to integrate and process emotional information leads to disadvantageous decision-making which can be described as risky, prone to short-term rewards, and insensitive to loss and long term consequences (Bechara, Damasio, Damasio, & Anderson, 1994).

A major assumption regarding somatic markers is that they operate not only consciously, when someone has accessible knowledge about the possible outcomes of a choice, but also unconsciously (Bechara, Damasio, & Damasio, 2000; Bechara, Damasio, Tranel, & Damasio, 1997). Specifically, in situations of high uncertainty, somatic markers can guide individuals to make advantageous decisions or avoid disadvantageous ones even when they are not explicitly aware of the quality or value of those decisions.

The question whether behaviour and decision-making can be influenced by unconscious "gut feelings" and "intuitive processes" has attracted considerable attention within psychological science. The different formulations of this distinction (e.g., System 1 vs. System 2, implicit vs. explicit) have different functional attributes and procedural features but also share some common characteristics. For example, System 1 is unconscious, associative, effortless, and fast whereas System 2 is conscious, deliberative, and rule-governed (see Kahneman, 2011). An important feature of System 1 is its reliance on affective information or signals (Slovic, Finucane, Peters, & MacGregor, 2002). Hence somatic markers (or emotional/affective biasing signals) can be seen as manifestations of System 1, guiding people to make advantageous decisions in situations of uncertainty and outside of awareness. Thus the main assumptions of the SMH fit readily within this dichotomy of reasoning systems.

1.2 The Iowa Gambling Task

In order to test the SMH empirically, Damasio and colleagues developed a gambling task (the Iowa Gambling Task, IGT) which attempts to simulate real-life decision-making in a laboratory setting in the way it employs uncertainty, rewards and punishments (Bechara et al., 1994; Bechara, Damasio, Damasio, & Lee, 1999). The IGT is one of the most popular and frequently used paradigms in decision-making under uncertainty and has become a standard screening tool for assessing decision-making deficits in a variety of clinical populations (Buelow & Suhr, 2009; Steingroever, Wetzels, Horstmann, Neumann, & Wagenmakers, 2013). It is a partial-information experience-based task (see Figure 1.1–c) where participants receive feedback only from the options they select and this feedback counts towards the total sum in the game. Participants choose repeatedly from a number of options without having any prior knowledge about the magnitude and the distribution of the outcomes. The ultimate goal in the task is to maximise the total winnings.

The original structure of the IGT consists of 4 decks of cards (labelled A-D) from which 100 cards with different monetary payoffs are chosen without replacement. Participants are given \$2000 facsimile money as a loan and are instructed to pick one card at a time from any deck they choose. They must learn that turning each card carries an immediate reward: Selecting a card from the first two decks (A and B) yields \$100 every trial, whereas selecting from the other two decks (C and D) yields \$50 (see Table 1.1). Unpredictably, the turning of some cards also carries a penalty which is large in the high reward decks A and B and small in the low reward decks C and D. Sampling from decks A and B (*bad* or *disadvantageous* decks) leads to an overall loss (a net loss of −\$25 per card), whereas playing from decks C and D (*good* or *advantageous* decks) leads to an overall gain (a net gain of +\$25 per card). Another feature of the task is that the probability of losses varies from deck to deck. In a selection sequence of 10 trials from deck A, the loss of \$1250 is distributed over 5 cards (loss probability 0.5; punishments from \$150 to \$350). In deck B the punishment of \$1250 occurs once, with the selection of one card (loss probability 0.1). A similar pattern of losses is reflected in the other two decks. Specifically, in deck C the \$250 loss is divided across 5 cards (punishments from \$25 to \$75) whereas in deck D, it occurs only once (Bechara et al., 1994).

Table 1.1: Payoff scheme of the Iowa Gambling Task.

	Deck A	**Deck B**	**Deck C**	**Deck D**
Gain per Card	\$100	\$100	\$50	\$50
Loss per 10 Cards	\$1250	\$1250	\$250	\$250
Net per Card	−\$25	−\$25	+\$25	+\$25
Loss Probability	0.5	0.1	0.5	0.1

The IGT assesses decision making under uncertainty (or ambiguity), in the sense that at the outset of the task participants are ignorant of the probabilities of gains and losses (risks) associated with each deck. Experimental and neuroscience studies (e.g., Kahneman & Tversky, 1979; Lee, 2013) have yielded considerable insight into the basic mechanisms of decision making under risk (where the probabilities are known a priori, as in studies of description-based decision-making). Recently, research on decision-making under uncertainty has provided insightful evidence about how people behave when the probabilities and associated payoffs have to be learned by repeated sampling (i.e., experience-based decision-making) and has identified significant behavioural differences between decisions based on description and experience (the description-experience "gap", see Erev & Barron, 2005; Hertwig et al., 2004; Hertwig & Erev, 2009). In fact the IGT may be best conceptualised as a hybrid of the two, in the sense that repeated choices permit the payoff probabilities to be learned (Brand, Recknor, Grabenhorst, & Bechara, 2007). Recent work has suggested that the brain systems engaged in decision-making under risk and uncertainty may be largely overlapping (Levy, Snell, Nelson, Rustichini, & Glimcher, 2010).

In one of the first studies using the IGT, Bechara et al. (1994) showed that VMPFC patients performed significantly worse than healthy participants. Over time, the healthy control group learned to consistently select more cards from the good decks, whereas the patient group continued to select from the disadvantageous decks for the duration of the task. The behaviour of the patient group was guided by the higher immediate rewards rather than the delayed punishments available on the disadvantageous decks. Bechara et al. (1994) proposed the term *myopia for the future* to describe this behaviour. Their interpretation relied on the somatic marker hypothesis: because of the inability of the VMPFC patients to activate the somatic marker system, and integrate and process affective signals and information, their representations of future outcomes could not be marked with positive or negative emotional value (valence) and thus could not be effectively accepted or rejected.

Findings supporting the hypothesised role of somatic markers in decision-making come from several studies employing the IGT with VMPFC patients and normal controls whose electrodermal responses were measured via skin conductance responses (SCRs) as an index of emotional arousal or somatic markers (e.g., Bar-On, Tranel, Denburg, & Bechara, 2003; Bechara, Tranel, & Damasio, 2000; Bechara, Tranel, Damasio, & Damasio, 1996). It has been argued that the VMPFC region involved in the processing of emotion controls the modulation and generation of SCRs (Critchley, Elliott, Mathias, & Dolan, 2000). SCRs generated during the task were divided into three categories: Reward SCRs generated after selection of cards which yielded a reward (win), punishment SCRs after cards which carried a punishment (loss), and anticipatory SCRs (aSCRs) prior to any deck selection. Both patients and controls showed reward and punishment SCRs. However, after a number of card selections the control group started to generate aSCRs which were larger in anticipation of selections from the bad decks, while the lesion group did not develop these responses. The main conclusion was that failure to activate the somatic marker system leads to impaired task performance, consistent with the idea that somatic markers play an important role in guiding decision-making in normal individuals (see Damasio, Adolphs, & Damasio, 2003). Following this, Carter and Smith Pasqualini (2004) reported that the stronger the aSCRs prior to disadvantageous choices, the greater the success of participants in acquiring the advantageous strategy in the IGT (see also Guillaume et al., 2009; Oya et al., 2005).

1.2.1 Deck selection patterns

Bechara et al. (1994) initially suggested that healthy participants begin to consistently select the good decks after an exploration phase. The driving force for this selection pattern is the overall expected values associated with each deck. In other words, decks with a positive total outcome are selected more often compared to the ones with a negative outcome, regardless of other deck features that may affect decisions such as the probability and magnitude of losses associated with each deck (see Figure 1.2–A). Regarding the latter, it was assumed that choices within good and bad decks were uniform, suggesting that the frequency of losses is not important. The crucial point of these observations is that healthy participants will always learn the advantageous strategy (i.e., select decks C and D) in a canonical and predictive manner.

However, it has also been observed that participants select cards from the decks which produce rare losses. Many studies using the IGT have shown that participants favoured the decks with infrequent losses (decks B and D; 0.1 probability of a loss compared to 0.5 on decks A and C) despite the fact that these losses are of greater magnitude (see Figure 1.2–B; Crone & van der Molen, 2004; Lin, Chiu, Lee, & Hsieh, 2007; Yechiam & Busemeyer, 2005). This effect has been observed not only in studies which have employed the traditional payoff scheme of the IGT but in other experience-based tasks where participants underweight the occurrence of rare outcomes (e.g., Barron & Erev, 2003; Hertwig et al., 2004). It has also been found that participants select equally often the good decks C and D and the bad deck B (B, C, and D pattern, Figure 1.2–C; e.g., Premkumar et al., 2008; Wood, Busemeyer, Koling, Cox, & Davis, 2005). This preference in deck selections can be seen as a combination of the previous two patterns: participants select decks on the basis of their overall expected value (good decks) and the rare losses associated with deck B. Overall, the bad deck A is selected less often as its negative expected value and the frequent losses it produces make it unattractive to participants.

Steingroever, Wetzels, Horstmann, et al. (2013) conducted a literature review on IGT studies and concluded that the major assumptions relating to decision-making in healthy participants are essentially invalid. Specifically, there is often no clear preference for good over bad decks, choice behaviour is not uniform, and the usual exploration-exploitation trade-off is rarely observed. Instead people seem to prefer the decks with infrequent losses (decks B and D) with no explicit tendency to exploit the most rewarding options (decks C and D). Similarly, Horstmann, Villringer, and Neumann (2012) concluded that the factors that influence performance on the IGT (in descending order of importance) are: gain frequency, loss frequency, and overall expected value.

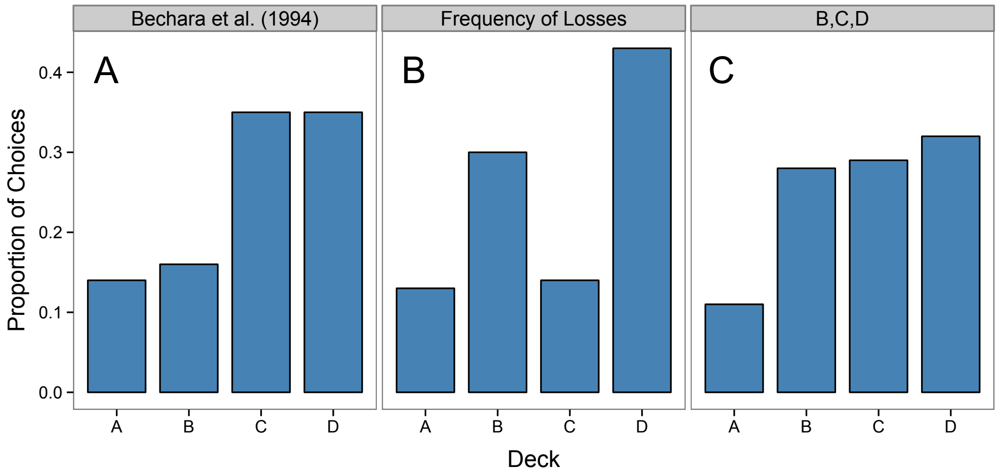

Figure 1.2: Mean proportion of choices from each deck for different selection patterns.

1.3 IGT and conscious knowledge

1.3.1 Implicit learning and cognition

Implicit learning is a term coined by Reber (1967) to describe a kind of learning or knowledge that one acquires without being conscious of it. How do we determine that learning or knowledge is unconscious and what kinds of measures are used to discriminate between mental states? In general, measures of awareness fall into two broad categories: *objective* and *subjective* measures. Objective measures ask participants to make some forced-choice discrimination related to the experimental task (Seth, Dienes, Cleeremans, Overgaard, & Pessoa, 2008; Shanks, 2005), whereas subjective measures require participants to report their internal state of awareness (Tunney & Shanks, 2003). Both tests assess the degree to which a measure of performance on a task is related to reports of mental states, meaning that a dissociation between performance and awareness indicates unconscious learning.

The most common subjective measure is qualitative or descriptive verbal reports. However, a problem that arises when participants are asked to verbally report their awareness is that they may withhold information about the knowledge they possess based on a criterion they have already set. For example, if participants have set a conservative criterion, then they would only choose to report knowledge held with high confidence and therefore knowledge with low confidence may be withheld. Also, participants may report knowledge unrelated to key features of the experimental task and their performance.

A commonly used alternative to verbal reports is confidence ratings (CR), in which participants express their awareness in terms of how confident they feel that they have given a correct answer (e.g., Dienes, Altmann, Kwan, & Goode, 1995; Sandberg, Timmermans, Overgaard, & Cleeremans, 2010). Confidence ratings can be used with different scales including binary, continuous, and percentage scales. In most variations "guess" or "no confidence" are used to denote the lowest rating. Other subjective measures include the perceptual awareness scale (PAS; Ramsøy & Overgaard, 2004; Sandberg et al., 2010), in which participants respond on a 4-point scale about their subjective experience with the task, and post-decision wagering (Persaud, Mcleod, & Cowey, 2007), in which participants place monetary gambles on the correctness of their decisions.

Dienes et al. (1995) proposed two criteria for the assessment of conscious and unconscious processing when using subjective measures of awareness: the *guessing criterion* and the *zero-correlation criterion*. According to the guessing criterion, participants possess some unconscious knowledge if their performance on a task is above baseline while they claim to be guessing or having no knowledge (Dienes, 2004). In other words, if participants' performance is random, no kind of knowledge (conscious or unconscious) affects their performance, but if they perform better than chance when claiming to be guessing, then they possess some knowledge that they are not aware of (unconscious knowledge). The zero-correlation criterion looks for correlations between performance and awareness in a way that if any positive relationship emerges, it means that there is some conscious knowledge which influences performance. If no relationship exists, knowledge is assumed to be unconscious.

1.3.2 Measures of conscious knowledge in the IGT

A major issue concerning the IGT is at what stage in the task participants learn the advantageous strategy and whether this knowledge is assisted by implicit or unconscious biasing signals. In a highly influential study, Bechara et al. (1997) proposed that normal participants decide before knowing the advantageous strategy, meaning that they start to select cards from the good decks before they have conscious knowledge that those decks are the best. Tranel, Bechara, and Damasio (1999) suggested that conscious knowledge alone is insufficient to explain advantageous performance in the IGT. Similarly, Peters and Slovic (2000) used a variation of the IGT and concluded that affective processes have an important role in decision-making and can influence choice independently of conscious knowledge.

To assess participants' knowledge about the task, Bechara et al. (1997) halted participants after 20 trials, and then after every 10 trials and asked them, "Tell me all you know about what is going on in this game" and "Tell me how you feel about the game". Analysis of their responses revealed that participants went through three periods before they reached the *conceptual* period where they had a firm and explicit understanding of the properties of each deck. In the *hunch* period participants developed a preference for the good over the bad decks and generated aSCRs prior to selecting from the bad decks but their verbal responses showed no confidence about this preference. The importance of Bechara et al.'s claim about the existence of unconscious signals comes from the *pre-hunch* period in which participants had experienced some losses but without any conscious insights about what was going on in the task (in the earliest period, *pre-punishment*, participants showed a preference for the bad decks before experiencing any losses from them). The key finding was that aSCRs and card selections from the good decks began in the pre-hunch period (though in fact this was not statistically significant) and were sustained throughout the task indicating that implicit learning was taking place prior to explicit understanding of the reward and punishment schedule for each deck. In other words, Bechara et al. claimed that participants behave advantageously even when their knowledge is still at the pre-hunch period, when their explicit conceptualisation of which were the good and the bad decks had not yet developed.

This proposal about the role of unconscious influences in guiding behaviour in the IGT has been extensively criticised on the basis of weaknesses in the method that Bechara and colleagues used to assess their participants' knowledge. For example, many studies in the implicit learning literature have shown that such broad questions as the ones they employed often fail to identify all of the conscious knowledge that participants have acquired in performing a task (Shanks, 2005; Shanks & St. John, 1994). Several criteria that a reliable measure of awareness must satisfy have been elaborated, such as reliability, relevance, immediacy, and sensitivity (Lovibond & Shanks, 2002; Newell & Shanks, 2014). Bechara et al.'s (1997) assessment does not fulfil any of these criteria. For instance, it is unlikely to be either sensitive (as participants may adopt a very conservative reporting criterion) or relevant (as participants may concentrate on reporting task features unrelated to deck value).

For the aforementioned reasons, Maia and McClelland (2004) developed a more sensitive test of awareness in the form of a structured quantitative questionnaire. After the first 20 card selections and then after every further 10 card selections, Maia and McClelland asked their participants a number of questions in which they had to give ratings on a scale from −10 to +10 concerning how good or bad they thought each deck was and to provide justifications for their ratings. They also asked participants specific questions about the expected wins and losses associated with each deck and their level of confidence that they were aware of the best strategy to win in the game. Also, participants were asked to report which deck they would choose if they could only select cards from one of the decks for the rest of the game.

Using this assessment, Maia and McClelland (2004) found that advantageous performance on the task was accompanied by accurate reports about the values of the decks. They concluded that there is no support for the claims of Bechara et al. (1997) that unconscious biases guide behaviour before conscious knowledge is acquired or that the activation of unconscious somatic markers is necessary in order to perform advantageously. Instead, deck selections in the IGT are driven by conscious knowledge about the decks and by conscious strategies about how to maximise payoffs. Also, the early awareness of the goodness and badness of each deck that Maia and McClelland observed (after only 20 trials) means that aSCRs obtained on the IGT could have been generated by conscious knowledge of the deck payoffs rather than being causally involved in the decision-making process (Dunn et al., 2006). Another interpretation of the high aSCRs before disadvantageous card selections lies in the reward and punishment schedule. Because the amount of money both gained and lost for each card is on average much greater for bad than for good decks, participants' aSCRs would have been higher for bad decks if they were expecting an immediate higher-magnitude reward (Tomb, Hauser, Deldin, & Caramazza, 2002). The possibility that unconscious somatic biases are activated during the task cannot be ruled out, but as Maia and McClelland pointed out, "there is no need to invoke such biases to explain participants' behaviour: verbal reports reflect consciously accessible knowledge of the advantageous strategy more reliably and at least as early as behaviour itself" (p. 16079).

Another divergence between the two studies concerns the trial at which the onset of awareness occurred. Bechara et al. (1997) reported that participants started to have some conscious knowledge on trial 50 on average (range 30-60) and the same finding was reported by Maia and McClelland (2004) in their replication using Bechara et al.'s assessment. However, using the more detailed quantitative questions described above, Maia and McClelland's participants were classified as aware of the difference between good and bad decks even after the first 20 trials. This divergence suggests that the measure employed by Maia and McClelland was considerably more sensitive in revealing the conscious knowledge that participants acquired.

Similar findings have been reported from other studies which employed quantitative and focused questions (e.g., Bowman, Evans, & Turnbull, 2005; Cella, Dymond, Cooper, & Turnbull, 2007; Evans, Bowman, & Turnbull, 2005; Fernie & Tunney, 2013; Wagar & Dixon, 2006). Bowman et al. (2005) assessed participants' knowledge by asking them to rate each deck in terms of how good or bad they felt it was. After the first 20 trials, participants showed substantial awareness of which decks were good and bad and their awareness discriminated the good from the bad decks better than their behavioural performance did, replicating the results of Maia and McClelland (2004).

1.3.3 Evidence for conscious processing

Dunn et al. (2006) suggested that there is little evidence to support the view that deck contingencies are consciously impenetrable and what needs to be tested is whether participants have an explicit understanding of the reward and punishment schedule or whether they can merely discriminate the quality of the decks by attributing positive or negative valances to each one. The important role of conscious knowledge in the IGT is also supported by a study (Gutbrod et al., 2006) with amnesic patients whose deck selections were no better than chance indicating that explicit task knowledge is essential for shaping a behavioural preference towards the advantageous decks. Gutbrod et al. argued that the causal link between SCRs and behaviour "might not be straightforward and that a lack of explicit task knowledge may be sufficient to explain why most of our patients failed to acquire a behavioural preference in the IGT" (p. 1323). Similar findings were also reported by Gupta et al. (2009) who suggested that declarative memory plays a significant role in forming and updating the representation of rewards and punishments associated with each deck. In addition, Stout, Rodawalt, and Siemers (2001) found that IGT impairments in Huntington's disease patients were significantly correlated with explicit memory deficits.

In a recent study, Fernie and Tunney (2013) presented evidence suggesting that autonomic activity or somatic markers are not important determinants of successful performance on the IGT. They observed that conscious knowledge developed after approximately 40 trials and was correlated with advantageous deck selections. Their main results showed that aSCRs did not discriminate between decks prior to the emergence of explicit knowledge, while reward SCRs differentiated between good and bad decks only for those participants who had already acquired some knowledge about the decks' quality. Another interesting finding relates to the punishment SCRs; these were found to be of greater magnitude following larger losses from the bad decks in the initial stages of the task but not after the emergence of knowledge, indicating that participants became aware that the bad decks produce big losses.

The previous findings highlight the importance of cognitive processes underlying performance on the IGT and offer support for the view that emotional or affective signals may not be as important as previously believed. Even though the involvement of emotion-driven learning of the task structure and deck contingencies cannot be entirely ruled out (see Wagar & Dixon, 2006), many studies have pointed out that the contribution of emotional information is rather limited. For instance, the decision-making impairments of VMPFC patients on the IGT can be explained by cognitive deficits (e.g., reversal learning) rather than by any inability to generate emotional or somatic markers (Maia & McClelland, 2005). When the reversal learning component is removed, VMPFC patients' performance on the IGT is comparable to that of normal controls (Fellows & Farah, 2005).

1.3.4 Post-Decision Wagering

In order to avoid some of the complications associated with verbal reports, Persaud et al. (2007) developed a novel non-verbal method of assessing awareness in the IGT in which participants are required to place wagers after their card selections. Persaud et al. characterised post-decision wagering as an objective and direct measure of awareness. When a participant maximises her earnings through advantageous wagering (that is, bets high after a correct decision and low after an incorrect one), this is taken to indicate conscious knowledge about the task.

In Persaud et al.'s (2007) variation of the IGT, participants were asked to make a wager £10 or £20 after each deck selection. The amount of reward, or of reward and punishment, was expressed as a multiple of the chosen wager. The reward and punishment schedule of each deck was modified in order to be dependent on wagering. Selections from decks A and B (bad decks) yielded a win of 2 times the wager whereas selections from decks C and D (good decks) returned the amount of the wager. The frequency of losses was identical to the structure of the original IGT whereas the magnitude was adjusted to reflect the ratio of losses to wins of the original IGT. The net outcome of choosing from the bad decks was a loss of 5 times the average wager per 10 cards, and the net outcome from the good decks was a gain of 5 times the average wager per 10 cards. Thus, the net outcome was either a win of £75 (good decks) or a loss of £75 (bad decks) $[((20+10)/2)\times 5]$ per 10 cards if participants randomly allocated their wagers (50% high, 50% low).

Persaud et al. (2007) investigated the influence of different modes of questioning in parallel with deck selections and wagering in three different groups. The first group was only asked to place a wager, whereas the second and third groups were also given the verbal assessments of Bechara et al. (1997) and Maia and McClelland (2004), respectively. Persaud et al. measured on which trial good deck selection and advantageous wagering began and conjectured that if a significant difference (lag) between these measures emerged, with deck selections revealing a preference for good decks before advantageous wagering emerged, this would indicate an unconscious influence on decision-making. In the first group (wagering only), good deck selection began on trial 40 and advantageous wagering on trial 70. The difference between these was statistically significant and indicated that participants showed a preference for the good decks while failing to maximise their winnings by advantageous wagering. The same pattern was observed in the second group in which participants were asked the open-ended questions used by Bechara et al. Good deck selection started on trial 46 and advantageous wagering on trial 76. However, using the quantitative questions of Maia and McClelland, there was an effect on wagering even though performance on the task in terms of deck selections was similar to the other two groups. Specifically, good deck selection began on trial 36 and advantageous wagering at almost the same time (trial 38).

Persaud et al. (2007) interpreted these findings as demonstrating that the assessment method can affect the knowledge that participants acquire during the IGT. While performance (selecting the good decks) was unaffected, participants gained earlier insight (as measured by wagering) about the reward and punishment schedule and the quality of each deck when they were concurrently asked more specific quantitative questions about the nature of the game. Indeed the onset of advantageous wagering was brought forward by over 30 trials in the group that was periodically asked the Maia and McClelland (2004) quantitative questions. Persaud et al. proposed that Maia and McClelland's assessment method was intrusive and altered participants' awareness and that performance on the IGT is primarily affected by unconscious processes which are masked if the measure of awareness itself makes participants aware of the nature of the task (also see Koch & Preuschoff, 2007; Reimann & Bechara, 2010; Wang, Krajbich, Adolphs, & Tsuchiya, 2012).

Persaud et al. (2007) noted that "Simply asking people might seem a straightforward method, but they may deny awareness if the question asked does not relate to the method they think they used to reach the decision" (p. 257) which is a reasonable critique of the open-ended questions used by Bechara et al. (1997). However, an intriguing issue that arises from their own results is to examine the trial on which participants first demonstrated awareness of the reward and punishment schedule. When no questions were asked or when participants' awareness was assessed by open-ended questions, advantageous wagering - putatively a measure of awareness - appeared quite late in the task (not before trial 70). This pattern, which Persaud et al. did not comment on, is strikingly inconsistent with the studies described above which showed that higher awareness ratings were given for the good decks even in the first 20 trials. Although some minor property of the way they implemented the task or of their participants might have induced this late sensitivity, it raises the important possibility that post-decision wagering is not as sensitive and direct as Persaud et al. claimed.

1.3.4.1 Criticisms of post-decision wagering

Although post-decision wagering seems a well-grounded objective method, it has been the subject of a number of methodological criticisms. First of all, the dichotomous nature of post-decision wagering seems to presuppose that conscious experience is dichotomous as well (see Overgaard, Rote, Mouridsen, & Ramsøy, 2006; Sandberg et al., 2010). If conscious experience does not have this binary character then it is difficult to ascertain when a participant is aware, as a low wager may not imply absence of awareness (Wierzchoń, Asanowicz, Paulewicz, & Cleeremans, 2012). Overgaard (2011) noted that continuous scales or measures which have multiple response options may be more sensitive than binary selections or ratings (but see Tunney & Shanks, 2003). Also, the type of analysis that Persaud et al. (2007) used to demonstrate a dissociation between performance and awareness has been criticised. Szczepanowski (2010) showed that absence of advantageous wagering does not constitute absence of conscious awareness.

Another issue is the influence of loss aversion in wagering strategies. According to prospect theory, humans have an asymmetric utility function; for example, the prospect of losing £5 has greater subjective magnitude than that of winning the same amount of money (Kahneman & Tversky, 1979; Schurger & Sher, 2008). Empirical studies have shown that losses are evaluated roughly twice as much as gains (e.g., De Martino, Camerer, & Adolphs, 2010; Thaler, Tversky, Kahneman, & Schwartz, 1997; Tom, Fox, Trepel, & Poldrack, 2007). Behavioural measures of awareness, such as post-decision wagering, require participants to place a criterion about whether to wager high or low. Hence, any response criterion may be modulated or affective by cognitive biases such as loss aversion (Seth et al., 2008). This consideration is important to allow us to determine whether or not an individual is consciously aware. Specifically, the individual could place a low wager in order to minimise loss even though she has some confidence in her decision. In an artificial grammar study, Dienes and Seth (2010) employed two different measures of conscious knowledge, confidence ratings and post-decision wagering. They found that wagering was affected by loss aversion and that confidence ratings comprised a more sensitive measure of awareness. Dienes and Seth concluded that wagering strategies depend on how loss averse the individual is, meaning that post-decision wagering is not an unbiased measure of conscious content (for a similar conclusion see Fleming & Dolan, 2010). Moreover, the type of reinforcer, real or facsimile, may influence wagering strategies. Persaud et al. (2007) used real money wagers and claimed that real money can increase motivation but it may also increase the influence of loss aversion on wagering and as result decrease sensitivity to conscious knowledge (Dienes & Seth, 2010).

A final issue regarding post-decision wagering is that the optimal strategy for wagering in the experiments of Persaud et al. (2007) is, paradoxically, always to wager high, as this strategy will give the same outcome as wagering low if good *vs.* bad deck discrimination is at chance, but will increase winnings if it is greater than chance. In this sense wagering high can be said to be a *weakly dominant* strategy with Persaud et al.'s payoff matrix as it is either no worse than wagering low, or better. A rational participant would always wager high regardless of her knowledge about the task (Clifford, Arabzadeh, & Harris, 2008). This leads to the question: "How can a failure of a subject to wager optimally be a measure of lack of awareness of the sensory evidence when the optimal strategy is independent of that evidence?" (Clifford et al., 2008, p. 58). Clifford et al. (2008) proposed a solution to this issue by modifying the original payoff matrix used by Persaud et al.

1.4 Cognitive modelling in the IGT

The study of human decision-making has benefited from the application of computational models. The importance of cognitive models lies in the fact that they offer valuable insights into the cognitive or psychological processes under investigation and provide a way to identify and quantify underlying processes which guide performance in cognitive tasks. They also provide formal connections between experimental evidence and theories, make assumptions about psychological phenomena which can be empirically examined, and generate quantitative predictions - instead of abstract verbal descriptions - about how these phenomena would develop in new or different situations (e.g., Ahn, Busemeyer, Wagenmakers, & Stout, 2008; Busemeyer & Stout, 2002; Shiffrin, Lee, Kim, & Wagenmakers, 2008).

In the case of experience-based tasks and in particular the IGT, cognitive modelling analysis constitutes a way to examine and identify the psychological processes which contribute to performance on the task (e.g., Steingroever, Wetzels, & Wagenmakers, in press; Yechiam, Busemeyer, Stout, & Bechara, 2005). The IGT is a complex task in which the interplay between cognitive, motivational, and response processes gives rise to deck selection on each trial (Busemeyer & Stout, 2002). Thus, a formal-quantitative assessment is needed in order to measure and understand the interactions between the psychological processes underlying decision-making in the IGT.

Importantly, the IGT has become a standard screening tool for decision-making deficits in clinical populations. Hence, it is very important to provide a formal modelling framework which could accommodate the behavioural results and make psychologically interesting observations regarding the underlying decision processes. This will lead researchers to identify key differences between clinical groups and healthy controls and make strong connections between neurophysiology and behaviour, leading to a better characterisation of the psychological symptoms of the disorder under investigation (Stout, Busemeyer, Lin, Grant, & Bonson, 2004). A main goal of clinical research is to establish an explanatory framework for particular clinical groups and hence connect pathological behaviour to patterns of behaviour on clinical and cognitive experimental tasks, and to map these deficits to neurophysiological mechanisms. Computational cognitive models serve as the intermediate step between the brain and observed behaviour where performance on a task can be decomposed into its constituent cognitive processes and mapped to neural mechanisms (Busemeyer, Stout, & Finn, 2003).

To this end, several studies have used cognitive modelling analysis to measure and characterise differences in behaviour found between clinical populations and healthy controls (e.g., Busemeyer et al., 2003; Cella, Dymond, Cooper, & Turnbull, 2012; Fridberg et al., 2010; Kjome et al., 2010; Lane, Yechiam, & Busemeyer, 2006; Yechiam et al., 2005). For example, Yechiam et al. (2005) applied a reinforcement-learning (RL) model to discriminate and find key differences among 10 groups of people with neuropsychological disorders and addictions (e.g., VMPFC patients, Parkinson's, Asperger's, cocaine addicts etc.). Although these groups perform very poorly on the IGT, their decision-making deficits could be attributed to different psychological impairments. Using this analysis, the different impairments and psychological processes were quantified (using the parameters of the RL model) and compared to those of healthy controls. Similarly, Fridberg et al. (2010) employed a cognitive modelling analysis to identify differences between healthy controls and chronic cannabis users and Premkumar et al. (2008) used the same analysis to characterise decision-making deficits in schizophrenia.

1.4.1 Modelling approaches to the IGT

1.4.1.1 Reinforcement Learning

A very successful computational approach to decomposing and explaining decision-making processes in the IGT comes from the reinforcement-learning (RL) framework (see Sutton & Barto, 1998). The core idea of the RL framework is that agents interact with the environment and make decisions between possible courses of actions/alternatives based on the learned experience they have had with each alternative up to the point of the decision. Thus, RL is an ideal framework to study decision-making and learning processes in dynamic experience-based tasks (Gureckis & Love, 2009). Optimal decision-making is defined as a process of maximising *rewards* and minimising *punishments* (Niv, 2009). Following this definition, the alternative with the highest learned experience will be chosen with a certain probability on the next decision time-point (i.e., trial). These ideas have been formulated into computational algorithms which describe how agents learn and update the expected value of each alternative (in the RL terminology, the alternatives are called *bandits*) and how they choose an alternative based on these learned values. The choice between bandits is not only governed by the learned experience of each option (that is, the bandit with the highest value will always be selected) but also by a trade-off between *exploration* and *exploitation*. In other words, agents have to first explore the environment and the available bandits and then exploit the most profitable of them. Importantly, in dynamically changing environments (i.e., the value of each bandit changes over time) agents do not perform a one-off switch from exploration to exploitation but are instead engaged in constant switching between these two response states in an attempt to maximise the overall rewards (see Cohen, McClure, & Yu, 2007).

Computational modelling of choice behaviour in the IGT using the RL framework has relied on three assumptions about the underlying psychological processes that drive performance on the task. These assumptions are thought to be directly mapped onto distinct motivational, cognitive, and response processes and translated into computational algorithms (Busemeyer & Stout, 2002, see Chapter 5 for details about the mathematical form of these algorithms). The first assumption relates to the formation of a *utility* or *valence* for the wins and losses experienced after selecting a deck. Because the IGT involves both positive and negative payoffs, agents may give more weight to losses compared to wins (*motivational* processes). These utilities are used to form expectancies for each deck which are learned or updated by an adaptive learning mechanism (*cognitive processes*). Finally, a choice is made which is a probabilistic function of the formed expectancies associated with each deck (*response* processes). Integrating these assumptions into computational algorithms allows the decomposition of the processes involved in complex decision-making and the examination of individual differences in choice strategies.

1.4.1.2 Heuristic, Rule-Based and Instance-Based Learning Models

Alternative modelling approaches for the IGT assume heuristic strategies or simple rules in which the level of computations or cognitive demands is minimal. According to these approaches, agents employ heuristics and simple strategies based on the environment and task demands (Busemeyer & Stout, 2002). One of these approaches is the *Win-Stay-Lose-Shift* (WSLS) strategy which is commonly employed in experience-based tasks (see Steyvers, Lee, & Wagenmakers, 2009; Worthy & Maddox, 2014). In its simplest form this heuristic dictates that participants "stay" with the same option/deck on the next trial if they receive a reward but "switch" randomly to another deck if they do not receive a reward. In the context of the IGT, participants stay with the same deck if the net outcome (wins − losses) on each trial is equal to or greater than zero and shift randomly to one of the three remaining decks if the net outcome is negative (Worthy, Hawthorne, & Otto, 2013). This model does not pose any cognitive demands as participants do not have to learn or keep in memory expected values associated with each deck. A more sophisticated version of the WSLS model requires participants to remember the payoff received on the previous trial (or the average payoff of n previous trials) and stay with the same option on the next trial if the received payoff is equal or greater than the payoff of the previous trial, or switch to other options otherwise (Worthy & Maddox, 2014).

Another model which falls into the same category is the *Strategy-Switching Heuristic Choice Model* proposed by Busemeyer and Stout (2002). This model attempts to capture the reversal learning component observed in IGT choice behaviour. At the beginning of the task, participants assume that the best strategy is to select cards from the bad decks because they return high rewards. However, after experiencing large losses from the bad decks, participants switch strategies and select more cards from the good decks. The choice mechanism of this model is a simple decision tree: the first stage represents the choice of a strategy (good or bad decks), and in the second stage a single deck is selected on the basis of the selected strategy.

The list of models that can be applied to the IGT and other experience-based tasks is very large. As Lejarraga, Dutt, and Gonzalez (2010) pointed out "a common approach in the study of decision making involves observing human performance in a choice task followed by the development of a cognitive model that reproduces that behaviour and predicts new unobserved behaviour within the same task" (p. 143). This has led to the development of many computational models which differ slightly in the way they integrate and test assumptions about behaviour in decision-making under uncertainty. Gonzalez, Lerch, and Lebiere (2003) developed a general theory of dynamic decision-making, *Instance-Based Learning Theory* (IBLT), that can accommodate and explain behaviour in a range of choice tasks. A computational model based on the principles of IBLT, the IBL model, assumes that people choose between options by learning and recalling *instances* from memory produced by experienced outcomes from each option (see Dutt & Gonzalez, 2012a). These instances are *activated* depending on the frequency and recency of experienced outcomes, i.e., more recent and frequent outcomes are more likely to be activated than distant and infrequent outcomes. Finally, a choice is made by selecting the option which has the highest *blended* value. The blended value of an option is defined as the sum of all experienced outcomes from this option multiplied by their probability of recalling each outcome from memory (also see Gonzalez & Dutt, 2011).

1.5 Overview of thesis

The main theme of the present thesis is to examine the role played by unconscious or implicit influences on experience-based decision-making. Our approach involves the investigation of choice behaviour in variations of a classic experience-based paradigm, the IGT, in combination with subjective measures of awareness such as post-decision wagering and confidence ratings (Chapters 2-4). In addition, in order to measure and understand the cognitive and psychological processes underlying performance on the IGT, computational modelling analysis is employed to provide deeper insights into the dynamics of decision-making. Reinforcement-learning models are evaluated using different model comparison techniques and a computational model of confidence ratings in decision-making under uncertainty is developed (Chapters 5-6).

Experiments 1 and 2 in Chapter 2 were designed as direct replications of Persaud et al.'s (2007) IGT study with post-decision wagering. Both experiments demonstrate that learning to make advantageous decisions in the IGT is not dissociable from awareness, which is at odds with the main claims of Persaud et al. Comparison of participants' advantageous wagering performance against their numerical reports in the quantitative questions of Maia and McClelland (2004) suggested that wagering is not a sensitive measure of conscious content.

Chapter 3 (Experiments 3-5) addresses two main response biases of post-decision wagering which depend on the design of the pay-off matrix: the definition of the optimal strategy and loss aversion. Chapter 4 (Experiments 6-7) compares binary and 4-point post-decision wagering against confidence ratings, showing that confidence ratings are indeed more sensitive and exhaustive compared to wagering. Experiment 8 uses confidence ratings to evaluate participants awareness in two experimental conditions: a standard 200-trial version of the IGT and a dynamic alternative, in which decks' payoffs and quality (good or bad) change periodically (every 50 trials). In both groups, participants' metacognitive judgments closely track advantageous deck selections, providing further evidence of the direct and close relationship between learning and awareness.

The final chapters of this thesis are devoted to computational cognitive modelling analyses. Chapter 5 provides a thorough examination and comparison of the proposed choice RL models for the IGT by employing comparison methods in addition to traditional model fitting such as simulation and recovery (model and parameter) methods. In Chapter 6, we develop a novel computational model as an extension to the RL models in order to account for participants' confidence ratings. The main results suggest that the same mechanism which is responsible for choice behaviour can give rise to confidence judgments. Finally, Chapter 7 summarises the main findings and provides suggestions for future research.

2 Replication of Persaud et al. (2007)

Our first experiments take a closer look at the study of Persaud et al. (2007) and their claims regarding the influence of implicit or unconscious processes on decision-making under uncertainty. This is an important starting point to gain better insights into post-decision wagering. It is also important to check whether the results in Persaud et al. are reproducible. To anticipate the results of our replication study, we only partially replicated the findings of Persaud et al. in Experiment 1. Thus, Experiment 2 was designed to provide further evidence of whether the difference in the pattern of results between Experiment 1 and that of Persaud et al. (2007) persists with a different group of participants.

2.1 Experiment 1

Experiment 1 included two groups in an attempt to reproduce the key findings reported by Persaud et al. (2007). Both groups performed the IGT and made post-decision wagers. In the questionnaire group participants were also regularly asked a subset of Maia and McClelland's (2004) quantitative questions, while those in the control group were not. Comparisons between these groups allow a number of issues to be addressed: First, in the control group, is there evidence that deck selections begin to discriminate good from bad decks before the trial at which advantageous wagering first occurs? This is the key piece of evidence for an unconscious influence on decision-making. Secondly, does quantitative questioning bring forward the point at which advantageous wagering occurs, as Persaud et al. suggested? Thirdly, what is the comparison between wagering and quantitative judgments in the questionnaire group? Although they included such a group, Persaud et al. did not report the quantitative judgments their participants made. Even if making these judgments has the effect of focusing participants' attention on the task and rendering them more rapidly aware of the task structure (and hence improving wagering), it is still of considerable interest to examine such data. Importantly, we can ask whether the quantitative assessments participants make at their first assessment (trial 20) - when questioning cannot have had any effect on task awareness - reveals awareness which is undetected by the wagering measure.

Experiment 1 thus aimed to replicate the design and methodology employed by Persaud et al. (2007). One difference between our experiment and Persaud et al.'s lies in the format of the IGT. We used a computerised version of the IGT whereas Persaud et al. used a classic manual format (see Bechara et al., 1994). Previous studies observed no differences in the pattern of deck selections between format types, however (Bechara, Tranel, & Damasio, 2000; Bowman et al., 2005).

2.1.1 Method

Participants

Thirty volunteers (20 females) between the ages of 19 and 30 years ($M = 22.66, SD = 2.96$) were recruited from the University College London subject pool. All participants received £3 for their participation.

Design

Participants were randomly assigned to one of two groups (no questioning [control], quantitative questioning [questionnaire]). The control group made a high or a low wager following each deck choice whereas the questionnaire group, in addition to wagering, were asked a subset of the Maia and McClelland (2004) questions every 20 trials. In the original study, the questionnaire was given to participants every 10 trials after the first administration on trial 20. We reduced the frequency of administering the questionnaire to limit fatigue.

Task

A computerised variant of the IGT was employed. There were four decks of cards with labels A, B, C, and D. The rewards and punishments were the same as in Persaud et al. (2007) and these were dependent on the quality of the deck (Good or Bad) and the wager (High: £20 or Low: £10). Specifically, selecting a card from the bad decks (A and B) yielded a win of two times (2 ×) the wager (High: £40, Low: £20) whereas selecting a card from the good decks (C and D) returned the wagered amount (1 ×; High: £20, Low: £10). Also, some trials carried a punishment: the distribution and frequency of the punishments were as for the original IGT whereas the magnitude was adjusted to reflect the ratio of losses to wins of the original IGT (see Appendix A for the reward and punishment schedule).

Participants in both groups were given an initial endowment of £400 of play money and were asked to maximise their earnings. The task comprised 100 card selections. After each card selection, a frame appeared on the screen with two alternative choices, "High (£20)" and "Low (£10)", allowing participants to place a wager on their card selection. Along with wagering the questionnaire group was administered a modified version of Maia and McClelland's (2004) questionnaire (see Appendix B). The qualitative parts of the questionnaire were omitted and it was administered every 20 trials. Participants were asked to provide ratings of the "goodness" of each deck, to report or calculate amounts of money related to the decks' payoffs, and to indicate which deck they would select cards from for the rest of the task if they could only choose from one deck.

Instructions were presented on the screen before the experiment started. At the top of the display was a green bar that expanded or contracted according to the amount of money won or lost after each deck selection and wager. Every time a participant clicked on a deck to pick a card, the deck was highlighted and the wagering frame appeared on the screen. After the wagering selection, the face of that card appeared on the top of the deck showing the amount of money behind the card and a message was displayed on the screen indicating the amount of money won or lost. Once the money had been added or subtracted, the face of the card disappeared and the participant could select another card.

Procedure

Participants sat in front of a PC display. They were then asked to read the on-screen instructions about the task. In brief, participants were told that the game was about learning to gamble on card selections, that all of the cards would yield some money but some would lose money, that their objective was to win as much as money as possible, and that they were free to switch from one deck to another at any time. Additionally, participants were presented with instructions about wagering. Specifically, they were told that if they were confident that their choice would give them some net winnings, then they should wager high, otherwise, they should make a low wager. The questionnaire group was presented with instructions about the quantitative questions. Each session ended after 100 trials.

2.1.2 Results

Choice and Wagering

Advantageous wagering was defined as either a high wager after choosing a good deck or a low wager after choosing a bad deck. Our analyses employed the average proportion of good deck selections (choice) and advantageous wagers (wagering) across subjects over successive blocks of 10 trials to investigate any differences between the two groups (control, questionnaire) and to locate the onset of learning and awareness (see Figure 2.1).

The onset of deck discrimination, as revealed by the first block in which choice of the good decks was significantly above chance (0.50), was at block 4 for both conditions (Control: $M = 0.63, t(14) = 3.25, p = .006, d = 0.84$, Questionnaire: $M = 0.70, t(14) = 5.29, p < .001, d = 1.37$). Advantageous wagering exceeded chance level (0.50) at the same time as choice, also in block 4 (Control: $M = 0.63, t(14) = 3.08, p = .008, d = 0.80$, Questionnaire: $M = 0.66, t(14) = 3.43, p = .004, d = 0.89$). These results indicate that there was no advancement in the onset of advantageous wagering in the quantitative

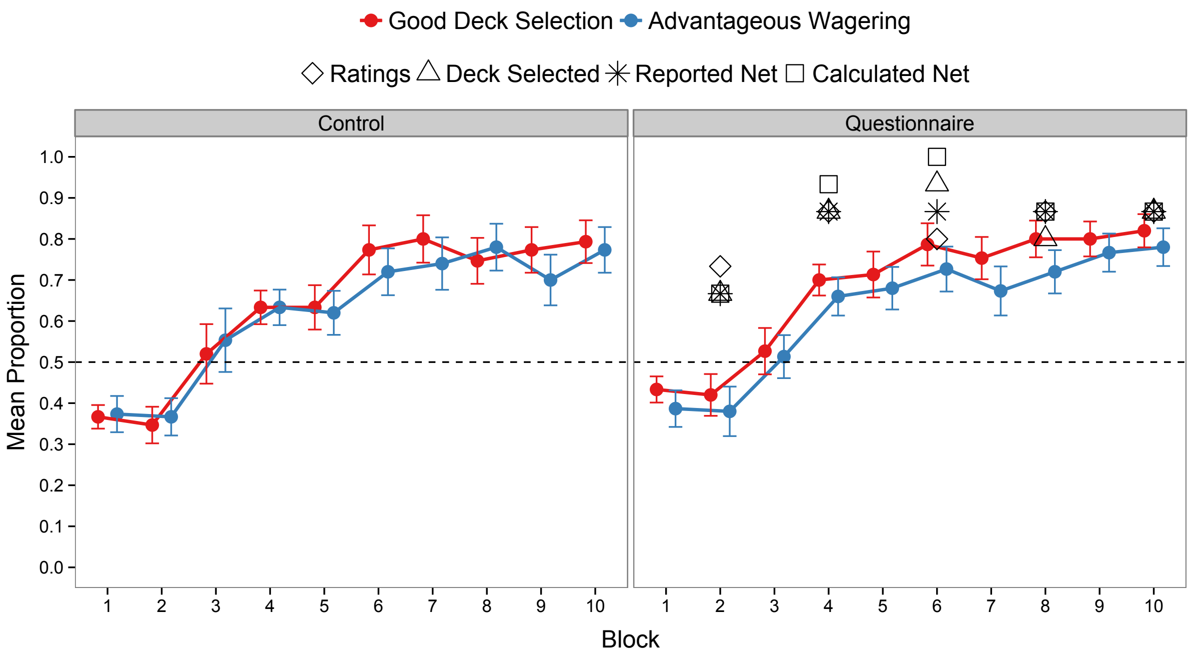

Figure 2.1: Proportion of good deck selections and advantageous wagering across blocks of trials for the control and questionnaire groups in Experiment 1 (lines). The grey diamond and the triangle markers represent the proportion of participants who gave higher ratings to one of the two best decks and the proportion of participants who selected one of the two best decks as their choice if they were allowed to select only one deck. The star and the square markers represent the proportions of participants who gave the highest reported expected net and calculated net to one of the two best decks. Points are offset horizontally so that error bars ($\pm$ 1 SEM) are visible.

questioning group who made explicit judgments about the deck payoffs.

Two separate mixed ANOVAs were performed on the proportion of good deck selections and advantageous wagers across blocks of 10 trials. It is important to note that even though they use the same scale, the two measures cannot be compared directly because advantageous wagering is dependent on the first-order decision (e.g., deck selection) and this creates the possibility of functional differences between the measurement scales. For example, if a participant always chooses a good deck (with the proportion of good deck selections therefore being 1.0), but decides to make both high and low wagers because she is more confident on some trials than others, then advantageous wagering cannot attain a value of 1.0. Its maximum value under such circumstances would be equal to the proportion of high wagers. Like Persaud et al. (2007), our contrast between deck selection and wagering is therefore an indirect one, based on estimating the trial block at which each reaches a level significantly above chance. This contrast is likely, if anything, to be biased in favour of obtaining evidence of learning without awareness. Both measures could be numerically above chance, but deck selection might be significantly so and wagering not (it might be a noisier measure, for instance).

A 2 (group [control, questionnaire]; between) $\times$ 10 (block: 10 trials each; within) mixed ANOVA was performed to assess group differences on good deck selections. For the main effect of block, polynomial contrasts were also applied. The analysis revealed a significant main effect of block, $F(9, 252) = 26.80, MSE = 3.00, p < .001, \eta_g^2 = 0.39$ (for generalised eta squared, η_g^2, see Bakeman, 2005), indicating that participants learned about the quality of each deck as there was a tendency for choice to increase across time (significant linear and quadratic effects, $p < .001$). The main effect of group did not reach significance, $F(1, 28) = 0.92, MSE = 10.94, p = .35, \eta_g^2 = 0.01$, and the interaction between group and block was not significant, $F(9, 252) = 0.39, MSE = 3.00, p = .94, \eta_g^2 = 0.01$, suggesting that the mean proportion of good deck selections across blocks was similar in the two conditions. This finding is in accordance with Persaud et al.'s (2007) results: in their study more detailed questioning did not affect participants' deck-selection strategies. In other words, when awareness is probed by more "invasive" methods, no effect is observed in the application of this knowledge to decision-making behaviour.

Analysis of the proportion of advantageous wagers revealed a similar pattern of results. Again, the main effect of group was not significant, $F(1, 28) < 1, MSE = 13.12, p = .86$. Participants were able to maximise their winnings as the proportion of advantageous wagers increased across blocks, $F(9, 252) = 19.35, MSE = 3.44, p < .001, \eta_g^2 = 0.32$. The interaction between group and block was not significant, $F(9, 252) = 0.46, MSE = 3.44, p = .90, \eta_g^2 = 0.01$.

These results indicate that participants favoured the good decks and became gradually capable of maximising their winnings by placing appropriate wagers. The estimated onsets of good deck selections are consistent with those reported by Persaud et al. (2007; also see Turnbull, Evans, Bunce, Carzolio, & O'Connor, 2005; Wagar & Dixon, 2006), namely on trials 40 and 36 (block 4) for their control and questionnaire groups, respectively. The key result though is that advantageous wagering developed according to approximately the same time-course as choice behaviour. The extra requirement to rate the quality of the decks and answer questions about the payoffs did not affect participants' decision-making or wagering strategies. Since choice and wagering displayed similar patterns in both groups there is no evidence of a dissociation between learning and awareness of the optimal strategy, assuming that wagering is indeed a valid index of awareness. The simultaneous onset of awareness in the two groups also contradicts the main claim of Persaud et al. about learning without awareness in the IGT. Specifically, Persaud et al. reported that in their control group, where no quantitative questions were asked, advantageous wagering lagged behind deck selections whereas this was not the case in their quantitative questioning group. This pattern was not observed here.

Signal Detection Theory Analysis

Another approach to investigate participants' awareness is to use the well-established framework of signal detection theory (SDT). Type 2 SDT provides valuable insights into participants' metacognitive sensitivity and wagering strategies by measuring sensitivity and bias independently (Clifford et al., 2008; Fleming & Dolan, 2010; Higham, 2007; Kunimoto, Miller, & Pashler, 2001). In applying SDT to the IGT and post-decision wagering, a *hit* is a high wager after a good deck selection and a *false alarm* a high wager after a bad deck selection. A constant of 0.5 was added to the counts of hits, false alarms, misses and correct rejections in order to prevent infinite values for the calculation of d' (metacognitive sensitivity) and ln β (metacognitive bias; Barrett, Dienes, & Seth, 2013; Snodgrass & Corwin, 1988).

A 2 (group [control, questionnaire]; between) $\times$ 10 (block: 10 trials each; within) mixed ANOVA was performed to investigate whether there were any differences in d' between the two conditions (see Figure 2.2). The analysis revealed no main effect of condition, $F(1, 28) = 0.02, MSE = 1.73, p = .90$. Figure 2.2 shows a tendency for d' to increase across blocks, resulting in a significant main effect of block, $F(9, 252) = 14.26, MSE = 0.60, p < .001, \eta_g^2 = 0.28$. As expected, the group $\times$ block interaction was not significant indicating that the questionnaire did not increase participants' metacognitive sensitivity, $F(9, 252) = 0.39, MSE = 0.60, p = .94$. Moreover, the analysis is consistent with the previous analyses as the first point at which d' significantly exceeds chance ($d' = 0$) is block 4 in both conditions (Control: $t(14) = 2.17, p = .04$, Questionnaire: $t(14) = 2.24, p = .04$).

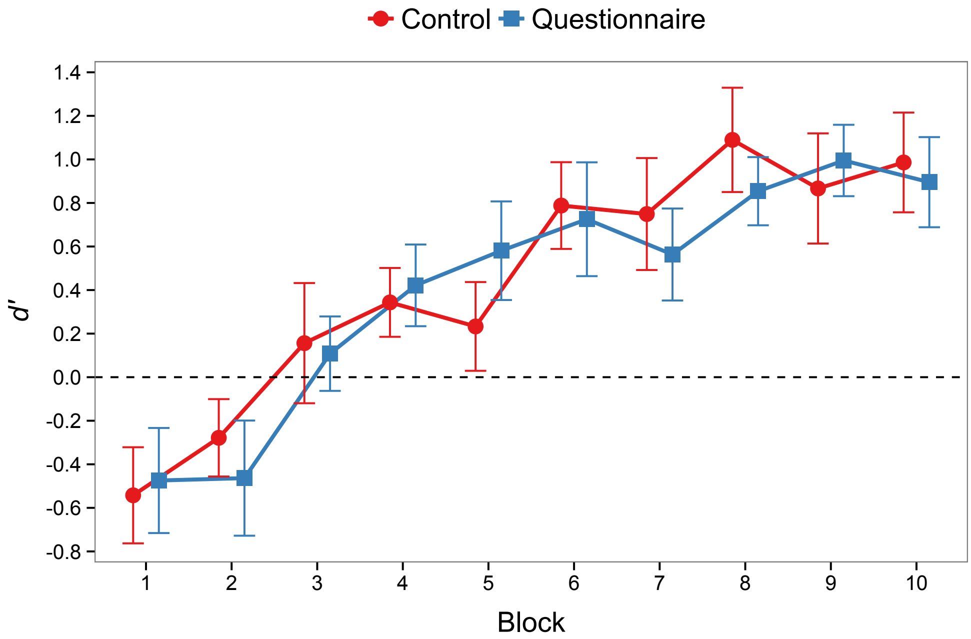

Figure 2.2: Participants' metacognitive sensitivity (d') as a function of questioning group and block of trials in Experiment 1. Error bars represent ± 1 SEM.

Moreover, we can obtain useful insights into participants' wagering strategies by examining the bias measure ln β (ln β = 1 if no bias; ln β > 1 if conservative; ln β < 1 if liberal; Higham, 2007; Macmillan & Creelman, 2005). Analysis of variance revealed that neither the main effect of condition, $F(1, 28) < 1, MSE = 1.00, p = .84$, nor the interaction (group × block), $F(9, 252) < 1, MSE = 0.29, p = .94$, were significant. There was a significant main effect of block, $F(9, 252) = 7.63, MSE = 0.29, p < .001, \eta_g^2 = 0.17$, as participants became more liberal across blocks (mean ln β ranged from 0.04 (block 1) to -0.58 (block 10) in the control group and from 0.26 (block 1) to -0.54 (block 10) in the questionnaire group).

Questionnaire

Participants' knowledge regarding the advantageous strategy in the questionnaire group was explored. Two of the measures reflect knowledge about the general quality of each deck and the remaining two about the actual payoffs. For questions 1 and 3 (see Appendix B) if a participant gives the highest rating to one of the two best decks and selects one of the two best decks to pick cards until the end of the experiment, that means the participant possesses accurate knowledge about the task. In the same manner, when the highest reported (Question 2.a) and calculated net (Questions 2.b, 2.c and 2.d) is attributed to one of the best decks, this indicates high levels of awareness. The calculated net (CN) for each participant, deck, and question period is obtained using the following equation: $\text{CN} = \text{Q2.b} + (\text{Q2.c}/10) \times \text{Q2.d}$.

Figure 2.1 (Questionnaire group) shows the proportion of participants whose answers favoured the good decks on each of the questionnaire measures. Participants whose verbal responses did not discriminate between good and bad decks (i.e., they give the same ratings or the same reported net for all decks) do not count towards this proportion. Inspection of the figure shows that participants exhibited substantial knowledge about the quality of each deck even in the first assessment period (trial 20). Not only did they rate the good decks higher than the bad decks, but also they had a firm basis for such an attribution as revealed by their reported and calculated net payoffs. Table 2.1 also shows the mean deck ratings (-10 very bad, +10 very good) for each deck and the proportion of selections throughout the task for both groups. The results show a clear trend, that is the more positive the rating for a deck, the more likely the deck was to be selected. This correlation between ratings and selections adds further support to the view that decision strategies in the IGT develop in parallel with explicit knowledge.

The use of the questionnaire allowed us to explore differences between the two measures (i.e., post-decision wagering and questionnaire) in terms of how sensitive each is in assessing participants' awareness. It is important to check whether the quantitative questions reveal more knowledge about the task than wagering in the first assessment (trial 20). However, the two measures are not directly comparable due to the fact that

Table 2.1: Mean ratings and proportion of selections for each deck in Experiment 1.

		Proportion of Selections	
Deck	Mean Ratings (SD)	Questionnaire	Control
A	-3.01 (3.69)	0.12	0.12
B	-2.71 (4.86)	0.20	0.24
C	0.67 (3.37)	0.33	0.30
D	2.17 (3.33)	0.34	0.34

the questionnaire was administered once at trial 20 whereas participants placed wagers after each deck selection. To overcome this problem, we classified each participant as aware or unaware based on the average proportion of advantageous wagers placed across trials 16-25. If the average was equal to or greater than .5, then participants were identified as aware of the advantageous strategy. The proportion of participants classified as aware by the wagering measure was then compared against the proportion of participants who favoured one of the two good decks in Question 3 (i.e., deck-selected measure). We used the deck-selected measure because it requires only one response and is therefore similar to wagering. To see whether the two proportions were significantly different (deck-selected: 0.67, wagering: 0.33) we used the McNemar test for dependent proportions (see Agresti, 2002; Wild & Seber, 1993) which was found to be non-significant, $\chi^2(1) = 2.78, p = .096$, possibly due to the small sample size (N = 15). Nevertheless, the numerical difference suggests that wagering underestimates participants' acquired knowledge possibly due to the effects of biases in participants' wagering strategies.

2.1.3 Discussion

We draw three principal conclusions from Experiment 1. First, under the conditions tested here awareness as measured by wagering tracked deck selections quite closely. We found no indication that wagering lagged behind the selection of good decks, with both measures becoming reliably better than chance fairly early in the task, between trials 30 and 40. Secondly, the results of the explicit questions revealed that wagering, if anything, underestimates task insight. As early as trial 20, the majority of participants were able to give accurate reports about the quality of the different decks. Thirdly, there was no evidence that eliciting explicit reports in the questionnaire group altered participants' wagering strategy. Persaud et al. (2007) did report such a bias, but it was not observed here. Regardless of whether they explicitly reported their task knowledge, participants began to wager advantageously in block 4 (and this is about the same point at which they began to reliably select the better decks).

2.2 Experiment 2

Experiment 1 revealed no evidence of a dissociation between learning and awareness. In both experimental groups, awareness of the advantageous strategy emerged early in the task as shown by advantageous wagering and participants' responses following the administration of the quantitative questions every 20 trials. This pattern of results is at odds with the claims of Persaud et al. (2007) about learning without conscious knowledge in the IGT and the intrusive nature of the questionnaire of Maia and McClelland (2004), which (they argued) alters participants' awareness and makes them more aware of the task structure. Experiment 2 was designed as an exact replication of the control group of Experiment 1. Since the pattern of results deviated from what was observed in the no-questioning group of Persaud et al., the purpose of Experiment 2 was to ascertain whether the effect we obtained in Experiment 1 was reliable and consistent. In other words, the question is whether advantageous wagering will closely track good deck selections and thus provide additional evidence regarding the close connection between learning and awareness in the IGT.

2.2.1 Method

Participants

Twenty undergraduate psychology students participated (10 females, age $M = 20.25, SD = 0.97$) and they received £2 for their participation.

Task/Procedure

The task and procedure of Experiment 2 was identical to the control group of Experiment 1.

2.2.2 Results

Choice and Wagering

The method for identifying the onset of good deck selections and advantageous wagering was the same as in Experiment 1. Performance exceeded the chance level at block 6 for both measures (Choice: $M = 0.67, t(19) = 2.93, p = .009, d = 0.66$, Wagering: $M = 0.63, t(19) = 2.05, p = .05, d = 0.46$) (see Figure 2.3).

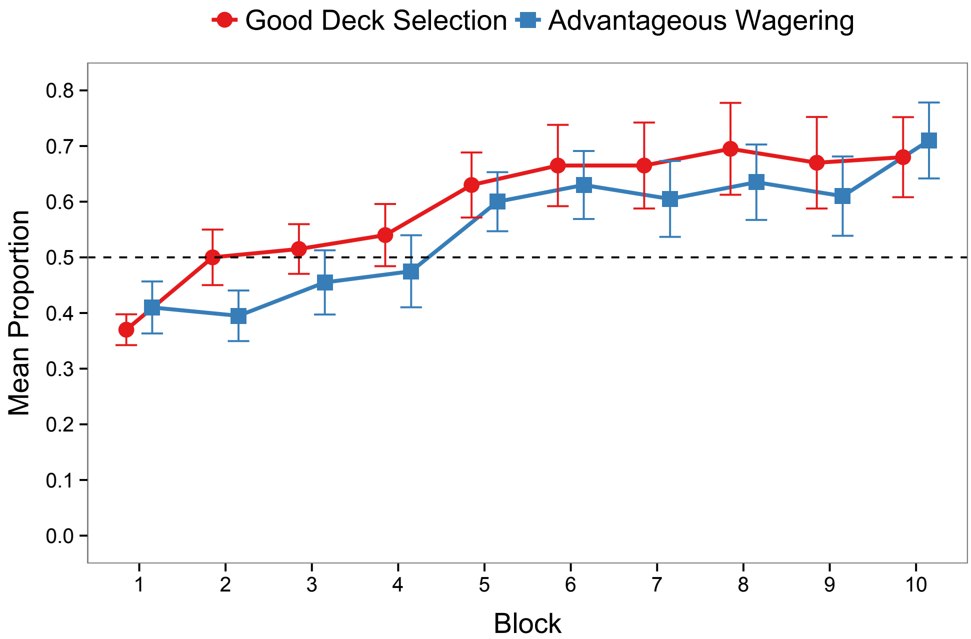

Figure 2.3: Proportion of good deck selections and advantageous wagering across blocks of trials in Experiment 2. Error bars represent ± 1 SEM.

A repeated-measures ANOVA revealed a significant effect of block on good deck selections, $F(9, 171) = 4.70, MSE = 4.87, p < .001, \eta_g^2 = 0.13$. Wagering performance closely followed the optimal decision-making strategy as demonstrated by a main effect of block, $F(9, 171) = 4.62, MSE = 5.11, p < .001, \eta_g^2 = 0.14$. These findings provide another piece of evidence that when participants start to consistently select cards from the good decks, they also possess conscious knowledge of their decisions (in this case, as indexed by wagering). In other words, the results from this experiment indicate once again that there is no dissociation between learning and awareness of the optimal strategy, as claimed by Persaud et al. (2007). In addition, the pattern of this experiment resembles the one observed in Experiment 1; that is, advantageous wagering did not lag behind good deck selections.

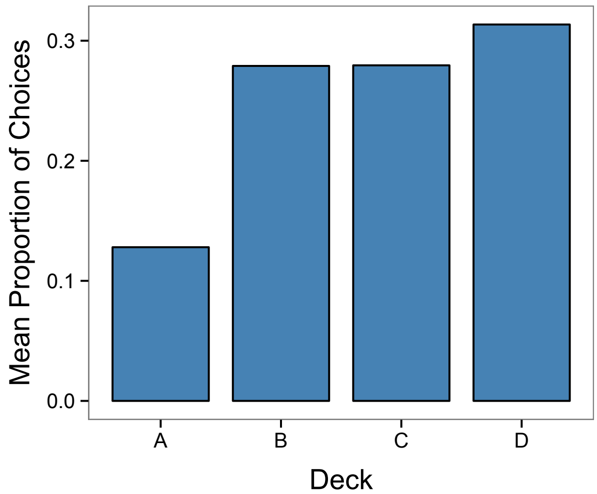

Figure 2.4: Mean proportion of choices from bad decks (A and B) and good decks (C and D) in Experiment 2.

However, there is a difference of about 2 blocks of trials in the onset of both measures (i.e., good deck selections and advantageous wagering) compared with the control group of Experiment 1. One possible explanation for this difference lies in the nature of the IGT as a probe of decision-making. In their literature review on IGT studies, Steingroever, Wetzels, Horstmann, et al. (2013) observed high variability in performance between studies but also within participants in the same study. Specifically, performance on the IGT ranged between .43 and .70 (proportion of selections from the good decks) with most studies reporting only a weak preference for the good decks, between .50 and .60. Similarly, Fernie and Tunney (2008) associated difficulties with manipulations of the reward and punishment schedule of the IGT; because wins, loss probability and magnitude, and overall expected values are all confounded with each other it is difficult to ascertain which aspect of the schedule has a bigger effect on choice behaviour. Thus, it is not surprising that participants differ in the way they perform on the IGT across studies and that individual differences exist in a complicated task which depends on each individual-difference variables such as risk propensity, general motivation, and working memory capacity. For example, some participants never learn to select the good decks whereas for others a few trials suffice in order to start behaving advantageously. In the present experiment, the percentage of good deck selections (C and D) across blocks was .593, indicating a weaker learning effect compared to the control group of Experiment 1 (.638; see Figure 2.4). The pattern of individual deck choices resembles the “B,C, and D” pattern described in the Introduction (see Figure 1.2). Participants’ choices are driven by the expected values (or overall goodness) of each deck but also by a loss-frequency effect; that is, even though deck B is a disadvantageous deck it is selected as often as the good decks C and D because it produces rare losses (0.1 probability of loss, 1 per 10 trials)

2.3 Discussion

The results of Experiments 1 and 2 do not offer any support for the claims of Persaud et al. (2007) that learning to make advantageous decisions can occur in the absence of awareness. We only replicated the results relating to the quantitative questioning group where deck selection and advantageous wagering exceeded chance at the same time. In contrast to Persaud et al.’s results, the same pattern was observed in the group that was asked only to make a wager after their deck selections, suggesting no dissociation between choice and wagering.

Even though the pattern of results between the two experiments is qualitatively the same (i.e., advantageous wagering did not lag behind good deck selections), a difference was observed in the time that both measures exceeded chance, which could be due to the high variability observed in performance on the IGT.

3 Optimal Strategy, Loss Aversion, and Pay-off Matrices

In the present chapter, we will focus on two problematic and rather contradictory aspects of post-decision wagering: the definition of the optimal strategy and loss aversion. As mentioned in the Introduction, the optimal wagering strategy under the pay-off matrix of Persaud et al. (2007; Table 3.1) is always to wager high, irrespective of the acquired task knowledge. This is because the same amount of money can be won or lost on any given wager (Clifford et al., 2008). If participants perform randomly (i.e., have no knowledge about the quality of each deck) then their expected gain from betting either high or low is zero. However, when they start to realise which are the good decks, betting high will increase their winnings. Following this, wagering high constitutes a *weakly dominant strategy* as it is never worse than wagering low. Consequently, this optimal strategy offers no insight about the knowledge participants have acquired during the task.

Table 3.1: Payoff matrices for the different combinations of deck selection and wager.

	Persaud et al.		Clifford et al.		Schurger & Sher	
	Deck Selection					
Wager	Good	Bad	Good	Bad	Good	Bad
Low	+1	-1	+2	-1	+1	-2
High	+2	-2	+5	-5	+10	-10

A second important issue is the influence of loss aversion on wagering strategies. The placement of a high or low wager can be strategy dependent, reflecting people's greater sensitivity to losses than equivalent wins. In other words, people may wager conservatively (i.e., low wagers) in order to minimise losses even if they have some confidence in their decisions. Many studies have shown that post-decision wagering is susceptible to loss aversion biases which makes it less sensitive and exhaustive than other measures of awareness (e.g., Fleming & Dolan, 2010; Szczepanowski, Traczyk, Wierzchoń, & Cleeremans, 2013; Wierzchoń et al., 2012). To overcome issues associated with loss aversion, Schurger and Sher (2008) proposed a different pay-off matrix which encourages participants to wager high (Table 3.1). Under this matrix low wagers are meant to avoid large losses so that participants would bet high even in situations when their task knowledge is not fully developed.

3.1 Experiment 3

Experiment 3 further examines the utility of wagering as a valid alternative to verbal reports for assessing awareness by applying two modifications to Persaud et al.'s (2007) procedure. First, the original reward and punishment schedule of the IGT was used, and secondly we tested the modified pay-off matrix proposed by Clifford et al. (2008). As noted earlier, the pay-off matrix used by Persaud et al. encourages rational participants to employ the weakly-dominant strategy of making high wagers all the time, irrespective of the knowledge they possess about the decks (see Table 3.1). The modified version of the pay-off matrix, in contrast, encourages participants to wager low under uncertainty and to wager high when they have acquired some knowledge about the decks.

Specifically, in the modified matrix participants are discouraged from wagering high until they feel confident that their decision is a good one. When discriminative knowledge about the decks is absent or low, it is advantageous to wager low. This can be shown by the expected payoff from wagering low which is $+1/2[=(+2-1)/2]$ compared to $0[=(+5-5)/2]$ from wagering high. However, when deck discrimination is better than chance, it is more rewarding to wager high due to a larger payoff with a good/high combination (+5) than a good/low one (+2). Based on this matrix a rational participant (i.e., a participant who seeks to maximise gains) would start to wager high only when her deck discrimination (i.e., probability of selecting a good deck is) 4/7 or .57. The latter can be computed from the differential loss of wagering on a bad decision $(5 - 1 = 4)$ divided by the sum of the differential loss and the differential gain of wagering on a good decision $(5 - 2 = 3$; Clifford et al., 2008).

Experiment 3 therefore asks two main questions. First, we have another opportunity to examine whether awareness as measured by wagering lags behind deck selection. Secondly, we can ask whether the modified payoff matrix locates the onset of awareness at an earlier point than the original matrix. Given that Experiment 1 suggested that wagering (under the original matrix) locates the onset of awareness far too late (in comparison to numerical reports on the values of the decks), it is possible that the modified matrix will yield a more appropriate, earlier, estimate.

3.1.1 Method

Participants

Sixty healthy volunteers participated (28 females, age $M = 22.32, SD = 3.02$). Thirty-five participants were recruited via the subject pool and the rest were undergraduate students who received course credit for participating. Participants were randomly assigned to the two conditions.

Design

The simple-wagering group participated in a replication of Persaud et al.'s (2007) IGT task with wagering. The differences between Persaud et al.'s study and this experiment are that we employed the original reward and punishment schedule of the IGT, and the wagers were divided by a factor of 10. The modified-wagering group was administered the IGT with wagering but using the pay-off matrix proposed by Clifford et al. (2008; see Table 3.1).

Task

The reward and punishment schedule used in this study was the same as in the original IGT. After selecting a card from decks A and B participants won £100 whereas on decks C and D they won £50. However, on some trials, there was a punishment which was larger on decks A and B compared to decks C and D. On deck A, 50% of the trials carried a punishment (varying from £150 to £350) leading to an overall loss of £250 every 10 trials. On deck B, the net outcome was the same as in deck A (−£25 per card) but there was one large loss (£1250) every 10 trials. The same pattern was present on decks C and D; on deck C, 5 out of 10 trials had a punishment (from £25 to £75) leading to an overall gain of £250 (+£25 per card) whereas on deck D, there was one loss (£250) every 10 trials.

After each card selection in the simple wagering group, a new frame appeared on the screen with two alternative choices, "High (£2)" and "Low (£1)", allowing participants to place a wager on their card selection. The amount behind the card was multiplied either by 2 or 1 according to wager selection. In the modified wagering group, the procedure was the same, except that the wagers were not expressed as amounts of money but simply as "High" and "Low". This is because the final amount of money won or lost after each card selection was multiplied by the appropriate weights in the modified pay-off matrix.

The task comprised 100 card selections. Because there were only 40 cards in each deck, it was possible to run out of cards from a given deck (as in the original IGT). When this happened, a message appeared on the screen instructing participants to stop choosing from that deck and to continue selecting from the remaining decks.

Procedure

The procedure of Experiment 3 was identical to that of previous experiments.

3.1.2 Results

Choice and Wagering

The method for identifying the onset of good deck selections and advantageous wagering was the same as in previous experiments (see Figure 3.1). Good deck selection commenced on block 5 for the simple wagering group, $M = 0.63, t(29) = 3.20, p = .003, d = 0.58$, but on block 4 for the modified wagering group, $M = 0.68, t(29) = 4.81, p < .001, d = 0.88$. There hence seems to be a small difference in the onset of learning between the two groups. Additionally, a difference was observed regarding the onset of advantageous wagering. Specifically, in the simple wagering group awareness arose relatively late in the task, on block 7, $M = 0.61, t(29) = 2.82, p = .008, d = 0.51$, whereas participants started to place appropriate wagers on block 4 in the modified wagering group, $M = 0.69, t(29) = 5.48, p < .001, d = 1.00$.

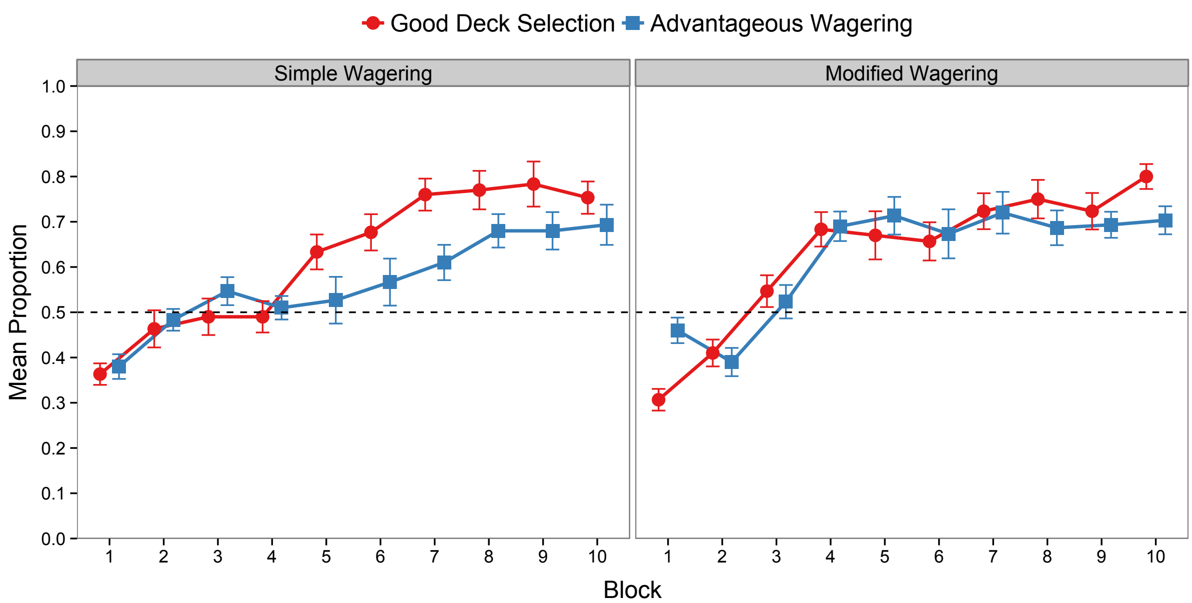

Figure 3.1: Proportion of good deck selections and advantageous wagering across blocks of trials in the simple wagering and modified wagering groups in Experiment 3. Points are offset horizontally so that error bars (± 1 SEM) are visible.

A 2 (group [simple wagering, modified wagering]; between) × 10 (block: 10 trials each; within) mixed ANOVA on the proportion of good deck selections revealed a non-significant main effect of group, $F(1, 58) = 0.22, MSE = 5.24, p = .65$. There was a significant main effect of block, $F(7.19, 416.22) = 32.59, MSE = 4.22, p < .001, \eta_g^2 = 0.33$ (significant linear and quadratic trends). The analysis also revealed a significant group × block interaction, $F(7.19, 416.22) = 2.18, MSE = 4.22, p = .022, \eta_g^2 = 0.03$ (Greenhouse-Geisser correction). Simple effects analyses showed a significant difference on block 4 ($F(1, 58) = 14.86, MSE = 3.77, p < .001$) which is consistent with the difference reported above in the onset of good deck selections between the two groups.

A similar analysis was performed on the proportion of advantageous wagers. The main effect of group was significant, $F(1, 58) = 4.15, MSE = 12.02, p = .046, \eta_g^2 = 0.02$, as participants in the modified wagering group demonstrated higher proportions of advantageous wagers across blocks. The main effect of block was significant, $F(9, 522) = 16.88, MSE = 3.65, p < .001, \eta_g^2 = 0.18$. The difference between the two groups in awareness was further supported by a significant group × block interaction, $F(9, 522) = 3.36, MSE = 3.65, p < .001, \eta_g^2 = 0.04$, with reliable differences in blocks 4, 5, 6, and 7 (simple effects comparisons, $p < .05$), reflecting the later onset of awareness in the simple wagering group. These results demonstrate that awareness lagged behind deck selections in the simple wagering group which is in accordance with the dissociation between the two measures observed by Persaud et al. (2007). In addition it appears that asymmetric weights in the pay-off matrix of the modified wagering group helped participants to perform advantageously earlier in the task.

Signal Detection Theory Analysis

We also investigated the confidence-accuracy relationship using Type 2 SDT (Figure 3.2). A 2 (group [simple wagering, modified wagering]; between) × 10 (block: 10 trials each; within) mixed ANOVA was computed. The analysis revealed a significant main effect of group, $F(1, 58) = 4.69, MSE = 2.43, p = .034, \eta_g^2 = 0.02$, indicating that the modified pay-off matrix was more sensitive in assessing participants' task knowledge (simple wagering: $M = 0.20, SEM = 0.06$; modified wagering: $M = 0.47, SEM = 0.06$). Also, there was a significant effect of block, $F(6.81, 394.91) = 12.52, MSE = 0.79, p < .001, \eta_g^2 = 0.14$. The interaction between block and group was significant,

$F(6.81, 394.91) = 2.77, MSE = 0.79, p = .009, \eta_g^2 = 0.03$. Simple effects analysis revealed significant differences between the two groups in blocks 4 and 5 (block 4: $F(1, 58) = 10.51, MSE = 0.74, p = .002$; block 5:$F(1, 58) = 9.13, MSE = 1.33, p = .004$), a pattern of results which resembles the differences found in advantageous wagering between the two groups.

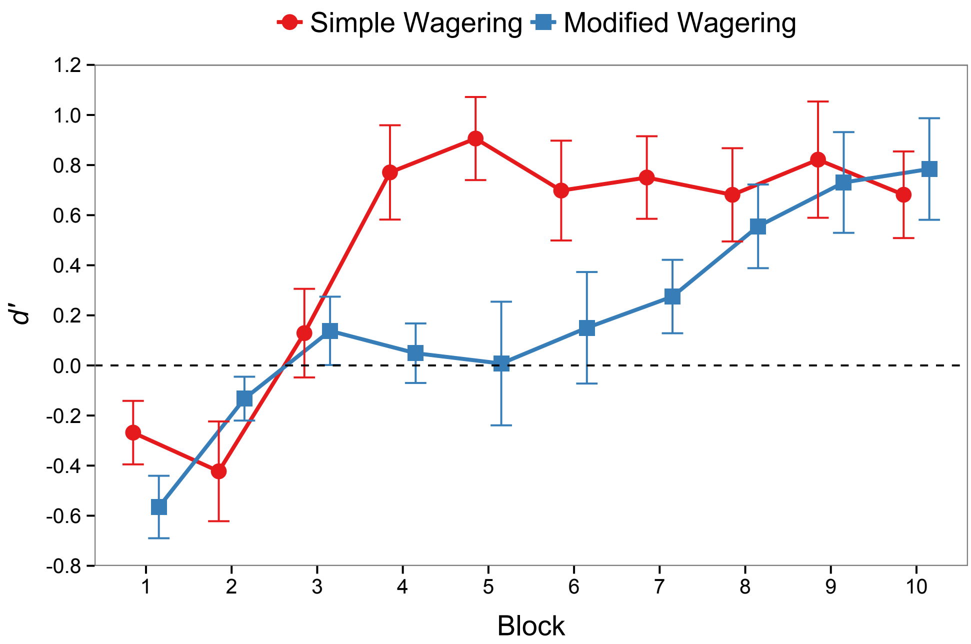

Figure 3.2: Participants' metacognitive sensitivity (d') as a function of wagering type and block of trials in Experiment 3. Error bars represent ± 1 SEM.

We also investigated the mean bias (ln β) in the two groups; in terms of loss aversion, we can ask whether the type of wagering matrix caused participants to develop a liberal or a conservative strategy about the wagers they placed. A 2×10 (group [simple wagering, modified wagering] $\times$ block) mixed ANOVA revealed that neither the group $\times$ block interaction, $F(9, 522) = 1.05, MSE = 0.22, p = .39$, nor the main effect of group, $F(1, 58) = 1.89, MSE = 0.39, p = .17$, reached significance indicating that, in general, the different pay-off matrices did not affect participants' wagering strategy. However, there was a significant effect of block, $F(9, 522) = 2.53, MSE = 0.22, p = .008, \eta_g^2 = 0.04$.

3.1.3 Discussion

Whereas Experiments 1 and 2 revealed no lag between deck selection and awareness - the latter measured by wagering - the present experiment did reveal such a lag in the simple wagering group, of approximately 2 blocks of trials. In this group, advantageous deck selections became reliable at block 5 whereas wagering only became significantly better than chance in block 7. Presumably one of the minor procedural changes between Experiments 1 and 2 on the one hand and Experiment 3 on the other led to the difference in findings.

The reward and punishment schedule used in Experiment 3 was the same as in the original IGT. In fact, the pattern of deck selections was slightly different compared to Experiment 1 (compare Table 2.1 and Figure 3.3). Deck B was selected more often which is in accordance with previous studies that used the original IGT and evaluated the perceived “badness” of deck B (Lin et al., 2007). Examination of choice behaviour in Experiment 1 (Table 2.1) and Persaud et al.’s (2007) study reveals a different pattern, which is not present in studies with the original IGT payoff schedule.

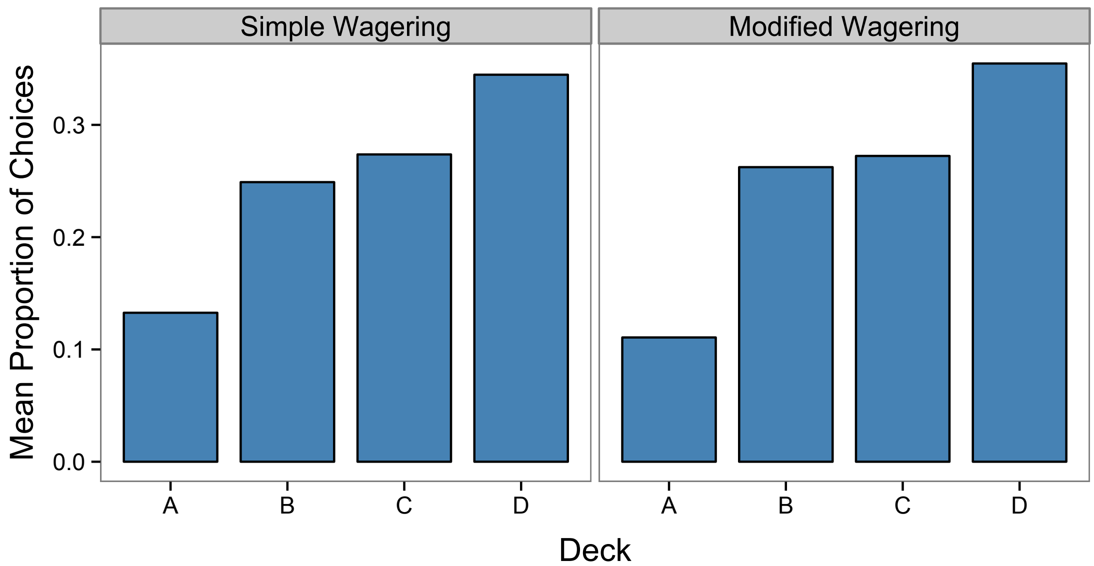

Figure 3.3: Mean proportion of choices from bad decks (A and B) and good decks (C and D) in Experiment 3.

Although the lag observed in the simple wagering group (which replicates what Persaud et al., 2007, found) might be taken as evidence that advantageous deck selection is driven initially by unconscious influences, the results from the modified wagering group suggest caution in drawing such a conclusion, because a relatively small change in the payoff matrix brought wagering back into line with deck selections (and led participants to select from the good decks slightly earlier than those exposed to the original matrix). Why might this have happened? One hypothesis is that it arises because the original payoff matrix discourages participants from thinking carefully about the wagers they place, especially before they have learned which are the best decks. There is a possibility that Persaud et al.'s matrix led participants to believe that, prior to learning, wagering had no overall effect on their winnings. As noted previously, it is indeed the case that with a symmetric matrix and random deck selection, it makes no difference how the participant wagers. Participants may therefore have stopped thinking carefully about their wagers. As the optimal weakly dominant strategy using the original matrix is always to wager high (Clifford et al., 2008), this means that the payoffs are independent of the wagers, and thus participants may have believed that their wagers were irrelevant. When they started to learn about the quality of each deck and discovered that their wagers might be relevant to the encountered payoffs, it may then have taken them longer to implement this new knowledge into their wagering strategy, leading to an apparent late onset of awareness. In contrast, the asymmetric payoffs of the modified matrix encourage participants to believe that it matters whether they wager high or low, even before they start to choose the good decks. In other words, the original pay-off matrix did not guide participants to express their knowledge as their wagering choices were random and not consistent with their deck selections.

3.2 Experiments 4A and 4B

Experiment 3 showed that the exact form of the pay-off matrix can affect participants' wagering strategy, with the Clifford et al. (2008) payoff matrix bringing forward by several blocks the point at which above-chance awareness was located. The sensitivity of wagering to small procedural changes undermines its reliability as a measure of awareness. Yet the results of Experiment 3 might nevertheless encourage the view that wagering under the modified matrix is an accurate measure (and the results of Experiment 1, in which wagering again developed early, might be interpreted in the same way). Even though wagering tracks choice under the modified matrix, does this mean that it is a reliable and sensitive measure of awareness? In Experiment 4A we address this question by measuring awareness both with Clifford et al.'s payoff matrix and simultaneously with Maia and McClelland's (2004) quantitative questions in a probabilistic alternative version of the IGT. Experiment 4B is a replication of Experiment 3's modified wagering condition with the inclusion of Maia and McClelland's questionnaire.

3.2.1 Experiment 4A

3.2.1.1 Method

Participants

Twenty-one volunteers participated (13 females, age $M = 23.45, SD = 3.56$), all of whom were recruited via the departmental subject pool. They were paid £2 for their participation and an additional amount between £0 and £3, depending on their performance in the task.

Task

A variation of the original IGT was employed in which the allocation of wins and losses on each trial was sampled at random from the overall distribution (for a similar task see Schönberg, Daw, Joel, & O'Doherty, 2007). This modification removes many of the complications that arise from using the typical IGT structure in which the disadvantageous decks are initially good (because losses do not occur early in the task), eliminating the predominant preference for the bad decks (see Fellows & Farah, 2005). The payoff structure of each deck was different from the original IGT; the pay-off matrix of Clifford et al. (2008) was used to determine the payoffs received by participants on each trial, in such a way that the amount won or lost was dependent on card selection and wagering.

For example, based on the contingencies of Table 3.1, a payoff of 2 is always associated with a good deck selection and a low wager. Whether this amount was a win or loss was defined by the distribution of outcomes associated with each deck. Specifically, for decks A and B, the probability of a loss was .75 and .60 respectively, whereas for decks C and D, the probability of a win was .75 and .60, respectively, resulting in different overall expected payoffs for each deck. In contrast to the original IGT (where the win on each trial could be coupled with a loss), the outcome on each trial was either a net win or a loss and participants could win or lose points, not real or facsimile money.

The task comprised 100 card selections. Each deck had 60 randomly predefined wins and losses based on the probabilities programmed for that deck. After each card selection, participants could place a wager, either High or Low, on their card selection. Based on the combination of deck selection and wagering, participants were presented with a single amount, either a win or a loss. Along with wagering, participants' conscious knowledge was assessed using a modified version of Maia and McClelland's (2004) questionnaire. The qualitative parts of the questionnaire were omitted and it was administered every 20 trials.

Procedure

The procedure of Experiment 4A was identical to that of previous experiments.

3.2.1.2 Results

Choice and Wagering

Performance exceeded the chance level on block 1 for both measures (Choice: $M = 0.59, t(20) = 2.83, p = .01, d = 0.62$, Wagering: $M = 0.65, t(20) = 3.80, p = .001, d = 0.83$) (see Figure 3.4).

A repeated-measures ANOVA revealed a significant effect of block on good deck selections, $F(9, 180) = 12.40, MSE = 2.32, p < .001, \eta_g^2 = 0.28$. Wagering performance closely followed the optimal decision-making strategy as demonstrated by a main effect of block, $F(4.92, 98.46) = 4.92, MSE = 2.23, p < .001, \eta_g^2 = 0.13$. These findings are

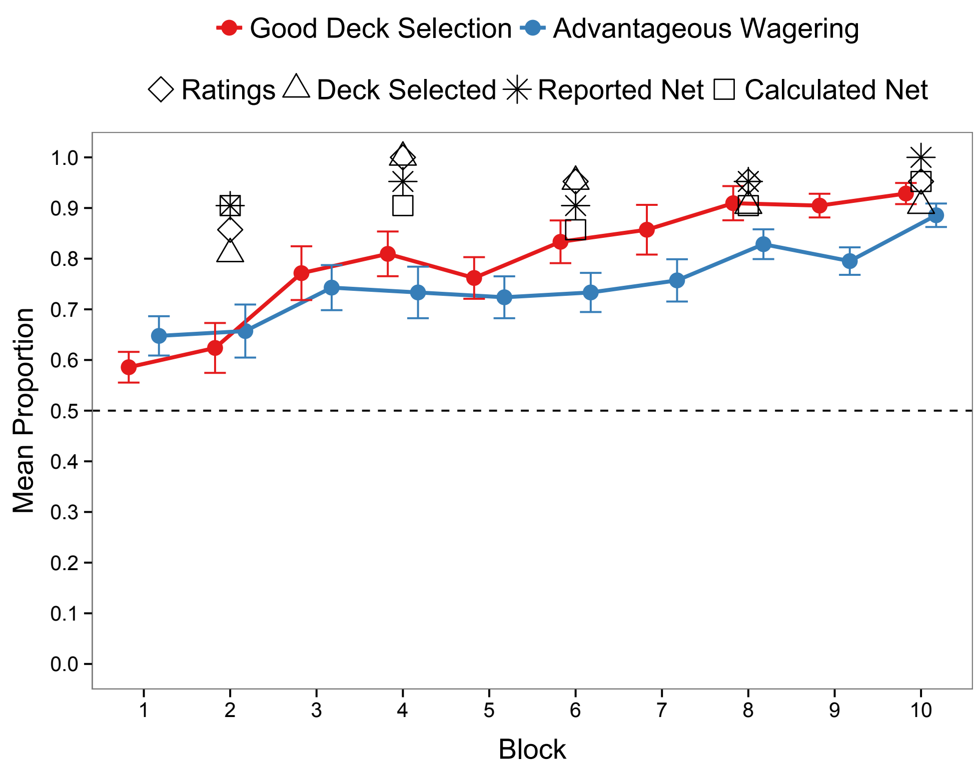

Figure 3.4: Proportion of good deck selections and advantageous wagering across blocks of trials in Experiment 4A (lines). The grey diamond and the triangle markers represent the proportion of participants who gave higher ratings to one of the two best decks and the proportion of participants who selected one of the two best decks as their choice if they were allowed to select only one deck. The star and the square markers represent the reported expected net and the calculated net, respectively.

consistent with the previous results relating to the modified payoff matrix, indicating no dissociation between performance and awareness. In fact, the pattern of both good deck selections and advantageous wagering is similar to the modified wagering condition in Experiment 2, albeit with accelerated learning.

Rapid learning can be explained by the probabilistic allocation of wins and losses on each trial. Fellows and Farah (2005) found that in their shuffled IGT version (the order of the decks was changed so that losses from the bad decks occurred at the start of the task) normal control participants selected more cards from the good decks even in the first 20 trials and they kept on choosing the good decks throughout the task. Our probabilistic version of the payoff schedule removes the reversal learning component (that is, to learn that the decks which yield higher rewards are disadvantageous in the long run) of the IGT which can be slow and delay learning of the optimal decisions.

Since each deck had different overall expected payoffs we investigated whether participants could discriminate not only between good and bad decks but also within each pair of decks (A vs B and C vs D). Participants selected more cards from the good decks in all blocks, $t(20) = 12.02, p < .001$, and this tendency increased from block 1 to block 5. Also, they selected more cards from deck C compared to deck D, $t(20) = 3.97, p < .001, d = 0.87$. No significant difference was observed between selections from decks A and B across blocks, although participants tended to select more cards from deck B.

Signal Detection Theory Analysis

Participants' confidence-accuracy levels were examined using Type 2 SDT. As shown in Figure 3.5 (circle markers), meta-cognitive sensitivity as measured by d' was significantly above chance ($d' = 0$) even in the first 10 trials indicating that even a few deck selections sufficed to acquire awareness of the advantageous strategy. In other words, participants were able to discriminate between good and bad decks and make an appropriate wager. Figure 3.5 shows a tendency for d' to increase across blocks, although the main effect was not significant, $F(9, 180) = 1.57, MSE = 0.43, p = .13$.

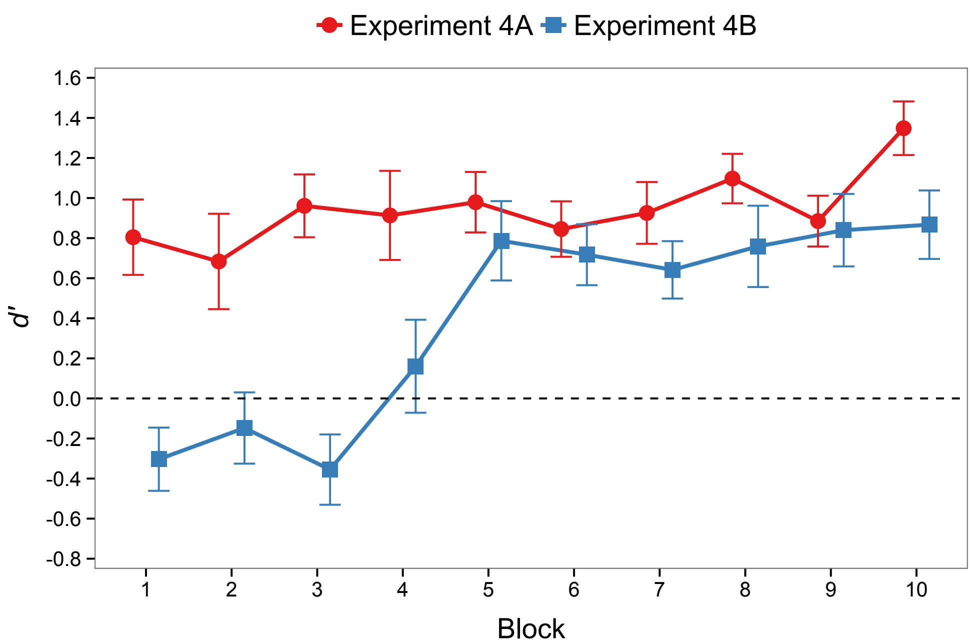

Figure 3.5: Participants' metacognitive sensitivity (d') in Experiment 4.

Questionnaire

Participants' knowledge regarding the advantageous strategy was further supported by the various measures included in the questionnaire. Figure 3.4 shows that they exhibited substantial knowledge about the quality of each deck, even in the first assessment of awareness (trial 20). In fact, the pattern is similar to that observed in Experiment 1. Importantly, the mean ratings for each deck give further support to the pattern of deck selections. Not only are the good decks selected more often than the bad decks, but also participants' ratings align with the expected value of each deck. Table 3.2 shows that deck C is evaluated more positively than deck D even though both decks are advantageous. In other words, knowledge about the quality of each deck led participants to select more cards from deck C. Similarly, deck A (which has a higher probability of loss compared to deck B) has the lowest mean rating.

However, the two measures of awareness are not directly comparable based on the information shown in Figure 3.4. We applied the same procedure as in Experiment 1 to test whether the proportion of participants who preferred a good deck in the deck-selected measure (.81) is significantly different from the proportion classified as aware of the optimal strategy based on wagering (.76) on trials 16-25. The McNemar test for dependent proportions was not significant, $\chi^2(1) = 0.2, p = .65$.

Table 3.2: Mean ratings and proportion of selections for each deck in Experiments 4A and 4B

	Experiment 4A		Experiment 4B	
Deck	Mean Rating (SD)	Choice	Mean Rating (SD)	Choice
A	-5.32 (3.61)	0.09	-2.46 (4.39)	0.15
B	-2.94 (4.40)	0.11	-1.47 (5.51)	0.25
C	5.60 (2.98)	0.42	1.58 (3.39)	0.25
D	4.31 (3.29)	0.38	3.52 (4.06)	0.35

Another way of examining the two measures is to look at participants' deck selection and wagering in the trials following the administration of the questionnaire (trials 21, 41, 61, 81; we also include trial 100 immediately prior to the final administration of the questionnaire). Specifically, we are interested in the verbal reports and wagers of those participants who behave advantageously (i.e., select good decks) in these trials. Figure 3.6 shows that the majority of participants demonstrate knowledge of the advantageous strategy in all the questionnaire items. However, wagering underestimates the acquired knowledge in all trials following the questionnaire compared to the verbal reports. Thus, it is evident that the detailed and structured questions reflected high levels of awareness compared to wagering.

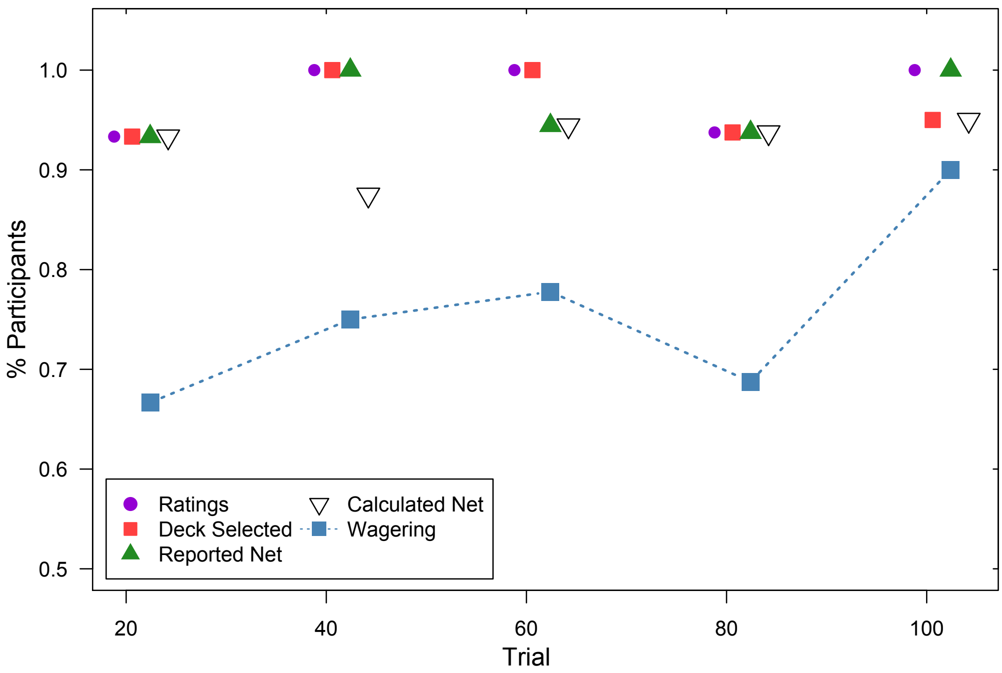

Figure 3.6: Percentage of participants who showed knowledge of the advantageous strategy in the questionnaire items versus in their wagers in Experiment 4A. Wagering indicates the percentage of participants who made an advantageous wager (high on a good deck choice) on the trial immediately following the administration of the questionnaire.

3.2.2 Experiment 4B

3.2.2.1 Method

Partcipants

Nineteen volunteers participated (10 females, age $M = 24.95, SD = 3.15$) from UCL's subject pool. As in Experiment 4A, they received £2 for participation and an additional fee up to £3 dependent on their performance in the task.

Task

The decision-making paradigm was identical to the modified wagering condition of Experiment 3, that is, the payoff schedule was the same as in the original IGT and wagering was expressed as a binary choice ("High" and "Low"). The extra component of this experiment was the questionnaire of Maia and McClelland (2004) which was administered every 20 trials.

3.2.2.2 Results

Choice and Wagering

The mean probability of selecting a good deck and making an advantageous wager exceeded chance on block 5 for both measures (Choice: $M = 0.67, t(18) = 3.21, p = .005, d = 0.70$, Wagering: $M = 0.65, t(18) = 3.51, p = .003, d = 0.77$) (see Figure 3.7). Compared to the onset of learning and awareness in Experiment 2, there seems to be a lag of one block. Despite the fact that both measures are numerically above chance on block 4 (M Choice = 0.57, M Wagering = 0.53), neither is significant.

A repeated-measures ANOVA showed significant main effects of block on choice, $F(9, 162) = 12.72, p < .001, \eta_g^2 = 0.32$, and wagering, $F(9, 162) = 8.81, p < .001, \eta_g^2 = 0.28$. These results agree with our previous experiments where we used the Clifford et al. (2008) matrix (Experiment 3, modified wagering group; Experiment 4A). Learning of the good decks and awareness progressed in the same manner and no dissociation was observed.

Signal Detection Theory Analysis

Figure 3.5 (square markers) shows that mean d' exceeded chance on block 5, $M = 0.78, t(18) = 3.97, p < .001$, in line with the advantageous wagering performance. There was a significant effect of block, $F(9, 162) = 8.47, MSE = 0.57, p < .001, \eta_g^2 = 0.28$, as metacognitive discrimination gradually increased over time. On the other hand, par-

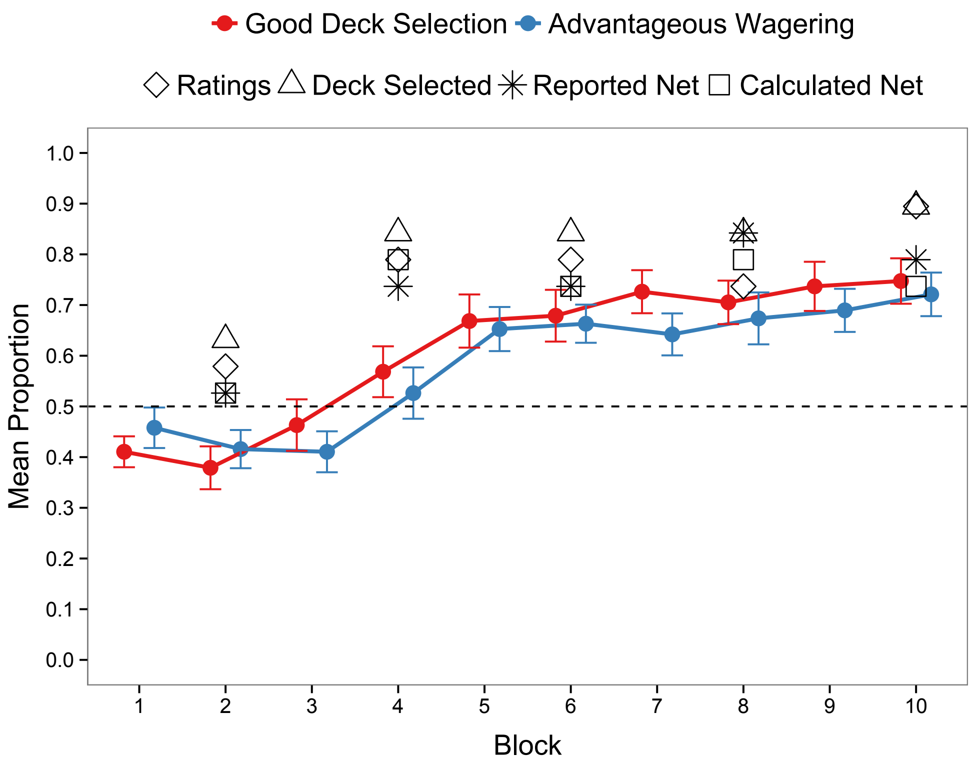

Figure 3.7: Proportion of good deck selections and advantageous wagering across blocks of trials in Experiment 4B (lines). The grey diamond and the triangle markers represent the proportion of participants who gave higher rating to one of the two best decks and the proportion of participants who selected one of the two best decks as their choice if they were allowed to select only one deck, respectively. The star and the square markers represent the reported expected net and the calculated net, respectively.

ticipants seemed not to change their wagering strategy across blocks as indicated by a non-significant effect of block on mean ln β, $F(9, 162) = 0.88, MSE = 0.25, p = .54$.

Questionnaire

The proportion of participants whose responses favoured the good decks is illustrated in Figure 3.7. The majority of participants showed a preference for the good decks in the verbal questions except at the first question period where the proportion was lower but still above chance.

The mean ratings for each deck (see Table 3.2) converge with the profile of deck selections. Deck D has the highest mean rating which explains why this deck is selected more often than the other decks. Even though decks C and D share the same overall expected values, the small probability of loss on deck D affects the perceived goodness of this deck. The same principle applies to deck B too; despite its overall negative appraisal, it is selected as often as deck C. Also, the high probability of loss on deck A in conjunction with its negative expected value led participants to negatively evaluate and to avoid selecting cards from this deck.

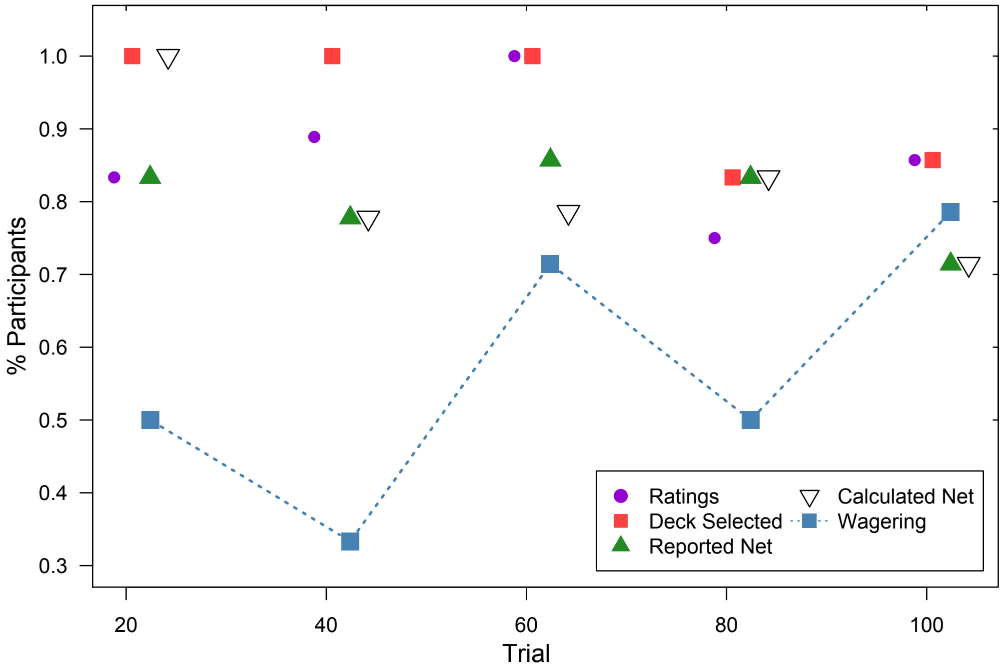

Figure 3.8: Percentage of participants who showed knowledge of the advantageous strategy in the questionnaire items versus in their wagers in Experiment 4B. Wagering indicates the percentage of participants who made an advantageous wager (high on a good deck choice) on the trial immediately following the administration of the questionnaire.

In order to compare how sensitive the two methods are in assessing conscious knowledge we again examined the proportion of participants who behaved advantageously in the trials following the administration of the questionnaire (we again include trial 100 which immediately preceded the final set of questions). Figure 3.8 demonstrates that in all question periods the proportion of participants who translated their knowledge into a high wager is less than the proportion who favoured the good decks in their verbal reports. This pattern suggests that wagering underestimated participants' acquired knowledge and that the more elaborated questions detected higher levels of awareness. Despite the fact that wagering closely tracks deck selections, it is not therefore an exhaustive and sensitive method to measure awareness. This conclusion is supported by a significant difference between the proportion of participants who opted for one of the good decks in the deck-selected measure (.63) and the proportion classified as aware based on wagering (.32) in the first administration of the questionnaire, $\chi^2(1) = 6, p = .014$.

3.2.3 Discussion

The key point of Experiments 4A and 4B is that even though wagering closely tracks deck selection and learning, it underestimates what participants have learned about the task and deck contingencies. This also applies to the results of Experiment 3 where we found that small procedural changes can affect the extent to which wagering tracks deck selection.

Finally, the analysis based on the trials following the administration of the questionnaire suggests that acquired knowledge is not automatically translated into an appropriate wager after a deck selection (Figures 3.6 and 3.8). Why is this? A possible reason is loss aversion. The prospect of losing more money/points even if knowledge is above guessing levels can be aversive.

3.3 Experiment 5

Depending on the setup of the pay-off matrix participants may employ different response criteria to place high or low wagers which makes the detection of the acquired knowledge very difficult. This leads to the possibility that the expression of awareness via wagering may be constrained by factors other than knowledge itself. For instance, many studies have shown that loss aversion affects awareness assessment as indexed by wagering (e.g., Dienes & Seth, 2010; Fleming & Dolan, 2010; Wang et al., 2012; Wierzchoń et al., 2012).

Schurger and Sher (2008) proposed that the design of a pay-off matrix should take into account the tendency of participants to evaluate losses worse than equivalent wins. Unlike Clifford et al.'s (2008) matrix which encourages low wagers when certainty is low, "subjects seem to need precisely the opposite sort of encouragement" (Schurger & Sher, 2008, p. 209). Table 3.1 shows the matrix devised by Schurger and Sher as a means to counter loss aversion. Specifically, looking at Table 3.1, when discrimination between good and bad decks is at chance it is more advantageous to wager high due to a negative expected payoff from wagering low $[(+1 - 2)/2 = -1/2]$ compared to a neutral payoff from wagering high $[(+10 - 10)/2 = 0]$. Following this, it can be shown that a rational participant would switch to high wagers even when her discrimination is below chance (50%), at 8/17 or 47%. Specifically, the differential loss of wagering on a bad decision is 8 $(= 10 - 2)$ divided by the sum of the differential loss and the differential gain of wagering on a good decision $(10 - 1 = 9)$.

Despite the fact that this matrix discourages participants from wagering low, its weights regarding high wagers are two times bigger compared to the matrix of Clifford et al. (2008). On the one hand, the larger loss following a low wager after an incorrect decision discourages participants from wagering low, thus overcoming the problem of loss aversion. On the other hand, the bigger weights for high wagers could discourage participants from wagering high, even when knowledge about the quality of the decks exists. Thus, the utilisation of this matrix might reveal that the remedy proposed to counter loss aversion cannot be achieved due to the increased weights associated with high wagering.

3.3.1 Method

Participants

We tested a total of 30 participants (24 females, age $M = 25.08, SD = 4.02$), recruited from UCL's psychology subject pool. Participants were rewarded between £1 and £5, proportional to their performance on the task.

Task

The payoffs of each deck were different to the original IGT, but their overall expected payoffs reflected the ratio of losses to wins of the original task (Table 3.3). There were four decks of cards each having 100 associated wins and losses, one for each trial. A randomly drawn value (win − loss) was then computed for each trial, which constituted the payoff on that deck for that trial. Decks A and B were bad decks, with an overall net outcome of −500 points (a net loss of −5 per card). These decks had high rewards (from 15 − 25 points), but large losses (from 25 − 75). Decks C and D were good decks, with an overall net outcome of +500 points (a net win of +5 per card). They had lower rewards (from 5 − 15), but their losses were smaller too. Decks A, B, and C had a loss on 50% of trials, whereas Deck D had a loss on 10% of trials. The characteristics of each deck matched the original IGT, including the probabilities and relative magnitudes of losses, except for deck B. The losses on deck B were distributed over 50 trials (as against originally 10 trials only). We did this to avoid a major loss if participants were unlucky enough to encounter the deck B loss with a high wager. The post-decision wagers comprised multipliers, with the payoff schedule as proposed by Schurger and Sher (2008). Accordingly, a given IGT trial payoff was multiplied by a factor of 2 when wagering low on decks A and B, and by 1 when wagering low on decks C and D. When wagering high, all deck payoffs were multiplied by a factor of 10.

Table 3.3: Payoff schedule in Experiment 5. The numbers in parentheses show the probability of the outcome.

Deck	A+	A−	B+	B−	C+	C−	D+	D−
	+15 (.33)	-25 (.17)	+15 (.33)	-50 (.50)	+5 (.33)	-5 (.17)	+5 (.33)	-50 (.10)
	+20 (.34)	-50 (.16)	+20 (.34)	0 (.50)	+10 (.34)	-10 (.16)	+10 (.34)	0 (.90)
	+25 (.33)	-75 (.17)	+25 (.33)		+15 (.33)	-15 (.17)	+15 (.33)	
		0 (.50)				0 (.50)		
Net	-5		-5		5		5	

3.3.2 Results

Choice and Wagering

Performance exceeded chance on block 1 for both measures (Choice: $t(29) = 2.39, p = .023$, Wagering: $t(29) = 2.52, p = .018$) (Figure 3.9). This result indicates that participants' optimal decision-making and learning about their selections occurred very early in the task, a pattern that is not observed in previous studies which have utilised a payoff schedule similar or identical to the original IGT.

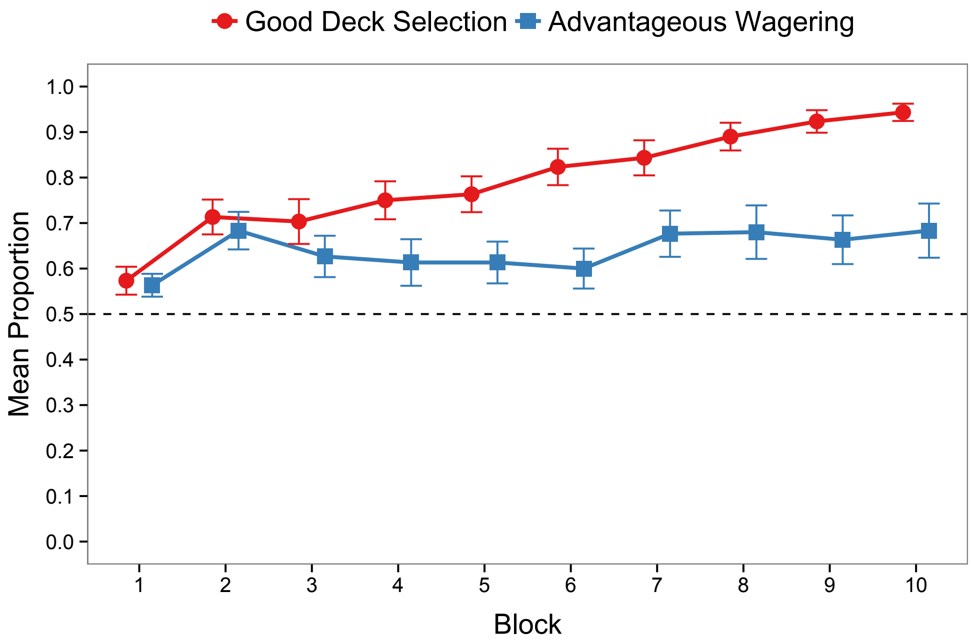

Figure 3.9: Proportion of good deck selections and advantageous wagering across blocks of trials in Experiment 5. Error bars represent ± 1 SEM.

Two separate within-subjects ANOVAs were conducted to investigate the progression of good deck selections and advantageous wagering across blocks. There was a significant main effect of block on the proportion of good deck selections, $F(5.54, 160.72) = 16.27, MSE = 2.41, p < .001, \eta_g^2 = 0.24$ (Figure 3.9). However, the same trend was not observed on the proportion of advantageous wagers as the main effect of block was not significant, $F(5.14, 149.04) = 2.06, MSE = 3.33, p = .10$. Even though wagering was above chance from block 1, it never exceeded 0.7. In a situation where high wagers have much greater stakes than low wagers, participants may wager conservatively throughout the task, independent of learning and awareness, due to an aversion to big losses. Additionally, advantageous wagering was above chance in all blocks of trials.

Figure 3.10 shows the proportion of deck selections throughout the task. Deck B was not selected as often as in previous studies using the IGT, a fact which reflects the change of the loss probability. When the occurrence of losses is more frequent (.5), the prominent deck B phenomenon is not observed. On the other hand, deck D (loss probability .1) was selected more often than deck C (loss probability .5) even though both decks have the same expected value.

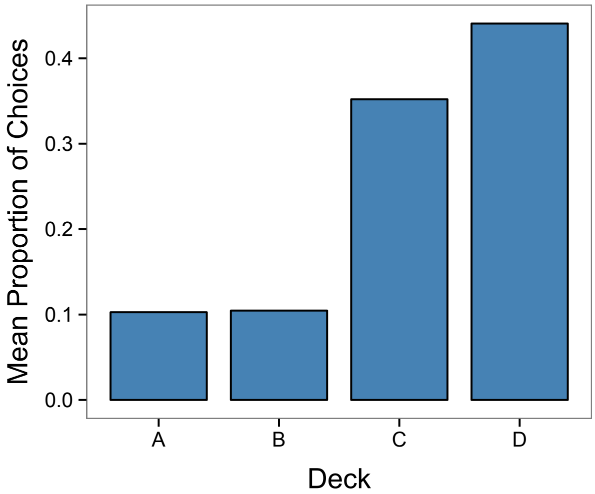

Figure 3.10: Mean proportion of choices from bad decks (A and B) and good decks (C and D) in Experiment 5.

3.3.3 Discussion

This experiment confirms the hypothesis that loss aversion modulates wagering strategies by making participants more sensitive to losses. While the payoff matrix we used encourages high wagering under uncertainty, the probabilistic IGT variant we employed was found to be easier to learn than the classic IGT and thus participants were able to grasp the payoff schedule in the first 10 trials, indicating that they did not go through a phase of exploration or uncertainty. Having learned the probabilistic structure of wins and losses early in the task, it might be expected that wagering would simultaneously follow the optimal choices. This was the case in the first 2 blocks where participants had learned about the good strategy and made high wagers. Yet a random loss which may occur from the selection of a good deck with a high wager ($\times 10$) would result in a large amount of points being deducted from the total sum. Hence, a "lose-less" strategy seems to overtake the tendency to maximise winnings, and in this particular case leads to suboptimal wagering. In other words, loss aversion constrains participants from wagering high on their good deck selections. This is indicative of a bias regulating wagering strategies and not lack of awareness. While good deck selections gradually increased to reach a maximum by the end of the task, it would be unreasonable to argue that this was the effect of an unconscious mechanism.

The present experiment also highlights the inadequacy of post-decision wagering to measure awareness objectively and directly. Small changes in the payoff matrix can dramatically change the expression of awareness as cognitive or response biases overtly influence participants' wagering strategies.

3.4 General Discussion

The experiments described in the current chapter examined two main response biases, dominance and loss aversion, which arise from the design of the pay-off matrix. In both cases, there was a direct effect of the design of the pay-off matrix on the wagering strategies that participants employed. In Experiment 3, despite the fact that there was a difference in the onset of learning and awareness in the simple wagering condition, no such difference was observed in the modified wagering group. Thus we were able conceptually to replicate Persaud et al.'s (2007) finding of a lag between choice and wagering, but a simple change in the weights of the pay-off matrix was sufficient to make wagering a more sensitive method.

The purpose of Experiment 4 was to measure wagering concurrently with explicit questioning. Experiments 1 and 3 (modified wagering condition) showed that wagering can closely track learning, but is that alone an adequate indicator of a robust method for measuring awareness? We employed the questionnaire of Maia and McClelland (2004) in order to examine how well wagering performs in comparison to another method of awareness. The results showed that even though wagering followed deck selections, it is not a sensitive index of awareness as it underestimates the knowledge that participants possess. We compared the proportions of participants classified as consciously aware by the two measures. No significant differences were observed in two of the experiments because of the small sample sizes, although more participants were identified as aware according to the questionnaire.

One possible criticism of the quantitative questions employed here is that they might have a reflexive effect on the very property they are attempting to measure, namely awareness. However, in Experiment 4A participants' wagering performance was better than chance even before the first administration of the questionnaire, indicating that the explicit nature of the questions did not make participants more aware of the decks' payoffs.

In Experiment 5, we tried to control for the effects of loss aversion on wagering strategies, mindful of the possibility that the high values in the wagering matrix could make participants reluctant to place high wagers. The matrix proposed by Schurger and Sher (2008) attempts to eliminate loss aversion in situations of uncertainty, that is when knowledge about a response option is weak. Participants were able to discriminate between the decks after a few trials. Although wagering performance was better than chance from the beginning of the task, it did not lead participants to maximise their earnings. One explanation lies in the design of the task: with random losses occurring even on selections from the good decks and wagers treated as multipliers of the actual payoffs, the prospect of losing a significant amount could inhibit the placement of high wagers.

4 Confidence Ratings and Wagering

In the previous experiments we examined two problematic aspects of the use of post-decision wagering as a measure of awareness and compared it with the quantitative questions of Maia and McClelland (2004). One possible problem with this comparison is that it is not a direct one as the questionnaire was administered every 20 trials whereas wagers had to be placed on every trial. In addition, wagering was used in a binary manner as participants could place either a high or a low wager on every trial. The first aim of this chapter was to compare wagering with confidence ratings, a common subjective measure of awareness which can be used on a trial-by-trial basis. The comparison between confidence ratings and wagering is a natural one as both measures assess the degree of certainty about one's judgments as opposed to other subjective measures such as the perceptual awareness scale (PAS; Ramsøy & Overgaard, 2004) and the continuous scale (CS; Sergent & Dehaene, 2004) which are direct and introspective measures of "pure" awareness (Wierzchoń, Paulewicz, Asanowicz, Timmermans, & Cleeremans, 2014). Also, wagering and confidence ratings were used in a gradual manner in the experiments of this chapter as including more categories could potentially allow for a finer investigation of awareness and acquired knowledge in the task.

The second aim of the present chapter was to examine the development of learning and awareness of the advantageous strategy in a dynamic deck-shifting version of the IGT (Experiment 8). Previous research on deck-shifting variants suggested that participants were able to adapt to the new environment and select advantageously after an exploration phase (see Dymond, Cella, Cooper, & Turnbull, 2010; Turnbull, Evans, Kemish, Park, & Bowman, 2006). While these studies indicated that people are good at detecting change and adapt their behaviour accordingly, it is not clear whether participants had conscious knowledge to support their decisions. It could be that people employ simple strategies or intuitive heuristics instead of having explicit knowledge of the decks' contingencies in order to deal with the increased uncertainty of the environment. In this experiment, conscious content was assessed using confidence ratings and two of the quantitative questions of Maia and McClelland (2004).

4.1 Experiment 6

The purpose of this experiment was to compare post-decision wagering with confidence ratings, the simplest and most commonly used subjective measure of awareness (for some examples see Dienes et al., 1995; Fleming, Weil, Nagy, Dolan, & Rees, 2010; Szczepanowski et al., 2013; Tunney & Shanks, 2003). Confidence ratings are metacognitive reports about having performed a judgment or discrimination accurately (e.g., perception of a subliminal visual stimulus) or having selected the best from a set of alternatives (e.g., selection of a good deck in the IGT). Confidence ratings can be expressed in a binary way such as "not confident" and "very confident" (different labels have been employed such as "guessing" and "knowing") or on a continuous Likert-like scale.

In this experiment we used a 2-point confidence scale in order to make a direct comparison with binary wagering, and also a 4-point scale to gain deeper insights into the confidence-performance relationship.

4.1.1 Method

Participants

There were 118 participants in the experiment (97 females, age $M = 18.73, SD = 0.90$), all of whom were psychology undergraduate students at University College London and took part in fulfilment of a course requirement. The 6 best performers on the task were awarded £15 each.

Design

The experiment consisted of three different conditions: binary wagering ($N = 40$), binary confidence ratings ($N = 40$) and 4-point confidence ratings ($N = 38$). Participants were randomly assigned to one of the three conditions.

Task

The original IGT payoff schedule was used across conditions. After each card selection, participants were asked to indicate their awareness of the deck payoffs using wagering or confidence ratings. In the binary wagering condition, participants had to place a wager, High (£2) or Low (£1), which multiplied the payoffs associated with each deck and trial (this condition was identical to the simple wagering condition of Experiment 3). In the binary confidence condition, participants were asked to express their confidence in having selected a good deck using the descriptions 1 = "I am not confident" and 2 = "I am very confident". The descriptions for the 4-point confidence scale were 1 = "I am guessing", 2 = "I am not confident", 3 = "I am quite confident", and 4 = "I am very confident".

Procedure

The procedure was identical to that of previous experiments, with the exception that a different set of instructions was presented for the confidence ratings measure.

4.1.2 Results

Choice and awareness

Evidence of conscious knowledge regarding the optimal strategy in the binary wagering condition was obtained using advantageous wagering (a high wager after a good deck and a low wager after a bad deck). The same principle was applied to the confidence ratings conditions so that the combinations good deck/high confidence and bad deck/low confidence were taken to indicate conscious knowledge. In this stage of the analysis the 4-point scale was dichotomised with confidence levels 1 and 2 collapsed to signify *low confidence* and 3 and 4 collapsed to give *high confidence*. We then identified the onset of choice and awareness as the first block at which performance was significantly above chance (0.5) for each of the three conditions (see Figure 4.1). With this method, good deck selections exceeded the chance level on block 5 for the binary confidence ratings, $M = 0.60; t(39) = 3.00, p = .004, d = 0.48$, and wagering groups, $M = 0.60; t(39) = 2.82, p = .007, d = 0.45$, and block 6 for the 4-point confidence group, $M = 0.62; t(37) = 3.73, p < .001, d = 0.61$. These results indicate slightly later (about 1 block) deck discrimination than in previous experiments.

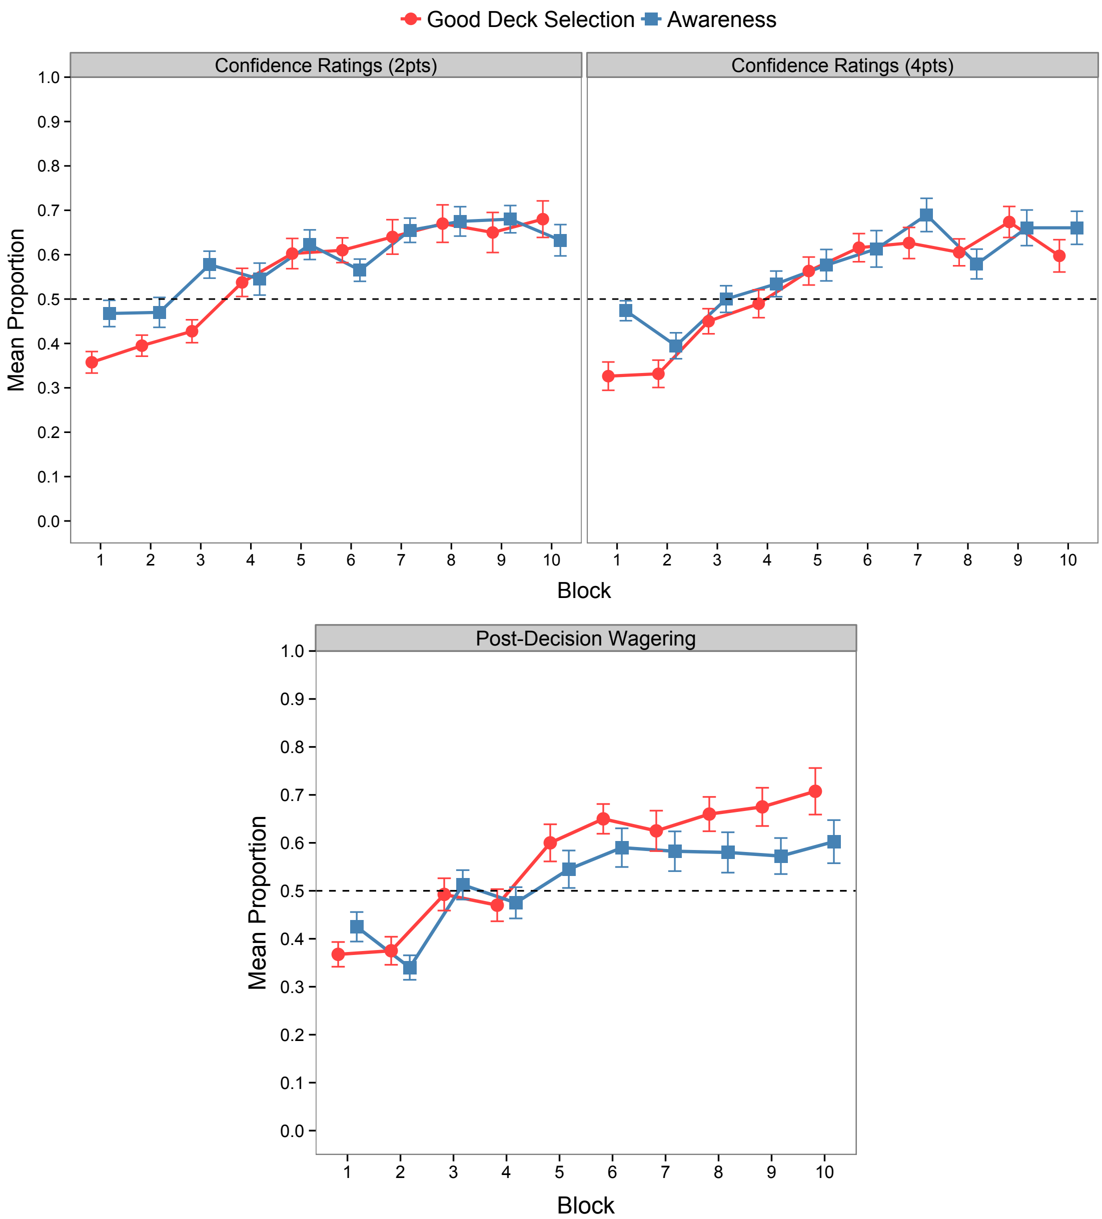

Figure 4.1: Proportion of good deck selections and awareness across blocks of trials for each group in Experiment 6.

Regarding conscious knowledge of the deck values, confidence ratings significantly exceeded chance at the same time as or earlier than choice in the confidence rating groups, namely at block 5 in both cases (2pts scale: $M = 0.62; t(39) = 3.89, p < .001, d = 0.62$, 4pts scale: $M = 0.58; t(37) = 2.28, p = .029, d = 0.37$), whereas there was a delay of (at least) one block in the onset of conscious knowledge as indexed by wagering (block 6), $M = 0.59; t(39) = 2.23, p = .031, d = 0.35$. This last result replicates what was observed in the simple wagering group of Experiment 3, namely a delay in the onset of awareness. In both groups, deck discrimination became significant at block 5, while advantageous wagering did not become significant until block 6 (Experiment 6, wagering) or block 7 (Experiment 3, simple wagering). In fact the data from the two groups are more similar still, as in the present wagering group advantageous wagering was not significantly greater than chance in blocks 7, 8, and 9.

A 3 (group) $\times$ 10 (block) mixed ANOVA on the mean proportion of good deck selections showed no main effect of group, $F(2, 115) = 0.88, MSE = 15.05, p = .42, \eta_g^2 = 0.004$, but there was a significant effect of block $F(7.38, 848.97) = 47.86, MSE = 3.65, p < .001, \eta_g^2 = 0.22$. Also, no significant interaction between group and block (main effect) was observed, $F(14.76, 848.97) = 0.73, MSE = 3.65, p = .75, \eta_g^2 = 0.008$ (Greenhouse-Geisser correction). In general, these results suggest that the acquisition of the advantageous strategy was not substantially affected by the different subjective measures of awareness and participants were able to learn to discriminate between the decks based on their overall expected values.

The second important conclusion from this analysis refers to the pattern of overall deck selections; despite the fact that participants learned to discriminate between the decks, this learning effect was rather weak. Figure 4.2 illustrates that participants did not take into account the infrequent but rather large losses in deck B as this deck was selected as often as deck D, showing a strong loss-frequency effect. Interestingly, deck B was the overall deck of choice in the 4-points confidence group.

The same type of analysis was applied to mean performance on the awareness measures. The analysis revealed a significant main effect of condition, $F(2, 115) = 3.96, MSE =$

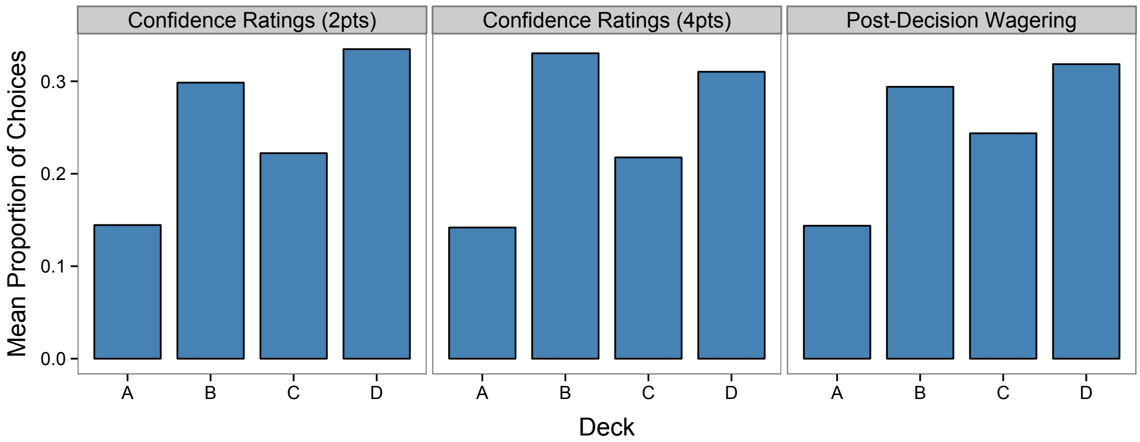

Figure 4.2: Mean proportion of choices from each deck in the three groups in Experiment 6.

$11.68, p = .022, \eta_g^2 = 0.02$, which was mainly driven by a significant difference between the overall means of binary confidence ratings ($M = 0.59$) and wagering ($M = 0.52$) (Tukey's HSD, $p = .016$). No other significant differences between the three measures were observed. The main effect of block was significant, $F(7.63, 877.57) = 21.41, MSE = 3.70, p < .001, \eta_g^2 = 0.12$, indicating that participants' responses in the confidence and wagering measures were consistent with learning of the advantageous strategy. Participants were able to demonstrate conscious knowledge which closely tracked their decisions. Also, the interaction between group and the block main effect did not reach significance, $F(15.26, 877.57) = 1.04, MSE = 3.70, p = .41$.

Signal Detection Theory Analysis

The confidence-accuracy relationship was examined using Type 2 SDT. Figure 4.3 confirms the pattern described above, namely that wagering underestimates awareness compared to confidence ratings. The mean d' exceeded chance in block 5 for both the confidence ratings scales (2pts: $M = 0.45, t(39) = 2.80, p = .007$, 4pts: $M = 0.57, t(37) = 4.17, p < .001$) although it was not significantly above chance for the 2-point scale in block 6 (see Figure 4.3). In contrast, the mean d' for wagering was only marginally above chance in block 6 ($M = 0.34, t(39) = 1.99, p = .05$), and never reliably exceeded chance for the rest of the task. A 3 (group) $\times$ 10 (block) mixed ANOVA on the mean d' confirmed a significant main effect of group, $F(2, 115) = 7.08, MSE = 1.83, p = .001, \eta_g^2 = 0.03$, due to significant differences between wagering and the confidence scales (Wagering $M = 0.02$, 2pts $M = 0.31$, 4pts $M = 0.34$) based on pairwise comparisons between wagering and confidence ratings using Tukey's HSD, $p = .005$ (2pts) and $p = .002$ (4pts). There was also a significant effect of block, $F(18, 1035) = 18.89, MSE = 0.78, p < .001$. The interaction between group and the block main effect, however, did not reach significance, $F(18, 1035) = 1.11, MSE = 0.78, p = .34$.

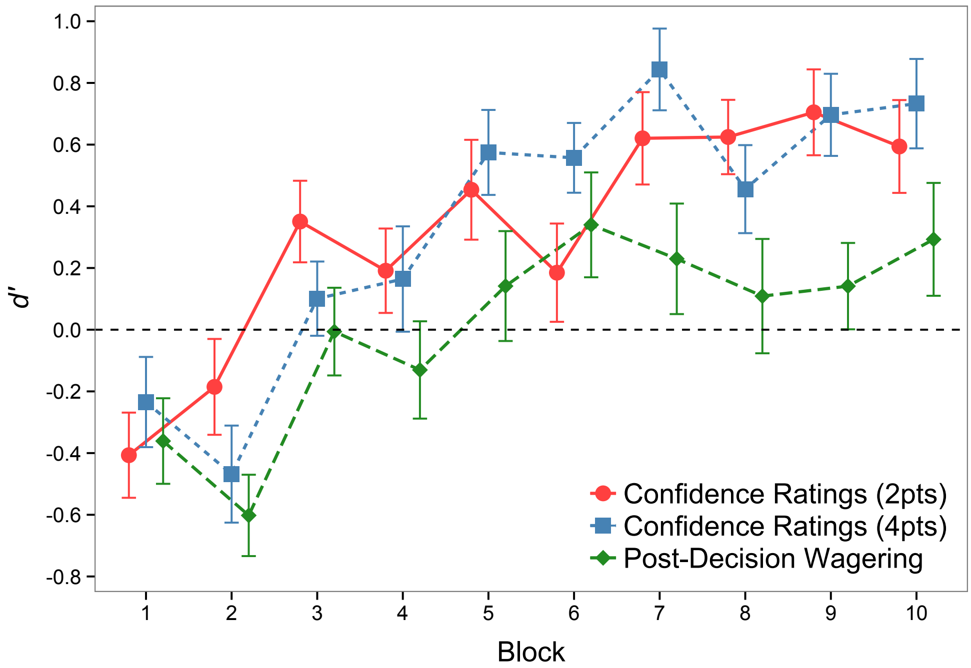

Figure 4.3: Participants' metacognitive sensitivity (d') in Experiment 6.

Analysis of the 4-point confidence ratings

We examined the 4-point confidence ratings in order to provide a more detailed assessment of conscious knowledge in the IGT by employing a nonparametric receiver operating characteristic (ROC) analysis. Two separate ROC curves were constructed, one before the onset of good deck selections (blocks 1-5) and one after (blocks 6-10). Deck selection performance did not significantly change across blocks 6-10, $F(4, 148) = 0.97, p = .43$, allowing for a finer examination of the respective ROC curve. Figure 4.4 shows that the probability of selecting a good deck gradually increases with confidence in blocks 6-10 whereas the straight ROC line for blocks 1-5 is indicative of a poor relationship between accuracy and confidence. The Type 2 sensitivity derived from these curves (A, the area under the ROC curve) indicated above-chance (.50) metacognitive discriminability for blocks 6-10, $A = .65$, 95% CI [.63, .68], but not for blocks 1-5, $A = .48$, 95% CI [.46, .51] (for the calculation of confidence intervals see DeLong, DeLong, & Clarke-Pearson, 1988). In addition, fitting the ROC model to each individual participant for blocks 6-10 revealed substantial variability across participants (see Figure 4.4-B).

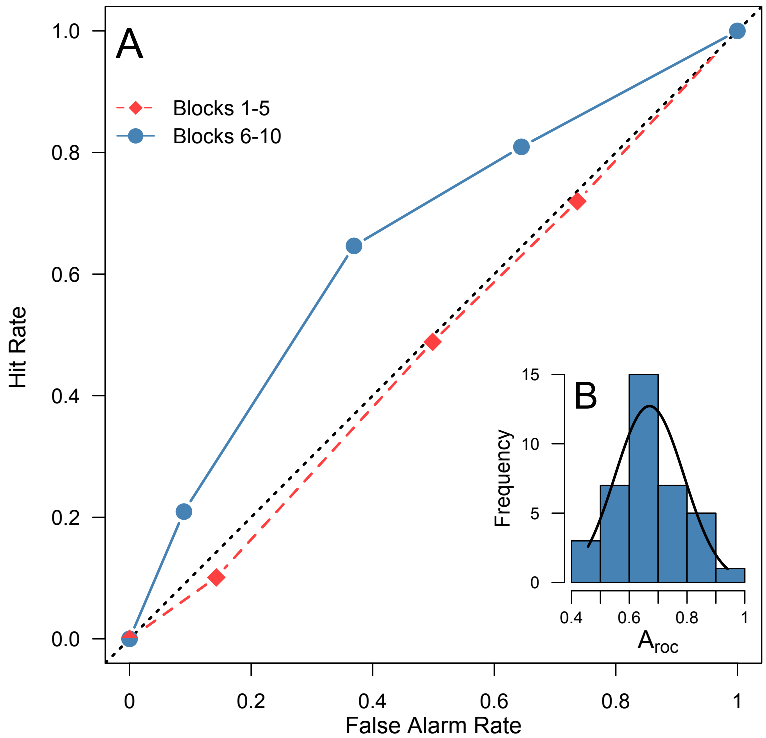

Figure 4.4: (A) Type 2 ROC curves for the blocks before (1-5) and after (6-10) the onset of good deck selection in Experiment 6. (B) Distribution of the area under the curve (A) when fitting the ROC model to each participant (blocks 6-10).

This fine-grained assessment of the confidence-accuracy relationship suggests that participants' decisions were accompanied by fairly accurate confidence reports. It is also important to investigate how participants utilised the confidence rating scale and whether there was any involvement of unconscious or implicit knowledge after the onset of good deck selection (blocks 6-10). The latter was assessed by using the *guessing criterion* (Dienes et al., 1995) according to which unconscious knowledge is present when participants can discriminate between good and bad decks at above chance levels when they are guessing or their confidence is low. We calculated the mean proportion of good deck selections for each confidence level. Importantly, at neither of the two low-confidence levels (i.e., 1 = "I am guessing" and 2 = "I am not confident") did deck selection significantly exceed chance (0.50) (Means for each level: 1 = 0.47, 2 = 0.49, 3 = 0.72, 4 = 0.78), indicating that good deck selections were not made under conditions of low confidence and that conscious knowledge strongly associated with above-chance performance on the IGT. Also, the mean confidence following good deck selections was 2.66 (SEM = 0.10) and for bad deck selections was 2.08 (SEM = 0.09). The difference between these values was significant, $t(37) = 8.50, p < .001$, suggesting the same conclusion as the guessing criterion. Participants were more confident when they made a good rather than a bad deck choice.

4.1.3 Discussion

The present experiment provides another demonstration of the involvement of conscious knowledge in the IGT. When participants started to consistently sample from the good decks, they were able to report their acquired knowledge through their confidence ratings. However, post-decision wagering showed a similar pattern as in the simple wagering condition of Experiment 3, namely a lag in the onset of wagering compared to deck selection. While this latter pattern might be indicative of unconscious processes in choice behaviour, the results from the confidence groups suggest a simpler explanation, namely that wagering is an insensitive measure of awareness.

The confidence rating scales produced similar results when the 4-point scale was collapsed into two categories. A more detailed examination of the continuous scale revealed that participants' deck selections were accompanied by accurate confidence ratings. Specifically, the ROC analysis showed increased metacognitive monitoring after the point at which performance on the IGT began to exceed chance. While the presence of conscious knowledge does not necessarily mean that unconscious or implicit processing is absent, data from the guessing criterion analysis suggest that deck selections and confidence ratings are highly related to one another.

4.2 Experiment 7

In all the experiments described so far, post-decision wagering was employed in a binary manner: participants could place a wager expressed as "Low" or "High". However, wagering can be used as a gradual scale, with discrete monetary wagers. According to Dienes (2008) a gradual scale with more than two categories is likely to be more sensitive than a binary scale as it allows finer discriminations of the acquired knowledge. In contrast, Tunney and Shanks (2003) showed that binary confidence scales were more sensitive in artificial grammar learning (AGL) tasks. Dienes (2008) used 6 different confidence scales in an AGL task: binary (high, low), binary (guess, sure), numerical (50%-100%) with and without detailed information about what the numbers should indicate, numerical categories in bins of 10 (e.g., 51-59, 60-69 etc.), and verbal categories ("complete guess", "more or less guessing", "somewhat sure", "fairly sure", "quite sure", "almost certain", "certain"). He surprisingly found no significant differences in terms of sensitivity as all of the scales demonstrated a similar measure of the confidence-accuracy relationship.

In this experiment, we employed two versions of post-decision wagering (binary and gradual with four monetary categories) and confidence ratings with four categories as in Experiment 6. Although many studies have shown that wagering (even with a gradual implementation) is no more sensitive and exhaustive than confidence ratings, it is possible that gradual wagering may be more appropriate than binary wagering.

4.2.1 Method

Participants

Fifty-three individuals (recruited from the UCL subject pool) participated in the study (33 females, age $M = 22.15, SD = 6.56$). They received £2 for their participation.

Design

There were three different conditions: binary wagering ($N = 18$), 4-point wagering ($N = 17$) and 4-point confidence ratings ($N = 18$). Participants were randomly allocated to one of three groups.

Task

The payoff schedule of Experiment 1 was used in this experiment (see Appendix A). The reason was to examine gradual wagering under the payoff schedule used by Persaud et al. (2007) and compare it against binary wagering. After their deck selection participants were asked to report their awareness in one of the three measures. In the binary wagering condition, participants chose between £10 (Low) and £20 (High) whereas in the 4-point wagering participants had 4 available wagers: £5, £10, £15, and £20. The selected wager multiplied the payoffs of each deck and trial. The descriptions for the 4-point confidence scale were as in Experiment 6: 1 = "I am guessing", 2 = "I am not confident", 3 = "I am quite confident", and 4 = "I am very confident". Since there was no change in the payoffs in the confidence ratings group, they were scaled up by a factor of 10 in order to maintain a similar level of payoffs across conditions.

4.2.2 Results

Choice and Awareness

In this stage, the 4-point scales (wagering and confidence ratings) were dichotomised with wagers £5 and £10 indicating low wagering and wagers £15 and £20 high wagering. Similarly, confidence levels 1 and 2 were collapsed to signify low confidence and levels 3 and 4 to give high confidence. Following the procedure of the previous experiments, neither of the measures (choice and awareness) was significantly above chance throughout the task (all one-sample t-tests, $p > .05$; Figure 4.5).

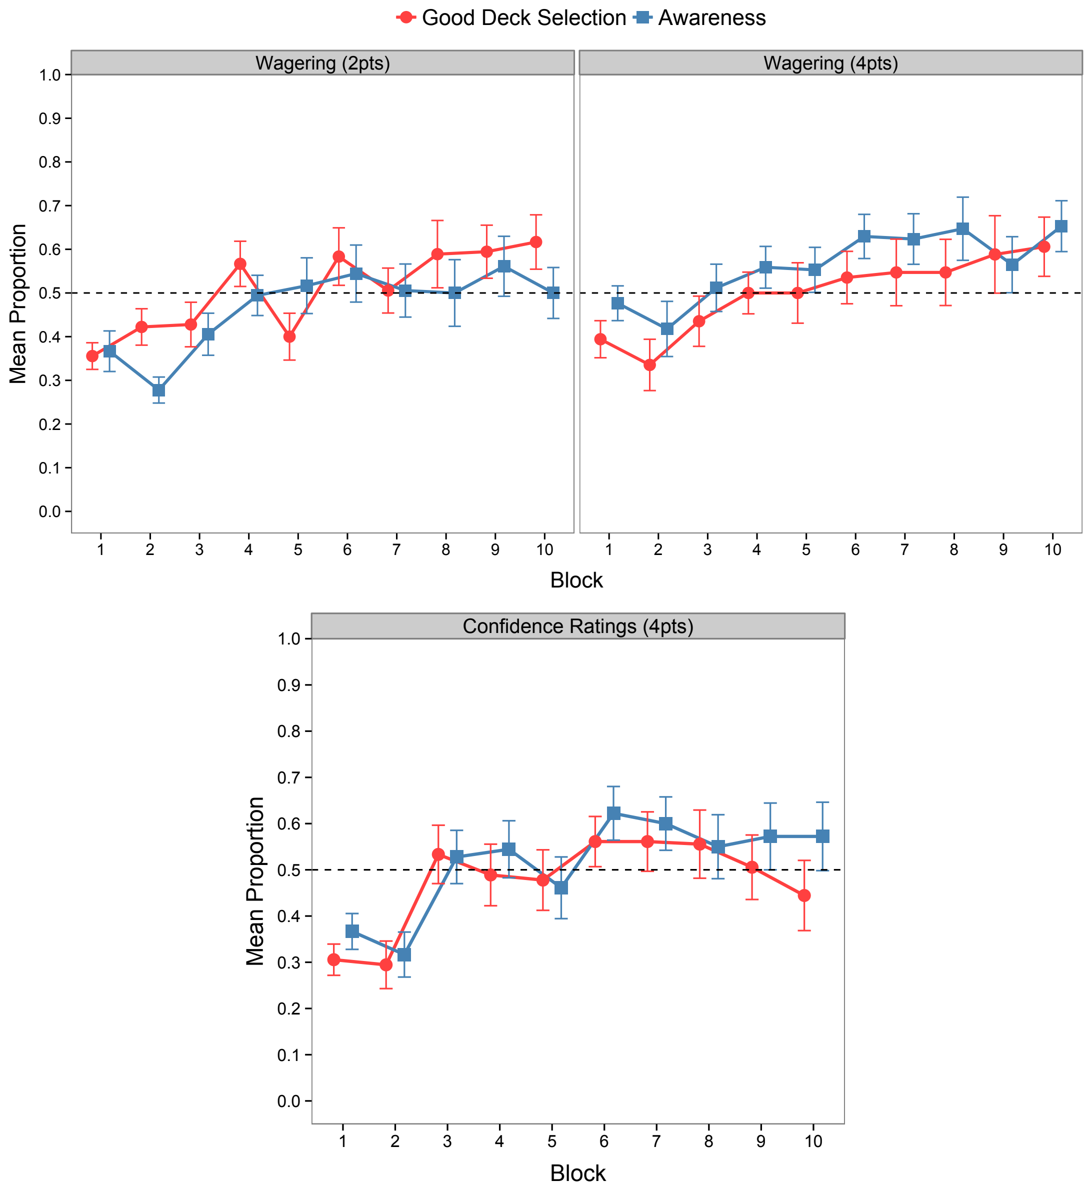

Figure 4.5: Proportion of good deck selections and awareness across blocks of trials for each group in Experiment 7.

Participants consistently selected more cards from the disadvantageous deck B, showing a strong loss-frequency effect and attention to immediate higher rewards (Figure 4.6). Deck B is the overall deck of choice across conditions followed by deck D. Participants never started to reliably choose from the good decks which is evident in the progression of the good deck selections measure in Figure 4.5. Selections from the good decks show an increasing trend across blocks of trials, $F(6.79, 339.64) = 7.63, MSE = 4.87, p < .001, \eta_g^2 = 0.09$, but this trend is rather weak as it never reached above chance performance. Neither the main effect of condition, $F(2, 50) = 0.24, MSE = 23.08, p = .79$, nor the interaction (condition × block), $F(13.59, 339.64) = 1.01, MSE = 4.87, p = 0.45$, reached significance. These results suggest that participants behaved similarly across conditions, favouring the decks with infrequent losses.

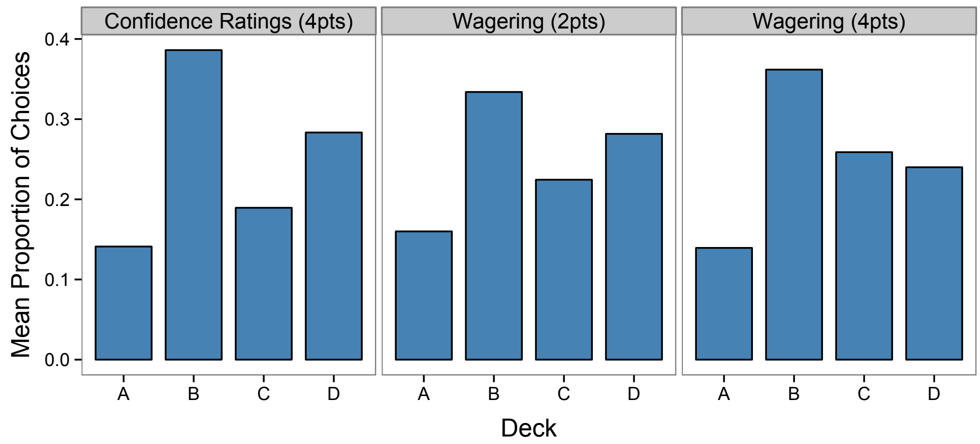

Figure 4.6: Mean proportion of choices from each deck and group in Experiment 7.

Mean awareness ratings across blocks revealed a similar picture: the main effect of block was significant, $F(6.82, 341.37) = 8.00, MSE = 4.81, p < .001, \eta_g^2 = 0.09$, but awareness was constrained due to the overall low performance in the task. The progression of awareness was similar across conditions as the main effect of condition, $F(2, 50) = 2.30, MSE = 17.67, p = .11$, and the interaction, $F(13.65, 341.37) = 0.46, MSE = 4.81, p = 0.97$, were not significant. At a first glance, Figure 4.5 suggests that awareness did not lag behind performance on the task, indicating no dissociation between the two measures. In fact, in the 4-point wagering group, awareness is higher than good deck selections throughout the task which suggests that participants wagered advantageously by placing low wagers (£5, £10) on bad deck selections. However, the observed pattern of the two measures is not overly suggestive that learning and awareness developed in parallel. It is possible that if participants had learned the advantageous strategy and selected more cards from the good decks, a dissociation could have emerged.

Although many IGT studies have found a clear preference for deck B (see Steingroever, Wetzels, Horstmann, et al., 2013), this preference is taken as evidence of poor performance on the task. To rule out the possibility that learning and awareness are not dissociable in the context of the IGT, we applied a median split on the overall proportion of selections from the good decks (C and D), in order to examine how awareness progressed in the group of participants who performed advantageously on the task (Figure 4.7).

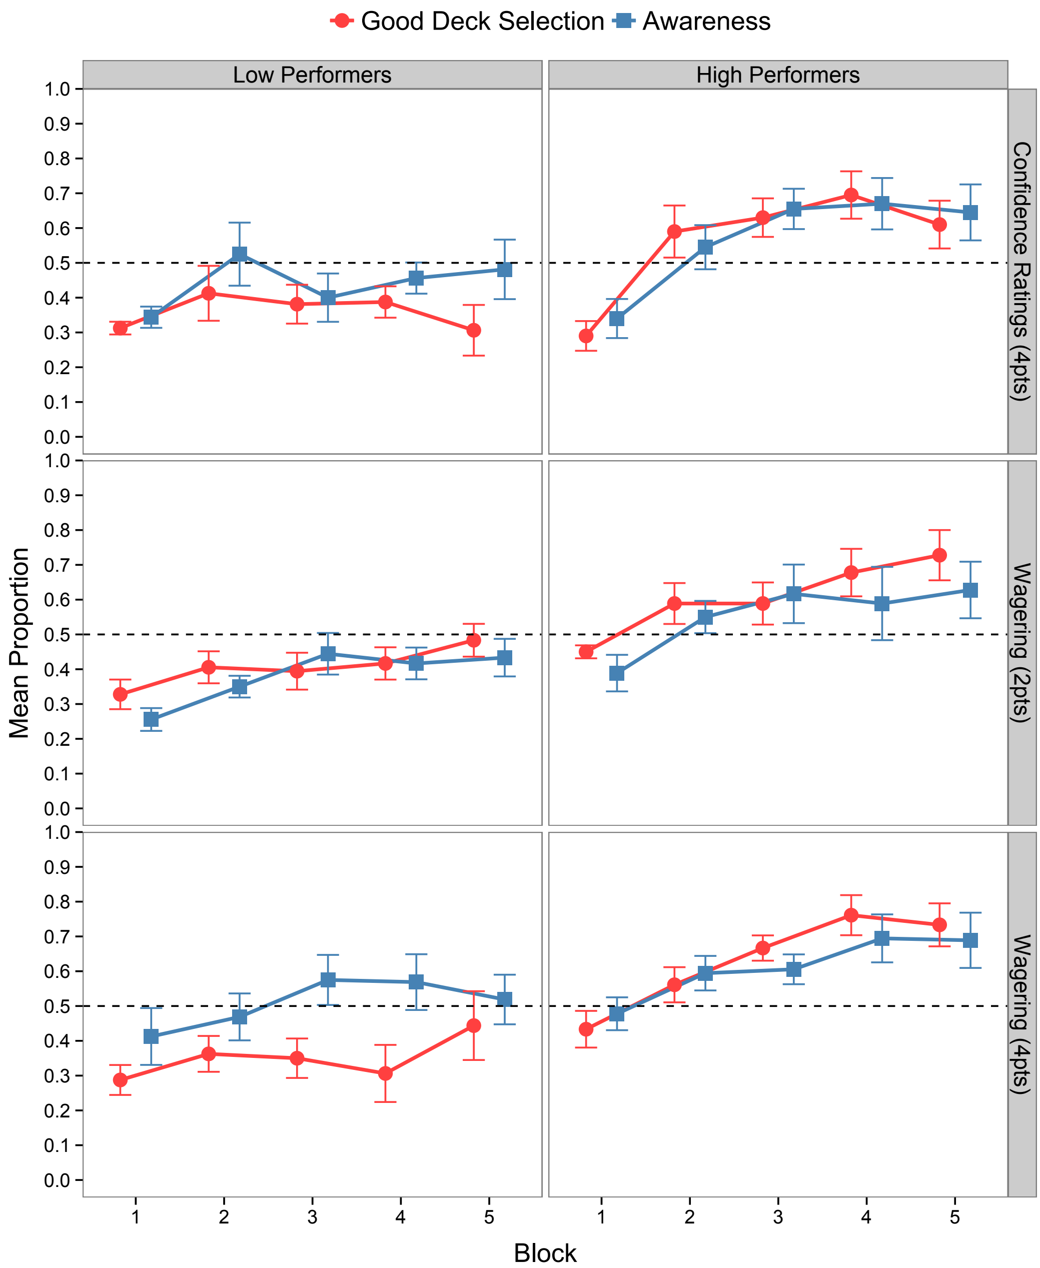

Figure 4.7: Proportion of good deck selections and awareness for low and high performers across blocks of trials and groups in Experiment 7.

Focusing on the high performing group, we identified the onset of choice and awareness for each of the three conditions. Good deck selection exceeded chance on block 3 (trials 41-60) for the 4-point confidence group, $M = 0.63, t(9) = 2.35, p = .043$, and the 4-point wagering, $M = 0.67, t(8) = 4.59, p = .002$, and block 4 (trials 61-80) for the binary wagering group, $M = 0.68, t(8) = 2.60, p = .031$. These results indicate that the onset of good deck selection for the high performers is similar to that of previous experiments. Regarding conscious knowledge of the advantageous strategy, the 4-point confidence and wagering scales significantly exceeded chance on block 3 (4-point CR: $M = 0.66, t(9) = 2.67, p = .025$, 4-point wagering: $M = 0.61, t(8) = 2.46, p = .039$), whereas the binary wagering never crossed the chance level (all blocks, $p > .05$), indicating a possible failure to detect acquired knowledge. The latter might be taken as evidence that the binary wagering employed in previous experiments and by Persaud et al. (2007) is an insensitive method. On the other hand, gradual wagering tracked closely good deck selection, suggesting that the use of more than two categories could potentially allow participants to use wagering adaptively and in accordance with their knowledge. However these are weak conclusions as the patterns in the two wagering groups are numerically very similar, as shown in Figure 4.7.

ROC Analysis

As in Experiment 6, we employed a nonparametric ROC analysis for a more detailed assessment of conscious knowledge and how participants utilised each awareness scale. We only applied the analysis on the high performing group of each condition and we constructed two ROC curves, one before the onset of good deck selections (blocks 1-2) and one after (blocks 3-5). Figure 4.8 shows good discriminability between good and bad decks after the onset of good deck selection (block 3) which increase with higher confidence or higher wagering amounts. All three scales demonstrated above chance metacognitive sensitivity (A, the area under the curve - AUC), 2-point wagering: $A = .55$, 95% CI $[.52, .60]$, 4-point wagering: $A = .64$, 95% CI $[.59, .68]$, 4-point confidence ratings: $A = .59$, 95% CI $[.54, .63]$. The analysis showed that the AUC for the 4-point wagering was numerically above .50 in blocks 1-2 but not significantly different from chance ($A = .53$, 95% CI $[.47, .59]$). A comparison of the AUCs showed that there was a significant difference between the 4-point wagering and the binary wagering scales, $z = 2.38, p =$

.017, but no other significant differences were observed.

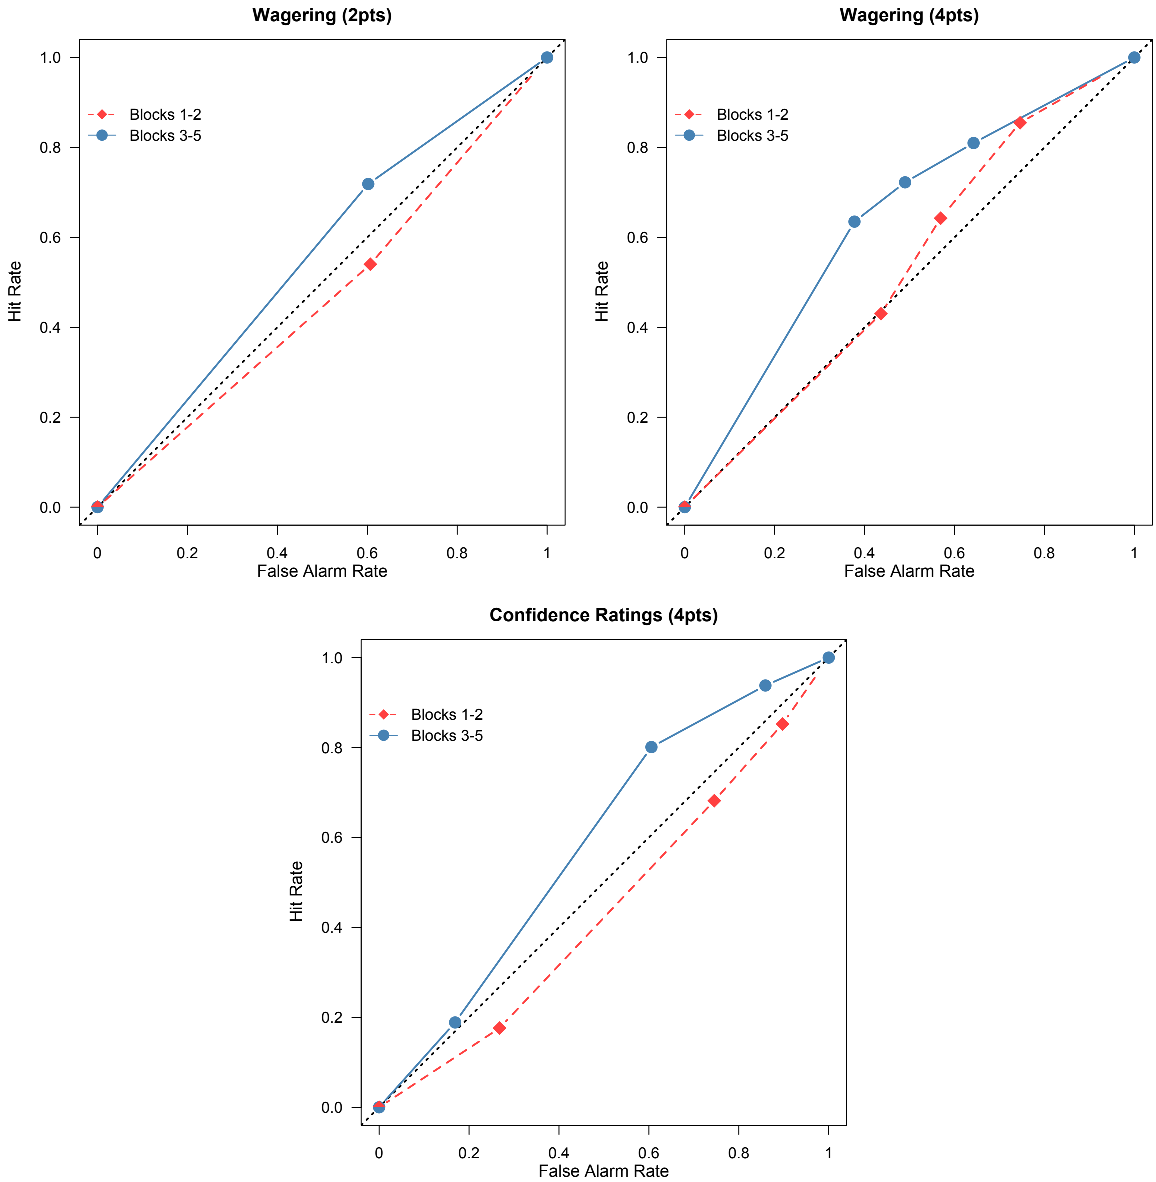

Figure 4.8: Type 2 ROC curves for the blocks before (1-2) and after (3-5) the onset of good deck selections for each group in Experiment 7.

We also examined how participants used the scales and the proportion of low ratings after selecting a good deck. According to the guessing criterion, if participants use the lowest ratings of a scale while their performance is above chance then this is an indicator of unconscious learning. The proportion of good deck selections for each confidence rating indicated that more good deck selections made under high confidence (Means for each level: $1 = 0.44, 2 = 0.49, 3 = 0.72, 4 = 0.67$), and a similar pattern was observed for the 4-point wagering scale (Means for each level: £5 $= 0.58$, £10 $= 0.59$, £15 $= 0.67$, £20

= 0.81). However, while there is a gradual increase from the lowest to the highest wager amount, the proportion of good deck selections at the lowest rating (£5) was above 0.50 which would indicate that above-chance performance is guided by implicit influences. However, this pattern can be explained by loss aversion: even when participants have acquired some conscious knowledge about their decisions and can discriminate between good and bad decks, the prospect of losing money/points may lead them to place low wagers.

4.2.3 Discussion

In this experiment we compared three different measures of awareness in the context of the IGT, binary wagering, 4-point wagering, and 4-point confidence ratings. The main results indicated that participants did not learn to significantly discriminate between good and bad decks and awareness was constrained due to this pattern. In order to examine whether learning dissociates from awareness, we performed a median split on the proportion of good deck selections and focused on those participants who performed advantageously. The results showed no dissociation between learning and awareness when participants started to select more cards from the good decks in the 4-point scales (wagering and confidence ratings). However, awareness as indexed by the binary wagering measure never reliably exceeded chance indicating a possible failure to detect conscious knowledge. In addition, the ROC analysis showed that there was a significant difference between binary and 4-point wagering in assessing the level of metacognitive monitoring suggesting that wagering can possibly be more sensitive if it is used with more categories.

The analysis based on the guessing criterion revealed a difference between 4-point wagering and confidence ratings in the way participants used the two scales. While confidence ratings showed no involvement of unconscious learning after the onset of good deck selections (i.e., the proportion of good deck selections at the lowest rating was below 0.50), inspection of the 4-point wagering measure revealed some influence of implicit knowledge (the proportion of good decks selected at the lowest rating was 0.58). As it seems rather unlikely that in the same version of the IGT the involvement of implicit learning is present or absent depending on the scale used, we conclude that confidence ratings and wagering differ in their sensitivity to low levels of consciousness (Wierzchoń et al., 2012). The most plausible explanation is that wagering is affected by loss aversion, leading participants to place low wagers even if they are consciously aware of the task structure. The effect of loss aversion on wagering strategies which forces participants to select more often the lowest rating has been observed in many studies which compared different subjective measures of awareness (see Sandberg, Bibby, & Overgaard, 2013; Szczepanowski et al., 2013; Wierzchoń et al., 2012). In a recent re-analysis of Szczepanowski et al.'s (2013) study which compared three subjective measures of awareness (post-decision wagering, confidence ratings, and perceptual awareness scale [PAS] ratings), Sandberg et al. (2013) found that confidence and PAS ratings were significantly more sensitive than post-decision wagering. Sandberg et al. suggested that "there in now convincing evidence that post-decision wagering should only be used when participants are unable to use other, more direct measures such as confidence ratings and PAS, and when doing this, analyses should take loss aversion into account" (p. 809).

4.3 Experiment 8

The payoffs that participants receive in experience-based tasks can depend on a series of dimensions: probability, magnitude, domain (wins, losses, or mixed), and relationship or association with the feedback from other alternatives. Another distinction refers to the overall nature of the task, ranging from complete stationary environments (i.e., actions are sequential and the environment does not change over time) to complete dynamic environments where the environment changes over time based on some inherent properties of the system but also as a function of the agent's previous decisions (Brehmer, 1992; Busemeyer, 2002; Edwards, 1962). The original IGT can be seen as a semi-stationary (or semi-dynamic) task because the payoff schedule for each deck is preprogrammed (i.e., the outcomes are not sampled from a probability distribution) but the payoffs that each deck returns are not identical on every trial. In addition, participants' selections have no effect on future outcomes.

The purpose of this experiment was to examine how learning and awareness develop in a dynamic variant of the IGT. Specifically, the decks swap their payoffs every 50 trials in such a way that the good decks become bad and vice-versa. It is of particular interest to investigate participants' choice behaviour and whether it adjusts to the environment of each deck-shift period. Importantly, if participants exhibit adaptive behaviour and select advantageously, then the question is whether this is accompanied by conscious knowledge of the decks' quality.

4.3.1 Method

Participants

Forty participants (21 females, age $M = 26.28, SD = 8.47$) were recruited from the UCL subject pool. All participants received a turn-up payment of £1 and an additional amount up to £2, depending on their performance in the task.

Task and Design

The experiment consisted of two conditions: the control ($N = 20$) and the switch groups ($N = 20$). The original IGT payoff schedule was used and participants had to select a card from any deck they chose for 200 trials. After their card selection and before receiving the outcome of their choice, participants indicated their confidence in having selected a good deck using a 4-point confidence scale. The descriptions of the scale were the same as in previous experiments (see Experiments 6 and 7). In addition, every 50 trials participants were asked to provide ratings about the perceived goodness of each deck (Ratings - Question 1 in Appendix B) and which of the four decks they would choose if they could only select cards from one deck for the remainder of the task (Deck Selected - Question 3 in Appendix B).

In the switch group, a new deck contingency period was introduced every 50 trials, immediately after the administration of the two questions. The deck positions remained the same on the screen but the payoffs of each deck changed. During these shift periods, the initially good decks C and D (Period 1: Trials 1-50) were replaced by decks A and B (Period 2: Trials 51-100). In shift Period 3 (trials 101-150) each deck had the same payoffs as in Period 1 whereas in Period 4 (Trials 151-200) each deck changed quality (good or bad) and the probability of the losses it produced (frequent or infrequent) (see Figure 4.9).

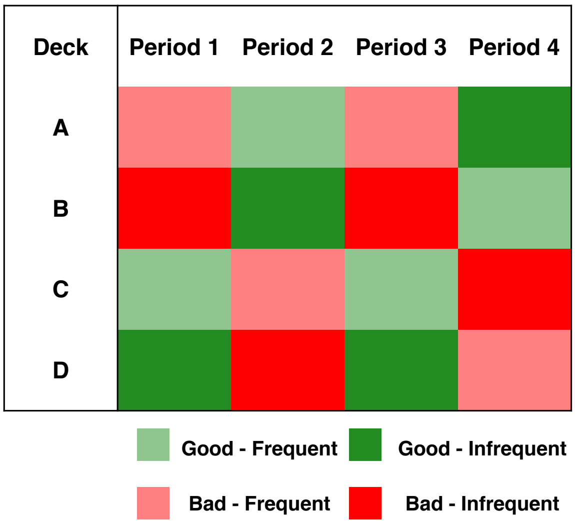

Figure 4.9: Schematic representation of the quality (good-bad) and frequency of losses (frequent-infrequent) of each deck for the four shift periods in the switch group of Experiment 8.

4.3.2 Results

Control Group

As in previous experiments, the 4-point confidence scale was dichotomised so that the lowest ratings (1 and 2) indicated low confidence whereas the highest ratings (3 and 4) indicated high confidence. We then identified the point at which both measures became significantly better than chance. Good deck selection exceeded chance on block 4 (trials 61-80), $M = 0.64, t(19) = 2.83, p = .01$, at the same point as awareness $M = 0.65, t(19) = 2.81, p = .011$ (Figure 4.10). A repeated-measures ANOVA showed significant main effects of block on good deck selection, $F(4.61, 87.64) = 10.26, MSE = 13.15, p < .001, \eta_g^2 = 0.21$, and awareness, $F(3.78, 71.84) = 2.73, MSE = 20.65, p = .03, \eta_g^2 = 0.08$. These results suggest a similar interpretation as in previous experiments: when participants started to consistently select more cards from the good decks, conscious knowledge about their choices closely followed. The slightly later onset of good deck selections compared to previous experiments can be attributed to the longer duration of the present task. It is possible that participants spent more time exploring the task before they switched to the decks which give better outcomes since they knew that the task comprised 200 trials.

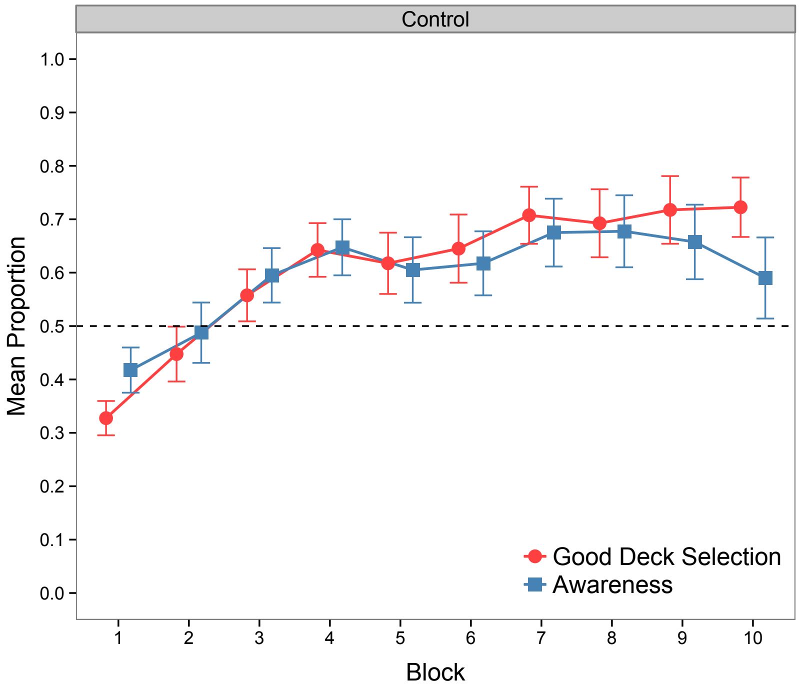

Figure 4.10: Proportion of good deck selections and awareness across blocks of trials in the control group of Experiment 8.

The next step in the analysis was to examine the proportion of choices from individual decks across blocks of trials in parallel with the confidence ratings and questionnaire responses. Figure 4.11-A shows the mean proportion of choices from each deck across blocks of 50 trials which is in accordance with the mean confidence ratings after each deck selection (Figure 4.11-B). During the first block, participants selected more cards from decks B and D and this is reflected in the confidence ratings. Both decks receive a higher rating compared to the decks with infrequent losses, A and C. However, the questionnaire responses revealed a different picture: even in the first administration of the questionnaire participants rated the good decks C and D higher than the bad decks A and B (Figure 4.11-C). Decks C and D receive positive ratings throughout the task whereas decks B and D are consistently on the negative side of the scale.

A 2 (deck [good, bad]) × 4 (block) repeated measures ANOVA showed a significant main effect of deck, $F(1,19) = 29.28, MSE = 41.3, p < .001, \eta_g^2 = 0.40$, with good decks receiving higher ratings than bad decks, and a significant deck × block interaction, $F(3,57) = 5.97, MSE = 7.55, p = .001, \eta_g^2 = 0.07$, indicating that the difference between good and bad decks increased across blocks. The main effect of block was not significant, $F(3,57) = 1.04, p = .38$. A similar picture emerges in the "deck selected" question (Figure 4.11-D): the majority (70%) of participants would select one of the good decks for the remainder of the task even when they were asked on trial 50 and this percentage increases in the remaining blocks (75%, 85%, 90%, respectively).

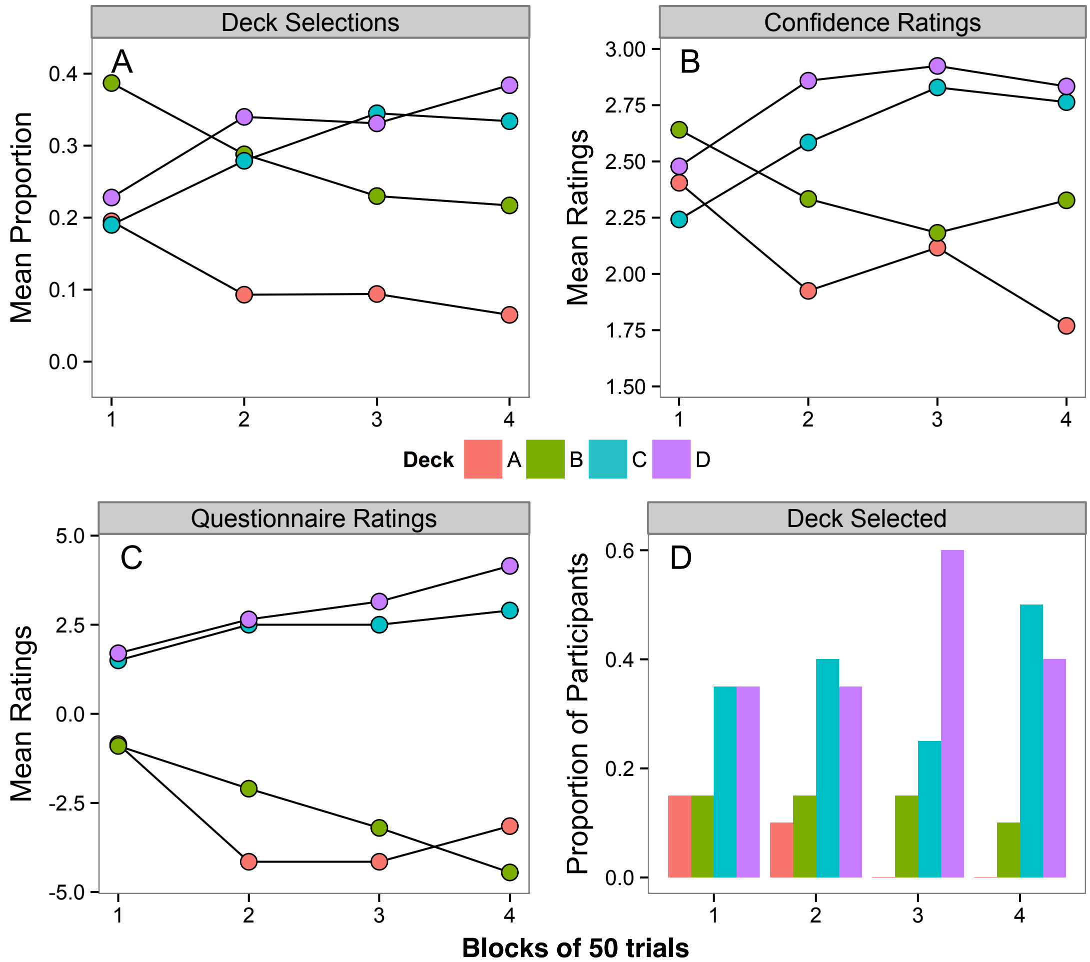

Figure 4.11: Deck selections and awareness measures across blocks of 50 trials in the control group of Experiment 8. A) Mean proportion of choices from each deck. B) Mean confidence ratings for each deck. C) Mean questionnaire ratings. D) Proportion of participants who would select each deck for the remainder of the task.

These results demonstrate that conscious knowledge of the deck contingencies is the driving force for participants to behave advantageously. Although knowledge of the advantageous strategy is present even before the first administration of the questionnaire, participants did not exclusively select cards from the good decks. In fact, our results are consistent with those of Maia and McClelland (2004) who noted that "the tendency is for knowledge of the advantageous strategy to be more evident in all of the verbal report measures than in behaviour (which may be due to exploration of the different decks or risk-taking by some participants)" (p. 16077).

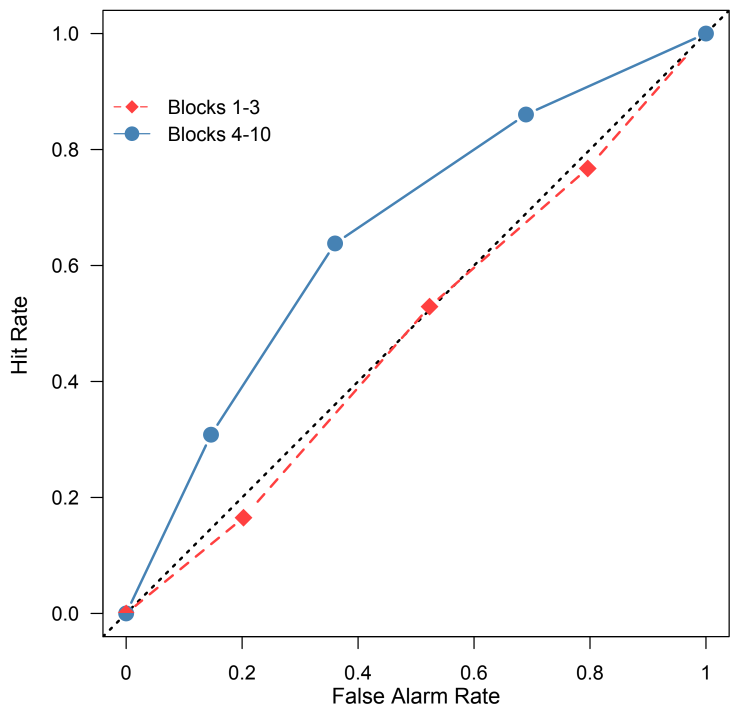

Figure 4.12: Type 2 ROC curves for the blocks before (1-3) and after (4-10) the onset of good deck selections in the control group of Experiment 8.

The confidence-accuracy relationship was examined using an ROC analysis. Two curves were constructed, one before (blocks 1-3) and one after the onset of good deck selections (blocks 4-10). Deck selection performance did not significantly change across blocks 4-10, $F(6, 114) = 1.25, MSE = 11.31, p = .28$. Figure 4.12 shows that there is above-chance metacognitive discriminability after the onset of choice (blocks 4-10), $A = 0.66$, 95% CI [0.64, 0.68] but not before (blocks 1-3), $A = 0.48$, 95% CI [0.45, 0.52]. The proportion of good deck selections for each confidence level showed that higher ratings were given to the good decks (levels: $1 = 0.48, 2 = 0.58, 3 = 0.76, 4 = 0.82$). According to the guessing criterion, the lowest confidence rating did not reveal above chance accuracy indicating no involvement of unconscious knowledge in advantageous selections (Dienes et al., 1995; Wierzchoń et al., 2012) which is consistent with the guessing criterion analysis in Experiments 6 and 7.

Switch Group

Figure 4.13 shows good deck selection and awareness in the switch group. As in the control group, the awareness measure is calculated after dichotomising the 4-point confidence scale. The good deck selection measure represents the proportion of selections from the advantageous decks across shift periods (dotted vertical lines). For example, in shift period 2 (blocks 6-10) the red line indicates the proportion of selections from decks A and B which are advantageous in this period. A repeated-measures ANOVA showed a significant main effect of shift period, $F(3, 57) = 4.26, MSE = 48.67, p = .008, \eta_g^2 = 0.09$, as participants progressively selected more cards from the advantageous decks. Individual repeated-measures ANOVAs were performed within each shift period. The main effect of block was significant only in period 3, $F(4, 76) = 3.16, MSE = 5.67, p = 0.02$, in which participants selected more cards from the good decks. Examination of the awareness measure revealed a similar picture. There was a significant main effect of shift period, $F(3, 57) = 4.75, MSE = 42.82, p = .005, \eta_g^2 = 0.10$, and significant effects of block in period 3, $F(4, 76) = 5.56, MSE = 3.22, p < .001$, and period 4, $F(4, 76) = 3.99, MSE = 4.91, p = .005$.

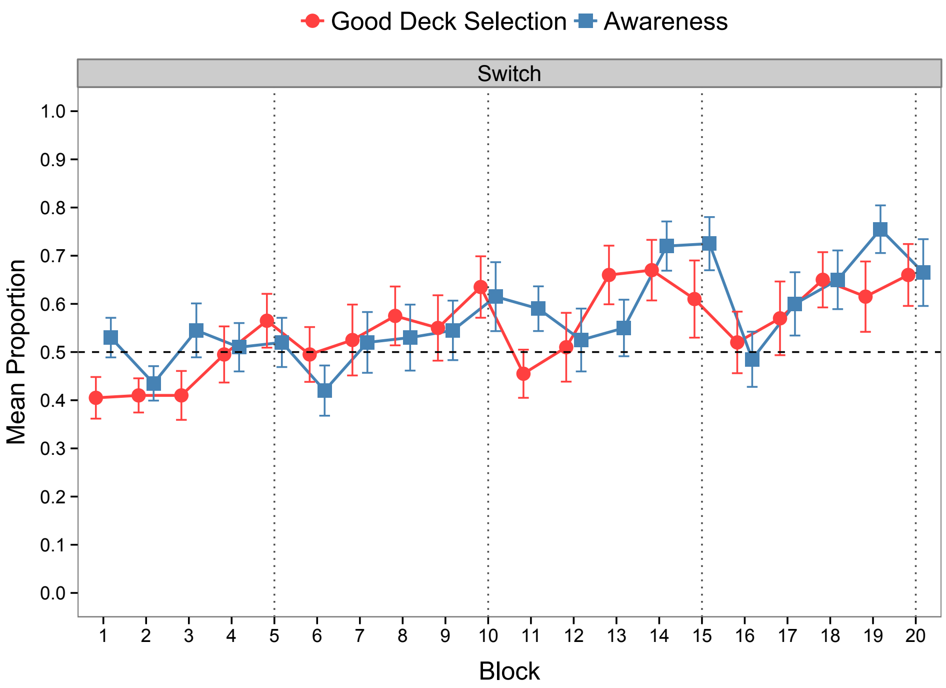

Figure 4.13: Proportion of good deck selections and awareness across blocks of trials and shift periods (vertical dotted lines) in the switch group of Experiment 8.

To further investigate participants' choice behaviour and the knowledge they acquired during the task, we analysed their responses in the questionnaire reports. The decks were categorised based on whether they were good (green colour) or bad (red colour) and whether they produced infrequent (solid colour) or frequent losses (transparent colour) in each shift period (Figure 4.14). Choice behaviour was similar to that in the control group in period 1 as participants selected more cards from the bad deck with infrequent losses (i.e., deck B) and their mean confidence ratings followed this tendency (Figure 4.14-B). When looking at the questionnaire responses on trial 50 (end of shift period 1) a slightly different pattern emerges: participants seem to be aware of the task structure as they gave higher ratings for the good decks (Figure 4.14-C) and opted for one of the good decks in the deck selected measure (Figure 4.14-D). A 2 (deck [good, bad]) $\times$ 4 (block) repeated measures ANOVA showed a significant main effect of deck, $F(1, 19) = 21.97, MSE = 12.96, p < .001, \eta_g^2 = 0.16$, with good decks receiving higher ratings than bad decks, but the interaction between deck and block was not significant, $F(3, 57) = 0.59, MSE = 6.58, p = .62$, suggesting that the difference between good and bad decks remained constant across blocks of trials. Interestingly, the good deck with frequent losses received a marginally higher rating and more participants would select it compared to the good deck with infrequent losses. In the following shift periods (2, 3, and 4) the pattern of questionnaire responses is similar: good decks attain higher ratings and participants would prefer these decks to select cards from for the remainder of the task. In other words, participants were able to discover which were the good decks by the end of each period and indicated their preference and knowledge in the quantitative questions.

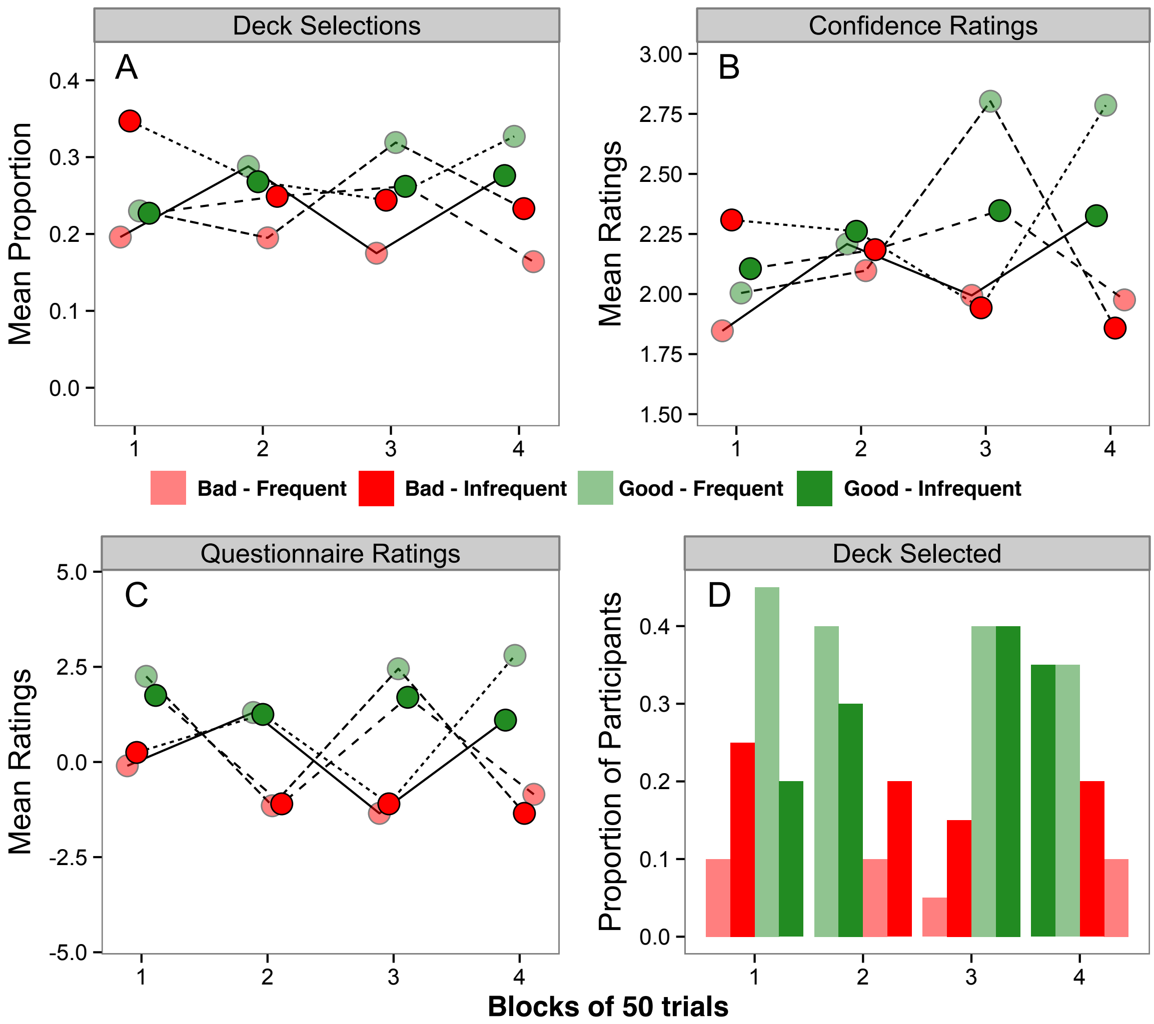

Figure 4.14: Deck selections and awareness measures across blocks of 50 trials in the switch group of Experiment 8. The green and red colours indicate good and back decks, respectively. Transparent and solid colours indicate decks with frequent and infrequent losses, respectively. A) Mean proportion of choices from each deck. B) Mean confidence ratings for each deck. C) Mean questionnaire ratings. D) Proportion of participants who would select each deck for the remainder of the task.

The mean proportion of good deck selections and confidence ratings in period 2 showed that participants did not translate their acquired knowledge into making more advantageous choices. Even though confidence ratings discriminate between good and bad decks by period 3, the effect of this knowledge on choice behaviour is rather weak as the difference between good and bad decks is small (Figure 4.14-A) even in the last 5 blocks of the task (period 4). However, this approach takes into account all the choices within one period and thus participants' choices within the first blocks of each period may reflect what they have learned about the task and the decks from the immediately preceding period. Thus, we looked at the choice behaviour and confidence ratings at the last block of each period. Figure 4.15 shows a pattern which is in accordance with the questionnaire responses. Apart from block 5 in which participants thought that deck B (bad deck-infrequent losses) was a good deck and gave higher confidence ratings after selections from this deck, the remaining blocks showed a consistent pattern: advantageous decks with high expected values are selected more often and their ratings suggested a strong relationship between selections and awareness. Importantly, participants seemed to effectively adapt to the shifts in contingencies and learned the advantageous strategy by the end of each shift period.

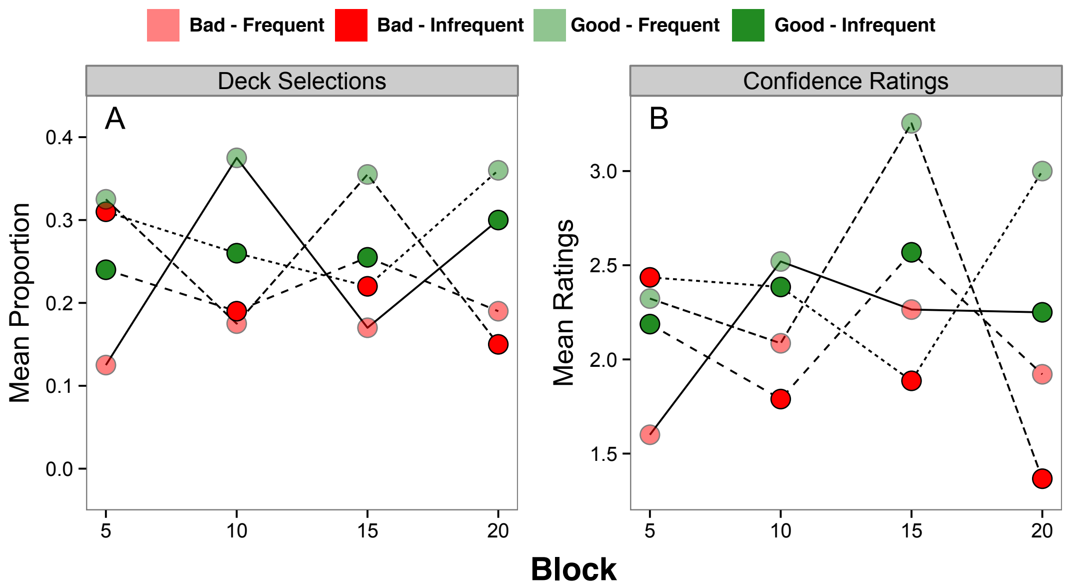

Figure 4.15: Mean proportion of choices from each deck (A) and mean confidence ratings (B) for the last block of each shift period in the switch group of Experiment 8.

Another way to examine whether participants learned the advantageous strategy in each period is to look at the number of switches between decks as a rough measure of exploration. It was hypothesised that participants would explore more in the first trials of each period but when they discovered which were the good decks they would start to exploit and select more cards from these decks. As expected, Figure 4.16 shows this pattern: participants switched between cards more often during the first blocks than later blocks of each period and this tendency was more pronounced in the last 2 shift periods of the task. The main effect of block was significant only in periods 3, $F(4, 76) = 3.94, MSE = 0.04, p = .006, \eta_g^2 = 0.05$, and 4, $F(4, 76) = 2.52, MSE = 0.04, p = .048, \eta_g^2 = 0.03$, indicating that participants progressively decreased the number of switches and focused more on a single deck to choose cards from.

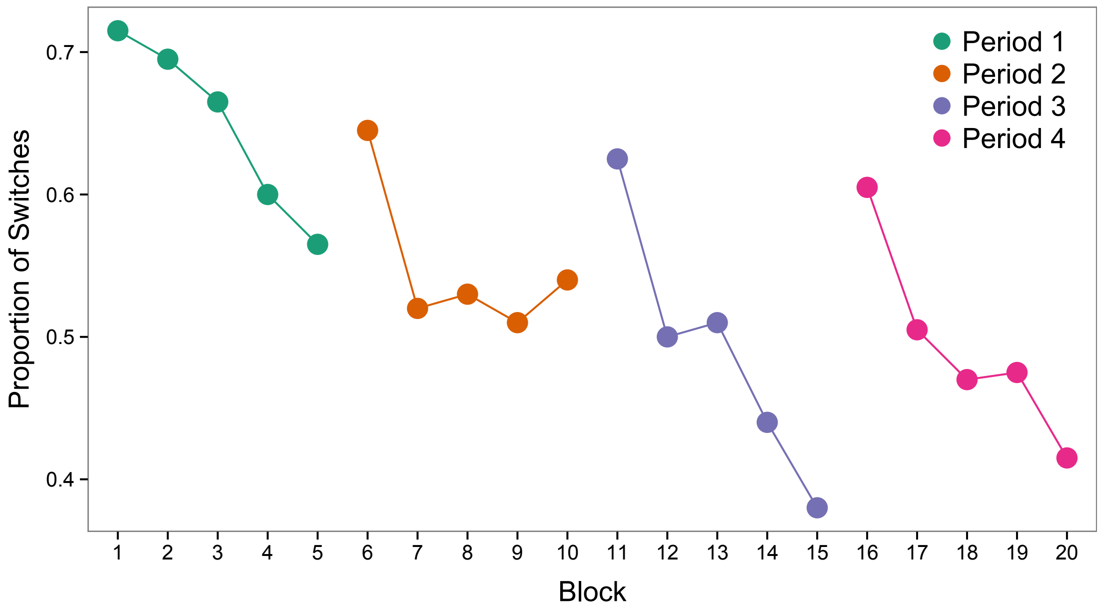

Figure 4.16: Mean proportion of switches across blocks of trials for each shift period in the switch group of Experiment 8.

The next step was to examine choice behaviour in the "exploration" trials and whether there was systematic switch behaviour. For example, it could be the case that participants switch from a good deck to the other good deck or to a bad deck and vice-versa, and from a bad deck to a bad deck or a good deck. We found no systematicity in exploration behaviour as shown in Figure 4.17. When participants switched decks, they did it in an exploratory manner and no apparent trend was observed. Individual chi-squared analyses for each shift period showed no significant effects apart from shift period 1, $\chi^2(1) = 8.29, p = .004$. The probability of switching to a good deck after a good deck selection is significantly smaller than the other switches. Participants prefer the decks with higher positive rewards hence are more likely to switch to a bad deck after a good deck selection.

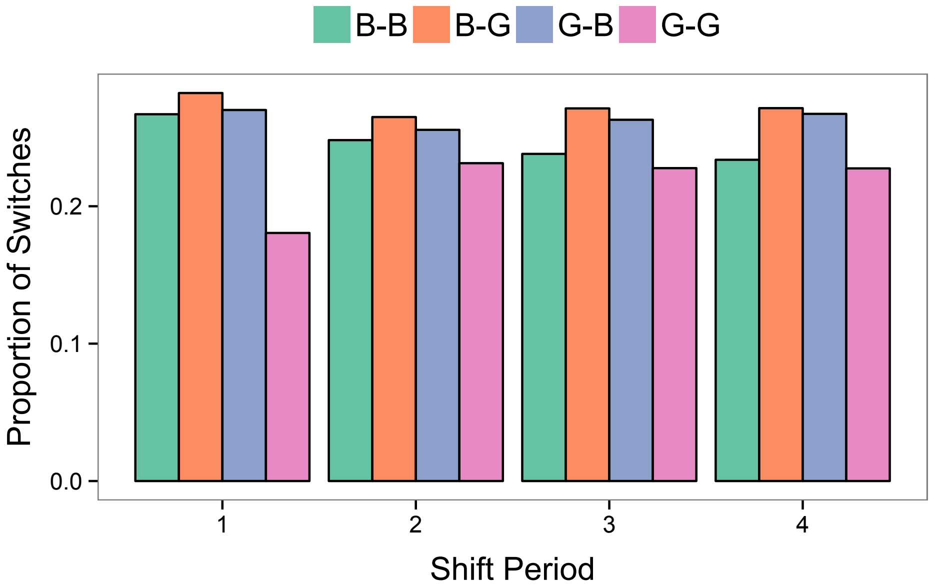

Figure 4.17: Mean proportion of switches for each category across shift periods in the switch group of Experiment 8. B: Bad deck, G: Good deck

An interesting finding of the present experiment relates to the choices from each deck separately. While most IGT studies have found a strong loss-frequency effect with participants preferring the decks which produce rare losses (decks B and D - solid coloured decks in Figures 4.14 and 4.15), the choice behaviour observed in this experiment suggests that participants had a slight preference for the good decks with frequent losses. This is evident not only in the proportion of choices from each deck but also in the confidence ratings and the questionnaire measures.

The confidence-accuracy relationship was further examined using an ROC analysis. Four curves were fitted to the data for each shift period. As shown in Figure 4.18-A participants exhibited good metacognitive monitoring in the last two periods of the task. The AUC was numerically above chance in period 2 but this was not significant, $A = .53$, 95% CI [.49, .57]. As in previous analyses, the problem is that each period also contains "transfer" trials. When the ROC analysis is performed only on the last block of each period, there is a stronger relationship between discrimination of good and bad decks and confidence. The AUC is significantly above chance even in the last block of period 2 $A = .67$, 95% CI [.60, .74], indicating the close link between learning of the advantageous strategy and awareness.

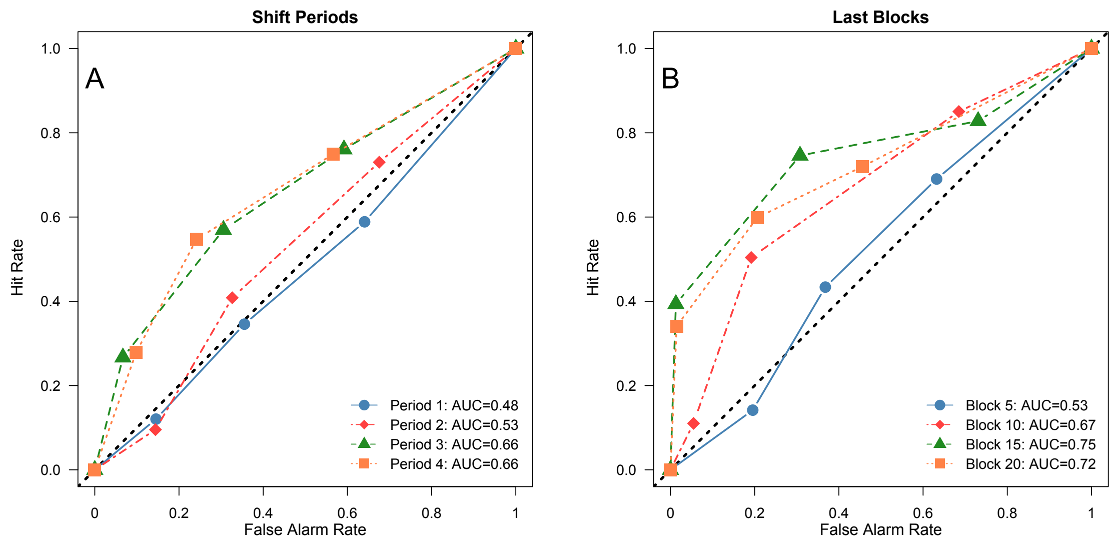

Figure 4.18: Type 2 ROC curves for each shift period (A) and the last block in each period (B) in the switch group of Experiment 8.

4.3.3 Discussion

The purpose of the present experiment was to examine learning and awareness in a controlled dynamic environment. In the control group, where there were no shifts in the decks' quality and payoffs, participants' choice behaviour and conscious knowledge developed in parallel, with no evidence of unconscious influences. Their questionnaire verbal reports showed that they acquired significant insights about the deck properties which drove their advantageous selections. Even in the first 50 trials in which they select more cards from the disadvantageous decks, their questionnaire responses showed that they possessed discriminative knowledge about the decks but they did not apply it to make more selections from the good decks. If anything, this suggests that decision-making in experience-based tasks is not only governed by explicit knowledge of the decks' values (which is substantial) but in addition involves a combination of exploration and risk taking behaviours. Also, participants' use of the confidence ratings scale was in alignment with their choices and never lagged behind advantageous selections. These findings support the view that decision-making is not dissociable from awareness and that conscious knowledge about the decks is sufficient to guide choice.

A similar interpretation applies to the findings from the switch group. Previous research has found that people are good at detecting change in the environment and adapt their choices based on new information (e.g., Dymond, Cella, et al., 2010; Turnbull et al., 2006). We extended these findings by investigating the level of conscious knowledge that people possess when they make decisions in dynamic environments. It can be argued that in such volatile environments in which uncertainty looms higher than in the original IGT, participants do not acquire conscious knowledge of their decisions but come to rely instead on their intuition and on simple strategies to cope with the uncertainty and make advantageous decisions. However, our results suggest the opposite interpretation: conscious knowledge of the task structure was present as shown in the trial-by-trial confidence ratings and the questionnaire responses at the end of each shift period. A clear preference for the good decks was more profound in the last blocks of each period after an initial stage of exploration which decreased across blocks of trials. As in the control group, knowledge of the advantageous strategy was more evident in the questionnaire measures and confidence ratings rather than choice behaviour.

The IGT in the switch group can be seen as a controlled dynamic task because each shift period did not introduce a new payoff scheme as decks only swapped positions and the overall expected values associated with good and bad decks remained constant. Participants had to learn first which were the good decks and then apply this knowledge on subsequent shifts periods. A simple strategy that they could have adopted is to select the decks with smaller rewards as these decks are advantageous in the long run. In addition, the administration of the questionnaire coincided with the start of a new shift period which could signal them about possible changes in the task. Participants seemed to employ this strategy in periods 3 and 4 as shown by the increase of good deck selections following an exploration period. However, exploration of the task lasted about 2 blocks of trials until they started to consistently select cards from the good decks which was guided by explicit knowledge of the decks' properties.

5 Cognitive Modelling

The previous chapters demonstrated that choice behaviour in the IGT is not consistent and uniform across experiments. Even though participants learned the advantageous strategy and selected more cards from the good decks in most experimental conditions, this outcome was more pronounced in some experiments compared to others. Also, the mean proportion of selections from each deck indicated significant differences between experiments. We identified three different choice patterns based on the overall proportion of choices from each deck: preference for the good decks (decks C and D, e.g., Experiment 1), preference for the decks with infrequent losses (decks B and D, e.g., Experiments 6 and 7) and preference for the good decks and the bad deck with infrequent losses (decks C, D, and B, e.g., Experiments 4B and 8). This variability in choice behaviour arises from a combination of factors, including learning long-term contingencies, sensitivity to losses and loss aversion, and the trade-off between exploration and exploitation. Inspection of choice profiles is not always adequate to identify which of these factors is most responsible for the observed pattern (Busemeyer & Stout, 2002).

Cognitive modelling provides a way to decompose and quantify the underlying psychological processes which drive performance on a task. In the case of the IGT, it is important to investigate the serial dependencies between successive choices and how a decision is made based on the interplay of a range of psychological factors. In this chapter, we will focus on cognitive models for the IGT with an emphasis on reinforcement-learning (RL) models.

5.1 Reinforcement learning and the IGT

The main objective of RL models of the IGT is to account for how future decisions are shaped by the experienced history of previous decisions and their associated payoffs. In other words, these models try to account for the sequential dependencies between each current choice and the previous choices and payoffs. Performance on the IGT can be explained based on the interplay of psychological processes which are quantified by each model's parameters (Steingroever, Wetzels, & Wagenmakers, 2013a). RL models employ three different assumptions about decision-making in the IGT. The first assumption relates to the objective evaluation of wins and losses (if any) after selecting a deck by a *utility* function. These utilities serve as input to a *learning* rule which updates expectancies about the utilities of each deck (second assumption). Hence, on a given trial each deck is associated with an expectancy which is the product of adjusting all the previous representations of utility from that deck. The third assumption is that a *choice* is made between decks based on these adjusted expectancies. It follows that the higher the expectancy of a deck, the more likely it will be chosen on a trial. However, this assumption is not always satisfied as participants' selections do not necessarily match the expectancies of each deck. Exploration of the task and risk-taking behaviours may lead to choices that do not follow the formed expectancies. In the next sections we will present different computational algorithms for each of the three assumptions and describe each model's core mechanisms.

5.1.1 Utility

5.1.1.1 Weighted utility function

The translation of the received wins and losses (if any) on trial t from deck j, into a single value utility $u_j(t)$, is governed by the following utility function:

$$u_j(t) = (1 - W) \cdot \mathrm{win}(t) + W \cdot \mathrm{loss}(t), \tag{5.1}$$

where win(t) and loss(t) represent the wins and losses on trial t respectively. The free parameter W is the *attention to losses* or *loss aversion* parameter which dictates how much attention is given to losses compared to wins. The value of the parameter ranges between 0 and 1 with higher values indicating more attention to losses (the middle point 0.5 assumes equal weighting of wins and losses). This function (also known as the expectancy utility function) assumes that evaluations of wins and losses are computed separately and then summed (Ahn et al., 2008; Busemeyer & Stout, 2002).

5.1.1.2 Prospect utility function

According to Ahn et al. (2008), the weighted linear utility function (Equation 5.1) cannot account for the gain-loss frequency effect observed in experience-based tasks. For example, the prospect of receiving -10 three times may seem worse for participants compared to receiving -30 just once, even though both events have the same expected value ($EV = -30$). In order to capture this tendency, Ahn et al. (2008) used the non-linear value function of Prospect Theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992):

$$u_j(t) = \begin{cases} x(t)^{\alpha} & \text{if } x(t) \geq 0 \\ -\lambda |x(t)|^{\alpha} & \text{if } x(t) < 0\,. \end{cases} \tag{5.2}$$

Equation 5.2 assumes that participants evaluate only the net outcome of wins and losses on trial t, $x(t) = \text{win}(t) - \text{loss}(t)$. This function has two free parameters: α $(0 \leq \alpha \leq 1)$ which defines the shape of the utility function, and λ $(0 \leq \lambda \leq 5)$ which is a loss aversion parameter. When α is close to zero then all positive net outcomes attain the same utility (i.e., 1) and negative net outcomes obtain the value of the loss aversion parameter (i.e., $-\lambda$). If α takes a value of 1, the utility of the positive outcome matches the objective value of the net outcome, $u_j(t) = x(t)$, and the negative outcome is weighted by λ. When α is between 0 and 1 then the shape of the utility function is curved. A value of λ larger than 1 indicates loss aversion, whereas values smaller than 1 suggest risk-seeking behaviour. As λ goes to zero, then losses are ignored in the estimation of subjective utility (Ahn et al., 2008; Ahn, Krawitz, Kim, Busemeyer, & Brown, 2011; Steingroever et al., in press).

5.1.1.3 Prospect utility with separate evaluation of payoffs

One problem with the prospect utility function (Equation 5.2) is that subjective utility is estimated based on the net outcome of each trial, $x(t)$. Thus, if participants receive the same amount of wins and losses on a given trial, the subjective utility will be zero regardless of the value of the parameters. However, this may not reflect what participants think, as losses can still loom larger than equivalent wins, leading to a negative subjective utility when the net outcome is zero. The latter is formulated in the following utility function, which maintains the main features of the prospect utility function (i.e., non-linearity and loss aversion) but assumes separate evaluation of wins and losses on each trial:

$$u_j(t) = [win(t)]^{\alpha} - \lambda[|loss|(t)]^{\alpha}. \tag{5.3}$$

5.1.2 Updating of expectancies

5.1.2.1 Delta learning rule

After the subjective utility $u_j(t)$ has been estimated, the expected utility of deck j on trial t, $E_j(t)$, is updated using the Delta learning rule, also known as the Rescorla-Wagner rule (Rescorla & Wagner, 1972):

$$E_j(t) = E_j(t-1) + A \cdot \delta_j(t) \cdot [u_j(t) - E_j(t-1)]. \tag{5.4}$$

This rule updates only the expectancies of the chosen deck j on trial t whereas the expectancies of the unselected decks remain unchanged. The dummy variable $\delta_j(t)$ takes a value of 1 if deck j is chosen on trial t and 0 otherwise. The free parameter A $(0 < A < 1)$ determines how much the old expectancy $E_j(t-1)$ is updated by the prediction error, $[u_j(t) - E_j(t-1)]$. Large values of A indicate strong recency effects and rapid forgetting (as $E_j(t-1)$ is fully modified by the prediction error, i.e., the new expectancy, $E_j(t)$, is strongly based on new payoffs/utilities), whereas values of A close to zero indicate weak recency effects and slow forgetting (Busemeyer & Stout, 2002; Yechiam & Busemeyer, 2005).

5.1.2.2 Decay reinforcement-learning rule

Another class of updating rules assumes that the expectancies of the unselected decks can decrease over time due to memory decay. The decay RL rule (Erev & Roth, 1998) takes this assumption into account:

$$E_j(t) = D \cdot E_j(t-1) + \delta_j(t) \cdot u_j(t). \tag{5.5}$$

Based on this rule, past expectancies are discounted and the expectancy of a chosen deck j on trial t is updated by its subjective utility $u_j(t)$. The expectancies of the unselected decks decay towards zero as a function of time. The recency (or decay) parameter D $(0 < D < 1)$ determines how much past expectancies are discounted. The interpretation of the D parameter is reversed in this rule as compared to the A parameter in the Delta learning rule: values close to 1 denote weak recency effects and less decay whereas values close to 0 indicate strong recency effects and rapid forgetting.

5.1.2.3 Mixed updating rule

One issue with the Decay RL rule is that the updated expectancy of the selected deck can be larger in value than both the previous expectancy, $E_j(t-1)$, and the subjective utility of the current outcome, $u_j(t)$. In other words, the expectancies formed by the Decay RL rule may not reflect what participants experience in the task. Imagine the following scenario: the expectancy of deck j on trial $t-1$ is $E_j(t-1) = 100$ and the utility from the same deck j on trial t is $u_j(t) = 150$. With a value of $D = 1$, the expectancy of deck j on trial t becomes $E_j(t) = 250$, which is not consistent with the payoffs that participants have received.

On the other hand, the updated expectancy of the Delta learning rule is constrained between the value of the previous expectancy and the utility of the most recent outcome. Dai, Kerestes, Upton, Busemeyer, and Stout (2014) proposed a new learning rule which implements the features of both learning rules:

$$E_j(t) = D \cdot E_j(t-1) + A \cdot \delta_j(t) \cdot [u_j(t) - D \cdot E_j(t-1)], \tag{5.6}$$

where D is the recency/decay parameter, A is the updating parameter, and $\delta_j(t)$ a dummy variable which denotes whether a deck was chosen on trial t.

5.1.3 Choice rule

After the formation of the expectancies, the next step involves making a choice between the decks, which is a probabilistic function of the expectancies of all decks. This is formulated using the softmax choice rule (or ratio-of-strength rule) which assumes that the probability of selecting a deck j on the next trial $t+1$, $Pr[D(t+1) = j]$, is proportional to the relative strength of this deck:

$$Pr[D(t+1) = j] = \frac{e^{\theta(t) \cdot E_j(t)}}{\sum_{k=1}^{4} e^{\theta(t) \cdot E_k(t)}}, \tag{5.7}$$

where θ is a sensitivity parameter which dictates the degree to which choice probabilities match the expectancies. When θ is zero, then choice behaviour is totally random ($Pr[D(t+1) = j] = 1/k$, where k is the number of decks) allowing for exploration behaviour. On the other hand, large values of θ allow for a complete match between expectancies and choice probabilities, indicating that the deck with the highest expectancy will be chosen more often.

The relationship between response sensitivity and time (i.e., trials in the task) can be trial-dependent or trial-independent. In other words, sensitivity to expectancies can increase over time or be independent of the trial number. A trial-dependent sensitivity is formulated with the following equation:

$$\theta(t) = (t/10)^c, \tag{5.8}$$

where c ($-5 < c < +5$) is a consistency parameter. If c is positive, choices over trials become more deterministic, matching the deck expectancies, whereas negative values of c indicate random choice between the decks.

The trial-independent sensitivity (Yechiam & Ert, 2007) assumes that the tendency to match the deck expectancies is constant across time and is given by the following equation:

$$\theta = 3^c - 1, \tag{5.9}$$

where c ranges between 0 and 5, with values close to 0 indicating random choice and values close to 5 suggesting deterministic choice.

5.1.4 Exemplar RL models

Following the previous section, a typical RL model in the context of the IGT consists of three different equations. Each of these equations captures one of the three general assumptions regarding decision-making in the IGT: utility, learning or updating of expectancies, and choice. Even though one can use any combination of these equations to arrive to a model, two specific combinations have been the most popular and extensively used in the literature: the Expectancy Valence Learning model (EVL; Busemeyer & Stout, 2002) and the Prospect Valence Learning model (PVL; Ahn et al., 2008, 2011). Table 5.1 shows the equations used by the EVL and PVL models. The EVL has been applied to numerous datasets to identify key components of decision-making processes and assess differences between healthy participants and clinical populations (e.g., Busemeyer & Stout, 2002; Kjome et al., 2010; Stout et al., 2004; Yechiam et al., 2005; Wood et al., 2005). After the PVL model was developed (Ahn et al., 2008), most studies started to use this model (e.g., Fridberg et al., 2010; Lorains et al., 2014). In some cases, hybrid versions of the two models provided a better description of the underlying psychological processes, such as the PVL-Delta model in which the decay RL rule of the PVL is replaced by the delta learning rule of the EVL (Fridberg et al., 2010; Steingroever, Wetzels, & Wagenmakers, 2013b). Table 5.1 also lists the PVL-PU2 model which was recently found to fit the IGT data better than its competitors (Dai et al., 2014). This model makes use of the prospect utility function with separate evaluation of wins and losses (second prospect utility function, see Equation 5.3; PU2) but the learning and sensitivity rules come from the PVL.

Table 5.1: Exemplar RL models for the IGT.

Model	Utility	Learning	Sensitivity
EVL	Weighted (5.1)	Delta (5.4)	Dep (5.7, 5.8)
PVL	Prospect (5.2)	Decay (5.5)	Ind (5.7, 5.9)
PVL-Delta	Prospect (5.2)	Delta (5.4)	Ind (5.7, 5.9)
PVL-PU2	Prospect 2 (5.3)	Decay (5.5)	Ind (5.7, 5.9)

Note: Numbers in parentheses indicate the equations in the text. Dep = trial-dependent, Ind = trial-independent.

5.2 Win-Stay/Lose-Shift Models

WSLS models make different assumptions from the RL models about decision-making processes in the IGT. The most simple form of the WSLS model assumes a heuristic strategy according to which participants stay with the same deck on trial $t+1$ if the net outcome on trial t is greater than 0, and switch randomly to one of the remaining three decks if the net outcome is negative. The WSLS has two parameters (Worthy, Hawthorne, & Otto, 2013). The first parameter defines the probability of staying with the same deck if the net outcome is equal to or greater than 0:

$$P[D(t+1) = j \mid choice(t) = x(D,t) \geq 0] = P(stay|win) \tag{5.10}$$

where $x(D,t)$ is the net outcome, $win(t) - loss(t)$, from deck D on trial t. The probability of switching after a trial with a positive net outcome is $P(switch|win) = 1 - P(stay|win)$ which is divided by 3 for each of the remaining three decks. In case of a negative net outcome, the probability of switching to any of the remaining decks, $P(shift|loss)$, is divided by 3 and assigned to the other three decks.

$$P[D(t+1) = j \mid choice(t) = x(D,t) < 0] = P(shift|loss) \tag{5.11}$$

The probability of staying with an option on a loss trial is $1 - P(shift|loss)$.

5.2.1 The Value-Plus-Perseveration (VPP) Model

The WSLS model introduced an important concept in experience-based decision-making: *perseveration* or *action inertia*. The probability of staying with the same option if the net outcome on a trial is positive can be seen as a manifestation of inertia, the tendency of participants to repeat their past selections (Erev & Haruvy, 2005). Similarly, the decay RL rule (Equation 5.5) incorporates inertia in the sense that the expectancies of the unselected options decay over time and hence the option which has been selected more often in the recent trials will be more likely to be selected as its expectancy will increase compared to the expectancies of the other decks. There are two problems, however, with the assumption of inertia in the WSLS model and the models which assume a decay RL rule: first, the WSLS model attributes no value to the expectancies of each deck (or their average value up to trial t) as participants make decisions only on the basis of the previous trial $t - 1$. However, this strategy assumes no learning of the task and participants surely have some expectations about the decks' payoffs when they select an option (Worthy, Pang, & Byrne, 2013). Second, the problem with the decay RL rule is that inertia/perseveration is confounded with the expectancy of each option. Since both of these dimensions are represented by a single numerical value, it is very difficult to ascertain which of these tendencies is responsible for the model's predictions.

Inertia is a manifestation of risk-taking behaviour (Dutt & Gonzalez, 2012b), and it has been implemented in various cognitive models of experience-based decision making (e.g., Biele, Erev, & Ert, 2009; Gonzalez & Dutt, 2011; Nevo & Erev, 2012). Worthy, Pang, and Byrne (2013) presented a model as an extension of the RL models in order to accommodate inertia. The Value-Plus-Perseveration (VPP) model incorporates two separate mechanisms to account for inertia and expected value as both of them are fundamental aspects of decision-making. The VPP model assumes three basic mechanisms for decision-making in the IGT as in the RL models described above: utility, learning, and choice. The only difference is that updating or learning of the expected values for each deck consists of two different components: expectancy and perseverance strength of each option. The idea is that participants assign values to an option which are dependent on the combination of the previous experience with that option (expectancy) and the tendency to stay with this option (perseveration). The VPP uses the prospect utility function (Equation 5.2) and the delta learning rule (Equation 5.4) for the estimation of utilities and expectancies, respectively. The following equation is used for the perseverance strengths of each option j on trial t, $P_j(t)$:

$$P_j(t) = \begin{cases} k \cdot P_j(t-1) + \delta_j(t) \cdot \epsilon_p & \text{if } x(t) \geq 0 \\ k \cdot P_j(t-1) + \delta_j(t) \cdot \epsilon_n & \text{if } x(t) < 0\,. \end{cases} \tag{5.12}$$

This rule adds 3 parameters to the model: k $(0 < k < 1)$ is a decay parameter similar to D in the decay RL rule, which dictates how much of the perseverance strength of each deck decays on each trial. ϵ_p and ϵ_n define the impact of positive net outcomes $(x(t) = win(t) - loss(t) \geq 0)$ and negative net outcomes $(x(t) = win(t) - loss(t) < 0)$, respectively. Positive values of ϵ_p and ϵ_n $(-1 < \epsilon_p, \epsilon_n < 1)$ reinforce the tendency to stay with an option on the next trial, whereas negative values reinforce switching to another option.

The overall expected value of a deck on a trial t, $V_j(t)$, is a weighted linear combination of its expectancy and perseverance strength:

$$V_j(t) = w_{Ej} \cdot E_j(t) + (1 - w_{Ej}) \cdot P_j(t). \tag{5.13}$$

Values of w greater than 0.5 indicate more weighting to the expectancy of each deck and values less than 0.5 indicate more weighting to the perseverative strength of each deck. A softmax rule (Equation 5.7) is used to define the probability of selecting a deck on the next trial, with a trial-independent sensitivity rule (Equation 5.9).

5.3 Model evaluation

In order to evaluate the candidate RL models, we constructed 18 different models from the factorial combination of 3 utility functions (Equations 5.1, 5.2, 5.3) $\times$ 3 learning rules (Equations 5.4, 5.5, 5.6) $\times$ 2 choices rules (Equations 5.8, 5.9). Using this design, we can identify which of all the combinations provides the best description of choice behaviour in the IGT. In addition, we applied the WSLS model and two versions of the VPP model: one with the Prospect Theory utility function and the second with a utility function based on Prospect Theory but with separate evaluation of wins and losses (PU2). We used two model comparison methods to contrast each model: the *one-step-ahead* prediction method (post hoc fit method) and a *simulation* method.

5.3.1 One-step-ahead prediction method

The one-step-ahead method investigates how well a model predicts the choice on the next trial, $t+1$, based on each individual's experience with the task up to and including trial t (i.e., previous sequence of choices, $Y_i(t)$, and payoffs associated with these choices, $X_i(t)$; i represents each individual). The model predicts the probability that an individual will select a deck D_j, $j \in \{A, B, C, D\}$, given his or her history of choices and payoffs, $Pr[D_j(t+1)|X_i(t), Y_i(t)]$. These probabilities are compared to the actual choices of each individual using the log likelihood (LL) criterion (Ahn et al., 2008; Busemeyer & Stout, 2002). Maximum-likelihood estimation was used to find the best parameter values for each model by maximising the log likelihood of an individual's sequence of choices:

$$LL_{M,i} = \sum_{t=1}^{t-1} \sum_{j=1}^{4} \ln(Pr[D_j(t+1) \mid X_i(t), Y_i(t)]) \cdot \delta_j(t+1), \tag{5.14}$$

where M indicates a model, n is the number of trials (in this case is $n-1$ because the first trial is excluded as choice is random), and $\delta_j(t+1)$ is a dummy variable which indicates whether a deck is chosen on trial $t+1$. A combination of grid-search with 60 different starting values for each parameter set and simplex search methods (Nelder & Mead,

1965) was used to find the best parameter values. The solution (i.e., the combination of parameter values) which maximised the LL across all starting points was selected.

The goodness of fit of each model was compared with a baseline statistical model which makes no assumptions about the underlying psychological processes. The Bernoulli baseline model assumes constant and independent probabilities across trials which reflect the overall proportion of selections from each deck. For example, if a participant's overall proportion of selections from deck A is 0.17, then the Bernoulli model assumes that the probability of selecting option A on each trial is 0.17. The baseline model is a good competitor to the cognitive models which can outperform it only if they can explain sequential dependencies and learning effects. It has three parameters, the probability of selecting from decks A, B, and C (p_A, p_B, p_C - estimated from each participant's choice behaviour); the probability of selecting from deck D is $p_D = 1-(p_A+p_B+p_C)$. To compare the models we used the Bayesian Information Criterion (BIC; G. Schwarz, 1978) which includes a penalty term for additional free parameters:

$$BIC = 2 \cdot (LL_M - LL_{Baseline}) - \Delta k \cdot \ln(n) \tag{5.15}$$

where Δk is the difference in the number of parameters between a cognitive model M and the baseline model, and n is the number of trials. A positive value of the BIC index indicates that a cognitive model performs better than the baseline model.

5.3.2 Simulation method

The simulation methods constitutes the least demanding model comparison technique because it is the test with the smallest difference from the original experiment (Steingroever et al., in press). This method evaluates each model's ability to reproduce the observed choice pattern under new payoff sequences by using the estimated parameters of each individual from one-step-ahead predictions (model fitting). The payoff schedule used by the simulation method remains the same as in the original experiment, but the ordering of wins and losses changes in each simulation (in the case of the IGT only the ordering of the losses differs as wins are constant for each deck). While the post hoc fit criterion uses the actual history of choices and payoffs of each participant, the simulation method attempts to predict the observed pattern by taking no account of participants' experience with the task. The predictions from the simulation are compared to the actual choice pattern using a deviance measure such as the mean square deviation (MSD):

$$MSD = \frac{1}{4 \cdot n} \sum_{t=1}^{n} \sum_{j=1}^{4} (\bar{D}_{exp,\ j}(t) - \bar{D}_{sim,\ j}(t))^2, \tag{5.16}$$

where n is the total number of trials, $\bar{D}_{exp,j}(t)$ is the mean proportion of choices from deck j on trial t across participants' experimental data, and $\bar{D}_{sim,j}(t)$ is the mean proportion of choices from deck j on trial t across all participants' simulated data. The smaller the value of the MSD, the better can a model reproduce participants' actual choices.

5.4 Model fitting

5.4.1 Datasets

We used three different datasets from experiments described in previous chapters to evaluate the candidate models. Specifically, we chose datasets in which the assessment of conscious knowledge was based on confidence ratings. We did not include any data from experiments where we used post-decision wagering as (1) the change in the payoffs after a wager may have had an effect on participants' choice behaviour, and (2) wagering was found to be a poor measure of awareness compared to confidence ratings. In addition, the differences between the three datasets (see Table 5.2) allow for a number of factors to be examined. First of all, they differ significantly in terms of overall task performance, even though performance in Dataset 3 may have been the result of the longer duration of the task (200 trials). Second, the 200 trials of Dataset 3 may reveal better relative fits as more trials can significantly affect model fitting, producing better predictions and more reliable parameter estimates.

Table 5.2: Description of the datasets used for model evaluations.

Dataset	Payoff Schedule	Trials	Performance	N
1: Experiment 6 - CR groups	IGT	100	52.70%	78
2: Experiment 7 - CR group	Appendix A	100	47.30%	18
3: Experiment 8 - Control Group	IGT	200	60.80%	20

Note: Performance indicates the overall proportion of choices from the good decks. CR = Confidence ratings; N = Sample size.

5.4.2 Model fitting results

First, we looked at the factorial combination of RL models to identify which of the utility, updating, and choice rules provides the best description of the experimental data. Table 5.3 shows the mean relative BIC scores (differences between a cognitive model and the baseline statistical model). Initial inspection of the table suggests that across datasets all the models outperform the baseline model, apart from the EVL model (EU, DEL, TD) in Dataset 2. These results indicate that all cognitive models were able to capture trial-by-trial dependencies and learning effects. Also there is a consistent pattern across datasets regarding the best model: a model which incorporates a prospect utility function with separate evaluation of wins and losses (PU2), a decay RL rule (DRL) and trial-independent sensitivity (TI) outperforms all the other models, even if the differences between the second and third best models are relatively small.

An interesting result relates to the EVL model which has been used in more than 35 studies to investigate the psychological processes underlying performance in the IGT. Importantly, most of these studies employed the EVL to assess differences in decision-making between healthy participants and clinical populations. Our results indicate that the EVL model's performance is very poor compared to the other models, suggesting that its parameters may not reflect the way that people make decisions in the IGT and thus many of the conclusions and inferences based on this model may be inaccurate. Future research should consider applying other models instead of the popular EVL.

The next step in our analysis was to collapse the mean BIC scores across utility,

Table 5.3: Summary of mean BIC scores of the 18 RL models relative to the baseline model in three datasets.

			Dataset 1		Dataset 2		Dataset 3	
Utility	Updating	Choice	M	%(BIC>0)	M	%(BIC>0)	M	%(BIC>0)
EU	DEL	TD	2.06	45	-0.61	28	11.20	55
		TI	3.35	55	0.62	33	12.85	65
	DRL	TD	10.38	68	16.01	67	81.59	75
		TI	**11.28**	78	**22.07**	72	75.88	85
	ML	TD	6.96	60	16.72	67	74.20	70
		TI	6.94	62	18.07	72	70.65	65
PU	DEL	TD	5.36	64	5.18	39	23.46	75
		TI	6.15	72	6.50	50	19.45	85
	DRL	TD	5.81	51	12.77	56	80.05	70
		TI	**11.84**	78	**20.81**	61	84.28	90
	ML	TD	7.91	62	15.44	56	77.37	90
		TI	7.96	69	18.12	61	79.65	85
PU2	DEL	TD	5.51	62	4.35	44	22.98	85
		TI	5.83	62	5.45	50	20.68	80
	DRL	TD	6.90	53	15.17	56	87.13	75
		TI	**12.92**	78	**22.52**	67	**89.39**	85
	ML	TD	9.71	64	18.34	56	**89.43**	90
		TI	9.12	72	19.78	56	**85.83**	90

Note: M = Mean relative BIC score; %(BIC>0) = Percentage of participants whose relative BIC score is greater than 0; EU = expectancy utility function; PU = prospect utility function; PU2 = alternative prospect utility function; DEL = delta learning rule; DRL = decay RL rule; ML = mixed learning rule; TD = trial-dependent sensitivity; TI = trial-independent sensitivity.

learning, and choice rules in order to identify which instantiation within each assumption provided the best fit. Table 5.4 shows consistent patterns across datasets regarding which instantiation of each assumption (i.e., utility, learning, and choice) should be used. Even though this demonstration does not take into account the interactions between assumptions, it highlights some of the results we observed by looking at the fit measure of each model in Table 5.3. Regarding the different utility rules, the alternative prospect utility function (PU2) outperforms its competitors no matter what learning and choice rules are used. This result is consistent across the three datasets. Looking at the learning rules, the delta learning rule is inferior to the other two learning rules which suggests that a learning rule which does not incorporate decay (and subsequently inertia or perseveration) may not be able to account for the sequential dependencies in the IGT. Finally, trial-independent sensitivity seems to provide better fits in Datasets 1 and 2 but not in

Dataset 3 in which the difference is small but in favour of trial-dependent sensitivity. This can be explained based on the number of trials (200) in Dataset 3: participants' sensitivity to the expectancies of each deck may increase across trials as they learn about the values of each deck and make more choices based on the learned expectancies.

Table 5.4: Summary of collapsed mean BIC scores for each utility, learning, and choice rule in three datasets.

		Dataset 1	Dataset 2	Dataset 3
Utility	EU	6.82	12.14	54.39
	PU	7.49	13.12	60.70
	PU2	8.31	14.25	65.90
Learning	DEL	4.70	3.58	18.43
	DEC	9.85	18.22	83.05
	ML	8.08	17.73	79.51
Choice	TD	6.72	11.48	60.82
	TI	8.37	14.87	59.85

Note: EU = expectancy utility function; PU = prospect utility function; PU2 = alternative prospect utility function; DEL = delta learning rule; DRL = decay RL rule; ML = mixed learning rule; TD = trial-dependent sensitivity; TI = trial-independent sensitivity.

We then focused on the "exemplar" models that have been proposed for the IGT. The models we compared were the two popular models, EVL and PVL, the PVL-Delta which has been suggested as a better alternative (Steingroever et al., in press; Steingroever, Wetzels, & Wagenmakers, 2013b), and the model we found to perform best across our datasets, the PVL-PU2 model. In addition, we included the WSLS model and two versions of the VPP model, one with the prospect utility function, and the second with the alternative prospect utility function (PU2). To assess the goodness of fit of the candidate models we used two measures in addition to the mean BIC difference between the cognitive models and the statistical baseline model: BIC weights and the proportion of participants best fit by each model. Schwarz (BIC) weights approximate the posterior probability of the models (assuming equal prior probability). In other words, BIC weights can be interpreted as reflecting the probability, given the observed data, that a candidate model is the best model in the set of models under consideration (Wagenmakers & Farrell, 2004). Table 5.5 contains the different fit measures for each candidate model. Focusing on the BIC difference score, we see that the PVL-PU2 performs best across datasets. The WSLS and VPP models do not provide better fits than the RL models tested here which indicate that the penalty term of the VPP is large given its 8 free parameters. When we look at the BIC weights a slightly different picture arises: in Datasets 1 and 3 the PVL-PU2 is the best model but in Dataset 2 the WSLS outperforms its competitors. The discrepancy between the BIC and BIC weights in Dataset 2 is due to the fact than when the WSLS fits best, it fits significantly better than the other models, resulting in weights close to 1. This is also evident in the proportion of participants best fit by the WSLS in Dataset 2: almost half of participants' behaviour can be best explained based on a win-stay/lose-shift strategy. Furthermore, the WSLS is the second best performing model in Dataset 1 as indicated by its BIC weight and the proportion of participants best fit by it. However, in Dataset 3 there is strong evidence in favour of the PVL-PU2 model in all three fit measures. Only 10% of participants are best fit by the WSLS which suggests that participants in Dataset 3 focused more on the learned expectancies of each deck rather than employing a heuristic strategy to make decisions. Overall, it seems that the PVL-PU2 and the WSLS explain choice patterns better than the other candidate models and the VPP pays the price of complexity as 8 parameters may be too much given the limited number of trials in the IGT (usually 100).

Table 5.5: Measures of fit for each candidate model across three datasets.

	Dataset 1			Dataset 2			Dataset 3		
Model	BIC	*w*BIC	%Best fit	BIC	*w*BIC	%Best fit	BIC	*w*BIC	%Best fit
EVL	2.06	0.12	0.14	-0.61	0.07	0.11	11.20	0.06	0.05
PVL	11.84	0.16	0.14	20.81	0.13	0.06	84.28	0.14	0.05
PVL-Delta	6.15	0.14	0.10	6.50	0.06	0.11	19.45	0.11	0.10
PVL-PU2	12.92	0.31	0.35	22.52	0.17	0.11	89.39	0.45	0.60
WSLS	5.44	0.21	0.22	17.41	0.39	0.44	64.95	0.12	0.10
VPP	1.36	0.01	0.01	19.36	0.10	0.11	81.76	0.08	0.05
VPP-PU2	2.90	0.04	0.04	19.89	0.09	0.06	84.18	0.05	0.05

Note: *w*BIC = BIC weights.

Figure 5.1 shows the mean predicted choice probabilities (smoothed with a moving average window of 7 trials) from each model as compared to the observed choice proportions (Data) in Dataset 3. Here, we only included the results from Dataset 3 as the same interpretation applies to the remaining two datasets (see Appendix C.1). A first glance at Figure 5.1 suggests that all models capture the pattern of the observed choices, meaning the rank order of the decks across trials. Then, three distinct patterns can be identified related to the learning or choice rule of each model. It is evident that the models which assume decay (PVL, PVL-PU2) and decay + perseveration (VPP) track quite closely the sequential choices in the observed data pattern. On the other hand, the models using the delta learning rule (EVL, PVL-Delta) do not provide accurate trial-by-trial tracking of the observed choices as well as the decay models. The EVL model seems to fit the data worse than any other model as it predicts that each deck has the same probability of being selected in the first 50 trials. In the case of the WSLS model, predicted choice proportions for each deck do not discriminate adequately between the decks. For example, the probability of selecting deck A is almost the same as the probability of selecting deck B, a pattern which is not observed in the original data. In addition, no model can predict the predominant preference for deck B (up to 60%) in the first 20 trials of the task even though the models which use a utility function based on Prospect Theory (all models apart from the EVL and WSLS) approximate this initial tendency.

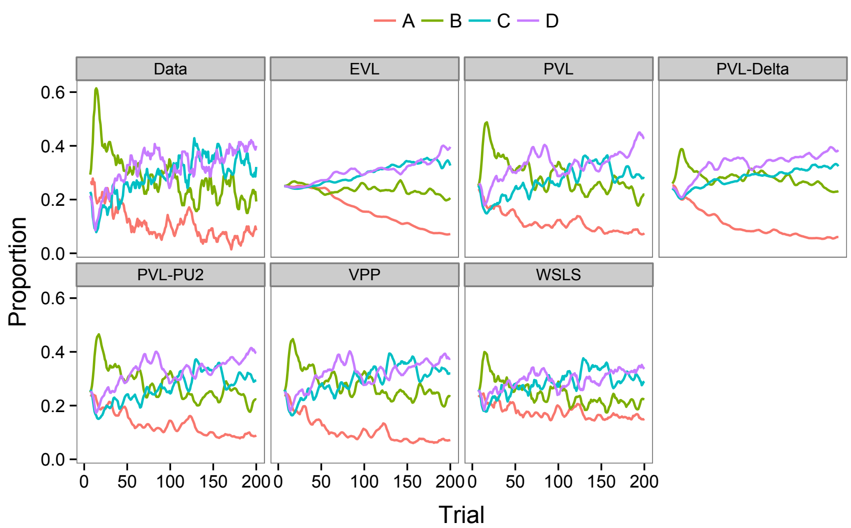

Figure 5.1: Mean predicted choice probabilities of each cognitive model as compared to experimental data in Dataset 3.

We then looked at the overall proportion of choices from each deck that the candidate models predict. This is a rough measure of goodness-of-fit as the candidate models are not evaluated on the overall proportion of choices but on their ability to predict inter-trial dependencies (the predictions of the statistical baseline match the overall proportion of choices from each deck). Figure 5.2 indicates that only the EVL and the WSLS models mispredict the overall proportion of choices from the bad decks. Specifically, both models assign higher overall probabilities to deck A compared to the observed proportion and underestimate the proportion of choices from deck B. The accuracy of the remaining models in predicting the overall pattern of choices is remarkably good and very close to the original proportions.

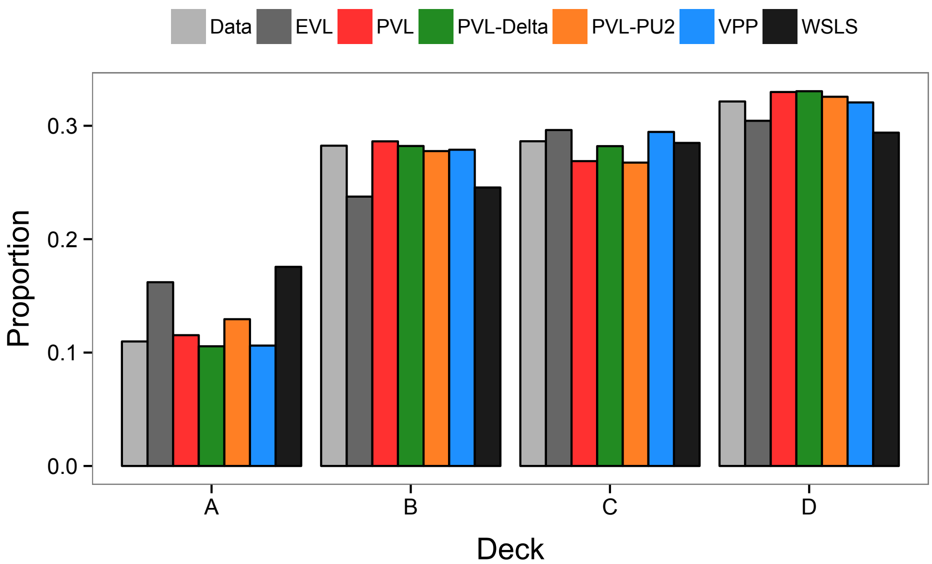

Figure 5.2: Overall proportion of predicted choices from each deck as compared to experimental data in Dataset 3.

5.4.3 Simulation performance

The second method we used to compare the adequacy and performance of each model was a simulation method (see Ahn et al., 2008; Steingroever et al., in press; Yechiam & Busemeyer, 2005, 2008). According to Steingroever et al. (in press) the simulation method is more indicative (compared to the post hoc fit method) of whether a model captures the psychological processes underlying the IGT (we will return to this issue in a later section of the present chapter). The simulation method uses the parameter estimates derived from one-step-ahead predictions to make a priori predictions for the entire sequence of choices. However, it does not take any feedback from participants' actual choices (and their intertrial dependency) but generates a choice (and its associated payoff) based on the model's predicted probabilities.

We generated 100 simulated datasets for each participant based on their best-fitting parameters according to each candidate model. These simulations were averaged across all participants to provide the predicted probability of choice on each trial. The predicted probabilities were then compared to the observed proportion of choices using the mean square deviation (MSD) and by visually inspecting models' predictions and observed choice behaviour.

Table 5.6 contains MSD scores for each candidate model across the three datasets. An inspection of the table suggests a different interpretation regarding which model should be preferred under the simulation method. Across datasets, the models which implement a delta learning rule (PVL-Delta and VPP) perform better under the simulation method. The PVL-PU2 and the WSLS models, which produced the best fits under one-step-ahead predictions, seem to not be able to match the observed choice proportions when no information about participants' actual choices are fed into the models. On the other hand, the PVL-Delta model, which produces worse post hoc fits (see Table 5.5), is able to reproduce the observed choice proportions by relying only on payoff-related feedback. This result demonstrates that the good performance of the decay models under one-step-ahead predictions may be due to their reliance on information about participants' past choices. Since the decks which are selected more often receive a "bonus" while the expectancies of the unselected decks decay towards zero over time, it can be inferred that decay models can track quite well people's choices based on observed choice histories. Consequently, if these models do not receive feedback from participants' actual choices and base their predictions only on the received payoffs, then their predictions can depart from the observed choice proportions. The discrepancy between one-step-ahead and simulation methods poses another challenge to the goal of model selection.

Table 5.6: MSD scores of each model between model predictions and experimental data in three datasets.

	Dataset 1	Dataset 2	Dataset 3
Model	MSD	MSD	MSD
EVL	108.56	222.42	147.72
PVL	42.40	136.77	147.57
PVL-Delta	38.38	123.08	92.21
PVL-PU2	48.22	148.77	146.57
VPP	40.32	136.38	91.78
WSLS	105.88	191.50	151.78

As can be seen in Figure 5.3 only the PVL-Delta and VPP models can closely reproduce the observed data pattern. Both models predict increasing preference for the good decks (C and D) after an exploration period and a steady decrease in selections from the bad decks (A and B). However, neither of these models can predict the initial preference for deck B. On the other hand, the PVL and PVL-PU2 models assume that more choices are made from the decks with infrequent losses (decks B and D) while they completely mispredict the preference for deck C. Even though the EVL model predicts more choices from the good decks, this almost appears after almost 2 blocks of trials in which predictions between decks are essentially random. Furthermore, it assumes that more choices are made from deck A compared to deck B. In the case of the WSLS model, choices between decks are constant and almost random across trials, indicating a failure to capture the observed pattern with no reliance on participants' actual history of choices.

Figure 5.4 shows the overall predicted proportion of choices from each deck as compared to the observed proportion. The best performing models under the simulation method (PVL-Delta and VPP) fail only to match the proportion of selections from deck B and this is due to the initial preference for this deck that is not captured by either model. The decay models fail to capture the observed tendency for more selections from deck C while all the models provide good predictions regarding choices from deck D. Both figures provide evidence against the decay RL and the WSLS models and highlight the inconsistency between the two model comparison methods (post hoc fit criterion and simulation method). We will discuss this discrepancy in more detail in the following section and explain why these two methods yield inconsistent results.

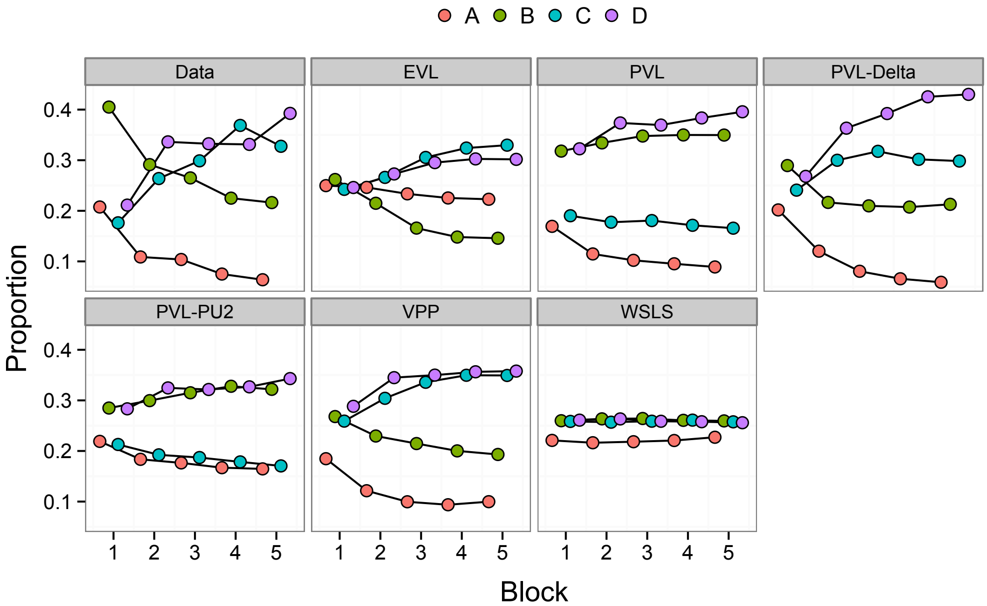

Figure 5.3: Simulation performance of the candidate models across blocks of trials with respect to the experimental data in Dataset 3.

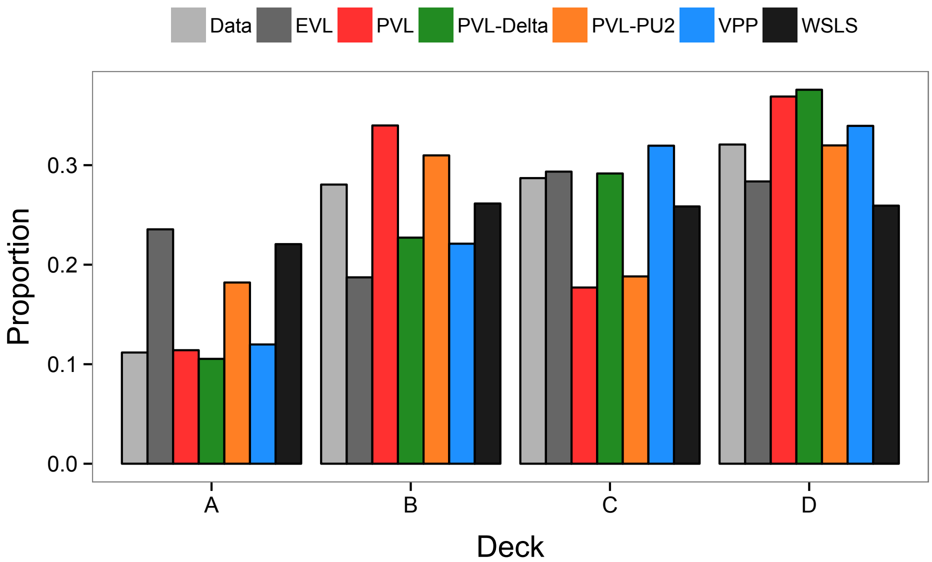

Figure 5.4: Overall proportion of choices from each deck under the simulation method in Dataset 3.

5.4.4 Parameter estimates

The ultimate objective of cognitive modelling is to decompose participants' performance into its constituent psychological and cognitive processes. The parameters of each model reflect these underlying processes, so the task of identifying a model which can provide descriptive and explanatory adequacy of the factors that drive performance becomes of primary importance. However, model comparison is not an easy challenge and the reliance on a statistical fit index (e.g., BIC) to draw conclusions about psychological mechanisms in a task may be inappropriate or at least insufficient. For example, we saw earlier (Table 5.5) that the performance of a large proportion of participants can be explained based on a simple win-stay/lose-shift strategy even though the BIC score of the WSLS model was smaller than that of the RL models. More importantly, the simulation performance of each candidate model revealed a new set of good models highlighting the discrepancy between the two selection methods employed here. Given these results, a researcher is left wondering which model should be used to characterise decision-making processes in the IGT and make inferences based on the estimated parameters. The task would be somewhat easier if the estimated parameters of each model were close to each other but this is often not the case, as Table 5.7 shows.

Table 5.7 contains the means and medians of the parameter estimates of four models across the three datasets. We did not include the EVL model, as it does not perform well in either of the two model comparison methods, or the WSLS model, as its parameters are different from those of the RL models and they do not invite a clear psychological interpretation. All models in the table assume a prospect utility function (PU or PU2) and trial independent sensitivity. Two of the models incorporate a decay RL learning rule (PVL and PVL-PU2) and the remaining two use a delta learning rule (PVL-Delta and VPP). In addition, the VPP model adds four extra parameters to decompose expected value and perseverative strength of each deck. Overall, the models can be categorised into two sets, decay and delta models, because the parameter estimates within each category are close. For example, in Dataset 2 the decay models (PVL and PVL-PU2) assume higher loss aversion ($\lambda = 1.35$ and 1.76, respectively) compared to the delta models (PVL-Delta and VPP, $\lambda = 0.85$ and 0.87, respectively). However, this pattern is reversed in Dataset 3, with decay models showing less weighting to losses ($\lambda = 0.69$ and 0.82) compared to the delta models ($\lambda = 1.25$ and 1.90).

With regards to the sensitivity parameter c the delta models assume that participants' choices are more deterministic (all values of $c > 2$) and consistent with the expectancies of each deck across datasets (higher values of c indicate that decks with higher expectancies are more likely to be selected) whereas decay models assume less deterministic behaviour (all values of $c < 1$). Similar differences between learning rules can be found in other parameter values such as the reward sensitivity parameter α. Given this inconsistency in parameter values across models, the conclusions about the underlying psychological mechanisms can be different depending on the model that one uses to explain performance in the IGT. For example, do people exhibit loss aversion behaviour or not? Similarly, do people's choices follow the learned expectancies or are their selections more random? These questions can take different answers depending on the parameters of the model that one chooses to use.

5.5 To simulate or not?

The inconsistency between the two model comparison methods, post hoc fit criterion and simulation, has been observed in many IGT model comparison studies (e.g., Ahn et al., 2008; Steingroever et al., in press; Yechiam & Busemeyer, 2008; Yechiam & Ert, 2007). Ahn et al. (2008) suggested that, given this discrepancy, one should select a model based on the objectives one wants to achieve. If the objective is to make short-term predictions and evaluate participants' decisions based on the experience they have had with the task (history of choices and associated payoffs) then the decay RL models should be chosen. On the other hand, if one wants to maximise long-term predictions (i.e., performance on a new gambling task that participants have never experienced) then models which implement a delta learning rule should be favoured.

Table 5.7: Parameter estimates from four different models across three datasets.

	Dataset 1		Dataset 2		Dataset 3	
PVL	M	Mdn	M	Mdn	M	Mdn
α	0.44	0.39	0.45	0.25	0.20	0.01
λ	1.44	0.96	1.35	0.61	0.69	0.01
D	0.74	0.89	0.61	0.63	0.73	0.79
c	0.38	0.22	0.46	0.32	0.53	0.61
PVL-Delta						
α	0.61	0.63	0.54	0.56	0.58	0.67
λ	1.47	0.91	0.85	0.33	1.25	0.88
A	0.22	0.04	0.31	0.10	0.08	0.02
c	2.59	1.80	1.39	1.02	2.72	2.21
PVL-PU2						
α	0.38	0.28	0.43	0.22	0.19	0.01
λ	1.96	1.31	1.76	0.86	0.82	0.56
D	0.73	0.87	0.58	0.62	0.68	0.75
c	0.38	0.26	0.46	0.35	0.55	0.57
VPP						
α	0.61	0.69	0.58	0.56	0.50	0.45
λ	1.59	0.81	0.87	0.14	1.90	1.37
A	0.17	0.04	0.32	0.16	0.27	0.02
k	0.43	0.38	0.37	0.33	0.55	0.61
e_p	0.20	0.16	0.17	0.19	0.57	0.79
e_n	-0.02	-0.09	-0.29	-0.42	0.20	0.42
V	0.72	0.84	0.49	0.44	0.72	0.82
c	2.33	2.05	1.69	1.70	2.02	1.95

Note: M = mean; Mdn = median.

Steingroever et al. (in press) suggested that "the simulation method seems to represent a more stringent and challenging test of absolute model performance than the post hoc absolute fit method because the simulation method relies on predicting the entire sequence of choices for another payoff sequence that could have been observed". They argue that the post hoc fit criterion is inferior in that it may only be able to reproduce the observed behavioural pattern by mimicking participants' previous choices, indicating that the estimated parameters may be biased by factors that do not invite a clear psychological interpretation. They also suggested that the models which show better simulation performance (e.g., PVL-Delta) should be preferred to decompose and analyse participants' choice behaviour and underlying psychological mechanisms.

In the remainder of this section, we argue against the idea that the inconsistency of model comparison methods is a failure (that is, one technique is better than the other - or more indicative - at disentangling the psychological processes that drive performance on the IGT). Instead, as we outline, such inconsistency in the results should be expected given the difference in approach of these two methods: The post hoc fit criterion uses participants' *previous history of choices and experienced payoffs* to make predictions about subsequent trials, whereas the simulation method relies on *the generation of new histories for choices and payoffs*. Second, while we agree that the simulation method is important, we argue that the post hoc fit criterion also has value in that its parameters are meaningful and provide useful psychological insight. Also, we outline and present theoretical and practical arguments for why relying solely on simulation performance in an experience-based decision-making setting may lead to misleading conclusions.

5.5.1 Sequential dependencies

The main objective of RL models of the IGT is to account for how future decisions are shaped by the experienced history of previous decisions and their associated payoffs. This is how learning occurs, and these models are evaluated on how well they capture learning effects throughout an individual's observed choice history. In other words, these models try to account for the sequential dependencies between each current choice and the previous choices and payoffs. In fact, all RL models share the assumption that the dependence of a current choice on previous choices is fully mediated by the payoffs (and not the past choices) experienced as a consequence of these previous choices. Thus, to predict a future choice, one only needs to consider the history of experienced payoffs.

In the simulation method, the parameters from the post hoc fit criterion, conditional upon observed choices and payoffs, are used to simulate a set of new choice and payoff sequences. Model selection is then based on the match between the marginal probabilities of the artificial choices in this simulated set and the proportions of observed choices. However, by marginalising over the simulated payoff sequences and focusing only on the choice probabilities, the method ignores the fundamental objective of RL models, which is to account for sequential dependencies between choices (i.e., how choices are shaped by experience). Furthermore, by considering their ability to reproduce (postdict) marginal probabilities in the *same* dataset as used for model estimation, the simulation method can favour models that are unlikely to generalise to new datasets. For example, consider a model with one parameter, which identifies a sequence of choices within the set of all possible choice sequences. The model predicts that a choice sequence displayed by a participant will be identical to the sequence identified by the parameter value. While, for a standard IGT with 100 trials and 4 choice alternatives, the parameter has 4^{100} possible values, estimation is easy once a participant's complete choice sequence is known. Because the model is not stochastic, it will reproduce, for each participant, their choice sequence exactly, regardless of the payoff sequence generated in the simulations. Averaging over participants, such simulations will perfectly reproduce the observed choice probabilities. Nevertheless, most people will agree that this is not an overly useful model: it predicts that people make exactly the same choices when performing the task another time and it can only *post*dict choices, not *pre*dict subsequent choices from the previous history. Because of this inability to make one-step-ahead predictions, the post hoc fit method, unlike a simulation method that estimates a model and evaluates simulation performance with the same data, will clearly disfavour this model.

5.5.2 Individual differences and research applied to clinical populations

The IGT has been used extensively as a neuropsychological test to assess decision-making in clinical populations. Thus, it is of great importance to derive a model that can offer informative conclusions regarding individual differences in the underlying psychological processes. This can enable researchers to identify key differences between clinical groups and healthy controls and make strong connections between neurophysiology and behaviour, leading to a better characterisation of the psychological symptoms of the disorder under investigation (Stout et al., 2004; Yechiam et al., 2005). A serious drawback of the simulation method is that it evaluates the average model prediction and ignores the fact that different models are required to fit different individuals. This was evident in the post hoc fit results where large proportions of participants were best fit by different models (see Table 5.5).

A fundamental goal of clinical research is to establish an explanatory framework for particular clinical groups, to connect pathological behaviour to patterns of behaviour on clinical and cognitive experimental tasks, and to map these deficits to neurophysiological mechanisms. Computational models serve as the intermediate step between the brain and observed behaviour where performance on a task can be decomposed into its constituent cognitive processes and mapped to neural mechanisms (Busemeyer et al., 2003). In this regard, model comparison should also be informed by psychological measures from relevant clinical assessments, which will facilitate the selection of the best models. In other words, a good model will not only perform well on a quantitative statistical index (e.g., goodness of fit) but will also provide explanations and make connections to findings or observations from the existing literature and to validated characteristics of a particular clinical sample. For example, if the model parameters are correlated with clinically relevant characteristics derived from psychometric scales and personality questionnaires, then the model can serve as a good representation of the underlying psychological processes. Also, another criterion to assess model performance is individual parameter consistency; that is, comparison of correlations of model parameters estimated from the same individual in more than one task. This is an important step toward the identification of "stable" internal characteristics that drive performance on different tasks and can be used as an extra assessment of model performance (Yechiam & Busemeyer, 2008). This stability of the internal characteristics guiding individual performance across somewhat disparate tasks helps to reinforce psychological explanations regarding underlying pathology in individuals.

5.5.3 Generalisation, biases, and inertia

One of the most fundamental properties of a good model is its ability "to make predictions about what will be observed in the future or generalisations about what would be observed under altered circumstances" (Shiffrin et al., 2008, p. 1249). Similar views have been expressed by other researchers who have discussed the importance of generalisation in identifying a good candidate model and who have introduced methods to assess a model's generalisability, such as the minimum description length (Pitt, Kim, & Myung, 2003; Pitt, Myung, & Zhang, 2002) and the generalization criterion (Busemeyer & Wang, 2000). If an (estimated) model is to reliably predict behaviour in new settings, it should apply independently of task-specific effects, biases that may arise from the experimental setup, and idiosyncratic strategies or heuristics individuals may adopt to deal with uncertainty. In experience-based tasks there are two sources of information that drive participants' choices: frequency of past choices from different options and payoffs experienced from sampling each of these options (Yechiam & Ert, 2007). A bias-free model (one that generalises efficiently) has to base its predictions only on past payoffs. We note that this is logical because the parameters of the RL models under investigation measure underlying psychological processes related to how participants respond to the payoffs they experience. However, we argue that the goal of predicting behaviour in different contexts is not well served by simulation methods.

Adding to this argument, Erev and Haruvy (2005) distinguish between two types of predictions in descriptive learning models: first, predictions for a task that are based on the interaction of a participant with that same task (*within-game* predictions), and second, predictions for different tasks with which participants are unfamiliar (*new-game* predictions). Following this distinction, the post hoc fit criterion is a within-game prediction, whereas the simulation method can be seen as a new-game prediction. Steingroever et al. (in press) argue that the simulation performance of a model should be of greater importance when our goal is to find the best model or assess the psychological processes underlying performance on the IGT. While the simulation method is indeed an important comparison technique, we present arguments why (a) simulation is a crude generalisation test, which (b) may not reflect stable psychological processes.

5.5.4 Not a direct generalisation test

Steingroever et al. (in press) suggest that the simulation method provides a good test of generalisability because it assesses models' predictions under new payoff sequences that participants have not encountered. To make these new predictions, the simulation method uses the best-fitting parameters from one-step-ahead predictions (post hoc fit criterion). Two possible problems may arise from the application of the simulation method: first, the use of parameters that optimise one-step-ahead predictions to predict new unobserved choice sequences (new-game predictions) might be far from ideal because these parameters carry information about participants' history of experience with the task and serial dependencies. In other words, the use of parameters estimated with the post hoc fit method could result in a considerable underestimation of a model's simulation performance. Future research can examine the benefits of the simulation method (i.e., model predictions under novel conditions/payoff sequences) by using simulated participants to estimate the model parameters at an individual or aggregate level (a common practice in other experience-based tasks, see Erev & Barron, 2005; Gonzalez & Dutt, 2011). Then, these parameters can be compared with those from the one-step-ahead method (and the model predictions of each method) in order to provide better inferences regarding the level of generalisability and utility of each model.

The second point is related to model complexity, a model's capacity to fit different patterns of data (Pitt & Myung, 2002). Complexity is dependent on two different factors that affect model fit: the free parameters of the model and its functional form; that is, how the parameters and the mathematical equations are combined (Myung, 2000). Complexity is a very important concept in model comparison and is related to the generalisability of the model. The simulation method does not take into account either of the two factors described earlier. For example, the EVL model has 3 free parameters, whereas the PVL and the PVL-Delta have 4. Even though the differences in simulation performance between the PVL and PVL-Delta models cannot be ascribed to differences in the number of parameters or their functional form (there are qualitative differences in the patterns of choices that the two models predict), it could be the case that more (or less) complex models can be applied to IGT data such as the VPP model which has 8 parameters.

The question is how does one deal with complexity in model comparison? While there are a few techniques to capture complexity (see Pitt et al., 2002), we propose the use of the generalization criterion (GC) method (Busemeyer & Wang, 2000). The reason is that the GC is a better and more sophisticated modification of the cross-validation method and is sensitive to both factors of complexity (Pitt & Myung, 2002). Also, it has already been applied for model comparison in experience-based tasks (see Ahn et al., 2008; Yechiam & Busemeyer, 2008). A critical difference between the GC and the simulation method is that the former uses the participants' actual choice and payoff histories (i.e., there is intermediate feedback from participants' choices), whereas the latter takes no input from what participants have actually selected and observed. Hence, when we assess model generalisability, the GC is a more appropriate method because it deals with the concept of model complexity and also takes into account how participants' future choices are dependent on their actual choice and payoff histories.

The application of the criterion is straightforward: there are two experimental conditions or datasets. The first serves as the calibration set and the second as the test set. A model is fitted on the calibration set, and the estimated parameters are used to predict the data in the test set. Finally, the predictions of each of the candidate models in the test set are compared to assess the empirical validity of these models. Ahn et al. (2008) used the GC to evaluate different RL models in the context of experience-based decision-making. Specifically, the same group of participants completed two tasks, the IGT and the Soochow gambling task (SGT; Chiu et al., 2008), and the models were compared on their ability to predict one task's choices based on the estimated parameters of the other task. Using the GC, Ahn et al. found that the same pattern of results emerged as in the one-step-ahead prediction method; that is, the PVL model with the decay RL rule was the best-performing model (but see Yechiam & Busemeyer, 2008).

One disadvantage of the use of the GC method is that it "places very demanding requirements on the researcher" (Busemeyer & Wang, 2000, p. 187), because it requires that the same participants perform a second independent task (similar to the first) based on a new experimental design (e.g., different payoff schedule). This is the reason why we did not apply the GC method in our model comparison assessment as participants in our experiments only performed the IGT.

5.6 Equal Payoff Series Extraction (EPSE)

In this section we return to the inconsistency in results between the post hoc fit criterion and the simulation method and we attempt to explain the reasons for this inconsistency. As previously mentioned, RL models derive one-step-ahead predictions from two different sources of information: past history of choices and past payoffs associated with choices. When models make predictions, they do not equally weigh these two sources of information: it is a possibility that certain models rely more on past choices to make accurate predictions while other models rely more on past payoffs. For this reason, the performance of models which rely least on past choices to predict the next choice will be superior under a method which does not take into account participants' actual choices such as the simulation method. For example, the decay models tested here may strongly base their predictions on past choices, because the expectancies of the unselected decks decay towards zero across time. This could be the reason why the decay models show poor simulation performance.

Yechiam and Ert (2007) addressed the issue of reliance on past choices by comparing two models in three different experience-based tasks (none of which was the IGT): the EVL and the EVL with a decay RL learning rule. They found that the model with the decay RL rule relied strongly on past choices which led to poor generalisability and inferior predictions under new payoff conditions. Here, we try to re-address the issue of reliance on past choices by extending Yechiam and Ert's analysis to include 2 additional utility functions (Prospect Theory utility functions: PU and PU2) and 1 additional learning rule (mixed learning rule: ML). We also tested the WSLS and VPP models. The technique we used was identical to Yechiam and Ert's *equal payoff series extraction* method (EPSE). The EPSE uses artificial participants to assess the degree of reliance on past choices. For each "real" participant performing the task, there is an artificial one whose choices are identical to those of the real participant but whose payoffs are made to be equal for all options/decks, matching the average amount of wins and losses of the real participant across all trials. Hence, if the models do not base their predictions on past choices, then they should not be able to predict an artificial participant's future choices under the EPSE method, because there is no discrimination in the received payoffs.

Each model's parameters were estimated for each artificial participant using the one-step-ahead prediction method and the measure fits (BIC) were compared to the fits of the real participants. Following this, we fit again all 18 RL models (i.e., the factorial combination of 3 utility functions, 3 learning rules, and 2 choice sensitivity rules), plus the WSLS and VPP models. While the results may reveal differences mainly between the learning rules (as in Yechiam and Ert's, 2007, study), it is possible that the reliance of each learning rule on past choices may depend on its interaction with different utility functions and choice rules.

Table 5.8 contains the mean relative BIC ($= BIC_M - BIC_{Baseline}$) for each model across the three datasets. Positive values indicate better fits of the cognitive model, while negative values suggest that the statistical baseline model provided better fits. Two main patterns can be observed: first, the delta learning rule (DEL) has poorer fits under the EPSE method no matter what utility function or choice rule it is combined with. The effect is more pronounced when it is combined with the prospect utility function (PU). The PU function produces worse fits than any other utility function while the expectancy utility (EU) and the alternative prospect utility (PU2) function perform better in the EPSE method than in the original post hoc fit when they are combined with a decay RL learning rule (DEC) and trial-dependent sensitivity (TD). This result indicates that these models can rely only on past choices to make predictions about participants' future choices and that information about the actual received payoffs does not improve their predictions. While moderate reliance on past choices can be beneficial for RL models to adjust their predictions and track participants' actual choices, strong or absolute dependency on the history of choices (or absence of discriminative ability between the decks based on their associated payoffs) could make such models inappropriate to measure stable internal characteristics of each individual. Instead, these models may only reflect mimicry of past choices.

Table 5.8: Summary of BIC scores of the 18 cognitive models relative to the baseline model under the post hoc fit and EPSE methods in three datasets.

			Dataset 1		Dataset 2		Dataset 3	
Utility	Updating	Choice	Orig. Fit	EPSE	Orig. Fit	EPSE	Orig. Fit	EPSE
EU	DEL	TD	2.06	-1.64	-0.61	-4.63	11.20	6.92
		TI	3.35	-3.84	0.62	-4.76	12.85	4.03
	DRL	TD	10.38	12.01	16.01	23.73	81.59	93.04
		TI	11.28	-3.84	22.07	21.00	75.88	88.39
	ML	TD	6.96	6.90	16.72	18.95	74.20	86.69
		TI	6.94	5.13	18.07	17.54	70.65	84.40
PU	DEL	TD	5.36	-13.36	5.18	-21.68	23.46	-20.09
		TI	6.15	-15.34	6.50	-22.11	19.45	-22.04
	DRL	TD	5.81	-7.12	12.77	-11.91	80.05	47.73
		TI	11.84	-11.05	20.81	-13.62	84.28	40.68
	ML	TD	7.91	-11.40	15.44	-15.95	77.37	42.76
		TI	7.96	-13.98	18.12	-17.63	79.65	36.99
PU2	DEL	TD	5.51	-6.09	4.35	-9.08	22.98	2.53
		TI	5.83	-8.44	5.45	-9.37	20.68	-1.26
	DRL	TD	6.90	7.45	15.17	19.67	87.13	88.87
		TI	12.92	4.28	22.52	16.40	89.39	83.09
	ML	TD	9.71	2.98	18.34	14.79	89.43	83.85
		TI	9.12	0.57	19.78	12.94	85.83	79.11

Note: EU = expectancy utility function; PU = prospect utility function; PU2 = alternative prospect utility function; DEL = delta learning rule; DRL = decay RL rule; ML = mixed learning rule; TD = trial-dependent sensitivity; TI = trial-independent sensitivity; Orig. Fit = Original fit, Fit based on received payoffs; EPSE = Equal payoff series extraction, Fit based on average payoffs.

We then collapsed the mean relative BIC measures for each utility, learning and choice rule under both methods (original fit and EPSE) and calculated their numerical difference across datasets (Table 5.9). High positive difference scores (Diff. column in Table 5.9) indicate worse fits under the EPSE compared to the original fit method, suggesting less reliance on past choices and weak choice mimicry effects. Regarding utility functions, the prospect utility function is strongly affected by the lack of information about real payoffs whereas the expectancy utility function can make accurate predictions when discriminative information about payoffs is eliminated. The latter does not mean that different utility functions reveal different levels of reliance on past choices but rather when they are combined with learning rules, their interaction produces weak or strong reliance on past choices. As expected, the decay learning rule produced the smallest difference score suggesting that its predictions relied more on previous choices compared to its competitors. This is in accordance with Yechiam and Ert's (2007) analysis in which they showed that the decay RL learning rule produced superior post hoc fits but poor simulation performance because the component in the post hoc fit that was dependent on past payoffs was smaller compared to the delta learning rule. Comparison of the choice sensitivity rules revealed that models which use the trial-dependent sensitivity rule tend to rely more on previous choices than models with the trial-independent sensitivity.

Table 5.9: Summary of collapsed mean BIC scores for each utility, learning, and choice rules under the post hoc fit and EPSE methods in three datasets.

		Dataset 1			Dataset 2			Dataset 3		
		Orig. Fit	EPSE	Diff.	Orig. Fit	EPSE	Diff.	Orig. Fit	EPSE	Diff.
Utility	EU	6.82	4.65	2.18	12.14	11.97	0.17	54.39	60.58	-6.19
	PU	7.49	-12.04	19.53	13.12	-17.15	30.27	60.70	21.01	39.70
	PU2	8.32	0.13	8.19	14.25	7.56	6.70	65.90	56.03	9.87
Learning	DEL	4.70	-8.12	12.82	3.58	-11.94	15.52	18.43	-4.98	23.42
	DEC	9.85	2.49	7.36	18.22	9.21	9.01	83.05	73.64	9.41
	ML	8.08	-1.63	9.72	17.73	5.11	12.62	79.51	68.97	10.54
Choice	TD	6.72	-1.14	7.86	11.48	1.54	9.93	60.82	48.03	12.78
	TI	8.37	-3.70	12.07	14.87	0.04	14.83	59.85	43.71	16.14

The results in Table 5.10 explain why the PVL-Delta model manifested good simulation performance, as it has the worst fit in the EPSE method across the three datasets. This indicates that PVL-Delta's reliance on past choices is minimal which explains its good performance in methods of generalisability (simulation method). On the other hand, the PVL-PU2 seems to base its predictions almost exclusively on the history of past choices as its fit under the EPSE method is only slightly worse than the original fit.

The same also applies to the WSLS model which can reproduce the observed choices without any access to information about received payoffs. An interesting finding of the present analysis relates to the pattern of results of the PVL and VPP models under both methods. While the PVL model has the second worst fit under the EPSE method, its simulation performance is no better than that of the VPP model which under the EPSE method seems to perform quite well. Based on the results of Table 5.10, one would expect a close relationship between performance in the simulation and the EPSE methods: the worse the fit under the EPSE method, the better the performance under the simulation method. However, this is not observed in the case of the VPP model. It performs well under both methods which may be due to the fact that it divides the expectancy of each deck into two different components: expected value and perseverative strength. Having this extra manipulation, the VPP model can adaptively predict choices under any model comparison method.

Table 5.10: Mean relative BIC scores for each candidate model under the post hoc fit and EPSE methods across three datasets.

	Dataset 1		Dataset 2		Dataset 3	
Model	Orig. Fit	EPSE	Orig. Fit	EPSE	Orig. Fit	EPSE
EVL	2.06	-1.64	-0.61	-4.63	11.20	6.92
PVL	11.83	-11.05	20.80	-13.62	84.28	40.68
PVL-Delta	6.14	-15.34	6.49	-22.12	19.45	-22.04
PVL-PU2	12.91	4.28	22.51	16.40	89.38	83.09
WSLS	5.45	2.43	17.42	13.56	64.96	60.57
VPP	1.36	-9.93	19.31	5.13	81.73	67.64

5.7 Model and parameter recovery

5.7.1 Model recovery

In order to test the general identifiability of each model and how flexible they can be in mimicking the data generated by a competing model, we conducted a model recovery analysis (for similar approaches see Berry, Shanks, Speekenbrink, & Henson, 2012; Steyvers et al., 2009; Wagenmakers, Ratcliff, Gomez, & Iverson, 2004; Worthy, Hawthorne, & Otto, 2013). Model mimicry is an important concept in model comparison studies as it assesses the degree to which a candidate model can be excessively adaptive and flexible in explaining data patterns that have been produced by a different model. This is best explained by using two different non-nested models: model A and model B. Assuming that model A is the "true" model which has generated the data, then it follows that this model should be able to provide the best fits for the data it generated (it successfully recovers the data that it produced) as compared to its competitor, model B. If, however, model B gives better fits than model A (given that model A is the true model), then model B is an overly flexible model which can mimic a wide range of data patterns. Model recovery analysis in not only useful in assessing a model's flexibility but it also provides a good means to ensure that the model selection measurement (e.g., BIC) is valid at discriminating and selecting between competing models.

In the present model recovery analysis, we compared 5 different models (EVL, PVL, PVL-Delta, PVL-PU2, and VPP). We used the best fitting parameters of each participant under each model to generate 100 simulated datasets. The next step is to fit the simulated datasets with all the models, including the one which generated the data. This procedure is applied to each model. Then, each model is assessed on how well it can recover the datasets that it generated. The procedure we followed is *local* or *data informed* as it uses real experimental data to assess model recovery (see Wagenmakers et al., 2004). The *global* approach is to generate simulated datasets based on random sampling from the parameter space (i.e., from within the boundaries of parameter values) of each model, which can go a step beyond idiosyncrasies found in experimental data and provide an unbiased assessment of model recovery and mimicry. However, the global approach risks basing conclusions about model recovery on implausible data patterns (i.e., data patterns that have not been observed in the literature).

We used the best fitting parameters from Dataset 3 which resulted in 2,000 simulated datasets for each model (20 participants $\times$ 100 replications). The assessment of each model's recovery ability is performed using confusion matrices, that is the percentage of correctly recovered models for simulated datasets. Table 5.11 contains two measures of model recovery. The first is the percentage of recovered datasets of the model that generated the data. Perfect recovery is achieved when the diagonal of the matrix is 100, which means that each model is able to fully recover the datasets that it generated. In the present analysis, this is only true for the EVL model as it recovers 97.65% of the datasets it produced. The VPP and the PVL models recover less than half of their generated datasets, indicating that other models can mimic the data patterns they produce. Specifically, the PVL-PU2 model is highly flexible and can explain a substantial percentage of datasets generated by the PVL and VPP models (31.18% and 43.97%, respectively). The poor model recovery of the VPP can also be ascribed to its complexity as it has 8 parameters and the penalty term is higher for this model. Overall, this analysis has two main conclusions: first, the EVL model is quite distinct from the other models and the data patterns it can explain. Second, recovery performance is not clear between the PVL, PVL-Delta and PVL-PU2 models because their assumptions are quite similar leading to significant predictive overlap.

When looking at the mean BIC of each model (second part of Table 5.11) a clearer picture emerges. Numbers in parentheses indicate the rank of the BIC values so a rank of 1 suggests that the model had the best BIC value for the datasets it generated. Apart from the VPP model, all the other models provided the best BIC for the datasets they generated indicating good recovery performance and also that the BIC is an adequate measure for selecting between models. It can also be seen that the PVL and PVL-PU2 models show similar behaviour and ability to fit data that were generated by each other because their mean BIC values are very close. However, the PVL-PU2 seems to be more flexible and can fit better datasets that were generated from the PVL model.

5.7.2 Parameter Recovery

The final step in cognitive modelling analysis is to use the estimated parameters of the "best" candidate model to make inferences about the underlying psychological and cognitive processes. Hence, it is important to examine whether these parameters can be

Table 5.11: Model recovery performance of each candidate model.

	Percentage of recovered simulated datasets				
True Model	EVL	PVL	PVL-Delta	PVL-PU2	VPP
EVL	**97.65**	0.40	0.75	1.20	0
PVL	8.56	**45.75**	14.46	31.18	0.05
PVL-Delta	13.20	15.70	**57.35**	13.55	0.02
PVL-PU2	9.50	9.40	5.35	**75.75**	0
VPP	6.60	15.75	10.66	43.97	**23.01**
	Mean BIC				
EVL	**467.76 (1)**	487.37 (3)	490.10 (4)	481.86 (2)	496.44 (5)
PVL	450.11 (5)	**375.35 (1)**	421.95 (4)	377.41 (2)	395.51 (3)
PVL-Delta	447.70 (4)	440.00 (2)	**437.01 (1)**	443.13 (3)	454.07 (5)
PVL-PU2	488.36 (5)	397.07 (2)	475.04 (4)	**387.48 (1)**	406.91 (3)
VPP	450.08 (5)	393.51 (3)	435.53 (4)	386.39 (1)	**390.93 (2)**

estimated accurately and reliably (Wetzels, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2010). In other words, the question is how well a certain model can recover the parameters that were used to generate simulated datasets. The procedure is identical to the model recovery analysis: the best fitting parameters of each participant under each model are used to generate 100 simulated datasets. Then, these datasets are fitted again using the same model and the estimated parameters of the simulated datasets are compared to the parameters that were used to generate the datasets in the first place.

Here, we will focus on the PVL-PU2 and the PVL-Delta models which showed the best performance under the post hoc fit criterion and the simulation method, respectively. Figure 5.5 shows recovery performance for each parameter of the PVL-PU2. Red dots indicate the parameter values that were used to generate the simulated datasets (for each participant), and the black dots indicate the mean parameter values of the 100 fitted simulated datasets. When black dots are not visible, it means that there is an overlap of the true and the simulated parameter values, suggesting good recovery performance. Participants (x-axis) are ranked on the basis of their parameter values. The solid horizontal line indicates the overall mean of the true parameter value, and the dashed horizontal line indicates the overall mean of the simulated parameter value. Figure 5.5 shows good recovery for the recency, D, and response consistency, c, parameters. Recovery for these two parameters is unbiased, as the simulated parameters closely match the true parameter values for each participant, and the overall means (horizontal lines) are almost identical. Similarly, the loss aversion parameter, λ, shows relatively good recovery. However, parameter recovery for the reward sensitivity parameter, α, was not as good as for the other parameters of the model. Not only are the parameters not well estimated, but the density plots were quite broad, indicating increased variability in the estimation of this parameter. One possible explanation of this pattern is that the method for estimating the parameter values (i.e., maximum-likelihood estimation, MLE) is inadequate. The true parameter value of the first 14 ranked participants is virtually zero which is the lower bound of this parameter ($0 < \alpha < 1$). This is common problem with maximum likelihood estimation, as estimated parameters "hit" the boundaries of the parameter space (Wetzels et al., 2010). To avoid this, other estimation techniques could be employed instead of MLE, such as hierarchical Bayesian estimation which overcomes such problems (see Ahn et al., 2011; Steingroever et al., in press).

Figure 5.6 shows parameter recovery performance of the PVL-Delta model. The payoff sensitivity parameter shows better recovery compared to the PVL-PU2 model, whereas the loss aversion parameter is recovered as well as in the PVL-PU2. On the other hand, the response consistency parameter shows increased variability (broad density plots) and poor recovery, as was the case for the PVL-PU2 model. Overall, the two models indicate relatively good parameter recovery suggesting that their parameters can be used to explain the psychological phenomena underlying the IGT. There might be a potential problem with the payoff sensitivity parameter in the PVL-PU2 model, which could potentially be overcome using a Bayesian estimation framework.

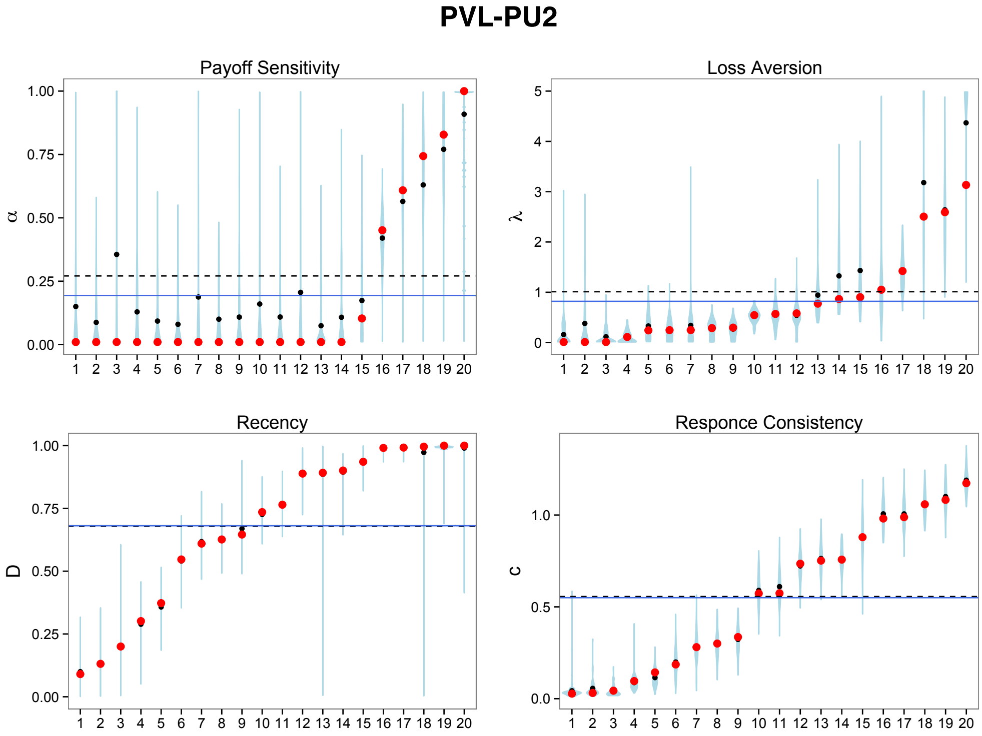

Figure 5.5: Parameter recovery performance of the PVL-PU2 model in Dataset 3. Red and black dots indicate true parameter values and means of the parameter values for the simulated datasets, respectively. Solid and dashed horizontal lines indicate the overall mean of the true and simulated parameter values, respectively.

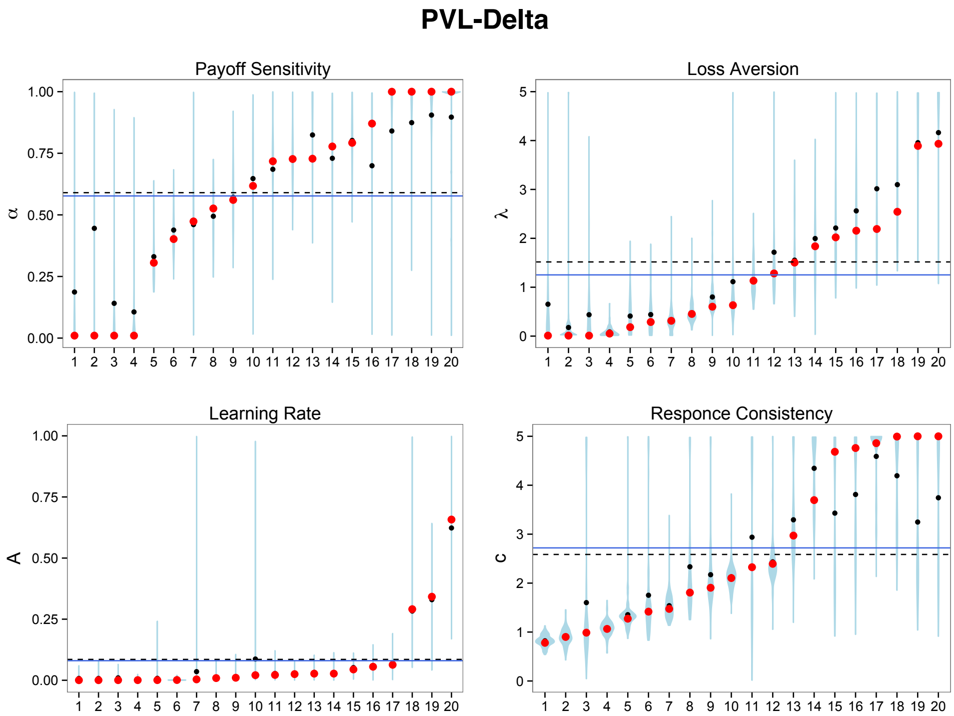

Figure 5.6: Parameter recovery performance of the PVL-Delta model in Dataset 3. Red and black dots indicate true parameter values and means of the parameter values for the simulated datasets, respectively. Solid and dashed horizontal lines indicate the overall mean of the true and simulated parameter values, respectively.

A potential issue with the previous analysis is that the number of trials (200 in Dataset 3) may have led to better parameter estimates, as the models have more information to adjust their parameter values. We applied the same analysis in Dataset 2 which consists of 100 trials (typical of the IGT) and we observed similar recovery performance. Figure 5.7 shows parameter recovery performance of the PVL-PU2 model in Dataset 2, which is quite similar to that observed in Dataset 1. Again, there are some problems with the payoff sensitivity parameter but overall the number of trial seems not to be a major factor in determining a model's parameter recovery performance.

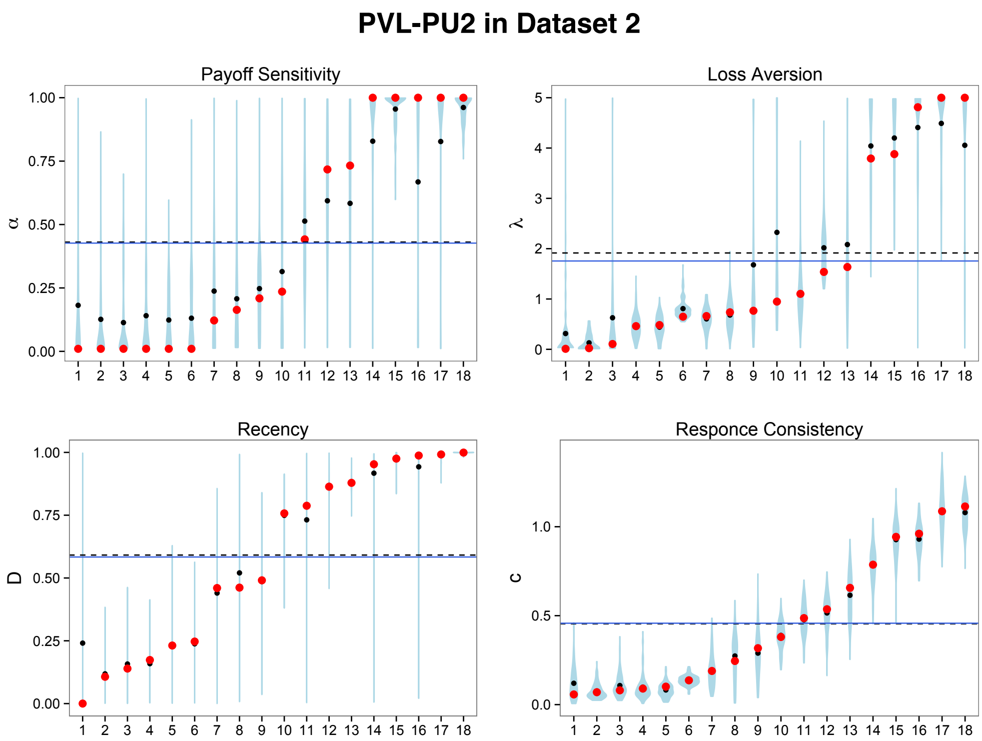

Figure 5.7: Parameter recovery performance of the PVL-PU2 model in Dataset 2. Red and black dots indicate true parameter values and means of the parameter values for the simulated datasets, respectively. Solid and dashed horizontal lines indicate the overall mean of the true and simulated parameter values, respectively.

6 Modelling Confidence

The relationship between decision accuracy, choice certainty or confidence, and reaction times has been extensively investigated since the early days of psychological research (e.g., Henmon, 1911; Volkmann, 1934). These early investigations showed that all three components of a decision process are closely related, for example, reaction times become faster as confidence about one's decision increases (also see Petrusic & Baranski, 2003), and choice accuracy is positively correlated with decision confidence (e.g., Baranski & Petrusic, 1998; Dougherty, 2001). Most of these studies examined psychological effects and decision processes in the domain of perceptual decision-making. However, little is known about the interplay of decision accuracy and confidence in sequential economic decisions as in the case of experience-based decision-making. Importantly, behaviour in decisions-from-experience is shaped by learning mechanisms and subsequent choices are dependent on the previous history of choices and associated payoffs. Hence, the level of confidence after each selection in experience-based tasks does not necessarily constitute a momentary evaluation of the quality of the selection, but instead reflects the accumulated experience (or learning) with this selection across time.

The purpose of the present chapter is to provide a formal description of confidence judgments in the context of the IGT. The results of previous experiments (e.g., Experiment 8) suggested a close link between choice behaviour and confidence ratings, meaning that decks with positive overall expected values (i.e., good decks) received higher confidence ratings compared to decks with negative overall expected values (i.e., bad decks). In addition, ROC analyses revealed good metacognitive discriminability indicating that good decisions were made under conditions of increased confidence. Despite the usefulness of these analyses, they can only be seen as *post hoc* or descriptive assessments of the relationship between choice and confidence. In other words, they do not inform us about the process that generates a certain confidence judgment or the psychological mechanism based on which participants assign confidence ratings to their first order decisions (i.e., deck selections).

In order to provide an explanatory framework for confidence judgments in the IGT, we developed a computational model that is dependent on the choice predictions generated by the RL models. Given the close relationship between choice and confidence, we hypothesised that there is a connection between the output of the RL models that we tested in the previous chapter and the mechanism that gives rise to confidence judgments. The results of the choice modelling analyses suggested that certain RL models can capture the dynamics of choice in the IGT and participants' preferences about the decks. Hence, if the same mechanism is responsible for both choice and confidence, then there should be a direct mapping between the predictions of RL models and confidence ratings.

6.1 The model

We constructed a model which allows for a direct mapping of the predictions of a certain RL model onto confidence ratings. RL models predict the probability of a deck j being selected on the next trial $t + 1$, $Pr[D(t + 1) = j]$, given the history of previous choices and associated payoffs. This probability reflects the relative expectancy of a deck, that is, its expectancy (or average adjusted expected value up to and including trial t) compared to the expectancies of the other decks. We used these probabilities as the relative strength of each deck upon which a confidence rating is based. The rationale behind our confidence model is that decks with high relative strength will receive high confidence ratings. However, the relationship between relative strength and confidence may not always be deterministic, thus, in order to account for uncertainty and variability in confidence judgments, we employed a normal distribution function over each selected deck's relative strength. This is graphically illustrated in Figure 6.1: the mean of this normal distribution is $\mu = Pr[D(t) = j]$ and the variance, σ^2, is a free parameter estimated from the data. The next parameter of the model is the position of the middle criterion, C_{2-3}, whereas the positions of the other two criteria, C_{1-2} and C_{3-4}, can be estimated given the last parameter of the model, d, which defines the constant distance between criteria. The probability of a confidence rating $C(t) = i$ on trial t, following selection of deck j and its relative strength is given as:

$$\begin{aligned} Pr(C(t) = i \mid Pr[D(t) = j]) = \Phi & \left(\frac{c_i - Pr[D(t) = j]}{\sigma^2} \right) \\ & - \Phi \left(\frac{c_{i-1} - Pr[D(t) = j]}{\sigma^2} \right), \end{aligned} \tag{6.1}$$

where Φ is the cumulative normal distribution function. The confidence ratings are defined as follows: $c_0 = -\infty$, $c_1 = C_{2-3} - d$, $c_2 = C_{2-3}$, $c_3 = C_{2-3} + d$, $c_4 = \infty$. The parameters of the model (σ^2, C_{2-3}, d) were determined by maximising the log likelihood of an individual's confidence ratings across all trials of the task:

$$LL_M = \sum_{t=1}^{t} \sum_{j=1}^{4} \ln(C(t) = i \mid D_j(t)]) \cdot \delta_i(t), \tag{6.2}$$

where δ_i is a dummy variable that is 1 if confidence rating i is placed on trial t and 0 otherwise.

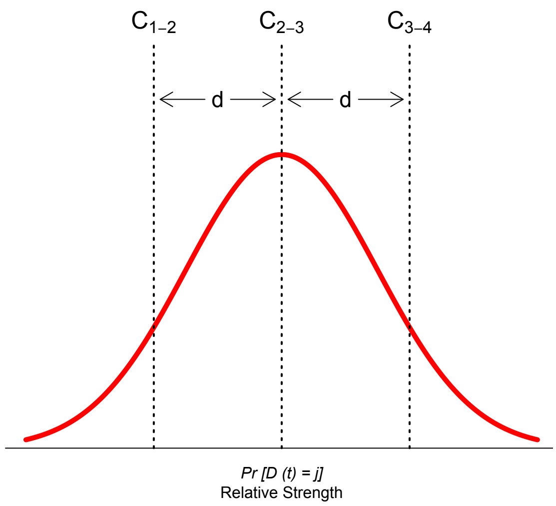

Figure 6.1: Graphical illustration of a hypothetical distribution based on which the probability of a certain confidence rating is estimated.

6.2 Model fitting results

In order to evaluate the model and its predictions we used the data from the two groups of Experiment 8 (Control and Switch). The use of the data from the switch group will allow us not only to evaluate the confidence model but also to examine how well the RL models perform in explaining choice behaviour in a controlled dynamic variation of the IGT. We used two RL models to generate the relative strength of each deck on each trial (using the best fitting parameters from each individual): the PVL-PU2 and the PVL-Delta models, which were found to provide the best descriptions of the choice data under the post-hoc fit and the simulation method, respectively. The reason for using the PVL-Delta model is that its predictions (i.e., relative strengths) rely more on payoff-related information compared to the PVL-PU2 which was found to strongly base its predictions on the history of past choices. Given that our hypothesis attempts to relate the value of each deck with its associated confidence, it could be the case that models that mimic past choices may distort the results regarding confidence by introducing biases (such as inertia or perseveration) which have nothing to do with the relative strength of each deck.

6.2.1 Control group

Table 6.1 contains the fit measures of choice and confidence models, and the parameters of the confidence model. The goodness of fit column shows the mean BIC scores relative to a statistical baseline model. As in the case of the choice models, the mean score for the confidence model represents the difference of its predictions from those of the baseline model, which assumes that the predicted probability of a confidence rating on each trial reflects the overall proportion of each confidence rating across all trials of the task. Thus, a positive score indicates that the cognitive confidence model is superior to the baseline model. It can be seen that the confidence model outperforms the baseline model under both RL models. However, the predictions from the PVL-Delta model seem to provide better input to the confidence model compared to PVL-PU2 model, as the BIC score of the confidence model given predictions from the PVL-Delta is higher (73.20

compared to 62.98).

Table 6.1: Mean relative BIC scores for the choice and confidence models and estimated parameters of the confidence model in the control group of Experiment 8.

Model	Goodness of Fit		Parameters		
	Choice	Confidence	C_{2-3}	d	σ^2
PVL-Delta	19.45	73.20	0.27	1.05	0.47
PVL-PU2	89.39	62.98	0.35	2.47	1.10

The question is why does this pattern emerge given that the PVL-PU2 model can perform better in predicting choice behaviour? A possible answer comes from inspecting the predicted probabilities of each model (PVL-Delta and PVL-PU2). Figure 6.2 shows that the predictions of the PVL-Delta model capture the rank ordering of the decks but they do not account for small changes in the predicted probabilities on a trial-by-trial basis. On the other hand, the PVL-PU2 accurately predicts the noisy changes in the observed choice probabilities and this is why it fits the data better than its competitor. However, participants' confidence ratings may depend on a more general representation of the value of each deck which does not conform to the trial-by-trial dynamics, which can also reflect exploratory behaviour. Therefore, the confidence model with relative strengths from the PVL-Delta provides a better description of participants' confidence judgments.

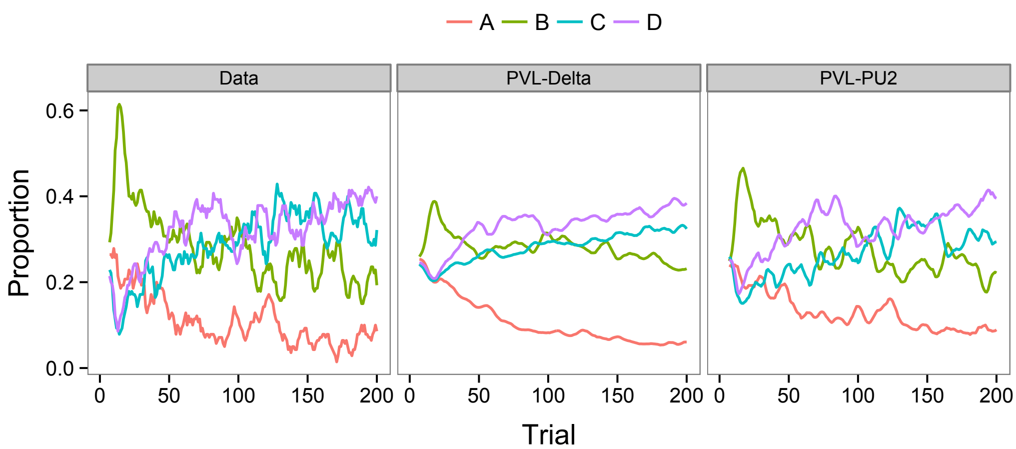

Figure 6.2: Mean predicted choice probabilities of the PVL-Delta and PVL-PU2 models in the control group of Experiment 8.

Figure 6.3 shows the confidence model predictions relative to participants' data regarding the proportion of each confidence rating after selection of each of the four decks. It is evident that the confidence model is able to capture the rank order of each confidence rating and its predictions are very close to the observed behavioural pattern. Figure 6.4 shows the mean confidence rating for each deck across blocks of 50 trials (solid colours) and the model's predictions (transparent colours), which yield a similar interpretation as in Figure 6.3: the model accurately predicts that good decks receive higher confidence ratings as participants accumulate more evidence about the value of these decks. Similarly, participants give lower ratings to the bad decks, and as their relative strength decreases over time, the confidence model adequately captures this tendency and assigns lower confidence ratings to the bad decks.

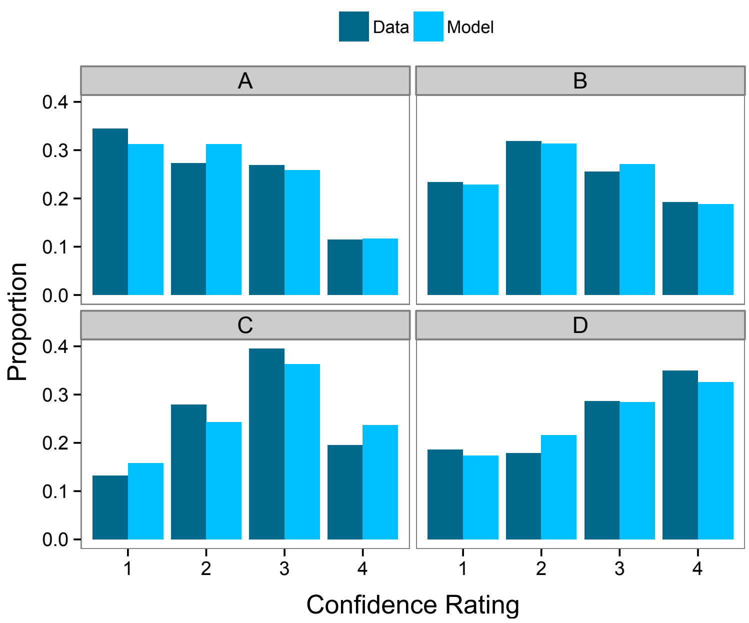

Figure 6.3: Observed proportions (Data) and predicted mean probabilities (Confidence Model) of confidence ratings after selecting each deck (A, B, C, and D).

Overall, the modelling results suggest that confidence ratings can be mapped onto and explained by the relative strength of each deck. This indicates that participants' knowledge about the task (as indexed by their confidence judgments) closely tracks the learned values of each deck and their relative strength, indicating a close connection between learning and awareness. By using a simple model to explain participants' metacognitive reports, we obtained evidence supporting the idea that confidence arises from the same mechanism that is responsible for learning the deck values and expectancies. Our results highlight the fact that choice behaviour and confidence are closely related and complementary components of the same underlying decision-making process.

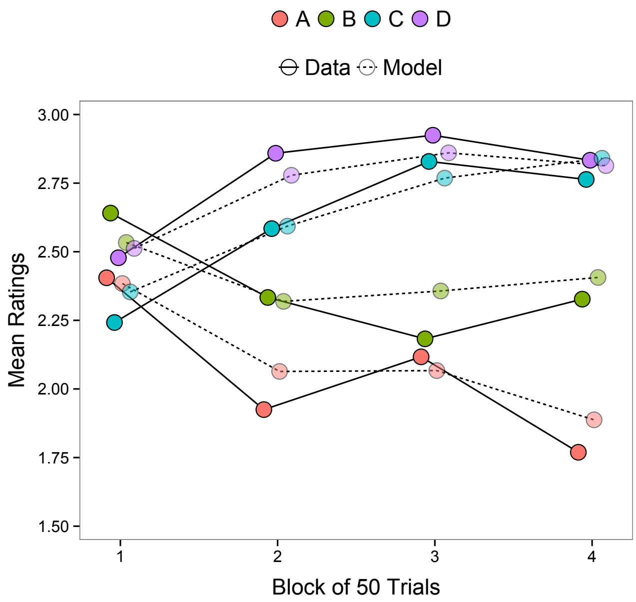

Figure 6.4: Observed (Data) and predicted (Confidence Model) mean confidence ratings for each deck across blocks of trials in the control group of Experiment 8. For each pair of colours, the darker one refers to the data and the lighter one to the model predictions.

6.2.2 Switch group

Examination of choice behaviour and confidence in the switch group poses two challenges: the first relates to whether RL models are able to capture the dynamics of the environment and adjust their predictions according to the new deck contingencies that each shift period introduces. The second challenge is whether the proposed confidence model can account for changes in confidence ratings across shift periods. The behavioural results of the switch group (see Figure 4.14: A and B) showed that participants discriminated in shift periods 3 and 4 between good and bad decks and gave higher confidence ratings to the good decks. The question is whether the cognitive models for choice behaviour and confidence ratings can make accurate predictions and capture these effects.

The first step was to fit the PVL-Delta and PVL-PU2 models to the choice data of the switch group. The procedure was identical to the one described in Section 5.3.1

(one-step-ahead predictions). Table 6.2 lists three different measures of fit (relative BIC score, BIC weights, and proportion of participants best fit) for the two RL models. The results resemble those of the control group, that is, the PVL-PU2 model provides better predictions compared to the PVL-Delta model regarding the observed choice pattern. However, the difference between the two models can no longer be ascribed to changes in the trial-by-trial micro-adjustments of the PVL-PU2 model to capture as accurately as possible the observed choice pattern (as was the case in the control group). Figure 6.5 shows differences between the two models regarding the rank ordering of decks. For example, the PVL-Delta model predicts more selections from the bad deck with infrequent losses in blocks 3 and 4 (original deck B) whereas the PVL-PU2 model correctly predicts the observed behavioural pattern (i.e., more selections from the good decks). It is possible that in such dynamic environments, some degree of reliance on previous choices can be beneficial as it can give a "boost" to the models to accurately predict subsequent choices. Hence, the PVL-PU2 model, which incorporates a decay RL rule, can closely track participants' choices, compared to the PVL-Delta model.

Table 6.2: Measures of fit for the choice and confidence models and estimated parameters of the confidence model in the switch group of Experiment 8.

Model	Goodness of Fit				Parameters		
	Choice			Confidence			
	BIC	wBIC	% Best Fit	BIC	C_{2-3}	d	σ^2
PVL-Delta	20.61	0.08	0.05	114.59	0.59	0.53	0.59
PVL-PU2	86.53	0.92	0.95	140.97	0.56	0.63	0.49

As expected, the confidence model fits the data better when it takes as input the relative strength of each deck as predicted by the PVL-PU2 model (see Table 6.2 - Confidence column). The difference in the choice pattern predicted by each of the two RL models is responsible for the better predictions of the confidence model under the PVL-PU2 (higher BIC score compared to the PVL-Delta). Figure 6.6 shows the behavioural data and model predictions regarding the mean confidence ratings following selections from each deck. The model accurately predicts the rank order of the mean ratings after selections from good and bad decks and its predictions are very similar to the ratings that participants assign to each deck. However, the model cannot reach the mean confidence rating of the good deck with frequent losses in blocks 3 and 4, as it predicts slightly lower ratings for this deck. Similarly, the model mispredicts the low mean rating following selections from the bad deck with infrequent losses in block 4, as it assigns a somewhat higher mean rating to it. Nevertheless, in the cases of mispredictions, the difference between observed data and model predictions are relatively small (≈ 0.25).

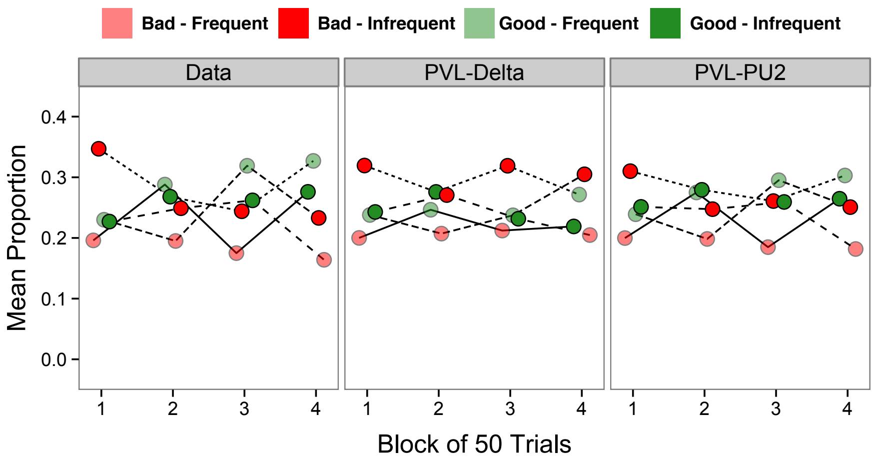

Figure 6.5: Mean predicted choice probabilities of the PVL-Delta and PVL-PU2 models in the switch group of Experiment 8.

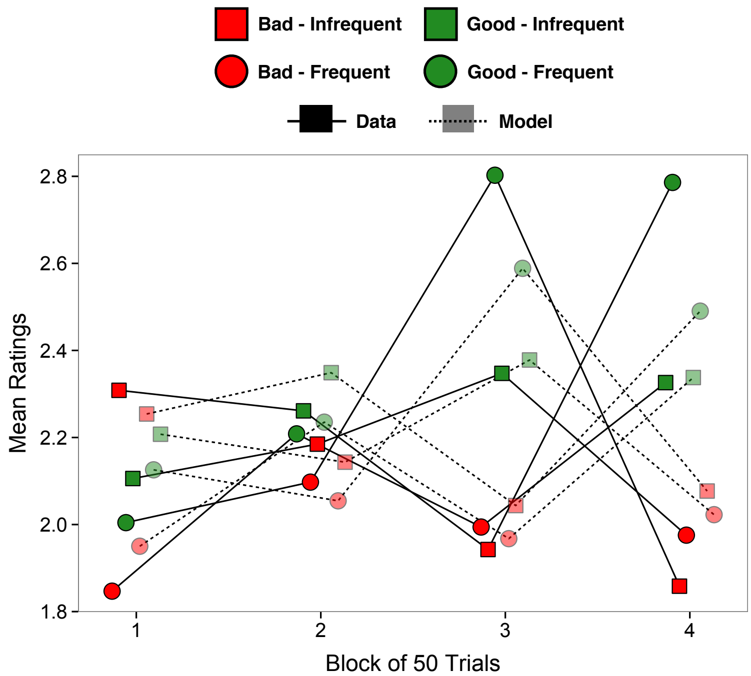

Figure 6.6: Observed (Data) and predicted (Model) mean confidence ratings for each deck across blocks of trials in the switch group of Experiment 8. For each pair of colours, the darker one refers to the data and the lighter one to the model predictions.

As in the control group, the confidence modelling results from the switch group suggest a very close link between choice behaviour and confidence ratings. More selections from the good decks are accompanied by higher confidence ratings and the model is able to capture this effect as it bases its predictions on the relative strength of each deck. In other words, decks with higher expectancies and relative strength are selected more often and receive higher ratings. Even in situations of high uncertainty (as in the switch group) where decks payoffs are not constant and deck quality (good or bad) changes over time, participants seem to rely on cognitive strategies to learn the values of the decks and make decisions. The same mechanism generates confidence ratings. This suggests that choice and awareness in the context of the IGT and other experience-based tasks are closely related, indicating that there is little evidence to support the view that learning to make advantageous decisions is dissociable from conscious awareness.

6.3 Discussion

The main objective of the present chapter was to develop a model which can account for confidence ratings in the context of the IGT. The general idea was based on the behavioural results of previous experiments which suggested that choice and confidence are closely related. Decks with positive overall expected values tend to be selected more often and participants' confidence judgments track this preference. The input of the confidence model was the relative strength of each deck as predicted by the choice RL models. The main results indicated that if the predictions of the RL models are in accordance with the observed choice pattern, then the confidence model can accurately predict the way that participants use the confidence scale.

Our results are in agreement with a range of studies in the decision-making literature that have established a close link between choice and subjective confidence. For example, Pleskac and Busemeyer's (2010) two-stage dynamic signal detection (2DSD) model suggests that a single dynamic and stochastic cognitive process (evidence accumulation as in a random walk/diffusion process) can give rise to the three components underlying a decision process: choice, reaction time, and confidence. In a recent neuroscience study, De Martino, Fleming, Garrett, and Dolan (2013) demonstrated that neural activity in the VMPFC brain area represents both choice and confidence, implying that these two components are separate behavioural aspects of the same underlying decision process. Similarly, other studies have shown that brain areas responsible for value computation also code neural signals related to choice certainty and confidence (e.g., Kepecs, Uchida, Zariwala, & Mainen, 2008; Kiani & Shadlen, 2009). Our confidence model is built on these ideas and the behavioural results from our experiments to provide a simple yet functional account of how confidence arises in the context of experience-based decision-making.

6.3.1 The importance of a reliable choice model

The modelling results of the present chapter highlight the importance of the model comparison analyses in Chapter 5 and underline the pitfalls of using an "erroneous" model to account for the observed choice pattern. The results from the PVL-Delta model in the switch group showed that this model cannot accurately predict the rank ordering of decks (see Figure 6.5) leading to inferior predictions regarding the confidence judgments compared to its competitor, the PVL-PU2 model. Had we used the PVL-Delta model to predict each deck's relative strength, the confidence model's predictions would have not accurately captured participants' confidence ratings.

As we mentioned earlier, good simulation performance does not always dictate that a model can reliably assess the mechanisms underlying choice. Learning to make decisions is dependent on the history of each individual's experience with the task and model comparison techniques which do not take this into account may fail to identify the best model. The inadequacy of the PVL-Delta model to correctly predict participants' choice behaviour in the dynamic variant of the IGT indicates that one should always look at a model's predictions before using it to make inferences about the underlying psychological and cognitive processes. The relative success of the PVL-PU2 may have been the result of reliance on the history of past choices to make accurate predictions, but this tendency could be beneficial in dynamically changing environments. Overall, the selection between candidate models must adhere to the participant's objectives and what she wants to achieve.

6.3.2 Limitations and future research

In the present formal analysis of participants' metacognitive judgments, we chose to examine the relationship between the relative strength of each deck and confidence ratings, because the latter was found to be sensitive in capturing participants' acquired knowledge of the task. A similar assessment could also be applied to wagering data and might shed light on the problematic aspects of this measure. For example, an explanation for why wagering is not as sensitive as confidence ratings is loss aversion. A possible way to tackle this and explain wagering judgments is to make the model's predictions dependent on some measure of the degree of loss aversion for each individual (for instance, the loss aversion parameter in the prospect utility function). In the context of our proposed model this could mean that the placement of the criteria for each wagering amount is more conservative.

A possible limitation of our modeling analysis is that model predictions and parameter estimates come from two separate model fitting procedures: first, the RL models are fitted to the choice data and then their predictions are used as input for the confidence model (second model fitting). We chose to apply this two-stage procedure in order to avoid any interactions between parameter estimates of the two models that could potentially distort model predictions for both choice behaviour and confidence ratings. Ideally, the parameter estimates of the RL and the confidence models would not interact as they measure different components of the decision-making process. Hence, our two-stage approach can be justified as an exploratory way to provide a simple formal account of confidence judgments.

Finally, a significant contribution to understanding choice and confidence is to include reaction times as part of the experimental design. Even though the relation between reaction times, choice, and confidence may not be easily identifiable in the context of experience-based decision-making (choices are not independent events/trials), it could still be useful to examine this and make use of a different modelling framework that accounts for all components of the decision-making process (e.g., modifications of Decision Field Theory, Busemeyer & Townsend, 1993, to account for confidence judgments).

7 General Discussion

The aim of this thesis was to examine the role played by unconscious or implicit influences on decision-making under uncertainty and specifically in the context of experience-based decision-making and the IGT. The results of the experiments reported here do not any offer any support for the claims that learning to make advantageous decisions can occur in the absence of awareness (e.g., Bechara et al., 1997; Persaud et al., 2007). As noted in the Introduction, research evaluating awareness in the IGT has formed a prominent and major element of the wider claim that unconscious thoughts or signals can influence choice, a fundamental theoretical idea most famously advocated and defended by Nisbett and Wilson (1977). The present work therefore bolsters recent suggestions (e.g., Newell & Shanks, 2014) that it is premature to assign a fundamental role to such processes in theories of decision-making.

The purpose of the behavioural experiments was twofold: first, to evaluate post-decision wagering as a sensitive and direct method of awareness, and secondly to investigate whether claims about implicit influences on decision-making are valid. A careful examination of wagering in comparison with other measures of awareness such as confidence ratings and quantitative questions also allowed us to explore the type of information that is essential for optimal decision-making and how participants use their acquired knowledge to make decisions in uncertain environments. The computational modelling analyses complemented the main behavioural results, showing a direct connection between learning of the advantageous strategy and participants' metacognitive reports.

7.1 Summary of results

Experiments 1 and 2 were near exact replications of Persaud et al.'s (2007) study. However, we only replicated the results relating to the quantitative questioning group of Experiment 1 where deck selection and advantageous wagering exceeded chance at the same time. In contrast to Persaud et al.'s results, the same pattern was observed in both experiments in groups that were asked only to make a wager after their deck selection, suggesting no dissociation between choice and wagering.

Chapter 3 examined the influence of different pay-off matrices on wagering strategies and two potential response biases: dominance of high wagers and loss aversion. Experiment 3 demonstrated that a simple change in the weights of the pay-off matrix can make wagering more sensitive (compared to the original pay-off matrix). Experiment 4 showed that even though wagering closely tracked learning of the advantageous strategy under the pay-off matrix proposed by Clifford et al. (2008), it underestimated participants' acquired knowledge as compared to their quantitative questions reports. The claim that loss aversion modulates wagering strategies was tested in Experiment 5. Even though a preference for the good decks appeared very early in the task, participants' wagers did not follow this tendency showing an effect of the high weights in the pay-off matrix.

The purpose of Chapter 4 was to compare wagering to another subjective measure of awareness, confidence ratings. In Experiment 6, we compared wagering with confidence ratings in an attempt to identify structural differences between the two measures and to provide a better examination of knowledge assessment in the IGT by employing a 4-point confidence scale. While both confidence scales (binary and continuous) showed conscious knowledge of the advantageous strategy in the IGT, this was not the case for wagering. Thus wagering is a less sensitive measure of awareness than confidence ratings. Also, knowledge in the IGT seems to be conscious: When we applied the guessing criterion (Dienes et al., 1995) to our data, there was no evidence of unconscious processing. Experiment 7 tested binary and continuous versions of wagering against confidence ratings. Confidence ratings and 4-point wagering indicated that participants made advantageous selections under situations of conscious knowledge but binary wagering failed to detect such a relationship, indicating that the use of wagering in a continuous manner may allow participants to use the scale more efficiently. However, the guessing criterion analysis on the 4-point wagering showed that participants used the lowest point of the scale more often compared to confidence ratings. While this might be taken as evidence that 4-point wagering can detect some unconscious learning in participants' decisions, an alternative explanation of this effect is that wagering is affected by loss aversion, leading participants to place low wagers even in situations of increased confidence. Experiment 8 demonstrated that conscious knowledge (as indexed by confidence ratings) is responsible for guiding decisions even in a dynamic version of the IGT where deck payoffs change periodically.

Chapters 5 and 6 provided a formal analysis of choice behaviour and confidence judgments, respectively. In Chapter 5, the computational modelling analysis of choice data using different model comparison techniques highlighted the importance of employing a reliable model to decompose the underlying psychological processes, and identified certain requirements that a good model needs to satisfy. Finally, the confidence model that we developed to account for participants' confidence ratings in Chapter 6 showed the strong relationship between choice and metacognition, suggesting that both processes are different manifestations of the same decision mechanism.

7.2 Measures of awareness

The task of validating measures and methods of assessing awareness is an important endeavour within psychological science as from the very beginnings of experimental psychology, researchers have been interested in the distinction between conscious and unconscious mental states (Dienes, 2008). This thesis put post-decision wagering under careful scrutiny because it is a method that supposedly removes biases and complications associated with verbal judgments of conscious knowledge. It has been extensively employed as a probe of conscious knowledge in several areas of experimental psychology such as perceptual decision-making and subliminal perception (e.g., Koriat, 2011; Nieuwenhuis & de Kleijn, 2011; Persaud et al., 2011; Persaud & McLeod, 2008; Sandberg, Bibby, Timmermans, Cleeremans, & Overgaard, 2011), implicit learning (e.g., Haider, Eichler, & Lange, 2011; Mealor & Dienes, 2012; Wierzchoń et al., 2012), and value-based decision-making (e.g., Lueddeke & Higham, 2011; Wang et al., 2012). Also, it has been used to study awareness in non-human animals (Kornell, Son, & Terrace, 2007; Middlebrooks & Sommer, 2011) and children (Miller, Brownell, & Zukier, 1977; Ruffman, Garnham, Import, & Connolly, 2001) where traditional confidence scales can be difficult to use.

Despite its extensive use, post-decision wagering has been found to be inferior to other traditional methods such as confidence ratings in terms of scale sensitivity and exhaustiveness. One major criticism refers to the fact that wagering is affected by loss aversion, leading participants to adopt more conservative criteria in the expression of their knowledge. Although that means foregoing large gains, wagering low minimises the likelihood of large losses, even when one has some confidence about which are the good decks. There is now widespread consensus that the application of wagering distorts measurements of conscious content and hinders expression of awareness (e.g., Sandberg et al., 2013; Szczepanowski et al., 2013; Wierzchoń et al., 2012). In addition, even when wagering is used in a continuous manner participants select the extreme points of the scale (i.e., the lowest and highest wagering amount) in a dichotomous manner, as this behaviour may allow for maximisation of earnings (Sandberg et al., 2010; Wierzchoń et al., 2012). If access to conscious knowledge has a graded character, then clearly post-decision wagering has a disadvantage in detecting changes in participants' metacognitive ability. The results of our experiments showed that indeed loss aversion affects participants' wagering behaviour. For example, in Experiment 5, even though participants demonstrated increased ability to discriminate between good and bad decks, they mainly placed low wagers in order to avoid large losses (which could occur even in selections from the good decks) given the high weights in the pay-off matrix associated with high wagering.

Comparisons between wagering and other measures of awareness such as confidence ratings and quantitative questions showed that wagering is a less sensitive method (e.g., Experiments 4, 6, and 7). It is possible that wagering is more sensitive to the emotional/motivational components of decision-making and thus a poor indicator of the acquisition of conscious knowledge. For example, it may assess the willingness of participants to engage in risk-taking behaviour which is not related to their acquired knowledge about the decks. Moreover, Pasquali, Timmermans, and Cleeremans (2010) argued that advantageous wagering can be acquired in the absence of awareness which would make it an unsuitable measure of awareness in the IGT. The overall utility of post-decision wagering as a reliable measure of awareness needs to be further examined under different settings (e.g., different pay-off matrices, no-loss gambling in order to remove the effect of loss aversion) and other experimental conditions, as it seems to be unsuitable in a context (the IGT) where the first-order task also involves gambling.

The results of the present thesis only extend to the IGT and further evaluations of post-decision wagering are needed using different behavioural tasks or populations. For instance, Persaud et al. (2007) reported a dissociation between awareness (as measured by wagering) and behaviour in three different tasks: artificial grammar learning, blindsight, and the IGT. However, our findings are consistent with other studies which have shown that wagering is no more reliable or exhaustive than confidence ratings (e.g., Dienes & Seth, 2010; Fleming & Dolan, 2010; Szczepanowski et al., 2013).

It is important to emphasise that although our conclusions are very different from those of Persaud et al. (2007), this is not because of any substantial disagreement about the fundamental data patterns. On the contrary, we were able to reproduce the key finding they reported - a lag between deck discrimination and the onset of advantageous wagering - in the simple wagering group of Experiment 3 and the wagering group of Experiment 6. It is true that the conditions in which we obtained this pattern were slightly different from those in which Persaud et al. obtained it (for example, in our studies it depended on using the original IGT payoff schedule) and that in the no questioning group of Experiment 1 we did not obtain it, despite the fact that this group comprised a near-exact replication of Persaud et al.'s experiment, suggesting that some subtle procedural factors influence whether or not a lag occurs. Where we are in disagreement is in the interpretation of this lag. Whereas Persaud et al. took it as evidence of unconscious influences in decision making that drive deck selections before participants become aware and able to wager adaptively, we take the lag as evidence of the insensitivity of wagering. Our case for this conclusion rests on the finding that the lag was eliminated or indeed reversed as a result of (1) a minor change in the payoff matrix in the modified wagering group of Experiment 3, (2) a switch from a binary wager response to either a binary or a 4-point confidence response in Experiments 6 and 7, and (3) employing explicit verbal questions such as "if you could only select cards from one of the decks until the end of the game... which of the four decks would you pick?" to assess awareness.

Confidence ratings, on the other hand, seem to reliably assess the level of knowledge about the task that participants have acquired. Even in situations where the payoffs and quality of the decks change (switch group in Experiment 8), participants' confidence ratings are in general agreement with their selections. The fact that confidence ratings do not change the received payoffs on each trial (as in the case of wagering) makes them a more reliable and sensitive measure in the context of the IGT, as participants can directly report their awareness without thinking of the consequences of the use of the scale. Also, our confidence model (Chapter 6) indicated a strong association between the relative strength of each deck and confidence judgments. It will be interesting for future research to apply the same formal analysis to wagering and identify what aspects of the decision-making process it measures.

7.3 Emotion, uncertainty, and awareness

The somatic marker hypothesis (SMH) re-introduced the idea that emotional signals are beneficial to reasoning processes and decision-making. The main claim was that these emotional signals or somatic markers help people make decisions in situations of uncertainty as they carry information about the quality of available choices and drive participants' decisions in the absence of conscious awareness. The *emotions-as-input* account (Davis et al., 2009) has been challenged by recent studies (e.g., Fernie & Tunney, 2013;

Gutbrod et al., 2006; see also section 1.3.3) which suggested that these covert somatic markers may be the result of acquired conscious knowledge about the task (*emotions-as-outcome* account) rather than its generating force (Newell & Shanks, 2014). Hence, somatic markers may perform overtly and consolidate or reinforce a cognitive representation about the value of an available choice. Another possibility is that the somatic markers (as indexed by SCRs) code the uncertainty or riskiness associated with an option (Davis et al., 2009; Fernie & Tunney, 2013; Tomb et al., 2002; Wagar & Dixon, 2006). Fernie and Tunney (2013) noted that explicit knowledge and awareness of risk for the bad decks (where the variance in the payoffs is higher compared to the good decks) may generate higher SCRs for these decks. This is in accord with Tomb et al. (2002) who found that it is the uncertainty/variance in the payoffs and not the overall quality of a decision (good or bad) that is responsible for higher somatic activity.

In the light of recent evidence, the relationship between SCRs and performance on the IGT (i.e., higher somatic activity leads to better performance on the task, e.g., Carter & Smith Pasqualini, 2004; Oya et al., 2005) may be seen as the result of the development of conscious knowledge about the task and not that anticipatory SCRs precede the emergence of knowledge. However, Wagar and Dixon (2006) presented evidence that affective information developed before participants behaviourally discriminated between good and bad decks and had conscious knowledge of the task. Even though the onsets of aSCRs (between trials 20-30) and good deck selections (between trials 30-40) in their study are very close to each other, Wagar and Dixon suggested that increased riskiness in the bad decks may have driven the early development of differential SCRs. But this surely means that participants had some knowledge (not yet conceptual) that the bad decks returned high magnitude losses. In another study, Guillaume et al. (2009) observed that conscious knowledge and aSCRs correlated with task performance, but not with each other. As these authors suggested, the lack of correlation between knowledge and somatic activity may have been the result of low statistical power. In addition, assessment of participants' knowledge was only assessed at the end of the IGT which did not allow for a direct comparison between the development of aSCRs and explicit knowledge and failed to satisfy the *immediacy* property of an accurate assessment of conscious knowledge (e.g., Lovibond & Shanks, 2002; Newell & Shanks, 2014).

In the experiments reported here we did not measure participants' somatic activity through SCRs, but we focused on whether differentiation between good and bad decks is accompanied by explicit knowledge of the deck contingencies. Emotional information may play a significant role in shaping decision strategies but this largely depends on the definition and functional properties of emotion and the methods that are used to measure somatic or emotional activity. It may be the case that conventional SCRs cannot fully assess the richness of emotional information (e.g., Newell & Shanks, 2014). In addition, more sophisticated model-based analyses of SCRs are needed in order to decompose and understand what this measurement reflects (e.g., Bach & Friston, 2013).

7.4 Decision-making in the IGT

Early observations regarding decision-making of healthy individuals in the IGT suggested that participants predominantly chose the good decks C and D after an exploration phase (e.g., Bechara et al., 1994, 1999). When losses started to accumulate from the bad decks A and B, which were initially thought as good because they returned high rewards, participants switched to the good decks and this preference continued until the end of the task. In addition, choice behaviour within good and bad decks was uniform, indicating that the preference for the good decks was driven by the positive overall expected values associated with these decks regardless of other factors such as the probability and magnitude of losses.

However, these early assumptions about the IGT have been questioned in light of recent experimental evidence (see also section 1.2.1). Many IGT studies have demonstrated that participants' choice behaviour is mainly guided by a loss-frequency effect, meaning that they prefer the decks which produce infrequent losses despite the fact that these losses are numerically high (e.g., Lin et al., 2007; Steingroever, Wetzels, Horstmann, et al., 2013). Also, there is no clear exploration-exploitation trade-off and participants do not exploit the most profitable decks but instead go back to select cards from the disadvantageous decks, especially deck B.

The results from the present experiments are broadly in line with these findings. In the experiments that used the standard IGT payoff schedule, participants' choices were predominantly guided by loss frequency, and deck D was selected more often than any other deck. The interesting finding is that deck B, which is a disadvantageous deck, was selected as often as deck C, even though the latter has a positive expected value, and in some cases was selected as often as deck D (e.g., see Experiments 6 and 7). In fact, a clear preference for the good decks was only observed in Experiment 1, where participants selected decks C and D more often than decks A and B.

Most IGT studies report only a weak overall preference for the good decks, between 50% and 60% (Steingroever, Wetzels, Horstmann, et al., 2013), with continued sampling from deck B. The results from our experiments, where we employed a payoff structure similar or identical to the original IGT, are in reasonable accordance with these percentages. However, in the experiments reported here, we assessed awareness concurrently (wagering, confidence ratings, and questionnaire) with decision-making and this may have had a reactive effect on deck-selections, making participants more attentive to the deck payoffs. For instance, participants may focus more on the task knowing that they will have to answer specific questions about the decks. Similarly, Cella et al. (2012) argued that the systematic assessment of participants' awareness may facilitate performance on the IGT.

The same also applies to wagering as the tendency to maximise winnings can increase participants' motivation to perform well in the task (Sandberg et al., 2010). In fact, Szczepanowski et al. (2013) found that performance on a cognitive task (detection of fearful faces) was increased when post-decision wagering was simultaneously used as a probe of conscious knowledge, suggesting that financial incentives can motivate participants to perform better on the primary task. Another aspect of using post-decision wagering is the magnitude change in payoffs, as wagers in our tasks were employed as multipliers of the actual deck payoffs. Better ability to discriminate between good and bad decks has also been observed in other IGT studies in which participants' awareness was assessed at the same time as their decision-making performance (e.g., Dymond, Bailey, Willner, & Parry, 2010; Evans et al., 2005; Maia & McClelland, 2004; Persaud et al., 2007; Wagar & Dixon, 2006).

Our results also show that participants can be sensitive to differences among the decks regarding their overall expected value. In Experiment 4A, the most profitable deck (deck C) was favoured and there was no overall difference between the decks with a negative total outcome. The key finding of this experiment was that participants were able to grasp the payoff structure very early in the task, which suggests that difficulties participants experience in the classic IGT may be associated with its idiosyncrasies. First, when participants encounter the initial loss in deck B on trial 9, they may think of it as a rare event and hence keep selecting cards from this deck. Secondly, the concurrent presentation of wins and losses might make it harder to acquire the optimal strategy. Thirdly, it has been shown that 100 trials are not sufficient for participants to learn and exploit the advantageous decks (Fernie & Tunney, 2008; Wetzels et al., 2010).

7.5 Insights from cognitive modelling

The purpose of the cognitive modelling analyses in Chapter 5 was twofold: first, to examine whether RL models can provide a good description of the observed choice patterns, and second, to compare the candidate models using different methods and assessments such as the post hoc fit method, the simulation method, the degree of reliance on past choices (EPSE) method, model recovery, and parameter recovery. These analyses provided a thorough examination of recently debated topics in cognitive modelling of IGT data including which model should be preferred (e.g., Steingroever, Wetzels, & Wagenmakers, 2013a; Worthy, Pang, & Byrne, 2013), the inconsistencies in model comparison techniques (e.g., Ahn et al., 2008; Steingroever et al., in press), and parameter recovery and identifiability (e.g., Wetzels et al., 2010). Even though we did not propose a single model which can be used to decompose and explain differences in performance on the IGT (e.g., between clinical and healthy groups of participants), we highlighted the importance of a thorough assessment of the candidate models and suggested additional tests of model adequacy such as parameter consistency and generalisation at the individual level.

7.5.1 Choice mimicry and reliance on past choices

Chapter 5 demonstrated that the selection between competing models for the IGT is not an easy task. Different model comparison techniques indicated that the use of a specific model depends to a substantial extent on one's objectives; that is, whether one wants to achieve short-term or long-term predictions about choice behaviour in the IGT. One possible explanation of the discrepancy between model comparison techniques (e.g., post hoc fit and simulation methods) lies in the degree of reliance of each model on the actual history of choices that participants make. Models which implement the decay RL rule provide better predictions under the post hoc fit method because they mimic participants' choices whereas models with the delta learning rule rely more on payoff related information and achieve better predictions under the simulation method. This hypothesis was tested using the EPSE method by Yechiam and Ert (2007) and the results revealed that this could explain the difference between the two methods.

In technical terms, the choice mimicry component of the decay RL rule lies in the fact that it favours decks which have been selected more often, while the expectancies of unselected decks decay towards 0. Hence, decks with higher expectancies are more likely to be selected on the next trial. The rate of decrease in expectancies depends on the decay parameter D (see Equation 5.5). If D has a value of 0, then the expectancy of the selected deck j takes the value of the subjective utility of the received payoffs, $E_j(t) = u_j$, whereas the expectancies of the unselected decks become 0. If choices are deterministic (i.e., decks with high expectancies are more likely to be selected) then a D of 0 suggests complete mimicry of previous choices. Inspection of the mean value of D in the PVL-PU2 model across the three datasets of Chapter 5 (0.73, 0.58, and 0.68, respectively; see also Table 5.7) suggests that complete mimicry of past choices was not observed. In addition, consider the case of the switch group of Experiment 8 where decks' payoffs changed every

50 trials. The PVL-PU2 (decay RL rule) model could accurately predict the rank ordering of the decks whereas the PVL-Delta (delta learning rule) failed to achieve a good level of prediction (see Figure 6.5). Hypothetically, this difference could be explained based on the strong reliance of the PVL-PU2 model on past choices. However, the value of D in the switch group was 0.50, indicating that the predictions of the PVL-PU2 model were not exclusively the result of choice mimicry.

Overall, choice mimicry or reliance on past choices may not be the only factor that drives accurate model predictions as in the case of the PVL-PU2 model. Future research can attempt to carefully assess this issue and tease apart the psychological processes that drive performance on the IGT.

7.5.2 Reinforcement-learning and alternative accounts

The results of the computational modelling analyses of choice data in Chapter 5 suggested that a simple reinforcement-learning strategy can account for choice behaviour in the IGT. The main mechanism of RL models is the formation of expectancies about the values of each deck which subsequently drive deck selection. Another assumption relates to the subjective experience of received payoffs which takes into account the asymmetric evaluation of wins and losses. The trade-off between exploration and exploitation is governed by a consistency/temperature parameter in the softmax rule. The fact that RL models provide a good account of choice dynamics in the IGT does not necessarily mean that people employ reinforcement-learning mechanisms to make decisions. For example, a simple heuristic strategy based on "win-stay/lose-shift" accounted for the performance of a large proportion of participants (see Table 5.5).

Alternative theories and modelling approaches have been proposed which depart from RL principles. For example, the Instance-Based Learning Theory (IBLT; Gonzalez et al., 2003; Lejarraga et al., 2010) assumes that people store instances of payoffs in their memory and that choice is based on the sum of all the instances of an option weighted by their probabilities of retrieval. A similar account of the way that people retrieve instances or experiences is the Explorative Sampler Model by Erev, Ert, and Yechiam (2008). On each trial, people enter one of two possible response states: exploration or exploitation where exploration implies random choice (Nevo & Erev, 2012, added inertia as a third response state). Choice in exploitation mode depends on each option's subjective value which is calculated as the average payoff of the recalled experiences. The average payoffs are subject to additional modification which reflects the level of payoff variability. Other approaches in experience-based decision-making have attempted to provide a more sophisticated view of the exploration-exploitation trade-off. For example, Speekenbrink and Konstantinidis (2014) found that explorative behaviour was dependent on the uncertainty regarding each option's average payoffs over time and that most people's choice behaviour was best explained based on a fine balance between exploration and exploitation by considering the probability that an option offers the maximum payoff out of all the available options.

As we noted in the Introduction, the list of models that can be applied to experience-based tasks including the IGT is very large. In addition, a good model fit may be misleading and the result of model properties that are unrelated to the psychological processes under investigation such as overfitting and flexibility (see Pitt & Myung, 2002). Cognitive models cannot stand in isolation but need to make informative connections to other measures of the same underlying process that they attempt to decompose and explain. Hence, future research on cognitive modelling of experience-based decision-making can address the issue of whether the proposed models capture basic psychological mechanisms and whether they reflect stable and generalisable assessments of the underlying cognitive and psychological processes.

7.5.3 Confidence and its relation to choice

The results of the behavioural experiments of this thesis (e.g., Experiments 6-8) suggested a close link between deck selections and confidence ratings, meaning that decks with positive expected values were more likely to be selected and received higher confidence ratings compared to the decks with negative expected values. Chapter 6 provided a formal analysis of participants' confidence judgments which were in accordance with the relative strength of each deck as predicted by the choice RL models. To the best of our knowledge, this is the first demonstration that choice and confidence in sequential economic decisions may constitute different manifestations of the same underlying decision mechanism. It also strengthens the claim that learning to make advantageous decisions is not dissociable from awareness in the IGT and offers little support to the view that unconscious influences play an important role in decision-making under uncertainty.

The confidence modelling results are in line with findings from studies in the perceptual and value-based decision-making literatures which have shown that choice and confidence are interrelated components of decision-making. As our approach was a first and exploratory attempt to relate and formally analyse confidence judgments, certain ideas can be taken forward by future research. For example, the investigation of reaction times and decision confidence in this context may shed more light into the psychological underpinnings of experience-based decision-making. Similar analyses can be applied to wagering and other assessments of awareness in order to examine the aspects of the decision process that these measures assess.

7.6 Concluding remarks

This thesis examined the influence of implicit or unconscious processes on experience-based decision-making and the IGT using behavioural experiments and computational modelling analyses. In the experiments reported here, we obtained results at odds with the predictions of the somatic marker hypothesis regarding the activation of an unconscious emotional system, which is assumed to provide information about the outcome of the decision-making process. Decision strategies in the IGT rely almost exclusively on acquired conscious knowledge about the properties of the decks. The second major point of the present thesis is that caution is advised when drawing conclusions about the existence of implicit influences in decision-making under uncertainty when unsuitable methods of measuring awareness are used. Persaud et al.'s (2007) conclusions seem ungrounded because of the pronounced failure of post-decision wagering to measure awareness with adequate sensitivity. We have shown that wagering underestimates awareness by comparison to other methods such as confidence ratings and quantitative questions, and that wagering strategies are affected by the design of the pay-off matrix and response biases. If we are to measure awareness as accurately and sensitively as possible, the results from different methods should be combined in order to provide a finer and deeper examination of claims involving implicit or unconscious influences.

The confidence model that we developed provided further support that the same psychological mechanism which is responsible for choice behaviour can explain the way participants place confidence ratings. Given the strong relationship between choice and confidence, the evidence in favour of unconscious influences is very tenuous.

Appendices

A Payoff schedule in Experiment 1

Table A.1: Payoff schedule in Experiment 1

	Deck					**Deck**			
Trial	**A**(+2)	**B** (+2)	**C** (+1)	**D** (+1)	**Trial**	**A** (+2)	**B** (+2)	**C** (+1)	**D** (+1)
1					**26**	-4		-1	
2					**27**	-5			
3	-3		-1		**28**	-3			
4					**29**			-1.5	-5
5	-6		-1		**30**			-1	
6					**31**	-7			
7	-4		-1		**32**	-4	-25		
8					**33**	-5			
9	-5	-25	-1		**34**			-0.5	
10	-7		-1	-5	**35**			-0.5	-5
11					**36**				
12	-7		-0.5		**37**	-3		-1.5	
13			-1.5		**38**	-6			
14	-5	-25			**39**			-1	
15	-4				**40**		-25	-1.5	
16					**41**				
17	-6		-0.5		**42**	-5		-1	
18	-3		-1.5		**43**				
19					**44**			-1	
20			-1	-5	**45**	-7		-1	
21		-25			**46**	-6			-5
22	-6				**47**			-1	
23					**48**	-3		-1	
24	-7		-1		**49**				
25			-0.5		**50**	-4			

Note: The payoff schedule was constructed based on the ratio of losses to wins of the original IGT. For example, in the original task deck A has a 50% probability of loss. The wins are always £100 and the losses range from £150 to £350. On trial 3 there is a loss of £150. The ratio of losses to wins is therefore 1.5. Since the coefficient of win is always 2 in Persaud et al.'s (2007) variation, we multiplied each ratio by 2. This gives us the schedule of the losses. The average wager is $(20 + 10)/2 = 15$. For example, for the first 10 trials on deck A the losses are $(3 + 6 + 4 + 5 + 7) \times 15 = 375$ and the wins are $(2 \times 10) \times 15 = 300$. The difference is five times the average wager per ten cards $(5 \times 15 = 75)$. This procedure was applied on the remaining decks. The payoff schedule is repeated twice for the 100 trials of the task.

B Adapted Questionnaire from Maia and McClelland (2004)

1. Please rate, on a scale from -10 to $+10$, how good or bad you think deck A is, where -10 means that it is very bad and $+10$ means that it is very good.

2. Okay, now suppose that you were to select 10 cards from deck A.

   a) What would you expect your average result to be?

   b) For those trials in which you would get a win, what would you expect your average winning amount to be?

   c) In how many of the 10 trials would you expect to get a loss?

   d) For those trials in which you would get a loss, what would you expect the average loss to be?

3. Now suppose I told you that you could only select cards from one of the decks until the end of the game, but that you were allowed to choose now the deck from which you would draw your cards. Which of the four decks would you pick?

Note: Questions 1 and 2 are repeated for the remaining decks (B, C, and D).

C Modelling Results

C.1 Predicted Choice Probabilities

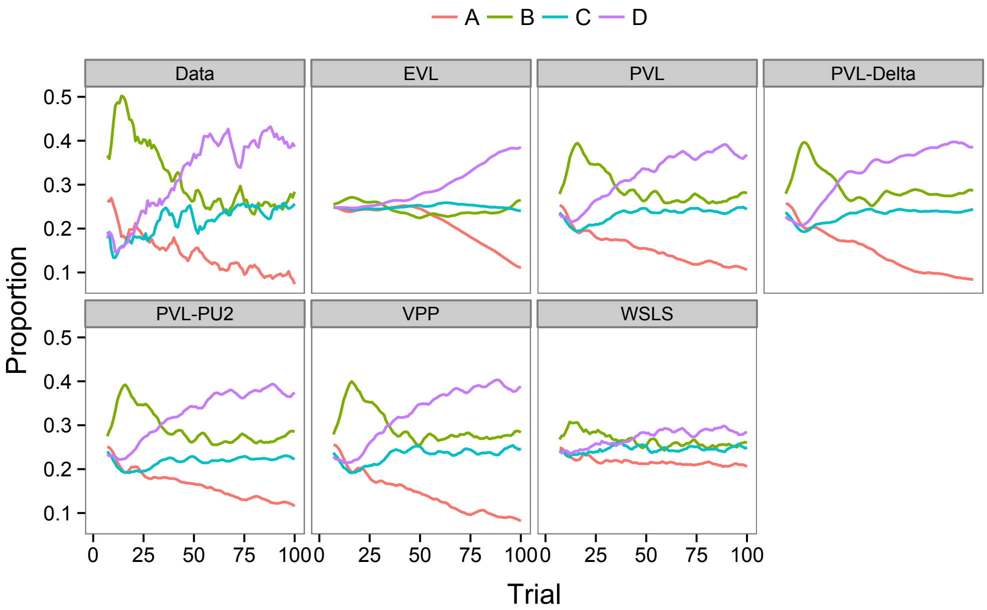

Figure C.1: Mean predicted choice probabilities of each cognitive model as compared to experimental data in Dataset 1

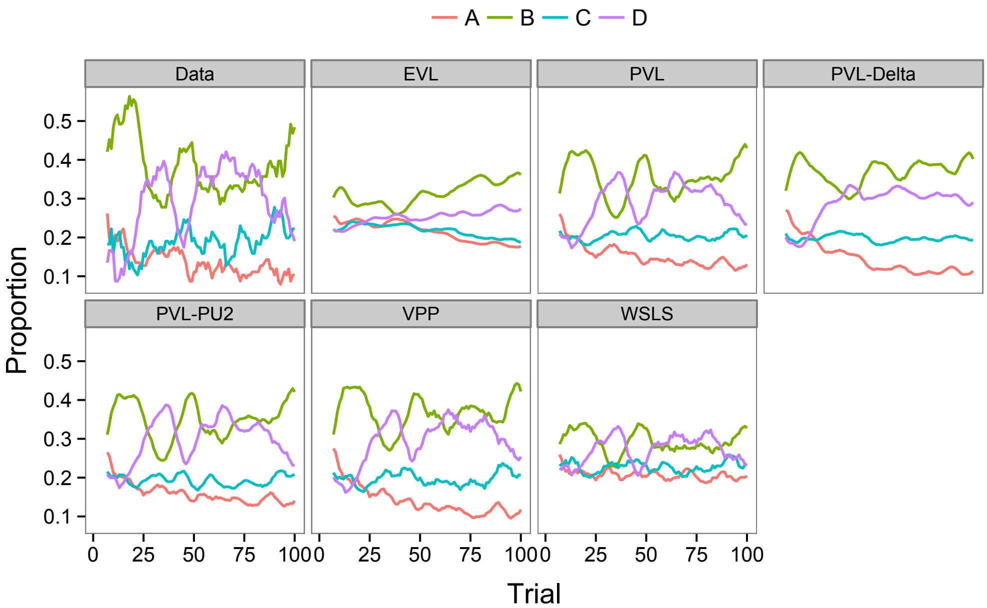

Figure C.2: Mean predicted choice probabilities of each cognitive model as compared to experimental data in Dataset 2

References

Agresti, A. (2002). *Categorical data analysis* (2nd ed.). Hoboken, NJ: Wiley-Interscience.

Ahn, W.-Y., Busemeyer, J. R., Wagenmakers, E.-J., & Stout, J. C. (2008). Comparison of decision learning models using the generalization criterion method. *Cognitive Science*, *32*, 1376–1402.

Ahn, W.-Y., Krawitz, A., Kim, W., Busemeyer, J. R., & Brown, J. W. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics*, *4*, 95–110.

Bach, D. R., & Friston, K. J. (2013). Model-based analysis of skin conductance responses: Towards causal models in psychophysiology. *Psychophysiology*, *50*, 15–22.

Bakeman, R. (2005). Recommended effect size statistics for repeated measures designs. *Behavior Research Methods*, *37*, 379–384.

Baranski, J. V., & Petrusic, W. M. (1998). Probing the locus of confidence judgments: Experiments on the time to determine confidence. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 929–945.

Bar-On, R., Tranel, D., Denburg, N. L., & Bechara, A. (2003). Exploring the neurological substrate of emotional and social intelligence. *Brain*, *126*, 1790–1800.

Barrett, A. B., Dienes, Z., & Seth, A. K. (2013). Measures of metacognition on signal-detection theoretic models. *Psychological methods*, *18*, 535–552.

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

Bechara, A. (2004). The role of emotion in decision-making: Evidence from neurological patients with orbitofrontal damage. *Brain and Cognition*, *55*, 30–40.

Bechara, A., & Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and Economic Behavior*, *52*, 336–372.

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

Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex*, *10*, 295–307.

Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *The Journal of Neuroscience*, *19*, 5473–5481.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*, 1293–1295.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (2005). The Iowa Gambling Task and the somatic marker hypothesis: Some questions and answers. *Trends in Cognitive Sciences*, *9*, 159–162.

Bechara, A., Tranel, D., & Damasio, H. (2000). Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions. *Brain*, *123*, 2189–2202.

Bechara, A., Tranel, D., Damasio, H., & Damasio, A. R. (1996). Failure to respond autonomically to anticipated future outcomes following damage to prefrontal cortex. *Cerebral Cortex*, *6*, 215–225.

Berry, C. J., Shanks, D. R., Speekenbrink, M., & Henson, R. N. A. (2012). Models of recognition, repetition priming, and fluency: Exploring a new framework. *Psychological Review*, *119*, 40–79.

Biele, G., Erev, I., & Ert, E. (2009). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology*, *53*, 155–167.

Bowman, C. H., Evans, C. E. Y., & Turnbull, O. H. (2005). Artificial time constraints on the Iowa Gambling Task: The effects on behavioural performance and subjective experience. *Brain and Cognition*, *57*, 21–25.

Brand, M., Recknor, E. C., Grabenhorst, F., & Bechara, A. (2007). Decisions under ambiguity and decisions under risk: Correlations with executive functions and comparisons of two different gambling tasks with implicit and explicit rules. *Journal of Clinical and Experimental Neuropsychology*, *29*, 86–99.

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, *81*, 211–241.

Buelow, M. T., & Suhr, J. A. (2009). Construct validity of the Iowa Gambling Task. *Neuropsychology Review*, *19*, 102–114.

Busemeyer, J. R. (2002). Dynamic decision making. In N. J. Smelser & P. B. Baltes (Eds.), *International encyclopedia of the social and behavioral sciences: Methodology, mathematics and computer science* (pp. 3903–3908). Oxford: Elsevier.

Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, *14*, 253–262.

Busemeyer, J. R., Stout, J. C., & Finn, P. (2003). Using computational models to help explain decision making processes of substance abusers. In D. Bargh (Ed.), *Cognitive and affective neuroscience of psychopathology.* Oxford, England: Oxford University Press.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision Field Theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432–459.

Busemeyer, J. R., & Wang, Y. M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, *44*, 171–189.

Carter, S., & Smith Pasqualini, M. C. (2004). Stronger autonomic response accompanies better learning: A test of Damasio's somatic marker hypothesis. *Cognition & Emotion*, *18*, 901–911.

Cella, M., Dymond, S., Cooper, A., & Turnbull, O. H. (2007). Effects of decision-phase time constraints on emotion-based learning in the Iowa Gambling Task. *Brain and Cognition*, *64*, 164–169.

Cella, M., Dymond, S., Cooper, A., & Turnbull, O. H. (2012). Cognitive decision modelling of emotion-based learning impairment in schizophrenia: The role of awareness. *Psychiatry Research*, *196*, 15–19.

Chiu, Y.-C., Lin, C.-H., Huang, J.-T., Lin, S., Lee, P.-L., & Hsieh, J.-C. (2008). Immediate gain is long-term loss: Are there foresighted decision makers in the Iowa Gambling Task? *Behavioral and Brain Functions*, *4*, Article 13.

Clifford, C. W. G., Arabzadeh, E., & Harris, J. A. (2008). Getting technical about awareness. *Trends in Cognitive Sciences*, *12*, 54–58.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*, 933–942.

Critchley, H. D., Elliott, R., Mathias, C. J., & Dolan, R. J. (2000). Neural activity relating to generation and representation of galvanic skin conductance responses: A functional magnetic resonance imaging study. *The Journal of Neuroscience*, *20*, 3033–3040.

Crone, E. A., & van der Molen, M. W. (2004). Developmental changes in real life decision making: Performance on a gambling task previously shown to depend on the ventromedial prefrontal cortex. *Developmental Neuropsychology*, *25*, 251–279.

Dai, J., Kerestes, R., Upton, D. J., Busemeyer, J. R., & Stout, J. C. (2014). *An improved cognitive model of the Iowa and Soochow gambling tasks with regard to model fitting performance and parameter consistency test.* Manuscript submitted for publication.

Damasio, A. R. (1994). *Descartes' error: Emotion, reason and the human brain.* New York, NY: Putnam.

Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *351*, 1413–1420.

Damasio, A. R., Adolphs, R., & Damasio, H. (2003). The contributions of the lesion method to the functional neuroanatomy of emotion. In R. J. Davidson, K. R. Scherer, & H. H. Goldsmith (Eds.), *Handbook of affective sciences* (pp. 66–92). Oxford: Oxford University Press.

Davis, T., Love, B. C., & Maddox, W. T. (2009). Anticipatory emotions in decision tasks: Covert markers of value or attentional processes? *Cognition*, *112*, 195–200.

DeLong, E. R., DeLong, D. M., & Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: A nonparametric approach. *Biometrics*, *44*, 837–845.

De Martino, B., Camerer, C. F., & Adolphs, R. (2010). Amygdala damage eliminates monetary loss aversion. *Proceedings of the National Academy of Sciences*, *107*, 3788–3792.

De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based choice. *Nature Neuroscience*, *16*, 105–110.

Dienes, Z. (2004). Assumptions of subjective measures of unconscious mental states: Higher order thoughts and bias. *Journal of Consciousness Studies*, *11*, 25–45.

Dienes, Z. (2008). Subjective measures of unconscious knowledge. *Progress in Brain Research*, *168*, 49–64.

Dienes, Z., Altmann, G. T. M., Kwan, L., & Goode, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1322–1338.

Dienes, Z., & Seth, A. K. (2010). Gambling on the unconscious: A comparison of wagering and confidence ratings as measures of awareness in an artificial grammar task. *Consciousness and Cognition*, *19*, 674–681.

Dolan, R. J. (2002). Emotion, cognition, and behavior. *Science*, *298*, 1191–1194.

Dougherty, M. R. (2001). Integration of the ecological and error models of overconfidence using a multiple-trace memory model. *Journal of Experimental Psychology: General*, *130*, 579–599.

Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis:

A critical evaluation. *Neuroscience and Biobehavioral Reviews*, *30*, 239–271.

Dutt, V., & Gonzalez, C. (2012a). Making Instance-based Learning Theory usable and understandable: The Instance-based Learning Tool. *Computers in Human Behavior*, *28*, 1227–1240.

Dutt, V., & Gonzalez, C. (2012b). The role of inertia in modeling decisions from experience with instance-based learning. *Frontiers in Psychology*, *3*, Article 177.

Dymond, S., Bailey, R., Willner, P., & Parry, R. (2010). Symbol labelling improves advantageous decision-making on the Iowa Gambling Task in people with intellectual disabilities. *Research in Developmental Disabilities*, *31*, 536–544.

Dymond, S., Cella, M., Cooper, A., & Turnbull, O. H. (2010). The contingency-shifting variant Iowa Gambling Task: An investigation with young adults. *Journal of Clinical and Experimental Neuropsychology*, *32*, 239–248.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processing. *Human Factors*, *4*, 59–73.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*, 912–931.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., . . . Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 15–47.

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, *21*, 575–597.

Erev, I., & Haruvy, E. (2005). Generality, repetition, and the role of descriptive learning models. *Journal of Mathematical Psychology*, *49*, 357–371.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, *88*, 848–881.

Evans, C. E. Y., Bowman, C. H., & Turnbull, O. H. (2005). Subjective awareness on the Iowa Gambling Task: The key role of emotional experience in schizophrenia.

Journal of Clinical and Experimental Neuropsychology, *27*, 656–664.

Fellows, L. K., & Farah, M. J. (2005). Different underlying impairments in decision-making following ventromedial and dorsolateral frontal lobe damage in humans. *Cerebral Cortex*, *15*, 58–63.

Fernie, G., & Tunney, R. J. (2008). Decision making in the iowa gambling task. In G. R. Burthold (Ed.), *Psychology of decision making in legal, health care and science settings* (p. 69-99). New York, NY: Nova Science Publishers.

Fernie, G., & Tunney, R. J. (2013). Learning on the IGT follows emergence of knowledge but not differential somatic activity. *Frontiers in Psychology*, *4*, Article 687.

Finucane, M. L., Alhakami, A., Slovic, P., & Johnson, S. M. (2000). The affect heuristic in judgments of risks and benefits. *Journal of Behavioral Decision Making*, *13*, 1–17.

Fleming, S. M., & Dolan, R. J. (2010). Effects of loss aversion on post-decision wagering: Implications for measures of awareness. *Consciousness and Cognition*, *19*, 352–363.

Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J., & Rees, G. (2010). Relating introspective accuracy to individual differences in brain structure. *Science*, *329*, 1541–1543.

Fridberg, D. J., Queller, S., Ahn, W.-Y., Kim, W., Bishara, A. J., Busemeyer, J. R., ... Stout, J. C. (2010). Cognitive mechanisms underlying risky decision-making in chronic cannabis users. *Journal of Mathematical Psychology*, *54*, 28–38.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*, 523–551.

Gonzalez, C., Lerch, F. J., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*, 591–635.

Guillaume, S., Jollant, F., Jaussent, I., Lawrence, N., Malafosse, A., & Courtet, P. (2009). Somatic markers and explicit knowledge are both involved in decision-making. *Neuropsychologia*, *47*, 2120–2124.

Gupta, R., Duff, M. C., Denburg, N. L., Cohen, N. J., Bechara, A., & Tranel, D. (2009).

Declarative memory is critical for sustained advantageous complex decision-making. *Neuropsychologia*, *47*, 1686–1693.

Gureckis, T. M., & Love, B. C. (2009). Short-term gains, long-term pains: How cues about state aid learning in dynamic environments. *Cognition*, *113*, 293–313.

Gutbrod, K., Kroužel, C., Hofer, H., Müri, R., Perrig, W., & Ptak, R. (2006). Decision-making in amnesia: Do advantageous decisions require conscious knowledge of previous behavioural choices? *Neuropsychologia*, *44*, 1315–1324.

Haider, H., Eichler, A., & Lange, T. (2011). An old problem: How can we distinguish between conscious and unconscious knowledge acquired in an implicit learning task? *Consciousness and Cognition*, *20*, 658–672.

Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological Review*, *18*, 186-201.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523.

Higham, P. A. (2007). No special K! A signal detection framework for the strategic regulation of memory accuracy. *Journal of Experimental Psychology: General*, *136*, 1–22.

Horstmann, A., Villringer, A., & Neumann, J. (2012). Iowa gambling task: There is more to consider than long-term outcome. Using a linear equation model to disentangle the impact of outcome and frequency of gains and losses. *Frontiers in Neuroscience*, *6*, Article 61.

Kahneman, D. (2011). *Thinking, fast and slow.* New York, NY: Macmillan.

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

Kepecs, A., Uchida, N., Zariwala, H., & Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, *455*, 227–231.

Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, *324*, 759–764.

Kjome, K. L., Lane, S. D., Schmitz, J. M., Green, C., Ma, L., Prasla, I., ... Moeller, F. G. (2010). Relationship between impulsivity and decision making in cocaine dependence. *Psychiatry Research*, *178*, 299–304.

Koch, C., & Preuschoff, K. (2007). Betting the house on consciousness. *Nature Neuroscience*, *10*, 140–141.

Koriat, A. (2011). Subjective confidence in perceptual judgments: A test of the self-consistency model. *Journal of Experimental Psychology: General*, *140*, 117–139.

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

Kornell, N., Son, L. K., & Terrace, H. S. (2007). Transfer of metacognitive skills and hint seeking in monkeys. *Psychological Science*, *18*, 64–71.

Kunimoto, C., Miller, J., & Pashler, H. (2001). Confidence and accuracy of near-threshold discrimination responses. *Consciousness and Cognition*, *10*, 294–340.

Lane, S. D., Yechiam, E., & Busemeyer, J. R. (2006). Application of a computational decision model to examine acute drug effects on human risk taking. *Experimental and Clinical Psychopharmacology*, *14*, 254–264.

Lee, D. (2013). Decision making: From neuroscience to psychiatry. *Neuron*, *78*, 233–248.

Lejarraga, T., Dutt, V., & Gonzalez, C. (2010). Instance based learning: A general model of repeated binary choice. *Journal of Behavioral Decision Making*, *25*, 143–153.

Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, *116*, 286–295.

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology*, *103*, 1036–1047.

Lin, C.-H., Chiu, Y.-C., Lee, P.-L., & Hsieh, J.-C. (2007). Is deck B a disadvantageous deck in the Iowa Gambling Task? *Behavioral and Brain Functions*, *3*, Article 16.

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, *127*, 267–286.

Lorains, F. K., Dowling, N. A., Enticott, P. G., Bradshaw, J. L., Trueblood, J. S., & Stout, J. C. (2014). Strategic and non-strategic problem gamblers differ on decision-making under risk and ambiguity. *Addiction*, *109*, 1128–1137.

Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 3–26.

Lueddeke, S. E., & Higham, P. A. (2011). Expertise and gambling: Using Type 2 signal detection theory to investigate differences between regular gamblers and nongamblers. *The Quarterly Journal of Experimental Psychology*, *64*, 1850–1871.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

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*, *101*, 16075–16080.

Maia, T. V., & McClelland, J. L. (2005). The somatic marker hypothesis: Still many questions but no answers. *Trends in Cognitive Sciences*, *9*, 162–164.

Mealor, A., & Dienes, Z. (2012). No-loss gambling shows the speed of the unconscious. *Consciousness and Cognition*, *21*, 228–237.

Mellers, B. A., & McGraw, A. P. (2001). Anticipated emotions as guides to choice. *Current Directions in Psychological Science*, *10*, 210–214.

Middlebrooks, P. G., & Sommer, M. A. (2011). Metacognition in monkeys during an oculomotor task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 325–337.

Miller, S. A., Brownell, C. A., & Zukier, H. (1977). Cognitive certainty in children: Effects of concept, developmental level, and method of assessment. *Developmental Psychology*, *13*, 236–245.

Myung, I. J. (2000). The importance of complexity in model selection. *Journal of*

Mathematical Psychology, *44*, 190–204.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, *7*, 308–313.

Nevo, I., & Erev, I. (2012). On surprise, change, and the effect of recent outcomes. *Frontiers in Psychology*, *3*, Article 24.

Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*, *37*, 1–61.

Nieuwenhuis, S., & de Kleijn, R. (2011). Consciousness of targets during the attentional blink: A gradual or all-or-none dimension? *Attention, Perception, & Psychophysics*, *73*, 364–373.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231–259.

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, *53*, 139–154.

Overgaard, M. (2011). Visual experience and blindsight: A methodological review. *Experimental Brain Research*, *209*, 473–479.

Overgaard, M., Rote, J., Mouridsen, K., & Ramsøy, T. Z. (2006). Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Consciousness and Cognition*, *15*, 700–708.

Oya, H., Adolphs, R., Kawasaki, H., Bechara, A., Damasio, A. R., & Howard, M. A. (2005). Electrophysiological correlates of reward prediction error recorded in the human prefrontal cortex. *Proceedings of the National Academy of Sciences*, *102*, 8351–8356.

Pasquali, A., Timmermans, B., & Cleeremans, A. (2010). Know thyself: Metacognitive networks and measures of consciousness. *Cognition*, *117*, 182–190.

Persaud, N., Davidson, M., Maniscalco, B., Mobbs, D., Passingham, R. E., Cowey, A., & Lau, H. (2011). Awareness-related activity in prefrontal and parietal cortices in blindsight reflects more than superior visual performance. *NeuroImage*, *58*, 605–611.

Persaud, N., & McLeod, P. (2008). Wagering demonstrates subconscious processing in a binary exclusion task. *Consciousness and Cognition*, *17*, 565–575.

Persaud, N., Mcleod, P., & Cowey, A. (2007). Post-decision wagering objectively measures awareness. *Nature Neuroscience*, *10*, 257–261.

Peters, E., & Slovic, P. (2000). The springs of action: Affective and analytical anformation processing in choice. *Personality and Social Psychology Bulletin*, *26*, 1465–1475.

Peters, E., Västfjäll, D., Gärling, T., & Slovic, P. (2006). Affect and decision making: A "hot" topic. *Journal of Behavioral Decision Making*, *19*, 79–85.

Petrusic, W. M., & Baranski, J. V. (2003). Judging confidence influences decision processing in comparative judgments. *Psychonomic Bulletin & Review*, *10*, 177–183.

Pitt, M. A., Kim, W., & Myung, I. J. (2003). Flexibility versus generalizability in model selection. *Psychonomic Bulletin & Review*, *10*, 29–44.

Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in Cognitive Sciences*, *6*, 421–425.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, *109*, 472–491.

Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, *117*, 864–901.

Premkumar, P., Fannon, D., Kuipers, E., Simmons, A., Frangou, S., & Kumari, V. (2008). Emotional decision-making and its dissociable components in schizophrenia and schizoaffective disorder: A behavioural and MRI investigation. *Neuropsychologia*, *46*, 2002–2012.

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, *23*, 1–14.

Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, *3*, 1–23.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning & Verbal Behavior*, *6*, 855–863.

Reimann, M., & Bechara, A. (2010). The somatic marker framework as a neurological theory of decision-making: Review, conceptual comparisons, and future neuroeconomics research. *Journal of Economic Psychology*, *31*, 767–776.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In H. Black & W. F. Prokasy (Eds.), *Classical conditioning II* (pp. 64–99). New York, NY: Appleton-Century-Crofts.

Ruffman, T., Garnham, W., Import, A., & Connolly, D. (2001). Does eye gaze indicate implicit knowledge of false belief? Charting transitions in knowledge. *Journal of Experimental Child Psychology*, *80*, 201–224.

Sandberg, K., Bibby, B. M., & Overgaard, M. (2013). Measuring and testing awareness of emotional face expressions. *Consciousness and Cognition*, *22*, 806–809.

Sandberg, K., Bibby, B. M., Timmermans, B., Cleeremans, A., & Overgaard, M. (2011). Measuring consciousness: Task accuracy and awareness as sigmoid functions of stimulus duration. *Consciousness and Cognition*, *20*, 1659–1675.

Sandberg, K., Timmermans, B., Overgaard, M., & Cleeremans, A. (2010). Measuring consciousness: Is one measure better than the other? *Consciousness and Cognition*, *19*, 1069–1078.

Saver, J. L., & Damasio, A. R. (1991). Preserved access and processing of social knowledge in a patient with acquired sociopathy due to ventromedial frontal damage. *Neuropsychologia*, *29*, 1241–1249.

Schönberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *The Journal of Neuroscience*, *27*, 12860–12867.

Schurger, A., & Sher, S. (2008). Awareness, loss aversion, and post-decision wagering. *Trends in Cognitive Sciences*, *12*, 209–210.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*, 461–464.

Schwarz, N. (2000). Emotion, cognition, and decision making. *Cognition & Emotion*, *14*, 433–440.

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon ? *Psychological Science*, *15*, 720–728.

Seth, A. K., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L. (2008). Measuring consciousness: Relating behavioural and neurophysiological approaches. *Trends in Cognitive Sciences*, *12*, 314–321.

Shanks, D. R. (2005). Implicit learning. In K. Lamberts & R. L. Goldstone (Eds.), *Handbook of cognition* (pp. 202–220). London: Sage.

Shanks, D. R., & St. John, M. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, *17*, 367–447.

Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical bayesian methods. *Cognitive Science*, *32*, 1248–1284.

Slovic, P., Finucane, M. L., Peters, E., & MacGregor, D. G. (2002). The affect heuristic. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 397–420). New York, NY: Cambridge University Press.

Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34–50.

Speekenbrink, M., & Konstantinidis, E. (2014). Uncertainty and exploration in a restless bandit task. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 1491–1496). Austin, TX: Cognitive Science Society.

Steingroever, H., Wetzels, R., Horstmann, A., Neumann, J., & Wagenmakers, E.-J. (2013). Performance of healthy participants on the Iowa Gambling Task. *Psychological Assessment*, *25*, 180–193.

Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2013a). A comparison of reinforcement learning models for the Iowa gambling task using parameter space partitioning. *The Journal of Problem Solving*, *5*, Article 2.

Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2013b). Validating the PVL-Delta model for the Iowa gambling task. *Frontiers in Psychology*, *4*, Article 898.

Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (in press). Absolute performance of reinforcement-learning models for the Iowa gambling task. *Decision*.

Steyvers, M., Lee, M. D., & Wagenmakers, E.-J. (2009). A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, *53*, 168–179.

Stout, J. C., Busemeyer, J. R., Lin, A., Grant, S. J., & Bonson, K. R. (2004). Cognitive modeling analysis of decision-making processes in cocaine abusers. *Psychonomic Bulletin & Review*, *11*, 742–747.

Stout, J. C., Rodawalt, W. C., & Siemers, E. R. (2001). Risky decision making in Huntington's disease. *Journal of the International Neuropsychological Society*, *7*, 92–101.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction.* Cambridge, MA: The MIT Press.

Szczepanowski, R. (2010). Absence of advantageous wagering does not mean that awareness is fully abolished. *Consciousness and Cognition*, *19*, 426–431.

Szczepanowski, R., Traczyk, J., Wierzchoń, M., & Cleeremans, A. (2013). The perception of visual emotion: Comparing different measures of awareness. *Consciousness and Cognition*, *22*, 212–220.

Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, *112*, 647–661.

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, *315*, 515–518.

Tomb, I., Hauser, M., Deldin, P., & Caramazza, A. (2002). Do somatic markers mediate decisions on the gambling task? *Nature Neuroscience*, *5*, 1103–1104.

Tranel, D., Bechara, A., & Damasio, A. R. (1999). Decision making and the somatic marker hypothesis. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences* (pp. 1047–1061). Cambridge, MA: MIT Press.

Tunney, R. J., & Shanks, D. R. (2003). Subjective measures of awareness and implicit cognition. *Memory & Cognition*, *31*, 1060–1071.

Turnbull, O. H., Evans, C. E. Y., Bunce, A., Carzolio, B., & O'Connor, J. (2005). Emotion-based learning and central executive resources: An investigation of intuition and the Iowa Gambling Task. *Brain and Cognition*, *57*, 244–247.

Turnbull, O. H., Evans, C. E. Y., Kemish, K., Park, S., & Bowman, C. H. (2006). A novel set-shifting modification of the Iowa Gambling Task: Flexible emotion-based learning in schizophrenia. *Neuropsychology*, *20*, 290–298.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

Volkmann, J. (1934). The relation of the time of judgment to the certainty of judgment. *Psychological Bulletin*, *31*, 672–673.

Wagar, B. M., & Dixon, M. (2006). Affective guidance in the Iowa gambling task. *Cognitive, Affective, & Behavioral Neuroscience*, *6*, 277–290.

Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, *11*, 192–196.

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & Iverson, G. (2004). Assessing model mimicry using the parametric bootstrap. *Journal of Mathematical Psychology*, *48*, 28–50.

Wang, S., Krajbich, I., Adolphs, R., & Tsuchiya, N. (2012). The role of risk aversion in non-conscious decision making. *Frontiers in Psychology*, *3*, Article 50.

Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, *60*, 53–85.

Wetzels, R., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E.-J. (2010). Bayesian parameter estimation in the Expectancy Valence model of the Iowa gambling task. *Journal of Mathematical Psychology*, *54*, 14–27.

Wierzchoń, M., Asanowicz, D., Paulewicz, B., & Cleeremans, A. (2012). Subjective measures of consciousness in artificial grammar learning task. *Consciousness and Cognition*, *21*, 1141–1153.

Wierzchoń, M., Paulewicz, B., Asanowicz, D., Timmermans, B., & Cleeremans, A. (2014). Different subjective awareness measures demonstrate the influence of visual identification on perceptual awareness ratings. *Consciousness and Cognition*, *27*, 109–120.

Wild, C. J., & Seber, G. A. F. (1993). Comparing two proportions from the same survey. *The American Statistician*, *47*, 178–181.

Wood, S., Busemeyer, J. R., Koling, A., Cox, C. R., & Davis, H. (2005). Older adults as adaptive decision makers: Evidence from the Iowa Gambling Task. *Psychology and Aging*, *20*, 220–225.

Worthy, D. A., Hawthorne, M. J., & Otto, A. R. (2013). Heterogeneity of strategy use in the Iowa gambling task: A comparison of win-stay/lose-shift and reinforcement learning models. *Psychonomic Bulletin & Review*, *20*, 364–371.

Worthy, D. A., & Maddox, W. T. (2014). A comparison model of reinforcement-learning and win-stay-lose-shift decision-making processes: A tribute to W.K. Estes. *Journal of Mathematical Psychology*, *59*, 41-49.

Worthy, D. A., Pang, B., & Byrne, K. A. (2013). Decomposing the roles of perseveration and expected value representation in models of the Iowa gambling task. *Frontiers in Psychology*, *4*, Article 640.

Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, *12*, 387–402.

Yechiam, E., & Busemeyer, J. R. (2008). Evaluating generalizability and parameter consistency in learning models. *Games and Economic Behavior*, *63*, 370–394.

Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005). Using cognitive models to map relations between neuropsychological disorders and human decision-making deficits. *Psychological Science*, *16*, 973–978.

Yechiam, E., & Ert, E. (2007). Evaluating the reliance on past choices in adaptive learning models. *Journal of Mathematical Psychology*, *51*, 75–84.