Awareness in Decision-Making: Implicit Influences or Measurement Biases?

Emmanouil Konstantinidis (emmanouil.konstantinidis.09@ucl.ac.uk) & David R. Shanks (d.shanks@ucl.ac.uk)

Department of Cognitive, Perceptual and Brain Sciences
University College London
26 Bedford Way, WC1H 0AP, London,UK

Abstract

Can our decisions be guided by unconscious or implicit influences? According to the somatic marker hypothesis, emotion-based signals guide our decisions in uncertain environments outside awareness. However, evidence for this claim can be questioned on the grounds of inadequate assessments of conscious knowledge. Post-decision wagering, in which participants make wagers on the correctness of their decisions, has been recently proposed as an objective and sensitive measure of conscious content. In the experiments reported here, we employed variations of a classic decision-making paradigm, the Iowa Gambling Task, in combination with wagering in order to investigate the role played by unconscious influences. We also examined biases that affect wagering strategies such as the definition of the optimal strategy and loss aversion. Our 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.

Keywords: Iowa Gambling Task; decision-making; post-decision wagering; awareness; implicit learning; loss aversion.

Introduction

One of the most influential paradigms in the study of decision-making under uncertainty is the Iowa Gambling Task (IGT), a four-armed bandit task which requires participants to sample from decks of cards (labeled A-D) with different monetary payoffs (Bechara et al., 1994). Each card selection yields either a win or a combination of a win and loss. Overall, two of the decks are *bad*, leading to higher immediate wins but long-term losses (a net loss of $-\$25$ per card) whereas the remaining two are *good*, leading to lower immediate rewards but long-term wins (a net win of $+\$25$ per card). Another feature of the task is that the decks differ in the probability of loss with two having infrequent ($p=.1$) and two decks having frequent ($p=.5$) losses.

The IGT was initially designed to test empirically the somatic marker hypothesis (SMH), according to which bodily affective signals or *markers* can assist decision-making processes by marking response outcomes with an emotional signal, thus facilitating the selection of the most rewarding options in situations of uncertainty (see Damasio, 1996). A major assumption regarding somatic markers is that they can inform advantageous decision-making even when participants are not aware of the quality of their decisions.

Measures of Conscious Knowledge in the IGT

In a highly influential study, Bechara et al. (1997) assessed participants' knowledge by pausing the task every 10 trials and asking them to report whatever they knew and felt about the task. In addition, participants' electrodermal skin conductance responses (SCRs) were measured as an index of emotional arousal. The crucial finding was that participants selected cards from the good decks before they had conscious knowledge that those decks were the best. Importantly, the SCRs were higher prior to selections from the bad decks, suggesting that non-conscious biases (or markers) helped participants to avoid disadvantageous selections.

The assertion that our decisions are guided by implicit influences has been extensively criticized on the basis of weaknesses in the method that Bechara et al. (1997) used to probe awareness. Broad questions often fail to capture the rich spectrum of conscious knowledge, and moreover participants may not fully report their knowledge due to a conservative response criterion (see Newell & Shanks, in press). Following these methodological considerations, Maia and McClelland (2004) developed a more sensitive test of awareness in the form of a quantitative questionnaire. They asked their participants to rate each deck on a numerical scale, to report the expected wins and losses from each deck, and to indicate which deck they would pick if they could only choose from one deck for the rest of the task. Using this assessment, Maia and McClelland found that optimal performance on the IGT was accompanied by accurate reports of the decks' payoffs, indicating that there is little evidence that implicit or emotional biases are essential for successful learning of the task structure.

Post-Decision Wagering

Despite the apparent evidence in favor of conscious processing in the IGT, Bechara et al. (2005) suggested that the method of probing awareness used by Maia and McClelland (2004) was intrusive and so may render participants more rapidly aware of the task structure and thus undermine the role of somatic markers. In order to avoid the methodological inconsistencies created by verbal reports and numerical confidence ratings, Persaud, McLeod, and Cowey (2007) developed a novel non-verbal method of assessing awareness in the IGT, in which participants are required to make wagers after their card selections (post-decision wagering; PDW). The rationale is that if a participant maximizes her winnings through advantageous wagering (high wager after a good deck choice or low after a bad deck), then this is taken to indicate conscious knowledge.

Persaud et al. (2007) examined the influence of different types of questioning in combination with wagering in three different groups. In their version of the IGT, participants were instructed to wager £10 or £20 after each deck selection. Selections from the bad decks yielded a win of two times the wager whereas selections from the good decks returned the amount of the wager. The net outcome was ei-

ther a win of £75 (good decks) or a loss of the same amount (bad decks) per 10 cards if participants randomly allocated their wagers. Persaud et al. measured on which trial good deck selection and advantageous wagering began and conjectured that if a significant difference between these measures emerged, with performance revealing a preference for good decks before advantageous wagering emerged, this would indicate an unconscious influence on decision-making. In fact, the groups who were asked no questions or the questions of Bechara et al. (1997) demonstrated a lag between good deck selection and advantageous wagering, indicating a dissociation between performance and awareness. On the other hand, when awareness was probed by the questionnaire of Maia and McClelland (2004), performance and advantageous wagering developed simultaneously. These findings led Persaud et al. to conclude that there are implicit influences on decision-making which are masked if the measure of awareness makes participants aware of the nature of the task.

Overview of the Experiments

Although post-decision wagering holds much promise as a method of measuring awareness, it has received many criticisms (e.g., Clifford, Arabzadeh, & Harris, 2008; Fleming & Dolan, 2010; Schurger & Sher, 2008). In the present study, we will focus on two problematic and rather contradictory aspects of the method: the definition of the optimal strategy and loss aversion. Paradoxically, the optimal strategy is always to wager high as this strategy will give the same outcome if 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 (2007) payoff matrix, as it is either no worse than wagering low, or better. Consequently, a rational participant would always wager high, regardless of her knowledge about the task.

A second important issue is the influence of loss aversion in wagering strategies. According to prospect theory (Kahneman & Tversky, 1979), people have an asymmetric conceptualisation of wins and losses; for example, the prospect of losing £5 is considered to be of greater psychological magnitude than that of winning the same amount. Subjective measures of awareness require participants to set a criterion about, for example, whether to wager high or low. Hence, any criterion may be modulated by cognitive biases such as loss aversion.

The present study examines the validity and sensitivity of post-decision wagering by looking closely at the two aforementioned problematic aspects and by measuring participants' awareness in combination with Maia and McClelland's (2004) quantitative questions.

Experiment 1: The Weakly Dominant Strategy

In order to overcome the problem that wagering high is the rational strategy irrespective of the acquired knowledge, Clifford et al. (2008) proposed a solution by modifying the original payoff matrix used by Persaud et al. (2007) (see Table 1). In this modified matrix participants are discouraged

Table 1: Payoff matrices for the different combinations of deck selection and wager.

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

from wagering high all the time. This can be shown by computing the expected payoff from wagering low which is $+1/2$ $[= (+2 - 1)/2]$ when deck selection is random 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 higher 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 maximize gains) would start to wager high only when her discrimination or level of awareness 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 ($5 - 2 = 3$).

We also addressed the question of how sensitive wagering is compared to another measure of awareness by using the quantitative questions of Maia and McClelland (2004).

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, dependent on their performance in the task.

Task A probabilistic variation of the original IGT was employed. There were four decks of cards with identical physical appearance, labelled A-D. The payoff structure of each deck was different to the original IGT; the matrix of Clifford et al. (2008) was used as the basis of 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. Based on the contingencies in Table 1, on the majority of trials a payoff of +2 was associated with a good deck selection and a low wager. Similarly, the combination of a bad deck and high wager produced a payoff of -5. On a minority of trials the signs of the payoffs were reversed (e.g., the payoff was -2 for good deck selection/low wager). For decks A and C, the majority outcome occurred on a randomly-chosen 75% of trials while for decks B and D, the majority outcome occurred on 60% of trials, resulting in different overall expected payoffs for each deck. In contrast to the original IGT, the outcome on each trial was either a net win or a loss and participants could win or lose points, which translated into money at the end of the experiment.

The task comprised 100 card selections. After each card selection, participants could place a wager, either High or Low.

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. 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, which was administered every 20 trials.

Results

Choice and Wagering In order to examine whether decision-making strategies in the IGT are dissociable from awareness we computed the average proportion of good deck selections (decks C and D) and advantageous wagering (either a high wager after a good deck or a low wager after a bad deck) over successive blocks of 10 trials. If learning of the good decks emerged earlier than awareness (as indexed by wagering) then that would indicate an unconscious influence on decision-making strategies in the IGT. However, performance exceeded the chance level on block 1 for both measures (Choice: $t(20) = 2.83, p = .01$, 95% CI [0.52, 0.65], Wagering: $t(20) = 3.80, p < .001$, 95% CI [0.57, 0.73]) (see Figure 1A). 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.

A repeated-measures ANOVA on good deck selections with polynomial contrasts revealed significant linear, $F(1, 20) = 72.38, MSE = 3.12, p < .001, \eta_g^2 = 0.41$, and quadratic trends, $F(1, 20) = 7.34, MSE = 1.83, p = .013, \eta_g^2 = 0.04$. Wagering performance closely followed the optimal decision-making strategy as demonstrated by a linear effect of block, $F(1, 20) = 19.57, MSE = 4.26, p < .001, \eta_g^2 = 0.19$. These results indicate that participants favored the good decks and became gradually capable of maximizing their winnings by placing appropriate wagers. Since choice and wagering displayed similar patterns there is no evidence of a dissociation between learning and awareness of the optimal strategy, a pattern which is at odds with the main claim of Persaud et al. (2007) about learning without awareness in the IGT.

An important feature of our task variant is that each deck yielded different overall expected outcomes. If participants had explicit insight about their choices, they could discriminate not only between good and bad decks but also within each pair of decks (A vs B and C vs D). More selections from the most rewarding of the good decks (deck C) would suggest that participants possess substantial knowledge about the payoff structure of the task. Figure 1B shows that participants indeed selected more cards from the deck with the highest win
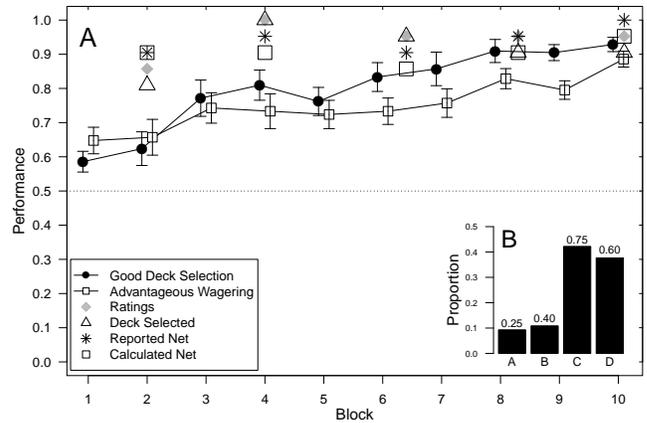

Figure 1: (A) Proportion of good deck selections and advantageous wagering across block of trials (lines) and proportion of participants who favored the good decks in the quantitative questions (markers). Error bars are ±1 SEM. (B) Overall proportion of deck selections. The win probability associated with each deck is depicted on the top of each bar.

probability (.75) and there was a significant preference over the second best deck, $t(20) = 3.98, p < .001, d = 0.87$.

Questionnaire Two of the questionnaire measures, in which participants provided a rating for each deck and indicated which deck they would select cards from for the remainder of the task (Ratings and Deck Selected in Figure 1), reflect knowledge about the general quality of the decks. The remaining two measures, in which participants reported the average net outcome from each deck (Reported Net), refer to the actual payoffs. Participants also reported the average win and loss and the frequency of losses, based on which a net amount for each deck is obtained (Calculated Net).

Figure 1 (markers) shows the proportion of participants whose answers favored the good decks in each of the questionnaire measures. Participants whose verbal responses did not discriminate between good and bad decks (i.e., they gave the same ratings or 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.

Another way of examining the two measures is to look at participants' deck selections and wagering in the trials following the administration of the questionnaire (trials 21, 41, 61, 81) in order to examine whether participants who behaved advantageously had knowledge of the advantageous strategy. Specifically, we are interested in the verbal reports and wagers of those participants who behaved advantageously (i.e., selected good decks) in these trials. Figure 2 shows that the majority of participants who made good deck selections also demonstrated knowledge of the advantageous strategy in all the questionnaire items. However, wagering underestimates the acquired knowledge in all trials following the question-

naire compared to the verbal reports (the figure also shows wagers on trial 100, immediately prior to the final set of questions). Thus, it is evident that the detailed and structured questions revealed higher levels of awareness compared to wagering.
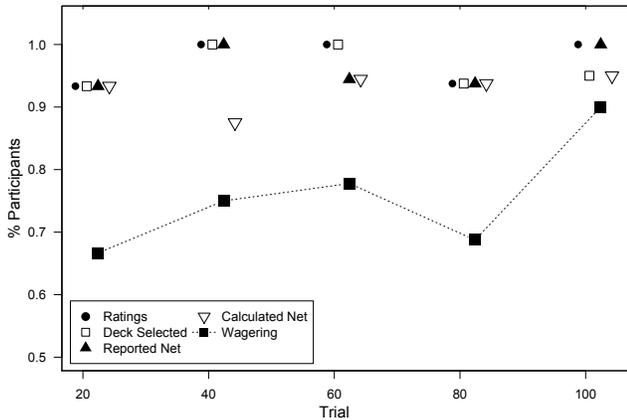

Figure 2: Percentage of participants who showed knowledge of the advantageous strategy in the questionnaire items versus in their wagers. Wagering indicates the percentage of participants who made an advantageous wager (high on a good deck choice) on the trial immediately following the administration of the questionnaire.

Discussion

The key point of Experiment 1 is that wagering did not lag behind the selection of good decks, with both measures becoming reliably better than chance very early in the task. Under the conditions tested here, awareness as measured by wagering tracked deck selections quite closely. In addition, the results of the quantitative questions revealed that there was actually conscious knowledge that was left undetected by wagering as participants' wagers were less sensitive than their verbal reports. This places a question mark over the validity of post-decision wagering as a valid and sensitive method of assessing awareness.

An interesting finding is the early onset of learning and awareness (Block 1) which can be primarily 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 immediately at the start of the task) normal controls selected more cards from the good decks even in the first 20 trials and they kept on choosing the good decks throughout the task. This setup of the payoff schedule removes the reversal learning component of the IGT which can delay the learning of the optimal decisions.

Experiment 2: Dealing with Loss Aversion

Depending on the payoff matrix, participants may employ different response criteria to place high or low wagers which can make detection of 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, several studies have shown that loss

aversion affects awareness as indexed by wagering (e.g., Dienes & Seth, 2010; Fleming & Dolan, 2010).

Schurger and Sher (2008) proposed that the design of the payoff matrix should take into account participants' tendency to evaluate losses worse than equivalent wins. Unlike Clifford et al.'s matrix, which encourages low wagers when certainty is low, "subjects seem to need precisely quite the opposite sort of encouragement" (Schurger & Sher, 2008, p. 209). Table 1 shows the matrix devised by Schurger and Sher as a means to counter loss aversion. Specifically 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 under uncertainty, its weights regarding high wagers are two times bigger compared to the matrix of Clifford et al. 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, employment 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.

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. There were four decks of cards each having 100 associated wins and losses, one for each trial. A randomly drawn (win or loss) value 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 dis-

tributed 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.

Results

Performance exceeded chance on block 1 for both measures (Choice: $t(29) = 2.39, p = .023$, 95% CI [0.51, 0.64], Wagering: $t(29) = 2.52, p = .018$, 95% CI [0.51, 0.61]) (see Figure 3). This result indicates that participants' optimal decision-making and learning 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.
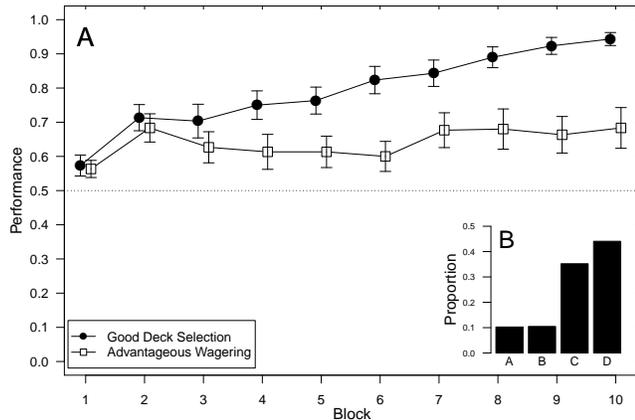

Figure 3: (A) Proportion of good deck selections and advantageous wagering across block of trials. Error bars are ± 1 SEM. (B) Overall proportion of deck selections.

Two separate repeated-measures ANOVAs were conducted to investigate the progression of good deck selections and advantageous wagering across blocks. Polynomial contrasts revealed a significant linear effect of block on the proportion of good deck selections, $F(1, 29) = 143.40, MSE = 2.34, p < .001, \eta_g^2 = 0.37$. However, the same trend was not observed on the proportion of advantageous wagers as the linear effect was not significant, $F(1, 29) = 2.06, MSE = 8.44, p = .16$. Even though wagering was above chance even 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 large losses. Additionally, advantageous wagering was above chance in all blocks of trials.

Figure 3B 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 ob-

served. 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.

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 maximize 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 the maximum point 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 the reasoning behind participants' wagering strategies.

General Discussion

The purpose of the present article was twofold: first, to evaluate how sensitive and direct post-decision wagering is as a measure of awareness, and secondly to investigate whether there are implicit influences on decision-making under uncertainty or whether past suggestive results have been the by-product of use of insensitive measures.

We examined two main response biases, dominance and loss aversion, which arise from the design of the payoff matrix. In both cases, there was a direct effect of the design of the payoff matrix on the wagering strategies that participants employed. In Experiment 1, no delay was observed in the onset of awareness relative to deck selection, as advantageous wagering closely followed learning of the good decks. The early onset of awareness can be accounted for by either of two factors: first, the values in the payoff matrix which may have encouraged low wagers, and secondly, the quantitative questions which might have influenced the development of awareness of the deck values. However, 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 2, 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 maximize 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.

Even though the decision-making paradigms employed in our experiments are not identical to the original IGT, they maintain its key characteristics such as sequential choices, rewards, punishments, and uncertainty, while removing some of its problematic features such as participants' tendency to focus on decks with infrequent losses, the lack of exploration-exploitation, and the dual presentation of the outcomes. These latter issues make the interpretation of IGT data quite difficult (see Steingroever et al., 2013). Nevertheless, future work could seek to determine whether advantageous deck selection and wagering develop in parallel when our probabilistic version includes the reversal learning element of the original IGT task.

The claim that unconscious biases are essential for successful performance in the IGT has not been confirmed in either of the experiments reported here. In fact, participants' advantageous selections developed in parallel with their conscious knowledge, as revealed by verbal reports and wagering. Learning of the task contingencies through emotional markers cannot be ruled out, but does not seem to precede conscious evaluation of the experienced outcomes. In other words, the activation of an unconscious emotional system which provides critical information for decision-making processes is not essential or sufficient to explain learning in the IGT. Dunn, Dalgleish, and Lawrence (2006) suggested that there is little evidence to support the view that deck contingencies are consciously impenetrable, and that what needs to be tested is whether participants have an explicit understanding of the payoff schedule or whether they can simply discriminate the quality of the decks. Our questionnaire results demonstrate that participants not only were able to show a general preference for the good decks but that in addition they could justify their preferences by accurately reporting the average wins and losses associated with each deck. This knowledge develops very early in the task, leaving little room for unconscious influences.

References

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., 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.

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

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.

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.

Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience & Biobehavioral Reviews*, *30*, 239-271.

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.

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.

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

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.

Newell, B. R., & Shanks, D. R. (in press). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*.

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

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

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.