

A reexamination of the evidence for the somatic marker hypothesis: What participants really know in the Iowa gambling task

Tiago V. Maia*[†] and James L. McClelland**

*Department of Psychology and [†]Center for the Neural Basis of Cognition, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213

Contributed by James L. McClelland, September 9, 2004

Bechara, Damasio, and coworkers [Bechara, A., Damasio, H., Tranel, D. & Damasio, A. R. (1997) *Science* 275, 1293–1295] have reported that normal participants decide advantageously before knowing the advantageous strategy in a simple card game designed to mimic real-life decision-making. Bechara *et al.* have used this result to support their view that nonconscious somatic markers can guide advantageous behavior. By using more sensitive methods, we show that participants have much more knowledge about the game than previously thought. In fact, participants report knowledge of the advantageous strategy more reliably than they behave advantageously. Furthermore, when they behave advantageously, their verbal reports nearly always reveal evidence of quantitative knowledge about the outcomes of the decks that would be sufficient to guide such advantageous behavior. In addition, there is evidence that participants also have access to more qualitative reportable knowledge. These results are compatible with the view that, in this task, both overt behavior and verbal reports reflect sampling from consciously accessible knowledge; there is no need to appeal to nonconscious somatic markers. We also discuss the findings of other studies that similarly suggest alternative interpretations of other evidence previously used to support a role for somatic markers in decision-making.

Damasio (1) has suggested that normal decision-making in humans is often assisted by somatic markers: bodily states (or brain representations thereof) that correspond to emotional reactions to possible courses of action, effectively reflecting the goodness or badness of the outcomes associated with each course of action. According to Damasio and coworkers (1–3), such markers can operate not only consciously (when one has a “gut feeling” about the goodness or badness of a possible course of action) but also nonconsciously. They further claim that in the latter case somatic markers can even lead people to make advantageous decisions before they are consciously aware of which decisions are advantageous (3).

A large part of the support for the somatic marker hypothesis has come from an influential study in which Bechara, Damasio, and colleagues (3) reported that normal participants “[decide] advantageously before knowing the advantageous strategy” in a simple card game designed to mimic real-life decision-making. According to Bechara *et al.* (3), normal participants started to make the right selections in that game before they had conscious knowledge that those were the best selections. Furthermore, when participants were about to make a bad selection, they exhibited higher skin conductance responses (SCRs) than when they were about to make a good selection, again seemingly before they had conscious knowledge about which were the good and bad selections. Bechara, Damasio, and colleagues (2, 3) took these findings to support the somatic marker hypothesis, claiming that the skin conductances reflect somatic markers that allow participants to make advantageous selections even before conscious knowledge is available. Here, we show that, in fact, players have extensive conscious knowledge about the game, as indicated in verbal reports obtained with a more sensitive questionnaire than that used by Bechara *et al.* (3). Indeed, participants’

verbal reports indicate knowledge of the advantageous strategy more reliably than their actual behavior does, and when they behave advantageously, they nearly always report knowledge about the outcomes of the decks that would be sufficient to guide such advantageous behavior. Thus, our data provide no reason to posit that nonconscious biases guide advantageous behavior in this task before knowledge that is consciously accessible does. We also provide evidence that supports the view that the contrary conclusion of Bechara *et al.* (3) arose from their use of methods that were not sufficiently powerful to uncover all of the knowledge that participants had about the game. Our findings undermine one of the main pillars of support for the somatic marker hypothesis. As we will describe below, additional recent findings in the literature raise questions about the other main lines of evidence that have been taken to support this hypothesis.

In the game used by Bechara *et al.* (3), henceforward referred to as the Iowa gambling task (IGT), participants must select, on each trial, a card from one of four decks (4). On every card, participants win some play money. For two of the decks, the winning amount is always \$100, and, for the other two, the winning amount is always \$50. However, on some cards, in addition to winning money, participants also lose money. The schedule of losses is such that, in the long run, the decks that give \$100 rewards produce net losses, whereas the decks that give \$50 rewards produce net gains. The long-term expected net loss for the two \$100 decks is the same (a loss of \$25 per trial); they differ only in terms of the frequency and magnitude of the losses, with one deck having larger but less frequent losses than the other. The situation for the two \$50 decks is similar: Both have the same expected net gain (\$25 per trial), but one has larger and less frequent losses than the other. The game ends after 100 trials. The contingencies and the duration of the game are not known by participants in advance; participants are told simply that their goal is to have the best possible net outcome in the game.

To assess participants’ knowledge about the game, Bechara *et al.* (3) interrupted it after the first 20 trials and then at 10-trial intervals and asked participants, “Tell me all you know about what is going on in this game,” and, “Tell me how you feel about this game.”

An extensive body of work in the implicit-learning literature has shown that such broad, open-ended questions often fail to identify all of the conscious knowledge that participants have acquired in performing a task (5, 6). There are many reasons for such failures. For example, the questions may not reliably cue recall of all relevant knowledge, or participants may fail to report knowledge that they still consider tentative. Furthermore, the amount of knowledge that participants volunteer in answer to these questions may depend heavily on such factors as person-

Abbreviations: IGT, Iowa gambling task; SCR, skin conductance response; VMPFC, ventromedial prefrontal cortex.

[†]To whom correspondence should be addressed. E-mail: tmaia@cmu.edu.

© 2004 by The National Academy of Sciences of the USA

ality or level of engagement. These observations raise the possibility that participants may actually have conscious knowledge of the advantageous strategy as early as they behave advantageously in this game and that such knowledge may be uncovered by a more sensitive test.

Another issue that requires reconsideration is what it means to decide advantageously in the game. The analyses of Bechara *et al.* (3) contrast the number of selections from “good decks” and “bad decks,” where the good decks are the decks that give only \$50 rewards but lead to positive net outcomes in the long run and the bad decks are the \$100 decks, which end up giving net losses. However, the sequence of wins and losses is fixed for each deck and is the same across participants, and the arrangement of cards is such that, early in the game, the decks that produce the best net outcome are the \$100 decks (4). Obviously, each participant has nothing but his or her own experience to go by, so we follow the general practice in the decision-making (7, 8) and computational reinforcement learning literature (9) in considering advantageous behavior to be behavior that follows the net outcomes experienced up until the trial in which a decision is being made. In our analyses, we therefore determine for each participant and each trial which two decks are best and which two decks are worst, based on the mean net outcomes the participant has had with each deck up until that trial. Note that, like Bechara *et al.* (3), we do not distinguish between the two best decks or the two worst decks, because, as mentioned above, the task is structured in such a way that, over the long run, the two best decks have the same expected net result, as do the two worst decks. Thus, we define advantageous behavior as choosing one of the two decks with the highest observed mean net outcome.

Finally, it is important to define precisely what is meant by “knowing the advantageous strategy.” We take knowing the advantageous strategy to mean having conscious knowledge that would support the choice of one of the two best decks (as determined by the observed mean net outcome). We operationalize conscious knowledge as knowledge that can be reported verbally, and we identify three possible levels of conscious knowledge of the advantageous strategy:

- Level 0. The participant does not have any conscious knowledge specifying a preference for one of the two best decks.
- Level 1. The participant has conscious knowledge specifying a preference for one of the two best decks but does not have conscious knowledge about the outcomes of the decks that could provide a basis for that preference.
- Level 2. The participant has conscious knowledge specifying a preference for one of the two best decks and has conscious knowledge about the outcomes of the decks that could provide a basis for that preference.

These three levels are closely related to the “pre-hunch,” “hunch,” and “conceptual” periods, respectively, from Bechara *et al.* (3). The claim of Bechara *et al.* (2, 3) is that participants behave advantageously even when their knowledge is still at Level 0 [although, even in their paper (3), the data that would support this claim did not reach statistical significance]. In particular, they claim that “nonconscious biases guide behavior before conscious knowledge does” (3) and that “this biasing effect occurs even before the subject becomes aware of the goodness or badness of the choice s/he is about to make” (2). A weaker claim that would still be of interest would be that participants behave advantageously when their knowledge is at Level 1.

Materials and Methods

Assessment of Conscious Knowledge. To assess participants’ level of knowledge throughout the game, we developed a more sensitive test of awareness in the form of a structured questionnaire (Fig.

Q1. Rate, on a scale of –10 to +10, how good or bad you think deck 1 is, where –10 means that it is terrible and +10 means that it is excellent.

Q2. Okay, why did you rate deck 1 with ...?

[Repeat questions Q1 and Q2 for decks 2 through 4.]

Q3. In answering the questions that follow, consider the following definitions. Your “winning amount” for a trial is the amount you won on that trial. Your “loss” on a trial is the amount you lost on that trial. Your “net result” for a trial is the amount you won minus the amount you lost on that trial. Do you understand these definitions and the differences between the three terms? [If not, explain again using examples.]

Okay, now suppose you were to select 10 cards from deck 1.

Q3.1. What would you expect your average net result to be?

Q3.2. What would you expect your average winning amount to be?

Q3.3. In how many of the 10 trials would you expect to get a loss (not necessarily a net loss)?

Q3.4. For those trials in which you would get a loss, what would you expect the average loss to be?

[Repeat question Q3 for decks 2 through 4.]

Q4. Okay, now tell me, on a scale of 0 to 100, how much you think that you know what you should do in this game in order to win as much money as possible (or, if you can't win, to avoid losing money as much as possible). 0 means that you have no idea of what you should do and feel that you still need to explore the game more and 100 means that you know exactly what you should do and have no doubts that that would be the best strategy.

Q5. 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?

Fig. 1. Questionnaire.

1). As in the study of Bechara *et al.* (3), participants were asked these questions after the first 20 trials and then every 10 trials. Note that care has been exercised to attempt to minimize carry-over effects from question to question within a question period. Specifically, participants are asked about the ratings (question Q1) before they are asked about the expected outcomes for the decks (question Q3), to minimize the influence of the latter on the former. Similarly, question Q3.1 is asked before questions Q3.2–Q3.4, so that participants’ answers about the expected net result are not based on their answers to questions Q3.2–Q3.4.

If a participant’s knowledge is at Level 1, such knowledge should be reflected in the answers to questions Q1 and Q5. Q1 asks for a simple numerical rating of each deck. The response is taken to indicate Level 1 knowledge if the deck with the highest rating is one of the two best decks according to the participant’s experience up to that point. Q5 asks the participant to indicate the deck that he or she would choose if he or she could only select from that deck for the rest of the experiment. Again, the response is taken to indicate Level 1 knowledge if the participant names one of the two best decks. If a participant’s knowledge is at Level 2, in addition to having knowledge of which deck is best, he or she also should have additional reportable knowledge about the outcomes associated with the various decks that would justify that conclusion. A quantitative assessment of this knowledge is provided by the questions in Q3. Question Q3.1 directly assesses participants’ knowledge of the expected net for each deck. The response is taken to indicate Level 2 knowledge if the participant attributes the highest expected net to one of the two best decks. Questions Q3.2–Q3.4 assess participants’ knowledge of the outcomes of each deck in terms of each deck’s reward value, probability of getting a loss, and mean loss value. Note that questions Q3.2–Q3.4 allow one to calculate the mean net that a participant should expect, based on his or her knowledge about the outcomes of the decks. We call this value the “calculated

net”; for each question period q , participant p , and deck d , it is obtained from the participant’s answers to questions Q3.2, Q3.3, and Q3.4 by using the following formula [in which $CN(q, p, d)$ refers to the calculated net in question period q for participant p and deck d , and $Q3.2(q, p, d)$, $Q3.3(q, p, d)$, and $Q3.4(q, p, d)$ represent the answers of participant p to the corresponding questions for deck d , in question period q]:

$$CN(q, p, d) = Q3.2(q, p, d) + \frac{Q3.3(q, p, d)}{10} \times Q3.4(q, p, d).$$

The responses to questions Q3.2–Q3.4 are taken to indicate Level 2 knowledge if the highest calculated net is for one of the two best decks.

It is important to note that participants at Level 2 do not necessarily keep track throughout the game of the numerical estimates that they give us in answer to the questions in Q3. Participants may produce such estimates only at the time the questions are asked, possibly by sampling relevant exemplars from memory. We cannot assume that participants’ actual behavior is based directly on this type of quantitative information. Nevertheless, if whenever they behave advantageously, participants can show evidence of knowledge of the advantageous strategy in their answers to the Level 2 quantitative measures (the reported net and the calculated net), we can conclude that they have consciously accessible knowledge that would provide a sufficient basis for such advantageous behavior. Because participants’ estimates may be noisy, in practice the reported net and the calculated net do not necessarily coincide.

It is also important to note that our questionnaire focuses primarily on reportable quantitative knowledge. Unfortunately, such measures may fail to capture reportable qualitative knowledge (e.g., “this deck is bad because once in a while it gives me a very large loss”), which might nevertheless be sufficient to guide advantageous behavior. We included question Q2 in the questionnaire to attempt to uncover some of that qualitative knowledge, if that proved necessary. However, our focus is primarily on the quantitative questions. We relied on the qualitative knowledge conveyed by Q2 only to assess the knowledge of a specific participant who behaved advantageously but whose answers to Q3 did not reveal reportable quantitative knowledge of the advantageous strategy. We also included a question that assessed participants’ overall confidence in their understanding of the game (Q4) to address issues unrelated to those discussed in this paper.

Experimental Conditions. In addition to the condition with this questionnaire, our experiment also included a second condition that consisted of a direct replication of the study of Bechara *et al.* (3) by using their original questions. Participants were randomly assigned to either condition.

Participants. There were 20 participants per condition, all of whom were undergraduate students at Carnegie Mellon University. Participants received course credit for their participation.

Results

As a preliminary to our main analyses, we compared the overall performance of participants in our questionnaire condition and in the replication condition. Details are provided in Section 1 of the supporting information published on the PNAS web site. This analysis found that there was no statistically significant difference in overall performance in the two conditions, as assessed by the total number of cards selected from the \$100 and \$50 decks (two-sided $P = 0.28$). The overall profile of selections




Fig. 2. Participants’ knowledge that one of the two best decks is the best deck, as reflected in several verbal report measures, compared with participants’ tendency to behaviorally select from one of the two best decks. (As mentioned in the text, we define the two best decks to be the two decks with the highest observed mean net outcome, according to each individual participant’s sequence of observations up until the trial under consideration.) The green line shows how many participants actually picked one of the two best decks behaviorally. The red and cyan markers correspond, respectively, to the number of participants who gave the highest rating to one of the two best decks if they could only select from one deck. The light-brown marker corresponds to the number of participants who gave the highest expected net to one of the two best decks, and the dark-blue marker corresponds to the number of participants who had the highest calculated net for one of the two best decks. (Note that on trial 70, the light-brown marker is covered by the dark-blue marker.)

is also very similar across the two conditions (see Fig. 4 in Section 1 of the supporting information). These results suggest that our use of a more detailed questionnaire did not have a significant impact on participants’ acquisition of the advantageous strategy. Section 2 of the supporting information provides a detailed analysis of participants’ responses to the questions in the replication condition. This analysis shows that by using the methods of Bechara *et al.* (3), we replicated their statistically significant results; specifically, participants behaved advantageously when they were classified according to the criteria of Bechara *et al.* as being in either the hunch or conceptual periods (our Levels 1 and 2, respectively).

We now consider what our questionnaire condition reveals about participants’ knowledge of the advantageous strategy. Fig. 2 shows how many participants showed evidence of knowledge of the advantageous strategy in the different verbal report measures, as well as how many participants behaved advantageously, on a trial-by-trial basis. As can be seen, in every period in which participants’ knowledge was probed, all of the verbal report measures demonstrate knowledge of the advantageous strategy for the majority of participants. In fact, the tendency is for knowledge of the advantageous strategy to be more evident in all of the verbal report measures than in behavior (which may be due to exploration of the different decks or risk-taking by some participants).

Fig. 3 shows that virtually all of the participants who behaved advantageously on a given trial showed knowledge of the advantageous strategy in their answers to questions Q1 and Q5 on that trial. Further analysis revealed that a participant who behaved randomly, referred to as participant 41, accounted for




Fig. 3. Participants’ knowledge that one of the two best decks is the best deck, as reflected in several verbal report measures, among participants who behave advantageously. The markers have the same meaning as in Fig. 2, but, rather than referring to the total number of participants, they refer to the percentage of participants, among those who behaved advantageously on the corresponding trial, who showed evidence of knowledge of the advantageous strategy in each of the verbal report measures. (Note that on trials 30 and 50, the light-brown marker is covered by the dark-blue marker.)

two of the cases in each of these measures in which knowledge of the advantageous strategy was not shown, suggesting that in these cases the apparently advantageous behavior occurred by chance (see Section 3 of the supporting information). This observation leaves only two additional cases in each measure unexplained. However, these cases occurred in the early trials (20 and 30), in which typically there is still a fair amount of exploration and may well have been cases in which participants behaved advantageously by chance. In summary, in virtually all cases, participants were at least at Level 1 of knowledge when they behaved advantageously.

Fig. 3 further shows that, in the vast majority of cases, participants who behaved advantageously also showed knowledge of the advantageous strategy in both quantitative measures of Level 2 knowledge: the reported net and the calculated net. However, it is also apparent in the figure that, on several trials, a small number of participants failed to show evidence of knowledge of the advantageous strategy in at least one of these two measures. The first question of interest is whether there were participants who behaved advantageously but did not show evidence of knowledge of the advantageous strategy in either quantitative measure of Level 2 knowledge. We restrict this analysis to participants who exhibited Level 1 knowledge to exclude from consideration the very small number of cases mentioned above in which participants seem to have behaved advantageously by chance early in the game, when exploration of the different decks was still prominent.

It turns out that only two participants fulfill these conditions. One is the aforementioned participant 41, who behaved randomly during most of the game. The outcome of appearing to behave advantageously on some trials while not showing Level 2 knowledge of the advantageous strategy would be expected to occur occasionally by chance, even if all responses were essentially random, so this participant will not be discussed further. The other participant (participant 36) does not seem to have behaved randomly. This participant behaved advantageously and demonstrated Level 1 knowledge in three question periods (trials 40, 70, and 80) in which neither quantitative Level 2 measure reflected knowledge of the advantageous strategy. However, an

analysis of this participant’s answers to question Q2 revealed that she had qualitative Level 2 knowledge that was not reflected in her quantitative answers and which seems sufficient to have guided her advantageous behavior and her answers to Level 1 questions (see Section 4 of the supporting information).

In addition to participant 36, there were on some trials a small number of participants that behaved advantageously and demonstrated Level 1 knowledge, but showed knowledge of the advantageous strategy in just one of the two quantitative measures of Level 2 knowledge: the reported expected net (zero to three participants per question period; mean, 1.0) or the calculated net (zero to two participants per question period; mean, 1.11). These participants cannot be definitively classified as being at Level 1 or 2, because they demonstrated inconsistent conscious knowledge in the quantitative Level 2 measures. These inconsistencies could have resulted from estimates that are based on noisy memory-sampling processes. If behavioral selections also result from an incomplete sampling of the same knowledge, this could lead to the observed advantageous behavior (or to cases in which participants behave disadvantageously despite answering the questions in accordance with knowledge of the advantageous strategy). It also should be noted that these participants might have had qualitative knowledge similar to that of participant 36, which could also have guided their advantageous behavior. (See also Section 5 of the supporting information for additional discussion concerning Fig. 3.)

As mentioned above, question Q4 was included in the questionnaire to address issues unrelated to those discussed in this paper. Nevertheless, for completeness, Section 6 of the supporting information provides a brief analysis of the answers to this question. This analysis reveals that participants’ certainty about what they should do in the game increased gradually as the game progressed.

In summary, in the overwhelming majority of cases, when participants behaved advantageously, they exhibited Level 2 knowledge in both the reported and the calculated nets. In a small number of cases, participants only showed knowledge of the advantageous strategy in one of these two measures, effectively reporting inconsistent knowledge. Nevertheless, a sample from such knowledge could have provided the basis for their advantageous behavior. Finally, there was a single participant who, on any trial, behaved advantageously but did not show evidence of knowledge of the advantageous strategy in either the reported or the calculated net. However, this participant had qualitative Level 2 knowledge that also could have provided the basis for her behavior.

Discussion

We have found that when participants behave advantageously in the IGT, (i) they have conscious access to the relative goodness and badness of the decks, and (ii) they have explicit, reportable knowledge that could provide the basis for such judgments and behavior. We therefore have found no support for the claims of Bechara *et al.* that, in this task, “nonconscious biases guide behavior before conscious knowledge does” (3) and that such biases occur “before the subject becomes aware of the goodness or badness of the choice s/he is about to make” (2).

It is important to note that, even though we have shown that participants’ reportable knowledge is sufficient to explain their advantageous behavior, the extent to which the participants actually based their behavior on knowledge held in conscious form at the time of choice remains an open question. The fact that the participants generated conscious reports when asked to do so does not *per se* imply that such conscious knowledge played a causal role in their actual decisions. Many models are consistent with our results, including not only a model in which conscious knowledge guides behavior, but also, for example, one in which the same knowledge store that participants canvass to

generate verbal reports also can directly feed a response-selection mechanism without the need for conscious intermediation or, indeed, a model such as the one proposed by Bechara, Damasio, and colleagues (1–3), in which behavior and conscious knowledge result from partially separate mechanisms. Our point is therefore not to claim that we have ruled out nonconscious biases as possible contributors to behavior in the IGT but only to suggest that there is no need to invoke such biases to explain participants' behavior: Verbal reports reflect consciously accessible knowledge of the advantageous strategy more reliably and at least as early as behavior itself.

It is important to note that, in some situations, human choice behavior does appear to be influenced by nonconscious processes. The work of Nisbett and Wilson (10, 11), among others, shows that people sometimes provide explanations for their choices that contain no mention of the factors that actually seem to be governing their behavior. Also, studies in the implicit-learning literature have shown that, in a variety of tasks, participants are capable of achieving levels of performance above those predicted by their verbalizations (5, 12). There is even evidence that preferences may be influenced by subliminal presentation of stimuli (13). A characterization of the features of tasks that make them more amenable to explicit or implicit processing is beyond the scope of this study. Nevertheless, we would like to point out three characteristics of the IGT that seem to promote explicit reasoning. First, the IGT is self-paced, allowing plenty of time for reasoning. Second, the outcomes are presented in explicit numerical form. Third, it may be relatively easy to explicitly keep track of the approximate characteristics of each deck. For each deck, the winning amount is fixed, and the magnitude of the loss is either fixed (for two of the decks) or assumes only a small number of relatively similar values (for the other two). The frequency of loss presentations is also relatively stable for each deck, making the structure of the decks more readily apparent.

As stated above, the possibility that both conscious and nonconscious knowledge is acquired in the IGT presently cannot be ruled out. However, further research may be able to shed additional light on this issue. For example, participants in the IGT could be required to perform a concurrent secondary task thought to interfere with the use of consciously accessible information but not nonconscious processes (e.g., a task that imposed a load on working memory). If nonconscious knowledge were available to guide behavior in the IGT, the secondary task might have little, if any, effect on performance in the IGT. This and other sophisticated methods have been developed in the implicit-learning literature to try to determine when learning is implicit (5, 12); it would be worthwhile to bring the sophistication of these methods to bear on the issue of whether learning in the IGT involves an implicit component.

It is important to note that a role for nonconscious knowledge in some tasks does not imply that such knowledge is related to somatic markers. Indeed, many models of implicit learning explain nonconscious influences as arising from nonsomatic sources, such as the adjustment of connection weights in neural networks (14). There are, however, several other lines of evidence that have been interpreted as providing support for the somatic marker hypothesis. We now turn to a consideration of that evidence.

One line of such evidence is the pattern of SCRs exhibited by participants in the IGT. As mentioned in the introduction, Bechara *et al.* (3) recorded SCRs while participants played the game and found that participants had higher SCRs when they were about to select from one of the bad decks than when they were about to select from one of the good decks (respectively, the \$100 and \$50 decks), even before they were classified according to the criteria of Bechara *et al.* (3) as having a conceptual understanding of the advantageous strategy. Bechara *et al.* (2, 3)

interpreted these anticipatory SCRs as reflecting nonconscious somatic markers that can guide advantageous behavior, even in the absence of concomitant conscious knowledge. Our finding that participants have far more knowledge of the advantageous strategy than Bechara *et al.* (3) claimed raises the possibility that the SCRs reflect instead emotional responses that are elicited by knowledge that is consciously accessible and that also can guide behavior. There is therefore no need to think that the somatic markers indexed by the SCRs play a causal role in guiding participants' behavior.

In fact, there is additional evidence against the view that anticipatory SCRs reflect somatic markers that can guide advantageous behavior in the IGT. In the original IGT, the bad decks (i.e., the decks with a negative expected value) also had the highest variance. The higher anticipatory SCRs for the bad decks therefore could be related either to their negative outcomes, as suggested by Bechara *et al.* (3, 4), or to the higher uncertainty associated with selecting from those decks. Tomb *et al.* (15) have shown that, in a modified version of the IGT in which the good decks (i.e., the decks with positive expected values) have the highest variance, participants have higher SCRs when they are about to select from the good decks, even though they clearly select more from those decks. This finding suggests that the higher SCRs in the original IGT are related to the higher variance of the bad decks, not to their negative outcomes. In a reply to Tomb *et al.* (15), Bechara and the Damasio (16) have a different proposal. They suggest that perhaps in the original IGT the SCRs reflected a negative somatic state associated with bad outcomes and in the modified IGT they reflected a positive somatic state associated with good outcomes. Although such a possibility cannot be ruled out, for somatic markers to guide behavior there would then have to be additional somatic information distinguishing between positive and negative outcomes. However, Bechara, Damasio, and colleagues have not thus far produced evidence of other indices of bodily state that could, alone or in conjunction with the SCRs, indicate the goodness or badness of each deck. Furthermore, as pointed out above, our findings raise the possibility that the somatic states that participants experience in the task reflect consciously accessible knowledge.

In sum, one of the bases for proposing the somatic marker hypothesis was that, in the IGT, the SCRs seemed to indicate the goodness or badness of the decks even before conscious knowledge was available. However, considering our results, as well as those of Tomb *et al.* (15), we see no evidence to support this conclusion.

Another line of evidence that Bechara, Damasio, and colleagues (1–4) have used to support their somatic marker hypothesis comes from patients with damage to the ventromedial prefrontal cortex (VMPFC). These patients have problems with real-life decision-making, and they perform poorly on the IGT (1, 3, 4), continuing to select more from the bad decks throughout the game, even though some of them can correctly report which decks are good and bad. Furthermore, unlike controls, these patients fail to develop higher anticipatory SCRs for the bad decks. Bechara, Damasio, and colleagues (1–4) have taken these results to imply that lesions to the VMPFC impair the processing of somatic markers, which results in deficient decision-making. More specifically, they claim that VMPFC patients “are oblivious to the future consequences of their actions, and seem to be guided by immediate prospects only” (4). They term this behavior “cognitive impulsiveness,” stating that it is “related to an inability to delay gratification” (2). Such impulsiveness purportedly explains why these patients prefer the “immediate” \$100 reward from the bad decks, not taking into account the larger “future” loss associated with selecting from those decks. However, this proposal, as stated, does not appear to predict a deficit in the IGT. The structure of the IGT is such that what

determines the value of a particular selection is always the immediate outcome of that selection. There is no sense in which a selection done at a certain point in time is associated with a loss that only becomes apparent later in the game. If a participant selects from a bad deck, it is the probability of getting a bad outcome on that trial that is higher. Therefore, even participants who were guided only by immediate prospects should be able to play the game advantageously.

Fortunately, there is another interpretation of the failure of VMPFC patients to perform well in the IGT. This interpretation hinges on the fact that the \$100 decks, which turn out to be bad in the longer term, initially appear very good. In fact, in the case of one of these decks, the first negative outcome, a whopping loss of \$1,250, only comes after nine consecutive cards with a \$100 win and no loss. These observations raise the possibility that the problem for these patients is a difficulty in overcoming a response tendency established as a result of initial positive experiences with one or both of the \$100 decks. This view is supported by a number of studies. In an early study, Rolls *et al.* (17) showed that patients with ventral frontal damage had difficulty in a simple reversal task. In this task, one of two simple patterns was presented at a time on a touch screen. For one of the patterns, patients gained one point if they touched it and lost one point if they did not touch it. For the other pattern, patients lost one point if they touched it and gained one point if they did not touch it. After patients had learned these contingencies, the contingencies were reversed. Patients with ventral frontal damage could report that the contingencies had changed but did not adapt their behavior accordingly. The failure of these patients to adapt their behavior to reversals in contingencies is consistent with an extensive body of animal research (18–20). Bechara and colleagues (21) have argued that the implications of the study by Rolls *et al.* (17) for the performance of VMPFC patients in the IGT were clouded by the fact that the patients in the study by Rolls *et al.* had lesions that extended more laterally in orbito-frontal cortex. However, Fellows and Farah (22) have now demonstrated that patients with lesions restricted to VMPFC also show normal acquisition but impaired reversal in a simple reversal learning task. In a subsequent study, Fellows and Farah (23) used a shuffled version of the IGT, which was equal in every respect to the original IGT, except that it changed the order of card presentations within decks to avoid the initial apparent advantage for the \$100 decks created by the sequence in the original IGT. The performance of VMPFC patients on the shuffled IGT was overall indistinguishable from that of controls. The results of the two studies by Fellows and Farah (22, 23) and the earlier study by Rolls *et al.* (17) provide important support for the view that these patients' deficiency consists of a difficulty

in adapting their behavior (although not their knowledge) to reversals in contingencies. Furthermore, both Fellows and Farah (22) and Rolls *et al.* (17) have shown that the deficit in adapting to reversals in contingencies correlates with patients' level of impairment in real-life daily functioning.

The question of precisely why VMPFC patients can report the changes in contingencies but persevere in their behavior is a fascinating one that remains to be fully addressed. For our purposes here, we only want to point out that the dissociation observed by Bechara *et al.* (3) between these patients' conscious knowledge and both their behavior and their anticipatory SCRs could occur under many different models of the basis for behavior in the IGT, including models in which, in the normal case, conscious knowledge guides both behavior and autonomic responses in the task. For example, such a dissociation could occur if the VMPFC lesions caused a disconnect anywhere in the pathways from conscious knowledge to behavior and to the mechanisms that generate autonomic responses.

A final complication with the view that somatic markers implemented in VMPFC are sufficient to guide advantageous decision-making in the IGT is that recent results have shown that patients with lesions restricted to the dorsolateral prefrontal cortex also have deficits in the IGT (23, 24). This finding may reflect an involvement of executive functions, possibly including working memory, in the IGT. The possible involvement of working memory in the IGT would be inconsistent with the view previously proposed by Bechara and colleagues (25) that it is possible to doubly dissociate decision-making on the IGT (based on somatic markers), which would depend on VMPFC, from working memory, which would depend on the dorsolateral prefrontal cortex.

To conclude, we have shown that normal participants in the IGT have conscious knowledge of the advantageous strategy when they behave advantageously. This finding undercuts one of the main pillars of support for the somatic marker hypothesis. Other recent results in the literature raise serious questions about the additional evidence that has been used to support this hypothesis. Even though our findings, together with these other findings in the literature, do not prove that the somatic marker hypothesis is wrong, they do undercut virtually all sources of support for it. For the somatic marker hypothesis and, more generally, the theory of decision-making originally proposed by Damasio (1), to remain viable, new evidence to support it must be produced.

This work was supported by a fellowship from the Foundation for Science and Technology of Portugal (to T.V.M.) and by National Institute of Mental Health Grant P50-MH64445 (to J.L.M.).

1. Damasio, A. R. (1994) *Descartes' Error: Emotion, Reason, and the Human Brain* (Putnam, New York).
2. Bechara, A., Damasio, H. & Damasio, A. R. (2000) *Cereb. Cortex* **10**, 295–307.
3. Bechara, A., Damasio, H., Tranel, D. & Damasio, A. R. (1997) *Science* **275**, 1293–1295.
4. Bechara, A., Damasio, A. R., Damasio, H. & Anderson, S. W. (1994) *Cognition* **50**, 7–15.
5. Cleeremans, A., Destrebecqz, A. & Boyer, M. (1998) *Trends Cogn. Sci.* **2**, 406–416.
6. Shanks, D. R. & St. John, M. F. (1994) *Behav. Brain Sci.* **17**, 367–447.
7. Luce, R. D. (1959) *Individual Choice Behavior: A Theoretical Analysis* (Wiley, New York).
8. Busemeyer, J. R. & Stout, J. C. (2002) *Psychol. Assess.* **14**, 253–262.
9. Sutton, R. S. & Barto, A. G. (1998) *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA).
10. Nisbett, R. E. & Wilson, T. D. (1977) *Psychol. Rev.* **84**, 231–259.
11. Wilson, T. D. (2002) *Strangers to Ourselves: Discovering the Adaptive Unconscious* (Harvard Univ. Press, Cambridge, MA).
12. Cleeremans, A. (2001) in *International Encyclopedia of the Social & Behavioral Sciences*, eds. Smelser, N. J. & Baltes, P. B. (Elsevier, London), Vol. 4, pp. 2584–2589.
13. Zajonc, R. (1980) *Am. Psychol.* **35**, 151–175.
14. Cleeremans, A. (1993) *Mechanisms of Implicit Learning: Connectionist Models of Sequence Processing* (MIT Press, Cambridge, MA).
15. Tomb, I., Hauser, M., Deldin, P. & Caramazza, A. (2002) *Nat. Neurosci.* **5**, 1103–1104.
16. Damasio, H., Bechara, A. & Damasio, A. R. (2002) *Nat. Neurosci.* **5**, 1104.
17. Rolls, E. T., Hornak, J., Wade, D. & McGrath, J. (1994) *J. Neurol. Neurosurg. Psychiatry* **57**, 1518–1524.
18. Rolls, E. T. (1999) *The Brain and Emotion* (Oxford Univ. Press, Oxford).
19. Rolls, E. T., Critchley, H. D., Mason, R. & Wakeman, E. A. (1996) *J. Neurophysiol.* **75**, 1970–1981.
20. Schoenbaum, G., Nugent, S. L., Saddoris, M. P. & Setlow, B. (2002) *Neuro-Report* **13**, 885–890.
21. Bechara, A., Tranel, D. & Damasio, H. (2000) *Brain* **123**, 2189–2202.
22. Fellows, L. K. & Farah, M. J. (2003) *Brain* **126**, 1830–1837.
23. Fellows, L. K. & Farah, M. J. (June 24, 2004) *Cereb. Cortex*, 10.1093/cercor/bbh1108.
24. Manes, F., Sahakian, B., Clark, L., Rogers, R., Antoun, N., Aitken, M. & Robbins, T. (2002) *Brain* **125**, 624–639.
25. Bechara, A., Damasio, H., Tranel, D. & Anderson, S. W. (1998) *J. Neurosci.* **18**, 428–437.