

Intuitive decision making in complex situations: Somatic markers in an artificial grammar learning task

DICK J. BIERMAN

University of Amsterdam, Amsterdam, The Netherlands

and

ARNAUD DESTREBECQZ and AXEL CLEEREMANS

Université Libre de Bruxelles, Brussels, Belgium

In this article, we explore the extent to which implicit learning is subtended by somatic markers, as evidenced by skin conductance measures. On each trial, subjects were asked to decide which “word” from a pair of “words” was the “correct” one. Unknown to the subjects, each “word” of a pair was constructed using a different set of rules (Grammar A and Grammar B). A (monetary) reward was given if the subject chose the “word” from Grammar A. Choosing the Grammar B word resulted in (monetary) punishment. Skin conductance was measured during each of 100 trials. After each set of 10 trials, the subjects were asked how they selected the “correct word.” Task performance increased long before the subjects could even formulate a single relevant rule. In this *preconceptual* phase of the experiment, skin conductance was larger before incorrect than before correct choices. Thus, it was shown that artificial grammar learning is accompanied by a somatic marker, possibly “warning” the subject of the incorrect decision.

Everyday experience suggests that we often seem to know more than we can tell. Riding a bicycle or playing tennis, for instance, involves mastering complex sets of motor skills, yet we are at a loss when it comes to explaining exactly how we perform such physical feats. Such dissociations between our ability to report on cognitive processes and the corresponding behaviors are not limited to action but extend to higher level cognition as well. Most native speakers of a language are unable to articulate the grammatical rules that they nevertheless follow when uttering expressions of the language. Likewise, expertise in such domains as medical diagnosis or chess, as well as social or aesthetic judgments, all involve intuitive knowledge that one seems to have little introspective access to.

We also often seem to tell more than we can know. In a classic article, Nisbett and Wilson (1977) reported on many experimental demonstrations that verbal reports on our own behavior often reflect reconstructive and interpretative processes, rather than genuine introspection. Dissociations between behavior and verbal report also

form the basis of a large literature dedicated to implicit learning—broadly construed, learning without awareness (see Cleeremans, Destrebecqz, & Boyer, 1998, for a review). According to **Berry and Dienes (1993)**, learning is implicit when we acquire new information without intending to do so and in such a way that the resulting knowledge is difficult to express. Implicit learning thus contrasts with explicit learning (e.g., as when learning how to solve a problem or learning a concept), which is typically hypothesis driven and fully conscious. Implicit learning research has essentially focused on three experimental paradigms: artificial grammar learning, dynamic system control, and sequence learning. In Reber’s **(1967)** seminal study of artificial grammar learning, subjects were asked to memorize a set of meaningless letter strings generated by a simple set of rules embodied in a finite-state grammar. After this memorization phase, they were told that the strings followed the rules of a grammar and were asked to classify novel strings as grammatical or not. In this experiment and in many subsequent replications, subjects were able to perform this classification task better than chance would predict, despite remaining unable to describe the rules of the grammar in verbal reports. This dissociation is what prompted Reber to describe learning as implicit, for the subjects appeared sensitive to and could apply knowledge (the rules of the grammar) that they remained unable to describe and had had no intention to learn.

Although such findings suggest that unconscious influences on behavior are pervasive, it is important to note

This work was supported by European Commission Grant HPRN-CT-1999-00065 and by an institutional grant from the Université Libre de Bruxelles. A.D. and A.C. are, respectively, a scientific research worker and a senior research associate with the National Fund for Scientific Research (Belgium). Simon van Gaal ran the experiment as part of his master’s thesis. Correspondence concerning this article should be addressed to D. J. Bierman, University of Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands (e-mail: d.j.bierman@uva.nl).

that the relationships between learning and awareness continue to elicit controversy. Because there is no accepted operational definition of what it means for an agent to be conscious of something, difficult definitional, conceptual, and methodological challenges need to be overcome. One of the most difficult challenges in this respect is to determine which criterion one should use to determine whether processing was unconscious or not. Although it would be outside of the scope of this article to offer a detailed overview of the different methods one can deploy to assess the extent to which performance reflects implicit influences, any such method must necessarily rely on comparing two measures: a measure of awareness and a measure of performance. In this context, it might, therefore, be particularly useful to consider additional measures of performance, such as physiological responses, as an indication that implicit processes are involved in shaping performance. For instance, in a gambling situation, subjects might use a conscious strategy based on the idea that outcomes that have occurred recently become less probable in the future (the gambler's fallacy; Clotfelter & Cook, 1993). However, their decisions might also be driven by an implicit sensitivity to the actual probabilities of the different possible outcomes, and the results of this implicit sensitivity might preferentially express themselves through physiological measures. Although these issues also remain controversial, in general, it is fair to say that one might expect most decisions to be influenced by both implicit and explicit knowledge.

Damasio and colleagues explored performance in a gambling task while simultaneously measuring skin conductance (Bechara, Damasio, Trenz, & Damasio, 1997; Bechara, Trenz, Damasio, & Damasio, 1996). In this situation, subjects, after being given \$2,000 in play money, were asked, on each of a series of trials, to choose a card from one of four decks. Each choice resulted in a win or a loss. The subjects were told to play so as to maximize gains. Unbeknown to the subjects, the decks differed in their overall ultimate yield, with some decks being disadvantageous and others being advantageous. The subjects were free to choose cards from any of the four decks and did not know how many trials had to be performed before the experiment would end. The subjects were probed about their knowledge of the situation at regular intervals during the game.

The results of these experiments indicated (1) that the subjects started selecting cards from the advantageous decks before they were able to verbally motivate and explain their choices and (2) that they exhibited a larger skin conductance response (SCR) just before taking a card from a disadvantageous deck. Thus, differential SCRs to advantageous and disadvantageous decks appeared before the subjects were able to **motivate** their decisions, as if their body knew which decks were risky before the relevant knowledge was available for verbal reports. In contrast, patients with damage in the prefrontal cortex failed to exhibit anticipatory SCRs and tended to continue to

select cards from the bad decks even though some of them ended up being able to verbally describe the correct selection strategies. Damasio and colleagues interpreted these findings by proposing to formalize decision making as involving two parallel but interacting processes. The first involves mapping the currently experienced situation to knowledge about one's own emotional response in previously experienced similar situations. This process is assumed to be severely disturbed for the prefrontal patients. The relevant knowledge is assumed to be nondeclarative or implicit and to represent the agent's dispositions or biases. The second set of processes involves explicit recall of relevant facts pertaining to the consequences of previous choices and the activation of relevant reasoning strategies. In this case, the relevant knowledge and processes are assumed to be largely available to conscious awareness.

The role of intuition in decision making can thus be conceptualized as a two-step process in which (implicit) knowledge is first marked with a positive or negative valence, depending on the outcome of previous decisions, and then is used to shape further (explicit) decision making by means of the *somatic marker* (SM; the emotional valence) associated with the knowledge. According to Damasio (1996), "the hypothesis thus suggests that somatic markers normally help constrain the decision-making space by making that space manageable for logic-based, cost-benefit analyses. In situations in which there is remarkable uncertainty about the future and in which a decision should be influenced by previous individual experience, such constraints permit the organism to decide efficiently within short time intervals" (p. 1415).

However, although the gambling task is certainly an interesting choice situation, it might not be very representative of the sorts of situations faced by decision makers. Indeed, real-life choice situations often involve many interacting factors, as well as structured stimulus material. In this article, our main goal is, therefore, to explore whether we can replicate Damasio's (1996) results, using a more complex task based on Reber's (1967) artificial grammar learning task. In our adaptation, subjects were asked to select, on each trial, one of two letter strings presented concurrently on the screen. In each pair, one letter string had been generated using one finite-state grammar, whereas the other had been generated on the basis of another finite-state grammar. Correct decisions can, therefore, be reached only to the extent that people learn something about the structure of the stimulus material on the basis of the pattern of successive reinforcements to previous choices. SCR was monitored during all the trials.

This approach makes it possible to address several concerns with regard to the original gambling task. First, increasing the complexity of the stimulus material makes it possible to use a greater number of indirect questions to probe subjects' explicit knowledge about the task than could be used with the original task (e.g., "Which deck is the most advantageous?"). Pilot experi-



ments conducted in our lab led us to believe that subjects suspected that payout in the gambling task was driven by more complex rules involving, for instance, responses to earlier trials. Such subjects might, therefore, have possessed relevant explicit knowledge but failed to report on it, because they were still engaged in attempting to figure out the causes of the differences between the different decks. Our adaptation of the original gambling task addresses this issue to some extent by making it possible to ask subjects more open-ended questions, such as “On what basis do you make your decisions?”

Second, the original gambling task involved, at least for the noncomputerized studies, a fixed sequence of winning and losing cards within the decks that was far from being truly random. Some of the reported outcomes, especially the differences between healthy and patient groups, might thus be attributed to sequential *guessing* patterns specific for the healthy and the patient groups. The claim that the subject has “no way of predicting when a penalty will arise” (Bechara et al., 1996, p. 1293) does not appear to be justified under the assumption that subjects become sensitive to the statistical structure of the series of cards within the decks. Thus, in addition to introducing a more complex task to explore the SM hypothesis, true randomization with replacement was used to select the position of the correct “word” for each subject.

Because our main goal was to explore the extent to which Damasio’s (1996) findings generalize to a novel, more complex situation, we hypothesized that subjects would (1) perform above chance in their selection of correct strings before becoming able to verbally **motivate** their choices (*implicit learning hypothesis*), and (2) exhibit differential SCRs before becoming able to verbally motivate their choices (*SM hypothesis*). Validating the second hypothesis would lend support to the idea that above-chance selection performance depends on the availability of relevant SMs. This would not, however, exclude the possibility that the SM also plays a role even when knowledge about grammaticality has become explicit. Therefore, we also explored the role of the SM after the subjects had expressed explicit knowledge of one or more relevant rules.

METHOD

Subjects

Thirty volunteers (10 males and 20 females) 18 to 51 years of age ($M = 22.2$ years, $SD = 7.1$) participated in this study. The subjects were either acquaintances of the experimenters or freshman psychology students at the University of Amsterdam, who participated for course credit. All the subjects were paid €7 and a variable bonus (range, €1–€3), depending on their performance.

Materials and Apparatus

Stimuli were presented in Helvetica 18 on the screen of an iMac computer. The stimuli consisted of pairs of letter strings six elements long. Each element consisted of one of four possible symbols (#, *, and +). Each string of a given pair was generated on the basis of one of the two finite-state grammars depicted in Figure 1, so that all the pairs contained one string from Grammar A and one string from Grammar B. Each grammar involved the same set of four symbols and differed only by the transition probabilities associated to certain arcs. String generation proceeded as follows. A starting node was first selected at random. Next, subsequent elements were generated by randomly selecting among the arcs emanating from the current node according to the transition probabilities specified in the grammar. The symbol associated with the node pointed to by the selected arc was then recorded. Generation continued in this manner until six elements had been generated. The probability associated with the self-transition loops in either grammar was reduced to zero after one self-transition had occurred, so that no strings would contain runs of more than two identical symbols. Note that the grammars are orthogonal, except for the transitions between the * and the # symbols. As a result, string discrimination can be achieved exclusively on the basis of knowledge of certain bigrams that occur only in one of the two grammars. The task is, therefore, overall easier than in typical artificial grammar learning situations, in which overall surface similarity between the various categories of items (e.g., grammatical vs. ungrammatical strings) is carefully controlled so as to eliminate its influence on the subject’s decisions (see, e.g., Knowlton, Squire, & Gluck, 1994). However, our main purpose in using such simplified material in the context of this study was (1) to ensure that the subjects could verbalize the rules that they used in making their decisions and (2) to facilitate the scoring of the verbalizations. Hence, we make no claims about the extent to which subjects actually learn about rules or merely about the surface structure of the material. The location (left or right) at which strings generated from either Grammar A or Grammar B appeared on the screen was truly random.

Two Ag-AgCl electrodes were attached to the middle and index fingers of the nonpreferred hand. Isotonic paste was used. Skin con-




Figure 1. Transition probabilities for Grammars A and B. The transition of a symbol to itself was allowed only once.

ductance was measured with the Orion 4AD22, which determines skin conductance using a constant AC method (10 microamps, 100 Hz). The data were sampled continuously on an interrupt basis with a sample frequency of 5 samples/sec. After each trial, epochs were stored to disk, using a temporal window that began 4 sec before the choice and that extended 13 sec after the choice had been made (see Figure 2).

Procedure

The subjects received written instructions describing the goal of the experiment as a learning task. The instructions emphasized the possibility for the subjects to earn money. After attaching electrodes to the nonpreferred hand, this hand was positioned on a small pillow, and skin conductance was measured on a deep breath. The subjects were then given a practice trial so as to familiarize them with the experiment. The experiment itself was initiated after the experimenter had answered possible questions and the subjects had received €500 worth of play money as an initial amount. The experimenter remained in the room for the duration of the experiment but could not see the display. On each trial, the experimenter adjusted the pile of play money in front of the subject according to the auditory feedback given to him or her.

The entire experiment consisted of 100 trials presented in blocks of 10 trials each. Figure 2 illustrates the sequence of events taking place within a single trial. The subjects initiated each trial by pressing any key on the computer keyboard. Two strings were then presented together on the screen. The subjects had to indicate which string they thought belonged to a language spoken on "Planet A" by pressing one of two predefined keys. No time pressure was imposed on string selection, but reaction time was recorded. After the subjects had responded, a delay of 3 sec occurred, during which the strings continued to be displayed. Feedback was then provided (1) by highlighting the correct string in green, (2) by a digitized voice saying "prima" (correct) or "jammer" (incorrect), and (3) by displaying the cumulative amount of euros that had been won or lost so far. Correct choices yielded a reward of either €10 or €100. Incorrect choices incurred a penalty of either €10 or €100. For both correct and incorrect choices, the actual amounts were chosen at random (following the varying rewards in the original Damasio, 1996, gambling task). To enhance the *emotional* effect of success or failure, the experimenter physically removed play money from the pile on each trial. Feedback remained on the screen for 10 sec, after which a message indicating that the next trial could be initiated appeared on the screen.

To assess the subjects' explicit knowledge of the material, they were asked, after each set of 10 trials, to answer the following question displayed on the computer screen: "How do you come to a choice between the two words?" Responses were entered by the experimenter on a standardized scoring form. Knowledge of the grammar was scored as having become explicit when the subject (1) correctly formulated at least one correct decision rule and (2) mentioned the same rule again in answer to the next probe (i.e., after another set of 10 trials). For instance, if at Trial 30 the subject mentioned that for the correct choices the symbol # would always be followed by the symbol * (which was indeed the case) and the subject mentioned this rule again on Trial 40, this subject was scored as having acquired an explicit rule at Trial 25 (rather than at Trial 30, because this apparently stable rule could have been discovered anywhere between Trials 21 and 30).

RESULTS

Unresponsive subjects, defined to be subjects exhibiting skin conductance variability smaller than 10 μ mho over the course of the entire experiment, were eliminated. Five such subjects were thus removed from the original 39 participants. A further 4 subjects were removed due to equipment failure. These decisions were made before analysis of the remaining 30 subjects was initiated.

Data-Reduction

Baseline-corrected skin conductance values were averaged over the 7-sec period extending from the first (baseline) sample up until the point at which feedback was given. The resulting measurement therefore represents average SCR during the decision and anticipation phases. It corresponds to Damasio's (1996) SM. These "SM" values were averaged separately, on a subject-by-subject basis, for the correct and the incorrect choices. Only those trials for which it was determined that the subject had no explicit knowledge of the grammar rules were used. This analysis thus resulted in two dependent variables: *SM_correct* and *SM_incorrect*.

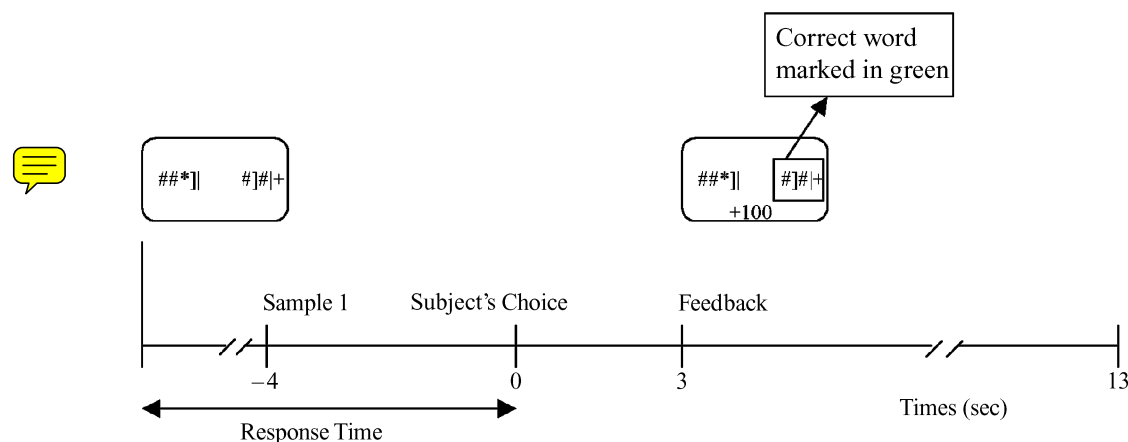


Figure 2. Timing of a single trial. Data are stored from 4 sec before until 13 sec after the choice between the two words.




Figure 3. The performance of the 25 subjects who failed to express any explicit knowledge before Trial 50.

Implicit Learning Hypothesis

For each subject, the start of the conceptual phase (explicit knowledge phase) was determined using the method described earlier. This was compared with their performance curve. For most of the subjects, performance started to increase long before they entered the conceptual phase. Only 5 subjects reported a correct explicit rule before Trial 50. Twelve subjects failed to formulate any rule before the end of the experiment. The average performance curve for the 25 subjects who did not formulate any rule before Trial 50 is shown in Figure 3. For each data point, the average percentage correct over that and the 9 subsequent trials is plotted. It can be seen that performance for these nonreporting subjects already increases very early in the experiment. Twenty-one of the 25 subjects had an average score over 50% between Trials 10 and 40. Their mean scoring rate was 72.9% ($t(24) = 7.39, p < .0001$). On the basis of these results, we can thus conclude that implicit learning, comparable to the learning found in Damasio's (1996) original gambling task, occurred in this experiment.

“Somatic Marker” Hypothesis

Figure 4 shows the time course of the average skin conductance over all the subjects, using only the trials on which no explicit knowledge was formulated.

The figure suggests that the average skin conductance value was larger before feedback for incorrect, rather than for correct, choices. Because the distribution of skin conductance measurements is known to be nonnormal, we used a binomial test to compare the number of subjects who (1) exhibited a larger SM before the incorrect trials than before the correct trials and (2) exhibited the reverse pattern. Two subjects formulated a correct rule at Trial 10 and were thus eliminated from this analysis be-

cause the preconceptual phase was too short. Of the remaining 28 subjects, 19 expressed a larger SM before incorrect choices than before correct choices, and 9 exhibited the reverse pattern. This difference was significant (binomial $p = .044$). Note that the corresponding within-subjects analyses produced nonsignificant results. In other words, the fact that 9 subjects exhibited a larger SM for correct than for incorrect choices should not be taken as suggesting that these subjects exhibited a reverse SM effect. Instead, they simply belong to the lower part of the distribution of responses.

Exploratory Analyses

SM over the whole experiment. In a subsequent analysis, we explored how skin conductance varied over the course of the entire experiment. We therefore replicated the analyses described above, this time also including the trials described by Damasio (1996) as characteristic of the *conceptual phase*—that is, all the trials for which the subjects had expressed at least one correct rule. Interestingly, this analysis indicated that the effect increased slightly from 67.9% to 73.3% (22 out of 30 subjects; binomial $p = .009$).

Correct versus incorrect decisions and response times. In most decision tasks, there is a tradeoff between response time and performance. However, in a number of tasks that involve nonconscious processes in the realm of perception, it has been found that using a pop-up strategy will improve performance (Snodgrass, Shevrin, & Kopka, 1993). This pop-up strategy basically consists of choosing the first alternative that comes to mind, thus preventing any further analysis. We compared the mean response times for correct and incorrect decisions in the preconceptual and conceptual phases of the experiment (Table 1).

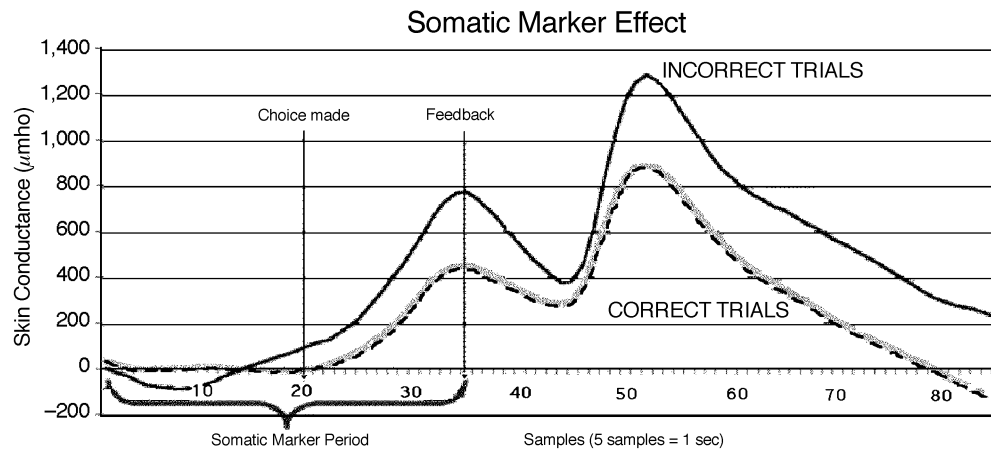


Figure 4. The skin conductance preceding, during, and after feedback for incorrect and correct decisions for all the subjects, averaged over their preconceptual trials.

It can be seen that response times for correct decisions were significantly shorter than those for incorrect ones. However, it is difficult to interpret this result, since the correlation between response times and performance can be attributed to a causal factor originating either in the speed of the response (resulting in a pop-up strategy with better performance) or in the difficulty of the specific item (resulting in a longer response time). In the conceptual phase, correct decisions required only 2.46 sec, but incorrect decisions took nearly twice that long, suggesting that these decisions concerned trials in which the subject's explicit knowledge was insufficient to solve the problem.

DISCUSSION

The major finding in this experiment is that somatic marking, as originally found by Damasio et al. in a gambling task, is also present in an artificial grammar learning task. This happened before the subjects could formulate any explicit knowledge of the grammar, during what Damasio called the *preconceptual* phase. We found that skin conductance was higher before incorrect choices than before correct choices. This was the case in both the preconceptual and the conceptual phases. These findings lend support to the suggestion that the SM process is important for everyday complex intuitive decisions in problem-solving situations that are underspecified or for which not enough time is available for a complete analytical solution.

Is Our Measure of Conscious Knowledge Sensitive Enough?

Our finding that subjects' classification performance improves well before they are able to verbalize their decision criteria (the implicit learning hypothesis) is, at first sight, rather convincing. After all, people were placed in a situation in which they were actually searching for rules; they were asked simple and direct questions every 10 trials, probing directly for any knowledge they might

use in making decisions. Even under these conditions of intentional learning accompanied by tangible rewards, in which the subjects were repeatedly prompted to verbalize any knowledge that they might consciously hold about their decision criteria, a substantial majority of them failed to verbalize anything before Trial 50—that is, after having been prompted to do so on five separate occasions. We believe that even critics of implicit learning will have to admit that this might, at least, represent a good indication that people make successful decisions on the basis of something other than reportable knowledge. Nevertheless, one might argue that our method of assessing the subjects' conscious knowledge (which closely followed that in Bechara et al.) was not sufficiently sensitive to the subjects' conscious knowledge. For instance, the subjects might fail to report their knowledge not because it is implicit, but rather because they might be reluctant to volunteer low-confidence knowledge or because they are anxious to avoid reporting erroneous knowledge. Thus, our first hypothesis, that implicit learning occurs, could be accepted falsely, because learning was established in a phase in which explicit knowledge was already available, yet not verbalized. Although our main goal was to shed more light on the role of the SM, and not so much on the hypothesis that implicit learning occurs in this situation, the issue of knowing the extent to which subjects possess explicit knowl-

Table 1
Mean Response Times (RTs, in Milliseconds) for Incorrect and Correct Decisions in Preconceptual and Conceptual Phases of the Experiment

Phase	Decision	RT	Wilcoxon z	p
Preconceptual	Incorrect	4,130	3.98	< .0001
	Correct	3,480		
	Difference	650		
Conceptual	Incorrect	4,257	3.42	< .0001
	Correct	2,459		
	Difference	798		

edge before, concurrently, or after SMs have become detectable is certainly relevant in the context of discussing the functional role of the latter.

As a case in point, Maia and McClelland (2004) recently raised exactly the same issues in the context of a study that replicated the original Bechara et al. findings. In a subsequent experiment in which a more elaborate probing scheme was used, however, Maia and McClelland found that subjects, in fact, turned out to possess much more explicit knowledge than had been revealed through the simpler knowledge elicitation method used by Bechara et al. and thus concluded that there was, in fact, no evidence for implicit learning in this situation—a conclusion that has itself been disputed by Damasio et al. (2005; Maia & McClelland, 2005).

We will discuss our own findings again in light of this debate at the end of this section but would first like to point out that devising an appropriate measure of awareness is a particularly challenging problem that has long been and continues to be controversial in fields as diverse as subliminal perception, memory, learning, and conditioning. Most experimental paradigms dedicated to exploring the relationships between conscious and unconscious processing have relied on a simple quantitative dissociation logic aimed at comparing the sensitivity of two different measures to some relevant information: a measure (*C*) of subjects' awareness of the information and a measure (*P*) of behavioral sensitivity to the same information in the context of performing some task. Unconscious processing, according to the quantitative dissociation logic, is then demonstrated whenever *P* exhibits sensitivity to some information in the absence of correlated sensitivity in *C*.

As has been noted by many authors, however, there are several important pitfalls to the simple dissociation logic. First, the measures *C* and *P* cannot typically be obtained concurrently. This *retrospective assessment* problem (Shanks & St. John, 1994) entails the conclusion that finding that *C* fails to be sensitive to the relevant information need not necessarily imply that the information was processed unconsciously during encoding, but that, for instance, it might have been forgotten or otherwise distorted before retrieval. This is unlikely to be the case in paradigms such as ours, however.

A second issue is ensuring that the information revealed through *C* is indeed relevant to performing the task. As Shanks and St. John (1994) have suggested, many studies of implicit learning have failed to respect this *information* criterion, also called the *relevance* principle by Lovibond and Shanks (2002). For instance, successful classification in an artificial grammar learning task need not necessarily be based on knowledge of the rules of the grammar but can, instead, involve knowledge of the similarity relationships between training and test items. Subjects asked about the rules of the grammar would then understandably fail to offer relevant explicit knowledge. The work of Dulany, Carlson, and Dewey (1984), for instance, clearly showed that probing sub-

jects, not about their knowledge of the rules, but simply about their knowledge of which letters made a string grammatical or not was a much more sensitive way of revealing the subjects' conscious knowledge. This concern applies to our experimental situation, for it is indeed the case that knowledge of bigram statistics is sufficient to ensure correct classification. Note, however, that we did not ask the subjects to verbalize rules but simply asked them to indicate how they came to make their decisions. In other words, they were free to report bigram knowledge, which was indeed the case.

A third issue is ensuring that *C* and *P* are equally sensitive to the relevant information. This is what Shanks and St. John (1994) called the *sensitivity criterion*. At first sight, verbal reports and other *subjective measures*, such as confidence ratings, would appear to offer the most direct way through which to assess the contents of subjective experience. The use of subjective measures to assess awareness was first advocated by Cheesman and Merikle (1984), who also introduced the notions of subjective and objective thresholds. Performance on a given task (i.e., identification) is said to be below the subjective threshold if one can show that performance is better than chance while subjects indicate that they are guessing (through confidence judgments, for instance). Performance is said to be below the objective threshold if it fails to differ from chance. According to this logic, unconscious perception, for instance, would thus be demonstrated whenever performance is below the subjective threshold and above the objective threshold. Dienes and Berry (1997) suggested that this logic could also be applied to the domain of implicit learning, and Dienes, Altmann, Kwan, and Goode (1995) operationalized it by proposing two criteria with which to assess the extent to which learning is implicit. The first criterion is the *guessing criterion*, which basically states that one can conclude that learning is implicit to the extent that people perform better than chance while believing that they are guessing. The second, first explored by Chan (1992), is the *zero-correlation criterion*, which states that one can conclude that learning was implicit if confidence judgments offered by subjects about their own performance fails to correlate with it. Several studies have now applied these ideas in the domains of artificial grammar learning (Dienes & Altmann, 1997) and sequence learning (Shanks & Johnstone, 1998). Overall, these studies indicate that the knowledge acquired by subjects in these empirical situations can, indeed, be implicit to the extent that it is *below the subjective threshold*.

However, as Reingold and Merikle (1990) have pointed out themselves, there are clear methodological shortcomings involved in the use of such subjective measures of conscious awareness. For instance, people might simply refrain from reporting on knowledge held with low confidence or might offer reports that are essentially reconstructive in nature, as Nisbett and Wilson's (1977) experiments indicate. For this reason, many authors have advocated using so-called *objective measures* of aware-

ness. Objective measures of awareness include forced choice tests, such as recognition, presence–absence decisions, or identification. Today, numerous studies have been conducted using objective measures.

Even if the different criteria briefly reviewed above are fulfilled, however, it might be elusive to hope to be able to obtain measures of awareness that are simultaneously *exclusive* and *exhaustive* (see Jiménez, 1997; Jiménez, Mendéz, & Cleeremans, 1996; Reingold & Merikle, 1988) with respect to knowledge held consciously. In other words, finding null sensitivity in *C*, as required by the dissociation paradigms for unconscious processing to be demonstrated, might simply be impossible because no such absolute measure exists. A significant implication of this conclusion is that, at least with normal subjects, it makes little sense to assume that conditions exist in which awareness can simply be “turned off.” It might, therefore, be more plausible to assume that any task is always sensitive to both conscious and unconscious influences. In other words, no task is *process pure*. We believe that this is precisely the case in our experimental situation, as we will suggest below.

Returning to our own findings now, it is important to realize that our paradigm is, in some respects at least, rather different from that used by Bechara et al. In particular, it should be much easier for subjects to verbalize their decision rules in the context of a grammar learning task than in the original Iowa gambling task. Hence, it is not clear to what extent further probing, as advocated by Maia and McClelland, would help. This impression is supported by the results of an unpublished experiment in which, at Trial 20, half of the subjects were probed using a much more elaborate method. This involved stressing the fact (1) that mentioning incorrect rules did not matter and (2) that a considerable extra reward of €50 (on top of the amount won by categorizing better than chance) could be obtained if *any* correct rule was mentioned. Even under these conditions, we found no difference in the number of correct rules formulated at Trial 20 between the two groups of subjects (Bierman, 2005). In fact, the subjects who had been probed extensively at Trial 20 eventually mentioned, on average, their first correct rule no earlier than did the subjects who had been probed using the standard method. The difference with Maia and McClelland’s findings can probably be explained by differences between the two tasks. In the gambling task, the probing is confusing because the word *advantageous* is used when questioning which deck of cards is the most advantageous. By making this concept more explicit, as was done by Maia and McClelland, the subject reveals more knowledge concerning the amount of the rewards and the relative frequency of the negative reinforcements.

In our task, however, people have clear and simple criteria with which to classify the strings. Regardless of the nature of their knowledge—the frequency of particular bigrams or more complex abstract rules—the requirement to verbalize *any knowledge used to make decisions* is un-

ambiguous and was easily understood by the subjects. All knowledge verbalized by the subjects, in fact, took the form of simple production rules, such as “IF there is a pair of + and * symbols THEN the word is an A word.”

One further aspect of our probing methodology deserves discussion. On each probing trial, the subjects who failed to repeat a rule that they had previously mentioned were reminded that they had done so before. When assessing the subjects’ explicit knowledge, we considered that if a rule was mentioned at Trial *X* and not mentioned again at Trial *X*+10, the rule was not included in the subjects’ explicit knowledge. Although this procedure seems conservative, it virtually never happened that a correct rule was subsequently abandoned. It did often happen, however, that incorrect rules were abandoned.

One can wonder whether subjects, once they formulate a rule, do adhere to their own rule. In order to check this, we estimated theoretically and with a simulation the mean percentage correct to be expected if a single rule is known. There are eight rules (i.e., eight possible transitions in each grammar). The relative probability for each of the transition rules to be **fired** is confirmed to be 1/8 in a simulation. Since each word contains six elements, we have five transitions, so logically there are five rules **fired** for each word. However, sometimes a rule is **fired** twice, so that only four rules were involved in the construction of the word. The probability that a rule will be **fired** twice in one word turns out to be 14%. The expected scoring rate, after one correct rule has been identified by the subject, therefore is $5/8 * 86 + 4/8 * 14 = \sim 60\%$. The remaining 40% of the words will be guessed at. Thus, the simulation yields **the conclusion** that the mean scoring rate (including guessing) will be 80% after a single rule has been identified.

Interestingly, the 25 subjects who failed to express any rule before Trial 50 eventually correctly classified slightly fewer than 80% of the strings—the rate that would be expected on the basis of knowledge of a single rule. We further explored how well the subjects who had expressed knowledge of a single rule performed on the classification task and found that the mean scoring rate of the 17 subjects who did express such knowledge was 90.58% over the 10 trials that immediately followed verbalization of a correct rule. When combined, the two observations that (1) the subjects who failed to verbalize any rule achieved an 80% correct classification rate and (2) the subjects who did express a single rule scored better than would be expected on the basis of knowledge of a single rule (80% after correction for guessing vs. 90.58%) suggest that classification performance is at least partly driven by implicit knowledge, even in the conceptual phase. Hence, although we cannot rule out **the possibility** that more sensitive measures would not have resulted in more knowledge being expressed on a direct test of the subjects’ knowledge, overall, the evidence, in our view, is supportive of the notion that implicit learning indeed took place in this situation. At the very least, the subjects’ performance appears to have been driven by

knowledge that they failed to verbalize, even in a fully intentional situation in which they were repeatedly probed about any knowledge that they might have been aware of.

It should be noted that the claim that implicit learning took place is dependent not only on the correct measurement of the transition to explicit knowledge, but also on a liberal interpretation of the contexts resulting in implicit learning. Although the subjects were not explicitly instructed to try to learn, the context of the experiment, with specific direct (nondelayed) reinforcement, could easily be interpreted by the subjects as a learning environment.

In artificial grammar experiments, there are generally no such reinforcements. In those experiments, the examples might be “stored” as instances, whereas the subjects do not really try to figure out underlying regularities. In the present implementation of the grammar task, by contrast, we expect that subjects do generate hypotheses about the nature of the underlying regularities. These hypotheses are then “tested” against the subsequent pairs of words presented during the next trials. In future work, the probing question therefore should ask also for current and, especially, previously held hypotheses. There is suggestive evidence in the present experiment for the notion that such decisions are better if they are made according to the first thing that comes to mind.

The SM, as measured by skin conductance, is assumed to reflect the emotion associated with a given decision. Although skin conductance can be seen as a correlate of arousal, it cannot be used to differentiate between positive and negative emotions. In general, however, negative emotions do generate larger arousal than do positive emotions, and one could thus assume that the larger SM preceding incorrect decisions reflects the negative emotions that were experienced in earlier instances when a similar decision was taken. This interpretation, however, runs counter to the idea that the SM is used as a warning signal, for we would then expect larger SMs before correct decisions. Addressing these issues would require further experiments that would make it possible to differentiate the early speculations of subjects from their final decisions. This is the goal of our future work, using faster physiological measures, such as pupil dilation, as a potential SM.

REFERENCES

- BECHARA, A., DAMASIO, H., TRANEL, D., & DAMASIO, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*, 1293-1295.
- BECHARA, A., TRANEL, D., DAMASIO, H., & DAMASIO, A. R. (1996). Failure to respond to anticipated future outcomes following damage to prefrontal cortex. *Cerebral Cortex*, *6*, 215-225.
- BIERMAN, D. J. (2005). *Using an eye tracker to measure the three factors underlying intuitive decision making*. Manuscript in preparation.
- CHAN, C. (1992). *Implicit cognitive processes: Theoretical issues and applications in computer systems design*. Unpublished doctoral thesis, University of Oxford.
- CHEESMAN, J., & MERIKLE, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, *36*, 387-395.
- CLEEREMANS, A., DESTREBECQZ, A., & BOYER, M. (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, *2*, 406-416.
- CLOTFELTER, C. T., & COOK, P. J. (1993). The “gambler’s fallacy” in lottery play. *Management Science*, *39*, 1521-1525.
- DAMASIO, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London: Series B*, *351*, 1413-1420.
- DIENES, Z., & ALTMANN, G. T. M. (1997). Transfer of implicit knowledge across domains: How implicit and how abstract? In D. C. Berry (Ed.), *How implicit is implicit learning?* (Vol. 5, pp. 107-123). Oxford: Oxford University Press.
- 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, & Cognition*, *21*, 1322-1338.
- DIENES, Z., & BERRY, D. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review*, *4*, 3-23.
- DULANY, D., CARLSON, R., & DEWEY, J. (1984). A case of syntactic learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology: General*, *113*, 541-555.
- ERICSSON, K. A., & SIMON, H. A. (1984). *Protocol analysis: Verbal reports as data*. Cambridge, MA: MIT Press.
- JIMÉNEZ, L. (1997). Implicit learning: Conceptual and methodological issues. *Psychologica Belgica*, *37*, 9-28.
- JIMÉNEZ, L., MENDÉZ, C., & CLEEREMANS, A. (1996). Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *22*, 948-969.
- KNOWLTON, B. J., SQUIRE, L. R., & GLUCK, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*, 106-120.
- 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.
- NISBETT, R. E., & WILSON, T. D. (1977). Telling more than we can do: Verbal reports on mental processes. *Psychological Review*, *84*, 231-259.
- REBER, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning & Verbal Behavior*, *6*, 855-863.
- REINGOLD, E. M., & MERIKLE, P. M. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, *44*, 563-575.
- REINGOLD, E. M., & MERIKLE, P. M. (1990). On the inter-relatedness of theory and measurement in the study of unconscious processes. *Mind & Language*, *5*, 9-28.
- SHANKS, D. R., & JOHNSTONE, T. (1998). Implicit knowledge in sequential learning tasks. In M. A. Stadler & P. A. Frensch (Eds.), *Handbook of implicit learning* (pp. 533-572). Thousand Oaks, CA: Sage.
- SHANKS, D. R., & ST. JOHN, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral & Brain Sciences*, *17*, 367-447.
- SNODGRASS, M., SHEVRIN, H., & KOPKA, M. (1993). The mediation of intentional judgments by unconscious perceptions: The influences of task strategy, task preference, word meaning, and motivation. *Consciousness & Cognition*, *2*, 169-193.

(Manuscript received September 14, 2004;
revision accepted for publication May 13, 2005.)