Running head: Somatic Markers

Word count: ~6600

Correspondence:
Dick J. Bierman
University of Amsterdam
Roetersstraat 15
1018 WB Amsterdam
The Netherlands
Email: d.j.bierman@uva.nl
Phone: (+31) 20 5256727
Fax: (+31) 20 6391656

Keywords: Intuition, Somatic Marker, Decision, Artificial Grammar, Implicit Learning.

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, Bruxelles, Belgium

In this paper 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’ word. Unknown to 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 choose 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 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 ‘pre-conceptual’ phase of the experiment, skin conductance was larger before incorrect than before correct choices. Thus it was shown that artificial grammar learning is accompagnied by a somatic marker, possibly ‘warning’ the subject for 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, involve 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 they nevertheless follow when uttering expressions of the language. Likewise, expertise in domains such 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 et al., 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 been focused on three experimental paradigms: Artificial grammar learning, dynamic system control, and sequence learning. In Reber's 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 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.

While such findings suggest that unconscious influences on behavior are pervasive, it is important to note 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. While it would be outside of the scope of this paper 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 to shape performance. For instance, in a gambling situation, subjects might use a conscious strategy based upon the idea that outcomes that have occurred recently become less probable in the future (the Gambler's fallacy, Clotfelder & 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. While these issues again 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 et al, 1996, 1997). In this situation, participants, after being given $2000 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 in a loss. Subjects were told to play so as to maximize gains. Unknown to participants, decks differed in their overall ultimate yield, with some decks being disadvantageous and others being advantageous. 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 stopped. Subjects were probed about their knowledge of the situation at regular intervals during the game.

The results of these experiments indicated (1) that 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 SCR responses to advantageous and disadvantageous decks appeared before subjects were able to motivate their decisions, as if their body knew which decks are risky before the relevant knowledge was available for verbal reports. In contrast, patients with damage in the prefrontal cortex failed to exhibit anticipatory SCR's 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 our own emotional response in previously experienced similar situations. This process is assumed to be severely disturbed for the pre-frontal 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 used to shape further (explicit) decision-making by means of the "somatic marker" (the emotional valence) associated with the knowledge. According to Damasio (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, while 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 paper, our main goal is therefore to explore whether we can replicate Damasio's results using a more complex task based on Reber's artificial grammar learning task. In our adaptation, participants 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 a finite-state grammar, while the other had been generated based on another finite-state grammar. Correct decisions can therefore only be reached to the extent that people learn something about the structure of the stimulus material based on the pattern of successive reinforcements to previous choices. SCR was monitored during all trials.

This approach makes it possible to address several concerns with the original gambling task. First, increasing the complexity of the stimulus material makes it possible to use more indirect questions to probe subjects' explicit knowledge about the task than made possible through the original questions (e.g., "Which deck is the most advantageous?"). Pilot experiments conducted in our lab led us to believe that participants suspected that payout in the gambling task was driven by more complex rules involving for instance responses to earlier trials. Such participants 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 participants more open-ended questions such as "On what basis do you make your decisions?".

Second, the original gambling task involved, at least for the non-computerized studies, a fixed sequence of winning and loosing 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 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 somatic marker 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 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 (*somatic marker hypothesis*). Validating the second hypothesis would lend support to the idea that above-chance selection performance depends on the availability of relevant somatic markers. This would not, however, exclude the possibility that the somatic marker also plays a role even when knowledge about grammaticality has become explicit. Therefore, we also explored the role of the somatic marker after subjects had expressed explicit knowledge of one or more relevant rules.

Method

Participants. Thirty volunteers (10 male and 20 female) aged 18 to 51 (mean = 22.2, sd = 7.1 years) participated in this study. Participants were either acquaintances of the experimenters, or freshman psychology students at the university of Amsterdam who participated for course credit. All participants 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 based on one of the two finite-state grammars depicted in Figure 1, so that all 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 6 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 | symbol. As a results, strings discrimination can be achieved based exclusively on 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 participant's decisions (see e.g., Knowlton et al, 1994). However, our main purpose in using such simplified material in the context of this study was (1) to ensure that participants could verbalize the rules 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 participants 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.

[INSERT Figure 1 about here]

Two Ag-AgCl electrodes were attached to the middle and index finger of the non preferred hand. Isotonic paste was used. Skin conductance was measured with the Orion 4AD22, which determines skin conductance using a constant AC current 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 seconds before the choice and that extended 13 seconds after the choice had been made (see fig.2)

Procedure. Subjects received written instructions describing the goal of the experiment as a learning task. The instructions emphasized the possibility for subjects to earn money. After attaching electrodes to the non-preferred hand, this hand was positioned on a small pillow and skin conductance was measured on a deep breath. 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 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 subjects.

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. Subjects initiated each trial by pressing on any key of the computer keyboard. Two strings were then presented together on the screen. Participants 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 subjects had responded, a delay of three seconds 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 euros. Incorrect choices incurred a penalty of either 10 or 100 euros. For both correct and incorrect choices, the actual amounts were chosen at random (following the varying rewards in the original Damasio 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 seconds, after which a message indicating that the next trial could be initiated appeared on the screen.

To assess 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 to have become explicit (1) when the subject 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 is indeed the case) and the subject mentioned this rule again on trial 40, then this subject was scored as having acquired an explicit rule on trial 25 (rather than at trial 30 because this apparent stable rule could have been discovered anywhere between trials 21 and 30).

[Insert Figure 2 about here]

Results

Unresponsive subjects, defined to be subjects exhibiting skin conductance variability smaller 10 microMho over the course of the entire experiment, were eliminated. Five such subjects were thus removed from the original 39 participants. A further four 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-seconds period extending from the first (baseline) sample up until the point that feedback was given. The resulting measurement therefore represents average skin conductance response during the decision and anticipation phases. It corresponds to Damasio's somatic marker (SM). These 'SM' values were averaged separately, on a subject-by-subject basis, for the correct and 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.*

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 subjects, performance started to increase long before they entered the conceptual phase. Only five 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 of 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 non reporting subjects already increases very early in the experiment. Twenty-one of the 25 subjects had an average score over 50% between trial 10 and 40. Their mean scoring rate was 72.9% ($t= 7.39$, $df=24$, $p < 0.0001$). Based on these results, we can thus conclude that ‘implicit learning’, comparable to the learning found in Damasio’s original gambling task, occurred in this experiment.

[Insert Figure 3 about here]

“Somatic Marker” hypothesis. Figure 4 shows the time course of the average skin conductance over all subjects using only the trials where no explicit knowledge was formulated.

[Insert Figure 4 about here]

The figure suggests that the average skin conductance value is larger before feedback for incorrect rather than for correct choices. Because the distribution of skin conductance measurements is known to be non-normal 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 because the pre-conceptual phase was too short. Of the remaining 28 subjects, 19 subjects expressed a larger SM before incorrect choices than before correct choices, and 9 exhibited the reverse pattern. This difference was significant (binomial $p = 0.044$). Note that the corresponding within-subjects analyses produced non-significant results. In other words, the fact that 9 participants exhibited larger SM for correct rather than for incorrect choices should not be taken as suggesting that these participants exhibited a reverse somatic marker effect. Instead, they simply belong to the lower part of the distribution of responses.

Exploratory analyses

a) 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 as characteristic of the "conceptual phase", that is, all the trials for which subjects had expressed at least one correct rule. Interestingly, this analysis indicated that the effect slightly increased from 67.9% to 73.3% (22 out of 30) subjects (binomial $p = 0.009$).

b) Correct versus incorrect decisions and response times. In most decision tasks there is a trade off between response time and performance. However in a number of tasks that involve non-conscious processes in the realm of perception it has been found that using a pop-up strategy will improve performance [Snodgrass et al, 1993]. This pop-up strategy basically consists of choosing the first alternative that comes to mind, so preventing any further analysis. We compared the mean response times for correct and incorrect decisions in the pre-conceptual and conceptual phases of the experiment (Table 1).

Insert Table 1 about here

It can be seen that response times before correct decisions are significantly smaller than for the incorrect ones. However, it is difficult to interpret this result since the correlation between response times and performance can be attributed either to a causal factor originating 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 larger 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 where 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 artifical grammar learning task. This happened before subjects could formulate any explicit knowledge of the grammar, during what Damasio called the pre-conceptual phase. We found that skin conductance was higher before incorrect choices than before correct choices. This was the case in both the pre-conceptual as well as in the conceptual phase. These findings lend support to the suggestion that the somatic marker process is important in everyday complex intuitive decisions in problem-solving situations that are under-specified, or for which not enough time is available for a complete analytical solution.

Is our measure of conscious knowledge sensitive enough?

Our finding that subject's 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 where 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 participants were repeatedly prompted to verbalize any knowledge they may 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 some good indication that people make successful decisions based on something else than reportable knowledge. Nevertheless, one might argue that our method of assessing participants' conscious knowledge (which closely followed that of Bechara et al.'s) was not sufficiently sensitive to participants's conscious knowledge. For instance, participants 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 falsely accepted because learning was established in a phase where explicit knowledge was already available, yet not verbalized. Although our main goal was to shed more light on the role of the somatic marker and not so much on the hypothesis that implicit learning occurs in this situation, the issue of knowing the extent to which participants possess explicit knowledge before, concurrently, or after somatic markers 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 the exact 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 participants in fact turned out to possess much more explicit knowledge than 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 is itself disputed by Damasio et al. (see Damasio et al., 2005; Maia & McClelland, 2005).

We 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, or 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 behavioural 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 noted by many authors, however, there are several important pitfalls with the simple dissociation logic. First, the measures C and P cannot typically be obtained concurrently. This "retrospective assessment" problem (Shanks & St. John, 1994) entails that finding that C fails to be sensitive to the relevant information need not necessarily imply that 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 like ours, however.

A second issue is to ensure that the information revealed through C is indeed relevant to perform the task. As Shanks & St. John (1994) have suggested, many studies of implicit learning have failed to respect this "information" criterion, also called “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 & Dewey (1984), for instance, clearly showed that probing subjects 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 subject’s 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 participants to verbalize rules, but simply asked 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 to ensure that C and P are both equally sensitive to the relevant information. This is what Shanks and StJohn (Shanks & 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 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 through 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 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 participants in these empirical situations can indeed be implicit to the extent that it is "below the subjective threshold".

However, as Reingold and Merikle (1990) point 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 experiments indicate (Nisbett & Wilson, 1977). For this reason, many authors have advocated using so-called objective measures of awareness. Objective measures of awareness include forced-choice tests such a recognition, presence-absence decisions, or identification. Today, numerous studies have been conducted using objective measures

Even if the different criteria briefly overviewed 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; E.M 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 participants, it makes little sense to assume that conditions exist where awareness can simply be "turned off". It might therefore instead 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 this is precisely the case in our experimental situation, as we 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 participants 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 & McClelland, would help. This impression is supported by the results of an unpublished experiment in which, at trial 20, half of the participants were probed using a much more elaborate method. This involved (1) stressing the fact 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 et al., in preparation). In fact, subjects who had been probed extensively at trial 20 eventually mentioned, on average, their first correct rule no earlier than 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 tasks, 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 McCelland, 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 unambiguous and was easily understood by subjects. All knowledge verbalized by subjects in fact took the form of simple production rules, such as "IF there is a pair of + and * symbol THEN the word is an A word".

One further aspect of our probing methodology deserves discussion. On each probing trial, participants who failed to repeat a rule they had previously mentioned were reminded that they had done so before. When assessing participants's explicit knowledge, we considered that if a rule was mentioned at trial X and not mentioned again at trial X+10, then the rule was not included in participants' 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 if subjects, once they formulate a rule, do adhere to their own rule. In order to check this, we estimated theoretically and by a simulation the mean percentage correct to be expected if a single rule is known. There are 8 rules (i.e., 8 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 6 elements, we have 5 transitions, so logically there are 5 rules fired for each word. However, sometimes a rule is fired twic, so that only 4 rules were involved in the construction of the word. The probability that a rule is fired twice in one word turns out to be 14%. The expected scoring rate, after one correct rule is identified by the subject, therefore is: 5/8 * 86 + 4/8 * 14 =~ 60%. The remaining 40% of the words will be guessed at. Thus the simulation yields that the mean scoring rate (including guessing) will be 80% after a single rule is identified.

Interestingly, the 25 subjects who failed to express any rule before trial 50 eventually correctly classified slightly less than 80% of the strings — the rate that would be expected based on knowledge a single rule. We further explored how well participants 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) subjects who failed to verbalize any rule achieve an 80% correct classification rate, and (2) the fact that those subjects who did express a single rule score better than would be expected based on 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, while we cannot rule out that more sensitive measures would not have resulted in more knowledge being expressed on a direct test of subject's 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, subject's performance appears to be driven by knowledge that they fail to verbalize, even in a fully intentional situation where they were repeatedly probed about any knowledge they might have been aware of.

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

In artificial grammar experiments there are generally no such reinforcements. In those experiments the examples might be ‘stored’ as instances while the participant does not really try to figure out underlying regularities. In the present implementation of the grammar task, by contrast, we expect that participants 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 also ask 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 somatic marker, as measured by skin conductance, is assumed to reflect the emotion associated with a given decision. While 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 positive emotions and one could thus assume that the larger somatic marker 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 somatic marker is used as a warning signal, for we would then expect larger somatic markers before correct decisions. Addressing these issues would require further experiments that make it possible to differentiate the early speculations of the subject from their final decisions. This is the goal of our future work using faster physiological measures, like pupil dilation, as a potential somatic marker.

Acknowledgments

This work was supported by European Commission Grant HPRN-CT-1999-00065 and by an institutional grant from the Université Libre de Bruxelles. Arnaud Destrebecqz and Axel Cleeremans are respectively Scientific Research Worker and Senior Research Associate with the National Fund for Scientific Research (Belgium). Simon van Gaal ran the experiment as part of his Master's Thesis.

REFERENCES

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**-2*,* 215-225.

Bierman, D..J. (in preparation). Using an eye tracker to measure the three factors underlying intuitive decision making.

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

Chan, C. (1992). Implicit cognitive processes: Theoretical issues andapplications in computer systems design. Unpublished Ph.D. Thesis,University of Oxford, Oxford.

Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception and Psychophysics*, **36**, 387-395.

Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: news from the front. *Trends in cognitive sciences,* **2-**10*,* 406-416.

Clotfelter, Charles T. and Cook , Philip J. (1993). The "Gambler's Fallacy" in Lottery Play. *Management Science*, December 1993.

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-biological sciences, **351**: (1346), 1413-1420.

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

Dienes, Z., & 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 University Press.

Dienes, Z., & Berry, D. C. (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 Psychol_ogy: General,* **113***,* 541-555.

Ericsson, K. A., & Siimon, 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 and Cognition*, **22**, 948-969.

Knowlton, B.J., Squire,L.R. and Gluck, M.A. (1994). Probalistic classification Learning in Amnesia. *Learning-and-Memory*. **1**-2, 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**-1, 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 and 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 and 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 Publications.

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

Snodgrass M., Shevrin H. and Kopka M. (1993). The Mediation of Intentional Judgments by Unconscious Perceptions: The Influences of Task Strategy, Task Preference, Word Meaning, and Motivation, *Consciousness and Cognition*, **2**-3, 169-193.

FIGURE CAPTIONS

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

Figure 2. Timing of a single trial. Data are stored from 4 seconds before, till 13 seconds 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

Figure 4. The skin conductance preceding, during and after feedback of incorrect and correct decisions for all subjects averaged over their pre-conceptual trials.

TOP

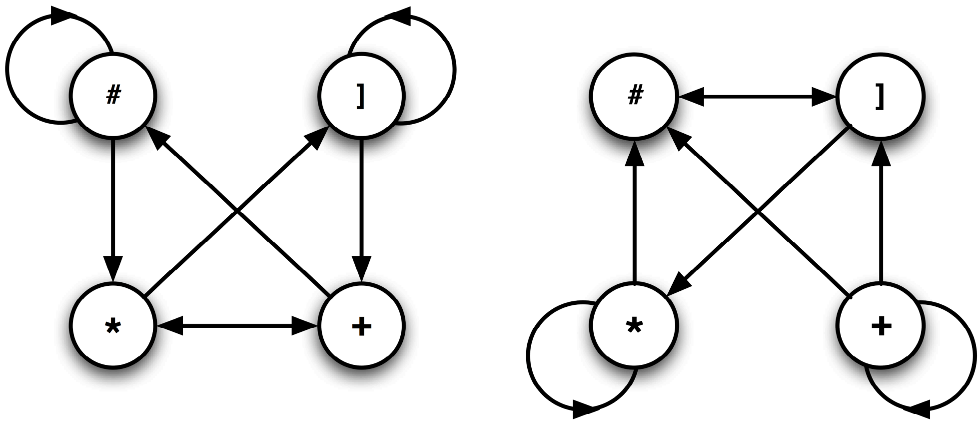

Figure 1

TOP

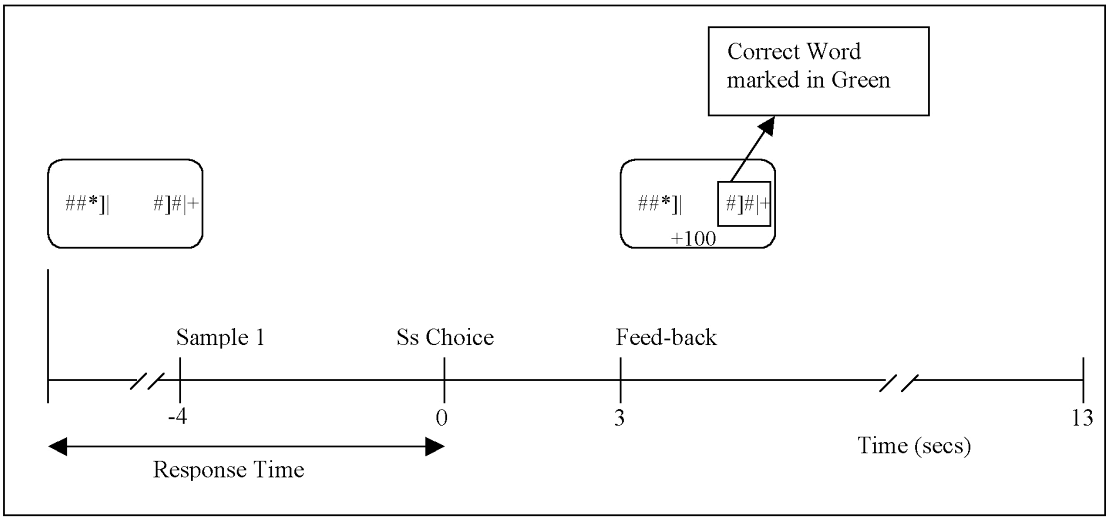
Correct Word
marked in Green
##*]|
#]#|+
##*]|
#]#|+
+100
Sample 1
Ss Choice
Feed-back
-4
0
3
13
Time (secs)
Response Time

FIGURE 2

TOP

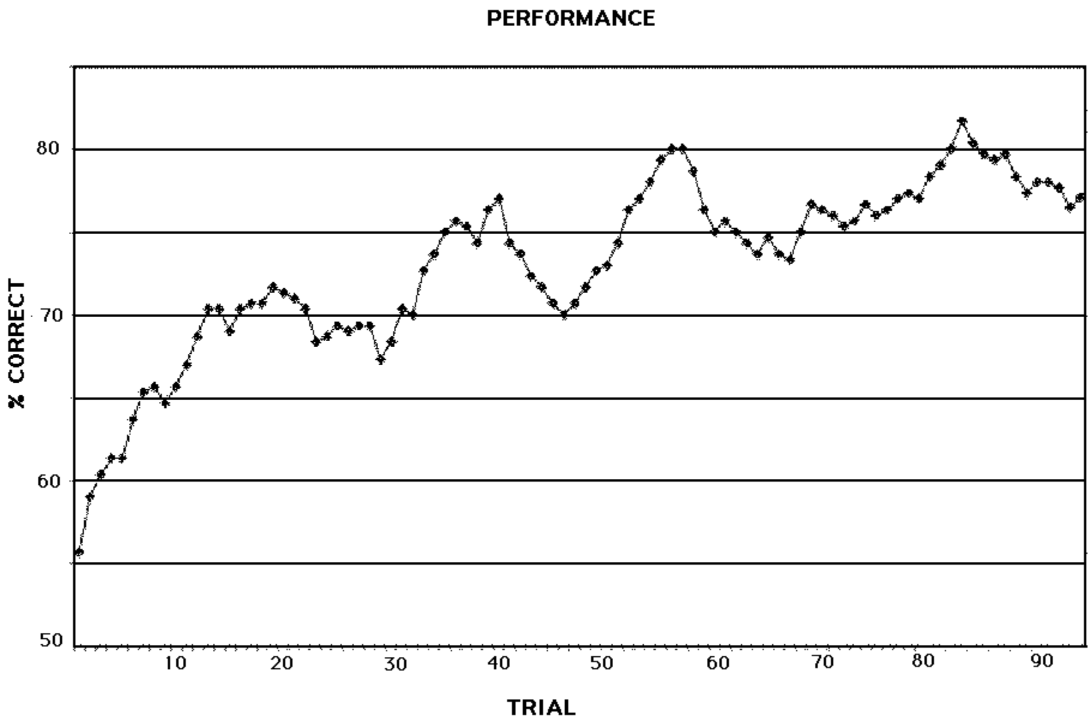

Figure 3

TOP

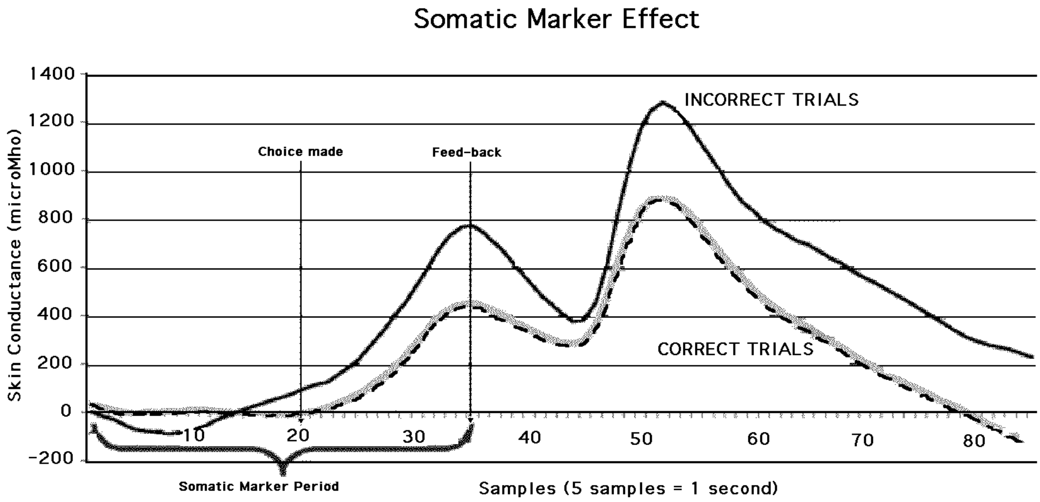
Somatic Marker Effect
1400
1200
1000
800
600
400
200
0
-200
Skin Conductance (microMho)
Choice made
Feed-back
INCORRECT TRIALS
CORRECT TRIALS
10
20
30
40
50
60
70
80
Somatic Marker Period
Samples (5 samples = 1 second)

Figure 4

Table 1: Mean response times for incorrect and correct decisions in pre-conceptual and conceptual phase of the experiment

	Mean (msec)	**Diff (msec)**	**Wilcoxon Z**	**p**
RT incorrect preconceptual	4130	650	3.98	<0.0001
RT correct preconceptual	3480			
RT incorrect conceptual	4257	798	3.42	<0.0001
RT correct conceptual	2459			