A

IMPLICIT PROCESSES AND THE MERE EXPOSURE EFFECT

by

DIANE M. ZIZAK

A dissertation submitted to the Graduate Faculty in Psychology
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy, The City University of New York

2002

UMI Number: 3047282

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

This manuscript has been read and accepted for the Graduate Faculty in Psychology in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

5/1/02 ______ [signature] ______
Chair of the Examining Committee

5/2/02 ______ [signature] ______
Executive Officer

Arthur S. Reber

Glen Hass

Laraine McDonough
Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

IMPLICIT PROCESSES AND THE MERE EXPOSURE EFFECT

by

Diane M. Zizak

Adviser: Professor Arthur Reber

This study involves six experiments which examine the ways in which the mere exposure effect can generalize to structurally similar items. The "classic" mere exposure effect refers to the tendency for previously seen items to be liked more than novel ones (Zajonc, 1968, 1993). The structural mere exposure effect reflects the tendency to show hedonic appreciation for novel stimuli which are structurally consistent with those previously encountered (Gordon & Holyoak, 1983; Manza & Bornstein, 1995). Using an artificial grammar paradigm typically employed in implicit learning studies (Reber, 1993), participants are asked to make both well-formedness and affective evaluations about test strings. During the acquisition phase, participants are presented exemplars of grammatical strings. They are then asked to evaluate test items which are either repeated from the learning set, new yet grammatical, or grammatically incorrect. In some experiments, highly unusual and unfamiliar symbols were used to instantiate the grammar (Chinese and Japanese characters). In others, more familiar and mundane elements (such as English consonants) were used. Whereas subjects were capable of making grammatical discriminations in all six experiments, the results indicate that the nature of the stimuli plays a key role in determining when classic and/or structural mere exposure effects are elicited. With the more commonplace elements, both the classic and structural mere exposure effects were found. However, the use of highly unfamiliar elements tends not to produce either classic or structural effects. In Experiment 5, native Russian readers were asked to evaluate Cyrillic letters. However, the results of this experiment were anomalous. Results discuss the way in which both induced and *a priori* familiarity relates to implicit learning theory, as well as the role of familiarity with respect to the structural mere exposure effect.

Acknowledgments

This dissertation is dedicated to Arthur

In looking back at all of the people who have been instrumental and influential, I find that it's not an easy task to capture them all. Some are more obvious, while others glimmer in my memories. Perhaps there is a message about the impact of implicit processes which applies here. For, without intentionally meaning to, my daughter, Dion Griffin has inspired me to continue with my education, to show her how intrinsically rewarding it can be, and to demonstrate a love of learning that should last a lifetime. My hope is that she is likewise sparked to pursue her intellectual passions.

Back when I was the age my daughter is now, a person whose path I was fortunate to have crossed introduced me to the thrill of the hunt, that is, the excitement of experimental pursuit. The late Charles Honorton was fascinated with the outer limits of the human mind, and was intrigued with its inner mysteries. In spite of my taste for the unconventional, he pointed me in the right direction by encouraging me to study experimental psychology in a solid and reputable program.

Brooklyn College has certainly exceeded all expectations in this regard. I consider myself extremely fortunate to have had the opportunity to study and work with world class professors in the psychology department, all of whom have contributed, both implicitly and explicitly, to the accomplishment of this degree. In that spirit, I could honestly acknowledge all of the faculty members of the department, but a few people stand as being of particular influence. I would like to thank especially my committee members and readers: Glen Hass, my chairman, who is the embodiment of grace under pressure and a pillar of calm within a storm. Laraine McDonough, whose attentive interest and compassionate encouragement has been instrumental in my finding my voice --literally--when I sometimes lost my nerve. Buzzy Chanowitz, who has always made me feel like he's on my side, especially when I needed support; and Stephen Lepore, whose display of enthusiasm for my work was much needed. All of the psychology faculty have contributed a hand in my reaching this goal, and so I would like to express my appreciation to Dorathea Halpert, David Owen, Don Scarborough, Louise Hainline, Izzy Abramov, Andy Delameter, Anthony Scalfani and, of course, Matthew Erdelyi. My special thanks goes to Neil Macmillan, who has helped me out of many a tight spot -- statistically and otherwise.

None of this would have even been conceivable, though, had it not been for Arthur Reber. Not only has he been my ever-patient adviser, but he has become a very special friend to me. Arthur was the one who convinced me to apply to the doctoral program, and it was Arthur who saw through my neuroses and believed in me from the start. I know I put him through a lot. So I dedicate this dissertation to Arthur Reber, whose work on implicit learning provided the foundation for my research, and whose enduring compassion has kept me going. Thank you.

Of course, a special 'thanks' for those friends who cheered, especially Paula and Phil, who both commiserated as fellow-graduate students and succeeded in convincing me that it was worth the struggle. I owe a special 'thank you' to my colleague and friend, Virginia, who has listened to me fret and worry over the years. She has heard an earful.

But these acknowledgments would not be complete unless I express a heartfelt debt of gratitude and affection for my secret inspiration, Macska. He has been there for me when the going got rough--very rough, and would not allow me to throw in the towel. He has provided a special space for me, both literally and metaphorically, in which I can be free to just be. It's a place I hope to know better.

Table of Contents

List of Tables and Figures

While the literature on nonconscious processing is remarkably wide ranging in scope and diverse in perspective, a consistent thread that does emerge suggests that we are capable of processing complex information about our environment generally with ease and efficiency, without really trying. This paper will attempt to develop conceptual links between these nonconscious processes, focusing most specifically on the ways in which implicit learning (Reber, 1993) is related to and overlaps with the mere exposure effect (Zajonc, 1968).

Implicit Learning

Implicit learning (IL) most generally refers to those processes by which information is absorbed and abstracted largely without conscious effort, intent, control or influence. In addition, these processes may operate with even greater efficiency and with more natural ease than some of the more conscious, task-driven analyses (Reber, 1967, 1993). Implicit learning has been shown to be robust, stable and relatively invariant across the general population when compared with performance on some explicit tasks which may alter with age, psychological disposition, and physiological impairments.

Laboratory studies involving IL have included various techniques and methodologies which have creatively constructed situations that are devoid of prior associations, retain arbitrariness, and are sufficiently complex so as to render the regularities nonobvious. For example, Lewicki (1986a,b) has explored IL by pairing personality traits as associations which covaried systematically with specific physical properties of people presented in photographs. Serial reaction time studies have investigated IL using frequency counts and the tracking of structured sequences of events (e.g., Hasher & Zacks, 1979; Nissen & Bullemer, 1987; Cleermans & McClelland, 1991; Hsiao & Reber, 2001). In a more real-life simulation of complex production systems, Broadbent and his colleagues have studied IL in the context of simulating the control of a manufacturing plant (Berry & Broadbent, 1984; Broadbent, FitzGerald, & Broadbent, 1986). And, not surprisingly, there have been reports also implicating implicit acquisition processes in the manner in which young children come to learn about the world (Small, 1990). While IL has been explored in a variety of experimental settings, the technique which has, in a sense, become the anchor design is that of the artificial grammar (AG) learning paradigm, first used in the 1960's (see Reber,1993 for a review).

Artificial Grammar and the Well-formedness Task

The typical AG study employs a series of letter strings, derived from a finite-state grammar that dictates letter order (See Figure 1). Exemplars of these 'grammatical strings' are usually presented to subjects for whom this is ostensibly a memory experiment. Based on the information gathered during the learning phase, subjects are asked to evaluate novel sequences with regard to their grammatical correctness. Although initially unaware of the existence of the grammar, subjects have been nonetheless able to classify which novel strings conformed to the rules of grammar at greater than chance levels. Moreover, the overt memorization of the learning strings is not required; apparently the mere observation of the exemplars is sufficient in eliciting appropriate well-formedness judgments (Brooks & Vokey, 1992; Reber & Allen, 1978).

Representational Form. Although the AG paradigm has been used now for several decades by many researchers, there still remain doubts as to whether subjects are truly forming abstract representations of the inherent grammar, or whether they are in some ways drawing on explicit knowledge about the learning strings. This question may be particularly germane when we explore the extent to which the abstracted information can be applied in a number of different circumstances. The structural display of the instantiated grammar affords a ripe opportunity to examine a host of cognitive applications. For example, the grammatical display could systematically vary across several dimensions, such as complexity, length, constitution and mode of representation. As will be discussed later, it also becomes an appropriate (and agile) tool with which to examine other types of nonconscious applications, such as the generalization of the mere exposure effect.

Figure One - Artificial Grammar

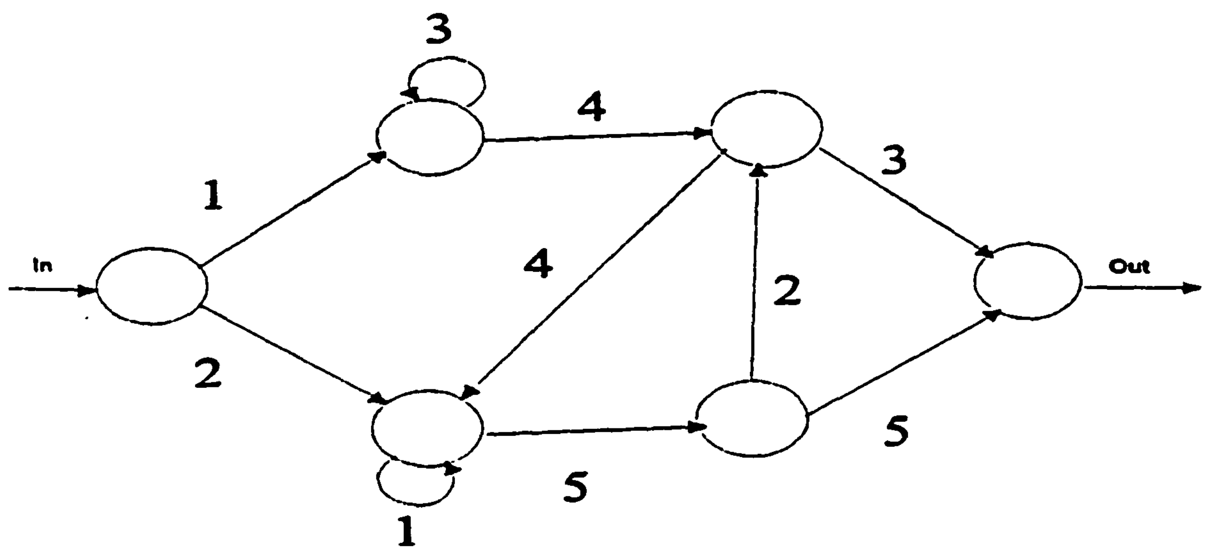

Grammatical strings are generated by entering the above diagram at the left, following the sequentially indicated arrows, until exiting at the right. This grammar can be instantiated using an infinite array of stimulus elements.

Japanese Kana characters:		ア	シ	ル	オ	セ
English:	Letter Set 1:	T	P	S	V	X
	Letter Set 2:	C	L	H	M	J
Chinese characters:		也	四	月	五	百
Keyboard symbols:		%	#	*	&	@
Russian (Cyrillic):		Ж	Ћ	Ч	Д	Љ

If it is indeed the case that an instantiated memory and not an abstract representation is operating during this task, then subjects ought, for example, to be able to indicate *which portions* of the test strings were relevant in their grammatical classifications. Dulaney, Carlson, and Dewey (1984) had subjects either cross out those letters which rendered the strings nongrammatical, or to underline those elements which made them well-formed. The ability for them to do this was taken as evidence for the explicit awareness of the subjects' grammatical judgments. However, this task is problematical. Whereas crossing out pieces of the strings where violations of the grammar render a correct response, a 'holistic impression' about the grammaticality of a string would mean that underlining *any portion* of that string would be considered a hit. Moreover, just because elicited responses, in the form of identifying grammatical/nongrammatical segments, might be viewed as an explicit task, it does not necessarily follow that the information was not *encoded* implicitly, or that this explicit knowledge was available during the classification task (Reber, Allen & Regan, 1985).

Similarly, Manza and Bornstein (1995) included a stem-completion task in their later experiments which was designed to tap whatever explicit knowledge subjects may have had by requiring them to 'fill in the blanks' for grammatical test strings. They found that those subjects who first made well-formedness evaluations performed better on the subsequent stem-completion task than did other groups of subjects. Interestingly, though, even those subjects who were in no way involved with grammaticality assessments were nonetheless able to provide accurate stem-completions at above-chance rates, suggesting an implicit acquisition of the structure in the absence of any explicit mention of the grammar.

On the other hand, if knowledge about only small pieces of the learning strings was sufficient to classify subsequent test items, then perhaps subjects trained only on permissible bi- and tri-grams would also be able to demonstrate their sense of the grammar (Perruchet & Pacteau, 1990). Indeed, subjects were able to do so with accuracy rates matching those who had the opportunity to study entire learning strings. These studies were viewed as evidence that neither implicit learning nor abstract representations need be invoked in order to render significant well-formedness judgments. However, when attempting to replicate this study, Manza and Reber (1997) found that although subjects trained using only bi- and trigrams were able to classify strings at above chance rates, they were unable to transfer this knowledge to test strings which were instantiated using a different letter set–a facility which those subjects trained on full strings did have. Although it may be possible to make certain types of classifications when trained only with pieces of the strings, the information acquired is less thorough than it would be when given the opportunity to learn from entire cohesive strings. Ironically, these studies therefore lend support for the extraction of abstracted knowledge as opposed to discrete memorial representations one would expect in episodic encoding.

Along these lines, a model of 'competitive chunking' suggests that fragmentary knowledge at the time of classification can explain subsequent performance without resorting to automatic abstraction, because fragmentary knowledge representations summarize information across separate learning episodes. Subjects then break letter strings into chunks of 2 and 3 letters which are later joined together into higher level chunks, the strength of which is modified according to frequency, thereby representing an 'abstraction', (without really being abstractions) because prototypical chunks become heavily weighted while the idiosyncratic details of each episode are lost. This 'competitive chunking' presupposes a certain value-laden marker to particularly salient clusters of letters (Servan-Schreiber & Anderson,1990).

On the other hand, it might be more parsimonious to suggest that, in contrast to the notion of an abstracted grammar, perhaps a more surficial, distributive paradigm might account for the success subjects have with grammaticality judgments (Vokey & Brooks, 1992). This view suggests that what is more likely happening during an AG study involves a literal account of the relevant qualities of the stimulus array that includes frequency counts, privileged positions, and other possibly relevant characteristic cues. It had previously been acknowledged that the letters which appear at the beginning and ends of strings appear to have particular memorial salience (Reber & Lewis, 1977). This subtype of abstract analogy is relational in that the exploited information concerns the position of 'X', regardless of 'X's physical incarnation. In this way, the notion of analogy in the absence of literal (physical) similarity could be examined. For example, a relational analogy would result in a sense of similarity between MCTTTM and PVXXXP because of the *pattern* of letters is maintained. Literal similarity is absent from this example because an entirely new letter set was used to instantiate the same grammatical string. If transfer studies did not control for the literal similarities across letter sets, then essential patterns that earmark a string as being grammatical remain available. The subtle distinction is that abstracted knowledge assumes that a knowledge base is acquired across the general pool of grammatical exemplars. The acquired knowledge then is not stimulus string-specific, but rather holistic in its comprehensive inclusion. However, it could be argued that this level of analogy, albeit quite simple and direct, is merely one form of abstraction that contributes to the notion of an underlying structure being tacitly acquired.

Reber and Winter (1994) prefer to focus the discussion of implicitly acquired information on the *representations*, and not the debate over encoding of instances vs. the abstraction of rules. "Rules...must be the result of abstractions...which are induced...and can easily be thought of as beginning with the encoding of the fragmentary elements that made up the initial instances. If the implicit learning of artificial grammars results in an abstract representation, the abstract representation itself is founded upon two and three-letter chunks that are present in the learning" (pp. 418-19).

As Whittlesea and Dorkin (1993) have shown, the nature of the memorial representation that subjects develop in IL experiments is functionally determined by the conditions of learning and testing. McAndrews and Moscovitch (1985) noted that the number of learning strings displayed during learning also effects the nature of the acquisition. When presented with relatively few strings during the learning phase, subjects have a tendency to form more concrete representations in the absence of sufficient exemplars with which to generate rules. There have been a number of demonstrations (i.e., Manza & Reber 1997; Brooks & Vokey, 1991; Mathews et. al., 1989) that knowledge acquired with stimuli instantiated in one physical form can be used to make decisions about stimuli instantiated in a different physical form--even to the point of using a different modality, that allow for the conclusion that learning involves the abstract memorial representations which are required for such transfer. Mathews (1997) argues that in order for experiments to demonstrate implicit abstraction, they must be designed to include tasks that demand it, as well as experience (in the form of a variety of practice situations) that will produce it. As will be discussed later, it appears that the demands of the encoding task to a large extent determine the quality of knowledge representation. Task performance, then, would be intimately determined by the consistency of encoding and retrieval conditions (Neal & Hesketh, 1997).

> "Knowledge representation appears to be critically dependent on the demands of the encoding task, and performance is strongly influenced by the match between encoding and retrieval conditions." (Neal & Hesketh, 1997; p. 29).

They would prefer to distinguish the task dissociation processes in terms of the participant's 'intention'. The distinction then, should be one of the intention during *both* the learning and testing, as well as to how this knowledge would then be applied strategically. This actually is quite consistent with Reber's (1993) assertion that implicitly learned knowledge is acquired passively (and incidentally). This information then enables us to behave adaptively to the global properties of the environmental demands.

If, as some claim, there are separate and distinct systems operating on instance vs. rule learning, we need to examine how the registration of instances can produce the effect of the learned relationships between those items. Shanks and St. John (1994) choose to parse the field in terms of both implicit/explicit acquisitions, as well as instance vs. abstractions. Stadler and French (1994) view both learning and memory retrieval as being either implicit or explicit. They offer the possibility of explicit learning with memory implicit. In this case, relevant stimuli/relationships are explicitly acquired (encoded) but their subsequent influence on behavior is without intention, as is seen with the automaticity of practiced behaviors. Likewise, they posit the occurrence of implicitly-learned material that is explicitly retrieved. Over time, subjects may become aware of the information that they are utilizing even though it was originally encoded implicitly. The learning of instances reflects an aspect of discrete memories against which there may be a number of plausible operations, such as prototype matches (Posner & Keele, 1968); similarity, (Vokey & Brooks, 1992); chunk strength, (Servan-Schreiber & Anderson, 1990); and fluency (Whittlesea, 1993).

Flexibility

The seemingly intractable inflexibility of implicit memories has often been used as a barometer with which to distinguish them from explicit memories (e.g., Dienes & Berry, 1997; Schacter, 1987). Analogously, implicitly acquired information has been argued to be similarly rigid and inflexible (Neal & Hesketh, 1997; Dienes & Berry, 1997). However, common sense shows that this parallel does not necessarily hold. Take, for example, our learning of and fluent use of language (Mathews, 1997) where not only young children, but certainly adults have highly flexible production and comprehension of novel utterances independent of explicit prescriptive knowledge. It appears that the more we look to real world evidence, the more compelling the position of the primacy, fluidity and universality of implicit processes.

In their 1982 paper, Jacoby and Witherspoon discuss how the distinction between episodic and semantic memory is congruent with the abstractionist view (e.g., Reber, 1989; Dienes & Berry, 1991). This position maintains that the rules abstracted from a stimulus array are done so in such a way as to render them separable from the stimulus. In this way, rules and concepts (semantic memories) derived from the experience are encoded separately from the actual memory of that event (episode). This characterization of how memory works suggests that the "priming of an existing concept provides a form of memory that is relatively context free." (Jacoby & Witherspoon, 1982, p. 311). This view differs from one which suggests that concept formation actually makes use of particular memory for an individual instance (e.g., Brooks, 1978) and would therefore be specific and predictably inflexible.

Those features of the stimulus which are salient to the processing of the information should be consistent with the kinds of operations expected to be performed later. First, the ways in which materials are invited to be encoded ought to match the subsequent demands made on this information, and it is just this congruence that predicts the flexibility and application of both implicit *and* explicit knowledge (Willingham, 1997). For example, when the font style of a word is a relevant feature of subsequent recognition/memory tests, then it ought to be the font style which is being manipulated. In this regard, if the semantic identification of the stimulus word is of relevance to the test situation, then deeper, semantic processing would be required during the learning phase. The application of such rules implies a degree of flexibility. Moreover, when information acquisition requires the abstraction and later application of rules and concepts to be applied, then it becomes unreasonable to suggest that this implicit knowledge is not flexible in that it can be applied to novel situations. This is relevant because of the nature of the stimuli utilized in the AG paradigm encourages flexible applications of knowledge--at least as measured by well-formedness judgments.

Evaluative Learning

Implicitly acquired information has been demonstrated in a number of other realms, as well. For example, the pairing of items in ways that foster associations between the two can successfully result in their shared associative value, even when the explicit recall of the items being paired is unavailable.

Back in 1960, Johnson, Thomson and Frincke found that when pairing nonsense syllables with words rated as high/medium/low on the word value association scale, subjects tended to like better those items which were linked with the higher associative values (on the 'good'/'bad' scales). Developing this concept, Baeyens, Elen, and Van den Bergh (1990) demonstrated how previously neutral-rated words change their affective valence as a result of having been paired with words which subjects had previously designated to be highly affective (see De Hower, Baeyens, & Hendrickx; 1997 for a review). Moreover, even when subjects were unable to indicate any explicit memory for which item was paired with which, or even which *type* of item (positive or negative) was paired with the neutral ones, they nonetheless demonstrated increased affect for those items which were paired with the highly positive words. This they were able to do even when they had reportedly low confidence in their decisions.

What is also particularly interesting about these studies is that when subjects were cued as to what the initial pairings were about, they tended to demonstrate a diminished effect relative to the naive subjects. In other studies, the pairing of neutral and positive-valenced words presented the strongest effect when subjects reported to be *totally unaware* that a second, subliminally presented word was being presented right after the neutral item.

In a more recent study, Murphy and Zajonc (1993) tested the effects of both suboptimal and optimal affective primes. Subjects in this study were asked to rate how much they liked Chinese ideographs that were preceded by a variety of primes, which included: faces that were clearly either "happy" or "angry", random polygons, or, in some cases, no primes presented at all. The ideographs were primed either suboptimally (at 4 msec. with a backward mask), or optimally (at 2 seconds). A 2AFC test of awareness confirmed that the suboptimally presented primes were unrecognizable, and subjects reported not having been aware of the primes presented. However, results showed that only under the suboptimal, and not the optimal condition, did affective priming become effective, in that the ideographs which were preceded with happy faces elicited significantly higher liking ratings, and those preceded by angry faces resulted in lower ratings. None of the control primes produced any such shift in affect.

These results speak to the issue of how conscious editing may actually inhibit the effects of certain implicit and nonconscious acquisitions. Spence and Holland noted in 1962 that impaired perception often resulted in stronger effects, a phenomenon that seems counter-intuitive. However, it might be argued that when perceived consciously, stimuli may reveal certain characteristics or features which could result in contradictory affective appraisals (Lazarus, 1982). In addition, such overt manipulations as the pairing of affectively laden primes with otherwise neutral stimuli could 'turn subjects off' to the associations which are blatantly being fostered. Participants are not receptive to these conspicuous attempts at shifting their judgments, and, in effect, reject the manipulation. Perhaps it is only when the source of the affective shifts remain a mystery to the subjects that evaluative learning is effective (Winkielman, Zajonc, & Schwarz, 1997). More of this will be discussed later.

The Mere Exposure Effect–A Short History

In 1968, Robert Zajonc noted that evaluative judgments were established and developed in ways other than associative pairings of apparently neutral items with highly affective ones. He instead examined the effects of word frequency on affect by deliberately manipulating the number of times each of his set of "Turkish words" was presented. These manufactured seven-letter words were presented for either one of five frequencies (1, 2, 5, 10 and 25), or not at all. He found that the subsequent 'goodness' ratings for the 'words' at each of the six frequencies mirrored the amount of exposure, thus providing evidence for the now-famous 'mere exposure effect'.

Concerned that the repeated pronunciation facilitated the subjects' articulations, rendering these items attractive by virtue of heightened proficiency (Wilson & Becknell, 1961), Zajonc (1968) also used a stimulus set that would retain ambiguity while providing less opportunity for repeated subvocal articulations and rehearsals, and so switched to using pseudo-Chinese ideograms in place of the alleged Turkish adjectives. Subjects were once again informed that they were participating in a study of foreign language acquisition, and were not required (nor expected) to pronounce any of the characters which were being presented. The data again revealed a clear increase in affective ratings for the "characters" which were presented at the greater frequencies.

Getting the Terms Straight

Zajonc does note that although subjects were asked to rate the test items in terms of 'goodness', they were never explicitly asked how much they 'liked' them. This might be a qualitatively different task. Although 'good/bad' scales seem to produce the mere exposure effect as do 'like/dislike' scales, Zajonc, Markus, and Wilson (1974) suggest that whereas 'good' scales measure a feeling of acceptance and affinity, 'like' scales tap curiosity as well. In rating novel stimuli for both 'goodness', (as in, "What does this paralog [nonsense syllable, character, symbol, etc.] mean?"), and 'liking', (as in "How much do you like each of these items?"), Saegert and Jellison (1970) found that both kinds of measures increased as a function of greater exposure frequency. And, as Bornstein's impressive meta-analysis of this huge body of research revealed (1989), although a number of affective terms such as 'like', 'prefer', 'goodness', and 'wellness' are all effective in eliciting the mere exposure effect, certain probe questions of this general nature did have a tendency to work better than others under certain circumstances.

There are indications that earlier researchers were intrigued with similar phenomena. In terms of the visual arts, Pepper (1919) noted that repeated exposures resulted in greater aesthetic appreciation for unusual color combinations. Maslow (1937) re-presented pieces by famous artists repeatedly for 15 days, after which he interspersed items from the same artists which had not been previously used in the study, finding that subjects indeed grew to like those specific paintings which they had previously encountered. Alpert (1953) found that repeatedly presenting subjects with unfamiliar rhythms increased their enjoyment of such patterns. Although initially complaining that these rhythms were actually unpleasant, subjects grew to appreciate and even enjoy the 'music'.

The implications for consumer behavior could not be missed here. Becknell, Wilson and Baird (1963) varied the frequency with which nonsense syllables were presented, and then offered subjects a choice of products labeled with these "brand names". Although preference-ranking of products and forced-choice preference decisions varied methodologically from the classical mere exposure studies, it is interesting that these earlier studies were nonetheless able to demonstrate product selection as a function of exposure frequency. In what will eventually come to be a poignant query, Zajonc asks "whether...repeated exposure to a given stimulus results in the enhancement of attitudes toward similar or related stimuli?" (1968, p. 25).

What Becomes Familiar. Zajonc duly notes that the "...impressive correlation coefficients...may not reflect the effect of frequency on attitude but the effect of attitude on frequency. Thus, it can be argued that many roses are grown because people like roses. But it can also be argued that people like roses because there are many roses growing." (1968; p. 3).

Studies also show an overall tendency for people to rate as more favorable those items on a list that appeared to be most *familiar*. For example, countries and cities were rated as a function of popularity. Trees, fruits, vegetables were 'preferred' most often in terms of their appearance within one's realm of immediate experience. Bonnano and Stillings (1986) found that judgments of both 'preference' and 'familiarity' behave in a parallel fashion when compared to recognition judgments which have been altered through experimental manipulations. The link between 'familiarity' and 'preference' has been explored elsewhere (Whittlsea, 1993), and becomes somewhat relevant when discussing fluency effects. However, this may be more of an indication of the ambiguous nature of the task demand in that it is not clear just how the probe question ('familiar?', 'prefer?') is interpreted by the participant.

Theoretical Models

It would be expected that familiarity would necessarily have high survival value. It makes reasonable sense for an animal to prefer (or be comfortable with) an item or situation that had previously proved to be non-toxic. Yet, much of the animal approach-avoidance literature seems to point to a tendency for animals (and humans) to display behavior that seems to challenge this assumption in that they often exhibit exploratory attention and approach behavior towards novel items and situations (e.g., Berlyne, 1950). Although seemingly contradictory, Zajonc points out that this apparently inquisitive response actually makes good ecological sense in that the novel items might very well turn out to be a foe rather than a friend. Without the necessary exploration (and thus, orientation), some animals could not identify and avoid danger. Likewise, through exploring and (repeatedly) experiencing a given situation in the absence of negative consequences, the environment then becomes user-friendly. Novelty is in this way associated with uncertainty, unfamiliarity, and therefore, possibly with conflict and danger.

Harrison's response-competition model (1977) explored the relationship between 'liking' and 'exploring' using nonsense syllables, Chinese characters, and photographs of men's faces to obtain 'liking' responses in one group of subjects, and measures of 'exploration' in another. He concluded that the correlations between 'exploration' and 'liking' were indeed negative. The function of orienting behavior might better be thought of as a means of trying to bring the unfamiliar into the realm of the familiar; and, if not necessarily *enhancing* affective attitudes, then *decreasing* initial negative biases.

Two-stage Models. Berlyne's 2-stage inverted-U model (1970) suggests that the initial response to repeated presentations is one of habituation, which accounts for increased affect; followed by tedium, or increasing disinterest. Repeated exposure to a stimulus provides opportunities for learning about it, thereby developing a sense of familiarity. However, after the learning is essentially complete, satiation occurs and interest wanes (Stang,1974).

Along these same lines, the initial attraction, and then subsequent decline, of affect/interest (as in the inverted-U curve that reflects much of the data) might also be explained in terms of levels of optimal arousal induced by the stimulus (Kagan, 1971). The disruption of already established visual percepts by a novel stimulus might evoke negative emotion (fear?). However, once the item is assimilated into the realm of the familiar through repeated exposure, it becomes less disruptive. The stimulation that is being received, or has in the past been received, results in a pattern of adaptation to that stimulus. Stimuli which deviate from this homeostatic level of ideal adaptation cause excitatory responses that could, in small increments, be considered positive (attractive? a curiosity?), and, in greater degrees, negative (a threat? unknown danger?). According to this model, those stimulus items which differ only slightly from the expected, but are still within the realm of the somewhat familiar, should elicit greater evaluative responses. (As an aside, Bornstein (1989) notes that there appears to be a tendency for young children to not only not display the predicted mere exposure effect, but rather to exhibit the opposite pattern of results: they tend to prefer novel to familiar stimuli. Circularly, one could offer that children tend to get bored more easily, and so the inverted-U is present from the beginning, but this is seen as Bornstein as further support for the evolutionary model in that youngsters have an external monitor, namely their care-taker, to discriminate potential dangers. Having such a screen, young children are free to explore and learn more efficiently about their environment.)

Harrison, (1977), however, noted that the problem with theories relating to the effect to adaptation levels is that much of Zajonc's work indicates greatest affect at the adaptive levels. Perhaps, as Berlyne later suggested (1970), it is the *change* in level of arousal that contributes to affective changes. Thus, it is the arousal potential, as determined by the stimulus's particular properties, that contribute collectively to the overall response level. By definition, the repeated encountering of a stimulus would necessarily reduce (negative) arousal attributions, resulting in an elevated response. Consequently, over-exposure to the affable stimulus characteristics would eventually render them less interesting (arousing?), eventually resulting in the declining curve. This interpretation implies that complex stimuli are required to be presented at greater frequencies before satiation--and subsequent reversal--of affect is to appear. "Deviations from an optimal level of arousal are experienced as 'unpleasant', and thus alter the evaluative judgment of a stimulus when it is initially encountered. The potential of a stimulus for producing arousal depends on both the conditions under which it is encountered and on the novelty, complexity, and incongruity of the stimulus itself". (Moreland & Zajonc, 1977, p. 172).

Stimulus Complexity. The complexity of the stimuli used during the exposure phase could predict the rapidness with which a subject might become 'bored' with it. Stimuli which are insufficiently stimulating might not only fail to elicit the expected initial rise in favorable attitude, but might actually reflect a degree of diminished interest. This pattern of behavior supports the notion of response-competition, as set forth above (Berlyne, 1970; Stang, 1974) and developed further by Saegert and Jellison

(1970), who found a main effect of stimulus complexity when presenting subjects with 'Chinese ideographs' at various frequencies and differing in levels of complexity. They found that 'liking' for complex stimuli increased as a function of frequency exposure; whereas 'liking' for simple stimuli reached asymptote after a relatively small number of exposures. However, Stang and O'Connell (1974) did not find any difference in 'liking' between simple and complex computer-generated line drawings. Zajonc et. al (1972) compared ratings of abstract (complex) paintings with small sections of these paintings (i.e., simple stimuli) and were only able to elicit the mere exposure effect for the simple items, but not for the complex. Musically, Heyduk (1975) found the reverse: subjects preferred the more complex melodies over the simpler tunes. The factor of stimulus complexity is tenable, at best, and needs to be examined.

However, as Bornstein (1989) notes, those studies which indicate heightened preferences for complex stimuli used within-subject designs. The one study by Zajonc et. al.'s (1972) which demonstrated the relative effectiveness of simple stimuli was the only one that used a between-subject design, suggesting that differential ratings of complex/simple stimuli rely on subjects' ability to draw comparative evaluations.

A Synthesized Approach. Bornstein (1989) proposed integrating a two-stage model of processing along with an evolutionary one to handle the mere exposure effect. It is assumed that, initially, it makes adaptive sense to proceed with caution when approaching something novel in the environment. Only after (repeated) non-punished exposure would it be survival-wise to acclimate (and come to like, to fear less, etc.) the environmental stimulus in question. Eventually, further exploration could prove counter-productive for exploratory and acquisitional benefits. As a result, the reaction would eventually become one of diminished interest, boredom, and lower affect with regard to that same item.

Zajonc recently articulated this formulation of the mere exposure effect by characterizing this type of nonconscious learning in terms of classical conditioning (Zajonc, 2001). The absence of the aversive consequences would function as the unconditioned stimulus (US). In other words, repeated exposures to non-toxic events, which are paired with the absence of harm (i.e., safety), result in a form of conditioned learning. The conditioned stimulus (CS) would be the benign repetition of the stimulus, in much the same way that the subsequent and consistent absence of harm might then be construed as the unconditioned stimulus (US). The conditioned response (CR), then, might be the approach tendency as expressed by increased preference. Zajonc has the CS-CS model of classical conditioning expressed by the presentation of the stimulus paired with neutral (non-toxic) events, resulting in the CR of approach behavior, in whatever form the task demands (i.e., liking).

This model also entails the interesting hypothesis that the mere exposure effect should be stronger when the affective evaluations are made after a delay–upwards to days–following the initial exposure (Seamon, Brody & Kauff, 1983; Zajonc, 2001). This temporal pattern makes ecological sense in that items which have been represented in more of a long-term memory trace would appear to be reliably less toxic, dangerous or threatening, and thereby more readily accepted. Affective responses would be the direct result of mere exposure, and not necessarily with the intervention of cognitive analyses. This model also accommodates the above-mentioned tendency for subliminally and marginally recognized items to elicit relatively stronger mere exposure responses than those which are readily recognized, a topic to be pursued in more detail below.

Importantly, both Harrison's (1977) response-competition model and Berlyne's (1970) optimal arousal theory necessarily imply the identification and/or recognition of the stimuli presented. Zajonc's (1980; 1984) approach, however, is predicated on the *separation from* and *primacy of* the affective response relative to cognitive processing. This conceptualization was a key feature of the position set forth by Zajonc back when research in this area was just beginning. Ironically, although left for dead, Zajonc's conceptualization actually anticipates some of the cogent questions which arise when the mere exposure effect is examined in terms of implicit learning. (For a thorough review of this debate, see Lazarus,1984 and Zajonc,1984.) Support for the position could also be found in the tendency for subliminally presented stimuli to have affective influences on subsequent evaluations. The affective primacy hypothesis "hinges on the assumption that the simple affective qualities of stimuli, such as good versus bad or positive versus negative, can be processed more readily than their non-affective attributes." (Murphy & Zajonc, 1993). Evidently, suboptimal primes succeed in altering a subject's response only insofar as emotionally-laden responses are concerned. This 'specialness' of affective responses was systematically examined by Murphy and Zajonc by presenting subjects with primes that were designed to influence a variety of other subsequent judgments, specifically, gender identification, symmetry and size.

What they found was that these primes were only effective when presented optimally, that is, when subjects were able to report seeing them. However, these (non-affective) primes were ineffective when presented suboptimally and had no bearing on the reported judgments of the target stimuli.

Bornstein (1989) maintains that the modified two-factor evolutionary model of exposure effects supports Zajonc's assertion that (sophisticated) cognitive activity is a relatively late-bloomer in that to think otherwise would necessarily be to imply that a total shift away from primary affective responses occurred at the point at which cognitive evaluations evolved as the mediating system.

The Paradox of Stimulus Recognizability

Zajonc, Shaver, Tavris and Van Kreveld (1972) further investigated the recognizability and discriminability of various pieces of modern art by presenting one group of subjects with slides of very similar paintings presented at varying frequencies, while presenting another group with pictures that differed greatly from one another. It was expected that subjects would be better able to discriminate (and hence recognize) those paintings viewed in the heterogenous group. The 'liking' ratings for the dissimilar group reflected the predicted mere exposure effect only insofar as they related to frequency of exposure in a curvilinear manner: first ascending, and then descending when recognition was achieved; whereas those ratings for the homogeneous group of paintings were related to frequency of exposure with a more regular and monotonically increasing slope. The suggestion here is that discriminability of the paintings (and hence, recognition) interfered with the predicted mere exposure effect of attenuated frequencies. Zajonc et. al. concluded that the recognition of a stimulus actually served to *lower* its arousal potential, resulting in quicker stimulus satiation and consequently less favorable 'liking' ratings. Evidently, recognition is an important consideration especially within the Harrison/Berlyne models which rely upon stimulus identification. On the other hand, recognition may not only not be required to elicit the mere exposure effect, it might even derail it.

Kunst-Wilson and Zajonc (1980) noted that if heightened affect is not the consequence of getting to know more about a stimulus after repeated exposures, then "overt affective responses may be unrelated to prior cognitive outcomes which result from stimulus exposure" (p. 558). Their technique was to deliberately impoverish the 'knowability', (i.e., recognition) of the items exposed. Even when the recognizability of the previously-seen stimuli was at chance levels, affective ratings nonetheless reflected a preference for those which had been previously presented. This finding represents an straightforward demonstration of the potency of the mere exposure effect--specifically in the absence of what was considered to be conscious awareness, as measured by verbal reports of recognizability.

In their experiment, Kunst-Wilson and Zajonc used 20 arbitrarily constructed and discriminable octagons which were tachistocopically flashed for 1-msec. Each stimulus was flashed five times in a random sequence to subjects who were simply instructed to pay attention to the flashes. The task then involved the viewing of pairs of octagons (new::old) for one full second. Subjects were asked to make both affective and recognition judgments as well as to provide confidence ratings. Although recognition performance approximated chance (48%), the previously-flashed stimuli were preferred over new stimuli significantly more often (60%). Overall confidence ratings were higher for affective than recognition judgments, although both types of responses were at chance levels when subjects were reportedly 'just guessing'. These results support the suggestion of separate cognitive and affective processing systems, as well as demonstrating that conscious awareness (recognition) is not a prerequisite for the mere exposure effect.

In a similar vein, Wilson (1979) asked subjects to verbally shadow a passage while simultaneously 'correcting' its written transcript. Using a dichotic listening paradigm which required subjects to track aural passages presented in one ear while ignoring a tonal sequence in the other, subjects were subjectively (reportedly) unaware of the tonal patterns that had been presented during the shadowing task, but nonetheless reported heightened preferences for tonal sequences which had been presented in the non-tracked ear. Together, the results of these two studies appear to contradict Harrison's response-competition model of the mere exposure effect, for in order for tension to be reduced through repeated exposures, it would necessarily follow that a certain amount of cognitive processing (i.e., identification) is required. Wilson argues that it may be the change in affect, brought about purely as a consequence of repeated exposure, that helps mold our cognitive percepts (learning), and not the other way around.

However, a more accommodating and contemporary view would suggest that we could not altogether eliminate the possibility of rather complex cognitive processing occurring in the absence of subjective awareness. The measure of (non)awareness in these studies was deduced by the chance levels of reported recognition. Issues abound in this area with respect to the appropriate assumptions made about the exhaustiveness of such an indicator (e.g., Merikle and Reingold, 1991; Holendar, 1986). If, on the other hand, we are not looking to these early studies exploring the mere exposure effect as proof positive of the existence of nonconscious processing, but rather as an interesting phenomenon in itself which is expressed (especially) in the absence of reported stimulus awareness, then these studies shed an important light on some of the ways in which we might develop affective responses.

Murphy, Zajonc and Monahan (1995) investigated the combined effects of both repeated exposure and of affective priming (the pairing of neutral stimuli with postively-valued items). First, what they found was evidence for the classic mere exposure effect, in that the effects of having stimuli repeated resulted in a degree of heightened affect, regardless of whether these presentations were presented suboptimally or were fully visible to the subjects. Interestingly, there was no added increase in affect when the visually accessible repeated stimuli were presented for longer durations. However, the additive effect of the affective primes was only found when these primes were presented suboptimally, resulting in a spilling over of the valence of the primes onto the target stimuli. These results are consistent with the notion that affective responses whose source remains unknown are less constrained than those whose source is apparent. Affect is deemed nonconscious "when the person is not aware of its source, its target, or both" (p. 589). This compensates nicely for the effectiveness of the mere exposure effect, even at the supra-optimal level, in that subjects are generally unaware of the *connection* between the repetition of the stimuli and the expectation of increased affect. In this way, repeated exposures, at whatever frequency and duration, fail to explicitly connect the manipulation with the result. This study furthers Zajonc's original position of the primacy and separateness of affective processing in that the notion of "cognition...even when nonconscious, appears to be dedicated and content specific, influencing only related concepts" (Murphy, Zajonc & Monahan; 1995, p. 590).

Conscious Awareness

Defining the parameters of subliminality as it has been used in the laboratory has been fraught with theoretical issues and methodological problems (see Holender, 1986 for a review). First of all, it becomes all but impossible to identify that magic boundary across which the threshold of conscious awareness is traversed. And, even if it were possible to calculate with precision what this might be, we then have the methodological handicap of having continuously to adjust this elusive parameter to accommodate such factors as individual differences concerning focal angles and ambient illumination, as well as internal variables to include motivation, divided attention, etc. Be that as it may, the notion of subliminality generally refers to that technique (whether tachistoscopically induced, or via various masking schemes) which renders the stimulus to be identifiable to the subject only at chance levels of accuracy. Reingold and Merikle (1988) characterize this zone as consisting of that window of exposure which includes the range that just exceeds the minimal exposure necessary to register that 'something' was presented (the objective threshold) on the one hand, yet remains below the exposure that when rendered allows for stimulus identification (subjective threshold). Anything presented below the objective threshold would, by definition, elude the possibility of perceptually registering a specific item; whereas presentations above the subjective threshold would garner above-chance recognition rates that would signify conscious awareness. The effectiveness of stimuli presented below the subjective threshold is of particular interest not only as it relates to the mere exposure effect, but also with regard to matters such as examining how affective associations are formed, investigating the parameters of priming effects, and studying how perceptual facilitation is rendered.

Several researchers have been suggesting that nonconscious information could be filtered from a cluster of responses (i.e., Erdelyi, 1985; Reingold & Merikle,1988). As was the case in the implicit learning studies, the mere exposure effect has been examined by monitoring the responses to characteristics of the stimulus domain which the subjects do not consciously register. As a means for separating the conscious and nonconscious contributions to response patterns, Reingold and Merickle suggest that, in addition to probing for information about the stimulus array in a manner that is specifically dedicated to the relevant dimensions, subjects may also be asked to provide indirect assessments of their experience. When these indirect measures exceed the magnitude of the more direct (explicit) knowledge which is elicited, it then becomes fairly evident that nonconscious (implicit) knowledge is being tapped. In addition to comparing the effects of supra- vs. subliminal stimuli presentations, several researchers have also provided reinforcement of the hypothesis that stimuli perceived without awareness can indeed influence subsequent behavior. Perception in the absence of overt awareness can be assumed when there is evidence of a stronger indirect measure to a stimulus, (in this case, affective judgments) than of a more direct measure (i.e., recognition).

The differential responses to a singular stimulus episode has also been discussed in terms of Erdelyi's dissociation paradigm (1985). Put simply, this model allows for the possibility that there is more information available than what might be accessible to the subject at a given point in time. By repeatedly probing for responses, it may become possible to increase the amount of retrievable information. Likewise, alternative conduits may succeed in accessing information that had been otherwise inaccessible through a single indicator.

Summarizing the general findings in the subliminal mere exposure experiments, Bornstein, reports that "subliminally presented material results in effects that are three to four times as strong as when stimuli are recognized" (1989, p. 196), and elsewhere simply reported as 'twice as strong' (Bornstein, & D'Agostino; 1994). Whatever the actual potency of the effects, it does appear that information acquired nonconsciously has the potential to have a very real impact on the way we respond.

Perceptual Fluency

Stimuli which have been processed before are re-processed with greater ease than those stimuli which had never before been encountered, even if the original encounter is not explicitly recalled (Whittlesea, 1993). Perceptual fluency predicts that (rapid) judgments on *any relevant dimension* should favor the distinct (recently activated) item (Jacoby, 1983). When subjects are explicitly aware of the recognizable stimuli that they are asked to evaluate, their responses could include adjustments (or recalculations) that would account for this fluency (and familiarity). However, when stimuli are presented subliminally, repeated exposures continue to enhance their perceptual fluency without contributing to their identifiability. Subjects would not then have the opportunity to make 'the corrections' in their evaluations to account for this heightened fluency, thereby misconstruing a sense of 'familiarity' as 'liking'. Therefore, stimulus awareness might actually be *inhibiting* the mere exposure effect, shedding light on why subliminally presented items demonstrate the effect more strongly.

To test experimentally the supposition that subjects are 'correcting' misattributed 'liking' for perceptually facilitated stimuli, Bornstein and D'Agostino (1994) manipulated the experimental set by altering the instructions under which subjects were asked to respond. Ten photographs were presented to subjects at 5 msec each. One third of the subjects were told that the test items they were about to rate (affectively) were all seen before, (providing them with the information that would enable them to recognize the connection between the sense of familiarity and any subsequent increases in 'liking'), one third were told that all of the test items were new to them, and the remaining third were not told anything about the stimuli at all. Those items which actually were previously presented elicited higher 'liking' ratings than their novel counterparts across all three conditions, thereby producing the exposure effect. However, the instructional set did have a significant effect on the pattern of results in the predicted way: the 'liking' ratings of the to-be-believed-as-old stimuli received *lower* ratings than either the allegedly-new or standardly presented stimuli. Moreover, in further support of the mere exposure effect, previously-presented stimuli were liked at greater than baseline rates even when thought to be new.

Evidently, perceptual fluency does not necessarily require the direct re-presentation of an entire item. R. Reber, Winkielman and Schwarz (1998) presented a variety of stimuli which were preceded by subliminal primes consisting of highly degraded contours. In some cases, these contours were consistent with the target stimulus, and in other cases they were not. Participants were then asked to make both recognition and affective judgments. They found that when items were preceded by consistent primes, subjects were not only more rapid in their recognition responses, but they also tended to rate these items as being 'prettier'. Although reported as a departure from the mere exposure effect, it nonetheless supports the notion that the highly similar items may be effective in inducing increased affect. Consistent primes are, in part, a form of repeated exposure.

In a related experiment, R. Reber et. al. (1998) varied the amount of contrast of a stimulus set. Black (or white) circles were presented for one second each on backgrounds of varying shades of contrasting grey. Subjects reported higher liking scores for those sets with the greatest contrast. This contrast effect was affect-specific in that when asked to rate how 'ugly' each stimulus was, the high contrast items received the lowest scores. Overall, this study was viewed as support for the notion that the increased affect of neutral stimuli can be accomplished by means of facilitating perceptual fluency. Moreover, these results could be seen as further evidence for what Zajonc has been claiming all along, in that there is something primary and special about affective responses that distinguishes them from a host of other kinds of responses.

Misattribution. In a compelling demonstration of the effects of fluency, Jacoby, Kelley, Brown and Jasechko (1989) had participants read a list of names of nonfamous people. Half of the subjects in the experiment were given the opportunity to concentrate on this task, while the other half was distracted during the reading. The test required that subjects were later to identify which of the names on the list were of famous people. Whereas those concentrating on the initial task misattributed 19% of the names incorrectly as famous, the group of subjects which were distracted misattributed 27% of the names as belonging to famous people. One day after having read a list of names which were identified as nonfamous, subjects surprisingly assumed that the now-familiar names were actually those of truly famous people. In this case, the sense of familiarity was enough to elicit an assumption of probable fame which occurred in spite of the fact that the actual source was a list of names that were *explicitly* labeled as nonfamous. This effect is not present, however, when the subjects are questioned about the fame of the names shortly after having read the original list. In such instances, it is believed the source of the familiarity is still cogent and readily available, disavowing the tendency to hypothesize about the apparent fame of the familiar name.

Jacoby et. al. assert that the enhanced perceptual fluency of having read the names, especially in the absence of full attention, resulted in an erroneous assumption that the sense of familiarity experienced during the test phase was due to the *celebrity* of the people. Bornstein points out that subliminally presented stimuli provide no such opportunity for reassessment. "Just as the divided-attention subjects in Jacoby et al.'s (1989) study attributed fluency to fame based on contextual cues provided by the experimenters, so subjects in subliminal mere exposure studies attribute fluency to 'liking' based on the available contextual cues." (Bornstein, 1992, p. 199). In the absence of the opportunity to attribute familiarity to fluency, subjects fail to make the necessary adjustments in their evaluations, and come to consider the repeated, albeit subliminal, stimuli as very likeable, (aka 'familiar').

The perceptual fluency/attribution model also accounts nicely for the tendency of mere exposure effects to increase as a function of test delay. The longer period of time between exposure and test, the greater the effect. It might very well be the case that a substantial test delay disinclines subjects to attribute perceptual fluency to the original familiarization phase (Bornstein, 1992). As with subliminally presented stimuli, the 'source memory' obfuscates the identification of the familiarization event. This again more reliably mimics a naturalistic setting in which there is no definitive familiarization occasion.

Apparent Familiarity. In attempting to separate the concepts of perceptual fluency and apparent familiarity, Whittlesea (1993) defines 'familiarity' as the feeling one gets when attributions to past events occur unconsciously, resulting in the sense that we have encountered this stimulus before. Although these types of cues are generally accurate in accounting for our perceptions, they are not always reliable. As examples, he has examined the extent to which the perceptual fluency of items might be manipulated by means of impoverishing the perceptual processing of the stimulus via various masking techniques. He found that subjects had a tendency to believe that words which were more easily readable had appeared on a previously viewed list, even though this was not the case. This finding suggests that the ease of perceptual processing, and not the actual memory, is what contributes to this sense of recognition and familiarity. If familiarity could be illusory, so then the mere exposure effect might be elicited with *apparent* exposures.

Whittlesea (1993) tested this notion by presenting subjects with lists of target words which varied semantically. Words that met certain predictions regarding their occurrence in a sentence were thought of as not only having been seen before, but actually to be better liked than those words which were semantically incongruent with the rest of the sentence. However, this outcome could speak more of the predictive nature of our language than it does to the effects of familiarity as they relate to the mere exposure effect. Direct prior exposure to a specific test item in this case appears to contribute less to its overall affective rating than does its semantic sense. This basically tells us that more than simple exposure to an item is necessary when it already is laden with contextual meaning.

Specificity of Effects

The perceptual fluency/attribution model is not inconsistent with some of the findings that were offered as criticisms and "counter-examples" of the mere exposure effect. For example, it has been suggested that in addition to generating elevated affective responses, subjects also tend to endorse familiarized stimuli as 'brighter', 'darker', etc. (Mandler, Nakamura, Shebo, & Van Zandt, 1987).

Mandler et. al. explored the extent to which the effects of fluency would generalize to responses other than 'liking', arguing that prior exposure to a stimulus generates representations which may then be related to *any* relevant assessment. They tachistoscopically presented randomly constructed octagons for 2 msec. to four groups of subjects who were asked to make judgments of 'recognition', 'liking', 'brightness, or 'darkness'. Although 'recognition' was at chance levels, judgments of 'preference', 'brightness', and 'darkness' *all* proved to be significantly greater than chance. However, one might argue that recognition levels for stimuli presented at this exposure might constitute a sense of 'familiarity', which, in turn, may equate with '*n*'-ness. It may also be worth noting that the octagons were solid black on white backgrounds, presented via computer screen, possibly resulting in after-images–a concern which was not addressed. Moreover, the two groups of subjects which were required to assess either the 'brighter' or' darker' octagons might have been responding to the same stimulus feature, namely, 'contrast'. In fact, the mean percentages of 'brighter' or 'darker' previously-seen octagons were 60.0 and 60.1 respectively.

Extending the notion of nonspecific activation, a follow-up group of subjects was asked to designate which of the octagon pairs they '*disliked*' more. In this case, results were somewhere between 'recognition' and 'liking' levels, not differing significantly from either of these two groups. Mandler et. al. do suggest, however, that rather than comparing the results of the 'disliking' judgments (53.3%) to those of the 'liking' (61.7%), they should instead compare them to the complement of the 'liking' judgments, which would be 38.3%. However, in this comparison, it becomes clearer that the results of asking subjects which stimulus they actively 'dislike' is not the same as deriving the 'disliked' stimulus from the preference judgment. This asymmetry is also evidenced in the non-complementarity of the 'brighter'/'darker' judgments. Mandler et. al. take these results as counter-examples of Zajonc's claim for the primacy of specifically affective judgments in the absence of cognitive processing (as measured by recognition levels). In the strict sense of the word, preferences could not arguably be responsible for the endorsements made on other dimensions, since judgments for both 'brightness' and 'darkness' were elevated as a result of prior exposure. However, the notion of 'preference' might be tapping into an increase in 'fluency', to be translated as whatever semantic manifestation is rendered.

Seamon, McKenna and Binder (1998) revisit Mandler's assertion that previously viewed items would elicit higher ratings for whatever relevant dimension is being sought by investigating the patterns of responses which would occur in a variety of tasks. They presented octagons to subjects for durations of 8 msec.each, followed by bright field energy masks. After completion of the study trials, three groups of subjects were asked either to identify (2AFC) which of the octagons they 'recognized', or which they considered to be 'brighter' or 'darker'. The groups which provided 'brighter' or 'darker' judgments were then asked either to choose which item they 'liked' or 'disliked' in order to see whether preference judgments could occur in the absence of 'brighter'/'darker' judgments, as well as allowing for a comparison of the 'liking' and 'disliking' judgments.

With 8 msec. presentations, subjects were performing above chance in their recognition rates (57%). Judgments of neither 'brighter' nor 'darker' revealed an exposure effect. Moreover, the 'liking' and 'disliking' judgments reflected mutually opposite responses. Whereas previously presented items were 'liked' 58% of the time, these same items were chosen to be 'disliked' only 35%. To test the reliability of the 'disliking' judgments, another group was tested at exposure levels of 100 msec., and asked to provide their 'liking' or 'disliking' judgments immediately after the study phase (without the interim 'darker'/'brighter' task). Again, subjects 'disliked' previously-seen items only 32% of the time. These results fail to support Mandler's proposed notion of non-specific activation of the mere exposure effect at *both* the sub- and supraliminal levels. However, this study employed somewhat recognizable stimuli, a feature which was not consistent with Mandler's work. In repeating the study using briefer, 6 msec. durations, Seamon et. al. found that when recognition levels were reduced to chance, judgments of 'brighter' or 'darker' still did *not* favor previously presented items. However, those stimuli which were previously presented were 'liked' more than the novel items. The authors conclude that although the mere exposure effect has been demonstrated across a wide range of experimental conditions (see Bornstein, 1989), this effect is particular to affective judgments only, and *not* to a host of other seemingly applicable probe questions, as was reported by Mandler et. al.

The specificity of affective judgments has been further supported by the aforementioned work done by R. Reber, Winkielman and Schwarz (1998) in that they were able to find fluency effects only for affective responses. Additionally, Zajonc and Murphy (1993) revisited some of the research involving evaluative learning whereby they suboptimally primed targets with a number of different types of items. They found affective changes only when the primes involved emotionally-related items, such as 'happy' or 'angry' faces. No response shifts were noted when other types of nonaffective judgments, such as gender, size, and symmetry, were investigated.

Generalizability

Insofar as the mere exposure effect appears to be affect-specific, it also seems to behave in a way that suggests the increase in affect may to some degree generalize to related stimuli. If the connection between repeated exposure and increased affect escapes the subjects' overt interpretation, then the mere exposure effect remains essentially 'nonconscious' in its manipulation. Perhaps a question that needs to be asked is whether relatively sophisticated subjects are immune to the mere exposure effect, in much the same way that naive subjects fail to make the connection between repeated exposures and increased liking.

Could the actual repetition of a stimulus generate a heightened appreciation for not only that specific item, but for similar items as well? Zajonc first speculated about this in his original 1968 paper, and has recently explored this issue systematically. Monahan, Murphy and Zajonc (2000) suspected that there might be something about the experience of having been repeatedly presented with something that may simply produce a generally heightened tonic mood. This 'better mood' might then be reflected in elevated liking judgments for whatever comes along.

They looked first at the differences in heightened affect resulting from the repeated presentation of stimuli versus the single presentation of a series of somewhat related items. Twenty-five Chinese characters were flashed at suboptimal exposures for two groups of subjects. For one group, the display consisted of 25 different characters (single-exposure), and for the other group of subjects, there were five presentations of five different characters (repeated-exposure). Participants were simply asked to report on their overall mood at the end of the presentations. What they found was that subjects in the repeated-exposure condition reliably reported to be in a better mood than those participating in the single-exposure group. Although this may arguably be viewed as a confound in the mere exposure effect, Monahan, Murphy and Zajonc instead concluded that this provides evidence that the very act of being repeatedly exposed to some 'thing' results somehow in a positive effect on one's temperament. It is this elevation in tonic mood that might then be contributing to a person's overall affective disposition. Naturally, being asked to provide liking judgments while in this 'better mood' would reflect this increase in positive attitude.

Taking this further, Monahan, Murphy and Zajonc then looked to see the extent to which this induced elevated mood might translate into greater liking not only for those items which were specifically presented, but also to new items which share conspicuous similarities to the ones presented. In this study, four different groups of subjects were suboptimally presented with a series of stimuli consisting of either

Chinese characters or arbitrary polygons. These were presented as either a series of 25 single-exposures or as groups of five items exposed five times each. Participants were then asked to rate how much they liked 15 test items which included five items from the exposure phase, five novel-yet-similar items (belonging to the same category, i.e., Chinese characters or polygons), and five novel-and-different items from the alternate stimulus group. In other words, for example, subjects exposed to Chinese characters would then be asked to evaluate five of those characters, five new Chinese characters, and five polygons.

The results showed that those subjects in the repeated-exposure condition gave overall higher liking ratings to *all three* stimulus categories than did those participating in the single-exposure groups. This result reinforces the assertion that the experience of having been repeatedly exposed to something (e.g., five times vs. one time) somehow elevates one's tonic mood in a manner that generalizes even to items which have not been seen before.

What is also interesting, (and which foreshadows the empirical work that will be reported in this paper) is that not only is there clear evidence for a 'generalized' mere exposure effect, but the relatively higher liking scores that were given to the novel-yet-similar items were virtually identical to those given to the items that had actually been seen before. In other words, the heightened affect was applied equally to all members of the same category, regardless of whether they had actually been previously presented. This was found to be the case in both the single-exposure and repeated-exposure conditions. Of course, in this experiment, those items which were considered 'novel-and-different' were actually *quite* different from those presented to the subjects during the exposure phase – belonging to the alternate set of either Chinese characters or polygons – and resulting in 1/3 of the test items conspicuously belonging to a totally different category from the rest. The relatively higher liking scores for those 2/3 might then be an expression of frequency sensitivity. Likewise, subjects might be cuing in on identifying consistent group membership, which might then be expressed as increased familiarity and affect. These results do not necessarily contradict the perceptual fluency explanation of the mere exposure effect in that the ease of processing which is a byproduct of prior exposure could apply to highly similar members of the same category.

If the mere exposure effect were solely due to an increase in elevated affect due to the experience of having been repeatedly exposed to an 'some item', then earlier studies using the 2AFC test situations would not have necessarily revealed any privileged preferences for before-seen items–but they do. Insofar as the literature does provide ample support for this finding (see Bornstein,1989), the original presupposition regarding the specific effects of familiarized items still has merit. However, what we do see here is that items which are 'similar enough' appear also to benefit from prior exposures.

An interesting finding in this study was that those subjects in the repeated-exposure condition rated even the very-different items higher than did the control group. This was not the case for subjects participating in the single exposure condition. Monahan et. al. take this as an indication of the generalizability of the heightened tonic mood induced by repeated exposures, and challenges how the perceptual fluency model might account for this finding. The authors also point out that this notion of generalized tonic disposition actually makes good evolutionary sense. In terms of survival, young animals tend to make positive adaptations to their environments, even new environments, which is a useful survival strategy. This positive emotional response to their surroundings may be the result of an overflow of parental imprinting. Of course this is quite speculative, but it does fit in nicely with has been observed in many species.

Neurophysiological Support

These results reinforce the notion of separate, and possibly independent, cognitive and affective processing systems. Neurophysiological evidence to support this supposition includes the research done with patients with Korsakoff Syndrome. These patients exhibit retrograde amnesia for the recent past and severe anterograde amnesia for new information due to damage sustained to the frontal lobes and thalamus (sometimes including the medial temporal lobes, the hypothalamus and the amygdala). Not only do these patients display severe cognitive impairments in the form of explicit memory deficits, but they also demonstrate flat emotional tone. In the absence of being able to recognize melodies which were presented to them, Korsakoff sufferers have nonetheless demonstrated the mere exposure effect by preferring those tunes which were previously presented (Johnson, Kim, & Risse 1985). These findings demonstrate not only a separable emotional/cognitive processing system (as suggested by Zajonc), but also supports the suggestion that the fluency of processing, even in the absence of explicit identification, serves to drive affective behavior. Morever, these findings are consistent with conceptualizing the mere exposure effect as a demonstration of implicit memory (Schacter, Delaney, & Cooper, 1990; Seamon, McKenna & Binder, 1998), as well as providing support for the separability of the implicit and explicit systems.

Murphy and Zajonc (1993) cite several studies which support separable emotional and cognitive systems. For example, patients who suffer from prosopagnosia are able to identify emotional expressions in faces which they nonetheless are unable to recognize (e.g. Ellis, 1986). In addition, there also exist patients with the reverse syndrome, prosopo-affective agnosia, who are able to identify faces without being able to read their emotional content (e.g., Kurucz, Feldmar, & Werner, 1979). These patients provide further neurological support for the existence of somewhat separate (although not necessarily independent) emotional and cognitive processing systems. The amygdala appears to play a role in the network of emotional processing, as is seen with amygdalectomies, performed to control rage; or with Urban-Wiethe disease, an inherited syndrome resulting in damage to the amygdala, resulting in selective memory impairment for emotionally laden stimuli. The role of the amygdala has been invoked by Murphy and Zajonc as evidence for the separable processing systems, (although the primacy of affect is not established).

More recently, Elliot and Dolan (1998) used positron emission topography (PET) scans while replicating the Kunst-Wilson and Zajonc 1980 classic experiment. Brain activity was recorded to identify the regions active in both recognition and preference judgments. What they found was that while left frontopolar and parietal activation was found for recognition judgments, preference judgments appeared to involve right lateral frontal activation. Subjects in this study were unable to explicitly distinguish old from new stimuli, and did not report any sense of subjective familiarity. Whether this is a counter-indication of the perceptual fluency model is yet to be concluded.

The Mere Exposure Effect: Summary

To summarize, the theories accounting for the mere exposure effect can be divided into three camps:

(a) The notion of the primacy of affect defended by Zajonc (1968, 1984, 2001) argues that a system of affective processing exists prior to and independent of cognitive processing. The claim here is that this system is evolutionarily old and naturally predates and precedes cognitive operations in that it has ecological survival value and may exist in animals as well as humans. This position can accommodate a generalized rise in tonic disposition, which is also congruent with an evolutionary perspective. Moreover, Zajonc has offered a model of classical conditioning to account for this phenomenon.

(b) The perceptual fluency/attribution model proposed independently by both Whittlesea (1994) and Bornstein (1989) and investigated by others (e.g. Bonanno & Stillings,1986; R.Reber, Winkielman & Schwarz, 1998) accounts for one's readiness to *like* a previously encountered stimulus as a function of its acquired familiarity. Such familiarity is assumed to facilitate reactivation of the perceptual process, rendering those familiarized items as more likeable.

(c) A non-specific activation theory (Mandler et. al, 1987) that differs from the two above essentially by maintaining that merely exposed items elicit heightened *n*-ness, whatever that parameter might be, providing it adheres to some constraints of relevancy. It is their contention that preferences are but one of a number of possible expressions of this hyper-salience.

What these theories do share, however, is that they acknowledge a non-conscious, or non-intentional retrieval of memorial material, even in, (or, as in the case of the perceptual fluency/attribution model, *especially* in), the absence of overt recognition.

The Overlap: The Mere Exposure Effect and Implicit Learning

Recently, Dienes and Berry (1997) noted that many of the issues surrounding the characterizations of implicit and explicit knowledge share an interesting feature with those regarding threshold issues in subliminal perception. Both areas of inquiry are concerned with the identification of accessible knowledge acquired outside of conscious awareness. As discussed earlier, there are concerns regarding whether implicitly acquired information tends to be contaminated by an explicit component; just as there are questions concerning the conscious influences on nonconscious processing. A possible solution to this empirical impasse might be to characterize as *'implicit/subliminal'* that knowledge which had been processed outside of intentional conscious awareness, and which nevertheless resulted in significant task performance.

Another way in which the implicit/explicit controversy relates to the questions raised by the mere exposure effect has to do with what the person believes to know about the stimuli. In the typical subliminal mere exposure experiment, subjects provide a direct indication of stimulus awareness through simple recognition tasks. In the AG studies, however, an assessment of this metaknowledge may be indirectly gauged through confidence levels. Although the literature is replete with a full range of results regarding the relationship between confidence ratings and accuracy, Chan (1992) found that subjects' confidence levels appear to have virtually *no correlation* with the accuracy of their grammatical judgments [zero correlation criterion]. In other words, subjects did not appear to know when it was they were correct in classifying new test strings. However, Chan used nonverbalizable elements (irregular geometric shapes) to instantiate the AG strings, again suggesting that the nature of the stimuli might play an important role. Others (e.g., Reber, 1989; Manza & Skypala, 1996), using familiar, pronounceable stimuli such as letters of the English alphabet, have found that confidence ratings were indeed highest when making accurate classifications. However, subjects were still performing above chance even when claiming to be 'just guessing' [the 'guessing' criterion] (Dienes, Altmann, Kwan, & Goode, 1995; Dienes & Berry, 1997).

These later findings are reminiscent of the ongoing discussion of subject awareness regarding the issues of subliminal perception, and fit in nicely with the

Cheesman and Merikle's (1986) definition of lack of metaknowledge. In both paradigms, subjects are generally not capable of articulating (pronouncing, accessing, etc.) all of the information which they have about the stimulus. This inability to fully access all available material at one point in time is fundamentally what Erdelyi (1985) describes as his dissociation paradigm. In the case of a typical AG study, participants are permitted to sample a string sufficiently to even memorize it, yet the actual *structure* of the string is not explicitly revealed. This structure, which is the functional stimulus, instead needs to be extrapolated. Acting on this abstracted knowledge in the absence of its overt acknowledgment is, in a sense, working below Cheesman and Merickle's notion of a psychologically real subjective threshold.

Retrofitting

Whittlesea and Wright (1997) argue that implicit learning is fundamentally similar to explicit learning in that numerous factors need to be taken into account which interact and contribute to the overall learning experience. These factors include the structural affordances of the stimuli, as well as perceptual and personal intentions and experiences of the subjects. "Until the subject encounters the stimuli and processes them in some way for some specific purpose, the effective structure of the set is not real but is only potential, and it has many potential states, some of which are catastrophically different" (Whittlesea & Wright, 1997; p. 183). In this way, implicit learning is viewed to be just like "ordinary" learning only without becoming aware of the implications for performing unanticipated activities. Much of the processing is actually *potential processing*, to be determined at some later point when the task then becomes defined (Whittlesea & Dworkin, 1993; Shanks & St. John 1994; Neal & Hesketh, 1997). This conceptualization suggests that 1) a variety of tasks can be later performed on previously acquired information which are sometimes differentially sensitive to the learning environment (as in the structural mere exposure effect), and 2) there may be a dissociation of different types of task responses to a given stimulus.

This formulization of implicitly acquired representations is consistent with the dissociative results found in many subliminal mere exposure effect studies. If the evaluations are required *after* the initial processing, then the judgments could later be applied separately and with different results. Asking subjects to make grammatical judgments about structured stimuli would not necessarily result in the same pattern of affective judgments tapped for the same stimulus set. In this way, the processing is determined by the specific task demand, and constrained by certain affordances of the stimuli, as Whittlesea and Wright describe:

> "....implicit learning consists of encoding whatever organization of a stimulus one experiences in encountering that stimulus in some context and for some purpose, and the application of implicitly acquired knowledge to novel stimulus structures depends as much on the organization that the subject imposes on the stimuli, given their purpose at that time, as it does on the objective structure of those stimuli." (Whittlsea & Wright, 1997, p.195).

On the other hand, perhaps affective valences are encoded immediately upon perception--whether they pertain to the deep or surface structure of the stimulus--whereas grammatical assessments are a process which is then later applied only when a subject is required to do so. This kind of retrofitting characterization of implicit learning is somewhat in keeping with Neal and Hesketh's (1997) position that implicit processing actually describes the potential application of acquired information to later be strategically summoned as the need arises.

Structural Mere Exposure Effect

In 1983, Gordon and Holyoak first noted and experimentally developed the striking similarities between the ways in which the mere exposure effect and IL are conceptualized and studied. In both paradigms, subjects are presented with materials about which they are subsequently asked to make evaluations about items which are related to an exposure phase. Implicit learning is characterized as a type of learning that occurs naturally when people are presented with structured stimulus displays. Likewise, the mere exposure effect reflects an affective appreciation of stimuli which emerges spontaneously as a function of repeated exposure. Prior to Gordon and Holyoak's study, the mere exposure effect had been investigated exclusively using individual items that were actually presented to the subjects, and about which affective ratings were made. For Zajonc, the inability of subjects to articulate their recognition of the briefly presented stimuli suggests a lack of cognitive processing. However, as Gordon and Holyoak rightly point out, in light of the research on subliminality that has emerged in the literature since, it is unreasonable to assume that just because stimulus recognition was not achieved at a "conscious" level, that no cognitive processing occurred (e.g., Jacoby 1991; Fowler, Wolford, Slade & Tassinary, 1981; Allen & Reber, 1998).

Gordon and Holyoak maintain Zajonc's position of the primacy and independence of affect by conservatively suggesting that although both grammatical and affective judgments are effective in discriminating stimuli, it is not necessarily the case that these are the *same* processes. Disparate results would provide further support of the notion of separate processing systems, or, at the very least, of different treatments of the once-processed stimuli. Recognizing the similarities between the mere exposure effect and IL, Gordon and Holyoak linked these two fields by investigating the extent to which the information abstracted in AG studies could generalize to produce what has later come to be called a "structural mere exposure effect" (Manza, 1998).

The confluence of these two areas presents an opportunity to examine a number of issues. First, the AG learning design provides a means of gauging the role which cognitive processing may play in the formation of affective judgments. Zajonc (1980) argued for two separate processing systems, maintaining that it was the affective, and not the cognitive, system which is highly sensitive and reacts more quickly. If it is possible for the mere exposure effect to generalize to structured stimuli, would both tasks produce identical patterns of results, or would one be a more sensitive indicator of the learned AG? However, the first step would be to see whether the mere exposure effect can be *generalized* to stimuli which have not been specifically presented before. Previous studies investigating the mere exposure effect have commissioned the testing of discrete items whose prior level of exposure (either frequency or subliminality) have been systematically manipulated. Gordon and Holyoak instead used strings of letters which instantiate a structured grammar that has been used in AG studies. These test strings were never-before-seen items which varied only in terms of grammatical correctness. Now instead of looking at whether previously seen items elicit elevated affect, the focus is on whether exposure to the familiarized *structure* of the strings will produce the mere exposure effect to novel-yet-structurally coherent strings.

In a key experiment, Gordon and Holyoak asked a group of subjects simply to view a series of learning strings which were presented one at a time. Before disclosing the existence of a grammar, subjects were then asked to view a series of new strings (which consisted of both grammatical and nongrammatical items), and to provide a 'liking' score on a scale of 1 to 7. Afterward, they were informed of the existence of the grammar and asked to classify these same test items in terms of their grammaticality, as well as providing confidence ratings. A control group which did not experience any prior exposure to the learning strings gave assurance that there was no potential inherent 'likeability' of grammatical strings.

Grammaticality and confidence ratings were converted to a 6-point scale so that a response of 'grammatical' and 'highly confident' was registered as a '6', and 'non-grammatical' and 'highly confident' became a '1'. As is routinely found in these experiments, the well-formedness judgments demonstrated the subjects' ability to discriminate between the grammatical and nongrammatical strings. The landmark finding, though, was that the 'liking' ratings were significantly higher for the grammatical than for the nongrammatical test strings. This was the first successful demonstration of the ability of subjects to generalize the mere exposure effect to structured stimuli that had not actually been seen before.

Although these results might appear to strengthen Zajonc's assertion of separate processing systems, it is possible that both the affective and cognitive operations rely on a single processing system. These results could suggest that the subjects did have a sense of the grammar, which in turn, could have contributed to their 'liking' ratings. This view would necessarily implicate cognitive applications as a prerequisite to affective judgments. Moreover, as Gordon and Holyoak point out, Zajonc's argument "...is dependent on the identification of cognition with conscious recognition, a link that does not seem justified." (p. 498).

As discussed earlier, much of the criticism which has been waged against the interpretation of AG studies has to do with the extent to which explicit information is applied during the grammaticality judgments. Since subjects necessarily need to be informed about the existence of the grammar, there is concern as to how we then go about either re-visiting previously acquired knowledge (Neal & Hesketh, 1997), or strategically applying explicit knowledge derived from the prior learning experience (e.g., Vokey & Brooks, 1992; Perruchet & Pacteau 1990). However, since subjects need not be informed of the grammar when making their affective judgments, the generalizability of the mere exposure effect might provide a better opportunity with which to examine the purely implicit nature of acquired knowledge (Manza & Bornstein, 1995) Since the grammar need never be mentioned, there is no reason to suspect that subjects are either attempting to impose a structure on the stimuli post hoc, or are necessarily reacting to any specifically remembered fragments. In this way, the link between the mere exposure effect and the classic AG paradigm provides a simple means with which to satisfy one of the criteria set by critics of implicit learning in that no explicit reference to the initial study phase of the experiment need be made (Schacter & Graf, 1986).

This structural mere exposure effect could be viewed as an indirect measure of information about the existence of the grammar which, should it prove to be more sensitive than the direct, well-formedness classification task, could be considered to be an indication of nonconscious processing (Manza & Bornstein, 1995; Merikle & Reingold, 1991). In this way the 'liking' task may be a *purer* indicator of implicitly held knowledge because it is performed without revealing the existence of the grammar, and it does not invite any reflection on the learning strings. Because the mere exposure paradigm does not require revealing the existence of the structured grammar, it becomes an appropriate candidate with which to examine more than one potential processing mode. This circumvents the possibility that subjects will be deliberately trying to remember or activate certain information acquired during the initial learning phase of the experiment.

Manza and Bornstein (1995) worked with two groups of subjects who were to participate in either the 'rules' or 'liking' tasks. All subjects were asked to attend to and attempt to memorize a series of strings which were actually grammatical exemplars. Subsequently, they were all presented with 64 new strings, half of which were grammatical. At the conclusion of their respective tasks, all subjects were told (or reminded) about the existence of the grammar, and were given a 'fill-in-the-blank' task, wherein they were required to complete 41 stem-completion fragments. Both groups of subjects scored equally above chance on this latter task, although those who participated in the well-formedness task outperformed those in the liking task.

Manza and Bornstein found that subjects were not only able to discriminate grammaticality via well-formedness responses, but that this knowledge was displayed when asked to provide liking ratings. Although the liking ratings failed to provide more significant discriminations than did the well-formedness judgments, they nonetheless retain a certain purity in that they were made in the absence of any overt reference to the variable being probed (i.e., grammaticality).

Whittlesea and Wright (1997) stretched the parameters of the previous two studies on the structural mere exposure effect by choosing to use test items which differed on a number of dimensions from those used during the learning phase. Their test items instantiated the grammar using either a completely different letter set from the learning strings, or, more radically, using color patches instead of letters. Moreover, they chose not to use a finite-state Markovian grammar, but instead used deterministic sequences which they felt was under greater experimental control. Their 'nongrammatical', (or 'illegal'), test strings were generated from an entirely different sequence than was used for the learning strings (and were, in fact, internally consistent in that all 'illegal strings' followed the 'illegal sequence').

Three groups of subjects were asked to make both 'pleasantness' and 'grammaticality' judgments on a series of test strings. One group saw test items instantiated with a different letter set, another was presented with color patches, and the third group saw color patches and was told of the potential relationship between the colors patches and the learning strings. Subjects were presented with the learning sequences and asked to recite and memorize what they saw. However, the recitations of these highly regular (and therefore rhythmic) strings may have further enhanced the differences between the 'grammars' used to produce what were considered 'legal' and 'illegal' test items. The 'legal' and 'illegal' sets each conformed to separate permissible patterns. The recitation of one of these patterns would inevitably enhance its (legal) salience over the other.

Subjects in all three groups were able to make successful well-formedness discriminations. However, only the group of subjects who gave 'pleasantness' ratings to the test strings which were instantiated with the new letter set successfully rated the 'legal' grammar as being more 'pleasant'. Whittlesea and Wright interpret these results as a failure to transfer affective evaluations to the color patches. However, it may be the case that a series of color patches differ so much qualitatively from strings of letters (which invite a linguistic–and holistic–assessment) that it might be inappropriate to assume that the mere exposure effect would tap into the underlying structure rather than the colorful pattern which emerges. In this experiment, the manipulations are so striking that conclusions about what the subjects were or were not doing seem premature. Moreover, Whittlesea and Wright ought not to de-emphasize the fact that subjects *did react* to the deeper structure of the sequences (as they represented different grammars) in that their grammaticality judgments demonstrated this discrimination in all three test conditions. It is also noteworthy that the structural mere exposure effect was successfully elicited when test items were instantiated with letters–even new letters, but not with colors. A split/half analysis of these tests would also be of interest especially since a second, internally consistent grammar was used to serve as the 'illegal' set. What might be concluded from this study is that perhaps the structural mere exposure effect may only be realized when the grammatical stimulus array invites an affective evaluation of the *structure*, and not of the *individual items*.

Implications and Questions

As described above, there is compelling overlap in the ways these two fields have characterized nonconscious learning, as well as in the techniques often used to explore these phenomena. This series of studies examines the ways in which the mere exposure effect could be generalized to the structured stimuli typically used in AG studies. Not only would test strings be evaluated in terms of grammaticality (well-formedness), but they would also be rated in terms of affect (how much they are liked). Exploring the parameters of the structural mere exposure effect would provide an opportunity to investigate the extent to which structural regularities could generalize sufficiently to elicit elevated affect to stimuli which have never actually been seen before. Likewise, the presence of heightened affect can be viewed as a more indirect indicator of accessing implicitly acquired information. The following experiments incorporate these modifications, specifically:

a) the inclusion of 'learning' strings as part of the test stimuli

b) the use of unusual stimulus elements with which to instantiate the grammar

c) using a within-subject design

a) Including learning strings in the test phase. Insofar as these three studies (Gordon & Holyoak, 1984; Manza & Bornstein, 1995; Whittlsea & Wright, 1997) were the first attempts at investigating the possible superimposition of the mere exposure effect within an AG study, it is surprising that no instances of the original learning set were included as any of the test items.

Although Gordon and Holyoak (1983) and Manza and Bornstein (1995) were successful in demonstrating higher liking judgments for well-formed strings, their aim was to examine the possible ways in which the (structural) mere exposure effect could be exploited to investigate implicit learning. They therefore tested subjects using *only* novel strings. However, the shift in emphasis in the current study focuses on the extent to which the mere exposure effect could be generalized to structured stimuli. In this way, both the well-formedness and liking tasks become of interest. For this reason, it becomes especially interesting to look specifically at the effect of using actual strings that were presented during the learning phase. These strings would not only be grammatical (by definition), but would also have the added advantage of having been previously seen by the participants. This would be more in line with the classic mere exposure effect which predicts that the learning strings should result in the highest liking scores in that they have actually been presented. However, it is possible that highly structured stimuli are processed in such a way that the item-specific identity of the learning strings is not as important as the underlying grammar. If this is the case, then the learning strings would not contribute any additional affective advantage over novel-yet-structurally coherent items.

This issue is reminiscent of concerns regarding the nature of the information acquired during an AG study. Is it the abstract representation of the grammar that contributes to the well-formedness judgments (Reber, 1993), or is it rather the actual physical instantiation of bi-grams, tri-grams, or highly cogent chunks that are repeated during the testing phase that allow for successful grammaticality judgments (e.g. Brooks & Vokey, 1991; Perruchet & Pacteau,1990)? This inclusion is therefore not only interesting in terms of the structural mere exposure effect, but it also speaks to relevant issues within the implicit learning paradigm itself.

If it is the case that, as suggested in studies of transfer of implicitly acquired knowledge, subjects can develop sufficiently abstract representations to enable them to make well-formedness judgments about items instantiated using a totally different stimulus set from that used during the learning phase (see Manza & Reber, 1997 for discussion of the circumstances under which such abstract representations tend to emerge), then a finding of equivalent liking scores for both old and new grammatical items would not be entirely unexpected or inexplicable.

In a related note, by exploring the extent to which the structural mere exposure effect (SMERE) can be generalized, these studies examine the issue of 'familiarity'. Is heightened affect towards grammatically coherent strings potentially transferable to a completely new stimulus set using different letters? Just as the inclusion of learning strings permits the examination of physically identical stimuli vs. newly constructed grammatical stimuli, transfer and generalizability could be assessed more directly by substituting the elements used to instantiate the strings. The inclusion of 'learning strings' takes on a different dimension when this time referring to the sequenced pattern, and not the actual letters.

These questions have no obvious answers, and so a systematic inquiry seems appropriate at this stage.

b) The use of 'unusual' elements. How important are the actual symbols used to instantiate these displays? Must there be a degree of *a priori* familiarity in order to produce the generalized mere exposure effect? Do the effects emerge only when items are instantiated with symbols already well known to the subjects--such as the English consonants used in the standard AG learning study--or can stimuli unlikely to be within the subjects' routine experiences also serve successfully in these studies?

On one hand it would appear that any arbitrary symbol set could be used, at least that is the indication from the traditional AG learning literature and studies of the classic mere exposure effect. On the other hand, there are reasons for suspecting that highly unfamiliar forms might pose problems. Despite the variety of stimulus forms that have been used, prototypical AG learning experiments use letters of the English alphabet to make up the stimulus strings, and as several researchers have shown, these stimuli can be easily integrated into chunks or Gestalt-like wholes (Brooks, 1978; Perruchet & Pacteau, 1990; Servan-Schreiber & Anderson, 1990).

Many of the early mere exposure studies involved arbitrarily constructed octagons or deliberately meaningless 'foreign words, such as 'Turkish' syllable and 'Chinese' ideograms (see Bornstein, 1989 for a review). It would be interesting to see whether there would be limitations to the structural mere exposure effect when *very* *unfamiliar* items are used. It is unclear whether replacing letters with unfamiliar (and unpronounceable) symbols would successfully produce the structural mere exposure effect, or whether the physical properties of the individual items might compete rather than play in concert with the grammatical features.

Although implicit learning studies involving AGs have typically used the highly familiar letters of the English alphabet to instantiate the grammars, Chan (1992) demonstrated that highly irregular and unpronounceable stimulus items result in successful grammatical classifications. This issue is relevant because the pronouncability of stimuli could result in covert rehearsal, possibly resulting in a more explicit memory for particular or similar items. Manza and Reber (1997) found robust learning when the stimuli were composed of sequences of tones of varying pitch, and Altmann, Dienes and Goode (1995) report successful learning with AGs instantiated as sequences of graphic symbols, musical notes, and nonsense syllables. In terms of well-formedness judgments, therefore, no problems are anticipated.

Whittlesea and Wright (1997), too, were interested in this aspect of the phenomenon and selectively altered sequenced testing stimuli by using either a different letter set, or, more radically, color patches. He found that the transfer to color patches disrupted the generalization of the mere exposure effect, whereas the direct substitution of letters did not. One might have expected that such an aesthetically laden stimulus might be a most appropriate trigger with which to trigger affective evaluations. But this was not so. These stimulus arrays were not generated from a complex grammar, but rather assembled as repetitive sequences. Color patches are not totally unpronounceable in that although they do not necessarily invite linguistic (and wholistic) analyses, and their names could be subvocalized. However, even under these circumstances where the letter-to-color transfer appears to be so salient, a generalized version of the mere exposure effect was not seen.

In this regard, although unusual items should predictably not pose any difficulty in terms of well-formedness judgments in this blended paradigm, they may present difficulties in terms of exposure effects. It is conceivable that the composition of the strings could impose different demands for the different tasks. For example, one of the explanatory models of the mere exposure effect suggests that it is the perceptual fluency of the before-seen items that result in their heightened affect (Bornstein & D'Agostino, 1994). It may be that unusual items would not be processed with as great an ease as highly familiar ones. In this way, the *a priori* familiarity of the individual stimulus elements might facilitate the way the entire stimulus *string* is assessed. It is, after all, the embedded grammar, and not the items themselves, which are the actual stimuli.

c) ***Within-Subjects Design.*** Manza and Bornstein (1995) used a between-subjects design where participants made judgments of either well-formedness or liking, but not both. Gordon and Holyoak (1983) ran all of their subjects through both procedures but always began with the liking task first. The concern in both cases was that knowledge of the existence of the grammar might have an influence on the subsequent liking ratings.

The current study uses a within-subjects design with the stimuli re-randomized for whichever task is second. This design provides the opportunity to examine whether knowledge of the existence of the grammar has an influence on affective ratings. It also provides a basis of comparison in that individual subjects' responses could be examined in terms of both their well-formedness accuracy as well as affective evaluations. It is possible that the structural mere exposure effect could emerge in the absence of apparent implicit learning about the grammar—and vice versa. In principle, subjects could acquire sufficient knowledge about the AG to make grammaticality judgments but still show no particular affection for the well-formed strings. And, on the reverse side of this coin, they might develop distinct preferences for particular stimuli but be unable to utilize these affect-based representations to make grammaticality decisions. The latter outcome would be somewhat surprising given the robustness of the standard implicit learning effect but the former possibility remains very real.

An intriguing issue raised by Manza and Bornstein (1995) is the possibility that the two measures of implicitly held knowledge are differentially sensitive. They suggested that the liking task might well be a more indirect and therefore better measure of implicitly held knowledge because it is less likely to be contaminated by consciously modulated strategies. However, their design did not permit a direct test of this hypothesis. By having all subjects make both kinds of judgments about all stimuli, direct comparisons between the two dependent variables can be made.

In addition to issues surrounding different response patterns to the different tasks, (some of) the experiments in the current study use highly unusual elements to construct the strings. Might the relative unfamiliarity of these items detract from one, but not the other of the two tasks? Are these two tasks correlated within the same subject, or do individual differences on one predict the results of the other?

From an evolutionary perspective, Zajonc has repeatedly argued not only for the separation of cognitive and emotional systems, but also for the primacy of affect (e.g., 1980, 1993, 2001). Testing for both kinds of responses from each subject is a way to address this assertion. It may be possible that implicit acquisition in general is a phylogenetically primitive strategy. Are these skills always necessarily linked?

Designed as a replication and extension of the Gordon and Holyoak (1983) and Manza and Bornstein (1995) studies, the following experiments, although varying on particular parameters where specified, retain certain methodological consistencies directed at the questions just raised. To summarize, by modifying critical components of the procedure, the plan is to look more closely at the questions entailed by earlier findings: (a) Will test stimuli which are physically identical to strings encountered during the acquisition phase be treated differently than new, structurally coherent strings? (b) Will the structural mere exposure effect emerge when unusual, unfamiliar stimuli, ones that are not easily encoded or pronounced, are used to instantiate the grammar? and finally, (c) Do subjects participating in both tasks demonstrate their knowledge of the grammar in a consistent manner?

The Experiments

Normative liking data. To ensure that there were no adventitious preferences for particular stimulus strings, normative liking data were collected on all stimuli used in these experiments. A group of 21 subjects was presented with the full set of the 36 randomly mixed test strings, both grammatical and nongrammatical, and asked to rate them using the six-point liking scale described below. Where appropriate in the experiments that follow, the liking data from experimental subjects are compared both with chance performance as well as with the data produced by participants in this control group. The only exception is in Experiment 5, where a specific sub-population was recruited for reasons which will be come clear.

It is noted here that in none of the six experiments were any systematic preferences for any stimulus type observed. Participants who had no prior experience with well-formed strings show no independent preference for grammatical strings over nongrammatical ones. These normative data from all stimulus types used in the six experiments reported in this paper are presented in Table 1.

General Procedures

Artificial Grammar Paradigm

The basic artificial grammar paradigm mentioned above provides the model for all of these experiments. More specifically, five elements are used to instantiate the string sequences which serve as the stimuli throughout the present study. The specific set of elements used in each of the experiments varies and is discussed further when appropriate. This particular AG, which has been used in a variety of studies of implicit learning (see, e.g., Reber, 1993) generates 43 grammatical strings of lengths three through eight (See Figure 1). Of these, 20 representative items are selected to serve as the learning set. For the testing phase of the experiment, a total of 36 strings are used. These included 12 borrowed from the learning strings, designated as 'Old Grammatical' (OldG) strings. An additional12 grammatical strings which were not used in the learning phase serve as the 'New Grammatical' (NewG) strings, and 12 strings containing a violation of the grammar in one position, become the 'Nongrammatical' (NonG) strings. All 36 items were presented in a random fashion, and are presented in Appendix A.

Table 1

Normative liking ratings (6-pt scale)

Control subjects who had no prior learning experience (N = 21)

Symbols used	Mean on 'G' Items	Mean on 'NG' items	**Overall Mean**
Letter Set PSTVX	3.37	3.33	**3.35**
Letter Set CHJLN	3.22	3.34	**3.28**
Japanese Kana Symbols	3.38	3.33	**3.36**
Chinese Ideographs	3.14	3.18	**3.16**
Keyboard symbols	3.48	3.62	**3.55**
Russian	3.45	3.57	**3.51**
Native Russian Readers -- Experiment 5 (N = 14)			
English (PSTVX)	3.32	3.14	**3.23**
Cyrillic	3.06	2.83	**2.95**

The Design

The experiments reported in this study, while varying in specific ways as indicated, maintain certain consistencies which are employed in all six of the experiments. This design entails two phases, a 'learning phase' during which subjects work with grammatical strings, and a two-part 'testing phase' during which they were asked to make both well-formedness and liking judgments regarding the 36 grammatical and nongrammatical test items. Participants were run in small groups ranging from 3 to 15 subjects each. The tasks were counter-balanced so that subjects in approximately half of the groups made all the well-formedness judgments first, followed by the liking task. The participants in the other group made liking judgments before making well-formedness judgments. This procedure is a departure from that used by Gordon and Holyoak (1983) and Manza and Bornstein (1995). Manza and Bornstein used a between-subjects design where participants rated test strings according to *either* well-formedness or liking, but not both. Gordon and Holyoak ran all subjects through both procedures but always with the liking task first. The design decisions made by these earlier researchers were made because of concerns regarding the overt knowledge of the existence of an underlying set of rules having an effect on the manner in which liking judgments are made. By running all subjects in both tasks, this possibility can be examined more methodically. Should it prove to be the case that knowledge of the existence of the grammar does have a systematic impact on how the liking judgments are made, it then becomes possible to explore just what that influence might be. Moreover, it does not preclude the possibility of then separating 'contaminated' data provided by those subjects who were no longer naive about the existence of the grammar.

Part 1 - Learning Phase

The learning phase of all the experiments in this study consisted of a simple copying procedure in which each learning string was presented to the subjects for approximately 10 seconds, during (or after) which subjects were expected to copy them as quickly and accurately as possible. The full set of 20 strings was presented three times with order quasi-randomized. As is typical in these implicit learning experiments, subjects were told only that they were participating in a simple learning and memory experiment.

Part 2 - Testing Phase

Testing. All participants made both well-formedness and affective judgments to all test items. Groups of subjects were randomly assigned to one task order or the other.

Well-formedness judgments. Participants were informed (for the first time) about the existence (but not the nature) of the rules used to generate the learning strings. They were then presented with each of the 36 test strings (12 OldG, 12 NewG and 12 NonG), one at a time in a randomized sequence, and asked to make well-formedness judgments about each one using whatever information they may have gathered from the initial learning phase. Participants were asked to judge each stimulus string in terms of whether they thought it could have been generated by the rules which produced the strings encountered during the learning phase, or whether it violated those rules and was therefore nongrammatical. Booklets containing printed test items were distributed. Participants were to circle either 'YES' ("It is grammatical") or 'NO' ("It is not grammatical") printed beneath each of the 36 test items, as seen in Appendix B. They were also asked to provide a measure of their confidence in each judgment by using a three-point scale ranging where '3' = "I am highly confident in my judgment about this string", '2' = "I am moderately confident about my judgment", to '1' = "I am very unsure about my judgment, indeed, I'm just guessing." This scale was printed and present as a reminder during this phase of the experiment, and subjects simply had to circle the appropriate number corresponding to their level of confidence. Participants were also told that roughly 2/3 did follow the rules, and that some of the test strings were identical to those from the learning phase. They were given approximately 15 seconds during which to make each judgment. A template made of 8 ½ "x 14" posterboard with a window cut out was provided for each of the participants. In this way, the experimenter was able to pace and track each specific test item as it was being evaluated by all of the subjects in the group. This technique is identical to that used by Manza and Bornstein (1995), and ensured that all participants were appropriating the requisite amount of time to each response, without either rushing or lingering. Response times were not recorded.

Affective evaluations. The same stimuli (re-randomized) and presentation conditions were used as in the well-formedness testing, but here participants were asked to evaluate each item in terms of how much they liked it. They were told that the notion of liking should be viewed as a purely personal judgment on their part and to simply rate each string on a six-point scale where '6' = "I really like this item," '5' = "I like this item," '4' = "I sort of like this item", "3" = "I sort of don't like this item", "2" = I don't like this item, and '1' = "I really do not like this item." This scale was also present to the subjects throughout this phase, and subjects had simply to circle the number (1-6) under each test string which reflected their affective attitude towards it and can be seen in Appendix C.

Data recoding. The traditional way to view the well-formedness judgments in an AG learning task is on a two-point scale in that a subject's assessment of the grammaticality of a test string is either correct or incorrect. On the other hand, the liking judgments naturally map onto a six-point scale. In order to make direct comparisons between subjects' behavior with the same stimuli under each task condition the data from each task were recoded as follows:

a) *The well-formedness data:* The well-formedness judgments were transformed into a six-point "strength of grammaticality" scale using confidence ratings such that a "grammatical" judgment with a confidence of "3" received the highest grammatical score of "6"; whereas a "nongrammatical" judgment with a confidence of "3" received the lowest score of "1". In other words, a score of "6" reflects greatest confidence that a string is indeed grammatical, and a score of "1" greatest confidence that a string is ungrammatical. This procedure was also used by Gordon and Holyoak (1983) and produces a six-point score that parallels the one used in the liking task.

b) *The liking data:* In a parallel way, the affective ratings were transformed from the straightforward six-point scale to a two-point scale. Ratings of "4", "5", or "6" (the high liking scores) were deemed to be "correct", "coordinate" or "hits" (as per the prediction) when given to OldG or NewG strings, and as an "error" or "false alarm" when given to a NonG string. Likewise, ratings of "1" , "2", or "3" (lower liking scores) given to a NonG string were regarded as being "coordinate" or "hits", and to grammatical strings (OldG or NewG strings) as an "error". In the analyses presented below, the data are given in both the straightforward and the transformed versions -- depending on which is appropriate.

Experiment One - Japanese Kana characters

In the first experiment of this series, Japanese Kana characters were selected to instantiate the learning and test strings particularly because of their non-pronouncability (see Figure One). Also, participants were told simply that this was a learning experiment, and the use of a 'foreign language' is consistent with this description.

Method

Participants. Thirty subjects were recruited from the Introductory Psychology subject pool, with the proviso that they had no knowledge of either the Japanese or Chinese language.

Stimuli. Five samples of the Japanese Kana syllabary acquired from Corel WordPerfect character set were used as the elements to assemble the strings. All stimuli were printed onto overhead transparencies.

Procedure. Each of the learning strings was projected one at a time, while participants were asked to copy what they saw as accurately and quickly as they could in the booklets provided. After viewing and copying all twenty learning strings three times each (in a quasi-random fashion), subjects were then asked to begin rating the test items for each of the two consecutive tasks. All responses were recorded in the test booklets.

Results

Well-formedness judgments: There was clear evidence of learning effects, as had been found in so many other implicit learning studies involving artificial grammars (i.e., Reber, 1989, 1993; Berry, 1997). As seen in the top panel of Table 2, participants are able to distinguish between strings that conform to the rules of the grammar from those do not. Test items were endorsed as being grammatical 75% of the time when they were OldG strings, 62% of the time when they were NewG strings, and only 45% of the time when they were NonG strings. Whereas earlier studies included quite conspicuous violations in the nongrammatical test strings, the NonG strings in this study contained only single-symbol violations and were actually quite subtle in the extent to which they violated they rules. The ability of the participants to be able to discriminate grammaticality successfully is consistent with Chan's (1992) and Altmann, Dienes and Goode's (1995) findings in that an AG instantiated with arbitrary and unfamiliar strings can be acquired using a standard learning protocol.

Moreover, there is also evidence for instantiated memories of individual strings used during the learning phase. The OldG strings borrowed from the learning set were significantly more likely to be classified as grammatical than the new-yet-consistently-formed (New G) strings (as per Tukey's HSD at .05 signficance level). This finding is not surprising since instantiated memorial representation is often observed in AG learning (Brooks & Vokey, 1991; Perruchet & Pacteau, 1990). As in shown in Table 2, these effects emerge in a similar pattern when the data are converted into the six-point scale where OldG items receive the highest well-formedness rating (4.28) followed by NewG (3.88) and NonG (3.42) strings.

A 2 x 3 (task order x stimulus type) ANOVA was performed of the 6-point well-formedness data. As expected, there was a strong effect of stimulus type, $F\,(2, 29) = 29.3$; $MSE = 2.72$; $p < .0001$. There was, however, no task order effect, although there was a task order by stimulus type interaction $F\,(2, 29) = 7.78$; $MSE = 19.31$; $p < .001$. This effect turned out to be due to subjects displaying a higher level of performance on the NewG and NonG items when the well-formedness task was run *before* rather than *after* the liking task. This was a somewhat anomalous finding with no obvious explanation. As will be seen in later studies, task order effects, when they do arise, appear to be quite unpredictable and unreliable. In that vein, all following analyses were run collapsed across order. Specific task order effects will be discussed in context, where they arise.

Confidence: There was no significant correlation between accuracy and confidence ratings ($r = .034$; n.s.). Moreover, participants were no more accurate in their well-formedness judgments on trials when most confident of their classification ($M = .67$, $SD = .47$), than when either only slightly sure ($M = .63$, $SD = .48$) or totally unsure ($M = .63$, $SD = .48$). This finding is consistent with Chan's (1992) findings when he employed nonpronounceable elements to instantiate the grammar and fits the 'zero correlation criterion' which Chan claims to be an indication of nonconscious acquisition. Note, however, that even when subjects claim that they are just guessing and have no confidence in their classifications, (confidence ratings of "1"), their accuracy level is nonetheless reliably above chance, $t(29) = 2.67, p < .05$. The accurate grammatical classifications in the absence of any confidence meets the 'guessing criterion' and suggests that participants have little in the way of metaknowledge of the underlying patterns in the displays they have been working with all along (Dienes, Altmann, Kwan, & Goode, 1995; Dienes & Berry, 1997). The finding that the confidence ratings are not predictive of classification accuracy points to the implicit nature of the grammatical acquisition.

Table 2

Experiment 1

Well-formedness judgments with Japanese Kana characters

N = 30

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.75	4.28 [a]
New Grammatical	.62	3.88 [b]
Non Grammatical	.45	3.42 [c]

Liking Judgments with Japanese Kana characters

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.49	3.48 [d]
New Grammatical	.55	3.72 [d]
Non Grammatical	.53	3.59 [d]

[abcd] Superscripts indicate which stimulus types differ as per Tukey's HSD at .05 level

The above pattern of findings suggests that both implicit and explicit elements are contributing to the participants' behavior in this experiment. While the increased accuracy observed with the familiar OldG items suggests an explicit memorial residue of the learning phase of the experiment, it does not manifest itself in any significant increase in confidence. In addition, when stimulus types (OldG, NewG, NonG) were looked at separately, no pattern emerged. Participants showed no higher confidence in their judgments of the OldG strings ($M = 1.90$) than they did when classifying NewG ($M = 1.95$), or NonG items ($M = 2.03$).

Affective ratings: While the well-formedness data were in keeping with the literature in implicit learning, the data on the liking ratings were not. When using these unfamiliar Japanese Kana characters to instantiate the grammar, no evidence was found for the structural mere exposure effect. As is shown in the bottom part of Table 2, participants showed no differential liking for any particular stimulus type, $F\ (2, 29) = 1.85$, $MSE = 4.86$; $p > .05$. Moreover, grammatical strings were not liked any better than the nongrammatical strings. What is even more surprising was that here was no evidence even for the standard (classic) mere exposure effect. Participants showed *no special liking* of the OldG strings that had been part of the original learning set. The overall mean liking scores given for OldG, NewG, and NonG strings were 3.48, 3.72, and 3.59, respectively, none of which were statistically distinguishable from the others, or, for that matter, differing significantly from chance expectations, $\chi^2(2) = 2.99$; $p > .20$.

Discussion

The results of this initial experiment are somewhat surprising in that while demonstrating the anticipated findings that are consistent with the large literature on implicit learning, it is nonetheless discoordinate with an equally impressive literature on the mere exposure effect. While replicating Chan's (1992) and Altmann, Dienes and Goode's (1995) finding of AG learning with novel, nonalphabetic stimuli and showed that subjects can work with these unfamiliar displays in ways that parallel what has often been reported using more prosaic displays such as English consonants (see Reber, 1993 and Reber, Allen & Reber, 2000 for overviews), it fails to replicate what Gordon and Holyoak (1983) and Manza and Bornstein (1995) had found, in that there was no evidence of an increased liking of well-formed novel strings (NewG) over those with grammatical violations (NonG). In fact, there wasn't even evidence for increased liking of already-seen stimuli

There really never has been any systematic investigation of the classic mere exposure effect when multiple elements are involved. In this current study, not only do the stimulus displays consist of more than one element at a time (i.e., strings of letters), but their structure is determined through a set of governing rules which are non-obvious.

Actually, then, there are two anomalous findings here: 1) the failure to find the structural mere exposure effect of Gordon and Holyoak and, even more surprisingly, 2) the failure to find the more pedestrian (classic) mere exposure effect of Zajonc and colleagues. This finding is unique, in part, because while Gordon and Holyoak (1983) and Manza and Bornstein (1995) both found evidence of the *structural* mere exposure effect, neither study ever re-presented OldG strings during testing and, hence, the possibility of finding the *mere* mere exposure effect had never really explored.

Another key difference between this experiment and the two earlier studies of the structural mere exposure effect is the use of highly unfamiliar Japanese Kana characters. However, while the unusual nature of these symbols seems to be playing a role in the formation of preferences, it evidently played little or no role in subjects' ability to acquire knowledge of the grammatical patterns exhibited. Well-formedness judgments were made accurately and readily to all stimulus types, and subjects appeared also to be sensitive to whether or not test items had been presented during learning.

These results invite an obvious follow-up. Consequently, the next experiment has a slightly retrograde quality to it. It makes sense to first replicate the earlier success enjoyed when researchers used English consonants to investigated the structural mere exposure effect. However, the addition of OldG items in the testing phase was still included to see whether the classic mere exposure effect also emerges under these conditions.

Experiment 2 - English Consonants

It seems prudent at this point to establish that the mere exposure effect can be generalized to structured stimuli, as has been shown by others. In essence, this study is a replication of Experiment 1 with the difference that highly familiar English consonants are used to instantiate the grammar (see Figure 1).

Method

Participants. Twenty-nine Brooklyn College undergraduates participated in the experiment to fulfill a requirement of an introductory psychology course. They were run in small groups of no more than three subjects at a time. This was done because there was some concern raised as to the level of attention paid by the subjects participating in large groups, as in Experiment 1.

Stimuli. The stimuli were the same sequences as in the previous experiment, except that the strings were instantiated using the letters "P, S, T, V, X" of the English alphabet. Each of the learning items was printed on individual 3" x 5" index cards (flash cards).

Procedure. There was a slight shift in the manner in which the stimuli were viewed. Participants looked at the individual flash cards for approximately 15 seconds each, after which they were asked to reproduce the strings as quickly and accurately as they could. All responses were recorded in the test booklets provided. After this initial learning phase, all subjects participated in both the well-formedness and the liking tasks, again with order counterbalanced. The rating systems for well-formedness and liking judgments were the same as before. The smaller groups hopefully facilitated the participants' attention, and precluded the need for an overhead projector.

Results

Well-formedness: As expected, the typical AG learning effects were observed. As is displayed in the top panel of Table 3, subjects were able to reliably discriminate grammatical from nongrammatical strings, with OldG strings endorsed as being grammatical 81% of the time, NewG strings being endorsed 62% of the time, and NonG strings only getting 37% endorsements. Moreover, subjects showed a distinct sensitivity to previously presented item in that the OldG strings were correctly endorsed significantly more often than were the NewG strings.

Table 3

Experiment 2

Well-formedness judgments with English letters

N = 29

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.81	4.85 [a]
New Grammatical	.62	4.02 [b]
Non Grammatical	.37	3.09 [c]

Liking Judgments with English letters

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.65	4.04 [d]
New Grammatical	.63	3.99 [d]
Non Grammatical	.51	3.56 [e]

A 2 (task order) x 3 (stimulus type) ANOVA of the well-formedness scores (using the transformed six-point scale) shows that participants were quite capable of making grammatical classifications, $F(2, 28) = 47.03$, $MSE = 5.99$; $p < .0001$. As in Experiment 1 there was an unusual task order effect. This time, rather than emerging as an interaction, the order effect showed up as a simple effect with subjects performing slightly *better* when the well-formedness task was run *after* the liking task ($M = .73$) than when it was run before ($M = .66$); $F(1, 28) = 5.06$, $MSE = 12.93, p < .05$.

At first, this finding seems to suggest that the concern articulated by Gordon and Holyoak (1983) that one task might contaminate the other has some justification. However, this order effect turns out to be in exactly the opposite direction of the one which concerned them. Here, experience with the liking task seems to be having an impact on the performance of the well-formedness task, and, as is discussed below, there were no order effects of any kind found in the liking data, contrary to any of the concerns of the earlier researchers.

Confidence: Unlike Experiment 1, accuracy and confidence ratings were modestly but significantly correlated, $r = .21, p < .01$, a degree of concordance observed by other researchers in similar situations (see Dienes and Berry, 1997). Subjects were more accurate on those trials where they reportedly had the highest confidence ($M = .82$) than when than when they reported either moderate ($M = .63$) or low confidence ratings ($M = .60$). In short, subjects demonstrated better metacognitive ability in this study than in the previous one (Japanese characters) in that they seemed to be more assured of their grammaticality judgments. However, unlike Experiment 1, these participants were no better than chance in making grammatical classifications when they reported their responses to be just guesses, $t\,(24) = .35; p > .05$.

To summarize, these subjects appeared to meet neither the 'zero correlation' criterion nor the 'just guessing' criterion proposed above. One possible interpretation of these data might be that subjects had a clearer (more explicit?) sense of what constitutes grammatical strings. However, it would be premature to draw any definitive conclusions from just these two experiments, although it is worth noting the discrepant results. Perhaps in an effort to dissuade possible critics from claiming the well-formedness judgments are tainted with explicit knowledge, this pattern of results might hint at the advantage of using non-pronounceable elements to instantiate the grammar. Evidently, even in the presence (or absence) of this metacognitive insight, grammaticality judgments were still significantly above chance.

Affective ratings: As the lower panel in Table 3 shows, the use of English consonants to instantiate the AG succeeds in eliciting the structural mere exposure effect. Participants in this study show significantly higher liking scores for well-formed strings than for those that violate the rules of the grammar, $F\,(2, 28) = 10.32$, $MSE = 24.38$; $p < .001$. This result successfully replicates both Gordon and Holyoak's (1983) and Manza and Bornstein's (1995) findings. Moreover, the results indicate no task order effects. Interestingly, while there is strong evidence for the *structural* mere exposure effect, there really is none for the *classic* mere exposure effect, a feature which was not examined by earlier researchers Subjects do not like the OldG strings any better than they do the novel grammatical ones (NewG) in that, at least under these conditions, the structural integrity of the stimulus displays appears to have at least as strong an impact on affective evaluations as does physical identity. This pattern of findings suggests that the processes used in making affective evaluations may be different from those engaged while making well-formedness judgments. It is clear from the well-formedness data that subjects are sensitive to the re-presentations of the stimuli from the learning set (OldG strings), but there is no such indication that they have come to prefer these stimuli more than those which had been seen before.

Comparing liking and grammaticality: Manza and Bornstein (1995) proposed that the liking task might function as a more sensitive measure of implicit knowledge than the usual well-formedness classification task. While their design did not quite permit a direct comparison of these two measures, this one does—and the results appear to raise more questions than they answer.

In essence both dependent measures (liking and well-formedness) are different ways of assessing subjects' sensitivity to the grammatical integrity of the strings. As Macmillan and Creelman (1991) show, signal detection theory can be used in AG learning studies to assess relative sensitivity across treatments which, in this case, would be the two measures of applied knowledge. Accordingly, two measures of d' were calculated using the six-point scales, one from the liking data and one from the well-formedness data. As the figures in Table 3 suggest, the well-formedness judgments reflect a clearly greater sensitivity to structural integrity ($d' = 1.49$) than the liking judgments ($d' = .94$).

There seem to be two reasons for this difference in expressed sensitivity. First, as noted, subjects did not show the classic mere exposure effect in the NewG strings attracted liking scores that were essentially equivalent to those given the OldG strings. Second, the liking scores assigned to the NonG strings tend to cluster around the neutral point on both the two-point (51%) and six-point scales (3.56). What this finding suggests is that in experiments using complex, structurally organized stimulus displays, subjects do not react to ill-formed strings with negative affect. Instead, they typically display more-or-less neutral responses. In other words, they don't *dislike* these nongrammatical strings; they just don't especially *like* them very much to speak of. As noted earlier, these same subjects are capable of making appropriate grammatical classifications regarding these nongrammatical strings so they are clearly sensitive to the structural features in the display. This result is not very different from what Zajonc would have predicted. The mere exposure effect is the resulting elevation of affect due to repeated exposure. There is no theoretical premise for expecting items which have not been repeated (i.e., NonG strings) to get *low* liking scores. These scores would just be *relatively lower* than those given to previously seen items (the grammar of the OldG and NewG strings).

The point is that these differences in sensitivity do not necessarily map onto or are expressed equivalently in all measures of experience. Manza and Bornstein may have been on the right track, in that the liking judgments may be a better *indirect* measure of implicit knowledge, or of memorial representations held largely outside of consciousness and unavailable for controlled recall. Indeed, the rough equivalence of the liking scores attracted by the OldG and NewG items suggests that this may be true. But it is doubtful that these responses are a *better* measure of the subjects' overall ability to respond to variations in structural integrity. As Merikle and Reingold (1991) note, it can be exceedingly difficult to tease apart direct and indirect measures of held knowledge. It would be presumptuous to assume that absolute task purity is controlled for in this design. That grammaticality judgments show a disparity given to OldG and NewG strings is suggestive of a combination of both explicit and implicit influences. In all likelihood, the greater *d'* reflected in the well-formedness data may well be due to a mix of these strategies. Accessing greater resources would more likely lead to greater classification success.

Discussion

As expected, the data reveal typical AG learning. In addition, by instantiating the AG with highly familiar letters, the structural mere exposure effect was now observed. These conditions and results resemble what Gordon and Holyoak (1983) and Manza and Bornstein (1995) found earlier. What is particularly interesting is that this structural effect seemed to overwhelm the commonly observed classic mere exposure effect. Since these latter findings were not observed in Experiment 1 (Japanese characters), but they were found in this and earlier studies, the obvious suspect appears to be to the specific elements used to instantiate the AG.

What exactly is it about an overall sense of familiarity that seemingly contributes differentially to these two exposure effects? Participants show a sensitivity to the physically familiar in that they endorse the grammaticality of the OldG items at a greater rate than the NewG's, but this sense of the familiar was not expressed in their affective judgments. Perhaps the exact physical representation of the stimuli is not contributing any advantage over the structural familiarity. Put another way, under these conditions the structural mere exposure effect may be overwhelming the "mere" mere exposure effect. It may be that the underlying familiarity of the structure is driving the subjective sense of familiarity resulting in the greater liking of the grammatical stimuli.

In short, by shifting back to the use of *a priori* familiar stimuli to instantiate the AG, the structural mere exposure effect was restored. The implication is that there may be two distinct kinds of familiarity operating here and that they are playing themselves out in different ways in these studies. The classic mere exposure effect, involving singular, discrete items, is apparently easily obtained with any arbitrary stimulus domain, whereas the *structural* mere exposure effect seems to require that the rule-governed displays be instantiated using stimulus items which already are within the participants' repertoire. The obvious question now is whether this sense of familiarity can be acquired in a relatively straightforward manner in a controlled laboratory setting by extending the exposure to and promoting involvement with the learning strings instantiated with unusual elements.

Experiment 3 - Chinese Calligraphy

If it is the case that participants' relative unfamiliarity with certain types of stimuli compromises the structural mere exposure effect, then one possible solution would be to extend and enrich the learning segment of the experiment. This may be done by exposing the subjects to more instances of the learning strings. However, this design already has 60 strings (each of 20 items presented three times), a number which should be adequate for learning and familiarization. The repetitious nature of copying strings could conceivably make the learning phase tedious and unpleasant. Another approach to this problem would be to engage subjects with the unusual stimuli in a manner that would not only increase their level of familiarity, but also heighten their involvement with them. This is the route chosen for the next experiment. Specifically, this design uses highly unusual and linguistically unfamiliar items (Chinese characters), which are virtually unpronounceable to non-Chinese speaking participants. Many participants had reported the 'liking' task to be almost somewhat amusing in that being asked to judge 'strings' seemed a little odd. The concern here is the possibility that subjects were not taking this task as seriously as they should. So, in an effort to increase the face validity of the task, it might be more appropriate to evaluate items which were, in fact, of an artistic nature. The Chinese characters used for this experiment were therefore reproductions of calligraphically produced brushwork of no more than five simple strokes each.

Method

Participants. Twenty-five Brooklyn College undergraduates participated in the experiment to fulfill a requirement for an introductory psychology course. Subjects were run in small groups of no more than three, as was the case in Experiment 2. All participants were screened to exclude those with any prior knowledge of Chinese or Japanese.

Stimuli. Participant were provided with a calligraphy pens and told that they were participating in an experiment in which they would learn to paint Chinese characters The actual characters used in this study are shown in Figure 1. These exemplars were 'artful' examples of Chinese characters reproduced from a Chinese language textbook which featured the characters with simplicity and elegance. From a structural point of view, the stimuli for all parts of the experiment were the same as used in the two earlier experiments, the only difference being that now the strings were instantiated by the five Chinese characters shown in Figure 1. As in Experiment 2, all learning strings were printed on individual 3" x 5" flash cards.

Procedure. Participants were presented with the 20 individual learning strings asked to try to reproduce each of them as best they could by using the calligraphy pens. They were given up to 30 seconds to copy each string and were invited to "have fun with this experience." The full set of 20 strings was presented three times with order randomized. The experimenter paced the subjects through the series of exemplars, and, as in the previous experiments, task order was counterbalanced by group.

Although this learning session took relatively longer to complete (approximately 45 minutes versus the approximately 25 minutes taken in each of the first two experiments), the participants appeared to be quite involved with the task and seemed actually to enjoy the activity.

Results

Well-formedness judgments: There were no task order effects found in that subjects' grammatical assessments were equivalent in terms of accuracy, regardless of which task was performed first. As is shown in the top panel of Table 4, participants reliably discriminated among the stimulus types, endorsing well-formed strings at a significantly higher rate than those that violated the rules of the grammar. OldG strings were endorsed as being grammatical 62% of the time and NewG strings were endorsed 58% of the time. NonG strings received grammatical endorsements 51% of the time, which is at chance. A one-way ANOVA (collapsed across order) of the 6-point well-formedness judgments showed the expected AG learning phenomenon, $F(2,24) = 4.86$, $MSE = 13.19, p < .01$. Although significant, these levels are somewhat below those found in Experiment 2.

Confidence: There was no correlation between the overall accuracy of the well-formedness responses and confidence levels in that subjects seemed not to be aware of when their judgments were correct ($r = .04$). And, as in Experiment 2, subjects were no better than chance on the trials where they claimed to be just guessing, $t(24) = .53$, *n.s.*

However, there are differences in the findings here compared with those from the two earlier studies. Specifically, there was surprisingly poor performance with the NonG items in that virtually half of them were erroneously endorsed as 'well-formed'. The intense practice with the strings of Chinese characters either did not yield greater induction of the underlying structure of the sequences, or, as suggested by the data regarding confidence levels, subjects were not as sure of their grammatical assessments and were therefore reluctant to declare something as being 'nongrammatical'.

Table 4

Experiment 3

Well-formedness judgments with Chinese characters

N = 25

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.62	3.95 [a]
New Grammatical	.58	3.70 [a,b]
Non Grammatical	.51	3.53 [b]

Liking Judgments with Chinese characters

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.44	3.32 [c]
New Grammatical	.44	3.27 [c]
Non Grammatical	.49	3.29 [c]

Affective ratings: These results were disappointing in that not only did the focused and extended practice with these stimulus displays produce somewhat poorer overall performance in grammatical classifications (as mentioned above), but they yielded essentially no evidence for either the structural or for the classic mere exposure effect. A one-way ANOVA of the liking data in the bottom panel of Table 4 showed virtually no differences in liking for any stimulus type, $F(2, 24) = .082$, $MSE = .168$, $p > .05$, suggesting that they were more or less liked equally. In addition, the affective ratings for the three stimulus types were neither significantly different from those gathered in the control group, nor were they different from what would have been expected by chance [$\chi^2(2) = .59$, *ns*].

Discussion

These results more closely parallel those from Experiment 1 (wherein unfamiliar Japanese Kana characters were used to instantiate the grammar) much more closely than they do those from Experiment 2 (where highly familiar English consonants were used). Evidently, group size did not matter, and the extended time spent with the items and degree of involvement with the learning strings did not offer any advantage in the expression of exposure effects.

Participants could reliably make well-formedness judgments but showed no particular liking for those strings which were grammatical--not even for the OldG strings with which they had worked so extensively during the learning phase. This last result is especially surprising since, in keeping with the standard findings on the mere exposure effect, it was expected that these OldG strings would garner relatively higher liking ratings than the other test items. The learning phase of the experiment did, after all, provide quite an intensive and extended exposure.

In passing, it is interesting to note that the elements of these stimulus displays were, at least superficially, quite similar to those used in a variety of other studies where the mere exposure effect was successfully obtained (e.g., Zajonc, 1968, 1993). However, there are two key differences between the current experiments and the earlier explorations. First, in these earlier studies the stimuli were always presented as single items and not as sequences. Second, the stimuli in the current experiments are sequences of *structurally* coherent elements. When these sequential stimuli (defined by the complex rules of the AG) are clothed in unfamiliar and unusual elements, it appears then that neither of the two mere exposure effects (classic and structural) seems to appear. Evidently, participants are learning something about the underlying structure, as seen by their grammaticality judgments, but they do not appear to display any particular shifts in their affective evaluations as a consequence.

Prompting subjects for grammatical assessments of strings necessarily invites participants to evaluate strings in their entirety. However, as seen in Experiments 1 and 3, it is entirely possible that focusing so intently on the individual elements of the strings may be disrupting the mere exposure effect. As this was not the case with the use of English consonants (Experiment 2), it might be assumed that the *type of familiarity* plays a greater role than was even suspected.

As suggested earlier, there has been up until now a confounding of the two meanings of *familiarity*. In the literature on the mere exposure effect, the term is used so that it refers to an induced or controlled familiarity. That is, the typical study takes initially unfamiliar, affectively neutral stimuli and examines how participants come to develop positive hedonic tone for those very stimuli as a function of the circumstances under which they have been confronted. In the classic mere exposure studies participants generally have no pre-experimental familiarity with the stimuli.

However, in the *structural* mere exposure effect, a second type of familiarity gets introduced. While the participants in the studies of Gordon and Holyoak (1983), Manza and Bornstein (1995), and Experiment 2 above were initially unfamiliar with the underlying rule-governed characteristics of the displays, the actual stimulus components used to instantiate these AGs were extremely familiar. In Experiments 1 and 3, however, both the structural characteristics of the displays *and* the physical stimulus elements used to instantiate them were unfamiliar to the subjects. English consonants are easily identifiable and quite readily pronounceable. The identification and pronunciation of the components may facilitate an affective evaluation of the overall strings. Readers naturally group letters which are strung together to form words, and the stimuli in this study certainly resemble word-like items.

One way of examining this further might be to use a set of symbols which were mundane, familiar, but nonlinguistic. These elements should be highly familiar to the testing population, yet iconographic enough to fit the criteria of being 'arbitrary shapes'. They should be somewhat less word-like and should not fall naturally into language-like groups and chunks; of course, any shape or squiggle could be assigned a name, and therefore render it pronounceable. However since it would be best not to facilitate subvocal pronunciations, a set of common keyboard symbols which seem to meet the criteria were used.

Experiment 4 - Keyboard Symbols

This experiment is essentially a replication of the previous three experiments, with the exception this time using common keyboard symbols.

Method

Participants. Thirty-four Brooklyn College undergraduates participated in this experiment as part of an introductory psychology course requirement. They were run in groups of approximately 8 to 10 students each.

Stimuli. The strings in this experiment were the same sequences as those sequences used in the first three experiments, with the exception that this time the sequences were instantiated using familiar symbols found on the common computer keyboard: *, #, $, %, @.

Procedure. The learning strings were presented in much the same way as the test stimuli had been viewed above and was the same presentation technique which was used by Manza and Bornstein (1995). All learning and test stimuli were printed in the booklets in which subjects also recorded their responses. During learning, individual exemplars were viewed by means of a template that was slipped across the page, exposing one string at a time. This allowed for slightly more subjects to participate per session, and kept the flow of the experiment moving at a reasonable pace. As in the earlier experiments, the full set of learning strings was presented three times in quasi-randomized order.

The testing phase similarly involved the viewing of each test item one at a time, again through the use of the template. All other stimulus properties, instructions and counterbalancing were the same as before.

Results

Well-formedness judgments: As anticipated, a 2 (order) x 3 (stimulus type) ANOVA showed the typical implicit learning effect with subjects reliably differentiating the stimulus types, $F\,(2, 33) = 68.04$, $MSE = 2.48$; $p < .0001$. As in Experiment 2 (English consonants), participants endorsed the previously seen OldG strings most often (.75), the NewG strings somewhat less often (.59), and endorsed the NonG infrequently (.38). These data are presented in the top panel of Table 5. There were no task order effects with regard to grammaticality judgments.

Confidence: There was a small but statistically significant correlation between participants' accuracy of grammatical assessments and their confidence, $r = .13$, $p < .01$. However, subjects were no better than chance on trials where they claimed to be "just guessing," $t\,(33) = .43$, *n.s.*

Overall, these findings are in keeping with those found in the first three experiments, and further supporting the notion that AG learning is highly general and can be demonstrated with a wide variety of stimulus forms. Varying the items used to display the grammar appears not to interfere with participants' ability to make well-formedness judgments.

Table 5

Experiment 4

Well-formedness judgments with keyboard symbols

N = 34

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.75	4.43 [a]
New Grammatical	.59	3.86 [b]
Non Grammatical	.38	3.13 [c]

Liking Judgments with keyboard symbols

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.71	4.27 [d]
New Grammatical	.59	3.71 [e]
Non Grammatical	.57	3.73 [e]

Affective ratings: An ANOVA on the affective ratings shown in the bottom panel of Table 5 revealed an exposure effect, $F(2, 33) = 16.57$, $MSE = 2.46$; $p < .001$ with well-formed strings being liked more than those that violate the rules of the grammar. However, here there is a somewhat different pattern of results from those found with the English consonants. While there is strong evidence for the classic mere exposure effect, support for a structural exposure effect has faded. The previously encountered (OldG) strings received the highest liking scores (4.27) compared with the novel NewG (3.71) and NonG (3.73) strings, showing that the effect of stimulus type is entirely contained in the higher liking scores for test strings which were repeated from the learning set (OldG). The NewG and NonG strings were regarded equivalently, suggesting that the exposure effects in this study are manifested primarily of the classic kind. So, in keeping with the classic mere exposure effect, only those strings which were actually seen before elicited higher liking ratings.

In addition, there was a main task order effect, $F(1, 33) = 15.01$, $MSE = 2.70$; $p < 001$ in which the overall exposure effect was more pronounced when the liking ratings were made after the well-formedness judgments. However, those subjects who nonetheless first provided liking ratings did still demonstrate increased affect for the OldG strings.

Discussion

The well-formedness data are totally unsurprising and consistent with the literature. As noted, AG learning is seen with all stimulus instantiations, although it appears that subjects working with the most familiar items, the English consonants of Experiment 2, are somewhat better at making grammatical distinctions than those working with other symbol sets.

Here again, the interesting findings are those from the liking task. The keyboard symbols used in this experiment seem to fall somewhere in the middle of a continuum of *a priori* familiarity. These symbols appear far more frequently within the participants' daily experiences than do the cryptic Chinese and Japanese characters of Experiments 1 and 3, but considerably less often than the overwhelmingly commonplace English consonants of Experiment 2.

The results so far reflect an intriguing relationship between performance and *a priori* familiarity. With the highly familiar English letters both the classic and the structural mere exposure effects are obtained; with the totally unfamiliar Japanese Kana characters and Chinese ideograms, neither is found. With the modestly familiar keyboard symbols, the classic mere exposure effect was obtained (OldG), but there was little evidence for heightened preferences based on solely on structural characteristics (NewG).

In some ways, this finding is a major first step in this overall study. Setting aside the existence of an underlying structure, this study could be viewed as being basically a test of the classic mere exposure effect using stimuli which consist of multiple symbols. Whereas the literature on exposure effects has traditionally defined stimuli as discrete items, participants in this experiment did demonstrate that serial stimuli (i.e., items consisting of many individual elements) do succeed in eliciting the classic mere exposure effect.

In passing, it should be noted that there were some interesting parallel findings from the normative ratings of the stimulus sets. As mentioned earlier, all stimuli were given affective evaluations by a group of control subjects which were asked simply to provide baseline liking ratings. This was done as a precaution against any particular attraction that certain strings or groups of strings might have. As seen in Table 1, the strings composed of keyboard symbols actually had a higher *base* liking rate than any other symbol set. In fact, the (unbeknownst to the participants) "nongrammatical" strings had the highest overall affective rating of any group of stimuli in the study. Consequently, it is not entirely surprising that the NonG keyboard strings received higher liking ratings than has been seen in the first three experiments. The baseline mean 3.73 liking rating of NonG strings from the experimental subjects is, in fact, not reliably higher than the 3.62 obtained from the control subjects, $t\ (33) = .84$.

The generalizability of the mere exposure effect to structured stimuli appears to be most robust when the stimuli contain highly familiar elements. Indeed, thus far, in four separate experiments with a variety of stimulus types, heightened affect for novel grammatical items (NewG) was only observed when the highly familiar English consonants were used to instantiate the grammar (Experiment 2). What would happen if stimulus familiarity was manipulated not by adjusting the stimulus elements, but rather by selectively using specific participants? The following study delves into this question.

Experiment 5: The Russian Studies

The first four experiments have demonstrated that, although well-formedness judgments do reveal grammatical discriminations regardless of the nature of the stimulus elements, it appears that the structural mere exposure effect requires a degree of *a priori* familiarity. The use of English consonants in Experiment 2 worked perfectly well in demonstrating exposure effects that reflected grammatical sensitivity to both OldG and NewG test items. Experiment 4 showed that when using highly familiar, yet non-linguistic symbols, a classic–but not structural–mere exposure effect could be obtained. Even more striking, though, was the utter lack of (affective) exposure effects of *any kind* when very unusual (albeit linguistic) elements were used to instantiate the grammar (as in Experiments 1 and 3). Since this inconsistency was not evident in terms of the well-formedness judgments, this suggests that the (implicitly) acquired knowledge of the structure did not necessarily contribute to affective evaluations.

The results seen so far implicate the unusual nature of the stimulus elements as the pivotal factor in this experiment. This suspicion is consistent with the data in that well-formedness judgments probe the structural integrity of the strings in their entirety. On the other hand, when making liking judgments, participants either failed to evaluate the strings holistically, or relied on alternative impulses. Everyday keyboard symbols were used to instantiate the stimulus strings (Experiment 4) to see whether the structural mere exposure effect requires items which are already within he subjects' repertoire. And, as discussed above, the results were somewhat promising yet ambiguous in that only the classic, but not the structural effect was found.

If the speculation about using highly familiar items is correct, it should then be possible to manipulate this factor by selecting a specific subpopulation from the general subject pool.

The Brooklyn College community offers just such a population. There are a substantial number of Russian immigrants enrolled in the introductory psychology classes. Bilingual Russian students could be recruited to evaluate test items instantiated with English and/or Russian letters of the alphabet. It is expected that both letter sets should appear fully familiar in terms of linguistic saliency and fluency.

Method

Stimuli. The same learning and test items that had been used in the first four experiments were also used in this study. The English consonants were instantiated just as they were in Experiment 2. In addition, five letters of the Cyrillic alphabet were selected from Word Perfect which represent typical consonants used in Russian (see Figure 1). A native Russian speaker/reader was consulted in this selection.

All learning and test items were printed in a booklet as was done in Experiment 4. A template with a cutout window provided the experimenter with a means with which to pace all subjects through the sessions. Participants were run in groups of between 10 and 15 subjects each. As before, task order was a between group factor, as was the type of alphabet used.

Participants. The bilingual Russian students present the opportunity to note Russian subjects' responses to stimuli presented using their native Russian alphabet. It is also of interest to note their responses to items instantiated with English letters, which are presumably within the scope of fluent readability. On the one hand, it is expected that the English consonants would certainly qualify as being 'familiar', since all college classes are conducted in English. On the other hand, native Russian speakers/readers might have relatively less familiarity with the English than the Russian letters. Would this disparate familiarity be mapped out in terms of performance? The current experiment is therefore actually comprised of four separate conditions: English speaking subjects rating English letters; Russian speaking subjects judging English letters; English speaking subjects rating Cyrillic (Russian) letters; and Russian speaking subjects judging Cyrillic letters.

English-speakers (non-Russian) and native-Russian speakers were randomly assigned to each of four groups so that there were English-speaking subjects tested with either English or Russian letters, and Russian subjects likewise tested with either English or Russian letters, thereby comprising the four conditions.

Normative data. A control group of 14 native-Russian students provided normative data in that they were asked to rate how much they liked each of the 36 test items. This was done in a way which paralleled the previous control group rated all the other sample test items. Russian participants in the control group were shown sets of strings instantiated with both English and Cyrillic letters. The results of these affective evaluations indicate that the strings were not differentiated along any grammatical lines (see bottom of Table 1).

Results and Discussion

Condition A: English speaking subjects tested with English letters

Well-formedness judgments: Thirty-four non-Russian subjects were tested with English consonants. This was, in effect, a virtual replication of Experiment 2. As can be seen from Table 6, the results were predictably what had been found before. As expected, these subjects were able to make well-formedness judgments in a way that reflected their learning of the grammar, $F(2, 33) = 46.31$, $MSE = 3.58$; $p < .001$. As was found in Experiment 2, subjects' accuracy judgments were correlated with confidence [$r = .13$, $p < .01$], but this time judgments were significantly more accurate than would be expected by chance even when subjects were reportedly 'just guessing', $t(30) = 2.19$, $p < .05$. There were no task order effects.

Affective ratings: Consistent with earlier studies and in Experiment 2, the structural mere exposure effect was found when using the English consonants, $F(2, 33) = 16.05$, $MSE = 2.94$; $p < .001$. Although there was a main task order effect, $F(1, 33) = 11.27$, $MSE = 2.66$; $p < .01$, there was no interaction with stimulus type. In this case, affective ratings turned out to be surprisingly greater across the board when provided *before* the well-formedness judgments were made. In other words, subjects provided higher liking ratings to all three stimulus types (OldG, NewG, NonG) when tested before making their grammatical classifications. No such task order effects were found in Experiment 2, and this was certainly a reversal of the anticipated effect of grammatical awareness on liking.

Table 6

Experiment 5 - Condition A

English speaking subjects evaluating English letters

N = 35

Well-formedness judgments

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.72	4.35 [a]
New Grammatical	.59	3.95 [b]
Non Grammatical	.37	3.10 [c]

Liking Judgments

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.58	3.89 [d]
New Grammatical	.54	3.82 [d]
Non Grammatical	.42	3.23 [e]

Condition B: Russian speaking subjects tested with English letters

Thirteen native-Russian speakers participated in this condition whereby letters of the English alphabet were used to instantiate the grammar. As has been the case all along, participants were recruited from the subject pool of introductory psychology students and it was assumed that their level of English language proficiency met the requirements of undergraduate demands at Brooklyn College. As fluent speakers and readers of English, it could be expected that their grammatical judgments and affective evaluations ought to mimic those of Experiment 2 (English Consonants) and likewise not be altogether different from those involving English speaking subjects in Condition A.

Well-formedness judgments: The Russian speaking participants were quite capable of making well-formedness discriminations, $F(2, 12) = 23.21$; $MSE = 3.4$; $p < .001$. Their overall responses tended to be significantly more accurate when well-formedness judgments were made *after* the liking judgments, with no interactions with task order and stimulus type. Classification accuracy was positively correlated with confidence ratings ($r = .20$; $p < .01$), but this time accuracy ratings did not exceed chance when subjects were reportedly just guessing, $t(11) = .89$; n.s..

Affective ratings: As seen in Table 7, participants' affective ratings revealed a structural exposure effect in that grammatical strings were liked more than the nongrammatical ones [$F(2, 24) = 10.74$; $MSE = 4.41$; $p < .001$]. Whereas OldG strings received the highest overall mean liking scores, the NewG strings received ratings that were statistically indistinguishable from those given the OldG strings, thereby providing evidence of the structural mere exposure effect.

There was a slight task order effect in that affective ratings were higher for all three stimulus types when the liking ratings were provided *after* the grammatical judgments [$F\,(1, 24 = 6.62$; $MSE = 3.24$; $p < .05$]. This effect occurred in the predicted direction in that the prior experience with grammatical classifications facilitated the exposure effect.

The importance of the role played by an *a priori* sense of familiarity has not been discredited by these results. It could easily be assumed that the Russian subjects were comfortable enough with the English alphabet to perform in ways that were consistent with their non-Russian counterparts in both the well-formedness and liking tasks. The next step, then, is to use Russian letters to instantiate the grammar, and to see how well both English and Russian subjects perform with this stimulus set.

Table 7

Experiment 5 - Condition B

Russian speaking subjects evaluating English consonants

N = 13

Well-formedness judgments

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.80	4.59 [a]
New Grammatical	.65	4.25 [a]
Non Grammatical	.43	3.22 [b]

Liking Judgments

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.62	3.99 [c]
New Grammatical	.62	3.87 [c]
Non Grammatical	.40	2.98 [d]

Condition C: English speaking subjects tested with Russian (Cyrillic) letters

Russian is a language which is prevalent in and around the Brooklyn College community, and so there were no clear predictions as to how non-Russian readers would respond to strings instantiated with the Cyrillic alphabet. For one thing, it is possible to assume that the Cyrillic letters would be regarded in much the same manner as the Japanese or Chinese characters were in that they are not within the realm of the highly familiar. However, it might also be possible that this time the level of familiarity might more closely resemble that of the keyboard symbols–that is–moderately familiar and yet not readily pronounceable. Thirty-five non-Russian reading subjects participated in this condition which employed Cyrillic letters printed from Word Perfect.

Well-formedness judgments: The findings here were consistent with the previous AG studies in that subjects were able to make well-formedness judgments about the stimuli (see Table 8). English speaking subjects tested with Cyrillic letters were quite capable of differentiating the grammaticality of the test items, $F(2, 34) = 11.39$, $MSE = 2.73$; $p < .001$. Accuracy was significantly correlated with confidence ratings ($r = .11$, $p < .01$). However, subjects did not exceed chance performance when they claimed to be just guessing, $t(34) = .36$; *n.s.* These results are concordant with those found in other AG studies in which the grammar is instantiated with elements other than English consonants (i.e., Experiments 1, 3 and 4).

Affective ratings: An ANOVA of the liking ratings showed a marginal, yet significant effect, $F(2, 34) = 4.23$; $MSE = 2.54$; $p < .05$. The rate of higher affective endorsement (scores ≥ 4) hovered around chance levels,[$\chi^2_{(2)} = 3.02$]. There was, however, a striking task order effect in that it was *only* when the affective ratings were provided *after* the well-formedness judgments that was there any evidence for elevated affect for grammatical stimuli, $F(1, 34) = 17.72$, $MSE = 2.43$; $p < .0001$, as seen in Table 9. The added advantage secured by first making grammaticality judgments parallels the trend which was found when the Russian speaking subjects were tested with English letters.

These results are relatively consistent with what has been found in earlier experiments. The relatively unfamiliar Cyrillic letters were successful in eliciting accurate well-formedness discriminations. However, as has been seen throughout, use of the unusual elements did compromise the structural mere exposure effect. According to these data, the additional experience of having first made grammaticality judgments was so critical that when looking at just those subjects who made their liking judgments first (prior to making their well-formedness judgments), no exposure effect would have been found, $F(2, 16) = 1.69$; $MSE = 1.53$; $p > .05$.

Table 8

Experiment 5 - Condition C

English speaking subject evaluating Cyrillic letters

Well-formedness judgments

N =35

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.63	3.90 [a]
New Grammatical	.58	3.80 [a]
Non Grammatical	.50	3.38 [b]

Liking Judgments

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.46	3.40 [c]
New Grammatical	.45	3.33 [cd]
Non Grammatical	.41	3.11 [d]

Table 9

Experiment 5 - Condition C

English speaking subject evaluating Cyrillic letters - Task Order

Liking ratings made after well-formedness judgments

N = 18

Liking Judgments

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.53	3.58 [a]
New Grammatical	.55	3.64 [a]
Non Grammatical	.45	3.16 [b]

In Experiments 1 and 3 using highly unusual elements to instantiate the grammar, no evidence for exposure effects was found, even when subjects had the benefit of the additional exposure during the well-formedness task. One possible conclusion that might be drawn from the task order effect found here is that the Russian letters were *less unfamiliar* to the participants than were the Japanese and Chinese characters. However, unlike Experiment 4, which also used 'somewhat familiar' elements, these subjects successfully produced the exposure effect *only* after having additional experience with the strings during the well-formedness task. The keyboard symbols used in Experiment 4 were able to elicit classic exposure effects, regardless of whether subjects first had a chance to make grammatical evaluations (although prior classification experience did help).

Condition D - Russian Subjects tested with Russian (Cyrillic) letters

Up until this point, *a priori* familiarity, or lack thereof, has been manipulated through the use of a variety of different types of elements used to instantiate the grammar. Another way of manipulating familiarity is to screen and select participants with prior experience with a particular symbol set. Recruiting native Russian speakers/readers to evaluate Cyrillic stimuli provides the opportunity to test whether results will mirror what has been seen when English speaking subjects evaluate English consonants, the prediction being that the personal histories and experiences with the Russian language would facilitate the structural mere exposure effect.

To this end, 26 native Russian speakers participated in this condition which used Cyrillic letters to instantiate the grammar. All other experimental parameters remained the same. The stimuli were those used in Condition C (English speaking subjects tested with Russian letters).

Well-formedness judgments: As expected, participants were able to make grammatical discriminations, F (2, 25) = 21.69; *MSE* = 1.08; $p < .001$. The response patterns reflected those seen in the other studies, where OldG strings were the most accurately judged and received the highest ratings of grammaticality (see Table 10). There was no correlation between confidence ratings and accuracy [r = .03], and subjects performed no better than chance when reportedly just guessing [t (23) = .75; *n.s.*].

Affective ratings: The results of this liking task were actually quite surprising. The overall pattern of affective judgments did not differ from chance, F (2, 25) = .44, *MSE* = 2.49; $p > .05$. There was no evidence for any exposure effect whatsoever–either classic or structural–and although grammatical strings did appear to garner higher liking scores than did the NonG strings (3.52 vs. 3.42), this trend was not significant. Moreover, these affective ratings did not differ significantly from what would be expected purely by chance [$\chi^2_{(3)} = 2.49$; n.s.].

Table 10

Experiment 5 - Condition D

Russian speaking subjects evaluating Cyrillic letters

N = 26

Well-formedness judgments

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.61	4.14 [a]
New Grammatical	.56	3.98 [a]
Non Grammatical	.41	3.61 [b]

Liking Judgments

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.50	3.52 [c]
New Grammatical	.55	3.52 [c]
Non Grammatical	.49	3.42 [c]

In trying to understand what was going on in this case, it might be helpful to look back at the normative data provided by the control group of native-Russian speakers (bottom of Table 1). Interestingly, Russian subjects provided liking ratings for the Cyrillic letters which were actually significantly *lower* than those they had given to the English letters, 2.95 vs. 3.23 [$t\,(88) = 3.07$; $p < .01$]. Although it is difficult to draw any definitive conclusions for this phenomenon, it is fair to say that the native Russian speakers in the control group apparently did not like the Cyrillic letters as much as they did the English consonants.

What is seen here is that even though the Russian speakers were able to make well-formedness discriminations, they were apparently not reacting to these grammatical sensitivities when providing their liking judgments. Although this is a pattern which has been seen before, it becomes somewhat striking in this case in that the role of *a priori* familiarity had been addressed by having subjects evaluate strings instantiated with their native alphabet. Moreover, the affective evaluations of strings of Cyrillic letters proved problematic, in both the control and experimental groups. The prediction that these participants would respond to the already highly-familiarized Cyrillic letters in much the same way that the English speakers respond to English consonants was not supported by these data. Perhaps there are personal associations (not all of which are necessarily positive) with the Russian language which may be contributing to these unpredicted affective evaluations. It may be that such emotionally charged experiences and associations have an overshadowing effect on how the Russian participants regard stimuli instantiated with the Cyrillic alphabet.

Unfortunately, this was not a consideration that was anticipated or explored in this study. Zajonc deliberately used stimuli which were emotionally neutral when testing for the mere exposure effect so as to avoid any earlier emotional associations (i.e., 1968; 1979; 1980). However, it was presumed that the artificial grammar and the subsequent stimulus strings were likewise neutral in affective tone. Since the specific elements used to assemble these strings had traditionally been English consonants (Gordon & Holyoak, 1983; Manza & Bornstein, 1995), this possible emotional wrinkle had not been fully anticipated. This clearly invites further investigation into the roles played by specific stimuli especially since the other symbol sets used in the study were likewise thought to be arbitrary and yet failed to elicit exposure effects.

In summary, this experiment has had both predicted and surprising results. The use of the Cyrillic alphabet presented no difficulties with respect to the grammaticality judgments, regardless of the subject population. Moreover, the structural mere exposure effect was elicited with both English and Russian speakers when using the dependable English consonants to instantiate the grammar. What was a bit surprising given the earlier experiments in this study was that non-Russian subjects demonstrated exposure effects when confronted with the Cyrillic letters. Just why these unusual symbols succeeded where the Chinese and Japanese characters failed is open to speculation and may have something to do with their being somewhat more common in and around the Brooklyn College community. But perhaps the most unexpected finding of all was that of the Russian participants' affective ratings of Cyrillic letters. Evidently there were mitigating factors which somehow interfered with an expression of exposure effects.

Although hardly crisp, the data presented thus far do point to a couple of interesting conclusions. First of all, it has been seen that well-formedness task is relatively robust and does not seem to be seriously compromised by the use of unusual stimuli. And, second, that when investigating the mere exposure effects (both classic and structural) selecting viable stimuli is extremely important and not always obvious.

Experiment 6 - Transfer

Having replicated and discovered something about the boundaries of the structural mere exposure effect, it becomes interesting to look at the extent to which these effects can be generalized even further. This last study is an examination of one of the more intriguing issues that has concerned investigators of implicit learning, transfer. Specifically, if participants have acquired knowledge of the underlying structure of a stimulus display (such as strings generated by an AG), can they then carry this knowledge over to stimuli which are instantiated in superficially different forms? In other words, can the sense of grammaticality gleaned implicitly during an AG study *transfer* to test stimuli comprised of an entirely new set of letters? And, if so, how well do these effects relate to both well-formedness judgments and exposure effects?

Early evidence for such generalizations came from a study by Reber (1969) when he showed that subjects' ability to memorize strings from an AG was not compromised by changing the letter set used to instantiate the grammar. On the other hand, modifications in the underlying rules of the grammar seriously disrupted performance. The many studies examining grammatical sensitivities during transfer since then have tended to use a somewhat different procedure. Altmann, Dienes, and Goode (1995), Dienes and Altmann (1997), Manza and Reber (1997), Mathews, Buss, Stanley, Blanchard-Fields, Cho and Druhan (1989), Vokey and Brooks (1992) have all reported that subjects can distinguish the well-formedness of strings composed of one set of letters after having been trained using strings made up with an entirely different letter set.

Interestingly, subjects in these studies are unaware of the precise mappings between stimulus sets. That is, even though they can reliably classify novel strings on the basis of underlying grammatical coherency, subjects are unable to match the letter-for-letter swapping that characterized the transfer stimulus set (Manza & Reber, 1997).

This transfer effect has also been extended to other stimulus domains. Altmann, Dienes, and Goode (1995) and Manza and Reber (1997) reported cross-modality transfer where learning strings were composed of auditory sequences and test strings were instantiated using visual displays (and vice versa). Such transfer effects further suggest that the subjects' knowledge is not intimately tied to the specific stimuli to which they were originally exposed.

The underlying theme in this overall study has been with the relationship between structural representations and the distribution of elevated affect. In the real world, this may perhaps contribute to an understanding of how such diffuse hedonic tone emerges and merges in a way that might play a critical component of any theory of aesthetics. For example, once people have come to appreciate the nuanced art forms of a particular school of painting, they almost always find themselves enchanted by the work of new artists working within the same conceptual domain—even if they have never encountered their work before. And, likewise, aficianados of a particular composer are not only able to identify new pieces, but are generally predisposed to appreciating the music upon first hearing it.

Experiment 6 looks at the extent to which the (structural) mere exposure effect can transfer acquired affect (SMERE). Do subjects show heightened affect for structured strings when all of the test stimuli are comprised of completely new stimulus items? This question expands the dimension of generalizability. In the earlier experiments, this generalizability was explored by looking at the responses to NewG strings. Now, even the specific elements used to instantiate the grammar in the test strings will be new to the subjects.

Since English consonants have been the reliably successful stimulus elements in producing the structural mere exposure effect, they were used for this experiment. Note that, not only will participants be confronted with novel grammatical sequences in this study, they will also be dealing with the added dimension that all test strings will be composed using a completely new set of letters. In this way, two kinds of novelty are explored, both the physical *and* structural forms are changed, retaining only the consistent rule-based architecture that promotes the structure.

Method

Participants. Thirty-eight Brooklyn College undergraduates participated in this experiment to fulfill a course requirement for introductory psychology. Groups consisted of approximately 10 subjects each.

Stimuli. One of the sets of English consonants used to instantiate the grammar consists of those letters used in Experiment 2. In addition, a second set of letter strings was made up using a different set of consonants. The two letter sets are given in Figure 1 and are the ones used by Manza and Reber (1997). As in the previous experiments, test stimuli were of three kinds, OldG, NewG, and NonG strings. However, in this case, because of the switch to a different letter set during testing, the notion of an "OldG"string needs to be qualified: The strings in this study that are designated as "*Old Grammatical*" are, in fact, "new" to the participants in the sense that they are composed of letters they have not encountered previously in the experiment. They are considered "OldG" strings because they are *structurally identical* to those strings seen in the learning set. For example, a learning string in this study, such as "PTVVS," may become an "OldG" test item incarnated as "CNLLJ." In the same vein, the "NewG" strings in this experiment are the as-yet-unseen grammatical *sequences* which are displayed using the second letter set. So both kinds of novelty are altered in the case of "NewG" strings. Both the actual physical appearances (letter set), as well as the specific grammatical sequences, are novel. Similarly, the "NonG" strings would simply be those NonG strings used in the prior experiment with the basic letter-to-letter substitutions. It should be emphasized that none of the actual learning strings are ever displayed during the testing phase.

Procedure. Participants were randomly assigned so that approximately half of the subjects were trained on Letter Set 1 and tested on Letter Set 2; the other half, the reverse. As in earlier experiments, all learning and test materials were printed in booklets. And, as before, participants were told that this was a basic memory experiment in which they were expected to reproduce strings of letters after a short study period. Again, all subjects were asked for both liking and well-formedness judgments, the order of which was counter-balanced across groups.

After completing both tasks, participants were given a 'matching test' to assess their explicit knowledge of the letter substitutions. The letter sets were printed in random order in two columns. The task was for subjects to draw lines matching up the letters which they thought were congruent.

Results

Well-formedness judgments: A 2 (order) x 2 (letter set) x 3 (stimulus type) ANOVA revealed a highly significant effect of stimulus type ("OldG", "NewG", "NonG") showing that participants were able to make accurate grammatical classifications of the strings which were instantiated with new letters, $F(2, 37) = 14.06$, $MSE = 47$, $p < .0001$. This classic transfer effect replicates the work of several others, specifically, Altmann, Dienes and Goode (1995), Manza and Reber (1997), and Vokey and Brooks (1992). Interestingly, as the top panel of Table 11 indicates, subjects showed a clear sensitivity to both the analogical links between the two letter sets and the grammatical infrastructure. "OldG" strings were liked more than "NewG" strings, indicating that while subjects reliably classified grammatical strings in the new letter set as well-formed, they were more likely to do so when the transfer items were congruent equivalents of strings seen during learning. This effect reflects the formation of "abstract analogies," to use Vokey and Brooks's (1992) term, and is a particularly interesting finding in light of fact that, as is discussed below, there is no evidence that subjects were aware of the letter-to-letter correspondences that prescribed the transfer.

Also note that subjects' ability to classify strings in their novel instantiations is captured entirely by their performance on the well-formed strings ("OldG" and "NewG"). Their ability to detect the syntactic violations in the NonG strings was only at chance level. This result has been observed in other transfer studies, specifically Manza and Reber (1997). Evidently, a transfer in letter set presents an obstacle in detecting the (slight) grammatical violations of the NonG strings.

Confidence: There was a small but significant correlation between reported confidence and the accuracy of well-formedness judgment ($r = .13, p < .01$). And, as seen before, subjects performed essentially at chance on those trials where they were "just guessing," $t\,(37) = .12$, $p >$ 05. There was no significant difference among the mean confidence ratings given to response to the three stimulus types, ("OldG" = 2.15; "NewG" = 2.09; "NonG" = 2.11).

Table 11

Experiment 6

Well-formedness judgments — transfer to novel English letter set

N = 38

Stimulus Type	Probability of saying 'Grammatical'	Mean on 6-pt scale
Old Grammatical	.66	4.18 [a]
New Grammatical	.58	3.90 [b]
Non Grammatical	.51	3.54 [c]

Liking Judgments — transfer to novel English letter set

Stimulus Type	Probability of Score ≥4	Mean on 6-pt scale
Old Grammatical	.58	3.82 [d]
New Grammatical	.58	3.73 [d]
Non Grammatical	.47	3.57 [d]

There were no effects of letter set. Participants responded similarly to either of the letter sets used to instantiate the learning and test strings. Although this result is not surprising at this juncture, it does become relevant when looking at the liking task, discussed below.

There were also no main effects of task order, although it interacted mildly with stimulus type, $F(2, 37) = 4.79$, $MSE = 1.1$, $p < .05$. There were no other differences between stimulus types and the order in which the tasks were performed. Evidently the additional opportunity of evaluating the "NewG" strings may have facilitated the grammatical "sense" of the these items.

Affective ratings: A 2 (order) x 2 (letter set) x 3 (stimulus type) ANOVA of the liking judgments given in the bottom panel of Table 10 revealed a significant effect of stimulus type, indicating that, even though they were instantiated with the new letter set, well-formed strings were liked more than those that violated the structure of the AG [$F(2, 37) = 4.36$, $MSE = 11.5$; $p < .05$]. The structural mere exposure effect which has been observed in earlier experiments using familiar stimuli shows up here even when the surface form of the stimulus displays has been systematically altered.

There is some hesitation in referring to this as simply a "classic" mere exposure effect because in this case even the "OldG" strings were technically being seen for the first time. The analogic correspondence, although quite basic in its one-to-one substitutions, nonetheless relies on structural coherency in the absence of physical re-presentation.

However, there are a number of other interesting features to this finding. Specifically, the order in which the two tasks were run and the specific letter set used to instantiate the grammar both turned out to be playing roles in the affective responses. Looking at each in turn:

Task order: The overall mean affective ratings for *all* three stimulus types were significantly higher when the liking task was run *after* the well-formedness task was completed, $F(1,37) = 4.55$, $MSE = 2.68$; $p < .05$.

The obvious interpretation of this finding is that the additional experience with the stimulus displays gained during the well-formedness task has a positive impact on subjects' liking judgments. While this order effect is the one which both Manza and Bornstein (1995) and Gordon and Holyoak (1983) were wary of, interestingly, it appeared only sporadically throughout this study. Moreover, when there were task order effects in the liking task, they have generally occurred when the AG was represented with relatively unusual elements, as in Experiment 4 (keyboard symbols) and Experiments 5B and 5C (English speaking subjects rating Russian letters, and vice versa). Of course, since there were no exposure effects of any kind found with the Chinese or Japanese characters, task order was not a factor. It is not clear how reliable this particular effect is. One possibility, which is pursued in more detail below, is that it is caused by a generalized elevation of positive affect for *any* stimulus domain to which an individual has had sufficient exposure (Monahan, Murphy & Zajonc, 2000). There were no task order interactions with any other factor.

Letter set: Surprisingly, the specific letters used to instantiate the AG made a difference on subjects' liking judgments. The overall ANOVA of the data in Table 11 reveals a strong main effect of letter set, $F(1,37) = 20.16$, $MSE = 2.63$, $p < .0001$. Subjects who were trained with Letter Set "C, H, J, L, N" and evaluated strings composed with Letter Set "P, S, T, V, X" (noted as "C H J L N → P S T V X" in Table 12) gave systematically higher liking scores to all stimulus types ("OldG", "NewG", "NonG") than those who learned with strings instantiated with "P, S, T, V, X" and judged strings made up using "C, H, J, L, N". It is worth emphasizing that *all three* stimulus types composed with the "P, S, T, V, X" letter set were rated with higher liking scores—even the "NonG" strings. This finding was especially puzzling since these two letter sets neither posed any difficulties with regard to the well-formedness judgments in this experiment, nor did they result in conflicting response patterns when used in earlier transfer studies (Manza & Reber, 1997). The first impulse would be to suspect that it could have something to do with the physical characteristics of the two letter sets. This is unlikely, however, in that the normative data show that the control subjects did not like strings made up using "P, S, T, V, X" significantly more than those made up using "C, H, J, L, N" [$t(125) = 1.81$; n.s.]. However, as seen in Table 12, when evaluating how much they liked specific test items, participants in this experiment evidently were more likely to express exposure effects with test items composed with "P, S, T, V, X" than with "C, H, J, L, N".

Letter matching: Participants were unaware of the letter-to-letter substitutions that underlay the transfer task. As Table 13 shows, only one subject was able to match up all of the corresponding letters successfully, and overall the subjects were no better than chance at this task, $\chi^2(4) = 5.63$, *n.s.* This result replicates what was found by Manza and Reber (1997) in that subjects have the ability to transfer knowledge from one stimulus realm to another while having essentially no conscious knowledge of the shared privileges of occurrence of elements in the two domains–even *after* having run through the entire experimental session (the learning phase and both tasks). It also supports the even more intriguing suggestion that preferences for items can be based on structural properties in the absence of conscious awareness of the specific elements used to instantiate that structure.

Table 12

Experiment 6

Mean liking scores (6-pt scale) from Experiment 5

Transfer stimuli composed with the two different letter sets.

Stimulus Type	PTSVX → CHJLN N = 21	CHJLN → PTSVX N = 17
Old Grammatical	3.66	4.02[a]
New Grammatical	3.50	4.01[a]
Non Grammatical	3.49	3.67[b]
Total	3.55	3.90

Superscripts indicate significant differences as per Dunnett's T @ .05 levels

Table 13

Experiment 6

Number of subjects who correctly matched congruent letters in the transfer study

Number correct	Expected by chance	Observed
5	.31	1
4	0	0
3	3.15	4
2	6.35	10
1	14.25	14
0	13.95	9

Discussion

There are two primary findings from this experiment. First, the results replicate the basic transfer effect in that participants showed an ability to classify grammatical and nongrammatical stimulus items instantiated in a novel letter set in an appropriate manner. As several researchers have noted (Altmann, Dienes & Goode, 1995; Manza & Reber, 1997; Mathews et al., 1989), such transfer effects can be taken as evidence that subjects in these experiments are establishing mental representations that reflect at least some abstract features.

Second, and more interesting, is the corresponding structural mere exposure effect in that participants show increased affective evaluations of novel grammatical strings (both "OldG"and "NewG") instantiated using a physically distinct set of symbols. Importantly, this provides evidence for the generalizability of the mere exposure effect not only to stimuli which share structural similarities, (which has been found elsewhere), but also to test items which differed substantively (iconographically) from items used during the learning phase of the experiment. This is the first demonstration of successful affective transfer using structurally complex stimuli (derived from an artificial grammar).

It is worth noting that the two measures of held knowledge entail different characterizations of the mental representations acquired during the learning phase. The well-formedness data in this experiment show that subjects are better at detecting grammatical strings when they are isomorphs of strings used during learning ("OldG" strings) than when they are novel sequences ("NewG" strings), suggesting the use of a kind of "abstract analogy"operation (Vokey & Brooks, 1992). While this finding implies that the initial inputs are important in establishing mental representations, it is worth emphasizing that subjects carry out this analogical mapping with literally no awareness of the shared privileges of occurrence of the letters in the two instantiations.

On the other hand, the liking data on the structural mere exposure effect imply that subjects are not making evaluative distinctions between strings that are perfect analogs ("OldG" strings) of learning strings and those that are well-structured but entirely novel ("NewG" strings). This latter finding supports a possibility hinted at earlier, in that the structural affordances of a stimulus domain might, on occasion, overwhelm any acquired hedonic aspect of specific stimuli. And, to link this with the issue of aesthetic appreciation hinted at earlier, if you have come to understand and appreciate impressionist painting, you might then be likely to feel positively toward the work of a completely new painter whose work is detectably within the impressionist school, or, of another painting by the same artist which retains certain familiarized features.

General Summary

This series of experiments has successfully demonstrated a generalized mere exposure effect of structured (grammatical) stimuli under certain circumstances. As was the case with Manza and Bornstein and Gordon and Holyoak, the structural mere exposure effect is found when using stimulus elements that are the highly familiar letters of our alphabet. Moreover, there is evidence that, under limited circumstances, the mere exposure effect not only generalizes to newly constructed grammatical strings, but also transfers to such strings that are instantiated in a completely new letter set.

However, when employing unusual elements that are not within the 'prior vocabulary' of the participants, the exposure effects become more elusive. In cases where the elements are extremely unfamiliar (The Japanese Kana of Experiment 1 and Chinese characters of Experiment 3), it becomes difficult to elicit even the classic mere exposure effect. In those cases, even previously presented (OldG strings) did not garner heightened affective ratings. But the story is not as simple as all that. As Experiment 5 suggests, the nature of the stimuli used to test exposure effects plays a very real and relevant role. Russian subjects somehow refrained from endorsing already highly familiar items with greater liking scores. The apparent universality of the mere exposure effect would lead one to expect that this population would be fully able to demonstrate such effects–which they do–under certain conditions (i.e., with the English consonants used in Condition 5C). Mitigating factors, perhaps emotional and/or social, can and do play a role in the ways in which affective evaluations are prompted, the extent and limitations of which need to be further explored to be understood.

What can safely be said from this set of experiments is that the structural mere exposure effect can be quite temperamental. On the one hand, the trends that were found involve a certain level of *a priori* familiarity with the elements used to instantiate the grammar. It appears that highly unusual or unfamiliar elements have a tendency to derail the exposure effects. Perhaps the message is that it is best not to complicate the task by employing stimuli which might compete with the structural integrity of the strings. Moreover, even when using highly familiar letters, it is apparent in the transfer study (Experiment 6) that some letter sets work better than others in eliciting these effects. Of course, other than the normative data collected on these letter sets, there was no empirical base for predicting how one letter set might perform versus another in a straightforward (non-transfer) implicit learning study. Just such a follow-up experiment seems to suggest itself by these results.

These data suggest that well-formedness judgments are, by and large, more reliable and robust in indicating grammatical discriminations. Manza and Bornstein's (1995) view of the mere exposure effect as a more sensitive measure of familiarity is therefore not supported by this research. Affective judgments are perhaps a purer indicator of implicitly acquired knowledge of the grammar specifically because the liking judgments do not require highlighting the existence of the grammar. But in terms of sensitivity, it appears that well-formedness judgments more reliably reveal acquired grammatical discriminations across the entire range of stimuli used here

The confidence ratings provided by the participants in this study were not systematically correlated with the accuracy of their well-formedness judgments.

Although the temptation is to say that subjects' confidence ratings were adequately reflecting task performance in those cases where familiar items were used, (as in Experiments 2 and 4 with English consonants and keyboard symbols), this also appeared in the Russian study in three of the four conditions. At best, this unreliable effect reflects what Dienes and Berry (1997) promoted as evidence for, at best, a limited awareness of explicit knowledge about the grammar. These results also provide evidence for Chan's (1992) 'zero-correlation' criterion which indicates that subjects are not always aware of when it is they know something.

In addition, in practically all of the experiments in this study, participants' accuracy was not significantly different from chance when subjects claimed to be just guessing (confidence ratings of '1'). This (disappointing) finding fails to support the 'guessing-criterion'. Actually, the only situation in which confidence ratings were correlated with classification accuracy, *and* where total guesses were at above-chance levels was in Condition A of Experiment 5. This involved English speaking subjects rating English consonants--an exact replication of Experiment 2 (where this was *not* the finding). So although this hints at some level of meta-cognitive insight, it is, at best, quite tenuous. It is reasonable to assert that the well-formedness task in these, and other AG studies, appears to recruit a fair mix of implicit and explicit resources.

The matter of task order was a tricky one to parse. With regard to the accuracy of the well-formedness judgments, on the whole, it was usually in the direction suggesting that well-formedness judgments were relatively more accurate when participants first had experienced the liking task. This, in itself, is not particularly surprising because any additional experience with the stimulus strings would likely increase their grammatical saliency.

It was the effect of task order on the affective ratings that were of particular interest here. Earlier researchers were concerned that by first participating in grammatical classifications, subjects would be cued in to the dynamic of the experiment. For this very reason, Gordon and Holyoak (1983) and Manza and Bornstein (1995) used between-subject designs. Whereas their concerns were that explicit knowledge of the existence of the grammar might influence subsequent affective ratings, this was not consistently found to be true in the current study. The only occasions in which a task order effect was found was (somewhat) in Experiment 4 (keyboard symbols), and (particularly) in Experiment 6 (transfer study), where it appeared that the additional experience of first making grammaticality judgments enhanced the overall ratings of *all* test items. This enhanced hedonic tone applied to all test strings–including the "NonG" strings. This phenomenon of overall increased affect generalizing for all test items has been noted and explored by Monahan, Murphy and Zajonc (2000).

The findings in this study are in line with the caveat of Whittlesea and colleagues when they caution to be wary of drawing strong conclusions about the underlying nature of memorial representations on the basis of limited data bases (see especially Whittlesea & Wright, 1997). Considerations of function and the demand characteristics of particular tasks are likely having an impact. Tacitly held knowledge about a stimulus domain may be differentially applied depending upon the nature of the probe questions and the perception of the task demand (Neal & Hesketh, 1997).

Conclusions

While the data base is still small, certain conclusions, however tentative, can legitimately be reached. Specifically:

1. The classic mere exposure effect is real and emerges even when the stimulus display consists of highly structured, complex stimulus sequences. However, when participants are presented with structured stimuli that are identical to those experienced during learning, they tend to like them only somewhat better than similarly well-formed, yet novel, stimuli (OldG vs. NewG strings). This result is an important extension of the classic findings of Zajonc and many others (see Bornstein, 1989 for a review and meta-analysis). Experiment 4 (keyboard symbols) is the first to demonstrate the classic mere exposure effect using strings of structured symbols (non-letters), and Experiment 2 (English consonants) appears to be the first account of successfully uncovering *both* the classic and structural mere exposure effects using multiple structured items as stimuli. It would be interesting to look at strings of letters, or other types of serial stimuli, which are specifically *not* derived from a structured AG, and are devoid of *any pattern* – to see whether the classic mere exposure effect can readily be found.

2. The structural mere exposure effect is also real. It may be a bit temperamental, but it can be found. Participants who work with complexly structured

displays to the point where they develop a reasonably robust representation of that domain will, under the right circumstances, display enhanced affective appreciation of novel stimuli that are consistent with that structure. This result replicates and extends the work of others, specifically Gordon and Holyoak (1983) and Manza and Bornstein (1995).

3. With both kinds of exposure effects (classic and structural), the familiarity of the superficial aspects of the stimuli appears to play a critical role. There are at least two distinct senses of familiarity operating and they appear to contribute to the effects in different ways. Stimuli can be familiar in a personal, *a priori* manner simply because participants have frequently encountered them in the course of day-to-day living (i.e., English consonants, keyboard symbols), or, stimuli can also be rendered familiar through mere repeated presentations. This form of induced familiarity is the technique traditionally used in the literature on the mere exposure effect. And, as scores of studies have shown, subjects come to have increased liking for once novel, hedonically neutral items after they have be rendered "familiar" through simple exposure.

However, this form of induced familiarity is not necessarily sufficient to produce either the classic or the structural exposure effects with stimuli created with an AG. When confronted with highly unusual displays (e.g., Japanese or Chinese characters) subjects showed no heightened appreciation for *any* of the items, not even for those that had been part of the learning set and which were seen several times in the first phase of the experiment. It seems that in order for participants to develop heightened appreciation for complex *structured* displays, those displays have to be instantiated using symbols with which they have had considerable *a priori* experience. Interestingly, in the one experiment that used moderately familiar stimuli (Experiment 4 - keyboard symbols), there was evidence for the *classic* mere exposure effect--but not for the more elusive *structural* mere exposure effect.

Perhaps the pronounceability of the stimuli makes a difference in whether or not structural exposure effects are found. In addition to a high level of familiarity found with English letters, they are subvocally pronounceable and thereby rehearsable. Whereas the keyboard symbols are quite commonplace, they are somewhat difficult to label. Of course, the natural assumption, then, would be that the Russian participants in Experiment 5D (Cyrillic letters) should have displayed the structural mere exposure effect, since those subjects were well acquainted with the Cyrillic alphabet. The results to the contrary were surprising, which leads one to suspect that there were also other, socially linked issues at play.

Note that these two distinct senses of "familiarity" do not seem to manifest themselves when subjects are asked to make grammaticality judgments. When working with the exact same stimulus displays, subjects are quite adept at distinguishing between strings that are well-formed and those that violate aspects of the structure even when they show no particular shifts in aesthetic appreciation of these items. This appears to be the case for all groups of subjects--regardless of the nature of the stimuli used.

Moreover, there is at least some evidence to reinforce the notion that liking judgments tap into an emotionally-laden expression, as may have been the case with the

Russian participants. If this is so, then the nature of the *a priori* experiences also becomes a contributing factor in the expression of the structural mere exposure effect, which again is why Zajonc and others took precautions to work with what he considered to be evaluatively-neutral items when investigating the mere exposure effect.

It would be premature to conclude that *a priori* familiarity is an absolute prerequisite for the structural mere exposure effect, only that the amount and kind of experience subjects have had with particular stimulus displays play an important role in its emergence. Certainly, different kinds of experiences, degrees of familiarity, and associations will produce different outcomes.

4. The structural mere exposure effect transfers to different stimulus domains (SMERE). Subjects show a small but reliable increased liking for structurally coherent items even when the superficial form of the stimulus displays has been changed. Interestingly, this effect appears to be entirely dependent on the underlying structure of the displays. As seen in Experiment 6, participants liked the "NewG" strings as much as they did the "OldG" strings suggesting that, at least under these conditions, structural integrity can overwhelm any residual effect of an instantiated analogous representation. Clearly, further work in this area needs to be done to explore the contributions of specific stimulus sets, and the circumstances under which such an affective transfer can occur. In some senses, this is one of the more exciting areas of implicit learning and exposure effects in that real-world investigations, insofar as learning, stereotyping, advertising, and cultural trends can be explored through this circuit.

5. It is worth noting that relatively little evidence was found to support the notion that knowledge of the structural characteristics of the displays had an impact on subjects' liking judgments. Moreover, when this did emerge, the pattern appeared to reflect the tonic elevation effect whereby repeated exposures to particular stimuli produce an overall increase in "tonic mood" which is then applied to all and other stimuli (Monahan, Murphy & Zajonc; 2000). In this case, the additional exposure to the stimulus displays that resulted from participation in the well-formedness task occasionally produced an overall increase in liking for *all* of the stimuli, regardless of their grammaticality. It may not, therefore, be that first participating in the well-formedness task provides additional grammatical information; but rather that the increase in affective evaluations might be due to the general heightened sense of well-being which is induced by yet another re-presentation of the stimuli.

Although not what was expected, this pattern of elevated tonic mood is actually quite consistent with what has most recently been reported in the literature regarding the mere exposure effect. In their systematic investigation of the exposure effects of re-presenting single items to subjects, Monahan, Murphy and Zajonc (2000) did find that the resulting increased tonic mood generalized across all test items. The effect was most pronounced for those items which were members of the same category as those originally presented, but there was very little difference between the ratings of the 'old stimuli' and 'new stimuli'. Although theirs was an experiment involving the subliminal presentation of conspicuously discriminable items, the pattern of elevated affect resembles some of the results found in this study. Monahan et. al.'s experiments used test stimuli that were: 1) members of the original set that were subliminally presented ["OldG strings"]; 2) members of the same category as those originally presented ["NewG strings"]; and 3) members of an entirely different category ["NonG strings"]. The results of the current study parallel their findings in that all members of the same category, (e.g., both the "OldG" and "NewG" strings) did have a tendency to receive very similar affective ratings in some cases (especially Experiment 6 - transfer). This pattern of findings is interesting especially since the stimuli used in this study were differentiated by very subtle distinctions. To extend the analogy, then, category membership in the present study might be understood in terms of grammaticality, a discrimination which was demonstrated reliably via the well-formedness task in all of the experiments reported herein.

Although generalizing to all categories, the heightened affect appears to be strongest for those members (both old and new) of the target stimulus group. This is consistent with the tendency for prototypical items to garner greater affect (e.g., Gordon & Holyoak, 1983). If prototypicality tends to increase with familiarity, and familiarity can be induced through repetition (Whitfield & Slatter, 1979), then it should come as no surprise then, that the inherent pleasantness experienced through (repeated) exposure to (familiar) items in the absence of adverse associations could elevate one's general disposition (Zajonc; 2000). The notion of a generalized increase in tonic disposition as a result of repeated exposure does not discount the mere exposure effect. Studies which have traditionally investigated the mere exposure effect had not accounted for the extent to which non-target items were similar to those previously exposed. This rather serendipitous finding may help to explain the mechanism behind the effect. This increase in generalized affect might even play a hand in the ways in which we develop our aesthetic tastes. Certainly advertisers and fashion moguls count on our ability to 'pick up on' pervasive styles and evolving trends.

The mere exposure effect appears to be but one manifestation of implicitly acquired information. The ways in which the classic mere exposure effect has been studied map nicely onto some of the ways in which implicit learning is examined. In the case of the mere exposure effect, the link between repeated exposures (either sub- or supraliminally presented) is never explicitly highlighted. In the AG paradigm, the inherent structure of the stimulus strings is never quite identified. In both cases, subjects appear to rely on some sort of nonconscious processing, done at some point during the task, in order to be able to respond appropriately.

We tend to like that which is familiar, as the mere exposure effect demonstrates; and, in some cases, this familiarity need not be a simple case of item identification. As seen with the structural mere exposure effect, it appears that the consistency of the structure of an item may be sufficient to render that item as being 'familiar enough' to be worthy of generalized heightened affect. These results speak to the issues surrounding the circumstances under which the structural mere exposure effect might be elicited. The effect appears to be somewhat elusive, and unlike the classic mere exposure effect, the structural mere exposure effect relies on the extraction and enhanced appreciation of complex and structurally coherent items.

The long-range implications of these experiments yield empirical evidence where it is currently missing, specifically at the intersection of the mere exposure effect, implicit learning of complexly structured domains, and the role of familiarity in the development of preferences for particular stimulus domains over others. One reason why there has been so little work on these issues comes, no doubt, from the fact that they have traditionally been viewed as belonging to different sub-areas within psychology and not recognized as having a common conceptual base. Hopefully, this study has helped to strengthen links which evidently exist between these fields.

Appendix A

Test strings used in all experiments

(Instantiated with English consonants for illustrative purposes)

Old Grammatical	New Grammatical	Non-Grammatical
PTVPS	TXXVV	PTTPS
PTTVV	PVPXTTVV	SXXVPS
PTVV	TSXXVV	TXPV
TXS	TSXXVPS	TXV
TSSSXS	TSSXVV	TTVV
PVPXVV	TSSXXVPS	SVPXTVV
TXXTVPS	TSSXS	PVXPVXPX
PVV	TXXTTVV	PXPVXVTT
PVPXVPS	PTTVPS	PTVPPPS
PTVPXVV	TSSSXXVV	PVTTTVV
PVPXTVPS	PTVPXTVV	TSSXXVSS
TXXTVV	PVPS	VPXTVV

Appendix B

Sample Protocol used in well-formedness task

1. PTTPS

   Yes No 1 2 3

2. SXXVPS

   Yes No 1 2 3

3. TXXVV

   Yes No 1 2 3

4. PTVPS

   Yes No 1 2 3

5. TXPV

   Yes No 1 2 3

6. TXV

   Yes No 1 2 3

07. TTVV

   Yes No 1 2 3

08. PTTVV

   Yes No 1 2 3

09. PTVV

   Yes No 1 2 3

10. PVPXTTVV

   Yes No 1 2 3

11. TSXXVV

   Yes No 1 2 3

12. TXS

   Yes No 1 2 3

Appendix C

Sample Protocol used in liking task

1. TXS

   1 2 3 4 5 6

2. TSSXXVV

   1 2 3 4 5 6

3. TXXTVPS

   1 2 3 4 5 6

4. PVTTTVV

   1 2 3 4 5 6

5. TXXTVV

   1 2 3 4 5 6

6. PTVV

   1 2 3 4 5 6

07. TXPV

   1 2 3 4 5 6

08. TSXXVPS

   1 2 3 4 5 6

09. PVXPVXPX

   1 2 3 4 5 6

10. TSSXXVPS

   1 2 3 4 5 6

11. TXXTTVV

   1 2 3 4 5 6

12. PTTPS

   1 2 3 4 5 6

References

Allen, R., & Reber, A. S. (1998) Unconscious intelligence. In W. Bechtel & G. Graham (Eds.) *A companion to cognitive science*. Malden, MA, Blackwell.

Alpert, R. (1953). Perceptual determinants of affect. Unpublished MA thesis, Wesleyan University.

Altmann, G.T.M., Dienes, Z., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 899-912.

Baeyens, F., Eelen, P., & Van den Bergh, O. (1990). Contingency awareness in evaluative conditioning: A case for unaware affective-evaluative learning. *Cognition and Emotion, 4*, 3-18.

Becknell, J. C., Wilson, W. R., & Baird, J. C. (1963). The effect of frequency of presentation on the choice of nonsense syllables. *Journal of Psychology, 56*, 165-170.

Berlyne, D. E. (1950). Novelty and curiosity as determinants of exploratory behavior. *British Journal of Psychology, 41*, 68-80.

Berlyne, D. E. (1970). Novelty, complexity and hedonic value. *Perception and Psychophysics, 8*, 279-286.

Berry, D. C., Broadbent, D. E. (1984). On the relationship between task performance and associated verbalizable knowledge. *Quarterly Journal of Experimental Psychology, 36*, 209-231.

Bonanno, G. A., Stillings, N. A. (1986). Preference, familiarity, and recognition after repeated brief exposures to geometric shapes. *American Journal of Psychology, 99*, 403-415.

Bornstein, R. F. (1989). Exposure and affect: Overview and meta-analysis of research, 1968-1987. *Psychological Bulletin, 106*, 265-289.

Bornstein, R. F., & D'Agostino, P. R. (1994). The attribution and discounting of perceptual fluency: preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. *Social Cognition, 12*, 2, 103-128.

Bornstein, R. F. (1992). Subliminal Mere Exposure Effects. In *Perception without awareness: Cognitive, clinical and social perspectives.* Bornstein, R. R., & Pittman, T. S. (Eds.) The Guilford Press, New York.

Broadbent, D. E., FitzGerald, P., & Broadbent, M. H. P. (1986). Implicit and explicit knowledge in the control of complex systems. *British Journal of Psychology, 77,* 33-55.

Brooks, L. R., (1978) Nonanalytic concept formation and memory for instances. In. E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization.* New York: Wiley.

Brooks, L .R., & Vokey, J. R. (1991). Abstract analogies and abstracted grammars: Comments on Reber and Mathews et al. (1989). *Journal of Experimental Psychology: General, 120,* 316-323.

Chan, C. (1992). Implicit cognitive processes: theoretical issues and applications in computer systems design. Unpublished dissertation, University of Oxford.

Cheesman, J., Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology, 40,* 4, 342-367.

Cleeremans, A., McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General, 120,* 235-253.

DeHouwer, J., Baeyens, F., & Hendrickx, H. (1997). Implicit learning of evaluative associations. *Psychologica Belgica, 37,* 115-130.

Dienes, Z. & Altmann, G. (1997). Transfer of implicit knowledge across domains: How implicit and how abstract? In Berry, D. C. (Ed.) *How implicit is implicit learning?* (pp. 107-123).

Dienes, Z., Altmann, G.T.M., Kwan, L., & Goode, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. Unpublished manuscript, University of Sussex.

Dienes, Z., & Berry, D. (1997). Implicit learning: Below the subjective threshold? *Psychonomic Bulletin and Review, 4,* 3-23.

Dienes, Z., Broadbent, D., & Berry, D. (1991). Implicit and explicit knowledge bases in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory and Cognition, 17,* 5, 875-887.

Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1984). A case of syntactical learning and judgment: How conscious and how abstract? *Journal of Experimental Psychology: General, 113,* 541-555.

Elliot, R., & Dolan, R. J. (1998). Neural response during preference and memory judgments for subliminally presented stimuli: A functional neuroimaging study. *The Journal of Neuroscience, 18,* 12, 4697-4704.

Ellis, H. D. (1986). Processes underlying face recognition. In R. Bruyer (Ed.), *The neuropsychology of face perception and facial expression.* Hillsdale, NJ: Erlbaum.

Erdelyi, M.H. (1985). *Psychoanalysis: Freud's cognitive psychology.* New York: Freeman & Co.

Fowler, C., Wolford, G., Slade, R., & Tassinary, L. (1981). Lexical access with and without awareness. *Journal of Experimental Psychology: General, 110,* 341-362.

Gordon, P. C., Holyoak, K. J. (1983). Implicit learning and generalization of the "mere exposure" effect. *Journal of Personality and Social Psychology, 45,* 3, 492-500.

Heyduk, R. G. (1975). Rated preference for musical compositions as it relates to preference and exposure frequency. *Perception and Psychophysics, 17,* 84-91.

Harrison, A.A. (1977). *Advances in experimental social psychology. Vol. 10.* New York Academic Press. L. Berkowitz, Ed.

Hasher, L., Zacks, R.T. (1979) Automatic and effortful processes in memory. *Journal of Experimental Psychology: General.* 108, 356-388.

Holendar, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral and Brain Sciences, 9,* 1-66.

Hsaio, A. & Reber, R. S. (2001). The dual-task SRT procedure: Fine-tuning the timing. *Psychonomic Bulletin & Review.* 8 (2), 336-342.

Jacoby, L., (1983). Perceptual enhancement: Persistent effects of an experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 21-38.

Jacoby, L. (1991). A process dissociation framework: Separating automatic from intentional use of memory. *Journal of Memory and Language, 30,* 513-541.

Jacoby, L., Kelley, C., Brown, J., Jasechko, J. (1989). Becoming Famous Overnight: Limits on the Ability to Avoid Unconscious Influences of the Past. *Journal of Personality and Social Psychology, 56,* 326-338.

Jacoby, L. & Witherspoon D. (1982). Remembering without awareness. *Canadian Journal of Psychology, 36,* 300-324.

Johnson, M. K., Kim, J. K., & Risse, G. (1985). Do alcoholic Korsakoff's Syndrome patients acquire affective reactions? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 22-36.

Johnson, R. C., Thomson, C. W. & Frincke, G. (1960). Word values, word frequencies, and visual duration thresholds. *Psychological Review, 67,* 332-342.

Kagan, J. (1971). *Change and continuity in infancy.* New York: John Wiley.

Knowlton, B. J., & Squire, L. R. (1994).The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1, 79-91.

Kunst-Wilson, W. R., Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science, 207,* 557-558.

Kurucz, J., Feldmar, G., & Werner, W. (1979). Prosopo-affective agnosia associated with chronic organic brain syndrome. *Journal of the American Geriatrics Society, 27,* 91-95.

Lazarus, R. S. (1982). Thoughts on the relations between emotion and cognition. *American Psychologist, 37,* 9, 1019-1024.

Lazarus, R. S. (1984). On the primacy of cognition. *American Psychologist, 39,* 2, 124-129.

Lewicki, P. (1986a). Processing information about covariations that cannot be articulated. *Journal of Personality and Social Psychology: Learning, Memory, and Cognition,12,* 135-146.

Lewicki, P. (1986b). *Nonconscious social information processing.* New York: Academic Press.

Mandler, G, Nakamura, Y., Shebo Van Zandt, B. J. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13,* 4, 646-648.

Manza, L. (1998). The structural mere exposure effect. Paper presented at the Annual Meeting of the Eastern Psychological Association, Boston, MA.

Manza, L., & Bornstein, R.F. (1995). Affective discrimination and the implicit learning process. *Consciousness and Cognition, 4,* 399-409.

Manza, L., & Skypala, D. (1994) Attention and artificial grammar learning: interdependence of implicit and explicit knowledge system. Unpublished manuscript. Gettysburg College.

Manza, L., & Reber, A. S. (1997) Representing artificial grammars: Transfer across stimulus forms and modalities. In *How implicit is implicit learning?* Berry, D. C. Oxford University Press.

Maslow, A. H. (1937). The influence of familiarization on preference. *Journal of Experimental Psychology, 21,* 162-180.

Mathews, R. C. (1997). Is research painting a biased picture of implicit learning? The dangers of methodological purity in scientific debate. (1997). *Psychonomic Bulletin & Review, 4, (1),* 38-42.

Mathews, R. C., Buss, R. R., Stanley, W. B., Blanchard-Fields, F., Cho, J. R., & Druhan, B. (1989). Role of implicit and explicit processes in learning from examples: a synergistic effect. *Journal of Experimental Psychology, 15,* 6, 1083-1100.

McAndrews, M. P., & Moscovitch, M. (1985). Rule-based and exemplar-based classification in artificial grammar learning. *Memory & Cognition, 13,* 469-475.

Macmillan, N. A. & Creelman, C. D. (1991) *Detection theory: A user's guide.* Cambridge University Press.

Merikle, P. M. (1992). Perception without awareness. *American Psychologist, 47,* 6, 792-795.

Merikle, P. M., & Reingold, E. M. (1991). Comparing direct (explicit) and indirect (implicit) measures to study unconscious memory. *Journal of Experimental Psychology: Learning, Memory and Cognition, 17,* 2, 224-233.

Monahan, J. L., Murphy, S. T., & Zajonc, R. B. (2000). Subliminal mere exposure: specific, general and diffuse effects. *Psychological Science, 11,* 6, 462-466.

Moreland, R. L.& Zajonc, R.B. (1977). Is stimulus recognition a necessary condition for the occurrence of exposure effects? *Journal of Personality and Social Psychology, 35,* 4, 191-199.

Murphy, S. T., Monahan, J. L., & Zajonc, R. B. (1995). Additivity of nonconscious affect: Combined effects of priming and exposure. *Journal of Personality and Social Psychology, 69,* 589-602.

Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology, 64,* 723-739.

Neal, A., & Hesketh, B. (1997). Episodic knowledge and implicit learning. *Psychonomic Bulletin and Review, 4,* 24-37.

Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: evidence from performance measures. *Cognitive Psychology,19,*1-32.

Pepper, S. C. (1919). Changes in appreciation for color combinations. *Psychological Review, 26,* 389-396.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General,119,* 264-275.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77,* 353-363.

Reber, A. S., (1967). Implicit learning of artificial grammars. *Journal of Experimental Psychology, 77,* 353-363.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General. 118,* 3, 219-235.

Reber, A.S. (1993). *Implicit Learning and Tacit Knowledge: An Essay on the Cognitive Unconscious.* Oxford University Press.

Reber, A. S., & Allen, R. (1978). Analogy and abstraction strategies in synthetic grammar learning: a functionalist interpretation. *Cognition, 6,* 189-221.

Reber, A.S., Allen, R. & Regan, S. (1985). Syntactical learning and judgment, still unconscious and still abstract: Comment on Dulany, Carlson, & Dewey. *Journal of Experimental Psychology: General, 114,* 17-24.

Reber, A. S., Allen, R. & Reber, P. J. (2000). Implicit versus explicit learning. In R. Sternberg (Ed.), *The nature of cognition.* (pp. 475-513). Cambridge, MA: MIT Press.

Reber, A. S. & Lewis, S. (1977). Toward a theory of implicit learning: The analysis of the form and structure of a body of tact knowledge. *Cognition, 5,* 333-361.

Reber, A. S., & Winter, W. (1994) What manner of mind is this? Commentary on Shanks & St. John: Dissociable learning systems. *Behavior and Brain Sciences, 17,* 418-419.

Reber, R., Winkielman, P. and Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science. 9,* 1, 45-48.

Reingold, E. M., & Merikle, P. M. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics, 44,* 6, 563-575.

Saegert, S. C., & Jellison J. M. (1970). Effects of initial level of response competition and frequency of exposure on liking and exploratory behavior. *Journal of Personality and Social Psychology, 16,* 234-242.

Seamon, J. G., Brody, N. & Kauff, D. M. (1983). Affective discrimination of stimuli that are not recognized:II. Effect of delay between study and test. *Bulletin of the Psychonomic Society, 21,* 187-189.

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 501-518.

Schacter, D. L., & Graf, P. (1986). Preserved memory in amnesic patients: Perspectives from research on direct priming. *Journal of Clinical and Experimental Neuropsychology, 8,* 727-743.

Schacter, D. L., Delaney, S. M., & Cooper, L. A. (1990). Implicit memory for unfamiliar objects depends on access to structural descriptions. *Journal of Experimental Psychology: General, 119,* 1, 5-24.

Seamon, J. G., McKenna, P. A., Binder, N. (1998). The mere exposure effect is differentially sensitive to different judgment tasks. *Consciousness and Cognition, 7,* 85-102.

Servan-Schreiber, E., & Anderson, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16, 592-608.*

Shanks, D. R., & St. John, M.F. (1994) Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367-448.

Small, M.Y. (1990) *Cognitive Development*, Harcourt, Brace, Javonivich. San Diego.

Spence, D., & Holland, (1962). The restricting effects of awareness: a paradox and an explanation. *Journal of Abnormal & Social Psychology. Vol. 64*, 163-174.

Stadler, M.A., & Frensch, P.A. (1994) Whither learning, whither memory? Commentary: Shanks & St. John: Dissociable learning systems. *Behavioral and Brain Sciences, 17,* pp. 423-424.

Stang, D. J. (1974). Methodological factors in mere exposure research. *Psychological Bulletin. 81,* 1014-1025.

Stang, D. J., & O'Connell, E. J. (1974). The computer as experimenter in social psychology research. *Behavior Research Methods and Instrumentation, 6,* 223-231.

Vokey, J. R., & Brooks, L. R., (1992). Salience of item knowledge in learning artificial grammars. *Journal of Experimental Psychology: Learning, Memory and Cognition. 18,* 2, 328-344.

Whitfield, T. W. A. & Slatter, P. E.(1979). The effects of categorization and prototypicality on aesthetic choice in a furniture selection task. *British Journal of Psychology, 70,* 65-75.

Whittlesea, B. W. A (1993). Illusions of Familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition. Vol. 19*, 6, 1235-1253.

Whittlesea, B. W. A., & Dorken, M. D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General. 122*, 227-48.

Whittlesea, B. W. A. & Dorken, M. D. (1997). Implicit learning: Indirect, not unconscious. *Psychonomic Bulletin & Review. 4,* (1) 65-67.

Whittlesea, B. W. A.; Wright, R.L. (1997). Implicit (and explicit) learning: acting adaptively without knowing the consequences. *Journal of Experimental Psychology: Learning, Memory and Cognition, 23,* 181-200.

Willingham, D. B. (1997) Implicit and explicit memory do not differ in flexibility: Reply to Dienes and Berry. *Psychonomic Bulletin and Review, 4,* 3-23.

Wilson, W. R. (1979). Feeling more than we can know: Exposure effects without learning. *Journal of Personality and Social Psychology, 37,* 811-821.

Wilson, W. R., & Becknell, J. C. (1961). The relation between the association value, pronounciability, and affectively of nonsense syllables. *Journal of Psychology, 52,* 47-49.

Winkielman, P., Zajonc, R..B., & Schwarz, N. (1997). Subliminal affective priming resists attributional interventions. *Cognition and Emotion, 11,* 4, 433-465.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology Monograph Supplement, 9,* 2, Part 2.

Zajonc, R. B., (1980). Feeling and thinking: preferences need no inferences. *American Psychologist. 35,* 151-175.

Zajonc, R. B. (1984). On the primacy of affect. *American Psychologist, 39,* 2,117-123.

Zajonc, R. B. (2001) Mere exposure: a gateway to the subliminal. *Current Directions in Psychological Science. 10,* 6, 224-228.

Zajonc, R. B., Markus, H., & Wilson, W.R. (1974). Exposure effects and associative learning. *Journal of Experimental Social Psychology, 10,* 248-263.

Zajonc, R.B., Shaver, P., Tavris, C., & Van Kreveld, D. (1972). Exposure, satiation, and stimulus discriminability. *Journal of Personality and Social Psychology, 21,* 270-280.