ELSEVIER

Acta Psychologica 98 (1998) 141–165

acta psychologica

Why do strangers feel familiar, but friends don't? A discrepancy-attribution account of feelings of familiarity

Bruce W.A. Whittlesea [*], Lisa D. Williams

Department of Psychology, Simon Fraser University, Burnaby, B.C., Canada V5A 1S6

Abstract

Recent articles on familiarity (e.g. Whittlesea, B.W.A., 1993. Journal of Experimental Psychology 19, 1235) have argued that the feeling of familiarity is produced by unconscious attribution of fluent processing to a source in the past. In this article, we refine that notion: We argue that it is not fluency per se, but rather fluent processing occurring under unexpected circumstances that produces the feeling. We demonstrate cases in which moderately fluent processing produces more familiarity than does highly fluent processing, at least when the former is surprising.

PsycINFO classification: 2300; 2340; 2343

Keywords: Remembering; Familiarity; Fluency; Attribution

1. Introduction

Intuitively, the feeling of familiarity depends on possession of a trace in memory, a trace laid down at the time one experienced the object of the feeling. [1] A re-encoun-

[*] Corresponding author. E-mail: bruce_whittlesea@sfu.ca.

[1] The term "familiarity" has a number of common senses. One is that a person has actually encountered a stimulus (or even one like it) previously. This sense pertains to the historical fact that a person has previous experience with an object, whether or not that experience influences current behavior and whether or not the person can report that experience. A second is that the person has knowledge about a stimulus that permits them to perform appropriately toward an object, without necessarily having an accompanying feeling of having experienced that stimulus previously. For example, in watching *Hamlet* for the 15th time, I know what to expect is coming next, but I have no pressing feeling of having seen it

PII S 0 0 0 1 - 6 9 1 8 (9 7) 0 0 0 4 0 - 1

ter with that object serves to activate the trace; the feeling of familiarity is the conscious perception of the resonance of that activated trace. According to this idea, the possession of a trace of the event is a necessary and sufficient cause of the feeling of familiarity.

This intuitive notion of familiarity was discredited by Jacoby and his associates (cf. Jacoby et al., 1989a, for a review). Jacoby and Dallas (1981) observed that latency of identification and recognition of test stimuli were correlated. That is, their subjects tended to recognize those stimuli that they could most easily identify in a tachistoscopic presentation. They concluded that people use a "fluency heuristic" to perform recognition. The idea was that prior presentation of a stimulus would facilitate naming of that stimulus on a later occasion, including during the recognition test: People could correctly judge that item to be old by attributing their fluency in naming the stimulus to a prior experience of the item. This act of attribution was thought to be unconscious; consciously, the person would only experience a feeling of familiarity.

There is now a great deal of direct experimental evidence that the fluency of processing influences the feeling of remembering, both in recognition (e.g., Jacoby and Whitehouse, 1989) and in recall (e.g., Roediger and McDermott, 1995). When the fluency of processing a novel test item is artificially enhanced through manipulation of perceptual (Whittlesea et al., 1990; Lindsay and Kelley, 1996) or semantic (Lindsay and Read, 1994) properties, people experience an illusion of remembering. For example, Whittlesea (1993) presented old and new test items within sentences. Some sentence stems were fairly predictive of the target word ("The stormy seas tossed the BOAT"); others were less so ("She saved her money and bought a LAMP"). Subjects named old words about 100 ms faster than new words, showing an effect of prior experience on the fluency of processing; they also claimed to recognize old words 16% more often than new words. More important, the test context also influenced processing fluency: Target words in predictive contexts were named about 130 ms faster than words in nonpredictive contexts. Moreover, subjects claimed to recognize novel targets 18% more often when they were predicted by the context than when they were not. That is, they suffered an illusion of remembering, due to the enhanced fluency of processing induced by the predictive context. It was concluded that the feeling of familiarity occurs when fluent performance is unconsciously attributed, rightly or wrongly, to a source in the past.

Despite the apparent force of such demonstrations, the idea that familiarity is the conscious feeling accompanying fluent processing left a number of questions. One is why there is no pressing feeling of familiarity when encountering well-known people. For example, encountering one's spouse in one's own kitchen must sponsor some of the most fluent processing that one can perform, yet the encounter is not accompanied by a feeling of familiarity. One KNOWS who that person is; but there is no feel-

Foot note 1 (*continued*)
before. A third is the subjective feeling of having encountered a stimulus on some previous occasion, whether one actually has or not. For example, I encounter a person at a bus-stop; I cannot name the person, or say from where I know her, but I'm sure I've seen her before; and I may be right or wrong. It is this subjective feeling, and its source, that we will discuss in this article.

ing of "I've seen you before". In contrast, people who one knows only moderately well, and who are presumably processed less fluently, do sponsor feelings of familiarity. For example, if one encounters the clerk from the corner store when riding a bus, the feeling of familiarity can be very strong.

These examples suggest that fluency of processing is actually inversely related to the feeling of familiarity. However, the examples are incomplete. If instead you encounter the clerk in the store, you then fail to experience any feeling of familiarity: You simply know who that person is, and begin to conduct appropriate business. In contrast, if you suddenly encounter your spouse in an unexpected place, for example sitting in the middle of your lecture class, then the subjective reaction can be very powerful indeed.

These examples suggest that the occurrence of feelings of familiarity are indeed linked to processing fluency in some way, but indirectly, mediated through the context in which the stimulus is encountered. More specifically, they seem to suggest that, to experience a feeling of familiarity, one must be surprised by the fluency of one's processing: If a stimulus is fluently processed in a context in which it is expected, it will not produce a feeling of familiarity.

This suggestion is nicely illustrated by an experiment conducted by a student in our lab. Cook (1996) required her subjects to memorize 80 words, in anticipation of a memory test. At test, she gave subjects a story, consisting of 10 paragraphs. Each paragraph contained four target words, two that had been in the training list and two had not. Her subjects were instructed to read the stories aloud, in anticipation of a test of comprehension (which did not actually occur). As a secondary task, they were to mark any words that had been in the training list. Cook's major manipulation was that half of each of the new and old words were presented in sentences that they fit both syntactically and semantically, such as "hungry" in a sentence beginning "The big hungry dog...". The other half violated both the syntax and meaning of the sentence in which they occurred, for example "ceiling" in "The lazy ceiling cat...".

Cook's subjects (1996) detected 67% of the old target words that were placed in incongruous contexts. This rate of report was not a simple bias-type reaction to the incongruity: The subjects marked only 5% of novel words occurring in incongruous contexts. In contrast, they detected only 27% of old words placed in congruous contexts, and false alarmed on none of the new words in those contexts. Moreover, Cook demonstrated that target words in congruous contexts were named about 150 ms faster than those in incongruous contexts; by that measure, congruous words were processed more fluently than incongruous words. However, in the context of congruous text, that greater fluency was unsurprising to the subjects. They thus failed to experience a feeling of familiarity for words that were presented in unsurprising contexts, although they could recognize those words as old when presented in a surprising context. This of course does not prove that surprise is the direct basis of the feeling of familiarity: Cook's subjects probably relied on a spontaneous feeling of familiarity to detect old words while reading congruous text, but stopped reading and tried to recognize incongruous words. However, it does demonstrate that feelings of familiarity often do not occur for factually familiar objects when fluent processing is not surprising, within the current context.

2. Experiments 1 and 2: Phlooensy effects

In the experiments of this article, we attempted a more direct assessment of the role of surprise in creating the feeling of familiarity. We induced surprise in two ways: First, by showing subjects pseudohomophones (e.g., PHRAWG), items that look like nonwords, but, when pronounced, correspond to natural English words (Experiments 1 and 2); second, by showing subjects orthographically regular nonwords (e.g., HENSION), items that are similar to many natural words in orthography and phonology, but have no meanings (Experiments 3–7).

As pointed out by Mandler (1980), people can claim to recognize a re-presented item on one of two bases, either a feeling of familiarity alone or in addition through recall of the context in which that item occurred. In these studies, we were exclusively interested in the source of the feeling of familiarity, in isolation from recall of a prior event. For that reason, although we presented both old and novel items in test, we primarily examined trials on which the target stimulus had not been presented earlier in the experiment. On those trials, recall of earlier processing or context could not influence recognition reports: Only guessing and feelings of familiarity for the stimuli could lead the subject to claim to recognize them.

2.1. Method

Subjects. Twelve Simon Fraser University undergraduates participated in Experiment 1, and 9 in Experiment 2, for course credit.

Stimuli. We collected a pool of 60 one- and two-syllable natural words (e.g., FROG, CANCER, BOTTLE). From these items, we generated 60 pseudohomophones (e.g., PHRAWG, KANSER, BAUTEL), by spelling the natural words in an unfamiliar way that produced the same sound when pronounced. In addition, we created 60 nonwords (e.g., CULSE, BELINT, LAFER) which did not correspond in either spelling or sound to any English word.

Procedure. The experiments consisted of two phases, a training phase followed by a recognition test. In the first phase, subjects were shown 40 natural words (all spelled in the natural way) and 40 nonwords, in random order. These items were selected at random from their respective pools for each subject. Subjects were instructed to memorize these items "for a later test of memory", by pronouncing each item aloud. Trials were subject-paced.

At test, 20 of the training words were shown again, presented in their natural spelling. The remaining 20 were also presented, but shown in their pseudohomophone versions. Twenty nonwords that had been presented in the training phase were also presented in this test. In addition to these 60 old items, we presented 20 natural words and 20 nonwords that had not been seen in training and also 20 pseudohomophones that did not correspond to any word shown in training. Test items were also presented in random order, re-randomized for each subject.

On each test trial, subjects were shown one test item and asked to perform three activities: To pronounce the item, to decide if it was a word, and then to decide if they had seen that item previously. The middle task was phonological lexical decision: an item was to be judged a word if, when pronounced, its sound corresponded to that of a natural word. Both naturally spelled words and pseudohomophones were thus to be classified as words. For the recognition decision, we instructed subjects to call an item old either if they had seen that exact item before or (in the case of pseudohomophones) they had seen a word corresponding to the sound of the current item. That is, if FROG was presented in training, then PHRAWG should be classified as old.

Subjects might base their recognition decisions on the act of pronunciation alone, or also the act of making a lexical decision. In Experiment 1, subjects pronounced the item and then hit one of two keys to indicate their lexical decision. This keystroke recorded the combined latency of pronunciation and lexical decision as well as the accuracy of lexical decision. Subjects then used the same keys to indicate their recognition decision. In Experiment 2, we instead recorded the latency of pronunciation alone, requiring subjects to strike a third key at the same time that they pronounced the item; they then used the other two keys to record their lexical and recognition judgements. All other aspects of the two experiments were identical.

The logic of the experiments depended on subjects being willing to claim to recognize some stimuli, with some accuracy. We imposed a criterion that subjects must claim to recognize at least 50% of the training items to be included in the final analysis (two subjects were rejected under this criterion). In addition, our predictions about surprise depended on subjects realizing that the pseudohomophones actually represented words. Therefore, in evaluating the recognition and reaction latency data, we used only those trials on which subjects had made correct lexical decisions. Subjects were generally accurate on this task: The minimum mean success in lexical decision in any cell was 85%.

2.2. Results and discussion

We examined latencies for the combined pronunciation and lexical decision gathered in Experiment 1 (see Table 1). As might be expected, old items of each type were pronounced and judged slightly faster than new items, although the differences were not reliable with such a small sample, $F < 1$. More important for current purposes, natural words, whether new or old, were pronounced and judged faster than pseudo,

Table 1
Experiment 1: Pronunciation, lexical decision and recognition

	Old items	New items
Combined pronunciation and lexical decision latency (ms)		
Word	1435	1468
Pseudohomophone	1936	1956
Nonword	2195	2215
Recognition decision: *p*(claim "old")		
Word	0.67	0.15
Pseudohomophone	0.66	0.34
Nonword	0.57	0.06

and pseudohomophones faster than nonwords. As discussed earlier, we will concentrate on the cells involving novel items. In those cases, natural words were pronounced and judged about 500 ms faster than pseudohomophones, $F(1,11) = 30.09$, MSE = 47 570, $p < 0.001$, and pseudohomophones about 250 ms faster than nonwords, $F(1,11) = 18.74$, MSE = 21 527, $p < 0.001$. That is, natural words were processed most efficiently, and nonwords least. If fluency per se is the basis of the feeling of familiarity, then we should observe that subjects would produce most false alarms to new natural words in the recognition decision, and least to new nonwords.

That hypothesis was disconfirmed. Our subjects committed 9% more false alarms on words than on nonwords, $F(1,11) = 5.69$, MSE = 0.01, $p < 0.036$, but also 19% more false alarms on pseudohomophones than on words, $F(1,11) = 19.73$, MSE = 0.03, $p < 0.001$. The greater illusory familiarity of words over nonwords could be explained through the superior fluency of processing words, but the extra 19% illusory familiarity of pseudohomophones over words cannot be.

In Experiment 2, we measured the latency of subjects' pronunciation of items, rather than a combined pronunciation–lexical decision latency (see Table 2). Novel words were pronounced about 90 ms faster than novel pseudohomophones, $F(1,8) = 5.94$, MSE = 6660, $p < 0.04$, and pseudohomophones about 60 ms faster than novel nonwords, although that difference was not reliable, $F(1,8) = 1.84$, MSE = 9991, $p > 0.2$. Taking pronunciation latency as an index of the relative fluency of processing each of these types of stimuli, novel words should suffer more false alarms than pseudohomophones, if fluency per se is responsible for feelings of familiarity.

As in Experiment 1, that prediction was disconfirmed. Novel pseudohomophones produced 24% more false alarms than words, $F(1,8) = 9.44$, MSE = 0.03, $p < 0.015$, and 29% more false alarms than nonwords, $F(1,8) = 17.27$, MSE = 0.02, $p < 0.003$. (The 5% difference between words and nonwords was not reliable, $F(1,8) < 1$.) Again, the relative fluency of processing the words and pseudohomophones could not predict subjects' feelings of familiarity for them.

We suggest instead that the feeling of familiarity results from an interpretation of the fluency of performance, an interpretation that takes into account the degree of fluency one could expect to experience in dealing with an item for some purpose.

Table 2
Experiment 2: Pronunciation and recognition

	Old items	New items
Pronunciation latency (ms)		
Word	671	695
Pseudohomophone	830	789
Nonword	758	853
Recognition decision: p(claim "old")		
Word	0.62	0.19
Pseudohomophone	0.66	0.43
Nonword	0.65	0.14

The expectation is produced on-line, in the act of processing the item. [2] We further suggest that a feeling of familiarity is produced when there is a discrepancy between what is experienced and what is initially expected. Orthographic and phonological processing of a word occurs fluently, supported by the mass of prior experience with that word, even before the person becomes aware that it is a word: this fluent processing produces the expectation that one will be able to produce a meaning for that stimulus. The coming to mind of a meaning attached to that stimulus is thus not surprising; the outcome matches the initial expectation. Initial orthographic and phonological processing of nonwords occurs less fluently, because it is less well supported by previous experience: but again, that is not surprising to the person. The low initial fluency causes an expectation that one will not be able to produce a meaning, even before one knows it is a nonword. The failure of a meaning to come to mind for that item is thus concordant with the initial expectation raised by the fluency of processing.

However, the novel pseudohomophones produce a surprise. Initially, they are processed less fluently than natural words, because the orthographic pattern of the item does not match any specific representations in memory that could assist pronunciation. This nonfluent processing leads to the same expectations aroused by the presentation of nonwords. However, in the act of pronouncing the item, the person produces another stimulus, namely the spoken sound pattern corresponding to the orthography, and that sound pattern does correspond to representations in memory of a real word. This cues the meaning of the word to come to mind, violating the person's expectation about the extent of processing they will be able to perform on the item. The item catastrophically re-organizes into a known word, producing a feeling state.

The surprise resulting from violation of initial expectations can produce several different feeling states. If people are shown a pseudohomophone, like KAFF, outside of a recognition experiment, they are surprised by it, but do not usually experience a feeling of familiarity. Instead, they frequently say "Oh!" or smile as they realize its meaning, as though they had just heard a joke. [3] The surprise violation of expectation produces a feeling, but it is expressed as something like "That's cute" rather than as "That feels familiar": That is, the surprise is correctly attributed to the orthographic form in which the word is presented rather than to some past experience. However, in the context of performing a recognition decision, in which people

[2] As discussed later, we argue that these "expectations" exist as unconscious readiness to perform certain comparisons, not as consciously held attitudes.

[3] In fact, we think that the basis of our familiarity effects with pseudohomophones has much in common with the redintegration basis of humor, in which the punchline suddenly changes the meaning of events preceding it, such as "I know a man who broke his arm in two places and swore he would never visit those places again". Clang associations, such as ROW, ROW, ROW YOUR GOAT, in which expectations aroused by early processing are violated by a nearly fitting termination, also produce giggles. We did not count giggles and smiles among our subjects, but we believe the affective reaction to such sudden violations is the basis not only of the feeling of familiarity, but of a very wide variety of evaluative feelings, such as pleasantness, beauty, interestingness, goodness, badness and so on.

attempt to use the quality of their present experience of an item to decide about their prior experience of that item, the surprise is misattributed to a prior experience.

We thus suggest that familiarity is not the product of fluency per se, but rather the product of fluency that is unexpected for that item, task and context. Fluency that matches expectation requires no explanation, and produces no feeling state. However, fluency that fails to match expectation will be attributed to some source, rightly or wrongly, and will produce some sort of feeling state, the type of feeling depending on the source to which the fluency is attributed.

We argue that attribution of unexpected fluency to a source in the past not only produces an illusion of familiarity when novel pseudohomophones redintegrate, but that is also the basis for false feelings of familiarity experienced with novel words and nonwords. We also argue that it is the basis of accurate feelings of familiarity toward items that have been encountered previously. A prior encounter with a word in a particular context does usually facilitate later processing of that same word in a similar context on a later occasion. This is the phenomenon of repetition priming, which has been extensively documented over the last 20 years (e.g., Jacoby and Dallas, 1981). The effect of that prior presentation is to raise the efficiency of processing that item in that context above the level that could ordinarily be expected for it. A feeling of familiarity is engendered when the item is encountered in test, not directly by the fluency of processing that word, but instead by the unexpectedly great efficiency of processing that word in that test context.

The cautious reader might wonder why, if our account is correct, the probabilities of correctly claiming to recognize words, nonwords and pseudohomophones that actually were presented in training are so similar (see Tables 1 and 2), whereas the rates of false claims of recognition differ so much between those types of stimulus. Our answer must be that false claims of recognition can only be based on guessing and feelings of familiarity, whereas true claims can also be based on recall of the event itself. In consequence, the true recognition of different types of stimuli depends on the distinctiveness and elaborateness of encoding those various types of item in the original training, and in the cue properties of natural words, pseudohomophones and nonwords to recall contextual properties of the original experience. These additional factors can complicate the pattern of hits across different stimulus types, making it difficult to study feelings of familiarity independent of the other basis of recognition, actual recall. However, those factors are irrelevant for novel test items. Curiously, the basis of the feeling of familiarity is most easily studied when that feeling occurs in the absence of a prior experience.

3. Experiment 3: Unexpected failure of processing

Experiments 1 and 2 demonstrated that people experience strong feelings of familiarity when they are unexpectedly able to produce a meaning for a stimulus that initially appears to be meaningless. We wondered whether surprising success is an essential ingredient of familiarity, or instead whether, more generally, any departure from the expected outcome could produce such a feeling. In Experiment 3, we set out to create stimuli which would cause an initial expectation of knowing the items, which would later turn out to be unjustified.

We created three new pools of stimuli, respectively consisting of common natural words (e.g., RAINBOW, TABLE, CANDLE), a set of orthographically regular nonwords that would be easy to pronounce (e.g., HENSION, FRAMBLE, BARDEN), and a set of orthographically irregular nonwords that would be hard to pronounce (e.g., JUFICT, STOFWUS, LICTPUB). Half of each of the three types of stimulus were exposed in a memorization training task, and then subjects were exposed to the entire set in a recognition task.

The logic of the design was that the regular nonwords would be processed fluently during the pronunciation task, producing a false initial expectation that the person could go on to produce their meanings, an idea corrected in the subsequent lexical decision task. In contrast, words and irregular nonwords would each produce realistic expectations, respectively that the person could go on to produce more information about the stimulus or that they could not. If a feeling of familiarity is only produced by the happy surprise that one can produce more information about a stimulus than one initially thinks, then we would not observe false familiarity for the regular nonwords; but if any kind of violation of expectation about one's ability to produce more information about a stimulus can support a feeling of familiarity, then we could expect high rates of false claims to recognize novel regular nonwords.

3.1. Method

Subjects. 10 Simon Fraser University undergraduates participated for course credit.

Procedure. We created three pools of items, consisting of 60 words (e.g., GAMBLE, TABLE, ANIMAL), 60 regular nonwords that would be easy to pronounce and 60 irregular nonwords that would be hard to pronounce. Regular nonwords (e.g., FISSEL, PLANDIT, MANIPER) tended to contain soft and blended consonants, smooth transitions between syllables, and letter combinations that occur with high frequency, whereas items intended to be irregular (e.g., BLAFTED, WICSTA, NOTIRGIN) tended to contain hard consonants, interruptions between syllables and unusual letter combinations.

In a training phase, subjects were shown 40 items of each type, randomly selected from their respective pools and shown in random order, the selection and order being freshly randomized for each subject. They were asked to memorize the items for a later recognition test. Each item was exposed until the subject struck a key to expose the next item. In the subsequent test, 20 of these items of each type were shown, together with the remaining 20 items of each type that had not been selected for training. As in Experiments 1 and 2, subjects were asked to perform three tasks on each test item: To pronounce it, to decide if it were a word or nonword, and to decide if it had been presented in training. (Subjects performed the standard lexical decision task, rather than the phonological decision required in Experiments 1 and 2.) Naming latency, lexical decision latency and accuracy, and recognition accuracy were recorded separately, through separate keystrokes. We do not claim that the recorded latencies are exclusive measurements of naming and lexical decision: Clearly, subjects may have been thinking about the lexical status of the word while trying to pronounce it. As in Experiments 1 and 2, we required that subjects claim at least 50% of old items to be old. Further, in scoring the data, we used only those trials on which subjects had made correct lexical decisions. The minimum mean accuracy on this judgement for any type of stimulus was about 90%.

3.2. Results and discussion

First, we examined the latency of pronouncing novel test items (see Table 3). Novel words were pronounced about 150 ms faster than regular nonwords, $F(1,9) = 10.01$, $MSE = 11\ 433$, $p < 0.011$, and regular nonwords about 300 ms faster than irregular nonwords, $F(1,9) = 21.65$, $MSE = 21\ 905$, $p < 0.001$. If pronunciation fluency directly determines feelings of familiarity, we should observe most false alarms for words and least for irregular nonwords.

The efficiency of performing the lexical decision task might also form part of the basis of familiarity. Novel words were judged about 170 ms faster than regular nonwords, although the difference was of marginal reliability, $F(1,9) = 4.21$, $MSE = 34\ 539$, $p < 0.07$, and regular nonwords about 100 ms faster than irregular nonwords, although the difference was not reliable, $F(1,9) = 1.92$, $MSE = 28\ 739$, $p > 0.2$. This ordering makes the same prediction for relative rates of false alarms as does the pronunciation data, that false alarms for regular nonwords should be intermediate between the rates for words and irregular nonwords.

However, as seen in Table 3, there were 21% more false alarms for regular nonwords than for words, $F(1,9) = 17.35$, $MSE = 0.01$, $p < 0.002$, and 28% more for regular than irregular nonwords, $F(1,9) = 31.7$, $MSE = 0.01$, $p < 0.001$. There were slightly more false alarms for words than for irregular nonwords, but that difference was not reliable, $F(1,9) = 1.53$, $MSE = 0.02$, $p > 0.25$.

Table 3
Experiment 3: Pronunciation, lexical decision and recognition

	Old items	New items
Pronunciation *rt*: (ms)		
Word	827	837
Regular nonword	940	988
Irregular nonword	1237	1296
Lexical decision *rt*: (ms)		
Word	1056	1059
Regular nonword	1308	1230
Irregular nonword	1181	1335
Recognition decision: *p*(claim "old")		
Word	0.74	0.16
Regular nonword	0.68	0.37
Irregular nonword	0.61	0.09

As in Experiments 1 and 2, we observed that feelings of familiarity, as revealed by false alarms, are not directly associated with differences in processing fluency, either in pronouncing the item or deciding whether it is meaningful. Instead, we again interpret the result to mean that feelings of familiarity are produced when the actual fluency of processing violates the person's expectation for that item, in that task and context. The initial nonfluent pronunciation of irregular nonwords produced the expectation that those items were not words: Failure of meaning to come to mind during the lexical decision was therefore consonant with the initial expectation. Similarly, the initial fluent processing of words produced the expectation that the subjects did know those items: That expectation was confirmed in the lexical decision task. However, pronunciation of regular nonwords was also fairly fluent: They were pronounced 300 ms faster than irregular nonwords, and only 150 ms slower than words. That is, initially, regular nonwords were processed more like words than like nonwords, producing the same initial expectations about regular nonwords as about words.

There are in effect three possible sources of fluent performance in the pronunciation task: Experience of that item prior to the experiment (as is the case for natural words), experience of that item in the training phase of the experiment (the case for old test items), and experience of one or more items in the past that are orthographically similar to the test item (the case for words and regular nonwords). The subject's problem in the recognition task is to discriminate among these sources when a test item is fluently pronounced. The true source of the fluent pronunciation of novel regular nonwords was of course their orthographic regularity. Normatively, those items should feel easy, not familiar. However, the subjects did not realize that their performance on regular nonwords had been selectively enhanced in this way relative to irregular nonwords. They therefore developed the same initial expectation that they did for words, that they knew these items from the remote past, and so could produce more information about them, such as their meaning. When this expectation was violated during the lexical decision task, they attributed the fluency instead to the other possible source of which they were aware, namely prior experience of that item once before, in the training phase of the experiment. The surprising failure to produce a meaning for the fluently pronounced regular nonwords thus produced a feeling of familiarity for those items.

The general result of this experiment is the same as that of Experiments 1 and 2: That less fluent items can produce greater feelings of familiarity than more fluently processed items, if the fluency is surprising. However, the basis of the surprise is opposite in this case from the earlier one. Whereas in Experiments 1 and 2 nonfluent pronunciation was followed by fluent meaning production, in the present case fluent pronunciation was followed by failure to produce a meaning. We conclude that surprise is a very general basis for the feeling of familiarity: A discrepancy between initial expectation and later performance on an item – whether fluent processing followed by surprising failure, or nonfluent processing followed by surprising success – will inevitably produce an unconscious attribution and a conscious feeling.

4. Experiment 4: Evaluating fluency

We have suggested that people initially encountering a stimulus for the purpose of pronouncing it are able to evaluate the fluency of their processing, and use it to develop an expectation about the course of further processing. We wondered how the person sets norms for deciding whether performance on a stimulus is fluent or not. The problem is this: Never having encountered HENSION before, how do people decide that their processing of that item is surprisingly fluent for a nonword? One possibility is that people infer their degree of fluency relative to their performance on other items in the same task and context: They actually learn about the fluency of processing nonwords such as STOFWUS and LICTPED, and their processing of HENSION is surprising in comparison to the norm established by those items.

To test this idea, we replicated Experiment 3 in detail, with the exception that we no longer presented irregular nonwords at training or test. This omission meant that subjects had lost a standard for evaluating the regular nonwords as being fluently processed. Instead, all that subjects would know about those items is that they were harder to process than the natural words. If people evaluate their fluency relative to their performance on other stimuli, then in this case they should evaluate their performance on regular nonwords as poor: They would not be surprised when no meaning came to mind for those items, and hence not experience strong feelings of familiarity for novel items.

4.1. Method

Subjects. Ten undergraduates participated for course credit.

Procedure. The procedure was identical to that of Exp. 3, except for the elimination of irregular nonwords from training and test.

4.2. Results and discussion

Novel words were pronounced about 90 ms faster than novel regular nonwords, $F(1,9) = 12.8$, $MSE = 3365$, $p < 0.006$ (see Table 4). Lexical decision was also about 150 ms faster for novel words than novel regular nonwords, although that difference was not reliable, $F(1,9) = 2.44$, $MSE = 43\ 790$, $p < 0.15$. These data match the similar comparisons in Experiment 3.

The data for the recognition test also matched the pattern observed in Experiment 3. Subjects committed 31% more false alarms on regular nonwords than words, $F(1,9) = 44.26$, $MSE = 0.01$, $p < 0.001$. This again demonstrates that fluency of processing, of itself, is not the basis of familiarity. We concluded once again that the subjects initially expected the regular nonwords to be words: When this expectation was violated, they experienced a feeling of familiarity.

More important, the fact that decisions about the regular nonwords were not influenced by the absence of irregular nonwords tells us that the development of initial expectations about a stimulus is independent of how they perform on other stimuli within the same context. People do not have to learn how fluently they process irreg-

Table 4
Experiment 4: Pronunciation, lexical decision and recognition

	Old items	New items
Pronunciation *rt*: (ms)		
Word	825	843
Regular nonword	908	936
Lexical decision *rt*: (ms)		
Word	911	1029
Regular nonword	1182	1175
Recognition decision: *p*(claim "old")		
Word	0.68	0.09
Regular nonword	0.75	0.40

ular nonwords to be surprised by their processing of regular nonwords. Instead, they employ absolute norms about processing fluency that lead them to expect that regular nonwords will turn out to be words. We suspect that these norms are embedded in the mass of past experience of processing words: Efficient processing of the regular nonwords is supported by cueing many traces containing similar orthographic components. The recruitment of these traces, and the control that they exert on current processing, may be the basis of the feeling that current processing is proceeding fluently. We discuss this issue further in Section 7.

5. Experiments 5 and 6: Attempts to debias

We have suggested that the false alarms produced by subjects in Experiments 3 and 4 resulted from surprise at finding a nonword to be fluently processed, followed by an attribution of the surprising fluency to a source in the past. We have presented the process that leads to a feeling of familiarity as a logical process of deduction; and so, in some sense, we believe it is. However, we do not mean to suggest that it is undertaken deliberately or consciously. Instead, we believe that the development of an expectation during the pronunciation task, the evaluation of that expectation after making the lexical decision and the attribution of surprising fluency to a source in the past are all conducted unconsciously and automatically. The only aspect of the process of which the person becomes aware is the feeling of familiarity.

The surprise experienced by our subjects was based on a false premise, that all nonwords will be processed nonfluently, so that any fluently processed item could be expected to be a word. The use of that heuristic might be conscious and deliberate. In that case, it could be corrected by informing subjects about the effects of orthographic regularity on performance. Alternatively, if only the end product of the comparison between initial and later processing is experienced consciously, then informing people about the conditions under which they operate might modify their experience very little.

We tested these alternatives in Experiments 5 and 6. In Experiments 5, we told subjects in advance of the test that some nonwords were structured to resemble

English words, and that would cause those items to be processed easily. We also warned them that this fluency might artificially make those items feel familiar, and that subjects in other experiments had fallen into this error. They were instructed to discount feelings resulting from an item's structure in making their recognition decisions, and to call an item old only if they could truly remember it.

Of course, those subjects had to realize that a nonword was regular before they could attempt to discount that source in making a recognition judgement. We removed that obstacle in Experiment 6. Subjects in that study were given the same general warning prior to the test. But in addition, on each test trial, in advance of displaying a stimulus, we put a message on the screen, indicating the type of stimulus that was about to occur, messages reading respectively WORD, EASY NONWORD and HARD NONWORD. Of course, the subjects had still to decide whether the item was new or repeated from the training phase. Knowing that the item was created to be a regular nonword left them with the problem of trying to discriminate feelings caused by the structure of the item from feelings resulting from having seen it before.

5.1. Method

Subjects. Eleven Simon Fraser University subjects participated in Experiments 5, and 13 in Experiment 6, for course credit.

Procedure. The procedure of Experiment 5 was the same as that of Experiment 3, with the exception that prior to the test subjects were warned about the false-familiarity effects of the regular nonwords. Subjects in Experiment 5 were required to pronounce and recognize the stimuli, but not to make a lexical decision; in Experiment 6 we performed the latter decision for the subjects, by providing a message indicating whether the subsequent item was a word or regular or irregular nonword. This message was displayed until a key-press by the subject. The subjects thus had ample time to prepare for the type of stimulus to be displayed.

5.2. Results and discussion

In Experiment 5, in which subjects were given a general warning about the enhanced fluency of processing regular nonwords, we observed data very similar to that of Experiment 3 (see Table 5). New words were pronounced about 200 ms faster than new regular nonwords, $F(1,10) = 13.13$, MSE = 15 676, $p < 0.005$, but regular nonwords were pronounced about 400 ms faster than new irregular nonwords, $F(1,10) = 33.33$, MSE = 26 846, $p < 0.001$. In contrast, false alarms of recognition for regular nonwords were about 20% greater for regular nonwords than for words, $F(1,10) = 10.36$, MSE = 0.02, $p < 0.009$, or for irregular nonwords, $F(1,10) = 8.63$, MSE = 0.03, $p < 0.015$. Clearly, the warning about the false-familiarity effects of regular nonwords did not assist subjects in discriminating fluency due to regularity from that due to actual prior exposure of an item.

In Experiment 6, in which subjects were also told on each trial whether the item would be regular or irregular, claims to recognize novel regular nonwords declined

Table 5
Experiment 5: Pronunciation and recognition

	Old items	New items
Pronunciation *rt*: (ms)		
Word	941	984
Regular nonword	1065	1177
Irregular nonword	1416	1580
Recognition decision: *p*(claim "old")		
Word	0.76	0.22
Regular nonword	0.86	0.41
Irregular nonword	0.86	0.21

by about 10% compared to Experiment 5 (see Table 6), although the between-experiments difference was not reliable, $t(22) = 1.34$, $p > 0.2$. In this study, new words were pronounced about 80 ms faster than new regular nonwords, $F(1,12) = 5.04$, $MSE = 2748$, $p < 0.045$, but regular nonwords were pronounced about 430 ms faster than new irregular nonwords, $F(1,12) = 34.70$, $MSE = 36\,881$, $p < 0.001$. [4] However, false alarms of recognition for regular nonwords were still about 12% greater for regular nonwords than for words, $F(1,12) = 4.72$, $MSE = 0.02$, $p < 0.05$, and 10% greater regular nonwords than irregular nonwords, $F(1,12) = 4.06$, $MSE = 0.01$, $p < 0.067$.

We concluded that the feelings created by the processing of a regular nonword are, to a great degree, cognitively impenetrable. That is, people cannot discount the extra fluency produced by orthographic regularity, even when made aware of the orthographic status of the item and its effect on their performance: regular nonwords still feel surprisingly fluent. We further concluded that the unexpected fluency of processing regular nonwords is consciously experienced simply as a feeling of familiarity, not as a feeling of fluency which is deliberately attributed to prior experience. In the context of a recognition judgement, surprising fluency just feels to be familiarity.

However, the product of surprise and attribution is not inevitably a feeling of familiarity. If a bomb goes off next door, or if one returns home to find the furniture re-arranged, one will certainly be surprised, but will not experience a feeling of familiarity: Instead, one may experience a feeling of shock or unease. Surprising processing is attributed to whatever source seems likely given the type of stimulus, task and context. In the context of a recognition experiment, prior experience of a stimulus in the training phase is an obvious possible source of surprising fluency: The same fluency will instead be attributed to the structure of the stimulus domain in the context of a classification test (e.g., Whittlesea and Dorken, 1993). Further, the fluency of performance can be attributed simultaneously to multiple sources. For example, Whittlesea (1993) asked subjects to rate the pleasantness of words as well as to judge

[4] The slower responding in Experiment 5 than 6 is due to the fact that subjects in the latter did not have to perform lexical decision. As discussed earlier, we do not claim that the pronunciation latency is a pure measure of processes involved in producing an utterance: Clearly, subjects in Experiment 5 were preparing to perform the lexical decision task while pronouncing the item.

Table 6
Experiment 6: Pronunciation and recognition

	Old items	New items
Pronunciation *rt*: (ms)		
Word	821	873
Regular nonword	905	919
Irregular nonword	1250	1363
Recognition decision: *p*(claim "old")		
Word	0.71	0.18
Regular nonword	0.68	0.30
Irregular nonword	0.58	0.20

them for recognition. A covert manipulation of the fluency of processing produced large effects on both of these judgements. Familiarity is thus only one of many feeling states that a person can experience following surprise at their processing: Which feeling the person experiences depends on which source of performance occurs to them to consider, given the stimulus, task and context involved. We explored this malleability of feeling in the final study.

6. Experiment 7: An illusion of duration

Witherspoon and Allan (1985) asked subjects to study a list of words. In a subsequent test, they flashed old and new words on a computer screen, at two different short durations. The subjects were not asked to recognize old items: Instead, they were asked to name each item and then identify whether it had been presented for a relatively long or short duration. They observed that old words were often falsely judged to be presented for the longer duration. Witherspoon and Allen concluded that people judge duration from the fluency they experience in identifying the stimuli.

We agree with this analysis. However, we wondered whether it was the relative fluency per se, or instead the surprise caused by the extra fluency in naming old items that was responsible for the mistaken judgements of duration. We also presented stimuli at two durations, asking subjects to discriminate them. However, instead of presenting some items in a training phase, we simply presented the stimuli from Experiment 3 in a test. There was thus no question of feelings of familiarity: The subjects knew that all words were items they had known in the past, and that all nonwords, whether regular or irregular, were novel in their experience. Instead, the question was whether the differences in fluency of pronunciation that we had observed in Experiment 3 would predict the subjects' ratings of duration, or whether the surprising fluency of regular nonwords would disorder that prediction.

6.1. Method

Subjects. Ten Simon Fraser University undergraduates participated for course credit.

Procedure. We showed subjects 40 words, 40 regular nonwords and 40 irregular nonwords, randomly selected from the pools of those items that had been created for Experiment 3. Half were presented for 100 ms, the other half for 200 ms, the sequence of stimulus type and duration being freshly randomized for each subject. Trials were subject-paced. Each presentation of a target stimulus was pre- and post-masked by a row of Xs, extending beyond the target stimulus at front and back. On each trial, after the presentation, subjects first pronounced the stimulus, then decided if it was a word or nonword, then decided if that presentation was relatively short or long. Reaction latencies were taken for both naming and lexical decision tasks.

6.2. Results and discussion

Actual differences in the duration of the stimuli affected subjects much as one might expect. Long presentations were pronounced about 80 ms faster than short presentations, and were judged long about 26% more often than short presentations (see Table 7). More important, the type of stimulus (word, regular or irregular nonword) also had a large effect on subjects' performance. The two-way ANOVA did not produce a reliable interaction $F(1,9) < 1$, so we collapsed across the presentations factor in analyzing the effect of different stimulus types on judged duration.

The subjects pronounced words about 150 ms faster than regular nonwords, $F(1,9) = 5.80$, MSE = 21 277, $p < 0.039$, and regular nonwords about 270 ms faster than irregular nonwords, $F(1,9) = 9.81$, MSE = 38 062, $p < 0.012$. They were also about 350 ms faster in performing lexical decision on words than on regular nonwords, $F(1,9) = 31.86$, MSE = 19 207, $p < 0.001$, and about 110 ms faster on regular nonwords than irregular nonwords, $F(1,9) = 1.22$, MSE = 47 390, $p > 0.29$. In accord with the simple fluency hypothesis, words were claimed to be of long duration 16% more often than irregular nonwords, $F(1,9) = 6.20$, MSE = 0.02, $p < 0.034$. However, regular nonwords were judged to be of long duration at the same rate

Table 7
Experiment 7: Pronunciation, lexical decision and duration judgement

	Short items	Long items
Pronunciation *rt*: (ms)		
Word	899	869
Regular nonword	1057	1025
Irregular nonword	1408	1221
Lexical decision *rt*: (ms)		
Word	1149	1174
Regular nonword	1497	1526
Irregular nonword	1650	1588
Duration decision: *p*(claim "long")		
Word	0.48	0.71
Regular nonword	0.46	0.76
Irregular nonword	0.33	0.56

as words, 16% more often than irregular nonwords, $F(1,9) = 7.73$, $\text{MSE} = 0.01$, $p < 0.021$.

In this study, unlike the previous ones, we did not observe an overshoot effect, such that regular nonwords produced greatly exaggerated claims of the quality to be judged, relative to the more fluently processed words. [5] Nonetheless, the data clearly indicate that fluency per se, as measured by speed of naming and lexical decision, is insufficient to predict the ordering of duration judgements for the three types of stimulus. Instead, we concluded that surprise about the fluency of processing regular nonwords, that we argued in the other studies to be the basis of feelings of familiarity, can also be the basis of feelings of duration. In fact, we suspect that it is a very general basis of experiencing feelings about processing just conducted: It is likely to be the basis of feelings of wrongness in producing an error in speech, of well-formedness in categorical judgements, as well as feelings of pleasantness, beauty or truth.

7. A discrepancy-comparison account of familiarity

These results lead us to propose a general account of familiarity, and more generally of feelings with respect to current processing, based on the idea that the feelings result from the perception of a discrepancy in the processing of two aspects of a stimulus. Most simply, this discrepancy-computation mechanism can be thought of as comparing an initial expectation and an outcome. However, we suspect that the discrepant aspects of processing need not be sequential: For example, if simultaneous processing of orthographic and phonological properties produce discrepant interpretations of the stimulus, then a feeling-state will probably occur.

In our experiments, the initial expectation about the nature of the stimulus was probably controlled by the fluency of orthographic and phonological processes involved in transforming a printed word into an utterance. Fluent processing of these aspects creates the expectation that one knows something about that stimulus. The precise detail of what comes to be expected probably depends on the demands of the situation. Differences in processing fluency led our subjects to expect that the stimulus would be a word or a nonword, because of the task requirement to perform lexical decision; however, in other cases, such as in comparison of high- versus low-frequency words, or emotive versus nonemotive words, the expectation may instead be about those aspects of the stimulus. The development of an initial expectation will depend on other sources in other cases. For example, in passing a person at a bus-stop, one may suddenly turn back, experiencing a feeling of familiarity. In that

[5] False claims of long duration for regular nonwords were only 15% greater than for irregular nonwords, compared to a 28% difference in false claims of recognition in Experiment 3. We have replicated both experiments several times, and are convinced that false feelings of stimulus quality are indeed greater in the context of recognition than in a duration judgement. The reason for the attenuation of the effect is not clear, except that it obviously has something to do with the fact that one judgement is about events that are remote in time, whereas the other is about an event that just occurred.

case, the initial expectation that the face is a known one probably results from the fluency of organizing the features of the face into a gestalt, supported by the similarity of that face, or some of its major components, to other faces that one knows. Alternatively, the stranger's posture or movements might be similar to those of someone you know, triggering a sudden turning of attention to that stimulus.

Thus the initial processing of a stimulus proceeds with some degree of efficiency on some dimensions of processing. This fluency is then subject to an interpretation, leading one to expect that one has some degree and some kind of knowledge about the stimulus, even before one yet knows what it is. This interpretation may be based on a valid heuristic, such as to expect more fluently processed words to be of greater frequency. Alternatively, the heuristic may be invalid, such as interpreting fluent processing to mean that the stimulus will be a meaningful word, rather than a regular nonword. (Of course, in the real world, that heuristic usually is valid; one rarely has to defend oneself against regular nonwords in text or speech.) The validity of this interpretation determines in part whether a feeling of familiarity that is experienced later is true or false. For example, the partially valid heuristic that efficient processing of a face means prior knowledge of that face leads to many false alarms of familiarity, as well as hits.

7.1. Norms on the fly

When initial processing of the stimulus is completed, the person has another chance to evaluate the stimulus and their processing of it. This evaluation may match the initial expectation, or differ from it to some degree. When initial and later processing are coherent, no feeling state is aroused: the person simply continues to interact with the stimulus object. The initial processing of one's spouse in one's kitchen, or the face of the clerk in the corner store, is highly fluent: however, it is also unsurprising. The fact that these contexts are well known produces the expectation that processing of stimuli within those contexts will be fluent: The subsequent fluent interactions with the spouse or clerk are coherent with that expectation. In consequence, one experiences no feeling state: one simply continues to interact with those stimuli, performing actions appropriate to that stimulus in that context. In contrast, the context of a bus-ride through an unfamiliar part of town engenders the expectation that no stimulus will be processed with especially high fluency. Encountering the clerk's face is therefore a surprising event in that context. The unexpectedly fluent processing requires an explanation; attributed to past experience, it produces a powerful feeling of familiarity.

The unexpected encounter with one's spouse in one's lecture theater appears to require an extension of this account. It is not the case that one does not expect to see familiar faces in that context: In fact, many of the faces of one's students come to be fluently identifiable over a term. Instead, this case suggests a more specific relationship between expectation and feelings. In this case, it is not the fact of fluent processing that is startling, but that the specific identity that is produced by the processing does not match expectations for that context. Similarly, it is possible to encounter the clerk on the bus and recognize her, immediately know who she is. The recognition does not prevent the occurrence of surprise and a consequent feeling state; but that state is not one of familiarity, but instead a feeling of "what's she doing here?", a feeling that is not aroused by any of the other passengers. In both of these cases, although identification processing proceeds smoothly to completion, there is a mismatch between the expectations generated by the context and the outcome of processing the stimulus, resulting in a feeling state.

This suggests that surprise may be engendered in many ways, and that a feeling state is the inevitable consequence of that surprise. The feeling might be one of fear, as in the case of sudden noise, or of hilarity, in the case of a surprising punchline to a joke, or unease or bewilderment, in coming home to find the furniture rearranged, or of familiarity, in the case of unexpectedly fluent processing of a face seen at a bus-stop. Which of these is felt depends, we suspect, on one's interpretation of the possible meaning of the surprising event, as being dangerous or pleasant: The interpretation depends on the intuitive theories one has about the possible causes of various types of experience. The malleability of people's interpretations of their state, and the consequent malleability of the feeling state, has been extensively documented in the social psychology literature (cf., Nisbett and Ross, 1980). Similarly, unexpectedly great processing fluency resulting from a prior exposure can produce a bewildering variety of feelings about current stimuli, including brightness OR darkness (Mandler et al., 1987), duration (Witherspoon and Allan, 1985), loudness (Jacoby et al., 1988), truth (Begg and Armour, 1991), pleasantness (Zajonc, 1980), understanding (Carroll and Masson, 1992), and knowing (Jacoby et al., 1989b). We conclude that familiarity is only one of many feeling states aroused through a common process, namely the attribution of surprising processing to some source.

7.2. Idiosyncratic norms

It is easy to understand how the discrepancy-computation mechanism assists people to feel true familiarity for low-frequency items, such as a face seen only once before: The initial expectation that one knows the person, based on the fluency produced by the earlier encounter, is strongly discrepant with the subsequent failure to produce a name. However, it is harder to understand how the mechanism works in helping a person to experience accurate feelings of familiarity about well-known items that are encountered in yet one more context, for example in assisting a person to decide that they saw TABLE but not WINDOW in an earlier list. Studying TABLE in the training list will produce facilitation in reading that word again at test (the repetition priming phenomenon); the high initial fluency of constructing a percept could lead to a variety of possible expectations, such as that the item is a word, is meaningful, of fairly high frequency and so on. The point is that, unlike in our previous examples, none of these broad categorical expectations will be violated when the word is eventually identified. Instead, the difference between the actual fluency of producing a percept and the fluency that normatively could be expected for that word is small. This suggests that the discrepancy-comparison mechanism must be quite sensitively tuned to differences in fluency. Moreover, the difference in fluency between two different words that have not been primed by prior presentation in a training list is likely to be considerably greater than the difference in fluency for a single word when it has and has not been presented in training. To enable a person to distinguish between new words that are easily processed without prior presentation from those that are harder to process but were shown earlier, the mechanism must have some means of computing a norm for each individual word as it appears in test. We will suggest one way in which that might be accomplished.

In a standard recognition study, subjects are asked to study a list of words for a limited period. The words are presented in some specific context (perhaps presented on a computer screen in some font, at some time of day, in some room with some examiner). In test, new and old words are shown, usually in the same context as training items. The initial processing of an old item benefits from its specific similarity to the earlier presentation, as well as its similarity to background experience of words: The processing of that item therefore proceeds with greater fluency than could be expected from general experience of that item. However, that facilitation is dependent on recapitulation of the perceptual characteristics of the earlier experience: case, font, context, etc. We speculate that, once processed through to a phonological or semantic code, the item again cues memory. On this occasion, the specific similarity to a prior event, which produced the priming advantage in naming the item, is lost. As a stimulus, the item now presents a cue (a phonological or semantic code) which is not specifically similar to that presented earlier in the experiment. In consequence, it receives the same support from the general background of experience, but does not receive as much support from the priming experience. This diminution of support for processing results in a discrepancy between expected and outcome fluency, producing a feeling of familiarity. Whether exactly right or not, the important point about this suggested mechanism is that it demonstrates a way in which the person might gain access to normative information about their knowledge of the item, separate from the experience of that item in naming it, that might be used as the basis of a discrepancy comparison that is specific to each particular item the person encounters. The difficulty in exploring this idea experimentally is that, when an item has been presented earlier, it is very difficult to separate pure feelings of familiarity from a recognition decision based on recall of some aspects of the context in which that item was presented.

The discrepancy-computation hypothesis also leads to some very interesting revisions of prior thinking. One has to do with the frequency effect in recognition. It has long been known that low-frequency words are easier to recognize in a test than are high frequency words (e.g., Kinsbourne and George, 1974). This effect has usually been attributed to a distinctiveness advantage in encoding for the lower frequency items: The idea is that the person is able to encode some meaningful attributes of those items that are very different from other items, and so are not confusable with other items in the test (e.g., Eysenck and Eyseck, 1980). Without denying that explanation, we can suggest another: That low-frequency items that are repeated from training actually feel more familiar than repeated high-frequency items. Low-frequency items benefit more in repetition priming than do higher frequency items (e.g., Forster and Davis, 1984; Kinoshita, 1995). In consequence, the discrepancy between initial expectation and later evaluation of the fluency with which items should be processed given that they are new would be greater for low- than high-frequency items, resulting in a greater feeling of familiarity, and consequent greater recognition. Similar arguments might also be made for the levels-of-processing and encoding-specificity paradigms.

7.3. Relationship to FOK

The feeling of knowing (usually referred to as FOK) is the feeling one has, when asked a general-knowledge question, that one could answer it given enough time; or even that production of the information is imminent (as in the tip-of the tongue state, or TOT). Research on this question seems to have paralleled that on the feeling of familiarity (or FOF, if you want to get that way). For example, early accounts of FOK suggested that people have an internal monitor which keeps track of what one knows, without itself providing access to the knowledge (e.g., Hart, 1965). This parallels the intuitive theory of the feeling of familiarity discussed in Section 1, that the feeling results from activation of a memory trace, whether that trace is recalled or not. In both accounts, the feeling state represents a way of knowing that the relevant information is in memory, without necessarily gaining access to it.

Recent investigations have presented a different explanation for the FOK, one similar to the second account of feelings of familiarity (the fluency-attribution account) discussed in Section 1. For example, Koriat (1993) and Reder and Ritter (1992) have produced accounts in which the FOK is thought to be the product of an inference about the processing conducted in attempting to produce an answer. If elements of the question itself are fluently processed, or if the person can retrieve well-articulated partial information about the answer, then they are likely to experience an FOK, whether or not they can later produce the answer.

We suggest that the latter explanations are very near the mark. However, we also suggest that they lack the same vital ingredient as the fluency account of familiarity. People do not often experience the FOK or TOT state: Usually, in using their general knowledge to speak, work, have ideas and so on, people perform efficiently, without experiencing any strong feeling of KNOWING what they are doing: they "just know", and carry on. Thus the FOK presents the same problem as does the feeling of familiarity. When processing of both the question and the answer is very fluent, as in answering "What is your name?", there is no pressing, vivid feeling of knowing the answer. Instead, the "feeling of knowing" occurs when one has not yet produced the answer; and the TOT occurs when one finds one cannot produce the answer.

We therefore suggest that the basis of the FOK will turn out to be the same as that of the feeling of familiarity. That is, it is the discrepancy between fluent processing of the question without being able to produce an answer, or between fluent generation of information related to the answer without being able to produce the answer, that produces the feeling state. If that speculation is true, then in future we ought to be able to produce illusions of knowing, based on artificially enhanced processing of peripheral aspects of knowledge that people do not possess. [6]

7.4. What is fluency?

In the experiments of this article, as in many other studies of familiarity, fluency has been indexed through the speed with which a person performs some task, such as naming. This is a convenient index, which can be measured directly through a clock. However, we have long suspected that this index is only correlated with the psychologically meaningful construct of fluency. For example, Whittlesea et al. (1990) used a masking manipulation to influence subjects' processing fluency: Lighter masking reduced subjects' reaction latencies by about 70 ms, but produced only a 4% increase in false alarms. In contrast, Whittlesea (1993) manipulated the semantic predictiveness of context: Predictive contexts reduced subjects' reaction latencies by about 100 ms, but produced an 18% increase in false alarms. That comparison suggests that reaction latency and false feelings of familiarity are correlated, but not that strongly: A slight decrease in RT was accompanied by a more than four-fold increase in false alarms. Probably more important, the two experiments differed in making subjects concentrate on the perceptual ease in identifying targets versus the ease with which they could generate an identity in a meaningful context.

We suspect that the psychologically effective aspect of initial processing that is involved in the discrepancy comparison is not directly the speed or rate at which a name can be produced for a stimulus, but instead the quality of that processing: For example, the degree to which the component features of a stimulus configure into an integral whole when processed, rather than directly the speed with which a response can be produced. A word like TABLE is easily processed as a single unit, because of the frequency with which the person has encountered that constellation of features. HENSION is also easily processed as a unit, because of its similarity to many known units. In contrast, LICTPUB and STOFWUS are not experienced as units, but instead as relatively unorganized percepts consisting of separate parts, that are difficult to integrate.

We have used response latencies to index the psychological construct of fluency in these studies because they are probably linearly related, and because we cannot yet imagine a useful index of the qualitative goodness of processing a stimulus.

[6] In fact, however, we believe that such an enterprise is unnecessary, because we believe that there is precisely no difference between the actual mechanisms of knowing and remembering (c.f., Whittlesea, in press). The only difference between those activities is that, by definition, one involves the production of information about the identity of stimuli ("What country is Stockholm the capital of?") whereas the other involves production of information about specific events involving those stimuli ("When did you last visit Stockholm?"). These questions of course involve the cueing of different information to satisfy their different specific demands. However, we argue, the mechanisms through which one "just knows" one's own name and "just remembers" being in a major car accident, without any pressing FEELING that one is knowing or remembering, occur through the same principles, just as the FOK and the feeling of familiarity both occur when expectations based on initial processing do not match one's later processing.

However, we do not want to claim that our measurements of response latency directly correspond to the actual, effective variable to which our subjects responded. We, and others interested in the study of familiarity, will have to work on precising indices of the qualititative aspects of psychological experience.

7.5. Summary

We observed that less fluently processed stimuli can produce greater feelings of familiarity than more fluently processed items, if the fluency of the former is unexpected. The idea that feelings of familiarity are produced by violations of expectations about fluency, rather than fluency per se, adds a new dimension to the problem of familiarity. The development of these expectations, and the interpretation placed on a violation, is apparently automatic and unconscious, with the result that people experience familiarity as a primitive feeling. Future research on this topic will have to address the means by which different expectations are aroused by various aspects of processing in various contexts, and how the discrepancy between such expectations is computed without conscious intervention.

Acknowledgements

This research was supported by a grant to Bruce Whittlesea from the Natural Sciences and Engineering Research Council of Canada.

References

Begg, I., Armour, V., 1991. Repetition and the ring of truth: biasing comments. Canadian Journal of Behavioral Science 23, 195–213.

Carroll, M., Masson, M.E.J., 1992. Reading fluency as a basis for judgments of text comprehension: Misattributions of the effects of past experience. Unpublished manuscript.

Cook, S., 1996. When we remember in directed versus non-directed remembering tasks. Unpublished honors thesis.

Eysenck, M.W., Eyseck, M.C., 1980. Effects of processing depth, distinctiveness and word frequency on retention. British Journal of Psychology 71, 263–274.

Forster, K.I., Davis, C., 1984. Repetition priming and frequency attenuation in lexical access. Journal of Experimental Psychology 10, 680–698.

Hart, J.T., 1965. Memory and the feeling-of-knowing experience. Journal of Educational Psychology 56, 208–216.

Jacoby, L.L., Allan, L.G., Collins, J.C., Larwill, L.K., 1988. Memory influences subjective experience: noise judgments. Journal of Experimental Psychology 14, 240–247.

Jacoby, L.L., Dallas, M., 1981. On the relationship between autobiographical memory and perceptual learning. Journal of Experimental Psychology 110, 306–340.

Jacoby, L.L., Kelley, C.M., Dywan, J., 1989. Memory attributions. In: Roediger, H.L., Craik, F.I.M. (Eds.), Varieties of Memory and Consciousness: Essays in Honor of Endel Tulving. Erlbaum, Hillsdale, NJ, pp. 391–422.

Jacoby, L.L., Whitehouse, K., 1989. An illusion of memory: false recognition influenced by unconscious perception. Journal of Experimental Psychology 118, 126–135.

Jacoby, L.L., Woloshyn, V., Kelley, C.M., 1989b. Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. Journal of Experimental Psychology 118, 115–125.

Kinsbourne, M., George, J., 1974. The mechanism of the word-frequency effect on recognition memory. Journal of Verbal Learning and Verbal Behavior 13, 63–69.

Kinoshita, S., 1995. The word frequency effect in recognition memory versus repetition priming. Memory and Cognition 23, 569–580.

Koriat, A., 1993. How do we know that we know? The accessibility model of the feeling of knowing. Psychological Review 100, 609–639.

Lindsay, D.S., Read, J.D., 1994. Psychotherapy and memory for childhood sexual abuse: a cognitive perspective. Applied Congnitive Psychology 8, 281–337.

Lindsay, D.S., Kelley, C.M., 1996. Creating illusions of familiarity in a cued recall remember/know paradigm. Journal of Memory and Language 35, 127–211.

Mandler, G., 1980. Recognizing: the judgement of previous occurrence. Psychological Review 87, 252–271.

Mandler, G., Nakamura, Y., Van Zandt, B.J.S., 1987. Nonspecific effects of exposure on stimuli that cannot be recognized. Journal of Experimental Psychology 13, 646–648.

Nisbett, R.E., Ross, L., 1980. Human inference: Strategies and shortcomings of human judgement. Prentice-Hall, Englewood Cliffs, NJ.

Reder, L.M., Ritter, F.E., 1992. What determines initial feeling of knowing? Familiarity with the question terms, not the answer. Journal of Experimental Psychology 18, 435–451.

Roediger, H.L., McDermott, K.B., 1995. Creating false memories: remembering words not presented in lists. Journal of Experimental Psychology 21, 803–814.

Whittlesea, B.W.A. Production, evaluation, and preservation of experiences: constructive processing in remembering and performance tasks. The psychology of Learning and Motivation, 37, 211–264.

Whittlesea, B.W.A., 1993. Illusions of familiarity. Journal of Experimental Psychology 19, 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 122, 227–248.

Whittlesea, B.W.A., Jacoby, L.L., Girard, K., 1990. Illusions of immediate memory: evidence of an attributional basis for feelings of familiarity and perceptual quality. Journal of Memory and Language 29, 716–732.

Witherspoon, D., Allan, L.G., 1985. The effects of a prior presentation on temporal judgements in a perceptual identification task. Memory and Cognition 13, 101–111.

Zajonc, R.B., 1980. Feeling and thinking: Preferences need no inferences. American Psychologist 35, 151–175.