ELSEVIER

Acta Psychologica 98 (1998) 267–290

acta psychologica

Predicting the future and reconstructing the past: A Bayesian characterization of the utility of subjective fluency

Aaron S. Benjamin [a,*], Robert A. Bjork [a], Elliot Hirshman [b]

[a] *University of California, Los Angeles, USA*

[b] *Psychology Department, University of North Carolina, Chapel Hill, NC 27599, USA*

Abstract

The subjective sense of fluency with which an item can be perceived or remembered is proposed to be a vital cue in making decisions about the future memorability and the nature of our past experience with that stimulus. We first outline a number of cases in which such perceptual or retrieval fluency influences judgments both about our own future performance and our likely past experience, and then present a Bayesian analysis of how judgments of recognition – deciding whether or not a currently viewed item was studied at a particular point in the past – may incorporate information about the perceptual fluency of that item. Using a simple mathematical model, we then provide an interpretation of certain enigmatic phenomena in recognition memory.

PsycINFO classification: 2343

Keywords: Subjective fluency; Perceptual fluency; Retrieval fluency; Recognition memory; Metamemory; Word-frequency effects

1. Introduction

An emerging theme in modern research on memory is that subjective fluency – the ease or speed with which a task is accomplished – is used to guide a variety of

* Corresponding author. Address: Department of Psychology, UCLA, 405 Hilgard Avenue, Los Angeles, CA 90095-1563, USA. Tel.: +1 310 206-8257; e-mail: benjamin@psych.ucla.edu.

PII S 0 0 0 1 - 6 9 1 8 (9 7) 0 0 0 4 6 - 2

decisions and judgments that are made about the objective state of one's memory. In this article, we first attempt to lend a new organization to some of the recent findings that support such a conclusion. We then present a mathematical framework in which perceptual fluency is incorporated into judgments of recognition and use that framework to reconceptualize certain phenomena in recognition memory.

As a starting point, it is important to emphasize that the efficient use of memory involves more than the storage and retrieval of information mediated by an elegant menu of control processes. To use our memories optimally, we must solve a variety of problems of inference. Frequently, for example, we need to make judgments about our prior experience with a stimulus on the basis of a current encounter with that stimulus. This task may range from evaluating whether or not we have foraged a particular food location already today to trying to decide if the face on the post office wall is one familiar from a recent personal encounter or from the popular media. In either case, it is critical to accurately infer the objective status of the information in question on the basis of our current experience, perceptual or phenomenological.

It may be equally crucial to use current experience as a basis for decisions about the future. While studying for an exam, for example, a student needs to evaluate his or her ability to access information tomorrow relevant to a particular topic being studied today. In that case, again, the skillful use of current cues depends crucially on the inferential process linking current experience to probable future performance.

The problems of inference faced by the human memory system have both a cognitive and a metacognitive element. The cognitive component involves reconstructing the past based on fragments of perceptual and memorial experience that we have available at the time of remembering, as exemplified by Bartlett's classic example of the War of the Ghosts (Bartlett, 1932). In this case, the interference is one of logical continuity: The pieces of the past that we do remember inform our guesses as to the actual entire nature of the past experience. Metacognitive information is, however, equally crucial in such inferences. For example, we may have greater confidence in retrieved information that was easily accessible (Kelley and Lindsay, 1994) and we may predict more probable future recall of information that is readily retrievable now (Benjamin et al., 1998); we may even infer that we have not experienced a particular event or stimulus based on our evaluation of the likelihood of being able to remember it if we had (Brown et al., 1977; Glucksberg and McClosky, 1981; Strack and Bless, 1994).

It is in this framework that we emphasize the importance of characterizing the role of subjective fluency in decisions about memory. In the same way that it is critical to understand how prototypes and gist memories shape our specific remembrances of the past, we must also understand the metacognitive heuristics (and biases) people employ in actively reconstructing their past or constructing their potential future.

Our view of subjective fluency as one basis for judgments about future performance and past experiences is entirely compatible with the attribution view proposed by Jacoby, Kelley, and their colleagues (e.g., Jacoby and Dallas, 1981; Jacoby et al., 1989a; Kelley and Jacoby, 1998). Their work and that of others has provided abundant examples of how attributions of the source of one's current subjective experience can lead to spurious inferences of various sorts, including illusions of auditory clarity

(Jacoby et al., 1988), illusions of visual clarity (Witherspoon and Allan, 1985), illusions of fame (Jacoby et al., 1989b), aesthetic preferences (Kunst-Wilson and Zajonc, 1980), illusions of truth (e.g., Begg and Armour, 1991), illusions of problem ease (Jacoby and Kelley, 1987), illusions of knowing (e.g., Reder, 1987), and illusions of memory or pastness (e.g., Jacoby and Whitehouse, 1989). We limit our discussion to the latter two cases, in which subjective fluency influences predictions of future performance and judgments about past experience.

In Kelley and Jacoby's terms, the perceptual fluency of a stimulus is attributed to "pastness" (i.e., its having been seen earlier) when the available cues specify no other probable source. Thus, the subject attributes the sense of fluency to different sources depending on his or her ability to localize the nature of that fluency. Our attempt here is to extend such thinking to inferences both about the future and the past, and to lend a mathematical formalization to these ideas.

2. Fluency as an inferential tool

We turn now to a selective review of evidence supporting the notion that subjective fluency (of various sorts) is a vital inferential tool. A plenary review of the relevant literature is outside the scope of this article, but some representative findings may serve to illustrate the ubiquity of subjective fluency as a heuristic in inferences. More thorough analyses are provided by Benjamin and Bjork (1996), Bjork (in press), Jacoby and Kelley (1987), and Schwartz (1994). We first discuss cases in which judgments about future performance are influenced by current fluency, and we then turn to the more complicated case of judgments about the past.

2.1. Fluency and inferences about the future: Metacognitive judgments

In part, the recent blossoming of interest in metacognition is a product of researchers reconsidering the basic nature of metacognitive processes. The direction of this change has been towards an inferential view. In contrast, some early views of metacognition (e.g., Hart, 1965) postulated "privileged access" of the metacognizer to the nature and strength of his/her memories. According to such theories, predictions of one's own performance on tasks involving memory were accurate because the subject had direct access to the to-be-tested trace and could make informed decisions about later access. As the following examples will demonstrate, however, metacognitive judgments are best viewed as inferential in origin (see also Schwartz et al., in press). We argue in the next sections that perceptual and retrieval fluency are crucial components of the phenomenological experience that serve such inferences.

2.1.1. Perceptual fluency

One clear basis for metacognitive inferences is the sense of perceptual fluency we have in attending to the presentation of to-be-remembered material. That is, better future retention is predicted for stimuli that are more easily perceived in the present. Using perceptual fluency as a heuristic basis for predicting future retention is usually a good idea: More well-known material will tend to be more readily perceived at the current point in time, and that same information is more likely to be recallable in the future.

A number of authors have shown ways in which manipulating the perceptual fluency of a cue affects metacognitive judgments. The feeling-of-knowing (FOK) phenomenon provides a series of particularly good examples. Briefly, the FOK phenomenon is that subjects who are unable to recall answers to questions or previously studied responses to cue stimuli can nevertheless provide reasonably accurate estimates of their likelihood of being able to recognize the appropriate response (e.g., Hart, 1965; Nelson et al., 1982).

A number of results (e.g., Metcalfe et al., 1993; Reder, 1987; Reder and Ritter, 1992; Schwartz and Metcalfe, 1992) indicate that the perceptual priming of a cue increases FOK judgments about the retrievability of the associated response. Moreover, it does so without affecting the actual retrievability of that response. For example, Schwartz and Metcalfe (1992) showed that priming the cue in a paired-associate pair prior to attempted cued recalls elevated FOK ratings but did not affect the rates at which the unrecalled items were recognized on a subsequent test of recognition. Similarly, in comparing FOK judgments between the conditions of the classic interference theory paradigm, Metcalfe et al. (1993) demonstrated that FOK ratings were related to the number of times that A (cue) word was shown and not to the actual retrievability of the B (target) term. Thus, subjects submitted higher FOK judgments for pairs in an A–B′, A–B condition than ones in a C–D, A–B condition, despite the eventual superiority of retrieval of B terms in the latter condition.

The role of perceptual fluency in motivating rapid predictions of retrievability has also been documented. Reder (1987) showed that priming certain key words (such as *golf* and *par*) prior to a general-knowledge question such as, "What is the term in golf for scoring one under par?" increased subjects' estimates of their ability to retrieve the answer to that question. Again, the actual retrievability of the answer was unaffected by the cue-word preexposure. Reder and Ritter (1992) demonstrated a similar effect: Subjects' predictions of answer retrievability (as opposed to computability) for arithmetic problems were strongly affected by preexposure of the problem operands, even when the problem operator – and thus the entire nature of the problem – was altered from that prior exposure.

Other types of metamnemonic judgments are also affected by perceptual fluency. One striking example is provided by Begg et al. (1989), who showed that predictions of recognition performance were higher for high- than low frequency words, despite the fact that, in actuality, low-frequency words are recognized more accurately than high-frequency words. Begg et al. postulated that this result derived from the greater "ease of processing", or perceptual fluency, for the high-frequency words.

It is evident from these and other examples that the perceptual fluency of a stimulus is a salient cue utilized in making inferences about the future retrievability of that stimulus or other related information. It is worth noting that such fluency may arise from a number of sources, some of which may not involve actual prior perceptual presentation of the stimulus (e.g., Masson and MacLeod, 1992). We discuss next the use of retrieval fluency in such judgments of future retrievability.

2.1.2. Retrieval fluency

The ease with which we can currently retrieve a piece of information from long-term memory also serves as a basis for judgments of future retrievability. As we argued in the case of perceptual fluency, this strategy is also generally useful, in that current retrievability is quite likely to be a good predictor of retention and thus future retrievability. In general, more well-known material will be more readily accessible now, and will also be so in the future. However, because fluent retrieval can derive from sources other than those that foster future remembering, the interpretation of subjective retrieval fluency is subject to the same attributional demands as those implicated in the case of perceptual fluency.

One example of how retrieval fluency supports FOK judgments is the recent work of Koriat (1993, 1995). In Koriat's theory, the FOK judgment is made on the basis of the raw amount of information that comes to mind in response to the original recall cue. The judgment of future recognizability is thus an inference made on the basis of current phenomenological experience. To test this hypothesis, Koriat developed a series of questions that differed both in the number of different answers that they evoked in a sample of subjects and the proportion of those answers that tended to be correct. Some questions (*consensually correct* items) tended to evoke primarily correct answers; others, such as "What is the capital of Holland?" evoked a high proportion of incorrect answers (*consensually wrong* items). Koriat showed that FOK judgments were accurate only for consensually correct and not consensually wrong items, thus supporting his argument that FOK are accurate only insofar as future recognizability does in fact correlate positively with the amount of information retrieved in response to a cue (Koriat, 1995).

This example demonstrates how predictions of future performance are not based on any kind of objective access to the memory trace in question, but are actually inferences based in part on current experience. Koriat has demonstrated convincingly that FOK judgments are highly related to the raw amount of information that "comes to mind", independent of whether that information is in fact related to the target of the memory search: Judgments are accurate for cases in which this raw amount of information retrieval is positively correlated with probability of successful target access, but not for cases when that retrieval is misleading.

In addition, Benjamin et al. (1998) have shown that judgments of future retrievability (also called judgments of learning; JOL) correlate highly with the current retrievability of that item, even in cases for which actual future retrievability is negatively related to that current access. Subjects in their experiments used this retrieval fluency heuristic even when their current retrieval was "enhanced" by fleeting factors such as recency effects, and also when predicting retrieval from a qualitatively different memory store than the one they were currently accessing (episodic versus semantic memory). As an example of this latter case, subjects attempting to predict their own later free recall of answers produced in response to general-information questions often used the ease with which they had accessed that answer in response to the question as a basis for their prediction, despite the fact that ease of access is negatively correlated between the two cases (Gardiner et al., 1973). In other words, subjects predicted probable later access for answers provided quickly, but less probable access for answers that took longer to provide. Consequently, they predicted high levels of recall for items that they did *not*, in fact, recall well later, and low levels for items that they *did* recall well later. This work demonstrates the degree to which using the retrieval fluency heuristic as a substitute for an accurate model of memory can mislead metacognition.

In a priming paradigm, Lee et al. (1993) showed that a subthreshold presentation of a to-be-recalled target word immediately prior to a JOL elevated the predictions of future retrievability, but – because of the short-lived nature of such priming – did not affect actual later retrievability. Our interpretation of this finding is that subjects utilized the momentary accessibility of the target word (read: retrieval fluency) as an inferential index of future accessibility. Subjects' inability to recognize the source of the elevated retrieval fluency leads them to attribute that easy access to a state of learning that they have not, in fact, achieved. Presumably, if the word had been pre-presented at a suprathreshold level, the appropriate attribution would have been possible, and the JOLs would have been misled. [1]

These examples illustrate the variety of ways in which current retrieval fluency is used in generating metacognitive predictions of future performance. While the utility of such a strategy is overall quite high, it is also evident that the strategy fails in situations in which fluent retrieval derives from some source other than well-learnedness.

The use of both perceptual and retrieval fluency in making accurate predictions about future performance thus hinges on the ability to specify appropriately the source of that subjective experience. In particular, to the degree that the fluency of a perceptual or memorial experience cannot be attributed to spurious factors unrelated to long-term retention, such as priming, that fluency is believed to reveal something about the degree of learning and thus serve as a predictor of future retention.

2.2. Fluency and inferences about the past

In this section, we outline examples in which similar reasoning is used in the generation of inferences about past experience. To the degree that we are unable to attribute the sense of subjective fluency in the processing of a stimulus to any of the "fleeting" factors mentioned above, we may also use that experience as a basis for deciding something about the nature of our experience with that stimulus in the past.

Many of the studies we cite here concern a particular type of judgment about the past, namely a judgment of *recognition*. Before turning to a formal conceptualization of the recognition process, we characterize cases in which perceptual and retrieval fluency appear to play a role in recognition judgments.

[1] There may even be some reason to believe that JOL for superthreshold presented target primes would be lower than for the nonprimed items, in that the easy accessibility of those items that are in fact well-learned may be mistakenly attributed to the recent priming exposure (cf. Jacoby and Whitehouse, 1989).

2.2.1. Perceptual fluency

One potent example of how perceptual fluency affects judgments of recognition is provided by the work of Jacoby and Whitehouse (1989). In their experiments, as in the typical recognition memory experiment, subjects studied a list of words and were later required to discriminate between those and other nonstudied words by means of a recognition test. However, for some trials, a word was exposed immediately prior to the presentation of a given test item. This presentation was either a very rapid, subthreshold one (as in the Lee et al. experiment discussed above) or a suprathreshold one that the subjects could quite easily perceive. Furthermore, the presented word was either the same (match condition) as the to-be-presented recognition item or different (nonmatch condition).

In two separate experiments, Jacoby and Whitehouse found that false-alarm rates (the proportion of test trials on which subjects claim oldness, but the item is in fact new) were higher in the match than the nonmatch condition following a subthreshold presentation, but higher in the nonmatch than match condition when following a suprathreshold presentation (cf. Luo, 1993).

Jacoby and Whitehouse's interpretation of the findings is consistent with the attributional framework discussed earlier. In both subthreshold-match and suprathreshold-match presentation conditions, the perceptual fluency for the to-be-tested item is increased by its recent exposure. However, the subthreshold presentation did not allow subjects to localize accurately the source of that enhanced fluency, and it is instead attributed to oldness (i.e., having been studied). In the condition for which the prior presentation was suprathreshold, subjects were able to attribute the fluency to that prior exposure, and were not misled into mistakenly according that fluency to oldness. In fact, the tendency for subjects to false-alarm more in the suprathreshold-nonmatch than in the suprathreshold-match condition suggests that they are attributing fluency to that prior exposure even when it derives partially from actual prior study. Further support for this interpretation is provided by the fact that hit rates (the proportion of test trials on which subjects claim oldness for a truly old item) are higher in the suprathreshold-nonmatch than the suprathreshold-match condition.

Other data also bear on the issue of whether perceptual fluency is utilized in the judgment of recognition. Johnston et al. (1985) showed that words and nonwords that were quickly identified in a perceptual identification paradigm (and thus more perceptually fluent) were more likely to be judged as old than items that were not quickly identified. Moreover, this contingency held for both truly old and truly new items. Their claim (see also Johnston et al., 1991) has been that perceptual fluency is a source of information, but not an exclusive one, utilized in making recognition decisions.

One potential difficulty is to reconcile the findings of Johnston and his colleagues with results indicating dissociations between recognition memory and perceptual priming (Hayman and Tulving, 1989; Snodgrass et al., 1996) as well as failures of Johnston et al. (1991) to replicate their basic effect. The seemingly contradictory nature of these findings is addressed later in this analysis.

Kihlstrom et al. (1996) have suggested that perceptual fluency may indeed mediate recognition-memory judgments, and have provided additional evidence that encouraging such reliance may increase overall recognition accuracy. In a sample of subjects with ECT-induced amnesia, Dorfman et al. (1995) showed that recognition accuracy increases with instructions to adopt a liberal criterion for oldness. This direction presumably increased subjects' reliance on cues that were available despite their amnestic state, such as perceptual fluency.

As a final example of how perceptual fluency affects judgments about past experience and, in particular, recognition, Whittlesea, 1993; Whittlesea et al., 1990 has shown experimentally that manipulating the fluency with which a test item can be perceived – for example, by superimposing a variable-interference mask – directly affects the likelihood of calling that test item old. Words that were presented with greater clarity were more likely to be called old, regardless of whether that item was in fact old.

2.2.2. Retrieval fluency

The ease with which we retrieve information from memory also appears to be an index by which we gauge the nature of past experience with that information. One clear example is provided by Kelley and Lindsay (1994). In their experiment, subjects were asked to answer a number of general-information questions (e.g., "What was Buffalo Bill's last name?"). Immediately prior to this task, a number of words were pre-exposed to this subjects, some of which were correct answers to those upcoming questions (Cody), some of which were plausible but incorrect answers to those questions (Hickock), and some of which were unrelated words (Letterman). Subjects' rated confidence in the answers they provided was higher for pre-exposed than non-pre-exposed answers, regardless of whether those answers were correct or not! Our interpretation of this finding is that the elevated retrieval fluency for the pre-exposed words led subjects to accord those answers greater confidence.

Shaw (1996) provides a similar result in the domain of eyewitness memory. Subjects in his experiment viewed a series of slides that portrayed a mock crime scene, and were later questioned as to various details about the scene. Some subjects were then given an opportunity to carefully consider the answers that they provided, and other engaged in unrelated distractor tasks. After this period of consideration or distraction, subjects were then asked the same questions again and asked to rate their confidence in those answers. Subjects in both groups typically provided the same answers as they had initially, but those who had the opportunity to consider their answers tended to assign higher confidence to the answers. Presumably, the elevated retrieval fluency for those answers that were thought over more carefully mediated the greater confidence in the correctness of those answers.

Morris (1990) has also shown that confidence in recalled information varies with the accessibility of that information. In one of his experiments, subjects recalled information from expository passages. Confidence was affected by the experimentally - manipulated accessibility of the relevant information (by self- versus other-generated probes; see, e.g., McDaniel et al., 1988) and correlated significantly with retrieval latency. These findings further support the claim that retrieval fluency is a basis for confidence.

2.3. The conceptualization of fluency and inference

The studies we have cited here all demonstrate that subjective fluency is an important source of information for inferences about future performance and past experience. Critical to conceptualizing this "fluency heuristic" is understanding the manner in which fluency is incorporated into these judgments. In the section that follows, we outline one theory of how such judgments may be made. For the ease of exposition, we concentrate on only one portion of the taxonomy we have outlined here, the way in which perceptual fluency influences judgments of recognition. We attempt to illustrate the utility of our formalization in that context, but hope to shed light on the more general use of subjective fluency as an inferential tool.

3. Perceptual fluency and judgments of recognition

We have already cited a number of cases in which perceptual fluency appears to play a role in judgments of recognition. In this section, we attempt to operationalize the manner by which these judgements are made. To motivate our theory, we first digress to consider the nature of the task of recognition in general.

3.1. The task of recognition

The subject in the typical recognition-memory experiment is faced with the task of deciding, for each item in a series of presented stimuli, whether that item has been studied in a particular delimited episode in the past. In much of the recognition-memory research that has been conducted, the stimuli are words, and the task is to differentiate between words that were seen during a study phase in the experiment and those that were not. In that sense, every test of recognition involves an element of source monitoring (Johnson et al., 1993), in which the subject must not only remember having seen the item (or not), but must be able to localize it sufficiently within time and space as to make a decision about whether it was seen during the study phase of the experiment.

Certain historically prominent theories of recognition memory (e.g., Atkinson and Juola, 1974; Mandler, 1980) as well as some more current theories (e.g., Gardiner, 1988; Jacoby, 1991) postulate an inferential basis for the recognition decision. This position stands in contrast to a more "memory-based" view of the process, in which decisions simply reflect something about the nature of the prior experience (or lack thereof) with the to-be-recognized stimulus (as in recollection; e.g., Hintzman, 1988). One particular subclass of inferential theories (often termed dual-process theories) incorporate both familiarity – a current phenomenological sense of pastness – and recollective processes. We have avoided use of this nomenclature thus far in our analysis because of the potential ambiguity associated with the term.

The term *familiarity* is open to multiple interpretations, yet one seemingly undisputed aspect of the sense of familiarity is its rather coarse nature: While it may indeed provide for an accurate sense of pastness, or lack thereof, it fails to provide any information about the nature of past experience with that stimulus. It is phenomenologically impossible – perhaps only by a matter of definition – to have a sense of specific familiarity. If the sense of familiarity leads us to a particular place in space or time, then that information is essentially recollective in nature.

Thus, the utilization of familiarity in formulating a recognition decision – in which we are asked to localize the presentation of a stimulus in time and space – depends critically on the appropriate attributions of the source of that familiarity. Recollective information must be utilized in order to infer the source of the familiarity and, as in many of the examples cited earlier, failure to do so accurately quite often leads to failures of recognition.

Formally, the problem of distinguishing between new and old items on a test of recognition has been classified as a problem of *signal detection*. Because of past experience, items differ in their general levels of familiarity, and the distributions of new and old items overlap in their familiarity. Some items not viewed during the study episode may be highly familiar by virtue of prior experience, and some words that were studied are not as familiar because their prior familiarity was low.

In classical signal-detection theory (SDT; e.g., Green and Swets, 1966), the observer makes decisions about the presence or absence of a stimulus based not on a sensory threshold, but rather on a likelihood ratio. This ratio functionally compares the current strength of a sensory input to a putative distribution of strength given the to-be-detected signal and to a distribution of strength given only noise. The exact nature of such a comparison will be discussed more thoroughly in the following sections. The analysis of signal detection disconfounds true recognition accuracy – the standardized distance between the means of the old and new item familiarity distributions (d') – from response biases (β) that differ as a function of experimental conditions and individual differences (for a particularly clear exposition of signal detection theory and the role of likelihood ratios, see Coombs et al. (1970), Ch. 6).

Although the theory of signal detection has often been fruitfully applied to recognition memory, its usage has been somewhat limited to the standard recognition memory paradigm in which to-be-recognized test items truly differ in some underlying "strength" criterion. That is, because studied items presumably enjoy greater enhancement of familiarity than nonstudied items, the recognition test may indeed involve nothing more than an inference based on the likelihood ratio discussed earlier. The limitations of such an analysis are evident when considering data such as those from the suprathreshold presentation condition of Jacoby and Whitehouse (1989). Subjects are clearly able to reject certain unstudied recognition foils, high levels of familiarity notwithstanding, as long as the source of that familiarity can, via recollective processes, be attributed to exposures independent of the study episode.

We attempt here to provide an extension of the signal detection framework to situations where multiple recognition processes operate and jointly specify a recognition response. In doing so, we hope to sketch a theory that captures the frequently inseparable relationship between recollection and perceptual fluency. As we argue later in this analysis, recollective capacity is key in the appropriate attribution of perceptual fluency: Perceptual fluency may or may not bias our recognition decision depending on the believed source of that fluency.

One manner in which such an integration can be made is to adapt traditional SDT into a multidimensional framework (see, e.g., Ashby, 1992). Our approach is somewhat less general, but provides a method by which standard SDT can be applied to experimental circumstances in which multiple presentations of the stimuli are possible. In particular, it provides a way to incorporate critical recollective information concerning the source of an exposure (experimental or extraexperimental) into the evaluation of perceptual fluency and its consequent use in making recognition decisions. The theory thus provides a new conceptualization of the relationship between recollective processes and familiarity (or, in our terms, perceptual fluency).

A second salient aspect of the theory we propose lies in the nature of the likelyhood evaluation process. In standard SDT, a stimulus is compared to predetermined distributions of potential stimulus strength given signal and given noise (or target and foils, respectively, for the case of recognition). In the theory presented here, the current perceptual fluency of an item is compared to expected perceptual fluency given both unexposed and exposed status for that particular item. The variance of the distributions corresponds not to item differences, but rather exposure-to-exposure variations in perceptual fluency for that individual item.

In is general flavor, the theory we present is similar to the elegant model proposed and tested by Glanzer et al. (1993). In Glanzer et al.'s theory, the number of "marked" features that characterize a given test item is compared to the expected distribution of marked features given studied or unstudied status. The subject's expectations are not governed by an individual item analysis, however, but rather by experimental conditions. Different conditions, whether determined by item type, such as high versus low frequency, or by manipulations of the conditions of study, such as short versus long exposure time, are assumed to influence the sampling (marking) of features at study. In contrast to the present theory, the number of marked features that unstudied items have accumulated from sources other than the study episode is not assumed to vary by item type.

For our purposes, signal detection analysis is a useful tool for describing the manner in which subjects discriminate between old and new items on the basis of perceptual fluency. However, as we have noted before, information utilized in the recognition decision is not limited to perceptual fluency. Recognition also involves a host of factors unrelated to perceptual fluency, including (but not necessarily limited to) recollective information, knowledge about the experimental constraints on the study stimuli, and knowledge about the relative frequencies of studied versus unstudied items on the recognition test. What is needed, then, is a way of formulating how perceptual fluency is considered jointly with other information in the judgment of recognition. Our hope is that, by lending a formal interpretation to this one aspect of the recognition process, we can provide an interface to other potential sources for recognition judgments. Later in this article, we speculate on what these other bases for recognition may be.

As we noted earlier, the analogy that we have drawn with signal-detection theory is not complete. Most importantly, signal detection is typically utilized as a model of performance across many trials, and the parameters are consequently interpretable as meaningful characteristics, such as accuracy or bias. In our formulation, the critical process of estimating fluency given oldness and newness is done on an item-by-item basis. In that sense, we are using the signal-detection framework as a rough analogy for a process model.

4. A Bayesian account of recognition

The minimal theory that we have sketched here can be instantiated in the form of Bayesian inferences. Assuming that the task for the subject is to evaluate the status of an item given a particular level of perceptual fluency, the problem can be formalized as

$$p_i(s|F=f) = \frac{p_i(F=f|s)p_i(s)}{p_i(F=f)} \tag{1}$$

in which the probability of an item having been studied given a particular degree of fluency for an item $[p_i(s|F=f)]$ is equal to the product of the probability of that level of fluency given studied status $[p_i(F=f|s)]$ and the prior probability of study $[p_i(s)]$ divided by the baseline probability for the particular level of fluency of that item. The subscript i is used to denote that the analysis is carried out on an item-by-item basis, and that the identity of a particular test item matters in the evaluation of each of these terms.

This analysis, however, is incomplete; most critically, it lacks a mechanism to incorporate baserate information about fluency for the to-be-evaluated item. This information is carried in the equivalent formulation for the evaluation of unstudied status given a particular degree of fluency,

$$p_i(u|F=f) = \frac{p_i(F=f|u)p_i(u)}{p_i(F=f)}. \tag{2}$$

In Fig. 1 is seen a graphical depiction of the likelihood-estimation process represented in Eqs. (1) and (2). For every to-be-recognized item, an expected degree of fluency is estimated and is represented as the mean of the unstudied (left) distribution. Variations in fluency between encounters with an item and estimation error are incorporated by establishing a probability distribution about the estimated mean. How such a distribution is estimated is an empirical question, and the exact nature of this

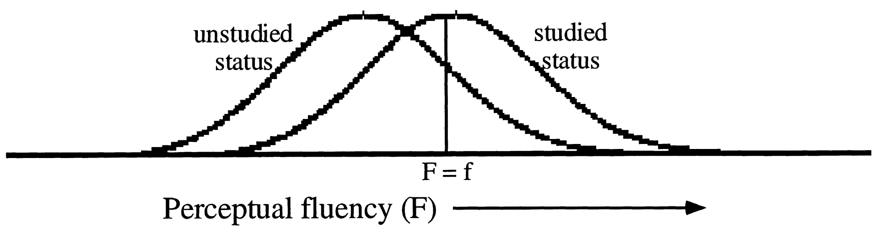

Fig. 1. Probability distribution functions of perceptual fluency for a given item given both putative unstudied and studied status.

process, as well as the form of the resulting distribution, are not specified by our theory. A probability distribution representing the fluency of the item if it was studied is estimated in an equivalent manner. These distributions are shown to be normal in form, as is often assumed within the signal-detection formulation, but they need not be. These distributions are then used to assess $p_i(F=f|s)$ and $p_i(F=f|u)$. The values are represented by the heights of their respective distributions at the point f.

The two evaluations of the probabilities of studied and unstudied status must then be combined into a decision rule that specifies which response is to be made. In particular, the Bayesian estimator can use the ratio of the two – a quantity called the *Bayes factor* – to choose between studied and unstudied status. This ratio is instantiated and simplified to yield the general model:

$$B(S{:}\,U)=\frac{p_i(F=f|s)p_i(s)}{p_i(F=f|u)p_i(u)}, \tag{3}$$

in which B is a scalar to be compared to some decision criterion, presumably 1.

However, assessment of unstudied status via Eq. (2) fails to account for cases in which the perceptual fluency of a stimulus is enhanced by an encounter from an episode other than the target episode. In other words, this formulation does not reflect the ability to discount enhanced perceptual fluency when it stems from sources other than exposure in the study list. To incorporate such a mechanism, the term representing the probability of a particular level of fluency given unstudied status must be expanded to include cases in which the item is truly new (i.e., unstudied and unexposed, u_{u}) and those in which it was exposed but not during the study episode (unstudied but exposed, u_{e}). This combination can be represented as

$$p_i(F=f|u)=p_i(F=f|u_{\mathrm{u}})p_i(u_{\mathrm{u}})+p_i(F=f|u_{\mathrm{e}})p_i(u_{\mathrm{e}}), \tag{4}$$

in which the probability of $p_i(F=f|u)$ is partitioned into $p_i(F=f|u_{\mathrm{u}})$ and $p_i(F=f|u_{\mathrm{e}})$.

Eq. (4) combines linearly the two sources of evidence for deciding that an item was not studied. First, the expected fluency given newness is evaluated; second, in cases in which a recent exposure other than the study list is recollected (a foil exposure), the expected fluency given that exposure is evaluated. Modifying the general model reflected in Eq. (3) to account for situations in which test items may have had a prior presentation in an episode distinct from the study episode yields

$$B(S{:}\,U)=\frac{p_i(F=f|s)p_i(s)}{[p_i(F=f|u_{\mathrm{u}})p_i(u_{\mathrm{u}})+p_i(F=f|u_{\mathrm{e}})p_i(u_{\mathrm{e}})]p_i(u)}. \tag{5}$$

Values larger than one indicate a preponderance of evidence for studied status; values less than one imply greater evidence for unstudied status of the item in question. This instantiation of the model can be simplified by assuming that a prior exposure – either during the study episode or during an unrelated "foil" exposure – has the same effect on later perceptual fluency. Thus, we equate $p_i(F=f|s)$ and $p_i(F=f|u_{\mathrm{e}})$, and alter the nomenclature to yield

$$B(S\colon U) = \frac{p_i(F = f|e)p_i(s)}{[p_i(F = f|u_\mathrm{u})p_i(u_\mathrm{u}) + p_i(F = f|e)p_i(u_\mathrm{e})]p_i(u)} \tag{6}$$

in which $p_i(F = f|e)$ represents the probability of a particular degree of fluency given an exposure, whether it derived for exposure during the study list or an unrelated episode. For ease of interpretation, note that Eqs. (5) and (6) are both subsumed under the more general model that combines the two contributions to the estimates of an item not having been studied, exemplified by Eq. (3).

To complete the analysis, provision must be made for the fact that a studied item may also have been exposed in an unrelated episode prior to the test. The case illustrated in Eqs. (5) and (6) only accounts for unrelated exposures of a future distractor item, but it is also possible for a future target to have an additional exposure prior to the recognition test. Thus, the term $p_i(F = f|s)$ must be expanded to incorporate this contingency:

$$p_i(F = f|s) = p_i(F = f|s_\mathrm{u})p_i(s_\mathrm{u}) + p_i(F = f|s_\mathrm{e})p_i(s_\mathrm{e}). \tag{7}$$

In this formulation, s_u indicates study items that were not exposed (other than during the study period), and s_e represents items that were both studied and exposed in an unrelated episode. This distinction is in some respects equivalent to the one drawn for the case of unstudied items, but the subscripts u and e now correspond to one and two total exposures, respectively. For unstudied items, those terms mapped onto zero and one exposure.

This discrepancy introduces the additional complication of having to consider the relationship between number of exposures and perceptual fluency enhancement. The simplest treatment assumes that an item either has its fluency enhanced by an exposure or not; multiple exposures are essentially treated as one. In this formulation, $p_i(F = f|s_\mathrm{u})$ is equivalent to $p_i(F = f|s_\mathrm{e})$, and the entire model reduces to Eq. (6). If perceptual fluency is assumed to bear some relationship (presumably positive monotonicity) to the number of exposures, then the model becomes

$$B(S\colon U) = \frac{[p_i(F = f|e)p_i(s_\mathrm{u}) + p_i(F = f|s_\mathrm{e})p_i(s_\mathrm{e})]p_i(s)}{[p_i(F = f|u_\mathrm{u})p_i(u_\mathrm{u}) + p_i(F = f|e)p_i(u_\mathrm{e})]p_i(u)}. \tag{8}$$

In this form, the characterization of the likelihood-estimation process illustrated in Fig. 1 must be augmented by an additional distribution corresponding to estimated fluency for two exposures. Again, because the model treats the effects of one presentation on perceptual fluency equivalently – whether it occurred during the study episode or an unrelated exposure period – three distributions (for 0, 1, and 2 exposures) characterize the four cases (unstudied and unexposed, unstudied but exposed, studied and unexposed, and studied and exposed).

Using the likelihood ratio as a basis for the decision rule is consistent with the signal-detection framework discussed earlier. A psychological interpretation of the model is that the decision between oldness and newness of a tested item is made as a joint function of a ratio of the two necessary factors described earlier. The first ratio contrasts the probabilities of fluency of the degree experienced for this item given studied and unstudied status of the involved item. This value incorporates the three basic sources considered earlier: expected "pure" baseline fluency, enhancements to fluency unrelated to the study episode, and enhancements as a function of study. The former two considerations are incorporated into the value of $p_i(F = f|u)$, and the latter into $p_i(F = f|s)$.

The other ratio, often referred to as the prior odds ratio, incorporates all recollective sources of evidence of oldness or newness of the item in question. Because our focus is on the nature by which perceptual fluency is utilized in judgments of recognition, and not on delineating a canonical model of recognition memory, we only briefly speculate as to the sources that might serve such a judgment.

4.1. Bases for the recognition judgment

By our cursory analysis, if a recognition decision is to incorporate perceptual fluency information, it must also incorporate the following factors.

4.1.1. Knowledge of the item set

A judgment of recognition may be based on information about the nature of the studied list. For example, if one knows that the studied list contained only names of flowers, and the current test item is "hamburger", then one can confidently reject that item in the absence of any additional information, including perceptual fluency.

4.1.2. Relative proportions of targets on the test

While not strictly necessary for making a recognition decision, it has been shown (e.g., Egan, 1956) that a subject's bias in recognition is affected by the proportion of targets on the test. In the limit (when the proportion of targets on the test is either 0 or 1), this heuristic strongly specifies a response even prior to the presentation of the to-be-judged item. In that case, the utility of the inference utilizing perceptual fluency is again obviated, this time by virtue of prior odds. To the degree that this proportion deviates from 0 or 1, the potential to use other sources of information – such as perceptual fluency or recollection – is increased.

4.1.3. Recollective information

Recollective information, as discussed earlier, involves specific episodic information relevant to the study phase that is retrieved from memory. While such information *may* support a recognition decision on its own, certain evidence suggests that recognition is not limited to recollection. The finding that the effects of word-frequency are opposite between tests of recall and tests of recognition is one such example (Gregg, 1976). Other evidence implicates perceptual fluency specifically as an additional factor (Johnston et al., 1985).

However, it is important to understand how recollective information is critical in the *interpretation* of perceptual fluency. As discussed earlier, perceptual fluency itself carries only evidence of pastness or not; it cannot specify a particular event as its cause. Specific recollections allow the sense of fluency to affect recognition decisions meaningfully. In Kelley and Jacoby's terms, the *attribution* of perceptual fluency – which can clearly affect the outcome on a recognition decision – depends on recollective ability. In the Jacoby and Whitehouse study, perceptual fluency is only attributed to pastness when no specific recollections support an alternative interpretation. When the recollective information supports an alternative source of perceptual fluency, this attribution is no longer made. The Bayesian characterization that we propose integrates the process of perceptual fluency information and the necessary underlying recollective processes.

4.2. Fluency estimation

If none of the factors above can support a recognition decision on its own, a decision strategy involving perceptual fluency may be used. First, the perceptual fluency of a recognition probe must be evaluated with respect to its expected fluency. That is, the current perceptual fluency can only be assessed relative to a baseline level of fluency expected for that particular stimulus. Thus, while the word "skep" may enjoy less perceptual fluency than the word "lift" in an absolute sense, the accurate incorporation of perceptual fluency in an informed judgment of recognition requires that the levels of fluency be evaluated relative to the fluency one might expect for that word if it had not been studied earlier.

This estimation involves considerations of word frequency and personal salience, as well as evaluating other recent encounters with the word. As Jacoby, Kelley and their associates have argued, and we have echoed throughout the course of this paper, one crucial aspect of such an inference is being able to discount accurately perceptual fluency to the degree that it derives from sources other than having occurred in the target episode. Thus, a consideration of those other encounters with the test word that might influence the current perceptual fluency allows for an increase in the expected baseline fluency for that word. In addition, the current level of perceptual fluency must also be evaluated relative to what one would expect for that word if it was studied earlier. In other words, subjects must possess some theory as to the effects of prior study on perceptual fluency. This estimate of expected fluency for an old item thus reflects a confluence of both baseline expected fluency and the predicted perceptual enhancement owing to study.

4.3. Recognition memory revisited

The theory of recognition that we have sketched here is not unlike other theories that have been proposed (e.g., Johnston et al., 1985; Mandler, 1980). However, we do postulate that the incorporation of perceptual fluency into such judgments takes the form of a Bayesian analysis. We believe that a theory must take this form in order to account for substantial differences in baseline fluency for to-be-tested items. This position stands in contrast to that of Glanzer et al. (1993), who propose that baseline differences influence encoding and attention at study.

In fact, the advantage in recognition for items of normatively low frequency makes clear that any model proposing a monotonic mapping between perceptual fluency and recognition decision is doomed from the outset. Mandler's demonstration of a nonmonotonicity between normative frequency and recognition accuracy (Mandler et al., 1982) further complicates the picture for any model that does not incorporate both an estimate of expected baseline fluency and current experienced fluency.

4.3.1. Psychological interpretations of the model

Eqs. (3), (6) and (8) instantiate the manner in which we propose that subjective perceptual fluency is incorporated into judgments of recognition. While the general form of Eq. (3) is appealing from the perspective of being able to separate the contributions of perceptual fluency from other information in a recognition decision, Eqs. (6) and (8) make clear that they are fundamentally inseparable, much in the spirit of the traditional Bayesian analysis.

The psychological interpretation Eqs. (6) and (8) thus necessarily involves both fluency estimation and other recollective information. Perceptual fluency supports a judgment of pastness to the degree that the experienced fluency is greater than expected fluency and cannot be attributed to enhancements other than from exposure during the study episode. Each of these components – experienced fluency, expected fluency, and potential enhancements – is jointly evaluated by estimation and recollection, that is, we estimate (based on experience and theories about the stimulus) what degree of perceptual fluency we are experiencing, what degree we should experience if we had studied the item in question, and what degree we would experience if it had been exposed in an unrelated recent episode. Each one of these factors is weighted by the extent to which recollected information supports a studied status, an unstudied and unseen status, or an unstudied but seen status for the item in question.

4.3.2. The recollection ratio

As mentioned earlier, remembering the presentation of a particular stimulus is enough to support an old response in the absence of perceptual fluency information. The task of estimating the relative contribution of this source has been attempted in recent years using both empirical (as in the remember/know paradigm; Gardiner, 1988; Tulving, 1985) and quantitative methodologies (as in the process-dissociation procedure; Jacoby, 1991). It is not our goal to argue the relative merits of these approaches, but we do wish to emphasize that, in both cases, familiarity is posited as the alternative source from which a recognition judgment can stem. We hope that our attempts to delineate how stimulus-driven estimates of perceptual fluency may support recognition complement these models of recognition memory. In particular, we have proposed a process in which recollective information supports the use of perceptual fluency in formulating a recognition judgment.

4.3.3. Two processes or one?

Much of the theoretical debate as to the nature of recognition memory has centered around whether one putative psychological process is sufficient to account for recognition performance, or if two are necessary. *Global memory models* (e.g.,Gillund and Schiffrin, 1984; Hintzman, 1988; Murdock, 1993) postulate a single process, although that process may be quite elaborate (e.g., Murdock, 1993). Dual-process models, as noted above, typically specify both a recollective and a familiarity basis for the recognition decision (e.g., Jacoby and Dallas, 1981; Mandler, 1980). Our contention is that two or more processes are probably at work, corresponding to the bases for recognition judgments described earlier. However, the model that we have outlined here also employs two distinct processes. The first involves the evaluation of expected and experienced perceptual fluency, as well as the computation of the corresponding distributions. The second process utilizes recollective information to weigh the outcomes of the first process, and corresponds to the notion of attribution. To reiterate an earlier point, a greater degree of perceptual fluency than would be expected is only diagnostic for recognition when the subject cannot recollect an extraexperimental even that supports such enhanced fluency. A prominent aspect of the model is thus the dual manner in which recollective information is used. First, recollecting the event in question can support a recognition judgment on its own, as when a subject episodically remembers studying a particular word. Second, recollection can also serve to discount the degree to which perceptual fluency is treated as evidence for prior study, as when a subject sees her own name as a test item and, despite the high degree of perceptual fluency for that name, does not endorse it as a studied item. (Brown et al., 1977).

4.4. Interpreting recognition phenomena using the model

In this section, we outline several phenomena in recognition memory literature and discuss the application of the Bayesian model towards understanding these effects.

4.4.1. Word-frequency effects

The model itself can, in principle, handle the superiority of low- over high-frequency item recognition in two distinct ways. The first involves the postulation that the recollection ratio supports recollection of low- over high-frequency stimuli. This interpretation is unlikely, however, given the quite distinct advantage high-frequency items possess in tests of recall. The questionable additional assumptions required to "force" the data upon the model make this particular explanation unappealing.

The alternate explanation involves the perceptual fluency ratio. While perceptual fluency itself is higher for high-frequency items, the ratio of experienced to expected fluency is quite likely to be higher for low-frequency items. Assuming that increments in perceptual fluency accrue in a negatively accelerated fashion with respect to baseline fluency, it is clear that the deviation from baseline (expected) fluency for studied items will be greater for those items of initially lower frequency. Furthermore, because of the greater discrepancy between studied and unstudied fluency for low- than high-frequency items, the model also predicts greater rates of false-alarms for high- than low-frequency items. This relationship can be seen in Fig. 2. Consider first the case of actually studied items. For those items that fall to the right of the intersection of their fluency distributions (as studied items are likely to), the ratio of $p_i(F = f|s)$ to $p_i(F = f|u)$ is larger for the low- than the high-frequency distributions. However, for items that fall to the left of that intersection (as unstudied items are likely to), the ratio of $p_i(F = f|s)$ to $p_i(F = f|u)$ is larger in the high- than the low-frequency case.

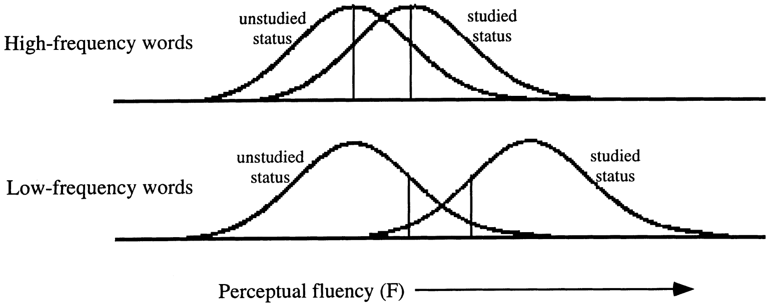

Fig. 2. Pairs of probability distribution functions of perceptual fluency corresponding to low- and high-frequency words.

In both cases, this ratio directly influences the probability of a "yes" response. Thus, this model can handle the mirror effect (e.g., Glanzer et al., 1993) simply by postulating the relatively uncontroversial assumption of a negatively accelerated relationship between baseline perceptual fluency and enhancement owing to exposure. This assumption is further supported by the finding that the perceptual identification of a low-frequency word is enhanced more than a high-frequency word by a recent study episode (Jacoby and Dallas, 1981).

The attention/likelihood theory of Glanzer and his colleagues explains the mirror effect in a qualitatively similar manner, but proposes that frequency differences arise as a function of the greater attention allocated to more novel stimuli (e.g., low-frequency words) at the time of study. Our explanation achieves such differences between low- and high-frequency items by assuming that items already high in perceptual fluency (e.g., high-frequency items) have less to gain from an additional exposure than do items lower in frequency, and thus perceptual fluency. These two explanations may prove difficult to distinguish in mathematical terms, but Glanzer et al.'s assumption appears contrary to the result that high-frequency words are more accurately *recalled* than low-frequency words.

Note also that the model makes the prediction that the relative advantage of low-frequency items should disappear when the perceptual fluency ratio is made nondiagnostic by a prior exposure of foil. Thus, recognition accuracy should not differ as greatly as a function of word frequency (either in hits or false alarms) when subjects are attempting to discriminate an exposure from one of several study lists, as opposed to just one.

4.4.2. Dissociations between recognition and perceptual identification

The ambiguous relationship between perceptual priming and recognition memory discussed earlier has a straightforward interpretation in the present framework. Because information about perceptual fluency is only utilized under circumstances in which recollective information can support accurate (or, occasionally, inaccurate) attributions of the source of that fluency, there is no reason to expect that those manipulations that produce priming will necessary lead to a greater proclivity to endorse those items on a test of recognition. It is only in cases in which such priming cannot be unambiguously attributed to events other than the study episode that such manipulations will necessarily affect recognition.

Jacoby and Dallas (1981) documented a number of manipulations that affected recognition, some of which influenced the task of perceptual identification (PI) similarly and some of which left PI unaffected. The task of PI involves the "recognition" of a stimulus presented for a very short exposure duration. The recognition in this sense is simply being able to perceive and read the word, and thus differs from recognition as discussed earlier, in which the attempt is made to discriminate whether or not that word was studied or not.

We presume that manipulations that affect the numerator of our fluency ratio in our model also affect the likelihood of PI for the stimulus. Moreover, because PI is typically evaluated using a measure of priming (the difference in rates of perception between the studied and unstudied items), this score corresponds in our analysis to the fluency ratio. However, the relationship between the two measures should be expected to be monotone but nonlinear, as the model uses a ratio score and the priming measure a difference score. Using this interpretation, we see that those manipulations that affect recognition but not PI – among them, depth of processing, reading as opposed to generating material, and study time – affect the recollective ratio but leave the perceptual fluency ratio unaffected. Jacoby and Dallas' finding that spacing of repetitions of a study stimulus affects recognition and PI similarly suggests that some of the recognition advantage deriving from spacing may have increased perceptual fluency as its basis.

4.4.3. Time-pressured recognition memory

Some authors have proposed that the familiarity process operates quickly and automatically, whereas recollection takes more time (e.g., Atkinson and Juola, 1974). Converting these terms into our own, it can be hypothesized that the perceptual fluency ratio is immediately available but the weighting factors corresponding to the prior odds are not. If this hypothesis is correct, then those manipulations that affect only recollection and not perceptual fluency should be lessened when the recognition decision is made under time pressure.

In fact, several experiments support just such a conclusion. The effects of depth of processing, which has been discussed here and elsewhere (e.g., Blaxton, 1989; Jacoby, 1983) to affect only recollection, have been shown to be attenuated when recognition decisions are made quickly (Mulligan and Hirshman, 1995). It is important to note, however, that this effect arises not from a difference in the *rates* of information accrual between retrieval of shallowly and deeply processed items, but rather from the different asymptotes between the two cases. Also, the advantages of elaborative over rote rehearsal appear to be lessened when recognition is made under time pressure (Benjamin et al., 1995, 1996). The emphasis that the model places on separating the putative constituents of the recognition decision as seen in Eq. (7) is consistent with the distinction drawn between components that are quickly accessed and those that require more effortful retrieval.

5. Concluding comments

In this article, we have argued that inferential judgments serve memory in a fundamental manner. Such inferences act not only to influence our decisions about the future, but also decisions about the past. In particular, we propose that inferring the objective status about a past event involves utilizing a host of subjective cues present at the time of judgment, including subjective fluency.

In both cases – whether we are attempting to make an informed decision about the past or the future – the inferential task is a Bayesian one. Given only subjective cues (like fluency), the subject is required to predict the objective status of our memory or our experience. In the preceding analysis, we presented a minimally mathematical model of how standard signal detection analysis (e.g., Green and Swets, 1966) on the dimension of perceptual fluency acts to specify certain parameters in the Bayesian estimation process of the pastness of a probed item. This model is intended to explain some of the peculiarities of recognition memory, including frequency effects, dissociations of recognition and perceptual identification, and interactions between time-pressured and standard recognition memory. While we have only sketched the beginnings of the application of such a model to recognition memory, we hope to have shown the utility of the Bayesian framework in formulating a recognition decision. We further hope to have shown that one particularly important index for the Bayesian observer, in both recognition and metacognition, is subjective fluency.

References

Ashby, G.F., 1992. Multidimensional Models of Perception and Cognition. Erlbaum, Hillsdale, NJ.

Atkinson, R.C., Juola, J.F., 1974. Search and decision processes in recognition memory. In: Krantz, D.H., Atkinson, R.C., Luce, R.D., Suppes, P. (Eds.), Contemporary Developments in Mathematical Psychology (Volume 1): Learning, Memory and Thinking. Freeman, San Francisco, CA, pp. 243–293.

Bartlett, F.C., 1932. Remembering: A Study in Experimental and Social Psychology. Cambridge Univ. Press, Cambridge.

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

Begg, I., Duft, S., LaLonde, P., Melnick, R., Sanvito, J., 1989. Memory predictions are based on ease of processing. Journal of Memory and Language 28, 610–632.

Benjamin, A.S., Bjork, R.A., 1996. Retrieval fluency as a metacognitive index. In: Reder, L. (Ed.), Implicit Memory and Metacognition. Erlbaum, New Jersey, pp. 309–338.

Benjamin, A.S., Bjork, R.A., Schwartz, B.L., 1998. The mismeasure of memory: When retrieval fluency is misleading as a metacognitive index. Journal of Experimental Psychology: General 127, 1–14.

Benjamin, A.S., Bjork, R.A., Wickens, T.D., 1996. Speeded and unspeeded access to the products of primary and secondary rehearsal during recognition. Paper presented at the 37th Annual Meeting of the Psychonomic Society, Chicago, IL.

Benjamin, A.S., Jacobs, S.L., Bjork, R.A., 1995. The time course of access to the products of primary and secondary rehearsal during recognition. Paper presented at the 36th Annual Meeting of the Psychonomic Society, Los Angeles, CA.

Bjork, R.A., (in press). Assuming our own competence: Heuristics and illusions. In: Gopher, D., Koriat, A. (Eds.), Attention and performance XVII. Cognitive Regulation of Performance: Interaction of Theory and Application. Cambridge, MA, MIT press.

Blaxton, T.A., 1989. Investigating dissociations among memory measures: Support for a transfer-appropriate processing framework. Journal of Experimental Psychology: Learning, Memory, & Cognition 15, 657–668.
Brown, J., Lewis, V.J., Monk, A.F., 1977. Memorability, word frequency, and negative recognition. Quarterly Journal of Experimental Psychology 29, 461–473.
Coombs, C.H., Dawes, R.M., Tversky, A., 1970. Mathematical Psychology: An Elementary Introduction. Prentice-Hall, Englewood Cliffs, NJ.
Dorfman, J., Kihlstrom, J.F., Cork, R.C., Misiaszek, J.J., 1995. Priming and recognition in ECT-induced amnesia. Psychonomic Bulletin & Review 2, 244–248.
Egan, J.P., 1956. Recognition memory and the operating characteristic. Indiana University Hearing and Communication Laboratory, AFCRC-TN-58-51.
Gardiner, J.M., 1988. Functional aspects of recollective experience. Memory & Cognition 16, 309–313.
Gardiner, J.M., Craik, F.I.M., Bleasdale, F.A., 1973. Retrieval difficulty and subsequent recall. Memory & Cognition 1, 213–216.
Gillund, G., Schiffrin, R.M., 1984. A retrieval model for both recognition and recall. Psychological Review 91, 1–67.
Glanzer, M., Adams, J.K., Iverson, G.J., Kim, K., 1993. The regularities of recognition memory. Psychological Review 100, 546–567.
Glucksberg, S., McClosky, M., 1981. Decisions about ignorance: Knowing that you don't know. Journal of Experimental Psychology: Human Learning and Memory 7, 311–325.
Green, D.M., Swets, J.A., 1966. Signal Detection Theory and Psychophysics. Wiley, New York.
Gregg, V., 1976. Word frequency, recognition, and recall. In: J. Brown (Ed.), Recall and Recognition. Wiley, London.
Hart, J.T., 1965. Memory and the feeling-of-knowing experience. Journal of Educational Psychology 56, 208–216.
Hayman, C.G., Tulving, E., 1989. Contingent dissociation between recognition and fragment completion: The method of triangulation. Journal of Experimental Psychology: Learning, Memory, & Cognition 15, 228–240.
Hintzman, D.L., 1988. Judgments of frequency and recognition memory in a multiple-trace memory model. Psychological Review 95, 528–551.
Jacoby, L.L., 1983. Remembering the data: Analyzing interactive processes in reading. Journal of Verbal Learning and Verbal behavior 22, 485–508.
Jacoby, L.L., 1991. A process-dissociation framework: Separating automatic from intentional uses of memory. Journal of Memory and Language 30, 513–541.
Jacoby, L.L., Allan, L.G., Collins, J.C., Larwill, L.K., 1988. Memory influences subjective experience: Noise judgments. Journal of Experimental Psychology: Learning, Memory, & Cognition 14, 240–247.
Jacoby, L.L., Dallas, M., 1981. On the relationship between autobiographical memory and perceptual learning. Journal of Experimental Psychology: General 110, 306–340.
Jacoby, L.L., Kelley, C.M., 1987. Unconscious influences of memory for a prior event. Personality & Social Psychology Bulletin 13, 314–336.
Jacoby, L.L., Kelley, C.M., Dwyan, J., 1989a. Memory attributions. In: Roediger, H.L., III, 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: General 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: General 118, 115–125.
Johnson, M.K., Hashtroudi, S., Lindsay, D.S., 1993. Source monitoring. Psychological Bulletin 114, 3–28.
Johnston, W.A., Dark, V.J., Jacoby, L.L., 1985. Perceptual fluency and recognition judgments. Journal of Experimental Psychology: Learning, Memory, & Cognition 11, 3–11.
Johnston, W.A., Hawley, J.J., Elliot, J.M.G., 1991. Contribution of perceptual fluency to recognition judgments. Journal of Experimental Psychology: Learning, Memory, & Cognition 17, 210–223.

Kelley, C.M., Jacoby, L.L., 1998. Subjective reports and process dissociation: Fluency, knowing, and feeling. Acta Psychologica 98, 7-20, this issue.

Kelley, C.M., Lindsay, D.S., 1994. Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. Journal of Memory and Language 32, 1–24.

Kihlstrom, J.F., Shames, V.A., Dorfman, J., 1996. Intimations of memory and thought. In: Reder, L. (Ed.), Implicit Memory and Metacognition. Erlbaum, New Jersey, pp. 1–24.

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

Koriat, A., 1995. Dissociating knowing and the feeling of knowing: Further evidence for the accessibility model. Journal of Experimental Psychology: General 124, 311–333.

Kunst-Wilson, W.R., Zajonc, R.B., 1980. Affective discrimination of stimuli that are not recognized. Science 207, 557–558.

Lee, V.A., Narens, L., Nelsen, T.O., 1993. Subthreshold priming and the judgment of learning (submitted).

Luo, C.R., 1993. Enhanced feeling of recognition: Effects of identifying and manipulating test items on recognition memory. Journal of Experimental Psychology: Learning, Memory, & Cognition 19, 405–413.

Mandler, G., 1980. Recognition: The judgment of previous occurrence. Psychological Review 87, 252–271.

Mandler, G., Goodman, G.O., Wilkes-Gibbs, D.L., 1982. The word-frequency paradox in recognition. Memory & Cognition 10, 33–42.

Masson, M.E.J., MacLeod, C.M., 1992. Reenacting the route to interpretation: Enhanced perceptual identification without prior perception. Journal of Experimental Psychology: General 121, 145–176.

McDaniel, M.A., Waddill, P.J., Einstein, G.O., 1988. A contextual account of the generation effect: A three-factor theory. Journal of Memory and Language 27, 521–536.

Metcalfe, J., Schwartz, B.L., Joaquim, S.G., 1993. The cue familiarity heuristic in metacognition. Journal of Experimental Psychology: Learning, Memory, & Cognition 19, 851–861.

Morris, C.C., 1990. Retrieval processes underlying confidence in comprehension judgments. Journal of Experimental Psychology: Learning, Memory, & Cognition 16, 223–232.

Mulligan, N., Hirshman, E., 1995. Speed-accuracy trade-offs and the dual process model of recognition memory. Journal of Memory and Language 34, 1–18.

Murdock, B.B., 1993. TODAM2: A model for the storage and retrieval of item, associative, and serial-order information. Psychological Review 100, 183–203.

Nelson, T.O., Leonesio, R.J., Shimamura, A.P., Landwehr, R.S., Narens, L., 1982. Overlearning and the feeling of knowing. Journal of Experimental Psychology: Learning, Memory, & Cognition 8, 279–288.

Reder, L.M., 1987. Selection strategies in question answering. Cognitive Psychology 19, 90–138.

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

Schwartz, B.L., 1994. Sources of information in metamemory: Judgments of learning and feeling of knowing. Psychonomic Bulletin & Review 1, 357–375.

Schwartz, B.L., Benjamin, A.S., Bjork, R.A. (in press). The inferential and experiential bases of metamemory. In: Current Directions in Psychological Sciences.

Schwartz, B.L., Metcalfe, J., 1992. Cue familiarity but not target retrievability enhances feeling of knowing judgments. Journal of Experimental Psychology: Learning, Memory, & Cognition 18, 1074–1083.

Shaw, J., 1996. Increases in eyewitness confidence resulting from postevent questioning. Journal of Experimental Psychology: Applied 2, 126–146.

Snodgrass, J.G., Hirshman, E., Fan, J., 1996. The sensory match effect in recognition memory: Perceptual fluency or episodic trace? Memory & Cognition 24, 367–383.

Strack, F., Bless, H., 1994. Memory for nonoccurences: Metacognitive and presuppositional strategies. Journal of Memory and Language 33, 203–217.

Tulving, E., 1985. Memory and consciousness. Canadian Psychologist 26, 1–12.

Whittlesea, B.W.A., 1993. Illusions of familiarity. Journal of Experimental Psychology: Learning, Memory, & Cognition 19, 1235–1253.

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 judgments in a perceptual identification task. Memory & Cognition 13, 101–111.