Journal of Experimental Psychology:
Learning, Memory, and Cognition
1995, Vol. 21, No. 3, 777–784


0278-7393/95/$3.00

# How Should Implicit Memory Phenomena Be Modeled?

Gail McKoon and Roger Ratcliff
Northwestern University

In a reply to R. Ratcliff and G. McKoon's (1995a) article on bias in the object decision task, D. L. Schacter and L. A. Cooper (1995) critiqued their theoretical arguments and presented an updated view of priming in the object decision task. In the present article, the updated view is examined in detail, and it is questioned whether Schacter and Cooper's explanation of the data is sufficiently articulated to be falsifiable. It is also argued in the present article that evidence from other research domains is not directly supportive of the memory systems hypothesis and that the statistical power available in data from object decision experiments is not great enough to test some relevant hypotheses. Finally, the bias hypothesis (Ratcliff & McKoon, 1995a) is elaborated to show that it requires a particular pattern of experimental results and that it serves as a target phenomenon for modeling.

Implicit memory phenomena have captured the interest of many cognitive psychologists. The phenomena are frequently said to provide insights into memory processes that occur without awareness—"unconscious" memory processes. They have given hope of significant steps forward in the understanding of how memory works.

We believe that progress in this domain can best occur through the articulation of specific, falsifiable models. In a series of experiments (Ratcliff & McKoon, 1995a), we examined object decision, an implicit task used frequently by Schacter and Cooper and their colleagues (Cooper, Schacter, Ballasteros, & Moore, 1992; Schacter & Cooper, 1993; Schacter, Cooper, & Delaney, 1990; Schacter, Cooper, Delaney, Peterson, & Tharan, 1991; Schacter, Cooper, Tharan, & Rubens, 1991). The results of our experiments did not fit the predictions of Schacter and Cooper's theory about how implicit memory operates in object decision. Schacter and Cooper (1995) have presented a more fully developed theory, and in this article we question whether their extended theory is falsifiable.

The main phenomenon toward which research on object decisions has been directed is repetition priming. In typical experiments (Schacter et al., 1990; Schacter, Cooper, Delaney, et al., 1991), participants are shown three-dimensional line drawings of objects that are either "possible," in that they could actually exist in the real world, or "impossible." In the study phase of the experiment, the participants perform a task designed to require that the drawings be encoded as objects; often, they are asked to decide whether each object is right or left facing. In the test phase of the experiment, each of a series of objects is flashed briefly and the participants decide whether it is possible or impossible. For a possible object, responses are more likely to be correct if the object was previously presented in the study phase. For an impossible object, Schacter and Cooper and their colleagues found that previous study had no significant effect on performance. However, Ratcliff and McKoon (1995a) found conditions under which previous study did affect decisions about impossible objects: Imposing a deadline on the amount of time allowed for a response and requiring participants to maintain digits in memory while they were responding both led to a *decrease* in accuracy for impossible objects as a function of prior study. Overall, performance exhibited a bias to respond "possible" to objects previously studied, a bias that led to increased probability of correct responses for possible objects and increased probability of incorrect responses for impossible objects.

Although the model with which Schacter and Cooper proposed to explain these patterns of results was described in part in earlier articles, their discussion of Ratcliff and McKoon's (1995a) experiments (Schacter & Cooper, 1995) provides a more complete presentation of their view. Their explanation of the complete set of relevant data appears to have been offered somewhat tentatively, but it was used in a critique of the alternative view we presented, and it merits serious consideration. In their view, information in memory is divided into separate systems. One main division is between presemantic perceptual representations and episodic information. The perceptual representations comprise several kinds of information, for example, auditory and visual representations of words (Schacter, 1994; Schacter & Tulving, 1994), in addition to the perceptual representations of objects that are relevant to the object decision task. In splitting episodic information in memory from perceptual representations (from "implicit information"), Schacter and Cooper's account of object decision data followed a frequently proposed account of dissociations between the effects of variables that tap episodic information and the effects of variables that tap implicit information. What distinguished Schacter and Cooper's account from others was their proposal about the perceptual representations that operate in the object decision task, to which we now turn.

Gail McKoon and Roger Ratcliff, Department of Psychology, Northwestern University.

This research was supported by National Institute of Mental Health Grants HD MH44640 and MH00871, by National Institute on Deafness and Other Communication Disorders Grant R01-DC01240, and by National Science Foundation Grant SBR-9221940.

We would like to thank Dan Schacter for his extensive comments on an earlier draft (though he might not like what they produced).

Correspondence concerning the article should be addressed to Gail McKoon, Department of Psychology, Northwestern University, Evanston, Illinois 60208.

## Structural Descriptions

According to Schacter and Cooper's (1995) account of data in the object decision task, encoded information about previously studied objects includes representations of parts of objects and representations of whole objects. Perceptual representations of the parts of both possible and impossible objects are stored, but only possible objects are stored as whole objects because only for them can three-dimensional "structural descriptions" be constructed. In terms of the object decision task, which was the task implemented in Schacter and Cooper's experiments, they hypothesized that only structural descriptions of previously encoded whole possible objects affect performance. Some aspect of the representation or processing of these structural descriptions prevents stored perceptual information about the parts of previously studied objects from affecting performance. A priori, the stored parts of objects might be expected to affect object decisions in either of two ways: Having the parts of an object already encoded from previous study might help later perception by giving a head start to processing when the object is flashed for the object decision. Alternatively, Schacter and Cooper (1995, p. 769) suggested that having the parts already in memory might lead to a tendency to respond "possible." However, according to Schacter and Cooper, both of these possibilities are ruled out by the finding that there is no effect of prior study on decisions about impossible objects. The stored parts of impossible objects neither increased nor decreased the probability of a correct object decision, and so it must be that the only previously encoded information that can affect object decision performance is structural descriptions of whole possible objects.

Schacter and Tulving (1994, p. 24) have said that perceptual representations play "an important role in making possible the identification of words and objects" and that they operate at a presemantic level. From such descriptions, we inferred that retrieval of structural descriptions of previously encoded objects was fast—fast enough to be intimately involved in the perception of objects in the real world (lions and tigers and drive-by shooters) and fast enough to be available in the first few hundred milliseconds of perceptual processing (Ratcliff and McKoon, 1995a, p. 759). However, Schacter and Cooper have now made it clear that, in the object decision task, a substantial amount of time (about 1,000 ms) is required before previously encoded structural descriptions can affect performance (Schacter & Cooper, 1995, p. 771).

The relatively large amount of time required for stored structural descriptions to affect performance in standard object decision experiments allowed Schacter and Cooper (1995) to explain Ratcliff and McKoon's (1995a) findings when the task was changed by imposing a deadline on response times. With a deadline, prior study affected performance on impossible objects as well as on possible objects. Schacter and Cooper (1995, p. 772) suggested that a deadline interferes with access to object representations from the structural description system. This allows other kinds of information, specifically, representations of the parts of previously studied objects and the episodic familiarity of the objects, to affect performance. Episodic information can represent both parts of objects and whole objects, for both possible and impossible objects. Both episodic information and information about parts of previously studied objects bias decisions toward a possible response, consistent with Ratcliff and McKoon's data.

The effects of several other variables are explained in the same way as the effect of a response time deadline: The elimination of structural descriptions of previously encoded whole possible objects from processing allows episodic information and encoded parts of previously studied objects to affect performance. Ratcliff and McKoon (1995a) showed that imposing a memory load during object decisions had the same effect as imposing a response time deadline. From Schacter and Cooper's (Schacter, Cooper, et al., 1990; Schacter, Cooper, Delaney, et al., 1991; Schacter, Cooper, Tharan, et al., 1991) earlier descriptions of their view, we had thought that a memory load should not impair the use of structural descriptions because their retrieval would be part of perceptual processing. However, in their reply Schacter and Cooper suggested that memory load "interferes with access to object representations from the structural description system" (1995, p. 772). Ratcliff and McKoon (1995a) also manipulated similarity by using test objects that were highly similar to studied objects, and that too is said to eliminate the use of structural descriptions (see also Williams, Crowder, & Tarr, 1994). Finally, Schacter and Cooper suggested that the use of structural descriptions is eliminated when the decision about a test object requires choosing whether it is the same as a possible object or a very similar impossible object (forced choice; Ratcliff & McKoon, 1995a).

Figure 1 shows two schemes that lay out Schacter and Cooper's (1995) proposals, as we understand them. Schacter and Cooper and their colleagues have never specified the processes involved in making object decisions (at least, not to our knowledge); they have been concerned only with the effects of previous encodings on these processes. However, we assume that making an object decision depends on the processes of building a three-dimensional structural description of the line drawing presented as the stimulus. Success in building a three-dimensional representation indicates a "*possible*" decision, and failure, an "*impossible*" decision. If an attempt was not made to construct a three-dimensional representation of the test stimulus, then there would be no basis for a decision about a drawing that had never been seen before. Therefore, both schemes in Figure 1 show a structural description construction process.

The key to both of the models in Figure 1 is that different sources of information about previously studied objects affect object decisions under different experimental conditions. Both models include the three sources of information about previously studied objects assumed by Schacter and Cooper (1995): perceptual representations of the parts of possible and impossible objects, perceptual representations of whole possible objects, and episodic information about both parts of objects and whole objects. In the standard task, without a response deadline or memory load, performance is affected by the perceptual representations of whole possible objects and not by episodic information nor by perceptual representations of parts of objects. With the task manipulations, the situation is exactly the opposite: Performance is not affected by perceptual representations of whole objects, but it is affected by episodic information and perceptual representations of parts of objects.

The difference between Models 1 and 2 lies in where the switch between the sources of information is located. In Model 1, it is located within the perceptual processing system that constructs three-dimensional representations. The construction process is either affected by memory for previously encountered whole possible objects or it is affected by memory for previously encountered parts and episodic information. To account for the data, the construction process cannot be affected by all three sources of information because then we would see effects of all three in all versions of the object decision task, and that does not happen. Instead, it must be that something about the representation or use of one of the sets of information blocks or inhibits the effects of the other. We are unclear about how this blocking could come about, that is, about why the use of one of the sets of information should inhibit the use of the other. We are also unclear about how switching occurs. Perceptual processes seem unlikely to be under the control of attentional processes, yet attention (or some other unspecified process under the control of task demands) would seem to be needed for switching from the use of one set of information to another, as proposed by Schacter and Cooper to explain the data.

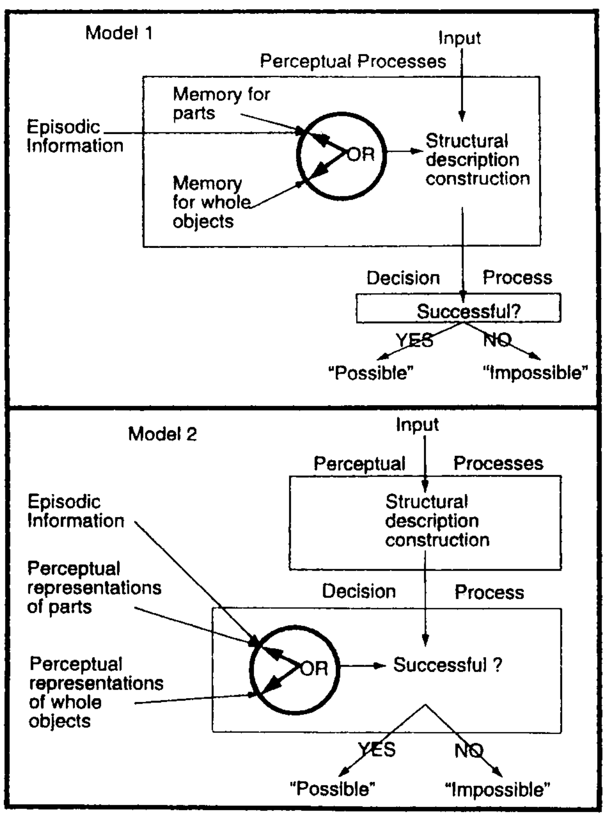


*Figure 1.* This figure shows the sources of information about previously studied objects that Schacter and Cooper (1995) assume affect performance on object decisions and how different task demands alter those effects. Perceptual processes attempt to construct a three-dimensional structural description of a line drawing, and their success or failure determines whether the decision is "possible" or "impossible."

In Model 2, the switch between different sources of information is located in the decision process. None of the sources of memory for previously studied objects affects the structural description construction processes (at least, not in any way that affects performance in the object decision task). Rather, they affect the decision process, with that process switching to be biased either by memory for whole possible objects or by memory for parts and episodic memory. Putting the switch in decision processing, rather than putting it in perceptual processing, seems more in accord with the fact that task demands affect performance. However, removing the influences of memory from perceptual processes seems an outright contradiction of the stated role of perceptual representations in "making possible the identification of words and objects" (Schacter & Tulving, 1994, p. 24).

Both of these schemes raise more questions than they answer. For example, exactly what mechanisms determine when one set of information affects priming and not the other? These mechanisms must be specified in order to understand processing in the object decision task, how priming affects that processing, and why some study-to-test changes in perceptual attributes of stimuli do not affect priming (Cooper et al., 1992).

## Problems, Questions, and Falsifiability

Regardless of which model in Figure 1 best describes Schacter and Cooper's (1995) proposals, a central question is immediately raised, and that is whether the structural description system is truly a perceptual system. It has been described as perceptual and as presemantic, and presumably it is preepisodic. It is part of the collection of implicit memory systems that "plays an important role in making possible the identification of words and objects, . . . operates at a presemantic level, and is typically involved in nonconscious or implicit expressions of memory, such as priming" (Schacter & Tulving, 1994, p. 24). Yet, as summarized above, previously encoded structural descriptions of whole possible objects affect object decisions only with a relatively substantial amount of processing time (about 1,000 ms), their effect on performance is subject to interference, and the effect is possibly not obligatory (as perceptual processes might be thought to be) because it only occurs with some, but not all, decisions about objects. With neither of the models in Figure 1 is it clear how these two sets of attributes are reconciled with each other.

There is already some evidence to suggest that they cannot be reconciled. Possible objects are said to be recognized better than impossible objects because they can be encoded three-dimensionally, whereas impossible objects cannot (Schacter & Cooper, 1995). A response time deadline should, by Schacter and Cooper's proposals, eliminate the advantage for possible objects because it eliminates the use of structural descriptions of the test objects.[1] If there is no structural description of a test

[1] We made this assumption because of Schacter and Cooper's (1995, p. 772) statement that deadlines "interfere with access to object representations from the structural description system." However, Schacter and Cooper have also said that their position "need not imply

object, then there is nothing to match against a previously encoded structural description, and the advantage given by the previous structural description should be eliminated. However, a 200-ms deadline does not eliminate the advantage (Schacter & Cooper, 1995). This result contradicts a prediction of their extended model. We expect that a second prediction would also be disconfirmed. Imposing a deadline on response times is supposed to allow episodic information to influence object decisions. This means that, under deadline, episodic variables should affect object decisions.

Do disconfirmations of predictions like these constitute falsifiability? Or can Schacter and Cooper's (1995) model change to accommodate problems like these? The central feature of Schacter and Cooper's view as it currently stands is the representation of structural descriptions of whole possible objects. Originally, these representations were put into the model to explain why prior exposure did not affect decisions about impossible objects, and now there are circumstances in which they are taken out of processing to explain why decisions about impossible objects sometimes are affected by prior exposure.

## Converging Evidence

Converging evidence is an important tool in psychology, especially in hypothesis and theory generation, but the question of when different pieces of research truly converge to support a single theory must be carefully considered. Schacter and Cooper defended their model by arguing that the notion of a structural description system is validated by work in computational vision and neuropsychology (Schacter & Cooper, 1995, p. 768). We believe that, currently, findings in these areas play a useful role in generating hypotheses, but they do not directly support each other; they are independent studies that offer only the potential of future connections.

Winston's (1977, chapter 3) presentation of research on object identification, for example, discusses computational algorithms for identifying objects from line drawings and for determining if a line drawing can represent a possible three-dimensional object. To provide direct converging evidence between proposed algorithms and human object identification, the algorithms would have to be taken seriously as a guide to experiments. Easy versus difficult objects for an algorithm would be compared to easy versus difficult objects for people. Objects mistakenly interpreted by an algorithm as being possible in the real world, real objects that an algorithm cannot interpret, and objects for which an algorithm provides ambiguous representations would all become the stimuli for psychological experiments. Without data from experiments like these, the relation between work in computational vision and Schacter and Cooper's structural description system is only a loose analogy.

Moreover, some computational vision research is inconsistent with Schacter and Cooper's (1995) model. Proposals that reflect a more biological perspective (Edelman & Weinshall, 1991; Poggio & Girosi, 1990) represent a three-dimensional object not as anything like Schacter and Cooper's three-dimensional structural description but instead as multiple two-dimensional representations. The question is whether there is any basis on which to choose between computational vision research that is consistent with Schacter and Cooper's view and computational vision research that is not consistent.

The situation is the same with research on neuropsychological deficits in object recognition. Riddoch and Humphreys (1987; see also Warrington, 1982), for example, studied a single participant with optic aphasia. The participant performed normally on semantic tasks, but was impaired in using pictured (well-known) objects to access semantic information about those objects. Riddoch and Humphreys hypothesized the existence of a link between a system that represents objects and the semantic system, and impairment of this link was suggested as the reason for the participant's poor performance with the objects. However, this hypothesis was aimed at accounting for 1 participant's (or class of participants') data, and it was geared toward representations of objects that were already well-known, not new three-dimensional line drawings with no semantic representation. To become converging evidence for Schacter and Cooper's (1995) structural description system, the domains of research would need to be tightly related, perhaps by conducting object decision experiments with patients with optic aphasia or by making predictions from Schacter and Cooper's model for performance by these patients on other tasks.

Another category of possibly converging evidence for a structural description system is the functional and stochastic independence that has been observed between performance on implicit tasks and performance on explicit tasks. Schacter and Cooper said that "no single kind of data provides conclusive support for dissociable memory systems" (1995, p. 773), but they have consistently cited this category of evidence in their discussions of implicit perceptual systems (e.g., Schacter, Cooper, et al., 1990; Schacter, Cooper, Delaney, et al., 1991; Schacter, Cooper, Tharan, et al., 1991). In Ratcliff and McKoon (1995a), we summarized research (e.g., Hintzman & Hartry, 1990; Ostergaard, 1992; see also McKoon, Ratcliff, & Dell, 1986) showing how functional independence and stochastic independence are basically flawed as sources of evidence

---

that no structural information can be extracted from an object when a deadline is used" (1995, p. 774). If it is the latter they intend, that is, that object representations *are* accessible under deadline conditions, then several problems are raised. The first problem is why, if structural information is extracted under deadline conditions, it does not affect object decisions. Schacter and Cooper said that deadlines "interfere with the slow decision process that is a feature of subjects' performance on the object decision task" (1995, p. 774), and lead subjects to rely on "fast-acting familiarity" (1995, p. 774). However, the 200 ms deadline speeds responses by about 200 or 300 ms (Ratcliff & McKoon, 1995a, Experiments 1–3), an increase in speed of about 20–30%, but the deadline changes accuracy by only about 2%. How such a large amount of decision time could be eliminated with such a small cost to accuracy is a problem for Schacter and Cooper. Familiarity alone would not allow discrimination between possible and impossible nonstudied objects. Another problem is that if both object decision and recognition depend on familiarity at fast deadlines (or under memory loads), then under those conditions they should not dissociate. Finally, a third problem is that to match the data, structural information would have to affect fast deadline recognition decisions while not affecting fast deadline object decisions; this seems odd and it would mean that recognition was at least as informative a task as object decision for investigating the use of information from the structural description system.

about separate memory systems. Here we reiterate that questionable sources of evidence plus loose analogies to computational algorithms and neuropsychological data do not converge into strong evidence. Computational and neuropsychological proposals should be taken seriously; research should be undertaken to relate them explicitly to the structural description hypothesis.

## An Alternative View: Bias

We claimed that there is no need to postulate implicit memory systems to explain performance in the object decision task (Ratcliff & McKoon, 1995a). We suggested instead that object decisions could be understood in terms of information processing models that do not distinguish between different memory systems, as in Hummel and Biederman's (1992) or Edelman and Weinshall's (1991) connectionist models. In Hummel and Biederman's model, for example, levels of representation correspond to different kinds of information, with lower levels processing perceptual features, higher levels processing parts of objects, and the highest levels encoding whole objects and semantic–associative links to other information in memory. Dissociations among variables occur because different variables affect processing at different levels of the system.

In an information processing framework, the patterns of data obtained with object decision have straightforward interpretations. Previous exposure to an object biases the system toward a positive, "possible" response to that object. This bias affects possible and impossible objects alike, and in the Hummel and Biederman (1992) or the Edelman and Weinshall (1991) models, it might come about in any number of ways, perhaps because of a reduction in some response criterion in the system or because of modifications to the weights in the neural network implementing the model.

Schacter and Cooper (1995, p. 773) criticized the "fuzzy notion of 'bias' " as being no more specific than the implicit memory systems view.[2] It is true that none of the information processing or neural network models were designed to predict performance on object decisions. However, the notion of bias is not fuzzy: First, it provides a specific target behavior that information processing, connectionist, or neural models of perception should produce,[3] and second, it imposes the specific prediction that the benefits of prior exposure will be offset by costs of about the same size as the benefits (as has been observed in a range of tasks by Ratcliff, Allbritton, & McKoon, 1994; Ratcliff & McKoon, 1995a, 1995b; Ratcliff, McKoon, & Verwoerd, 1989).

We hypothesized that the reason bias is observed only for previously studied possible objects in the typical object decision task and not for previously studied impossible objects is that for impossible objects, bias is overridden by explicit information about some feature or combination of features that cues that the object is impossible. This hypothesis leads to strong predictions: We know from considerable previous research that retrieval of explicit episodic information should be eliminated (or greatly reduced) when a time deadline is imposed on responses and when a concurrent memory load is imposed. Our tests of these predictions confirmed our hypothesis. With response time deadlines, a bias to respond "possible" to previously studied objects was found for impossible as well as for possible objects. The result was the same when participants had to remember seven digits while doing object decisions: a bias for both possible and impossible objects.

## The Data

For both the bias view and implicit memory theories, an obstacle that bedevils implicit memory research is the statistical power needed to determine whether there are real differences in performance among experimental conditions (see Ostergaard, 1992, for a related discussion). Across the experiments by Ratcliff and McKoon (1995a), Schacter, Cooper, Delaney, et al. (1991), and Schacter and Cooper (1993), the standard errors of the mean proportion correct values for object decisions typically fell in a range from 0.015 to 0.04. For a difference between two means to be significant (compared to the standard error in the means), the difference would have to be at least 0.042 (for standard error of 0.015) and perhaps as large as 0.113 (for standard error of 0.04). However, the total amount of facilitation due to prior study for possible objects averaged only about 0.09 (in Ratcliff & McKoon, 1995a) to 0.12 (in Cooper et al., 1992; Schacter et al., 1990; Schacter, Cooper, Delaney, et al., 1991; Schacter & Cooper, 1993). With standard errors so large relative to the amounts of facilitation, differences in the amount of facilitation across experimental conditions are going to be difficult to find, at best. For example, even a finding of twice the amount of facilitation in one condition versus another (e.g., 0.06 vs. 0.12) will have difficulty reaching statistical significance.

One place this problem becomes critical is in examining object decision performance by patients with amnesia, as illustrated by an experiment conducted by Schacter, Cooper, and Treadwell (1993). They attempted to compare performance for participants with amnesia versus that for matched controls. For the control group prior study facilitated decisions on possible objects by 0.075, with a standard error of about 0.035. It might be hypothesized that participants with amnesia would also show facilitation but less so than the controls. However, with such a large standard error, about half the size of the effect, it would be difficult to show both significant facilitation for the participants with amnesia and a significant decrement in this facilitation relative to the controls. In fact, in Schacter et al.'s experiment, even the difference in the amount of facilitation for possible versus impossible objects, a differ-

[2] Schacter and Cooper (1995) criticized our use of the language of implicit memory research. We believe we used the language in a way that provides reasonably unambiguous communication (and our usage is not uncommon). We deliberately did not conform to the strict definitions that have been given by some implicit memory theorists because those definitions can force a particular theoretical view on empirical phenomena. For example, some definitions might imply that if a phenomenon is not implicit, it must be explicit and there is no other possibility.

[3] It is likely that the mechanisms that produce bias will vary across different tasks and different response variables (e.g., accuracy or reaction time) and different response measures (e.g., forced choice or yes–no decisions or naming).

ence critical to their model. was apparently not significant. It appears that similar problems arose in another experiment with patients with amnesia (Schacter, Cooper, Tharan, et al., 1991). Some of the data in this experiment were consistent with Schacter, Cooper, and Treadwell's predictions, but some were not. For example, although participants with amnesia performed worse than college students on tests of recognition of possible objects, they performed about as well as matched controls. And, while the participants with amnesia showed benefits of prior study for possible but not impossible objects as Schacter, Cooper, Tharan, et al. would predict, the matched controls showed about the same benefits for possible and impossible objects. Thus, although Schacter and Cooper (1995) argued that some aspects of their data for participants with amnesia are problematical for our bias hypothesis, other aspects of their data are problematic for their view. A likely source of the inconsistencies in the data is insufficient statistical power.

## What Does Bias Fail to Explain?

Schacter and Cooper (1995) listed a number of reasons why they believed performance on the object decision task could not be explained as bias. Some of these implicated the issue of empirical power just discussed and some can be addressed by a deeper consideration of task requirements.

One of Schacter and Cooper's (1995) criticisms pertains to our hypothesis that episodic information overrides bias to produce the typical finding of no effect of prior study on impossible objects. They suggest that if episodic information is used in the object decision task, then object decision performance should be affected by variables that affect performance in episodic tasks (see Ostergaard, 1992). This does not necessarily follow, as we made clear in our original article, because the kind of episodic information that is useful in one task may be different than the kind that is useful in other tasks—correlations between performance even on different episodic tasks are not always high (Ratcliff & McKoon, 1995a). Even what might seem a straightforward manipulation like increasing the amount of learning (e.g., increasing the number of times an object is presented for a left vs. right facing judgment) may not have easily predictable consequences. Increasing the amount of learning could increase the amount of bias, it could increase the amount of encoded episodic information, or it could do both, to the same or different degrees.

However, suppose that increasing the amount of learning affected episodic information and only episodic information, and consider the processing that could make use of this information, as shown in Figure 2. When a test object is flashed, under the conditions of typical experiments, there are three possible outcomes of perceptual processing: On some trials, the processing system provides enough information to unambiguously determine whether the object is possible or impossible. On other trials, the processing system provides too little information to provide any pointer to a judgment about the object or to explicit memory about it. However, on some trials partial information is extracted from the flashed object, and this piece of information enables retrieval of relevant information from episodic memory to occur. It is only on this proportion of trials that episodic information can affect the object decision. This proportion of trials is probably small (10–20%), so that no matter how much episodic information is encoded, it cannot improve performance more than this small amount. Moreover, it may be that very little episodic information is needed to cue possibility versus impossibility, and the increase in episodic information that comes about with more learning is not relevant once there is some minimal amount sufficient to cue the possible–impossible decision. Thus, increasing the amount of learning could increase the amount of encoded episodic information without changing performance on object decisions that use episodic information. The important point, as this example makes clear, is that the effect of any variable on performance can be predicted only in the context of detailed processing assumptions.

In a second criticism of our hypothesis that episodic information is used in object decisions, Schacter and Cooper (1995, p. 770) suggested that the hypothesis might apply only to impossible objects. That is not what we intended. They made this suggestion because they believed, incorrectly, that the bias view could not otherwise accommodate the data from our experiments. They reasoned as follows: If episodic information cues whether an object is possible or impossible, and it does so for both possible and impossible objects, then imposition of a response time deadline or a memory load should affect performance on possible objects as well as on impossible objects; in particular, performance on possible objects should suffer because of the absence of episodic information. However, Schacter and Cooper claimed, performance on possible objects was not affected by the deadline or memory load manipulations.

What Schacter and Cooper (1995) did not note about our results was that, under conditions that would be expected to eliminate episodic retrieval (our Experiments 2, 3, 4, and 5), there was a greater bias effect when the participants were given a very short deadline (200 ms) than when they were given a longer deadline (800 ms) or no deadline (in the memory load condition). At the 200-ms deadline, the size of the bias effect

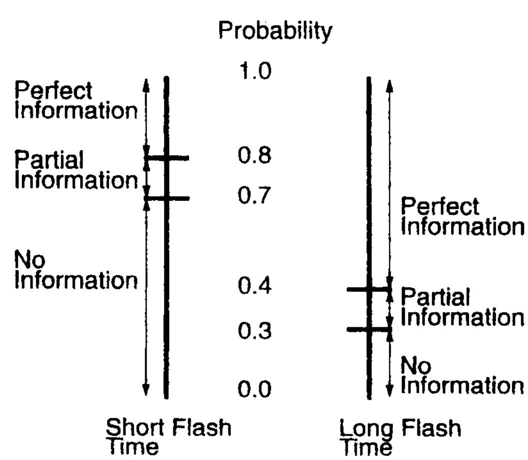


*Figure 2.* Possible states of information for a task in which perceptual information is limited (e.g., when a word or a three-dimensional object is flashed briefly).

averaged 0.13. Without such a short deadline, the bias effect averaged 0.06. With this change in the amount of influence from bias across conditions, direct comparisons of performance on possible objects across the conditions are unwarranted. To give an example, the proportion of possible responses to a possible object was 0.60 without a deadline and 0.56 with a deadline in our Experiment 4. These numbers are not directly comparable because the latter reflects a larger amount of bias than the former. The change in bias across conditions should not be surprising; the idea that participants can shift their reliance away from one kind of processing and toward another as a function of experimental conditions is a familiar one.

If we take the average amounts of bias seriously, we can add back in the effect of episodic information. Under conditions that were thought to eliminate episodic retrieval while otherwise being close to the standard task (i.e., conditions that did not use a short deadline), the size of the bias effect averaged 0.06, as mentioned above; that is, the likelihood of responding possible to a previously seen object was 0.06 greater than the likelihood of responding possible to a new object. For this .06 bias effect to be canceled by episodic information for impossible objects, the probability of episodic information leading to an "impossible" response would also have to be about 0.06. Then, for possible objects, we assume that the contribution of episodic memory would be the same (contrary to what Schacter and Cooper, 1995, suggested). Because both episodic information and bias increase the probability of a "possible" response for possible objects, the bias effect of 0.06 would be added to the episodic effect of 0.06, giving an advantage to previously studied possible objects of 0.12. This is about the same size as the effect usually obtained.

The point is that the data show a shift toward greater reliance on bias at a shorter deadline. Once this shift is taken into account, there is nothing in our data to contradict our assumption that retrieval of episodic information contributes about equally to performance on possible and impossible objects, and in fact, our assumption gives a good account of the complete data.

## What Do the Experiments in Ratcliff and McKoon (1995a) Refute?

Schacter and Cooper (1995, p. 773) expressed some confusion about what we thought was refuted by the results of the experiments in Ratcliff and McKoon (1995a). Most specific, the data refuted the notion that object decisions about impossible objects cannot be affected by prior study. Most important, the data refuted the notion that implicit memory systems provide useful explanations of performance on the object decision task.

The alternative we proposed, that prior exposure biases later processing, provides one strong and testable prediction: The benefit of prior exposure to possible objects should be about the same size as the cost of prior exposure to impossible objects. This prediction was confirmed by the data in our experiments, and support is also provided by bias patterns of data for a number of other implicit tasks (Ratcliff & McKoon, 1995b; Ratcliff et al., 1989). In contrast, it is difficult to think of any strong prediction that implicit memory theorists have made about object decisions for which disconfirmation would result in the theory being rejected or radically modified. Even the most basic of the ideas, the split between the episodic system and implicit systems that is central to their explanation of dissociations between explicit and implicit tasks, now seems at risk because episodic information might, under response time deadline, determine performance on object decisions.

The agenda for research that is mapped out by information processing theories is to begin to take apart the mechanics of how object decisions are made, to find subprocesses that might be biased, and to formulate qualitative and quantitative tests of different possible mechanisms. The agenda for research set by memory system theories should be to describe how information is represented in the different subsystems and how that information aids perceptual processes. These are compatible agendas, so long as both provide specific and testable predictions. Schacter and Cooper (1995, p. 773) listed a number of different views of implicit memory, but for the most part none of these were designed to make specific quantitative predictions about performance in the tasks toward which they were addressed. Schacter and Cooper pointed to a model proposed by McClelland, McNaughton, and O'Reilly (1994) as an example of a theory using multiple memory systems. McClelland et al.'s model is designed to explain retrograde amnesia and consolidation in learning and proposes a rapid hippocampal system operating as a teacher for a slower neocortical system. These two systems carry out *similar* computations in an interactive way; they do not represent the kind of independent systems each directed toward *different* computations that are the hallmark of the implicit memory systems approach (e.g., Schacter & Tulving, 1994).

What is needed for development of the implicit systems approach are detailed proposals about how implicit tasks are performed. Currently, the enterprise revolves around repetition priming—the facilitation of some decision about a test item by prior exposure to the item. However, understanding how that facilitation comes about must depend on understanding how the task is performed. That, in turn, can lead to an account of priming within an account of performance as a whole, and to critical tests among models.

## References

Cooper, L. A., Schacter, D. L., Ballesteros, S., & Moore, C. (1992). Priming and recognition of transformed three-dimensional objects: Effects of size and reflection. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 43–57.

Edelman, S., & Weinshall, D. (1991). A self-organizing multiple-view representation of 3D objects. *Biological Cybernetics, 64,* 209–219.

Hintzman, D. L., & Hartry, A. L. (1990). Item effects in recognition and fragment completion: Contingency relations vary for different subsets of words. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 955–969.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99,* 480–517.

McClelland, J. L., McNaughton, B., & O'Reilly, R. C. (1994). *Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist* *models of learning and memory* (Tech. Rep. No. PDP.CNS.94.1). Pittsburgh, PA: Carnegie-Mellon University, Department of Psychology.

McKoon, G., Ratcliff, R., & Dell, G. S. (1986). A critical evaluation of the semantic–episodic distinction. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 295–306.

Ostergaard, A. L. (1992). A method for judging measures of stochastic dependence: Further comments on the current controversy. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 413–420.

Poggio, T., & Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science, 247,* 978–982.

Ratcliff, R., Allbritton, D. W., & McKoon, G. (1994). *Bias in auditory priming.* Manuscript submitted for publication.

Ratcliff, R., & McKoon, G. (1995a). Bias in the priming of object decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 754–767.

Ratcliff, R., & McKoon, G. (1995b). *Biases in implicit memory tasks.* Manuscript submitted for publication.

Ratcliff, R., McKoon, G., & Verwoerd, M. (1989). A bias interpretation of facilitation in perceptual identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 378–387.

Riddoch, M. J., & Humphreys, G. W. (1987). Visual object processing in optic aphasia. A case of semantic access agnosia. *Cognitive Neuropsychology, 4,* 131–186.

Schacter, D. L. (1994). Priming and multiple memory systems: Perceptual mechanisms of implicit memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 233–268). Cambridge, MA: MIT Press.

Schacter, D. L., & Cooper, L. A. (1993). Implicit and explicit memory for novel visual objects: Structure and function. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 995–1009.

Schacter, D. L., & Cooper, L. A. (1995). Bias in the priming of object decisions: Logic, assumption, and data. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 768–776.

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

Schacter, D. L., Cooper, L. A., Delaney, S. M., Peterson, M. A., & Tharan, M. (1991). Implicit memory for possible and impossible objects: Constraints on the construction of structural descriptions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 3–19.

Schacter, D. L., Cooper, L. A., Tharan, M., & Rubens, A. B. (1991). Preserved priming of novel objects in patients with memory disorders. *Journal of Cognitive Neuroscience, 3,* 118–131.

Schacter, D. L., Cooper, L. A., & Treadwell, J. (1993). Preserved priming of novel objects across size transformation in amnesic patients. *Psychological Science, 4,* 331–335.

Schacter, D. L., & Tulving, E. (1994). What are the memory systems of 1994? In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 1–38). Cambridge, MA: MIT Press.

Warrington, E. K. (1982). Neuropsychological studies of object recognition. *Philosophical Transactions of the Royal Society of London, B 298,* 15–33.

Williams, P., Crowder, R. G., & Tarr, M. J. (1994, November). *The basis of the "bias" effect in object decision priming.* Paper presented at the 35th annual meeting of the Psychonomic Society, St. Louis, MO.

Winston, P. H. (1977). *Artificial intelligence.* Reading, MA: Addison-Wesley.

Received December 5, 1994
Revision received January 23, 1995
Accepted February 2, 1995 ■