Social Cognition, Vol. 27, No. 2, 2009, pp. 283–319

EVOLUTION AND EPISODIC MEMORY: AN ANALYSIS AND DEMONSTRATION OF A SOCIAL FUNCTION OF EPISODIC RECOLLECTION

Stanley B. Klein, Leda Cosmides, Cynthia E. Gangi, Betsy Jackson, and John Tooby
University of California, Santa Barbara

Kristi A. Costabile
Ohio State University

Over the past two decades, an abundance of evidence has shown that individuals typically rely on semantic summary knowledge when making trait judgments about self and others (for reviews, see Klein, 2004; Klein, Robertson, Gangi, & Loftus, 2008). But why form trait summaries if one can consult the original episodes on which the summary was based? Conversely, why retain episodes after having abstracted a summary representation from them? Are there functional reasons to have trait information represented in two different, independently retrievable databases? Evolution does not produce new phenotypic systems that are complex and functionally organized by chance. Such systems acquire their functional organization because they solved some evolutionarily recurrent problems for the organism. In this article we explore some of the functional properties of episodic memory. Specifically, in a series of studies we demonstrate that maintaining a database of episodic memories enables its owner to reevaluate an individual's past behavior in light of new information, sometimes drastically changing one's impression in the process. We conclude that some of the most important functions of episodic memory have to do with its role in human social interaction.

Anatomists dissect organs of the body. Dissection does not imply random cutting; it is a theoretically driven attempt to divide the body's parts into functional units. By contrast, psychologists rarely dissect the brain physically. Rather, we dissect it

We wish to thank Mike Miller, Nancy Collins, and Judith Loftus for their insightful comments on earlier versions of this article. We also want to thank Jeff Sherman for his expert assistance in helping shape an initially unwieldy manuscript into a far more concise, clearly stated set of arguments. Finally, the first author would like to acknowledge the NSF, whose opinion that the present research was uninteresting and possibly undoable provided considerable motivation for completing the project described herein.

Correspondence concerning this article should be addressed to Stan Klein, Department of Psychology, UCSB, Santa Barbara, CA 93106. E-mail: klein@psych.ucsb.edu.

conceptually. Most research in psychology aims to characterize the information-processing architecture of the brain, to dissect the mind into *functional units*. Although this requires theories of function, psychologists often define function in a rather impoverished way.

Take for example, the construct "memory," which has been traditionally defined as a cognitive system whose operation enables the encoding, storage and retrieval of information (e.g., Crowder, 1976; Gregg, 1986; Haberlandt, 1999; Hunt & Ellis, 1999; Lachman, Lachman, & Butterfield, 1979; Roediger, Dudai, & Fitzpatrick, 2007; Willingham, 2001). Psychologists have spent many years studying the processes that encode, store, and retrieve information, assuming they are equal across all domains—as if it makes no difference whether the words you are hearing are unrelated items on a list or your spouse explaining that s/he's fallen in love with someone else.

Psychologists have been cataloging samples of the seemingly inexhaustible set of things a memory system can do, without asking what it was designed to do. The current research shifts the focus away from capability by dissecting memory into a set of highly ordered, interlocking functional units. In doing so, we address some of the fundamental issues ignored by previous researchers and reveal that memory is not just a storage unit but an adaptive tool for sociality.

CAPABILITY VS. FUNCTION

One way to study the functional design of naturally selected systems is to think of them as components of a machine, and then distinguish the machine's capabilities from its functions (e.g., Anderson, 1991; Cosmides & Tooby, 1992; Dawkins, 1986; Klein, 2007; Klein, Cosmides, Tooby, & Chance, 2002; Williams, 1966). To specify a machine's function is to specify what it was *designed* to do. A three-hole punch, for example, is designed to put holes in writing paper so it can be stored in a three-ring binder. Knowing this function allows one to understand why its parts exist in their present form: Why it has elements sharp enough to cut paper, why there are exactly three of them, why they form a straight line, and so on. These elements are *design features* —aspects of the machine that are there because they contribute to its function.

Yet every machine is capable of doing an endless series of things that it was not designed to do. As many children discover, if you shake a well-used three-hole punch, confetti comes out. But confetti making does not explain the presence or arrangement of the punch's parts (e.g., why aren't there more punches to make more confetti? Why are they evenly spaced? Why always round?). Nor do any of the punch's other capabilities—its usefulness for weighting down paper, for example. These capabilities are arbitrary with respect to its intended function, by-products of the machine's design (e.g., Klein, 2007; Klein, Cosmides, Tooby, & Chance, 2002).

In this article we undertake a functional analysis of systems of memory by considering some of the adaptive problems they were designed to solve. Specifically, we focus on the role of episodic and semantic memory in the representation and utilization of trait knowledge about persons. The domain of trait judgment seems

well suited to studying the functional design of memory. Extracting accurate personality information and using it to predict behavior is a phylogenetically ancient problem with large fitness consequences—the kind of ability for which one might expect selection to have created cognitive adaptations (e.g., Byrne, 1995; Cosmides & Tooby, 1997; Donald, 1991; Humphrey, 1984, 1986; Mithen, 1996; Rushton, Bons, & Hur, in press; Whiten & Byrne, 1997; for a recent review, see Nettle, 2006). Studies of modern human hunter-gatherers converge on a picture of ancestral hunter-gatherers as long-lived, highly social, and living in relatively stable social groups—conditions that would favor the evolution of machinery that is good at extracting personality information and using it to predict behavior (e.g., Byrne, 1995; Donald, 1991; Dunbar, 1988, 1996; Mithen, 1996). Primate studies suggest this ability is phylogenetically ancient: Chimpanzees and bonobos adjust their behavior in ways that reflect the personality of different interactants (e.g., de Waal, 1982; Donald, 1991; Tomasello & Call, 1997). The human lineage is thought to have split from the chimpanzee lineage 5 to 7 million years ago (e.g., Mithen, 1996; Passingham, 1982; Takahata, Satta, & Klein, 1995), so there has been a long period of time during which selection could have specialized and improved the mechanisms that generate these judgments.

Humans, like other social animals, interact with each other in different types of situations (joint foraging, mating, zero-sum bargaining, parenting, friendship, aggression). Models that allow for the accurate prediction of how others will behave enable humans to maximize their payoffs from interactions (e.g., avoiding cheaters). These models are essentially what we call "personality." Categorizing people (both ourselves and others) along personality trait dimensions is a pervasive aspect of human life (e.g., Funder, 1995; Hampson, 1982; Mischel, 1968; Schneider, Hastorf, & Ellsworth, 1979). We automatically condense the rich complexity of human actions into a limited number of personality dimensions that have predictive validity (e.g., Funder & Sneed, 1993), and use these trait categorizations when deciding how to interact with others (e.g., Hampson, 1982).

There is widespread agreement among person memory theorists that both semantic and episodic memory store information about a person's traits (e.g., Craik et al., 1999; Kihlstrom & Klein, 1994; Klein, 2004; Klein & Kihlstrom, 1998; Klein & Loftus, 1993a; Klein, Sherman, & Loftus, 1996; Linville & Carlston, 1994; Sherman, 1996; Sherman & Klein, 1994; Tulving, 1993a; Wyer & Carlston, 1994). Over the past two decades, Klein and Loftus and colleagues have proposed and tested a model of the relation between semantic and episodic memory and trait conceptions of self and other (for reviews, see Klein, Cosmides, Tooby, & Chance, 2002; Kihlstrom & Klein, 1994, 1997; Klein, 1999, 2001, 2004; Klein & Kihlstrom, 1998; Klein & Loftus, 1993a; Klein, Robertson, Gangi, & Loftus, 2008). Their model incorporated three features. First, long-term knowledge of a person's traits is abstracted from multiple experiences with trait-relevant behavioral information and represented in semantic memory in summary form. Second, trait judgments about a person are made by accessing these summary representations without reference to the specific behavioral experiences from which they presumably were derived. And third, summary trait representations are functionally independent of memories of trait-relevant behavioral experiences.

WHY TWO MEMORY SYSTEMS?

Social interaction often requires rapid decisions, some of which are best made by taking into account a diverse array of information (e.g., Brothers, 1997; Byrne & Whiten, 1988; Donald, 1991; Humphrey, 1984; Klein, Cosmides, Tooby, & Chance, 2002). But searching for, retrieving, and integrating the information needed to make a judgment would be a slow process if the only database with pertinent information was the episodic store. Faster decisions can be made when answers have been precomputed and are therefore available when needed. Trait generalizations are precomputed summaries of the dispositions an individual has manifested in various behavioral episodes. These trait summaries form a fast access database, which provides quick answers to decision processes that require trait judgments.

Retrieving a summary may be faster than constructing a judgment on-line from the database provided by episodic memory, but it necessarily is less accurate. A trait summary lacks information present in the original learning experiences from which it was derived. Accordingly, maintaining a database of episodic memories solves several problems that a trait summary cannot (for discussion, see Klein, 2004; Klein, Cosmides, Tooby, & Chance, 2002). Trait summaries give information about behavior under "average" circumstances. It does not tell you under what circumstances a behavior deviates from average. Memories of behavioral episodes can provide boundary conditions on the scope of generalizations. Thus an excellent package of speed plus accuracy can be engineered into a decision system by jointly activating a trait summary and episodic memories that are *inconsistent* with it (e.g., Klein, Cosmides, Tooby, & Chance, 2002; for the relevant empirical findings, see Babey, Queller, & Klein, 1998 and Klein, Cosmides, Tooby, & Chance, 2001).

We proposed in Klein, Cosmides, Tooby, and Chance (2002) that another potential benefit of episodic memory is reevaluating the conclusions we draw about others. For example, your impression of Bob as a friendly person, based on his past willingness to help you with household repairs, may take on different significance if you subsequently learn that he is attracted to your wife. If the original learning experiences were lost after they had been analyzed to form a summary judgment of Bob's helpful character, reevaluating his past actions in light of new information about his intentions and values would be impossible. Examining this process—which we call the impression reevaluation hypothesis—is the focus of the research described in this article. At a more general level, our goal is to shed light on a possible social function of episodic memory and, in the process, gain a better understanding of how semantic memory both works with and compliments episodic function, enabling their possessor to navigate the complex web of relations that characterize human social interaction.

TESTING THE IMPRESSION REEVALUATION HYPOTHESIS

A good way to discover what something is useful for is to look at situations in which it no longer is present (e.g., Weiskrantz, 1997). In our first study, we drew on this logic to test the hypothesized role of episodic memory in reevaluating conclu-

sions we draw about an individual's personality traits. Our goal was to examine what happens to our impression of a person when we encounter new information calling our initial impression into question, but no longer are able to recall the original experiences on which our impression was based.

Participants were shown two paragraphs describing a target individual named Susan Bower. The paragraphs are described below:

Paragraph 1. Susan Bower is a 35-year-old woman who lives in Chicago. For the past five years, she has worked as an interior designer, specializing in business office design. Susan has a 10-year-old daughter who often accompanies her during business travel out of town. For the past four months, Susan has been seeing another interior designer named Ray Lisker. Susan and Ray have spent many romantic evenings together. Lately, Susan has been visiting Ray's apartment two or three times per week, and she would like to visit him more frequently. Ray finds Susan a very attractive and pleasant companion. He is interested in developing a serious relationship with Susan.

Paragraph 2. Susan Bower is married to a man named Henry Bower. Henry is a successful architect who is admired and respected by his colleagues. Susan and Henry have been happily married for 14 years and they have a 10-year-old daughter.

As we show below, pretesting revealed that each paragraph, considered separately, conveyed a positive impression of Susan Bower. However, when the first paragraph was reconsidered in light of information in the second paragraph, the implication of infidelity led participants to change their initially positive impression of Susan Bower to a negative evaluation.

There are two questions of interest: (i) are detailed trait summaries spontaneously computed on the basis of a single episode? and (ii) does reevaluation of these trait summaries require access to the facts presented in the original episode from which these initial impressions were derived?[1] To explore these questions, we manipulated participants' access to the original learning episode at the time they read paragraph 2. This was done by varying the delay between presentation of the first and second paragraphs about Susan Bower. Half the participants were shown the second paragraph one hour after reading the first paragraph (brief delay), and half saw the second paragraph one month after reading the first (long delay). As pretesting showed, people can recall a great deal about Susan Bower after a one hour delay, but very little after a one month delay—participants in the

1. To put this somewhat differently, what we are trying to show is that, absent episodic recollection, semantic summary knowledge has limits (in the present case, on judgment made in the person domain). Things did not have to be this way: As one reviewer has argued, it could be the case that semantic memory for traits stores enough information about specifics to permit reevaluation to take place without a need for episodic recollection.

2. It is important to keep in mind that we are not arguing that episodic memory for the specific acts performed by an individual are likely to be forgotten after a month. This demonstrably is not the case. What we are arguing is that, in the context of our study, participants are unlikely to recollect the specific details of a paragraph describing a fictitious other after passage of one month—especially when participants are given no reason to believe they will need to maintain access to this information at the completion of the initial testing session. In short, our delay manipulation was not designed to reflect the temporal course of episodic recollection in natural circumstances, but rather to provide an artificial context in which we could examine how judgments are affected when episodic recollection no longer is an option.

one month condition are functionally amnesic.[2] We predicted that if reevaluating one's impression of a target in light of new information requires access to detailed information present in the original episode, then two things should happen. First, participants in the one hour delay condition should form a negative impression of Susan Bower following presentation of paragraph 2. This is because these participants will have access to many of the facts presented during the original episode and thus be able to reinterpret that information in light of new information presented in paragraph 2.

Second, participants in the one month delay condition should maintain their positive impression of Susan Bower following presentation of the positive information in paragraph 2. This is because they will have forgotten most of the specific information on which their initial impression of the target was based, making it impossible to reevaluate her past actions in light of new information about her behavior and values.

By contrast, if memory of the original learning experiences is not a critical component of the impression reevaluation process, then (by extrapolation from the pretest data) participants in both the one hour and the one month delay conditions should see Susan Bower in a negative light following presentation of the material in paragraph 2.

To test these hypotheses, participants made two sets of personality judgments about Susan Bower. The first set of judgments was made either one hour or one month following presentation of paragraph 1. Judgments included participants' general impression of Susan Bower and ratings of Susan Bower on 14 trait dimensions (e.g., polite, honest, mature, intelligent). Following completion of the first set of personality judgments, participants read the second paragraph describing Susan Bower and completed a second set of personality judgments.

PRETESTING THE MATERIAL

Pretesting had two purposes. First, we wanted to be certain that each of the Susan Bower paragraphs, when tested individually, promoted a positive impression of the target; but when tested sequentially, would result in a negative impression. Second, we wanted to be sure that the passage of one month would be sufficient to produce a greatly impoverished memory for the specific statements about the target presented in the text.

METHOD

PARTICIPANTS

Seventy-five undergraduates enrolled in an introductory psychology course at the University of California, Santa Barbara participated as part of their course requirements. They were tested individually in sessions lasting approximately one hour.

MATERIALS, DESIGN, AND PROCEDURE

The stimulus items were two paragraphs (see above), each of which described a target individual named Susan Bower. Material was presented in booklet form. Participants shown only a single paragraph were given two minutes to read the text. Participants who viewed both paragraphs sequentially were allotted three minutes. Earlier testing indicated that these time periods were sufficient to allow our participants to complete their reading at a relaxed pace.

Participants were randomly assigned to one of five experimental groups. The first three groups allowed us to assess the validity of our assumption that paragraphs 1 and 2, taken individually, portray Susan Bower in a positive light, but considered together produce a negative evaluation of the target. Groups 4 and 5 enabled us to assess the efficacy of the one month delay condition as a means of impairing recall.

IMPRESSION FORMATION CONDITIONS

Group 1 was shown a picture of Susan Bower along with the text of Paragraph 1. Following completion of that task, participants were asked to do a series of unrelated tasks (e.g., word jumbles, Fibonacci puzzles) in a session lasting one hour. After completion of the unrelated task set, participants completed a questionnaire asking them to rate Susan Bower on a series of 14 personality traits and to provide their overall impression of her. To aide recollection, participants were re-presented the picture of Susan Bower that accompanied Paragraph 1.

The second group was treated identically to the first group with the following exception: The text accompanying the photograph of Susan Bower was that of Paragraph 2.

A third group was treated identically to groups 1 and 2, with the exception that they were asked to read paragraphs 1 and 2 in succession.

MEMORY CONDITIONS

A fourth group of participants were treated identically to group 1 with one important exception: Instead of asking them to rate Susan Bower at the end of the session, we asked them to recall as much as they could of the information they had read about her during the initial part of the study. Participants were allotted 2 minutes for recall and informed that they were free to remember the material in any order it came to mind.

Group 5 was treated identically to group 4, except that participants in this condition waited one month before being asked to recall what they had read about Susan Bower during the initial testing session.

The rating sheet consisted of a series of scales asking participants to rate Susan Bower on a 9-point scale anchored by -4 (less than most people) and +4 (more than most people). The questions included one global evaluation of the target: "What is your general impression of Susan Bower?" followed by 14 trait ratings (brave, honest, kind, friendly, polite, mature, moral, intelligent, cautious, popular, faithful,

independent, practical, and considerate). The traits were selected from the norms provided by Anderson (1968) and Kirby and Gardner (1972), and represented a range of attributes that an independent group of raters (N = 27) felt would be applicable to the target based on their reading of paragraphs 1 and 2.

RESULTS AND DISCUSSION

IMPRESSION FORMATION

Pretesting revealed that each paragraph, considered separately, conveyed a positive impression of the target (*M*s = 2.47 and 1.93, for paragraphs 1 and 2, respectively). One-tailed t-tests revealed that both means were significantly above the zero point on the rating scale, both ts > 2.11, $p < .01$. However, when the first paragraph was reconsidered in light of information provided by the second paragraph (group 3), the participants' impression of the target changed from initially positive (one assumes, given the results just presented) to a decidedly negative evaluation ($M = -1.47$, a value reliably below the scale midpoint point, $t(14) = 2.21$, $p < .05$, one-tailed).

Trait ratings showed a similar pattern of effects. For statistical computation, the mean rating across participants for each of the 14 traits served as the unit of analysis. Participants viewing the positive paragraphs describing Susan Bower (groups 1 and 2) rated her positively along all 14 trait dimensions (overall *M*s = 1.79 and 1.19, for groups 1 and 2, respectively; both means reliably greater than the scale midpoint, ts > 2.02, $p < .05$, one-tailed). By contrast, participants in group 3 held a negative opinion of the target, with 9 out of 14 traits rated negative in valence ($M = -.67$). Although this value was in the intended direction, its distance from the scale midpoint fell just short of significance ($.05 < p < .10$, one tailed): More importantly, however, the value obtained was reliably lower than the mean trait ratings for either group 1 or 2 (both ts > 2.47, $p < .01$, one-tailed).

MEMORY

An independent set of raters (N = 4) identified 13 separate facts that could be recalled from paragraph 1 (e.g., target's name, age, gender, facts about her relationship with Ray). Two additional raters, blind to the hypotheses of the study, scored participants' responses for accuracy of recall. Agreement between raters was high ($r = .91$, $p < .01$), with the few disagreements ($M = 7$) resolved by mutual consent.

As expected, long versus short delay had a pronounced effect on participants' ability to recall the specific information described in paragraph 1. Analysis of recall data revealed that participants in the one hour delay condition remembered a considerable amount of the information in the first paragraph ($M = 7.9$), whereas participants in the one month delay condition recalled hardly any of the facts they had read one month earlier ($M = 1.4$, $t(28) = 15.17$, $p < .0001$). Approximately half the participants in the one month delay condition recalled nothing, and the remainder were most likely to produce vague details such as "she had a good job

in an office," "She's about 30 years old." Of particular relevance for our proposed studies, not a single respondent in the one month delay condition remembered information pertaining to Susan's relationship with Ray.

In sum, pretesting confirmed that the stimulus material works as required: (a) Each paragraph, considered separately, conveyed a positive impression of the target. When considered together, by contrast, they produced a negative evaluation, and (b) Following a one hour delay, participants showed substantial recollection for the stimulus items (approximately 8 out of 13 facts); by contrast following a month delay, memory for text was close to zero.

STUDY 1: WHEN TWO RIGHTS MAKE A WRONG

The purpose of our first study was to test (i) whether presentation of a single episode about Susan Bower will trigger the computation of detailed trait summaries that can be retrieved even after a long delay, and (ii) whether the ability to revise these trait summaries in light of new information depends on access to facts presented in that single, initial episode.

In phase 1, participants were shown a photo of Susan Bower, asked to read paragraph 1, and then spent 60 minutes on a series of tasks unrelated to the study. In phase 2, they were shown the photo of Susan Bower once again, and asked to evaluate her on the series of personality and impression questions described above. They were then shown her photo again and asked to read paragraph 2, after which they were asked to evaluate Susan Bower once again on the same series of personality and impression questions.

Half the participants completed phase 2 after a delay of only one hour. Based on pretesting, these participants should still be able to recall many facts about Susan Bower. The other half of the participants completed phase 2 after a delay of one month—pretesting showed that with a delay this long, participants will be able to recall almost no facts about Susan Bower.

With this design, we can determine (i) whether detailed trait summaries are computed in response to a single episode about a novel person, (ii) whether people who are functionally amnesic for facts presented in the original episode can, nevertheless, access detailed trait summaries after a long delay, and (iii) whether the inferences that would lead to a reevaluation of these summaries can be made by people who are episodically amnesic for the details of the original episode.

METHOD

PARTICIPANTS

Forty undergraduates enrolled in an introductory psychology course at the University of California, Santa Barbara participated as part of their course requirements. They were tested individually in sessions lasting approximately one hour.

MATERIAL, DESIGN, AND PROCEDURE

The stimulus material was the same as that used during pretesting. Participants randomly were assigned to either a one hour delay or a one month delay condition. In the one hour delay condition, participants were shown a photograph of a woman named Susan Bower below which was printed the text of paragraph 1. Two minutes were allotted to read the text. Following completion of that task, participants were asked to do a series of unrelated tasks in a session lasting one hour. On completion of these distracter tasks, participants were asked to rate Susan Bower on a series of 14 personality traits and to provide their overall impression of her. To aide recollection, participants were re-presented the picture of Susan Bower that had accompanied Paragraph 1. Immediately following the rating task, participants were again shown the picture of Susan Bower, this time accompanied by Paragraph 2. Two minutes were provided to read the text. Participants were then asked once again to complete the rating questionnaire for Susan Bower. A picture of the target was present to aide recollection.

The one month delay condition was identical to the one hour condition, with one change. After completion of the distracter tasks, participants were dismissed. Testing was resumed and completed one month later.

Two versions of the trait and impression questionnaire were prepared. They were identical in content, but differed in the order in which the questions were presented. Each participant was randomly assigned one of these two versions when performing the initial rating task. The other member of the pair served as the questionnaire for the final rating task.

It is important to note that we did not test memory for factual information found in the text. Our reason for abstaining from this procedure was to ensure trait and impression ratings would not be colored by episodic recollection. (It has been repeatedly shown that participants can, under certain circumstances, use either the content of retrieval; e.g., Klein, Cosmides, Tooby, & Chance, 2001; Klein & Loftus, 1993a; Klein, Sherman, & Loftus, 1996, or the ease of retrieval; e.g., Caruso, 2008; Schwarz, 1998, 2004, to inform a subsequent trait judgment.) Informal post-test interviews showed that participants in the one month delay condition were unable to recall any information about Susan Bower's relation with Ray, confirming what we learned from pretesting: participants lack access to facts from the original learning episode following a one month delay.

RESULTS AND DISCUSSION

Are general impressions of Susan Bower revised in light of new information, and does this happen only when people can recall facts from the original episode?

This can be tested by comparing the two general impression ratings (after reading the initial paragraph, #1, and then after reading the final paragraph, #2) for participants in the one hour and one month delay conditions. A 2 (rating session:

TABLE 1. General Impressions and Mean Trait Ratings of Susan Bower

General Impression	Delay	
Session	**One Hour**	**One Month**
Initial	2.65	2.05
Final	-1.65	2.40

Trait Ratings	Delay	
Session	**One Hour**	**One Month**
Initial	1.60	1.51
Final	-.79	1.54

initial vs. final) X 2 (delay: one hour vs. one month) mixed analysis of variance (ANOVA) was conducted on participants' general impressions of Susan Bower (Table 1, top panel). The two significant main effects were qualified by the predicted interaction between rating session and delay—Interaction: $F(1, 38) = 81.55$, $p < .001$; Main effects: one month versus one hour, $F(1, 38) = 20.11$, $p < .001$, initial versus final session, $F(1, 38) = 58.85$, $p < .001$.

Importantly, Tukey tests ($p < .05$) revealed that participants' general impression of Susan Bower varied reliably across repeated testing in the one hour delay condition, from positive ($M = 2.65$) to negative ($M = -1.65$) for the initial and final rating sessions, respectively. This shows that people who can recall facts about Susan Bower from the original episode do in fact revise their general impression of her in light of new information. By contrast, in the one month delay condition, no reliable differences across sessions were found (*M*s = 2.05 and 2.40, for the initial and final testing, respectively). That is, participants who could not recall facts about Susan Bower from the initial episode did not revise their initial, positive impression of her in light of what they learned in the second paragraph.

Are trait ratings of Susan Bower revised in light of new information, and does this happen only when people can recall facts from the original episode?

To answer this question, we extracted a positivity rating for Susan Bower by averaging across all the trait ratings, and conducted a 2 (rating: initial vs.final) X 2(delay: one hour vs.one month) mixed ANOVA on them (Table 1, bottom panel). The trait ratings showed the same pattern as the general impression ratings: The two significant main effects were qualified by the predicted interaction—Interaction: $F(1, 26) = 51.47$, $p < .001$; Main effects: one month versus one hour, $F(1, 26) = 17.14$, $p < .01$, initial versus final session, $F(1, 26) = 45.98$, $p < .001$.

As can be seen in the bottom panel of Table 1, participants in the one hour delay condition rated Susan Bower more positively during the initial rating session than during the final session (*M*s = 1.60 and -.79, for the initial and final rating, respectively), indicating that they were able to revise their trait ratings in light of the new information presented in paragraph 2. By contrast, participants in the one month delay condition, who were unable to recall facts from paragraph 1, remained positive in their trait ratings: no difference across sessions was observed (*M*s = 1.51

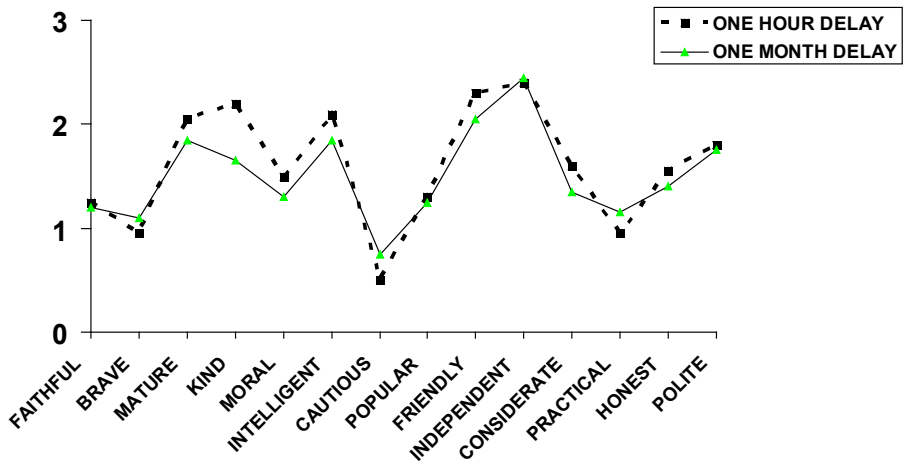

FIGURE 1. Mean trait ratings across delay conditions following presentation of only paragraph 1 (Susan Bower, Study1).

and 1.54, for the initial and final ratings, respectively). These observations were confirmed by Tukey tests ($p < .05$).

Is a single episode about a novel person sufficient to trigger computation of detailed and stable trait summaries?

It is interesting to note that participants' semantic summary knowledge of the target, even following a delay of one month, was both detailed and nuanced. This can be seen by examining the correlation between the initial task trait ratings produced in response to paragraph 1 (i.e., ratings collected prior to presentation of paragraph 2) for participants in the one hour and one month delay conditions. As shown in Figure 1, participants in the one month delay condition produced ratings that tracked almost perfectly those provided by participants in the one hour delay condition, $r = .94$, $p < .001$. We know the trait ratings produced by participants in the one month condition are not based on retrieval of facts from the original episode because, after a one month delay, people can recall virtually none of these facts. The high correlation between the trait ratings made by participants in the one hour and one month condition therefore indicates that participants in the one month condition were basing their ratings on precomputed trait summaries, not on facts from a retrieved episode.

This means that a single episode about Susan Bower was sufficient to trigger the computation of very detailed trait summaries that were still accessible one month later. In keeping with previous findings of functional independence between episodic and semantic trait knowledge (e.g., Babey et al., 1998; Klein & Loftus, 1993a; Klein, Loftus, & Kihlstrom, 1996; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman, 1996; Sherman & Klein, 1994), the high correlation between the trait ratings made by participants in the one hour and one month delay conditions indicates that participants could access detailed semantic trait summaries of the target, even when they were unable to remember the specific information from which that summary knowledge was derived.

"I'VE JUST SEEN A FACE:" CONTROLLING FOR THE EFFECT OF FACE CUES

We included the photograph of Susan Bower during rating sessions to serve as a cue to help participants identify and remember the target (pretesting suggested this would be particularly beneficial to participants in the one month delay condition). However, adoption of this procedure may have had an unintended consequence. Specifically, one could argue that participants were not really using a detailed semantic summary to inform their ratings, but rather were making guesses based on the picture of Susan Bower. Since participants in both conditions had access to the same photograph, this could explain the high correlation reported above.

To examine this possibility, an additional group of 20 participants were shown the picture of Susan Bower, unaccompanied by text, and asked to evaluate her on the same 14 traits used in the main study. We then correlated their responses with the initial task trait ratings (i.e., ratings made prior to presentation of paragraph 2) provided by participants in the one hour and one month delay conditions. Both values obtained (rs = .41 and .37, for the one hour and one month delay groups, respectively) were reliably less than the correlation computed between the two delay groups (r = .94; both ps < .05). This indicates that trait summaries for participants in both delay groups were based on the facts presented in the original episode, and not merely on the photograph of Susan Bower.

As an additional check on the possibility that the photograph of Susan Bower might, by itself, account for the high level of agreement across delays in participants' general impressions following presentation of paragraph 1, the same control participants were asked to rate the target's picture, absent text, on the 9-point general impression scale used in the main study. T-tests ($p < .05$) revealed that their mean rating (M = .55), though above the scale midpoint, was significantly less positive than initial task ratings produced by either the one hour (M = 2.65) or one month (M = 2.05) delay groups.

In sum, it appears that, despite their inability to recall specific textual material, participants in the one month delay condition retained access to a detailed, well-differentiated summary representation of the target. It's just that, lacking episodic recollection, they were unable to modify this abstract knowledge in the presence of specific factual information provided by paragraph 2.

Summary for Study 1

Taken together, these findings are consistent with the proposition that, absent episodic recollection (i.e., the one month delay group), participants ability to reevaluate an existing impression in light of new information is limited. Specifically, while the one hour delay group was able to reconsider their initially positive impression of the target in light of the information provided in paragraph 2, participants in the one month delay group, lacking recollection of the specifics that led to their positive evaluation, were unable to revise and correct their existing impression in light of information made available by paragraph 2. Absent specific episodic memories of the original learning event, participants were forced to rely on the semantic summary abstracted at the time of learning. And that summary, by itself, did not signal a need for reevaluation.

Study 2. Generalizing the Susan Bower Study Findings: When Two Wrongs Make a Right

The Susan Bower study showed that joint consideration of two evaluatively positive descriptions of a target could yield a negative impression, provided participants had access to facts present in the original learning episode. Although this finding supports the hypothesized role of episodic memory in impression reevaluation, it is important to demonstrate the generality of the effect by showing it is not dependent on the evaluative valence of the material presented. In the present study we replicated the Susan Bower findings using two paragraphs that each convey a negative impression of the target when considered alone but, when considered together, support a positive impression of his character (i.e., the target's behavior is seen as both heroic and moral when the original paragraph is reconsidered in light of the material contained in the second paragraph). The paragraphs are shown below:

Paragraph 1. Tom West is a 35-year-old man who lives in Dallas. He has been married for 11 years and has three young children. Because Tom's wife has chosen to at stay home and raise their children, the family is dependent on Tom for their financial well-being. Today, without warning, Tom quit his job, leaving the company when they needed him most. Tom knows that his sudden decision also leaves his family without any means of support. Nonetheless, despite the hardships his decision will bring for his family and the company, he is convinced that quitting is something he wants to do.

Paragraph 2. Tom West worked as Public Relations Director for ExoTech, an international manufacturing company based in Dallas. Tom's duties were to present a positive corporate image of ExoTech to the public. ExoTech recently was accused of dumping dangerous pollutants into a river that serves as the primary water source for a small community. The company publicly denied any responsibility in the matter. However, in a meeting held yesterday, the company privately admitted to Tom that it was guilty of wrongdoing and asked him to coordinate a massive corporate cover-up.

Pretesting (M = 60) showed that each paragraph, considered separately, conveyed a negative impression of the target: On a scale ranging from -4 (negative impression) to +4 (positive impression), we obtained mean ratings of -1.00 and -1.50, for paragraphs 1 and 2, respectively. However, when the first paragraph was reconsidered in light of information in the second paragraph, the initially negative impression of the target became positive (M = +2.20). Of interest, a separate group of 10 participants asked to give their general impression of Tom West exclusively on the basis of his picture (i.e., no text presented) rated him +.80—that is, positively rather than negatively. Apparently, the content of each paragraph, not the affective reaction to his photograph, largely determined participant ratings.

With regard to memory, recall of the specific information in Paragraph 1 was significantly higher following a one hour delay (M = 7.4 out of the 11 facts identified by 4 independent raters) than following a one month delay (M = .85). Analyses performed on the impression and recall values yielded the same pattern of statistically reliable effects ($p < .05$, one-tailed) as those reported in Study 1.

The stimulus material performed as expected. Armed with this knowledge, we turned attention to the main study. If the results of the Susan Bower study reflect the hypothesized role of episodic memory in impression reevaluation (and are not

simply an artifact of the stimulus material used), we should be able to replicate those findings with the material describing Tom West. Specifically, participants in the one hour delay condition should form a positive impression of Tom West after reading paragraph 2. This is because their recollections of the material contained in paragraph 1 provide a context within which to reinterpret the material in paragraph 2.

By contrast, participants in the one month delay condition should maintain a negative impression of Tom West following presentation of paragraph 2. This is because they will have forgotten most of the material on which their initial impression of the target was based, and without access to this information, they will be unable to reconsider the original learning episode in light of new information provided by the second paragraph.

METHOD

PARTICIPANTS

Forty undergraduates enrolled in an introductory psychology course at the University of California, Santa Barbara participated as part of their course requirements. They were tested individually in sessions lasting approximately one hour.

MATERIAL, DESIGN, AND PROCEDURE

The material, design and procedure were identical to those used in Study 1 with one important change: Text and pictures now referred to Tom West.

RESULTS AND DISCUSSION

Are general impressions of Tom West revised in light of new information, and does this happen only when people can recall facts from the original episode? As before, this can be tested by comparing the two general impression ratings (after reading the initial paragraph, #1, and then after reading the final paragraph, #2) for participants in the one hour and one month delay conditions. A 2 (rating session: initial vs. final) X 2 (delay: one hour vs. one month) mixed ANOVA was conducted on participants' general impressions of Tom West (Table 2, top panel). Paralleling the results of Study 1, the two significant main effects were qualified by the predicted interaction—Interaction: $F(1, 38) = 16.84$, $p < .01$; Main effects: one month versus one hour, $F(1, 38) = 41.33$, $p < .001$; initial versus final session: $F(1, 38) = 6.25$, $p < .05$.

As anticipated, participants in the one hour delay condition changed their initially negative impression of the target (M = -1.15) to a positive evaluation (M = 2.70) after reading the content of paragraph 2. This shows that those who can recall facts about Tom West presented in the initial episode can revise their impression of his character. This was not true of participants who could not recall these facts:

TABLE 2. General Impressions and Mean Trait Ratings of Tom West

General Impression	Delay	
Session	**One Hour**	**One Month**
Initial	-1.15	-.60
Final	2.70	.25
Trait Ratings	**Delay**	
Session	**One Hour**	**One Month**
Initial	-.91	-.33
Final	1.76	.55

Participants in the one month delay showed no reliable change in their impression of Tom West after exposure to the content of paragraph 2 (*M*s = -.60 and .25 for the initial and final ratings, respectively). These observations were confirmed by Tukey tests, $p < .05$.

Are trait ratings of Tom West revised in light of new information, and does this happen only when people can recall facts from the original episode?

To answer this question, we extracted a positivity rating for Tom West by averaging across all the trait ratings. The results reveal a pattern very similar to that found for the general impressions (Table 2, bottom panel). As before, a 2 (rating session: initial vs.final) X 2 (delay: one hour vs. one month) mixed ANOVA showed the predicted interaction between delay and session—Interaction: $F(1, 26) = 22.23$, $p < .01$; Main effects: initial versus final session, $F(1, 26) = 54.28$, $p < .001$; delay condition did not produce a reliable main effect, $F(1, 26) = 2.69$, $p > .05$.

As predicted, participants in the one hour condition changed their initially negative evaluation ($M = -.91$) to a positive one following exposure to the second paragraph ($M = 1.76$). By contrast, participants in the one month delay condition did not vary reliably in their trait assessments across sessions (*M*s = -.33 and .55 for the initial and final ratings, respectively). These observations, confirmed by Tukey tests ($p < .05$), show that those who can recall facts about Tom West from the initial episode are able to revise their trait summaries in light of new evidence, whereas those who are amnesic for those facts cannot.

Is a single episode about a novel person sufficient to trigger computation of detailed and stable trait summaries?

Yes. As in Study 1, we compared trait ratings for Tom West following presentation of paragraph 1 alone for participants in the one hour and one month conditions. To assess the possibility that any high correlation is due merely to Tom West's photograph rather than to the information participants read about him, we also obtained trait ratings from a control group ($M = 20$) who read nothing about Tom West but did see his photo. With regard to general impression, the photo-only group rated Tom West more positively ($M = .90$) than did participants in either the one hour ($M = -1.15$) or the one month ($M = -.60$) delay conditions (*T*-tests, $p < .05$). This shows

that trait ratings in the delay conditions were based on information contained in the paragraphs, and not merely on Tom West's photo.

Trait ratings for Tom West were detailed and stable, revealing the same striking pattern as those for Susan Bower. As can be seen in Figure 2, initial trait ratings following exposure to paragraph 1 (but before exposure to paragraph 2) were highly correlated in the one hour and one month delay groups (r = .92). By contrast, the correlations between the ratings of photo-only participants and participants in the one hour delay (r = -.14) and one month delay (r = .14) groups were not statistically reliable (ps > .50), showing that the stability in trait summary knowledge is based on exposure to the facts in paragraph 1, and is not an artifact produced by participants having just seen the photograph of Tom West. In sum, participants in the one month condition had formed a nuanced and stable set of trait summaries for Tom West, even though they had been exposed to only a single episode and could recall virtually no facts about him.

Summary for Study 2

All of the findings from our initial study, where Susan Bower was the target, replicated when the target was changed to Tom West; the pattern of results was the same when the evaluative valence of the stimulus material was varied from "two rights make a wrong" to "two wrongs make a right." Based on a single episode, participants formed detailed trait summaries for Tom West. When presented with new evidence, memory was searched for target-relevant information. Participants in the one hour condition, but not those in the one month condition, had available the resources (i.e., episodic memories of the specific facts contained in paragraph 1) to reevaluate their original impression of the target in light of new information calling into question the factual bases of their initial impression. Accordingly, only the former group was able to revise their impression of Tom West.

STUDY 3: TESTING THE IMPRESSION REEVALUATION HYPOTHESIS BY EXAMINING RATINGS PROVIDED BY AN AMNESIC PATIENT

The findings presented appear to provide strong support for the impression reevalaution hypothesis. A question might be raised, however, concerning the extent to which the outcomes from the preceding studies shed light on the importance of episodic memory, per se, for reevaluation. That is, even though participants were able to reevaluate their impressions only in the one hour delay condition, it remains possible that semantic memory, rather than the hypothesized episodic system, mediated this effect. For example, if a semantic memory system were able to preserve, for a brief period of time (e.g., 1 hour), a record of the facts present in the original learning event, along with a more enduring (e.g., 1 month) summary impression derived from those events, the differential effects of delay on reevaluation could be explained without needing to invoke episodic recollection. This possibility must be taken seriously given that facts about one's own life are sometimes stored as semantic knowledge, and can be retrieved by amnesics who cannot recollect any particular experiences that they have had (for review, see Klein, 2001; Klein, German, Cosmides, & Gabriel, 2004).

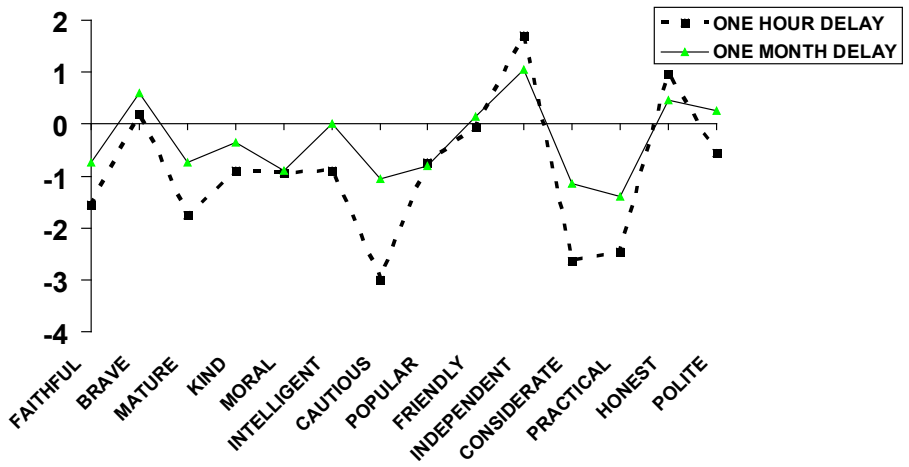

FIGURE 2. Mean trait ratings across delay conditions following presentation of only paragraph 1 (Tom West, Study 2).

In Study 3, we tested this possibility by examining a patient suffering from episodic amnesia. An amnesic individual who can access a semantic trait summary should also be able to access other semantic information—including semantic records of information in the original learning event, if these exist. If impression reevaluation depends on the ability to access a semantic record of facts presented in the original episode, rather than on episodic memory, then a patient suffering from episodic amnesia should be able to reevaluate his impression of a target. This should be the case when the delay before testing is short, as it was for participants in the one-hour condition of Studies 1 and 2. But if episodic memory is required for the impression reevaluation process, as we have argued, then such an individual should be able to form and access an initial impression of Susan Bower (relying on his intact semantic memory) yet be unable to reevaluate that impression in light of new evidence.

For this study, we enlisted the cooperation of D.B., an amnesic patient who, as a result of cardiac arrest with presumed hypoxic brain damage, was left incapable of consciously bringing to mind a single personal experience that transpired prior to his heart attack. In addition to his profound retrograde amnesia, D.B. also suffers from dense anterograde amnesia, leaving him incapable of remembering events that transpired only moments earlier (for review, see Klein, Rozendal, & Cosmides, 2002).

In contrast to his profound episodic impairment, D.B.'s semantic knowledge is largely intact. His speech is fluent and his general level of intelligence well preserved. His knowledge of word meanings is intact, as is his ability to understand and respond to questions. He knows a variety of facts about public figures and events, but can not consciously bring to mind a single experience involving any of those facts.

Although D.B.'s ability to learn new information is somewhat compromised by his episodic memory impairment, he can acquire new information semantically,

provided he receives multiple exposures to the to-be-remembered material. For example, shortly after becoming amnesic, D.B. inquired about his wife. When told she had been hospitalized, he became very distraught. A few days later he again asked about his wife, and again became emotional on hearing about her condition, although his upset was noticeably less than it was following his initial query. The next day we asked D.B. if he knew where his wife was. He immediately responded that she was "in the hospital." However, when asked how he knew this to be the case, he was unable to identify the source of that knowledge.

D.B. was run through the one hour delay condition of the Susan Bower study with two modifications. First, a second exposure to the text describing Susan Bower was added to ensure D.B. had sufficient exposure to the stimulus material to form a semantic summary of her personality. Second, to provide a conservative test of the episodic memory-impression reevaluation hypothesis, we reduced the delay between presentation of the first and second paragraphs describing Susan Bower from one hour to 15 minutes. If the episodic reevaluation hypothesis is correct, even this greatly reduced delay should prove too long to allow D.B. to reevaluate his impression in light of new information. This is because he will have lost access to information from the original learning episode within minutes after having read the first paragraph. By contrast, if semantic memory mediates reevaluation, D.B. should be able to revise his opinion of the target by accessing semantic-based memory for the specific details in Paragraph 1. By hypothesis, these details should be available in memory following a brief delay, just as available as are semantic summaries of the target's traits.

METHODS

PARTICIPANT

Patient D.B. was a 79-year-old man who became profoundly amnesic as a result of anoxia following cardiac arrest. Both informal questioning and psychological testing revealed that D.B. was unable to consciously recollect a single thing he had ever done or experienced from any period of his life. Semantic memory, by contrast, was largely spared. Case details can be found in Klein, Rozendal, and Cosmides (2002).

MATERIALS, DESIGN, AND PROCEDURE

The materials were the same as those used in Study 1 (Susan Bower). The design and procedure were altered to accommodate the patient's learning limitations (for discussion of the need for multiple presentations when working with amnesic patients, see Damasio, Tranel, & Damasio, 1989).

At the start of the session 1, D.B. was shown the picture of Susan Bower and asked to listen to the content of paragraph 1, which was read aloud by the experimenter. On completion, he filled out the impression/personality questionnaire used in Studies 1 and 2. Following a delay of 15 minutes (during each 15

minute delay D.B. performed simple tasks unrelated to the main study), session 2 began. We again presented the photograph of Susan Bower and read aloud the text of paragraph 1. He once again rated the target using the impression/personality questionnaire. After another 15 minutes had transpired, session 3 commenced: D.B. was again shown the photograph of Susan Bower, this time without the text, and asked for his impressions using the questionnaire. The final testing session (session 4) followed immediately on completion of another 15 minute delay. D.B. again was shown the picture of Susan Bower, but this time he was read the text of paragraph 2. He then rated her using the impression/personality questionnaire.

CONTROL PARTICIPANTS

To evaluate D.B.'s performance, 31 neurologically healthy control participants (mean age = 28.7 years) were tested using the same materials and procedures administered to D.B.

RESULTS AND DISCUSSION

The results are presented in Figures 3 and 4. For the purpose of analyses, we used the variance of the control group as an estimate of the population variance, from which we computed the estimated standard error of the mean (e.g., Rosenthal & Rosnow, 1991).

Can D.B. revise his general impression of Susan Bower in light of new information about her?

No. On the general impression question, both D.B. and control participants showed a gradual increase in the positivity of their ratings across the first three trials (i.e., prior to presentation of paragraph 2). Statistical testing revealed that although D.B. and the controls showed no reliable change in ratings between sessions 1 and 2, their impressions of the target became significantly more positive between sessions 2 and 3 (both ts > 4.7, $p < .05$). So repeated exposure to the first paragraph about Susan Bower improved their positive impression of her.

Like participants in the one hour condition of Study 1, the control participants of Study 3 were able to reevaluate their impression of Susan Bower in light of new information: the positive impression of her that they had expressed in session 3 changed to a decidedly negative impression in session 4, after they had heard the information contained in paragraph 2 (*M*s = 2.64 and -1.61, for sessions 3 and 4, respectively, t(30) = 13.39, $p < .001$). D.B., by contrast, showed no change in his evaluation of Susan Bower across sessions 3 and 4; the new information contained in paragraph 2 did not lead him to reevaluate the positive impression he had formed of her based on the information contained in paragraph 1.

Can D.B. revise his trait ratings of Susan Bower in light of new information about her?

No. As in previous studies, we computed a positivity score from the trait ratings that D.B. provided. Once again, D.B. and the controls evidenced no reliable change

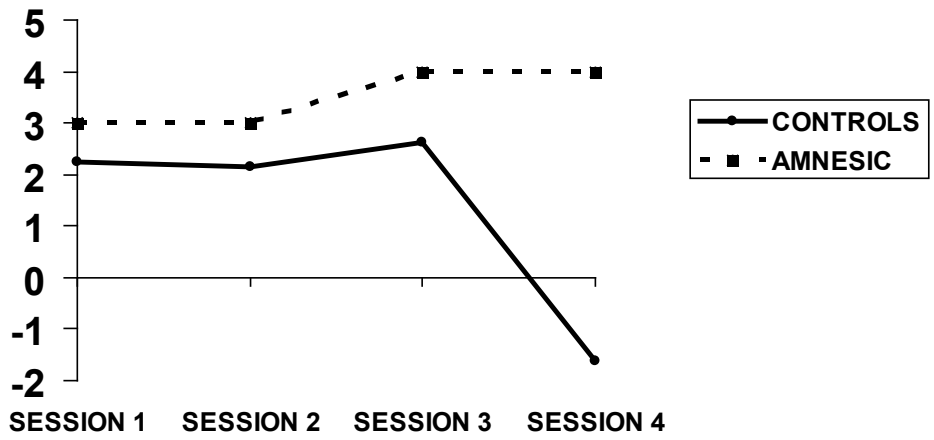

FIGURE 3. Change in general impressions across testing sessions in study 3 (D.B. and controls).

in ratings between sessions 1 and 2, and their evaluation of Susan Bower became more positive between sessions 2 and 3 (both ts > 6.12, $p < .05$). As expected, the second paragraph led the control participants to change their evaluation of Susan Bower's traits, from positive in session 3 to negative in session 4 (*M*s = 2.01 and -.84, for sessions 3 and 4, respectively, t(30) = 10.83, $p < .001$). D.B., by contrast, was unable to reverse his positive evaluation of Susan Bower in light of the new information about her husband; indeed, his trait ratings became slightly, though not significantly, more positive after hearing this information (*M*s = 3.50 and 3.71 for sessions 3 and 4, respectively, $p > .40$).

Because D.B.'s initial impression of the target was positive, it is difficult to be certain whether his impression was based exclusively on knowledge contained in semantic memory. For example, it is well-known that multiple target exposure leads to enhanced feelings of familiarity, which, in turn, foster a positive evaluation of the target; for review, see Zajonc, 1980). However, evidence presented by Damasio, Tranel, and Damasio (1989) suggests that it is the evaluative content of the text that plays the primary role in impression formation. In their study, repeated presentation of a paragraph describing a target in an unflattering light led their amnesic patient to form a negative impression of the target (for a classic demonstration of negative impressions issuing from repeated stimulus presentation with amnesic patients, see Claparede, 1911).

Was D.B. able to compute trait summaries for Susan Bower?

Further evidence that D.B.'s evaluations were based on the information he heard comes from an analysis of the detailed profile of trait ratings he provided—ratings that indicate he had computed nuanced trait summaries that were similar to those computed by the control participants.

D.B. did not give the same rating for every trait he was asked about; like control participants, his ratings varied across traits. They also became increasingly stable the more times he heard about Susan Bower: although the correlation between the

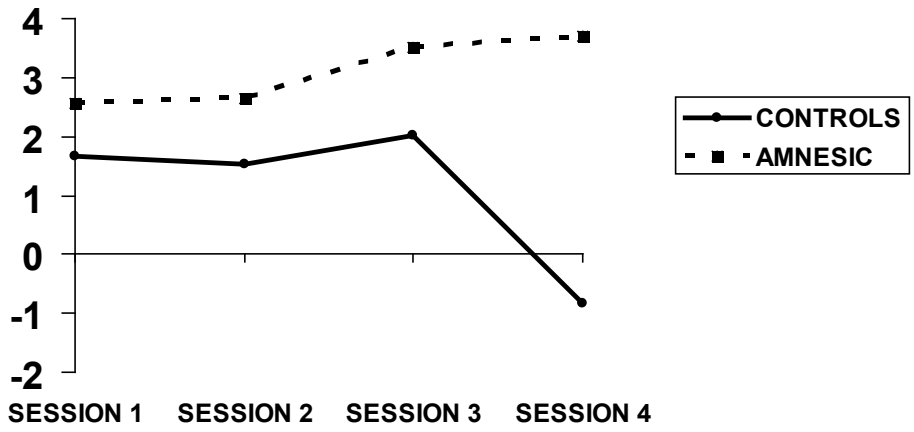

FIGURE 4. Change in mean trait ratings across testing sessions in study 3 (D.B. and controls).

ratings he provided in sessions 1 and 2 was low, it increased to $r = .36$ for sessions 2 and 3, and still further to $r = .59$ for sessions 3 and 4. Correlations for sessions 2 and 3 are the last that compare ratings based only on paragraph 1, and D.B.'s session 2-3 correlation of $r = .36$ compares favorably with those of control participants ($r = .44$). That is, the trait ratings D.B. provided were at least as stable as those provided by brain-normal participants.

Moreover, D.B.'s trait ratings were similar to those provided by control participants. We compared the trait ratings made by D.B. and by control participants for the third session—that is, after multiple presentations of paragraph 1, but prior to presentation of paragraph 2. Despite the fact that D.B. was 79 years old and brain-damaged, the correlation between his trait ratings and those of control participants in their 20s was reliable, $r = .47$, $p < .05$. That is, D.B. formed detailed trait impressions of the target that were similar to those formed by neurologically healthy participants.

We next compared D.B.'s trait judgments with those obtained from the photo-only control participants in Study 1. This enabled us to examine the possibility that D.B.'s trait impressions might be based on his affective responses to the photograph of Susan Bower, rather than semantic knowledge of her derived from material presented in paragraph 1. The correlation obtained was essentially zero ($r = .07$, *ns*). Thus, the fact that D.B.'s trait ratings correlated at $r = .47$ with those of the control participants cannot be explained by exposure to the photo alone. This provides further evidence that D.B.'s impressions of Susan Bower were influenced in an important way by the biographical material contained in the text.

Is reevaluation based on retrieval of facts from episodic or semantic memory?
The goal of Study 3 was to test the hypothesis that impression reevaluation rests on access to semantic rather than episodic memory. According to the semantic counterhypothesis, Study 1 and 2 participants could reevaluate their impressions

in the one hour, but not the one month, condition because the relevant facts about Susan Bower were briefly present in their semantic memory.

Prior testing of D.B. had demonstrated that he can still acquire semantic knowledge, especially after it has been repeated several times, as it was in Study 3. Moreover, he was able to use this information to form semantic trait summaries, reflected in ratings that were consistent with those provided by the control participants. If the ability to reevaluate impressions were based on access to facts held briefly in semantic memory, then D.B. should have been able to reevaluate his impression of Susan Bower when he heard paragraph 2. Yet he was not able to do so: Despite repeated presentation of the text of paragraph 1, D.B. failed to revise his positive evaluation of the target when tested immediately after presentation of paragraph 2. By contrast, the control participants showed clear evidence of reevaluating their initially positive impression of Susan Bower following presentation of the information contained in paragraph 2.

These findings are easily accommodated by the episodic memory explanation of impression reevaluation: Lacking the ability to retrieve facts from his episodic store, D.B. was unable to reevaluate Susan Bower's character in light of new information. They are, however, difficult to reconcile with the semantic memory hypothesis, which posits reevaluation is based on a briefly held record in semantic memory of the original learning events.

Summary of Study 3

Although D.B., a patient with episodic amnesia, could form a positive impression following multiple presentations of paragraph 1 (for related findings, see Damasio et al., 1989; Johnson, Kim, & Risse, 1985), his representation of the target was in the form of abstract trait summaries, not the specific biographical facts on which those summaries were based. Absent the biographical details normally provided by episodic memory, D.B. was unable to reconsider his positive impression in light of new information about the target's behavior.

STUDY 4: TESTING HYPOTHESIZED MECHANISMS

The impression reevaluation hypothesis argues that people, faced with new evidence about a target, will retrieve from memory both their abstract impression of the target along with specific details that served as the basis for that abstraction. The findings from our first three studies are consistent with this proposal. They do not, however, provide direct evidence for our claim that impression reevaluation entails recollection of the original learning episodes.

The purpose of the Study 4 was to provide such evidence by seeing whether facts from paragraph 2 prime recall of facts about Ray from paragraph 1—that is, facts that would provide the basis for reevaluating Susan Bower's character. Participants were treated identically to those in the one hour delay group in Studies 1 and 2, with the following changes. One hour after reading paragraph 1, participants were randomly assigned to one of three groups:

(a) Participants in the *priming* group were asked to read the second paragraph describing Susan Bower. Immediately after reading paragraph 2, participants were asked to recall as much of the material in paragraph 1 as possible.
(b) Participants in the *judgment* group were asked to provide their general impression of Susan Bower and to rate her on the 14 trait dimensions. After completing their judgments, participants were asked to recall as much of the material in paragraph 1 as possible.
(c) Participants in the *control* group were not asked to perform any tasks between presentation of paragraph 1 and the request to recall the material in paragraph 1.

If the episodic reevaluation hypothesis is correct, participants in the *priming* condition should differ from those in the *judgment* and *control* conditions in an important way. Specifically, the type of material recalled early in the free recall period should differ systematically between conditions: (a) Participants in the *priming* condition should preferentially retrieve material pertaining to Ray and his affair with Susan during the early stages of recall. This is because impression-relevant episodes (e.g., Ray and Susan's relationship) are more likely to be accessed during impression reevaluation than are impression-irrelevant episodes (e.g., Susan Bower's age, job, place of residence); thus they will be primed for recollection when systematic retrieval is requested (for a related view, see Srull, 1981). (b) By contrast, participants in the *judgment* condition should not show a preference for impression-relevant material at the beginning of the recall trial. This is because judgments about a target's personality typically do not activate the original learning episodes on which they were based (for reviews, see Kihlstrom & Klein, 1994; Klein, 2004; Klein & Loftus, 1993a; Klein, Robertson, Gangi, & Loftus, 2008). (c) Finally, participants in the *control* condition were expected to perform in a manner similar to that predicted for participants in the *judgment* condition (i.e., we had no basis for predicting preferential retrieval during early recall of facts about Ray).

METHOD

PARTICIPANTS

Seventy-two undergraduates enrolled in an introductory psychology course at the University of California, Santa Barbara participated as part of their course requirements. They were tested individually in sessions lasting approximately one hour.

MATERIALS, DESIGN, AND PROCEDURE

The material and design were identical to those for the one hour delay group in Study 1 with the following changes: First, following presentation of paragraph 1, participants were randomly placed in one of three conditions.

1. The episodic *priming* group. After reading the first paragraph describing Susan Bower, participants performed a series of unrelated distracter tasks (e.g., num-

ber puzzles) for one hour. Following delay, they were given 2 minutes to read the second paragraph describing Susan Bower. On completion of this task, they were requested to recall as much material as possible from paragraph 1.

2. The personality *judgment* group. After reading the first paragraph describing Susan Bower, participants performed a series of distracter tasks for one hour. At the end of the delay period, they were asked to provide their general impression of the target and rate her on 14 trait dimensions using the 9-point scale from Study 1. After making their ratings, they were asked to recall as much material as possible from paragraph 1.
3. The *control* group. After reading the first paragraph describing Susan Bower, participants performed a series of distracter tasks for one hour. Following the delay, they were requested to recall as much material as possible from paragraph 1.

For the purpose of memory testing, participants were given a lined sheet of paper and asked to remember as many facts about Susan Bower as possible from the material in the paragraph presented at the start of the study (i.e., paragraph 1). Participants were allowed 2 minutes for recall and told to write one fact per line. They were told they were free to recall the facts in any order that came to mind. Thirty seconds into the recall period, participants were asked to place a mark next to the most recently remembered entry on their recall sheet.

RESULTS AND DISCUSSION

Two raters, blind to the hypotheses of the study, scored participants' responses for accuracy of recall. Agreement between raters was high ($r = .96$, $p < .01$), with the few disagreements ($M = 4$) resolved by mutual consent.

Does overall recall differ across the three conditions?
No. A one-way ANOVA was performed on both overall recall and amount recalled during the first 30 seconds of testing. As can be seen in Table 3, the three experimental conditions did not vary reliably in the average amount of information recalled by each participant at either the 30 second mark or the 2 minute mark (both $Fs < 2.00$, $p > .10$). Because the instruction was to recall as much information as possible from paragraph 1, participants in the *priming* condition, who read the content of both paragraphs, had a greater potential for memory confusion than participants in either the *judgment* or *control* conditions. It is therefore interesting that they recalled just as much overall from paragraph 1 as participants in the other conditions.

Are facts about Ray primed when participants learn that Susan is married to Henry?
Yes. As predicted, the conditions differed significantly in mentions of Ray during the first 30 seconds of recall. Participants in the *priming* group were more likely to recall evidence about Susan's affair with Ray during the first 30 seconds of recall (13 out of 24 participants mentioned Ray at least once) than were participants in

TABLE 3. Mean Recall After 30 Seconds, After 2 Minutes, and Number of Participants who Mentioned Ray at Least Once in the First 30 Seconds of Recall

Condition	Mean Recall After 30 Seconds	Mean Recall After 2 Minutes	Number of Participants Who Mentioned Ray in First 30 Seconds
Priming	2.62	7.16	13
Judgment	2.21	7.92	3
Control	2.08	7.46	2

either of the other two conditions (*judgment*: 3 out of 24; *control*: 2 out of 24; $\chi^2(2) = 19.44$, $p < .01$).

Interestingly, by the end of the recall period mentions of Ray no longer differed across conditions (*M*s = 3.75, 3.21, and 3.46, for the *priming*, *judgment*, and *control* conditions, respectively), $F(2, 69) = 2.01$, $p > .10$. This means that information about Ray was available to participants in all three conditions. Therefore, the difference in the first 30 seconds means that facts about Ray were recollected faster after participants learned that Susan Bower is married to Henry Bower. In combination with the new information that Susan is married to a man named Henry, facts about her relationship with Ray imply that she is having an illicit affair. That these impression-relevant facts are retrieved first lends support to the claim that new information about a person activates trait-relevant content within past episodes that might lead to a reevaluation of that person's character.

GENERAL DISCUSSION

Endel Tulving once sent us a packet of reprints on episodic memory, with a post-it note on top saying "And what is it for?" (see also Tulving, 1985). We have argued that the answer to this question depends, in part, on what semantic memory is for. To gain purchase on this question, we have been using trait generalizations as a test system within semantic memory (e.g., Klein, 2004; Klein, Cosmides, Tooby, & Chance, 2002). The results presented herein establish several points about trait summaries and the function of maintaining a database of episodes even after they are formed.

TRAIT SUMMARIES

We have argued that trait generalizations are precomputed summaries of the behavioral dispositions an individual has manifested across many similar situations. They form a fast access database for social decision procedures that require trait judgments. This database of semantic knowledge is remarkably resilient in the face of brain damage: It has been possible to find neuropsychological patients like D.B. (Experiment 3; see also, Tulving, 1993a), who no longer can retrieve behavioral episodes from their personal past yet have preserved access to accurate trait generalizations. Behavioral episodes provide the data from which trait summaries

are computed (e.g., Klein & Loftus, 1993a; Sherman & Klein, 1994), but cases like D.B.'s show that social decision processes can access trait summaries independently of the episodes that led to their construction. The present results add to the emerging picture of how trait summaries are computed and their relationship to the episodic database.

HOW MUCH EPISODIC DATA MUST ACCUMULATE BEFORE TRAIT SUMMARIES ARE COMPUTED?

Trait summaries are inferred from behavior, but when are the social inference systems that use behavioral episodes to compute these summaries activated? Based on past research (summarized in Klein, 2004 and Klein, Cosmides, Tooby, & Chance, 2002) it was not clear whether trait summaries for a given person are computed and updated as each new episode occurs, or whether the computation of trait summaries is triggered only after the accumulation of a number of episodes. When behavioral episodes are few in number, asking for trait judgments does speed retrieval of trait-relevant episodes (e.g., Klein & Loftus, 1990, 1993a, 1993b; Klein, Babey, & Sherman, 1997; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman & Klein, 1994; but, see Klein, Sherman, & Loftus, 1996 for an important qualification). But this does not rule out the possibility that trait summaries were being computed and updated as each episode occurred.

Delaying the computation of trait summaries until there is a large database of relevant episodes would produce more reliable trait summaries, but at the cost of having no fast-access summaries when information about an individual is sparse. By contrast, immediate computation would allow trait summaries to be continuously updated, but at the potential cost of misjudgments when episodes are few. This cost could be mitigated, however, by trait judgment procedures that are designed to preferentially search and analyze situations from the episodic store when the relevant trait summaries have a low credal value (for discussion, see Klein, Comsides, Tooby, & Chance, 2002: Decision rules could assess credal value by the strength of a memory trace, its retrieval time, or an associated confidence value).

If the computation of trait summaries is delayed until there is a large database of relevant episodes, then exposure to a single episode about a novel character will likely be insufficient to trigger the formation of trait summaries about that character; only people who can recall the episode should be able to make trait judgments. But if trait summaries are computed as each new episode is encoded, then exposure to a single episode should support trait judgments even among people who are no longer able to recall the facts present in that episode—that is, among people who are amnesic for those facts (for discussion, see Klein, Sherman, & Loftus, 1996; Sherman & Klein, 1994).

We were able to test between these alternatives by comparing trait judgments made by people who could recall most of the facts present in a single episode about Susan Bower or Tom West to the trait judgments made by people who could recall virtually none of the facts presented about the targets. By introducing a one month delay between the encoding event and the time at which trait judgments were requested, we created a group of participants who were functionally amnesic (see footnote 1); D.B. provided a point of comparison because he was amnesic due to brain damage. Remarkably, a two minute episode, in which people read a

single paragraph of ~100 words, was sufficient to trigger the automatic computation of very detailed and stable trait summaries about Susan Bower or Tom West, supporting the hypothesis that trait summaries are computed and updated as each new episode is experienced (for a similar view, see Tulving, 1993b).

That participants who were tested after a delay of one hour were able to generate detailed trait judgments tells us little: Pretesting had shown that these people would have been able to recall many details from the original episode (~7 to 8 of 13 facts), so their judgments could have been based on retrieved episodes or on trait summaries. What is very surprising is that one month *after* having read 100 or so words about Susan Bower or Tom West, people were able to provide trait judgments that precisely matched the detailed ones generated by people who could still recall many of the specific facts from the initial learning episode: Their trait judgments correlated with those in the one hour condition at rs > .90.

These extremely high correlations cannot be attributed to participant's ability to judge traits from the face alone because correlations were much lower (rs ranged from -.14 to .41) when the trait judgments of people in the one month delay condition were compared with those of people who saw the photos of Susan Bower and Tom West, but had read nothing about them. That the correlation with the photo-only participants is significant means that the photos may have contributed to the high correlation between the one month and one hour participants; but that the latter effect size is approximately twice as large means that seeing the photos is not the whole story. The trait judgments of participants in the one month condition must have been based on the facts in the story, and not merely on the photos of Susan Bower or Tom West. Yet the facts in the story could not be accessed by these participants: Pretesting showed that after a one month delay, people can recall virtually nothing from the original paragraph about Susan Bower or Tom West (about 1 of 13 facts)—indeed, half of the pretested participants could recall nothing at all about Susan Bower or Tom West after a month had passed. Functionally speaking, participants in the one month delay condition were amnesic. So on-line computation of trait judgments based on conscious access to facts presented in the original learning episode is ruled out for the one month participants.

The fact that the one month trait ratings were almost perfectly correlated with the one hour ratings strongly implies that these judgments were based on precomputed trait summaries. The fact that these participants were not asked to generate trait judgments until one month after they had read the paragraph further implies that these trait summaries were computed spontaneously, and shortly after exposure to the initial paragraphs describing the target persons—while participants still had conscious access to the facts relevant to these judgments.

The inference that trait judgments by the one month participants were based on summaries rather than on-line analysis of consciously accessible episodes receives further support from Experiment 3: The pattern produced by the one month participants was very similar to that produced by D.B., who is densely amnesic. Even though D.B. was 79 years old and had substantial brain damage, he produced trait ratings for Susan Bower that were well correlated (r = .47) with those of 20-year-old undergraduates who could remember a great deal about her and, like the one month participants, these ratings were based on the paragraph D.B. read and not just on her photo—his trait ratings did not correlate at all with those produced by participants who had only seen Susan Bower's photo (r = .07). As a result of his profound episodic memory impairment (detailed in Klein, Rozendal, & Cosmides,

2002) D.B. was unable to recall any of the specific facts we presented about Susan Bower. Yet by the third exposure to this material in a 45 minute period, his mind had extracted trait summaries that were very similar to those produced by people one-quarter his age who could recall a great deal about her. D.B.'s success by 45 mintues implies that trait summaries are computed very soon after exposure to information about a person—within an hour of hearing about her, if not sooner.

Exposure to a single episode was sufficient to trigger the computation of trait summaries for novel characters; there was no need to accumulate a large database of episodes for this to happen. Social decision procedures may preferentially access episodes when the credal value of trait summaries is low due to sparse evidence, but the results we have presented suggest that these trait summaries do, in fact, exist. Immediate computation and updating with each new episode is consistent with previous results (e.g., Klein, Loftus, & Kihlstrom, 1996), showing that people update their own trait summaries as each new episode unfolds. But, because these participants already had 20 years of experiences with their own behavior, those studies could not directly address whether continuous updating occurs only after the prior accumulation of a large database of episodes (some data consistent with this view can be found in Sherman & Klein, 1994).

CANDIDATE DESIGN FEATURES OF THE TRAIT SUMMARY MEMORY SYSTEM FOR PERSONS

If evolution has produced two functionally differentiated systems for person memory, then they should each have separate sets of design features that tailor them for their divergent functions, costs, and payoffs. The present study, with its simultaneous, parallel formation of an episodic memory about a person and a parallel trait summary about that person allows us to explore contrasts. The results reported here, together with other well-established findings, can be used not only to explore the design features of episodic memory, but also to piece together design features of the memory system that produces abstract trait summaries for persons.

We begin by presuming that a primary function of social memory is to use past observations to supply social decision making systems with useable inputs—inputs about stable individual-specific deviations from generic human modal responses (e.g., careful: this person will become angry with less provocation than normal). With this in mind, experiences of and memory for episodes appear to be closer representationally to the raw material out of which trait assessments are manufactured. In contrast, abstract traits appear to be (or be close to) the final representational format used as input into social decision making mechanisms. In consequence, abstract trait summaries may be the single most valuable product of the social memory system, in much the same way that, pound per pound, refined metal is worth more than the ore from which it was extracted. Indeed, the brain appears designed to treat them this way.

The creation of an episodic memory of a new individual, and the simultaneous formation of an abstract trait summary for that same individual allows us to compare their decay functions. This study allows us to see that the abstract trait summary of a previously unknown, personally unimportant, once-described individual is preserved over a month with remarkably little deterioration. In contrast, over the same month, almost all of the episodic information on which the trait

summary was based was lost. The trait summary is, indeed, a complexly detailed data structure, so it is all the more remarkable that it is remembered across large periods of time with so little distortion. The conclusion that the brain appears designed to differentially protect abstract trait summaries is broadly consistent with the neuropsychological data as well. Episodic memory appears to be far more easily lost (producing common amnesias) or disrupted by organic assaults than are abstract trait summaries (for reviews, see Klein, 2001, 2004; Klein, German, Cosmides, & Gabriel, 2004). This does not mean that episodes are worthless or are not remembered—we argue the reverse below. Rather, because the brain is finite, retention rates for various kinds of information are expected to track relative value per neural investment.

The idea that abstract trait summaries are required to feed into social decision making systems is supported by the fact that they are manufactured after only one encounter with an unfamiliar person. If abstract trait summaries were simply a compression/distillation process designed to render large unwieldy data sets tractable—like executive summaries on long reports, or abstracts on scientific articles—then they should only appear when data sets grow large. But from this study, we know that detailed trait summaries are formed immediately and spontaneously, on the basis of a single episode—which is exactly what would happen if decision making required input in this format, and the system perpetually prepares itself for immediate decisions emerging rapidly out of a situation. This finding is also consistent with the finding that a few seconds of visual exposure similarly triggers the formation of a personality representation (for review, see Ambady, Bernieri, & Richeson, 2000).

Not only are abstract trait summaries rapidly formed, but they also appear to be designed to be rapidly accessed—indeed, more rapidly than episodes (e.g., Klein & Loftus, 1993a). Since social decisions have important consequences, and must often be made under urgent time constraints, rapid access qualifies as another apparent design feature. An abstract trait summary is a processed capsule of decision making information precomputed into the form most usable for rapid responses in social contexts. Finally, if improved social decision making is the functional goal of the system, then it should be designed to be accurate and consistent with others' judgments rather than idiosyncratic. These data indicate that different subjects exposed to the same episodic information produce remarkably convergent abstract trait summaries.

In short, the system that evolved to produce abstract trait summaries about persons appears to have the following design features. It produces summaries spontaneously and rapidly; it produces person-representations that are detailed and nuanced; it preserves summaries across large periods of time with remarkably little loss, especially when compared to episodes; it differentially preserves them against organic assaults; it allows the trait summaries to be accessed very rapidly, to make judgments; and it produces summaries that are "objective"—or consensually convergent on others' evaluations.

WHY MAINTAIN EPISODES? THE LIMITATIONS OF TRAIT SUMMARIES

Behavioral episodes provide the data from which trait summaries are inferred, but that does not explain why so many of them are preserved after a summary has

been formed. Indeed, memory research is replete with cases in which the original learning episodes are jettisoned after a generalization has been reached: We remember that apples are edible without being able to recall who told us this fact or the circumstances under which we discovered it; perceptual memory contains a representation of the typical oak tree without our being able to recall all the individual oaks from which the features of the typical one were extracted (for an early treatment, see Tulving, 1972). By contrast, it is possible to recall many details of our social interactions with other people, and our social lives include many conversations in which the details of these behavioral episodes are recounted and analyzed.

Trait summaries are generalizations abstracted from specific episodes. By contrast, episodic memories are situation-specificity incarnate: They are records of how particular people behaved in specific circumstances, and preserve many details of the original situation. Trait summaries may provide fast access to useful generalizations about a person's dispositions to act, but they have an important drawback: they can be wrong.

The results above show that trait summaries may be computed even after exposure to a single episode (for reviews, see Ambady, Bernieri, & Richeson, 2000; Newman & Uleman, 1989). But the facts of this single event may be atypical—they may not provide a good basis for predicting a person's dispositions to act. Even worse, very different trait summaries may be inferred when new facts come to light (e.g., Babey et al., 1998).

To be useful for generating adaptive decisions, trait summaries need to represent a person's dispositions with some accuracy. To remain accurate, however, they must be revised and updated in light of new information. It's not just that people and circumstances change; if that were the only problem, impressions could be updated incrementally without consulting memories of past behavioral episodes: a trait summary based on 20 episodes in which Tom West was friendly could be incrementally weighted by each new episode in which he was unfriendly (e.g., Sherman, 1996; Sherman & Klein, 1994). The problem is that past episodes can take on entirely new meaning when new information comes to light. Susan Bower seems like a pleasant single mother engaged in a promising romantic relationship with Ray; but given the new information that she is happily married to a successful and respected man named Henry, her relationship with Ray is reinterpreted as an illicit and adulterous affair. Tom West seems selfish for quitting his job knowing his wife and children will suffer; but given the new information that his employers dumped dangerous pollutants into a community's water supply and want Tom to cover it up, his quitting is reinterpreted as an act of conscience.

Participants who got this new information within one hour were able to revise the trait summaries they had computed for Susan Bower and Tom West, and reverse their impressions from positive to negative (or vice versa). D.B. was not able to do so, nor were participants who received the new information one month later. The difference is that the participants who received the new information within one hour could still consciously recall many of the facts presented in the original paragraph, whereas D.B. and the participants in the one month condition could not. Without these facts, the new information supports no new inferences about Susan or Tom's character; only when both sets of facts are considered together can the inferences that lead to trait revision can be made.

We cannot know whether some trace of the original episode is preserved in the minds of participants in the one month delay condition (or, for that matter, in D.B.'s mind)—priming studies sometimes show that the memories of amnesics contain information that they are unable to recall (for reviews, see Baddeley, Wilson, & Watts, 1995; Mayes, 1988; Papanicolaou, 2006; Parkin & Leng, 1993; Tulving & Craik, 2000). What we do know is that, without conscious access to these facts, D.B. and our one month participants could not revise their impressions of Susan Bower or Tom West. Maintaining the accuracy of trait summaries depends on maintaining an accessible record of past episodes. Not every past episode need be maintained; there may be an optimal forgetting function that tracks frequency of interaction. For people we interact with daily over long periods, there may be a shorter time depth to episodic retention because many important facts about their lives will become stored as autobiographical facts (or a narrative) in semantic memory, allowing reevaluation. For people we interact with frequently but sporadically, there are long intervals when we don't have any record of what they were doing; to allow reevaluation when so much data is missing, the library of episodes should cover a longer period of time.

MECHANISMS FOR MAINTAINING ACCURACY

On this account, adaptive social decision making requires a database of accurate trait summaries, so natural selection should have designed social inference mechanisms whose adaptive function is to maintain the accuracy of this database. These inference mechanisms should initiate the updating of trait summaries as each new event occurs; in doing so, they should search the database of past episodes for facts whose content is relevant to the new event—Susan Bower's happy marriage to Henry (new information) should activate past episodes bearing on her love life (she is dating Ray); Tom West's being asked to cover up his employer's wrongdoing (new information) should activate past episodes bearing on his work (he quit). Not only will this information be preferentially retrieved, but it will be further processed by social inference mechanisms because, in conjunction with the new information, it supports a reevaluation of Susan's or Tom's character.

The priming results of Experiment 4 support this prediction. All participants read the first paragraph about Susan Bower, and their recall was tested after one hour. Immediately prior to the recall test, some participants were asked to rate Susan's traits (judgment group), some did nothing Susan-relevant (control group), and some were asked to read the second paragraph presenting new information about Susan Bower. In the first 30 seconds of recall, all participants remembered ~2 facts about Susan Bower. However, these facts were disproportionately about Ray for participants who had just learned that she is married to Henry: 54% of participants mentioned Ray in this condition, compared to only 8.3% (control) and 12.5% (judgment) in the other two conditions. This is not because information about Ray is inherently forgetable; by the end of the two minute period, participants were equally likely to have mentioned Ray, regardless of condition. Rather, information that Susan is married to Henry primed past information about her romantic attachments, and this was further processed because it leads to a reevaluation of her character—a fact we know from the previous experiments. As a result, Ray-

facts were more easily accessed when participants were asked to recall information about Susan Bower.

THE INTERRELATED FUNCTIONS OF SEMANTIC AND EPISODIC MEMORY

Semantic knowledge is constructed from information extracted from events or episodes we experience in life, but there are many domains in which the memory of these episodes is lost after the information has been extracted. We have argued that social interaction is one domain in which it is functional to retain a database of episodes that preserve a record of people's behavior, even after semantic summaries of their traits have been extracted. By retrieving trait summaries, social decision making procedures gain fast access to useful generalizations. But speed often comes at the expense of accuracy. Maintaining a database of behavioral episodes helps to minimize speed-accuracy trade-offs that would otherwise destroy the utility of trait summaries (e.g., Klein, Cosmides, Tooby, & Chance, 2001, 2002). But this requires inference procedures and decision rules that are designed to access the right behavioral episodes at the right time. To dissect the architecture of memory into functional units, we need to pay as much attention to the function of these inference processes and decision rules as to the databases they access.

The importance of having a database of trait summaries for each person is underlined by our finding that a single learning episode is sufficient to initiate the process by which they are computed. One month after learning about Susan Bower or Tom West, our participants were able to pull up a detailed and stable profile of the target's character traits.

But trait summaries are useful only to the extent that they faithfully represent a person's dispositions to act, and snap judgments based on single episodes often will be inaccurate. To remain useful, trait summaries need to be continually updated as new evidence comes in. According to the impression reevaluation hypothesis, maintaining a database of behavioral episodes is critical to this updating process. Thematically related content stored in past episodes is preferentially retrieved by the inference mechanisms that compute trait summaries, initiating a reevaluation process that can revise or even reverse the valence of a previously computed trait summary.

In this view, impression reevaluation manages one speed-accuracy trade-off: Its inference procedures use past episodes to maintain the accuracy of the trait generalizations stored in semantic memory. The scope hypothesis (e.g., Babey et al., 1998; Klein, Cosmides, Tooby, & Chance, 2002), however, is about managing a different speed-accuracy trade-off: One that occurs when social decision making rules access trait summaries for the purpose of making judgments and planning actions.

Retrieving a trait summary may be faster than constructing a judgment on-line from a database of episodes, but it can be less accurate: Trait summaries lack information that was present in the behavioral episodes from which they were derived. Episodic memories provide records of how particular people behaved in specific circumstances, but not all of these memories are equally useful for social decision making. Episodes that are consistent with a trait summary provide no information above and beyond that provided by the summary itself, so in making trait judg-

ments, nothing is gained by retrieving consistent episodes in tandem with a trait summary (for reviews, see Klein, 2004; Klein & Loftus, 1993a; Klein, Robertson, et al., 2008). But episodes that are inconsistent with a trait summary can provide boundary conditions on its scope—they provide information about the situations in which the generalization does *not* apply (e.g., Babey et al., 1998; Klein, 2004; Klein, Cosmides, Tooby, & Chance, 2001, 2002). According to the scope hypothesis, a system engineered to retrieve, along with a generalization, episodes that place boundary conditions on its scope would achieve the best of both worlds: speed courtesy of semantic memory, accuracy courtesy of episodic memory.

Klein, Cosmides, Tooby, and Chance (2002) provided evidence for the scope hypothesis through a priming paradigm: Retrieving a trait summary sped recall of inconsistent episodes, but not consistent ones. The data could not be explained by any theory proposing that priming is a functionless byproduct of engineering constraints, such as spreading activation in a neural net, fan effects, or connectionist architectures. Not only did the pattern of episodic priming precisely fit the unusual pattern predicted by the scope hypothesis, but it differed from the (also functional) patterns of priming found in other domains, such as object recognition. Taken together, the data supported the hypothesis that priming is the functional product of evolved adaptations, with different patterns occurring as solutions to different adaptive problems.

We do not claim to have exhausted the functions of maintaining a database of behavioral episodes alongside a semantic database of trait summaries. Episodes may be just as important for determining how much credence to put on other people's claims, for example, as they are in reevaluating their traits or making decisions about how to interact with them (e.g., Klein, Cosmides, Tooby, & Chance, 2002). As the examples provided show, however, progress in dissecting the architecture of memory can be made by analyzing some of the adaptive problems its various systems were designed to solve, and then looking for features that are well engineered for solving them.

REFERENCES

Ambady, N., Bernieri, F. J., & Richeson, J. A. (2000). Toward a histology of social behavior: Judgmental accuracy from thin slices of the behavioral stream. In M. P. Zanna (Ed.), *Advances in experimental social psychology, 32* (pp. 201-272). San Diego, CA: Academic.

Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, *14*, 471-517.

Anderson, N. H. (1968). Likeableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology, 9*, 272-279.

Babey, S. H., Queller, S., & Klein, S. B. (1998). The role of expectancy violating behaviors in the representation of trait knowledge: A summary-plus-exception model of social memory. *Social Cognition*, *16*, 287-339.

Baddeley, A.D., Wilson, B. A., & Watts, F. N. (1995). *Handbook of memory disorders*. New York: Wiley.

Brothers, L. (1997). *Friday's footprint: How society shapes the human mind*. New York: Oxford University Press.

Byrne, R.W. (1995). *The thinking ape: Evolutionary origins of intelligence.* Oxford, England: Oxford University Press.

Byrne, R. W., & Whiten, A. (1988). *Machiavellian intelligence: Social expertise and the evolution of intellect in monkeys, apes, and humans*. New York: Clarendon Press/ Oxford University Press.

Caruso, E. M. (2008). Use of experienced retrieval ease in self and social judgments. *Journal of Experimental Social Psychology, 44,* 148-155.

Claparede, E. (1911). Recognition and "me"ness. In D. Rapaport (Ed.), *Organisation and pathology of thought* (pp. 58-75). New York: Columbia University Press. (Reprinted from *Archives de Psychologies, 11*, 79-90)

Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163-228). New York: Oxford University Press.

Cosmides, L., & Tooby, J. (1997). Dissecting the computational architecture of social inference mechanisms. *Characterizing human psychological adaptations* (pp. 132-161). Hoboken, NJ: Wiley.

Craik, F. I. M., Moroz, T. M., Moscovitch, M., Stuss, D. T., Winocur, G., Tulving, E., & Kapur, S. (1999). In search of the self: A PET investigation of self-referential information. *Psychological Science, 10*, 26-34.

Crowder, R. G. (1976). *Principles of learning and memory.* Oxford, England: Erlbaum.

Damasio, A.R., Tranel, D., & Damasio, H. (1989). Amnesia caused by herpes simplex encephalitis, infarctions in basal forebrain Alzheimer's disease and anoxia/aschemia. In L. Squire & G. Gainottim (Eds.), *Handbook of neuropsychology vol. 3* (pp. 149-166). New York: Elsevier.

Dawkins, R. (1986). *The blind watchmaker*. New York: Norton.

de Waal, F.B.M. (1982). *Chimpanzee politics: Power and sex among apes.* London, England: Jonathon Cape.

Donald, M. (1991). *The origins of the human mind: Three stages in the evolution of culture and cognition.* Cambridge, MA: Harvard University Press.

Dunbar, R. I. M. (1988). *Primate social systems.* London: Croom Helm.

Dunbar, R. I. M. (1996). *Grooming, gossip, and the evolution of language*. Cambridge, MA: Harvard University Press.

Funder, D. C. (1995). On the accuracy of personality judgment: A realistic approach. *Psychological Review, 102*, 652-670.

Funder, D. C., & Sneed, C. D. (1993). Behavioral manifestations of personality: An ecological approach to judgmental accuracy. *Journal of Personality and Social Psychology*, *64*, 479-490.

Gregg, V. H. (1986). *Introduction to human memory.* Boston, MA: Routledge & Kegan Paul.

Haberlandt, K. (1999). *Human memory: Exploration and application.* Needham Heights, MA: Allyn & Bacon.

Hampson, S. E. (1982). *The construction of personality*. Boston, MA: Routledge & Kegan Paul.

Humphrey, N. (1984). *Consciousness regained: Chapters in the development of mind*. New York: Oxford University Press.

Humphrey, N. (1986). *The inner eye*. London: Faber & Faber.

Hunt, R. R., & Ellis, H. C. (1999). *Fundamentals of cognitive psychology* (6th ed.). New York: McGraw-Hill.

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

Kihlstrom, J. E., & Klein, S. B. (1994). The self as a knowledge system. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition, vol. 1, basic processes* (pp. 153-208). Hillsdale, NJ: Erlbaum.

Kihlstrom, J. F., & Klein, S. B. (1997). Self-knowledge and self-awareness. *The self across psychology: Self-recognition, self-awareness, and the self concept* (pp. 5-17*).* New York: New York Academy of Sciences.

Kirby, D. M., & Gardner, R. C. (1972). Ethnic stereotypes: Norms on 208 words typically used in their assessment. *Canadian Journal of Psychology, 26*, 140-154.

Klein, S.B. (1999). Memory and the self. In *McGraw-Hill 1999 yearbook of science and technology* (pp. 227-229). New York: McGraw-Hill.

Klein, S. B. (2001). A self to remember: A cognitive neuropsychological perspective on how self creates memory and memory creates self. In C. Sedikides & M. B. Brewer (Eds.), *Individual self, relational self, and collective self* (pp. 25-46). Philadelphia, PA: Psychology Press.

Klein, S. B. (2004). The cognitive neuroscience of knowing one's self. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (3rd ed., pp. 1077-1089*).* Cambridge, MA: MIT Press.

Klein, S. B. (2007). Phylogeny and evolution: Implications for understanding the nature of a memory system. In H.L. Roediger, Y. Dudai, & S. M. Fitzpatrick (Eds.), *Science of memory: Concepts* (pp. 377-381). Oxford, England: Oxford University Press.

Klein, S. B., Babey, S. H., & Sherman, J. W. (1997). The functional independence of trait and behavioral self-knowledge: Methodological considerations and new empirical findings. *Social Cognition, 15*, 183-203.

Klein, S. B., Cosmides, L., Tooby, J., & Chance, S. (2001). Priming exceptions: A test of the scope hypothesis in naturalistic trait judgments. *Social Cognition*, *19*, 443-468.

Klein, S. B., Cosmides, L., Tooby, J., & Chance, S. (2002). Decisions and the evolution of memory: Multiple systems, multiple functions. *Psychological Review*, *109*, 306-329.

Klein, S. B., Gabriel, R. H., Gangi, C. E., & Robertson, T. E. (2008). Reflections on the self: A case study of a prosopagnosic patient. *Social Cognition, 26*, 766-777.

Klein, S. B., German, T. P., Cosmides, L., & Gabriel, R. (2004). A theory of autobiographical memory: Necessary components and disorders resulting from their loss. *Social Cognition, 22*, 460-490.

Klein, S. B., & Kihlstrom, J. F. (1998). On bridging the gap between social-personality psychology and neuropsychology. *Personality and Social Psychology Review, 2*, 228-242.

Klein, S. B. & Loftus, J. (1990). Rethinking the role of organization in person memory: An independent trace storage model. *Journal of Personality and Social Psychology, 59*, 400-410.

Klein, S. B., & Loftus, J. (1993a). The mental representation of trait and autobiographical knowledge about the self. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition, vol. 5* (pp. 1-50). Hillsdale, NJ: Erlbaum.

Klein, S. B., & Loftus, J. (1993b). Behavioral experience and trait judgments about the self. *Personality and Social Psychology Bulletin, 19*, 740-745.

Klein, S. B., Loftus, J., & Kihlstrom, J. F. (1996). Self-knowledge of an amnesic patient: Toward a neuropsychology of personality and social psychology. *Journal of Experimental Psychology: General*, *125*, 250-260.

Klein, S. B., Loftus, J., & Kihlstrom, J. F. (2002). Memory and temporal experience: The effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *Social Cognition*, *20*, 353-379.

Klein, S. B., Loftus, J., Trafton, J. G., & Fuhrman, R. W. (1992). Use of exemplars and trait judgments: A model of trait knowledge about the self and others. *Journal of Personality and Social Psychology, 63*, 739-753.

Klein, S. B., Robertson, T. E., Gangi, C. E., & Loftus, J. (2008). The functional independence of trait self-knowledge: Commentary on Sakaki (2007). *Memory, 16*, 556-565.

Klein, S. B., Rozendal, K., & Cosmides, L. (2002). A social-cognitive neuroscience analysis of the self. *Social Cognition, 20*, 105-135.

Klein, S. B., Sherman, J. W., & Loftus, J. (1996). The role of episodic and semantic memory in the development of trait self-knowledge. *Social Cognition*, *14*, 277-291.

Lachman, R., Lachman, J.L., & Butterfield, E.C. (1979). *Cognitive psychology and information processing*. Hillsdale, NJ: Erlbaum

Linville, P. W., & Carlston, D. E. (1994). Social cognition of the self. In P. G. Devine, D. L. Hamilton, & T. M. Ostrom (Eds.), *Social cognition: Impact on social psychology* (pp. 143-193). San Diego, CA: Academic.

Mayes, A.R. (1988*). Human organic memory disorders*. New York: Cambridge University Press.

Mischel, W. (1968). *Personality and assessment*. Hoboken, NJ: Wiley.

Mithen, S. (1996). *The prehistory of the mind*. London, England: Thames and Hudson.

Nettle, D. (2006). The evolution of personality variation in humans and other animals. *American Psychologist, 61*, 622-631.

Newman, L., & Uleman, J. S. (1989). Spontaneous trait inferences. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 155-188). New York: Guilford.

Papanicolaou, A. C. (2006). *The amnesias: A clinical textbook of memory disorders*. Oxford, UK: Oxford University Press.

Parkin, A. J., & Leng, N. C. R. (1993). *Neuropsychology of the amnesic syndrome*. East Sussex, UK: Erlbaum.

Passingham, R. (1982). *The human primate*. San Francisco, CA: Freeman.

Roediger, H. L., Dudai, Y., & Fitzpatrick, S. M. (Eds.). (2007). *Science of memory: Concepts.* Oxford, England: Oxford University Press.

Rosenthal, R., & Rosnow, R. L. (1991). *Essentials of behavioral research: Methods and data analysis* (2nd ed.). New York: McGraw-Hill.

Rushton, J. P., Bons, T. A., & Hur, Y. M. (in press). The genetics and evolution of the general factor of personality. *Journal of Research in Personality*.

Schneider, D. J., Hastorf, A. H., & Ellsworth, P. C. (1979). *Person perception (2nd ed.)*. Reading, MA: Addison-Wesley.

Schwarz, N. (1998). Accessible content and accessibility experiences: The interplay of declarative and experiential information in judgment. *Personality and Social Psychology Review, 2*, 87-99.

Schwarz, N. (2004). Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology*, *14*, 332-348.

Sherman, J. W. (1996). Development and mental representation of stereotypes. *Journal of Personality and Social Psychology, 70*, 1126-1141.

Sherman, J. W., & Klein, S. B. (1994). Development and representation of personality impressions. *Journal of Personality and Social Psychology*, *67*, 972-983.

Srull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 440-463.

Takahata, N., Satta, Y., & Klein, J. (1995). Divergence time and population size in the lineage leading to modern humans. *Theoretical Population Biology, 48,* 198-221.

Tomasello, M. & Call, J. (1997). *Primate cognition*. New York, NY: Oxford University Press.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving and W. Donaldson (Eds.), *Organization of memory* (pp. 381-403). New York: Academic.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie Caadienne*, *26*, 1-12.

Tulving, E. (1993a). Self-knowledge of an amnesic individual is represented abstractly. In T. K. Srull & R. S. Wyer (Eds.), *The mental representation of trait and autobiographical knowledge about the self* (pp. 147-156). Hillsdale, NJ: Erlbaum.

Tulving, E. (1993b). Human memory. In P. Anderson, O. Hvalby, O. Paulson, & Hokfelt (Eds.), *Memory concepts-1993: Basic and clinical aspects* (pp. 27-45). London, UK: Elsevier.

Tulving, E., & Craik, F.I.M. (2000). *The Oxford handbook of memory*. New York: Oxford University Press.

Weiskrantz, L. (1997). *Consciousness lost and found: A neuropsychological exploration*. New York: Oxford University Press.

Whiten, A., & Byrne, R. W. (1997). *Machiavellian intelligence II: Extensions and evaluations*. New York: Cambridge University Press.

Williams, G. C. (1966). *Adaptation and natural selection*. Princeton, NJ: Princeton University Press.

Willingham, D. B. (2001). *Cognition: The thinking animal.* Englewood Cliffs, NJ: Prentice-Hall.

Wyer, R. S., & Carlston, D. E. (1994). The cognitive representation of persons and events. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition: Vol. 1. Basic processes* (2nd ed., pp. 41-98). Hillsdale, NJ: Erlbaum.

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