Social Cognition, Vol. 19, No. 4, 2001, pp. 443-468

PRIMING EXCEPTIONS: A TEST OF THE SCOPE HYPOTHESIS IN NATURALISTIC TRAIT JUDGMENTS

Stanley B. Klein, Leda Cosmides, John Tooby, and Sarah Chance
University of California, Santa Barbara

Trait judgments draw on two kinds of memory: (a) trait summaries, which provide information in the form of a generalization, and (b) memories of episodes in which a person behaved in ways that are relevant to the trait. According to the *scope hypothesis* (e.g., Cosmides & Tooby, 2000; Klein, Cosmides, Tooby, & Chance, in press), a trait summary is most useful when its scope is delimited (i.e., when it is accompanied by information specifying those situations in which it does not apply). Episodic memories that are inconsistent with a trait summary can serve this function, because they encode specific situations in which the generalization fails to predict the outcome. This suggests that judgment procedures should be designed to search for summary information in semantic memory and, upon retrieving it, also search for episodic memories that are inconsistent with that summary. This prediction has been tested and supported in previous experiments using artificial target persons (Babey, Queller, & Klein, 1998). Herein, we present the findings from two experiments supporting this prediction using trait judgments about real people for whom subjects have real world knowledge: the self (Experiment 1) and one's mother (Experiment 2). The experiments also test a subtle prediction of the scope hypothesis: that a trait summary must exist and be retrieved for trait-inconsistent episodes to be primed. The results show that in the absence of a trait summary, trait-inconsistent episodes are not primed, but trait-consistent ones are.

Preparation of this article was supported by an Academic Senate Research Grant to Stanley B. Klein from the University of California, Santa Barbara, by grants to John Tooby from the James S. McDonnell Foundation and the National Science Foundation (#BNS9157-449), and by the UCSB Research Across Disciplines Program.

The authors wish to thank Judith Loftus, Don Carlston, and an anonymous reviewer for their extremely helpful comments on an earlier version of this paper. A more theoretical treatment of the issues presented can be found in a companion article "Decisions and the evolution of memory: Multiple systems, multiple functions" by Klein, Cosmides, Tooby, and Chance, *Psychological Review* (in press).

Address correspondence to Stanley B. Klein, Department of Psychology, University of California, Santa Barbara, CA, 93106; E-mail: klein@psych.ucsb.edu.

Social interaction often requires subtle decisions about personality traits. Accordingly, human memory systems store vast amounts of information relevant to judging people's personalities. The mechanisms that render these personality judgments recruit information from two distinct types of memory: an *episodic store* containing representations of specific events and behaviors involving the target, and a semantic store containing representations that summarize the target's personality traits (e.g., Babey, Queller, & Klein, 1998; Hirshman & Lanning, 1999; Kihlstrom & Klein, 1994; Klein, 1999; Klein & Kihlstrom, 1998; Klein & Loftus, 1993a; Klein, Loftus, & Kihlstrom, 1996; Sherman, 1996; Tranel & Damasio, 1993; Tulving, 1993).

Summary representations in the semantic store are derived, at least in part, from those in the episodic store: A trait summary is abstracted from a set of events involving the target, and this happens only when the set of trait-relevant behaviors in the episodic store is large enough (e.g., Klein & Loftus, 1993a, 1993b; Klein, Babey, & Sherman, 1997; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein, Sherman, & Loftus, 1996; Sherman, 1996; Schell, Klein, & Babey, 1996). These summaries are computed automatically — that is, in advance of any particular request for a trait judgment (e.g., Babey et al., 1998; Sherman & Klein, 1994) — but they do not appear to overwrite the database of events over which they were computed. For example, one may have stored a number of generalizations about a friend — usually kind; sometimes flirty; rarely hostile — yet still be able to recall incidents in which she manifested these traits. Indeed, a large literature on the use of memory in trait judgment, including many neuropsychological studies with amnesic individuals, shows that trait summaries and trait-relevant behavioral episodes can be retrieved independently of one another (e.g., Damasio, Tranel, & Damasio, 1989; Johnson, Kim, & Risse, 1985; Klein, 2001; Klein, Chan, & Loftus, 1999; Klein, Loftus, & Kihlstrom, 1996; Tranel & Damasio, 1993; Tulving, 1993). This appears to be true whether the target is another person, the self, or a social group (e.g., Babey et al., 1998; Klein & Loftus, 1993a; Klein, Loftus, & Burton, 1989; Lord, 1993; Sherman, 1996; Sherman, Klein, Laskey, & Wyer, 1998).

Some episodic memories record situations in which the target exemplifed a particular personality trait (e.g., an event in which the target did something that was kind), whereas others record situations in which the target did not manifest that trait (e.g., an event in which the target did something that was unkind). This means that trait judgments could, in principle, be made by retrieving (1) a trait summary, (2) behavioral episodes that are consistent with the trait (*trait-consistent* episodes) or (3) behavioral episodes that are not consistent with the trait (*trait-inconsistent* episodes). The question is, which of these three types of information are used, and why?

To examine the role that trait summaries and trait-relevant behavioral episodes play in trait judgment, Klein and Loftus (1990, 1993a, 1993b; Klein et al., 1989; Klein, Loftus, Trafton, & Fuhrman, 1992) developed a priming paradigm comparing three tasks: (a) a *describe* task asked subjects to judge whether a trait adjective was consistent with their impression of a target person, (b) a *recall* task required subjects to retrieve a memory in which the target person displayed behavior relevant to the trait, and (c) a define task asked subjects to generate a definition of the stimulus trait. For each trait word, subjects perform two of these tasks in succession — an initial task and a target task — and the data of interest are their response latencies when performing a target recall task. It is particularly relevant to see how response times for the recall task vary, depending on whether the initial task was a describe task or a define task. The describe task asks the subject to make a trait judgment, whereas the define task does not (for evidence, see Babey et al., 1998; Klein & Loftus, 1993a, 1993c; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein et al., 1997). If one needs to retrieve behavioral episodes in making a trait judgment, then the time to perform a recall task should be faster when a describe task is performed first than when a define task is performed first. This is because the behavioral information required for a recall task will have been activated during the describe task but not during the define task (see Babey et al., 1998; Klein & Loftus, 1990, 1993a, 1993c; Klein et al., 1997; Sherman & Klein, 1994, for evidence in support of these assumptions). By contrast, if trait judgments can be made without activating behavioral episodes, then performing a describe task first should not reduce the time to perform a recall task any more than would result from performing a define task first.

Most studies using this priming paradigm have explored the role of trait-consistent behavioral memories, and they have revealed an interesting pattern. When asked, "Does trait T describe person P?," a summary representation for trait T and person P is retrieved — if one exists. But if a trait summary is retrieved, trait-*consistent* episodes are not. This is evidenced by the finding that the time it takes to recall a trait-consistent episode is no faster following a describe task than following a define control task. Indeed, when the amount of experience with a target and trait is high (conditions that lead to the formation of a semantic trait summary; e.g., Klein & Loftus, 1990, 1993a, 1993b; Klein et al., 1997; Klein, Loftus, Trafton, & Fuhrman, 1992; Schell et al., 1996) specific memories in which person P manifested behavior consistent with trait summary T appear to play little, if any, role in trait judgments (e.g., Buss & Craik, 1984; Chaiken & Baldwin, 1981; Epstein, 1979; Klein & Loftus, 1993a; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein, Loftus, & Sherman, 1993; Lord, 1993; Schell et al., 1996; but, see Klein, Sherman, & Loftus, 1996, for an important qualification). A very large body of research now exists es-

tablishing this pattern across a wide variety of social targets: other people, the self, and social groups (Babey et al., 1998; Budesheim & Bonnelle, 1998; Craik et al., 1999; Hirshman & Lanning, 1999; Kircher et al., 2000; Klein et al., 1997; Klein et al., 1989; Klein & Loftus, 1993a, 1993b; Klein, Sherman, & Loftus, 1996; Klein, Loftus, Trafton, & Fuhrman, 1992; Lord, 1993; Sherman & Klein, 1994; Sherman, et al., 1998; Tulving, 1993). Indeed, this literature indicates that trait-consistent episodes are retrieved in making trait judgments only when the subject has not previously formed a relevant trait summary for the target person (e.g., Klein & Loftus, 1993a, Klein, Loftus, Trafton & Fuhrman, 1992; Sherman, 1996; Sherman & Klein, 1994; Sherman et al., 1998).

This makes sense. The virtue of a summary representation is that it provides a fast answer to a trait judgment question. Social interaction often requires split second decisions, and time can be saved if one has a system that computes summary representations in advance and stores them for later use (for discussion, see Klein et al., in press). The alternative — retrieving and then evaluating a series of episodes online each and every time a trait judgment is needed — is more costly in both time and computation (e.g., Klein et al., 1997; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman & Klein, 1994). This also explains why retrieving a trait summary fails to prime recall of trait-consistent episodes. Because summaries are precomputed answers to trait judgment questions, there is no additional advantage to retrieving trait-consistent episodes in tandem with them — the information that trait-consistent episodes provide is redundant with the summary (Babey et al., 1998; Klein et al., in press).

The finding that trait summaries can be retrieved independently of trait-consistent episodes has been replicated so many times that one might wonder whether it applies to all episodes, regardless of their content. But in a recent series of experiments on trait judgment, Babey et al. (1998) showed that what appeared to be a general dissociation between retrieval from episodic and semantic memory is not general at all. Their results indicated that when a trait summary is retrieved, trait-*inconsistent* behavioral episodes are retrieved along with it. More specifically, the time it takes to recall a trait-inconsistent episode is faster following a describe task than following a define control task. In other words, asking a subject whether person P is kind co-activates memories of episodes in which that person did something unkind.

This result may be counter-intuitive, but it was predicted. The experiments in Babey et al. (1998) were carried out to test what Klein et al. (in press) and Cosmides and Tooby (2000) have called the *scope hypothesis*. According to this hypothesis, 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. The reasons are as follows.

Trait summaries allow fast access to relevant information. But a trait summary (e.g., "He is rarely honest" or "I am usually friendly") gives information about behavior under "average" circumstances. It does not tell you under what circumstances the person's behavior deviates from average. In deciding how to behave, one is always facing a *particular* situation. Let's say your semantic memory has an entry on Ashley: "Ashley is usually calm." You are planning what you hope will be a relaxed evening with some friends who are political activists of a different stripe than Ashley. Access to appropriate episodic memories can delimit the scope of your semantic summary. Recalling that "Ashley is usually calm — except those times we talked about abortion" may alter your decision about whom to invite. (Indeed, if there is a pattern to the exceptions, a summary of the exceptions might eventually be made as well, and stored as an if-then proposition about the conditions under which Ashley can be expected to become tense; Babey et al., 1998; Wright & Mischel, 1988.) The same reasoning applies to representations of self (e.g., Klein et al., in press; for a related discussion, see Wagenaar, 1994).

In other words, there is a function to maintaining a store of episodic memories, even after a trait summary has been formed: Memories of behavioral episodes *can provide boundary conditions on the scope of generalizations.* This leads to a counter-intuitive prediction about the relation between semantic and episodic memory when people are called upon to make trait judgments. When asked, "Does this describe you: friendly?," the judgment procedures activated should retrieve two kinds of representations: a trait summary from semantic memory and episodes that are *inconsistent* with the trait about which one was asked. For example, "Is Ashley a calm person?" might activate a procedure that retrieves "Usually calm (except: when the topic of abortion is raised)."

In the set of studies reported in Babey et al. (1998) the target persons were created in the lab by providing subjects with information about trait-relevant behaviors performed by the target. In other words, subjects had no prior, real world knowledge about the targets whose personalities they were asked to judge. This type of stimulus material allows a great deal of experimental control over the nature of the behavioral episodes (e.g., their number and their trait relevance). But to assess the scope hypothesis fully, one also needs to test it using naturalistic targets: real people, whom the subject has experienced first hand, and about whom the subject has first hand knowledge. Accordingly, the present studies attempted to extend the generality of the scope hypothesis by examining the representation of pre-experimental, long-term knowledge about real people: one's self and one's mother.

An additional goal of the present studies was to expand the base of stimulus traits used to test the scope hypothesis. In the Babey et al. (1998) study, the evidence in support of the scope hypothesis was based on

priming data collected over a small set of stimulus traits: friendly, unfriendly, intelligent and unintelligent. In the present studies, by contrast, subjects rendered judgments about 96 different traits.

In the experiments reported below, we used the same priming method as in the previously described experiments by Klein, Loftus, and colleagues (e.g., Babey et. al., 1998; Klein & Loftus, 1993a; Klein et al., 1989). Based on the scope hypothesis, we predict that retrieving a trait summary to answer a describe question will prime trait-inconsistent episodes, but not trait-consistent ones. In our first experiment, we probed knowledge about one's own traits, because we can be fairly certain that the typical adult has already formed a database of trait summaries about him- or herself (e.g., Klein & Loftus, 1993a; Klein, Sherman, & Loftus, 1996). In our second experiment, we probed knowledge about another person's traits — the subject's mother. Previous research has revealed conditions under which trait knowledge of others (e.g., one's mother) is unlikely to be represented in summary form (e.g., Klein & Loftus, 1993a; Klein, Loftus, Trafton, & Fuhrman, 1992). By examining these conditions, we were able to see whether the same results pertain when trait summaries can be shown to be lacking.

EXPERIMENT 1: SELF AS TARGET

The goal of our first study was to determine whether accessing a trait summary about the self would prime trait-inconsistent episodes, but not trait-consistent ones. We used a variant of Klein and Loftus's priming technique (e.g., Klein & Loftus, 1993a; Klein et al., 1989; Klein, Loftus, Trafton, & Fuhrman, 1992). We presented subjects with a list of trait adjectives and asked them to perform one of three tasks with them — a describe task, a recall task, or a define task. Each trial consisted of performing two of these tasks —an initial task and a target task — in succession on either the same trait word or on a trait and its antonym (e.g., rude and polite).

After subjects completed the experimental trials, they again were presented with the trait words appearing in the initial task and asked to indicate on a 9-point scale the extent to which each trait described them. These ratings were used to sort subjects' response latencies into three levels of trait self-descriptiveness (high, medium, and low).

PREDICTIONS

If deciding whether a trait is self-descriptive activates a summary representation and episodes *inconsistent* with the trait being asked about, then two things should occur. First, for trials on which the initial and target task trait words are antonyms, an initial describe task should be more fa-

cilitating than an initial define task to the subsequent performance of a recall task. Deciding whether a trait describes oneself activates a trait summary, whereas defining a trait does not. When a trait summary is activated, trait-inconsistent behavioral memories should be activated as well (Babey et al., 1998). These are required for performance of the subsequent recall task. Recall should be faster after a describe task, then, because memories can be retrieved faster if they were recently activated (e.g., Collins & Quillian, 1970; Klein & Loftus, 1990, 1993a; Klein et al., 1989; Malt, 1989).

By contrast, for trials on which the initial and target task trait words are the same, an initial describe task should be no more facilitating than an initial define task to the subsequent performance of a recall task. This is because trait-*consistent* episodes are not activated by retrieval of a trait summary, and thus not activated by a describe judgment (e.g., Babey et al., 1998; Kihlstrom & Klein, 1994; Klein & Loftus, 1990, 1993a, 1993b; Klein et al., 1997; Klein et al., 1989; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein, Loftus, & Sherman, 1993).

Finally, work by Klein and Loftus and their colleagues (e.g., Klein & Loftus, 1993a; Klein, Sherman, & Loftus, 1996; Klein et al., 1997; Klein, Loftus, Trafton, & Fuhrman, 1992; see also Sande, 1990; Sande, Goethals, & Radloff, 1988; Smith, 1990) has shown that trait self-knowledge is represented in summary form across levels of trait descriptiveness. We therefore predicted that the above effects should occur regardless of whether the trait being judged fell in the high, medium, or low self-descriptiveness category.[1]

METHOD

Subjects. Eighteen 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.

---

1. Because decisions regarding personality traits are not amenable to all-or-none answers (if they are to be made accurately), it would be reasonable to assume that the trait summary takes a quasi-quantitative form. Thus, when asked, "Does this describe you: friendly?," possible answers might be "usually," "sometimes," or "rarely." Boundary conditions would be episodes in which you were not friendly. This should be true for any level your "friendly" summary specifies — even "*rarely*". This is because the decision context calls for a generalization about friendliness; the function of activating episodes is to place boundary conditions on *that* generalization. The issue is "How often is person Y friendly, and under what conditions is s/he not?"

Materials and Design. The stimulus words were 96 trait adjectives chosen from the norms provided by Kirby and Gardner (1972) and Anderson (1968). For each trait adjective, we generated an antonym using *Webster's New World Dictionary (1970)* and L. Urdang's *The Basic Book of Synonyms and Antonyms* (1985). Subjects received 96 trials. A trial consisted of performing an initial task and a target task in succession on either the same trait or on a trait and its antonym.

1. For the describe task, subjects were asked to judge whether a stimulus trait was consistent with their impression of themselves (e.g., "Does the word kind describe you?").
2. For the define task, subjects were asked to generate a definition for the stimulus trait (e.g., "Think of the meaning of the word kind"). This task required the subject to access semantic information regarding the trait, without requiring that a trait judgment be made about any individual (e.g., Klein et al., 1997). It thus served as a control for the trait judgment task (i.e., the describe task).
3. The recall task asked subjects to retrieve from memory a specific incident in which they behaved in a way that was relevant to the stimulus trait. Depending on the experimental condition (see below), subjects were required to recall either a trait-*consistent* episode (e.g., "Recall a specific incident in which you behaved in a kind manner") or a trait-*inconsistent* episode (e.g., "Recall a specific incident in which you behaved in an unkind manner").[2]

Three initial tasks (describe, recall, and define) were factorially combined with two target tasks (describe and recall) and two initial task/target task trait pairings (same trait or trait and antonym) to create twelve experimental conditions.[3] On half the trials, within each experimental condition, the trait paired with the initial task was high in normative social desirability and on half the trials it was low in normative social desir-

2. We did not request that subjects report their responses during the experimental trials; rather, we instructed them to generate responses to the task questions in their heads. Klein and Loftus (1993a, 1993c) provide a detailed discussion of our reasons for adopting this procedure and present research demonstrating the efficacy of the technique.

3. Even though our hypothesis requires examination of only conditions in which initial describe and define tasks are followed by a recall target task, we also included conditions in which the recall task served as an initial task and conditions in which the describe task served as a target task. We hoped this would discourage subjects from developing expectancies for a particular initial task/target task pairing by making it more difficult to anticipate the initial or target task on any trial.

ability (social desirability was determined using norms provided by Kirby and Gardner [1972], and N. H. Anderson [1968]). Assignment of trait adjectives to initial task/target task pairs and the order in which the task pairs were presented were randomized across subjects.

Procedure. Subjects were told that we were investigating their ability to perform different tasks on trait adjectives. We told them that it was important that they perform the tasks accurately and that they should indicate immediately when they had completed each task. We then explained the experimental tasks and gave instructions for performing them.

A microcomputer presented the stimulus words and recorded response latencies for the initial and target tasks. Each trial began with the appearance on a computer screen of a cue for the initial task. The cue was either Describe, Recall, or Define. After one second, a trait adjective appeared below the cue. Both the cue and trait adjective remained on the screen until the subject indicated that he or she had completed the initial task by pressing a key. The initial task cue and stimulus trait were then removed from the screen. After a one-second pause, the cue for the target task (Describe or Recall) appeared on the screen above a stimulus trait (either the same one or its antonym) and a timer was activated. Again the cue and the trait adjective remained on the screen until the subject indicated that he or she had completed the target task. The timer then stopped and the target task response latency was recorded. There was a two-second delay before the beginning of the next trial.

In our instructions, we informed subjects that the ordering of the tasks would be random. We also informed them that on some trials the initial and target task would be paired with the same trait word, while on other trials they would be paired with a trait and its antonym. If on any trial a subject was unable to perform a task within 40 s, the task was terminated, the subject was told to proceed to the next task, and the data from that trial was discounted; less than 1% of the response latencies were discounted on this basis. Subjects received twelve practice trials, one for each possible experimental condition.

After subjects had completed the experimental trials, we again presented them with each trait adjective appearing in the initial task and asked them to rate it on a 9-point scale ranging from 1 = "extremely unlike me" to 9 = "extremely like me." These ratings allowed us to sort subjects' response latencies into three degrees of self-descriptiveness. For each of the twelve experimental conditions (eight traits per condition), the two traits receiving the highest ratings were placed in the high-descriptiveness category, the two traits receiving the lowest ratings were placed in the low-descriptiveness category, and the two traits receiving ratings falling nearest the middle of the scale were placed in the me-

dium-descriptiveness category. In the case of ties in which more than two traits could be assigned to a category (e.g., three traits receiving a rating of 9), random assignment was used.

RESULTS

Trait-Descriptiveness Manipulation Check. To ensure that subjects' variability in the use of the full range of the rating scale did not invalidate our assignment of traits to the three categories, we computed the mean ratings for the high, medium, and low trait-descriptiveness categories. These means did indeed reflect scale values consistent with the category headings (*M*s = 7.80, 5.32, and 2.17 for the high, medium, and low categories, respectively). A one-way (ANOVA) on these means was significant, $F(2, 34) = 542.29$, $p < .001$, with Newman-Keuls testing ($p < .01$) indicating that all three means differed reliably.

Recall Target Task Response Latencies. The data of primary interest, the joint effects of initial task, initial task trait descriptiveness, and trait pairings on recall target task mean response latencies are shown in Table 1. A $2 \times 3 \times 2$ (Initial Task [describe versus define] × Initial Task Trait Descriptiveness [high, medium, and low] × Trait Pairing [same versus antonyms]) repeated measures ANOVA on these latencies yielded an interaction between Initial Task and Trait Pairing, $F(1, 17) = 9.90$, $p < .01$. No other main effects or interactions reached significance.

A. *Same-Trait Pairings (Latencies for Consistent Episodes).* The results for same-trait pairings are shown in the top panel of Table 1. Replicating findings from previous studies on trait-consistent episodes (e.g., Klein et al., 1997; Klein & Loftus, 1990, 1993a; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein et al., 1989), simple effects tests showed that regardless of whether traits fell in the high, medium, or low descriptiveness categories, the time taken to perform a recall task was not differentially influenced by the previous performance of a describe task ($M = 6{,}298$ ms) or a define task ($M = 5{,}968$ ms), $F < 1.0$.

B. *Trait-Antonym Pairings (Latencies for Inconsistent Episodes).* The picture changes considerably, however, for trait-antonym pairings (seen in the bottom panel of Table 1). Here, we find strong evidence that trait judgments activate trait-inconsistent episodes: Simple effects tests revealed that regardless of level of trait-descriptiveness, subjects were faster to perform the recall task when the initial task was describe ($M = 5{,}258$ ms) than when it was define, ($M = 6{,}775$ ms), $F(1, 17) = 13.35$, $p < .005$.

TABLE 1. Self study, Experiment 1. Mean Recall Target-Task Response Latency as a Function of Initial Task, Initial Task Trait Descriptiveness, and Trait Pairing

		Initial task trait descriptiveness		
	Initial task	**High**	**Medium**	**Low**
Same trait pairings	Define	5,337	5,567	7,000
(trait-consistent episodes)	Describe	5,427	6,144	7,323
	Difference	-90	-577	-323
Trait antonym pairings	Define	7,101	6,705	6,520
(trait-inconsistent episodes)	Describe	5,107	5,302	5,363
	Difference	1,994	1,403	1,157

Note. Response latencies in milliseconds. High, medium, and low refer to the descriptiveness of the trait presented during the initial task, *not* the recall task. Thus subjects are recalling episodes consistent with their self-view in the low column of the trait antonym condition and the high column of the same trait condition. Similarly, subjects are recalling episodes inconsistent with their self-view in the high column of the trait antonym condition and the low column of the same trait condition (e.g., a friendly person would be recalling a friendly behavior in the high column for a same trait pairing, and in the low column for a trait antonym pairing.)

The above analyses compared recall latencies following a describe task to those following a define task. But an additional perspective on whether making trait judgments activates trait-inconsistent episodes in tandem with summaries can be obtained by seeing whether recall latencies following the describe task are shorter for inconsistent episodes than for consistent ones. According to the scope hypothesis, trait self-descriptiveness judgments should activate trait-inconsistent, but not trait-consistent, behavioral episodes. This leads to the counterintuitive prediction that subjects should recall a behavior more quickly when its trait implications are opposite those of the trait judged during an initial describe task than when they are the same as those of the trait judged initially. Consistent with this prediction, simple effects tests revealed that subjects took significantly less time to recall a behavior when it exemplified the antonym of the trait judged during an initial describe task (M = 5,258 ms) than when it exemplified the same trait as that judged during an initial describe task (M = 6,298 ms), $F(1, 17) = 6.28$, $p < .05$.

ADDITIONAL ANALYSES

Trait Social Desirability. The finding that trait self-descriptiveness judgments facilitated retrieval of trait-inconsistent behaviors, regardless of whether the trait judged fell in the high, medium, or low descriptiveness categories, suggests that retrieval latencies were unaffected by

the social desirability of the trait in question. Nonetheless, to obtain a clearer picture of potential effects of social desirability on retrieval, we reanalyzed the recall target task data, this time segregating latencies on the basis of normative social desirability. For each experimental condition, we placed the four adjectives rated high in normative social desirability in the high-desirability category and the four rated low in social desirability in the low-desirability category (social desirability ratings were determined by the Kirby and Gardner [1972], and Anderson [1968] norms).

A 2 × 2 × 2 (Initial Task: describe and define × Initial Task Trait Social Desirability: high and low × Trait Pairing: same and antonyms) revealed no effect of social desirability on recall task latencies (all Fs < 1.0). The only effect to reach significance was an interaction between Initial Task and Trait Pairing, $F(1, 17) = 6.84$, $p < .05$. Replicating the pattern of findings from the main analysis, subjects were faster to recall a trait-inconsistent behavior when the initial task was describe than when it was define, $F(1, 17) = 4.81$, $p < .05$. By contrast, the time taken to recall a trait-consistent behavior was not differentially influenced by the previous performance of a describe task or a define task, $F(1, 17) = .36$, $p > .50$. It thus appears that normative social desirability was not responsible for the pattern of facilitation obtained in Experiment 1.

Replicating Klein and Loftus. Although our analyses of the same-trait pairing data focus on the predictions of the scope hypothesis, it is worth noting that the pattern of latencies we obtained is consistent with that reported by Klein and Loftus in other ways as well. For example, Klein, Loftus and colleagues repeatedly have shown that participants are quickest to retrieve behaviors exemplifying high-descriptive traits and slowest to retrieve behaviors exemplifying low-descriptive traits (e.g., Klein & Loftus, 1993a; Klein et al., 1997; Klein et al., 1989; Klein, Loftus, Trafton, & Fuhrman, 1992). Replicating these findings, an analysis of recall target task latencies as a function of trait self-descriptiveness revealed a monotonically increasing function; retrieval latencies were shortest for behaviors exemplifying high-descriptive traits (M = 5,382 ms), and longest for behaviors exemplifying low-descriptive traits (M = 7,161 ms), with latencies for behaviors exemplifying medium-descriptive traits falling in between (M = 5,856 ms). A test for linear trend was significant, $F(1, 17) = 12.20$, $p < .01$.[4]

4. Our interpretation of the results from Experiment 1 assumes that the information retrieved during an initial task will facilitate retrieval of related information on a subsequent target task. Although there is much support for this assumption (e.g., Babey et al., 1998; Collins & Quillian, 1970; Klein & Loftus, 1993a; Klein et al., 1997; Loftus & Loftus, 1974; Malt, 1989; Schell et al., 1996; Sherman & Klein, 1994), there also is evidence suggesting

DISCUSSION

In Experiment 1, we found that deciding whether a trait describes oneself facilitates retrieval of trait-inconsistent episodes, which suggests these judgments activate memories of behaviors inconsistent with the trait being judged. By contrast, we found no evidence that trait self-descriptiveness judgments facilitated retrieval of trait-consistent episodes. From this we infer that decisions about whether a trait describes oneself do not activate trait-consistent behavioral memories. That trait-inconsistent (but not trait-consistent) episodic memories are primed by retrieval of a trait summary from semantic memory is just what one would expect if the scope hypothesis were correct.

EXPERIMENT 2: MOTHER AS TARGET

Experiment 2 had two purposes. First, we wanted to extend the generality of the scope hypothesis by testing whether judgments about a well-known other also conform to its predictions.

Second, we wanted to address one interesting result from our first study — the finding that activation of trait-inconsistent behaviors occurred regardless of whether the trait judged fell in the high, medium, or low self-descriptiveness categories. This finding can be accommodated

---

that, under certain conditions, an initial task actually may inhibit a subsequent retrieval (e.g., Blaxton & Neely, 1983; Dagenbach & Carr, 1994; Roediger & Neely, 1982). This leaves open the possibility that the priming found in the trait-antonym pairing condition might have been due to the inhibitory effects of the initial define task rather than the facilitating effects of the initial describe task.

To this end, we conducted a second study comparing recall task latencies following an initial describe task to those obtained with no initial task at all. If the apparent facilitation of trait antonyms found in Experiment 1 was merely a byproduct of the inhibitory effects of an initial define task, then this facilitation should disappear when there is no initial task to inhibit recall: Subjects should be no faster to recall an inconsistent behavioral episode following a describe task than following no task. In contrast, if our original interpretation was correct, then subjects should be faster at recalling inconsistent episodes when a describe task is performed first than when there is no initial task — just as they were in Experiment 1. This is because a summary is being retrieved in the describe condition, but not in the no initial task condition and, according to the scope hypothesis, inconsistent episodes are retrieved to bound the scope of a summary.

A $2 \times 3 \times 2$ (Initial Task [describe and no initial task] $\times$ Initial Task Trait Descriptiveness [high, medium, and low] $\times$ Trait Pairing [same and antonyms]) repeated measures ANOVA conducted on the recall target task mean response latencies revealed an interaction between Initial Task and Trait Pairing, $F(1, 13) = 5.30$, $p < .05$. Replicating the pattern of findings obtained in Study 1, simple effects tests revealed that for same-trait pairings, subjects were no faster to perform the recall target task when it was preceded by a describe task (M = 6,415 ms) than when it was preceded by no initial task (M = 6,357 ms), $F < 1.0$.

within the scope hypothesis by assuming that trait knowledge about the self is represented in summary form across levels of trait self-descriptiveness. Although there is a body of evidence consistent with this assumption (e.g., Klein et al., 1997; Klein & Loftus, 1993a; Klein, Loftus, Trafton, & Fuhrman, 1992; Sande, 1990; Sande et al., 1988), we nonetheless felt it appropriate to conduct a more stringent test by examining recall task latencies under circumstances where there are strong a priori reasons for assuming the presence or absence of trait summaries. After all, the scope hypothesis asserts that the function of priming inconsistent episodes is to place boundary conditions on the scope of a semantic summary. If no trait summary exists, then there is no generalization whose scope needs to be delimited. As discussed in the next section, trait judgments about a well-known other provide exactly the conditions we need to test this prediction.

THE MENTAL REPRESENTATION OF TRAIT KNOWLEDGE ABOUT OTHER PERSONS

In a recent series of papers, Klein and Loftus have proposed a model of how trait knowledge about others is represented in memory (e.g., Babey et al., 1998; Klein & Loftus, 1993a; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman, 1996; Sherman & Klein, 1994). According to the model, one's representation of a person's traits varies with the amount of experience one has had with that person. If the amount of experience is not sufficient to support abstraction, then trait knowledge will be represented only at the level of behavioral memories. Trait judgments about the person, therefore, must be based on behavioral memories (i.e., on episodes). However, as the amount of experience becomes sufficiently large, trait knowledge is increasingly likely to be abstracted and represented in summary form. When this happens, trait judgments may be made by directly accessing the appropriate semantic memory representation (i.e., a trait summary). Thus, the model proposes that the more

---

Also replicating the results of Experiment 1 — and in contrast to the predictions of the inhibition hypothesis — subjects in the trait-antonym condition required less time to perform the recall target task when it was preceded by a describe task (M = 5,047 ms) than when it was preceded by *no* initial task (M = 6,397 ms), $F(1, 13) = 9.74$, $p < .05$. As in Experiment 1, simple effects tests revealed that subjects required significantly less time to recall a behavior when its trait implications were opposite those of the trait judged during an initial describe task (M = 5,047 ms) than when they were the same as those of the trait judged initially (M = 6,415 ms), $F(1, 13) = 10.00$, $p < .05$.

Taken together, these findings are consistent with the argument that the differential pattern of facilitation found in Experiment 1 was due to the facilitating effects of an initial describe task rather than the inhibitory effects of an initial define task (see also, Babey et al., 1998; Klein & Loftus, 1993a; Klein et al., 1997; Klein, Loftus, & Plog, 1992).

knowledge one has about a person's behavior, the more likely one is to have formed summary representations and the less likely one is to base judgments of that person on memories of specific behavioral episodes (for a similar view, see N. H. Anderson, 1989; Park, 1986).[5]

A series of studies by Klein and Loftus and their colleagues has provided evidence largely consistent with their model (Babey et al., 1998; Klein & Loftus, 1993a; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman & Klein, 1994). In one study (Klein, Loftus, Trafton, & Fuhrman, 1992; Experiment 1), in which the target person was the subject's mother, subjects first completed a series of priming trials and then rated each stimulus trait for the degree to which it described their mothers. Klein, Loftus, Trafton, and Fuhrman (1992) found no evidence that a describe task facilitated retrieval of trait-consistent behavioral episodes when the trait in question was highly descriptive of one's mother. However, they found considerable evidence of facilitation when traits were rated medium in mother-descriptiveness. To explain these findings, Klein, Loftus, Trafton, and Fuhrman (1992; see also Klein & Loftus, 1993a) proposed that traits rated highly descriptive of another person are those that he or she has manifested most often, and hence are those for which subjects have observed a relatively large number of behaviors. Accordingly, highly descriptive traits are likely to be those for which subjects have created a trait summary, which can be accessed to perform the describe task. Because memories of trait-consistent behavioral episodes would not be activated, the describe task would be no more beneficial than a define task to the subsequent performance of a recall task.

In contrast, medium-descriptive traits are likely to be those for which subjects have observed fewer behavioral exemplars. This means it is less likely that subjects will have a trait summary. In the absence of a trait summary, subjects would have to retrieve trait-relevant behavioral memories to make descriptiveness judgments. Consequently, a describe task would be more beneficial than a define task in performing a subse-

5. Given the literature on spontaneous trait inferences (e.g., Uleman, 1989; Winter, Uleman, & Cunniff, 1985; but see Carlston & Skowronski, 1986), it might seem that summary trait representations of a target should be formed following exposure to a single trait-relevant behavioral experience. However, there are important differences between the experimental paradigms employed to test the spontaneous trait inference hypothesis and the conditions that characterize the present studies. In most studies of spontaneous trait inference, the behaviors attributed to a target are unambiguously diagnostic of a particular trait. By contrast, in more naturalistic settings (i.e., those in which the subjects in our studies acquired knowledge of the target's behavior), behaviors often are amenable to multiple trait-relevant interpretations. As Sherman & Klein (1994) have shown, under conditions of attributional uncertainty, subjects typically require an accumulation of trait-relevant behavioral evidence before they form a summary representation of a target's traits.

quent recall task, and one would observe faster retrieval of trait-consistent episodes following a describe task.

Thus, consistent with the predictions of the Klein, Loftus, Trafton, and Fuhrman (1992) model, for judgments about mother, trait-consistent behavioral episodes appear to be required when the trait being judged is medium in mother-descriptiveness, but not when the trait being judged is highly descriptive of mother. (No clear predictions can be made for low-descriptive traits; for reasons, see Klein & Loftus, 1993a; Klein, Loftus, Trafton, & Fuhrman, 1992).[6]

THE MOTHER STUDY

In this experiment, we tested whether trait-inconsistent episodes are activated when making trait judgments about one's mother and, if so, under what conditions. We repeated the procedure described in Experiment 1 with one change: This time a subject's mother, rather than the self, served as the target person.

An additional goal of this experiment was to test the predictions of the scope hypothesis under conditions where we could be reasonably certain of the presence or absence of trait summaries. According to the research presented by Klein, Loftus, Trafton, and Fuhrman (1992; see also Klein & Loftus, 1993a), these requirements can be satisfied by examining judgments about traits rated high or medium in mother-descriptiveness, but not traits rated low in mother-descriptiveness. We therefore limited our predictions to the effects of judgments about high- and medium-descriptive traits on recall task latencies.

PREDICTIONS

1. *For High-Descriptive traits*:

   A. Recall task latencies in the same-trait pairing condition should be the same regardless of whether the initial task was describe or define. This is because subjects are likely to have summary represen-

---

6. Note that in Experiment 1, we found no evidence that trait-consistent episodes were involved in judgments about the self, regardless of the level of trait descriptiveness. The reason for this apparent exception of the self from the predictions of the model is that, compared with experience with others, experience with oneself is vast (e.g., Baxter & Goldberg, 1987; Kihlstrom & Klein, 1994; Klein & Loftus, 1993a; Sande, 1990; Sande et al., 1988; Smith, 1990). Accordingly, there are few, if any, trait dimensions for which we have not had sufficient experience to form abstract trait knowledge about ourselves (for reviews, see Sande et al., 1988; Kihlstrom & Klein, 1994; Klein & Loftus, 1993a). No matter what the trait is, a person will have a summary of how well it describes himself or herself.

tations of these traits, which they access to make trait judgments in the initial describe task. As before, the scope hypothesis predicts that trait-consistent behavioral memories should not be activated by retrieval of a trait summary in the describe condition.

B. In contrast, recall task latencies in the trait-antonym pairing condition should be shorter when the initial task is describe than when it is define. By the scope hypothesis, when a trait summary is activated in the course of making a trait judgment, memories of trait-inconsistent behavioral episodes are also activated to place boundary conditions on the summary's scope.

2. *For Medium-Descriptive traits*:

A. Recall task latencies in the same-trait pairing condition should be shorter when the initial task is describe than when it is define. This is because subjects are less likely to have summary representations for medium-descriptive traits, and thus are more likely to retrieve episodic memories of trait-consistent behaviors to make judgments. Our reason for assuming that retrieval will be biased toward trait-consistent behaviors follows from research showing that people rely on "positive-test" strategies (e.g., Klayman & Ha, 1987; Wason, 1968) to answer questions about the trait characteristics of others: They are more likely to search memory for examples of the trait being judged than for examples of its opposite (e.g., Devine, Hirt, & Gehrke, 1990; Fong & Markus, 1982; Snyder & Cantor, 1979).

B. In contrast, latencies in the trait-antonym pairing condition should be the same regardless of whether the initial task was describe or define. When trait judgments are not based on access to a trait summary, they are unlikely to activate episodic memories of trait-inconsistent behaviors. If no summary has been activated, there is no need to delimit its scope.

METHOD

Subjects. Eighteen undergraduates from the University of California, Santa Barbara participated for course credit. They were tested individually in sessions lasting approximately one hour.

Materials, Design, and Procedure. The stimulus words were the same as in Experiment 1. The design and procedure also were identical to that of Experiment 1, with two changes. First, the referent for the describe and recall task was changed from self to mother. Thus, describe task instructions became "decide whether the presented adjective describes your mother" and recall task instructions became "recall a specific incident in which your mother's behavior exemplified the presented trait." The de-

fine task was unchanged. Second, the scale for the trait descriptiveness ratings was changed to reflect the change in referent: The scale now ranged from 1 = "extremely unlike mother" to 9 = "extremely like mother."

RESULTS

If a subject could not perform a task within 40 s, the data from that trial was removed from the data set. This resulted in removal of 1.7% of the response latencies.

Trait-Descriptiveness Manipulation Check. A one-way ANOVA on the mean ratings for the high ($M = 8.12$), medium ($M = 5.30$), and low ($M = 1.72$) trait-descriptiveness categories was significant, $F(2, 34) = 642.32$, $p < .001$, with Newman-Keuls testing ($p < .01$) revealing reliable differences between all three means.

Recall Target Task Response Latencies. A $2 \times 3 \times 2$ (Initial Task [describe versus define] × Initial Task Trait Descriptiveness [high, medium, and low] × Trait Pairing [same versus antonyms]) repeated measures ANOVA on recall target task latencies yielded a significant three-way interaction, $F(2, 34) = 3.94$, $p < .05$.

1. *High-Descriptive Traits*.

   A. *Same-trait Pairings (Latencies for Consistent Episodes).* The latencies for same-trait pairings (trait-consistent episodes) are shown in the top panel of Table 2. Simple effects tests revealed that for traits falling in the high category, the time to perform a recall task was not differentially influenced by the previous performance of a describe task or a define task ($F < 1.0$).

   B. *Trait-antonym Pairings (Latencies for Inconsistent Episodes)*. Latencies for the trait-antonym pairings (trait-inconsistent episodes) are shown in the bottom panel of Table 2. For traits falling in the high category, there was considerable facilitation: Subjects were much faster to remember their mother's trait-inconsistent behaviors when they first performed a describe task than when they first performed a define task: Latencies were shorter by 1,956 ms when the initial task was describe than when it was define, $F(1, 17) = 5.01$, $p < .05$. In other words, mother's highly descriptive traits elicit results parallel to those for the self: Inconsistent episodes are primed, but consistent ones are not. This is what one would expect if subjects had formed a summary representation for those traits that are highly characteristic of their mothers, and if the system is designed to re-

TABLE 2. Mother study, Experiment 2. Mean Recall Target-Task Response Latency as a Function of Initial Task, Initial Task Trait Descriptiveness, and Trait Pairing

		Initial task trait descriptiveness	
	Initial task	**High**	**Medium**
Same trait pairings	Define	5,844	8,521
(trait-consistent episodes)	Describe	6,116	6,228
	Difference	-272	2,293
Trait antonym pairings	Define	10,081	7,253
(trait-inconsistent episodes)	Describe	8,125	8,174
	Difference	1,956	-921

Note. Response latencies in milliseconds.

trieve inconsistent episodes in tandem with these summaries to delimit their scope.

2. *Medium-Descriptive Traits*. The effects of initial task on latencies was quite different for medium- than for high-descriptive traits.

A. *Same-trait Pairings (Latencies for Consistent Episodes)*. Unlike the high-descriptive traits, mother's medium-descriptive traits did facilitate recall of trait-consistent episodes. For traits falling in the medium category, subjects performed a recall task considerably faster when it was preceded by a describe task than when it was preceded by a define task: Latencies were shorter by 2,293 ms when the initial task was describe than when it was define, $F(1, 17) = 4.77$, $p < .05$. This is what one would expect if subjects had not formed summary representations for their mothers' medium-descriptive traits. This pattern replicates that reported by Klein, Loftus, Trafton, and Fuhrman (1992; see also, Klein & Loftus, 1993a).

B. *Trait-antonym Pairings (Latencies for Inconsistent Episodes)*. Whereas high-descriptive traits did facilitate recall of trait-inconsistent episodes, medium-descriptive traits did not. Simple effects tests on these latencies revealed no differential facilitation for traits falling in the medium category ($F < 1.0$). This is what one would expect if the scope hypothesis were correct: Subjects have not formed trait summaries for mother's medium-descriptive traits, so there is no generalization whose scope needs to be delimited by retrieving episodes inconsistent with it.

DISCUSSION

For trait-consistent behavioral episodes, the data from judgments about high- and medium-descriptive traits in the same trait-pairing condition conform nicely to our predictions. We found no evidence of facilitation for judgments about traits rated "high" in mother descriptiveness, from which we infer that these judgments were accomplished by accessing summary trait knowledge about mother from semantic memory. Activating a trait summary does not require the simultaneous activation of consistent episodes, according to the scope hypothesis (e.g., Klein et al., in press).

For judgments about traits rated "medium" in mother-descriptiveness, we found considerable facilitation of trait-consistent episodes, which suggests that subjects do not have summary representations of these traits, and therefore require episodic memories to make a judgment. In other words, the mental representation of mother's traits varies with trait-descriptiveness: Knowledge of highly descriptive traits consists of summary representations, whereas knowledge of medium-descriptive traits consists primarily of memories of specific behavioral episodes.

These findings provide a backdrop against which to evaluate the predicted effects of trait-descriptiveness on recall task latencies in the trait-antonym pairing condition. According to the scope hypothesis, when a summary is activated, trait-inconsistent episodes are also activated to provide boundary conditions on its scope. If traits rated "high" in mother-descriptiveness are represented in summary form, performance of an initial describe task should facilitate the recall of trait-inconsistent behaviors. In contrast, if traits rated "medium" in mother-descriptiveness are not represented in summary form, as the data on trait-consistent facilitation suggests, an initial describe task should not facilitate the recall of trait-inconsistent behaviors: When there is no summary, there is no need to delimit its scope. Our data confirm both predictions. Trait-inconsistent episodes were primed only for highly descriptive traits. This provides further evidence for the scope hypothesis.

GENERAL DISCUSSION

We have argued that generalizations from semantic memory allow speedy decisions, but at the cost of accuracy, whereas episodic memories provide accurate — that is, situationally specific — information, but at the cost of speed (Cosmides & Tooby, 2000; Klein et al., in press). Both speed and accuracy can be engineered into a decision-making system, however, if it is designed to retrieve both kinds of information in the

right combination. A generalization is most useful when its scope is delimited: when it is accompanied by information specifying those situations in which it does not apply. Episodic memories that are inconsistent with the generalization can serve this function, because they encode specific situations in which the generalization fails to predict the outcome. This task analysis suggests the following: To render judgments that are both fast and accurate, judgment and decision procedures should be designed to search for summary information in semantic memory and, upon retrieving it, also search for episodic memories that are inconsistent with that summary — ones that place boundary conditions on the summary's scope. We called this the scope hypothesis (Babey et al., 1998; Cosmides & Tooby, 2000; Klein et al., in press).

In support of this argument, we presented data from experiments on trait judgments about naturalistic targets: self and mother. It was already known that when asked to decide whether a trait describes a person (e.g., "Does this describe you: friendly?"), subjects retrieve a trait summary from semantic memory (if one exists). It was also known that, in this decision context, retrieving a trait summary does not activate episodic memories in which the target manifested behavior that is consistent with the trait under consideration. So when it comes to behavioral episodes whose content is trait-consistent, retrieval from semantic and episodic memory is functionally independent. Retrieving a trait summary from semantic memory does not prime trait-consistent episodic memories.

However, the results reported by Babey et al. (1998) demonstrated that this result does not generalize to trait-*inconsistent* episodes — at least with artificial targets. Based on the scope hypothesis, they made the following prediction: When procedures designed to make trait judgments about a person succeed in retrieving a trait summary from semantic memory, they will also retrieve episodic memories of incidents in which that person exhibited behaviors that are inconsistent with that trait. They tested this prediction by using a priming paradigm, and a judgment task in which subjects were asked whether a trait describes a target about whom they had been given behavioral information. As the scope hypothesis predicts, retrieving a trait summary primed trait-inconsistent episodes, but not trait-consistent ones. This was not the case in other judgment contexts: Asking a subject simply to define a trait — a context that does not cause retrieval of a trait summary (Klein et al., 1997; Klein & Loftus, 1993a; Klein et al., 1989; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992: Schell et al., 1996) — primed neither trait-consistent nor trait-inconsistent episodes.

In the experiments reported herein, we tested the generality of the scope hypothesis by seeing whether it extends to judgments of naturalis-

tic targets: the self and one's mother. Our first experiment explored representations about the self. Consistent with the predictions of the scope hypothesis, retrieving a trait summary about the self primed trait-inconsistent episodes but not trait-consistent ones. Moreover, this effect was found regardless of whether the trait being judged was one the subject classified as high, medium, or low in self-descriptiveness. We argued that this was because subjects have so much experience of their own behavior that they are likely to have summary trait representations across an extremely wide variety of traits. Thus, representations such as rarely rude (low self-descriptive), usually trustworthy (high self-descriptive), and sometimes emotional (medium self-descriptive) may all co-exist in semantic memory.

The self studies are relevant for testing the scope hypothesis not only because they involve a naturalistic target, but also because they allow a test in which one can be reasonably confident that the subject has a pre-computed summary representation of the traits being judged (e.g., Klein, Sherman, & Loftus, 1996). However, because the scope hypothesis claims that inconsistent episodes are primed in order to bound the scope of generalizations, it is also necessary to test it under conditions in which generalizations are known to be absent. After all, there is no need to activate inconsistent episodes if no summary exists whose scope needs to be delimited. That is why we conducted the mother study (Experiment 2).

Previous research indicates that subjects do have summary representations for traits they consider highly descriptive of their mothers, but not of ones they consider medium-descriptive. Thus, by comparing episode retrieval times for highly descriptive versus medium-descriptive traits, we were able to test the hypothesis that trait-inconsistent episodes are primed only if a trait summary is retrieved. In accordance with the scope hypothesis, trait-inconsistent episodes were primed for traits highly descriptive of mother, but not for medium-descriptive traits. By contrast, trait-*consistent* episodes were primed only for medium-descriptive traits — ones for which the subject appears to lack a summary representation. This last finding is consistent with the hypothesis that, in the absence of a trait summary, subjects make trait judgments by consulting episodic memories, following a positive test strategy (as they do in many domains; Klayman & Ha, 1987).

In short, by considering when a judgment task would access a generalization whose scope needs to be bound, we were able to predict when retrieval of relevant knowledge from semantic memory would be functionally independent of retrieval from episodic memory, and when it would not.

A view that emphasizes the adaptive function of judgment procedures requires that care be taken in analyzing what these functions

might be. Restricting our attention to trait judgment, we would like to point out that scope problems are only one class of adaptive problem that can be solved by a system that primes trait-inconsistent episodes. There may be many other social inference tasks that would produce similar — or contrasting — phenomena. It would be particularly interesting to investigate episode priming in the context of discourse processing, especially in arguments. Imagine, for example, that you believe yourself to be honest. In the experiments reported above, one would expect retrieval of that summary to prime trait-inconsistent episodes. But suppose you are engaged in an argument with someone who has just accused you of being dishonest. Episodes in which you were honest may leap to mind, even though these are consistent with your own self-assessment. After all, they are *inconsistent* with your accuser's assessment of you, and that assessment is the topic that your antagonist has raised.

CONCLUSIONS

The results we have reported offer strong convergent evidence in support of the scope hypothesis. According to this hypothesis, maintaining a database of episodes even after summary representations have been formed is functional, as is the phenomenon of priming itself. This hypothesis, in turn, was derived from an evolutionary task analysis, which assumes that the human cognitive architecture has components that were designed by natural selection for solving certain problems in information management (e.g., Cosmides & Tooby, 2000; Klein et al., in press). Taking this perspective allowed us to predict and confirm a highly specific and articulated pattern of results, one not easily explained by alternative theories of social memory and judgment (for discussion, see Babey et al., 1998; Klein et al., in press). We believe framing questions about cognitive processing in functional terms has considerable potential to reveal the intricacies of the human cognitive architecture, and we hope our results encourage our colleagues to consider this kind of analysis when testing models within personality and social psychology.

REFERENCES

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

Anderson, N. H. (1989). Functional memory and on-line attribution. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 175-220). Hillsdale, NJ: Erlbaum.

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.

Blaxton, T. A. & Neely, J. H. (1983). Inhibition from semantically related primes: Evidence of a category-specific inhibition. *Memory & Cognition*, *11*, 500-510.

Budesheim, T. L., & Bonnelle, K. (1998). The use of abstract trait knowledge and behavioral exemplars in causal explanations of behavior. *Personality and Social Psychology Bulletin*, *24*, 575-587.

Buss & Craik (1984). Acts, dispositions, and personality. In B. A. Maher & W. B. Maher (Eds.), *Progress in experimental personality research* (pp. 241-301). San Diego: Academic Press.

Carlston, D. E., & Skowronski, J. J. (1986). Trait memory and behavior memory: The effects of alternative pathways on impression judgment response times. *Journal of Personality and Social Psychology*, *50*, 5-13.

Chaiken, S., & Baldwin, M. W. (1981). Affective-cognitive consistency and the effect of salient behavioral information on the self-perception of attitudes. *Journal of Personality and Social Psychology*, *41*, 1-12.

Collins, A. M., & Quillian, M. R. (1970). Facilitating retrieval from semantic memory: The effect of repeating part of an inference. *Acta Psychologica*, *33*, 304-314.

Cosmides, L. & Tooby, J. (2000). Consider the source: The evolution of adaptations for decoupling and metarepresentations. In D. Sperber (Ed.), *Metarepresentations: A multidisciplinary perspective* (pp. 53-115). Vancouver Studies in Cognitive Science. NY: Oxford University Press.

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.

Dagenbach, D., & Carr, T. H. (1994). *Inhibitory processes in attention, memory and language.* San Diego: Academic Press.

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.

Devine, P. G., Hirt, E. R., & Gehrke, E. M. (1990). Diagnostic and confirmation strategies in trait hypothesis testing. *Journal of Personality and Social Psychology*, *58*, 952-963.

Epstein, S. (1979). The stability of behavior: I. On predicting most of the people much of the time. *Journal of Personality and Social Psychology*, *37*, 1097-1126.

Fong, G. T., Markus, H. (1982). Self-schemas and judgments about others. *Social Cognition*, *1*, 191-204.

Hirshman, E., & Lanning, K. (1999). Evaluating the role of self-judgments in conscious recollection. *Applied Cognitive Psychology*, *13*, 29-42.

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

Kihlstrom, J. F., & Klein, S. B. (1994). The self as a knowledge structure. in R. S. Wyer Jr. & T. K. Srull (Eds.). *Handbook of social cognition, Vol. 1: Basic processes, Vol. 2: Applications* (2nd ed., pp. 153-208). Hillsdale, NJ: Erlbaum.

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.

Kircher, T. T. J., Senior, C., Phillips, M. L., Benson, P. J., Bullmore, E. T., Brammer, M., Simmons, A., Williams, S. C. R., Bartels, M., & David, A. S. (2000). Towards a functional neuroanatomy of self processing: Effects of faces and words. *Cognitive Brain Research*, *10*, 133-144.

Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, *94*, 211-228.

Klein, S. B. (1999). Memory and the self. *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: Psychology Press.

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

Klein, S. B., Chan, R. L., & Loftus, J. (1999). Independence of episodic and semantic self-knowledge: The case from autism. *Social Cognition*, *17*, 413-436.

Klein, S. B., Cosmides, L., Tooby, J. & Chance, S. S. (in press). Decisions and the evolution of memory: Multiple systems, multiple functions. *Psychological Review*.

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). The role of abstract and exemplar-based knowledge in self-judgments: Implications for a cognitive model of the self. In T. K. Srull, R. S. Wyer Jr. (Eds.), *Content and process specificity in the effects of prior experiences. Advances in social cognition* (Vol. 3., pp. 131-139). Hillsdale, NJ: Erlbaum.

Klein, S., & Loftus, J. (1993a). The mental representation of trait and autobiographical knowledge about the self. In T. K. Srull & R. S. Wyer (Eds.), *The mental representation of trait and autobiographical knowledge about the self. Advances in Social Cognition* (Vol. 5, pp. 1-49). 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 (1993c). Some lingering self-doubts: Reply to commentaries. In T. K. Srull, R. S. Wyer Jr. (Eds.), *The mental representation of trait and autobiographical knowledge about the self. Advances in social cognition* (Vol. 5, pp.171-180). Hillsdale, NJ: Erlbaum.

Klein, S. B., Loftus, J., & Burton, H. A. (1989). Two self-reference effects: The importance of distinguishing between self-descriptiveness judgments and autobiographical retrieval in self-referent encoding. *Journal of Personality and Social Psychology*, *56*, 853-865.

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*, *125*, 250-260.

Klein, S. B., Loftus, J., & Plog, A. E. (1992). Trait judgments about the self: Evidence from the encoding specificity paradigm. *Personality and Social Psychology Bulletin, 18,* 730-735.

Klein, S. B., Loftus, J., & Sherman, J.W. (1993) The role of summary and specific behavioral memories in trait judgments about the self. *Personality and Social Psychology Bulletin*, *19*, 305-311.

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

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.

Loftus, G. R., & Loftus, E. F. (1974). The influence of one memory retrieval on a subsequent memory retrieval. *Memory and Cognition, 2*, 467-471.

Lord, C. G. (1993). The "social self" component of trait knowledge about the self. In T. K. Srull, R. S. Wyer Jr. (Eds.), *The mental representation of trait and autobiographical knowl-*

edge about the self. Advances in social cognition (Vol. 5, pp. 91-100). Hillsdale, NJ: Erlbaum.

Malt, B. C., (1989). An on-line investigation of prototype and exemplar strategies in classification. *Journal of Experimental Psychology: Learning Memory and Cognition*, *15*, 539-555.

Park, B. (1986). A method for studying the development of impressions of real people. *Journal of Personality & Social Psychology*, *51*, 907-917.

Roediger, H. & Neely, J. H. (1982). Retrieval blocks in episodic and semantic memory. *Canadian Journal of Psychology*, *36*, 213-242.

Sande, G. N. (1990). The multifaceted self. In J. M. Olson & M. P. Zanna (Eds.), *Self-inference processes: The Ontario symposium* (Vol. 6, pp. 1-16). Hillsdale, NJ: Erlbaum.

Sande, G. N., Goethals, G. R., & Radloff, C. E. (1988). Perceiving one's own traits and others': The multifaceted self. *Journal of Personality and Social Psychology*, *54*, 13-20.

Schell, T. L., Klein, S. B., & Babey, S. H. (1996). Testing a hierarchical model of self-knowledge. *Psychological Science*, *7*, 170-173.

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.

Sherman, J. W., Klein, S. B., Laskey, A., & Wyer, N. A. (1998). Intergroup bias in group judgment processes: The role of behavioral memories. *Journal of Experimental Social Psychology, 34*, 51-65.

Smith, E. R. (1990). Content and process specificity in the effects of prior experience. In T. K. Srull & R. S. Wyer (Eds.), *Advances in Social Cognition* (Vol. 3., pp. 1-59). Hillsdale, NJ: Erlbaum.

Snyder, M., & Cantor, N. (1979). Testing hypotheses about other people: The use of historical knowledge. *Journal of Experimental Social Psychology*, *15*, 330-342.

Tranel, D. & Damasio, A. R. (1993). The covert learning of affective valence does not require structures in the hippocampal system or amygdala. *Journal of Cognitive Neuroscience*, *5*, 79-88.

Tulving, E. (1993). 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. Advances in Social Cognition* (Vol. 5, pp.147-156). Hillsdale, NJ: Erlbaum.

Uleman, J. S. (1989). A framework for thinking intentionally about unintended thoughts. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 425-449). New York: Guilford.

Urdang, L. (1985). *The basic book of synonyms and antonyms*. New York: New American Library.

Wagenaar, W. A. (1994). Is memory self-serving? In U. Neisser & R. Fivush (Eds.), *The remembering self: Construction and accuracy in the self-narrative* (pp. 191-204). New York: Cambridge University Press.

Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, *20*, 273-281.

Webster's new world dictionary (2nd college ed.). (1970). New York: The World Publishing Company.

Winter, L., Uleman, J. S., & Cunniff, C. (1985). How automatic are social judgments? *Journal of Personality and Social Psychology*, *49*, 904-917.

Wright, J. C., & Mischel, W. (1988). Conditional hedges and the intuitive psychology of traits. *Journal of Personality and Social Psychology, 55*, 454-469.