

Decisions and the Evolution of Memory: Multiple Systems, Multiple Functions

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

Memory evolved to supply useful, timely information to the organism's decision-making systems. Therefore, decision rules, multiple memory systems, and the search engines that link them should have coevolved to mesh in a coadapted, functionally interlocking way. This adaptationist perspective suggested the scope hypothesis: When a generalization is retrieved from semantic memory, episodic memories that are inconsistent with it should be retrieved in tandem to place boundary conditions on the scope of the generalization. Using a priming paradigm and a decision task involving person memory, the authors tested and confirmed this hypothesis. The results support the view that priming is an evolved adaptation. They further show that dissociations between memory systems are not—and should not be—absolute: Independence exists for some tasks but not others.

Memory is a gift of nature, the ability of living organisms to retain and to utilize acquired information or knowledge. . . . Owners of biological memory systems are capable of behaving more appropriately at a later time because of their experiences at an earlier time, a feat not possible for organisms without memory. (Tulving, 1995a, p. 751)

If there is one proposition on which all psychologists seem to agree, it is that memory is useful. Memory allows organisms to adjust their behavior on the basis of information they acquire ontogenetically, through their experiences with the world.

This connection between personal experience, memory, and behavior implies a close relationship between learning and memory: “Memory in biological systems always entails learning (the acquisition of information) and . . . learning implies retention (memory) of such information” (Tulving, 1995a, p. 751). Recognizing this, psychologists have long exploited learning paradigms to study the properties of memory (e.g., by list learning to probe free recall, cued recall, and recognition).

The same connection also implies a close relationship between decision rules¹ and memory. An organism cannot behave “more

appropriately”—that is, more adaptively—at a later time because of experiences at an earlier time unless it is equipped with rules that use ontogenetically acquired information to make decisions. Indeed, memory systems must have evolved their structure in response to the informational needs of the decision rules guiding behavior. This is because memory properties that have no impact on an organism's decisions will not be visible to—and hence will not be shaped by—selection. Moreover, because decision rules often differ in what information they require, different sets of decision rules may activate different retrieval processes or search engines and may access different memory systems or sets of memory systems. Without engines that can search for and retrieve the right information, supplying it to the right decision rule at the right time, an organ designed to store ontogenetically acquired information—that is, a memory system—would be a pointless appendage.

Hence, adaptive behavior over time depends on a functional, coadapted relationship between three components: (a) the decision rules that guide the organism's behavior, (b) the memory systems that store the data required by the decision rules, and (c) the search engines that identify and route the stored information dynamically to the decision rules that are active. This implies that decision tasks can be experimentally manipulated to illuminate the designs and differentiated functions of multiple memory systems, just as learning paradigms have in the past.²

Stanley B. Klein and Sarah Chance, Department of Psychology, University of California, Santa Barbara (UCSB); Leda Cosmides, Department of Psychology and Center for Evolutionary Psychology, UCSB; John Tooby, Department of Anthropology and Center for Evolutionary Psychology, UCSB.

Preparation of this article was supported by an Academic Senate Research Grant to Stanley B. Klein from UCSB, by grants to John Tooby from the James S. McDonnell Foundation and the National Science Foundation (#BNS9157–449), and by the UCSB Office of Research (Research Across Disciplines Program: Evolution and the Social Mind).

We thank Endel Tulving and Michael S. Gazzaniga for their insightful and helpful comments on this article and for organizing the conference that inspired it (Miniconference on Evolution, Memory, and Consciousness, University of California Davis Center for Neuroscience, 1995). We also thank Steve Pinker and Art Glenberg for very useful comments and suggestions.

Correspondence concerning this article should be addressed to Stanley B. Klein, Department of Psychology, University of California, Santa Barbara, California 93106. E-mail: klein@psych.ucsb.edu

¹ We intend that the term *decision rule* be interpreted broadly to include any mechanism that alters behavior on the basis of information (thus it could apply to an aplysia or pigeon as well as to a human being). Our use of this term is not meant to endorse or reject any particular proposals about whether the implementation is symbolic and rule following or subsymbolic and rule approximating.

² What do we mean by *decision rule* and *search engine*? Researchers observe organisms making decisions, and making them in a way that is patterned rather than random. Similarly, researchers also commonly obtain evidence that information was retrieved from memory. To avoid confusion, we use the terms *decision rule* and *search engine* to distinguish between observed phenomena (decisions, retrieval) and the hypothesized computational entities (decision rules, search engines) that are proposed to account

The search engine component of the triad can be explored in this way as well, allowing the development and experimental evaluation of alternative theories for why priming was engineered into the cognitive architecture. So far, priming has been defined as an experimental phenomenon: Information is said to have been *primed* when performing one task changes the availability of that information for use in a subsequent task. At the most general level, one family of theories would explain priming as a by-product of the way the cognitive architecture happens to be organized, whereas another family of theories would explain priming as having been specifically engineered into the cognitive architecture by evolution because of the functional contributions it made to behaving adaptively. According to this functionalist or adaptationist view, search engines should have evolved to accelerate the delivery of appropriate information to an activated decision rule.³ Search engines are designed to accomplish this by using current decision-making operations and events to predict which information is likely to be required by subsequent decision tasks, and to facilitate its delivery. As a result, the information that these procedures predict will be more relevant will be retrieved faster than other information, and therefore will be more available for use by activated decision rules. Experimentally, this would produce the observed phenomenon of priming. By seeing which patterns of priming are deployed by alternative decision rules, different hypotheses about the designs of search engines, and of the nature of priming itself, can be tested.

An Adaptationist Approach to Studying Memory

It is often productive to explore experimental results with the working hypothesis that they are, to some significant extent, the expression of a functional computational design (e.g., J. R. Anderson, 1989, 1991; J. R. Anderson & Milson, 1989; Cosmides & Tooby, 1987; Sherry & Schacter, 1987). That is, the human cognitive architecture can be interpreted as if it were a relatively well engineered system, analyzable through the application of such

for the patterns in these observations. This distinction allows us to investigate hypotheses about the properties of these computational entities without misleading readers into thinking that we are referring to the observed phenomena instead. After all, a claim about a design feature of one search engine (a computational entity) may not apply to all cases of retrieval (an observed phenomenon): Different search engines may exist, each with a different effect on retrieval. A decision rule is any system of computational organization in the cognitive architecture that takes information as input and produces behavioral decisions as output. A search engine is any set of stable procedures in the cognitive architecture that monitors the information requirements that an activated decision rule needs to operate, searches memory, and delivers information from memory to the activated decision rule. It may be that there is only one integrated search engine, or there could be many search engines—one for each kind of decision rule, for example. Making any more specific claims at this point would be premature and unwarranted and would be unnecessary for deriving and testing the hypotheses in this article. Similarly, we are not making any strong claims about whether the decision rules we are experimentally activating, or the search engines that supply them with information, are “innate” versus “acquired”—or more properly, reliably developing versus contingently developing. Our claim is merely that the procedures that cause the reliable development or contingent acquisition of decision rules and search engines are (a) evolved, (b) shaped by selection so that they generate functional products, and (c) designed to produce decision rule–search engine pairs that fit together to operate adaptively.

concepts as function, design features, task analyses, goal of computation, efficiency, reliability, and so on. This is warranted because, over the course of evolution, modifications in the design of the universal human cognitive architecture were likely to be incorporated to the extent that they improved the functional operation of the architecture—that is, increased the rate that the architecture successfully solved adaptive information-processing problems (e.g., Cosmides & Tooby, 1987; Sherry & Schacter, 1987). Chance and natural selection are the only two families of process that drive evolutionary change, and together over evolutionary time they constructed the human cognitive architecture. Researchers who apply functionalist concepts or analyses to cognitive problems should welcome the integration of modern evolutionary biology into cognitive science because, outside of the operation of natural selection on our ancestors, there is no logical reason the brain should include any functionally organized elements beyond what random processes would produce.⁴ For this reason, in this article we apply to the study of memory systems the logic of adaptationism⁵—the principles for studying the functional

³ More fully, a search engine ought to be designed to deliver information as a function of how its procedures predict its potential appropriateness, so that information predicted to be most relevant should be accelerated in comparison with information predicted to be less relevant. Moreover, information that meets criteria for appropriateness, but that has features that predict that it might disrupt decision making, is expected to be retarded or inhibited.

⁴ One of the benefits of injecting modern evolutionary biology into psychology is a more stringent, rigorous set of criteria for understanding the origins and nature of functional organization in biological systems such as the brain. Biologists recognize that all functional organization is present ultimately because of the prior operation of natural selection (see, e.g., Dawkins, 1982, 1986). Selection constructs proximate machinery whose operation then may interact with the environment ontogenetically to construct individual-specific organization, such as a skill, many of whose details were not prespecified by selection. However, absent natural selection, there would be no well-engineered proximate machinery that responds to environmental structure (or the organism’s goals) by orchestrating the construction of the functional organization. For this reason, every time a psychologist makes a functionalist argument, she is (knowingly or unknowingly) making an evolutionary argument. To improve the quality of functionalist reasoning in psychology, it is better to recognize that functionalism is evolutionary functionality, and to explicitly ground functionalist arguments in the only body of knowledge that makes them legitimate.

⁵ *Adaptationism* refers to the evolutionary principles used to distinguish functional organization from the by-products of function and from noise introduced into designs by the stochastic components of evolution. The rapid adaptationist advances in evolutionary biology that began in the 1960s were rooted in a widespread reaction against the practice of over-attributing adaptation (Williams, 1966). Indeed, what defined the emergence of modern adaptationism as a community and set of theories was its anti-Panglossian empirical and theoretical practices, which have been central to evolutionary biologists ever since. For example, in the book *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*, Williams (1966) began his summary by saying, “Evolutionary adaptation is a special and onerous concept that should not be used unnecessarily, and an effect should not be called a function unless it is clearly produced by design and not by chance. When recognized, adaptation should be attributed to no higher a level of organization than is demanded by the evidence” (p. v). Some terminological confusion has emerged for nonbiologists stemming from Gould and Lewontin’s idiosyncratic use of the word *adaptationism* to mean Panglossian overattributions of adaptation (see, e.g., Gould & Lewontin, 1979).

design of naturally selected systems (Dawkins, 1986; Williams, 1966; for application to computational problems, see Cosmides & Tooby, 1987, 1992, 1997; Marr, 1982). Although this way of thinking is commonplace and is the root of many discoveries in other fields such as physiology, medicine, and animal behavior (e.g., Daly & Wilson, 1995; Gallistel, 1995; Mayr, 1983; Williams & Nesse, 1991), it has not yet been widely applied in memory research (with several notable exceptions, e.g., Nadel, 1994; Sherry & Schacter, 1987; Tulving, 1995b). We think that functional analysis may contribute significantly to the understanding of the designs of memory systems and their relationships to other aspects of the cognitive architecture, and that recent data on multiple memory systems invites such an analysis. In this article, we test hypotheses arising from only the most general functional considerations of the interrelationships between decision rules, search engines, and multiple memory systems, deferring the analysis of domain-specific effects on memory to the future.

Multiple Systems, Multiple Functions

When learning was believed to be a unitary process (either associationistic or inferential, depending on the decade) with a single function, a unitarian view of memory prevailed as well. The same processes of encoding, storage, and retrieval were thought to operate, whether the information represented was a fact about the world, a skill, or a personal episode. Memory was a single system with a single function: to store information for later use, no matter what its content or kind (for reviews, see Foster & Jelicic, 1999; Nadel, 1994; Polster, Nadel, & Schacter, 1991; Schacter, 1995; Schacter & Tulving, 1994; Tulving, 1995a).

During the last two decades, however, unitarian views of both learning and memory have been criticized on two grounds: (a) They have trouble accounting for an expanding catalog of results that document dissociations—both neural and functional—in how information from different domains is acquired, represented, and retrieved (for reviews, see Foster & Jelicic, 1999; Polster et al., 1991; Nadel, 1994; Roediger & Craik, 1989; Schacter & Tulving, 1994; Sherry & Schacter, 1987; Squire, 1987); and (b) some information-processing problems cannot be solved by a single system of learning and/or memory because the computational demands of one problem are incompatible with the computational demands of the other (e.g., McClelland, McNaughton, & O'Reilly, 1995; Nadel, 1994; Pinker, 1979, 1984; Sherry & Schacter, 1987; Wexler & Culicover, 1980). In the study of memory, considerations of this kind led a number of scientists to propose that the human cognitive architecture contains up to five different, functionally isolable memory systems: procedural, perceptual-representational, primary (working memory), semantic, and episodic (for reviews, see Markowitsch, 1995; Squire & Knowlton, 1995; Tulving, 1995b; Tulving & Schacter, 1990; Weiskrantz, 1987). Each system mediates the acquisition, representation, and retrieval of a different kind of knowledge or skill, and at least four of them are thought to be composed of several functionally specialized subsystems.

Like all theories that question unitarian views, the proposal that one brain contains multiple memory systems is somewhat controversial. Nevertheless, studies from neuropsychology and experimental psychology have produced enough evidence for several neurally and functionally dissociable memory systems that this hypothesis must be taken seriously (e.g., Campbell & Conway,

1995; Cohen, 1984; Foster & Jelicic, 1999; Hodges, Spatt, & Patterson, 1999; R. Johnson, Kreiter, Zhu, & Russo, 1998; Knowlton, Mangels, & Squire, 1996; Markowitsch et al., 1993; Markowitsch, Calabrese, Neufeld, Gehlen, & Durwen, 1999; McCarthy & Warrington, 1990; McClelland et al., 1995; Perani et al., 1993; Peretz, 1996; Reber, Knowlton, & Squire, 1996; Schacter & Tulving, 1994; Shallice, 1979; Squire, 1987, 1992; Tulving, 1983, 1985, 1987, 1993a; Tulving & Schacter, 1990; Van der Linden, Bredart, Depoorter, & Coyette, 1996; Vargha-Khadem et al., 1997; Warrington, 1975; Weiskrantz, 1987, 1990).

If there are multiple memory systems—systems that are functionally isolable—then it is important to ask why they exist. 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. Understanding the problems each memory system solves will help guide a more systematic mapping of the design features of that memory system. In short, memory research could benefit from developing more nuanced, differentiated theories about the functions of memory. This article is offered as a step in that direction.

The first step in applying adaptationist logic to the study of memory is to think of memory as a component of a machine, the second step is to distinguish a component's capabilities from its functions, the third is to identify the functions of the component being investigated, and the fourth is to identify the causal relationship between the functions of the component and its structure.

Principles of Organic Design

Function Versus Capability

For almost a century, psychologists have been exploring the capabilities of human memory systems, which are staggering. These systems can encode, store, and retrieve a vast array of information, including much that is adaptively arbitrary and evolutionarily novel, from nonsense syllables, jingles, and batting averages to chess moves, equations, and the triple Lutz. However, agnostically cataloging arbitrary samples out of the inexhaustible set of everything a memory system is capable of doing is not likely to lead to knowledge of its function.

A machine is a system designed to solve problems—that is, a system whose parts exist in their present form because that arrangement solves a problem. Memory is a component of a neural machine designed to use information acquired in the past to coordinate an organism's behavior in the present.

To understand the design of organic machines (as well as machines made by humans), one needs to make a sharp conceptual distinction between a machine's capabilities and its function. 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 the three-hole punch's parts exist in their present form: why it has elements sharp enough to cut paper, why the elements are cylindrical (to cut out holes for the rings, not slits or lacy patterns), why their diameter is 1/4–1/2 in., why there are three of them, why they are spread out, why they form a straight line, why one handpress lowers all three, and so on. These elements are *design features*—aspects of the machine that are there because they contribute to its function.

Yet every machine, by virtue of having a particular causal structure, 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. The production of small circles of paper is a by-product of the machine's design. None of the parts named above exist because that arrangement makes confetti (and, in fact, they are not particularly well designed for doing so; a confetti maker should have many more punchers, these should be as close together as possible so as not to waste paper, they do not need to be cylindrical, etc.). Confetti making does not explain the presence or arrangement of the punch's parts. Nor do any of the punch's other capabilities—its usefulness for weighting down paper or for whacking indolent graduate students, for example. These capabilities are arbitrary with respect to its intended function, mere by-products of the machine's design.

The tradition of studying memory by seeing what it is capable of doing—without asking what it was designed to do—is like studying a three-hole punch as if it were a confetti maker or a paperweight. It is not an efficient method for homing in on the small set of highly ordered, interlocking elements that embody the functional design of a system. Indeed, some research principles may actively prevent the discovery of such central elements. In graduate school, for example, we (the authors) were taught the Ebbinghaus tradition: that in constructing memory experiments, one should avoid using any stimuli that are emotionally charged or drawn from adaptively important domains. But this prejudices the issue. What if there are memory systems that are designed to store information about one's own children or about food locations or about dangerous situations? Such systems have been found in other species (for offspring, Beecher, 1990; Kendrick, Levy, & Keverne, 1992; McCracken & Gustin, 1992; for food locations, Hampton et al., 1995; for danger, Mineka & Cook, 1988). What if functionally specialized, content-specific memory systems exist in the human mind as well? (as indeed they do; see Jacobs & Nadel, 1985; Ohman, Dimberg, & Ost, 1985, Pittman & Orr, 1995; Sherry & Schacter, 1987; Silverman & Eals, 1992). These cannot be discovered if one's experimental repertoire is limited to stimuli chosen because they are functionally arbitrary, or if one's pretheoretical commitments mandate that all a system's capabilities are co-equal. They are not. A system's function provides a privileged frame of reference: the only way of dividing a machine into parts that explains why those parts exist and take the form that they do. As a secondary matter, a functional description will, because it is causal, allow one to derive what the system's other, nonfunctional capabilities are as well.

Decision Rules, Search Engines, and Memory

Recently, some memory researchers have been investigating the relationship between the design of a given memory system and its adaptive function. Scientists taking this approach have been finding that many seemingly arbitrary design features of memory systems can be understood by considering what problems they were designed to solve (J. R. Anderson, 1989, 1991; J. R. Anderson & Milson, 1989; Babey, Queller, & Klein, 1998; Glenberg, 1997; Hampton et al., 1995; O'Keefe & Nadel, 1978; Nadel, 1994; Rozin, 1976; Sherry, 1997; Sherry & Schacter, 1987; Suddendorf & Corballis, 1997; Tulving, 1995b). We argue that further progress can be made by considering what problems the procedures that use data from memory systems were designed to solve. After all,

whether an organism can profit from its past depends not only on the design of its memory systems, but also on those aspects of the cognitive architecture that interact with these systems: learning mechanisms, inference systems, decision rules, and search engines. Here, we focus on cognitive processes that use stored information to make judgments and decisions. From a design point of view, there is an intimate relationship between decision rules, retrieval processes, and memory systems. For a decision rule that uses stored information to work, it needs search engines: programs designed to respond to the activation of a decision rule by searching for and identifying information in the organism's memory systems that is potentially relevant to the decision being made, and delivering what they find to the input buffer that feeds that decision rule.

However, the more information an organism encodes, the more difficult it is to engineer a system that can deliver the right information to the right decision rule at the right time while not simultaneously flooding it with irrelevant information. What is important is to engineer a match between the content of the stored information and the mechanism that requires it (Nadel, 1994). Decision rules need search engines that "surf the net" for appropriate content. Which content is appropriate depends on the adaptive problem that that decision rule was designed to solve: Mate choice requires different information than food choice or predator avoidance. Different decision rules will need different search engines, and one would expect these search engines to be designed to pull up data structures by their content. Whereas some tasks may require content-specific information from episodic memory alone or from semantic memory alone, other tasks may require content-specific information from both of these memory systems. Therefore, when it comes to retrieval, dissociations between memory systems may not be absolute; as a matter of design, one might find independence for some tasks and joint activation for others. This suggests that, when it comes to retrieval, the extent to which one finds functional independence between memory systems will reflect the informational requirements of the decision rule activated.

In this article, we concentrate on (some of) the functions of episodic and semantic memory, and the decision rules that access them. In doing so, we focus on research by Klein and colleagues that explored the conditions under which episodic and semantic memories are retrieved by decision rules that make trait judgments. These experiments used a priming paradigm and a decision task that requires retrieval of trait knowledge. The results showed that retrieval from episodic and semantic memory is independent in some decision contexts but not in others. They further showed that whether retrieval is independent depends in lawful ways on the content of the information being retrieved. The situations under which independence breaks down suggest that priming of person knowledge is not merely a by-product of spreading neural activation or some other incidental feature of the system's design: Rather, it is a functional component of a decision mechanism designed to use information from one memory system to place boundary conditions on information retrieved from the other (the scope hypothesis).

Why Trait Judgment?

We believe that episodic and semantic memory systems evolved to solve many different problems and that they are accessed by many different kinds of decision rules. Nevertheless, in testing a

theory, one must start with a single, specific problem domain. We began with the problem of judging personality traits.

Trait judgment is a good place to begin for five reasons. First, making trait judgments involves categorization and, as Jackendoff (1987) put it, "The ability to categorize is what makes it possible to use previous experience to guide the interpretation of new experience, for without categorization, memory is virtually useless" (p. 135). If memory systems and the mechanisms that use them have interlocking design features, then deciding whether a trait category adequately describes a person is the kind of task in which one should find these features.

Second, categorizing people—both oneself and others—as "kind," "cruel," "friendly," "aggressive," and so on, is a pervasive aspect of human life. People automatically condense the rich complexity of human actions into a limited number of personality dimensions that have predictive validity (Funder, 1995; Funder & Sneed, 1993; Kolar, Funder, & Colvin, 1996) and use these judgments in deciding how to interact with others. These judgments are highly accurate, even when they are made quickly and on the basis of very limited information (Funder & Sneed, 1993). Moreover, one's competence in this regard develops in the absence of explicit instruction (much as learning a first language does), despite the inherent complexity of the task (consider what it would take to create a computer program that could infer motivations and intentions from actions and use these inferences to make trait judgments). Taken together, these facts—ubiquity, complexity, speed, accuracy, automaticity, and competence without tuition—suggest that the human cognitive architecture may have evolved computational machinery specialized for performing personality trait judgments.

Third, evidence from primatology, paleoanthropology, and studies of modern hunter-gatherers converge on a picture of human hunter-gatherer ancestors 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 accurate personality information and using it to predict behavior. Primate studies suggest this ability is phylogenetically ancient: Chimpanzees and bonobos also adjust their social behavior in ways that reflect the personalities of different interactants (de Waal, 1982; de Waal & Lanting, 1997). The hominid and great ape lines are thought to have diverged 5–10 million years ago, so there has been a very long period of time during which selection could have specialized and improved the mechanisms that generate these judgments. Given this time depth and the positive fitness consequences that would accrue from any design alteration that improves its bearer's ability to predict behavior, one would expect the decision rules that make trait judgments, as well as the memorial databases that store the information used by these decision rules, to have been targets of selection. Indeed, we suspect that some of the most important functions of episodic memory have to do with social interaction, which depends (in part) on judging traits.

Fourth, there is a large literature on the use of memory in trait judgment, including many neuropsychological studies showing dissociations in retrieval of episodic and semantic trait knowledge (e.g., Damasio, Tranel, & Damasio, 1989; Klein, Loftus, & Kihlstrom, 1996; Tranel & Damasio, 1993; Tulving, 1993a). In humans, the decision rules that make trait judgments are known to access both the episodic and the semantic memory systems (see below). Priming paradigms allow one to sensitively probe the circumstances under which this occurs, and a good model system

describing these circumstances in some detail has already been developed by Klein and Loftus (1993b). This makes trait judgment a good domain for examining how decision rules interact with episodic and semantic memory systems and for testing hypotheses about the conditions under which independence will, and will not, be found.

Fifth and last, the computational problems that arise in trait judgment are an instance of a broader category of problems involving the adaptive functions of what we call *inceptive* and *derived* memories. Inceptive memories are representations of the world stored in the way they were encoded at their inception (the time at which they were first experienced). Many of the episodic memories retrieved during trait judgment appear to be inceptive. A derived memory is a higher level representation that was derived from inceptive memory stores but was computationally transformed to supply information in a form that minimizes the need for further processing by the decision rules that use it. Semantic memory includes many databases of derived memories, such as the trait summaries that are retrieved by the decision rules that judge whether a person is friendly, cruel, or kind. We argue that inceptive and derived memories exist because both are useful, but for different purposes.

Because a great deal is already known about how inceptive and derived memories interact in trait judgment, this domain provides a concrete point of entry for considering what we believe to be a more general class of information-management problems involving trade-offs between speed and accuracy. According to the scope hypothesis, an excellent package of speed plus accuracy can be engineered into a decision system by jointly activating a derived memory and the inceptive memories that contradict it. Therefore, although we tested the scope hypothesis in the domain of trait judgment, we believe it applies more generally and close this article with a review of evidence suggesting that this design feature is found in decision rules in a number of different domains.

Judging Oneself

There is a particularly rich literature on how people judge their own traits (for review, see Babey et al., 1998; Bem, 1972; Hampson, 1982; Kihlstrom et al., 1988; Kihlstrom & Klein, 1994; Klein & Loftus, 1993b; Schneider, Hastorf, & Ellsworth, 1979; Wyer & Gordon, 1984). If one wishes to investigate cognitive processes that access episodic and semantic memory, this is a propitious area. Most theories of trait judgment assume that the individual has (or had) a database of experiences from which trait judgments can be made. Every potential subject has a large database of experiences about themselves, which includes knowledge of their own reactions to a wide variety of social situations. For this reason, one can test hypotheses without worrying that the subject merely lacks memories relevant to the topic.

We examine below the relation among episodic memory, semantic memory, and a particular type of categorization/decision process—judgments of one's own traits (e.g., Kihlstrom & Klein, 1994, 1997; Klein & Loftus, 1990, 1993a, 1993b; Klein, Babey, & Sherman, 1997; Klein, Loftus, & Burton, 1989; Klein, Loftus, & Kihlstrom, 1996; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein, Sherman, & Loftus, 1996; Lord, 1993; Schell, Klein, & Babey, 1996). For example, when someone asks, "Are you kind?" do categorization algorithms consult a library of personal memories, computing answers on the fly from

whatever episodes happen to be activated by the situation at hand? Or have many of the answers to this class of questions been previously computed and stored as a database of abstract summary traits? That is, does a “look-up” procedure activate a representation in summary form—an entry under the heading “self” in an encyclopedic semantic memory? Or do both go on in parallel? If a summary representation is activated, are related episodes from an autobiographical memory system activated in its wake? If so, what kind of episodes does it entrain, supporting evidence or counterexamples? For example, when you are asked, “Are you kind?” does this prime recall of episodes that represent specific incidents in which you were kind—or unkind?

We first review findings from studies of both normal and clinical populations that converge on the conclusion that semantic knowledge of our own traits and the episodic memories that bear on it are typically stored and retrieved independently of one another. We then describe circumstances in which one might expect this independence to break down, and report experiments showing that it does.

Episodic and Semantic Personal Memory

Each of the five memory systems so far proposed is thought to mediate the acquisition, representation, and retrieval of a different kind of knowledge or skill. By hypothesis, the procedural system mediates conditioning and the acquisition of skills, the perceptual-representational system supports the identification of objects based only on modality-specific information about their form and structure, primary (working, short-term) memory is a buffer that makes information temporarily accessible to mechanisms that cause language comprehension and certain forms of inference, the semantic memory system contains general knowledge about the world, and the episodic memory system creates a personal past (e.g., Markowitsch, 1995; Schacter & Tulving, 1994; Tulving, 1983, 1995b; Tulving & Schacter, 1990; Weiskrantz, 1987).

Episodic memory records events as having been experienced by the self within a subjective space-time matrix; when retrieved, these events are reexperienced in a quasi-perceptual way, with conscious awareness that “this happened to me” (e.g., Tulving, 1972, 1983, 1985, 1989, 1993b). Episodic recollection thus has a self-referential quality, which was thought to be missing from other types of memorial experiences (e.g., procedural, perceptual, semantic; e.g., Tulving, 1984, 1987). Over the last few years, however, evidence from clinical neuropsychology has raised the possibility that other types of memory experiences might be self-referential as well. For example, Klein, Loftus, and Kihlstrom (1996) report the case of W.J., who suffered a concussive blow to the head shortly after completing her first quarter in college. As a result of her injury, W.J. showed profound amnesia for personal events and experiences over the 6 months immediately prior to her accident. However, despite her dense retrograde amnesia, W.J.’s memory for general facts about her life during that period seemed largely intact. For example, she knew which classes she had attended, although she could not recall a specific occasion when she attended class or a specific event that happened during a class; she knew the names of teachers and friends from college, but could not remember particular experiences shared with them.

A similar dissociation between event and factual self-knowledge is seen in the case of patient K.C. (e.g., Tulving, 1989; Tulving, Schacter, McLachlan, & Moscovitch, 1988). K.C., as a result of a

severe head injury, was no longer able to consciously bring to mind a single personal experience from any point in his life. Despite his total loss of episodic personal memory, K.C. still knew a variety of facts about himself. He knew, for example, that his family owned a summer house and where it was located, but he could not recall a single occasion when he was at the house, nor could he recall a single event that occurred there (for related findings, see Cermak & O’Connor, 1983; Evans, Breen, Antoun, & Hodges, 1996; Evans, Wilson, Wraight, & Hodges, 1993; Klein, Chan, & Loftus, 1999; Klein, Cosmides, Costablie, & Mei, in press; Klein, Rozendal, & Cosmides, in press).

To explain these dissociations between event and factual self-knowledge, Tulving has suggested broadening the concept of self-referential memory to include a semantic as well as an episodic component (e.g., Tulving, 1989, 1993a; Tulving et al., 1988; see also Brewer, 1986; Cermak & O’Connor, 1983; Conway, 1992; Craik et al., 1999; Kihlstrom et al., 1988; Kihlstrom & Klein, 1994, 1997; Klein, 1999; Klein, Sherman, & Loftus, 1996; Kopelman, Wilson, & Baddeley, 1989). Semantic personal memory enables people to retrieve general facts about their personal past. It is self-knowledge in summary form. Although these summary representations are about the self and have been abstracted from memories of the self in a set of events, they do not preserve a record of the events from which they were derived. The fact that the patients reported by Klein, Loftus, and Kihlstrom (1996) and Tulving (1989, 1993a) had access to semantic personal knowledge but not the particular episodes on which this knowledge was based can be interpreted as evidence in support of a separation between these two types of self-referential memories.

The Evidence for Independence

Not only can one have semantic knowledge about the circumstances of one’s life, but one can also have semantic knowledge about one’s own personality traits (for reviews, see Kihlstrom & Klein, 1994; Klein, 2001; Klein & Kihlstrom, 1998; Klein & Loftus, 1993b; Klein, Loftus, & Kihlstrom, 1996). Klein and Loftus and their colleagues recently proposed and tested a model of the relation between memories of personal experiences and trait conceptions of self (e.g., Klein et al., 1989, 1997, 1999; Klein & Loftus, 1990, 1993a, 1993b, 1993c; Klein, Loftus, & Plog, 1992; Klein, Loftus, & Sherman, 1993; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein, Sherman, & Loftus, 1996; Kihlstrom & Klein, 1994, 1997; Schell et al., 1996). Their model incorporated the following three features. First, long-term knowledge of one’s traits is abstracted from multiple experiences with trait-relevant behaviors and represented in semantic memory in summary form. Second, trait judgments about the self are made by accessing these summary representations without reference to the behavioral episodes from which they were derived. And third, summary trait representations are functionally independent of memories of trait-relevant behavioral episodes.

There now is ample evidence that people form summary representations of their own personality traits and retrieve these when asked to decide whether a particular trait describes them (e.g., Craik et al., 1999; Hirshman & Lanning, 1999; Kircher et al., 2000; Klein et al., 1989, 1993, 1997, 1999; Klein & Loftus, 1990, 1993a, 1993b; Klein, Loftus, & Kihlstrom, 1996; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992; Lord, 1993; Schell et al., 1996). Equally important, a large body of evidence

shows that people maintain detailed stores of behavioral episodes they have experienced (e.g., Barclay, 1996; Conway & Bekerian, 1987; Klein et al., 1989, 1997; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992; Nelson, 1993b, 1996; Rubin, 1986; Singer & Salovey, 1993; Strauman, 1990).

Klein and Loftus have provided evidence that episodic and semantic trait self-knowledge is independently represented and used. For example, using a priming paradigm, they examined whether behavioral memories are activated when one decides whether a trait describes oneself (*describe* task). If they are activated, then one should be able to retrieve those memories faster after performing a describe task than after performing a recall task. This was not the case: When subjects were asked to recall a specific behavioral incident in which they manifested a particular trait (*recall* task), those who had first made a self-descriptiveness judgment were no faster than those who had not (e.g., Klein et al., 1989; Klein & Loftus, 1990, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992). Yet the procedure that was used is known to be sensitive enough to detect episodic priming when it occurs (e.g., Babey et al., 1998; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman, 1996; Sherman & Klein, 1994; Sherman, Klein, Laskey, & Wyer, 1998). Klein and Loftus concluded from this that the semantic personal knowledge required for a self-descriptiveness judgment was accessed without activating memories of episodes in which that person's behavior exemplified the trait.

Subsequent experiments that eliminated alternative interpretations of these results provided even clearer evidence that trait self-descriptiveness judgments can be made without reference to behavioral evidence (e.g., Klein et al., 1997; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992). It has been suggested, for example, that some traits are more central to one's self-concept than others and that noncentral traits are less likely to be represented in summary form. If so, it may be necessary to activate and evaluate specific behavioral memories to decide whether a noncentral trait is self-descriptive, even though such "evidence" is unnecessary when the trait in question is central (e.g., Bower & Gilligan, 1979; Kihlstrom & Cantor, 1984; Kihlstrom et al., 1988). In this scenario, behavioral memories would be primed when people judge the self-descriptiveness of noncentral traits, but not central ones. If such a difference exists, it would not have been apparent in the data reported above because those studies combined all traits without regard to level of self-descriptiveness.

To address this concern, Klein, Loftus, Trafton, and Fuhrman (1992) repeated the Klein and Loftus priming procedure and afterward had subjects rate each trait for self-descriptiveness. These ratings were used to sort subjects' response latencies into three levels of trait self-descriptiveness (high, medium, and low). Regardless of the self-descriptiveness of the trait being judged, no priming was observed between the describe and recall tasks (see also Kihlstrom & Klein, 1994; Klein et al., 1997; Klein & Loftus, 1990, 1993b). Klein and Loftus interpreted this result as indicating that there are few, if any, trait dimensions for which people have not had sufficient experience to form summary trait self-knowledge: When it comes to our own personalities, we are all "experts."

Results from studies using other methods also support the independence of episodic and semantic trait knowledge about the self. Klein, Loftus, and Plog (1992), for example, made use of the phenomenon of transfer-appropriate processing (e.g., Roediger &

Blaxton, 1987; Roediger, Weldon, & Challis, 1989) in a study of recognition memory for traits to show that different processes are involved in accessing the two types of knowledge. Klein et al. (1989, Experiment 4) applied the principle of encoding variability (e.g., Bower, 1972; Martin, 1971) in a study of recall for traits and found that the type of information made available when making trait judgments was different from that made available when retrieving trait-relevant behaviors.

Additional support for the independence of episodic and semantic trait self-knowledge recently was presented by Craik et al. (1999). Using positron emission tomography, these investigators discovered that requiring participants to judge trait adjectives for self-descriptiveness produced activation of cortical areas associated with semantic memory retrieval (left frontal regions) but not those associated with episodic memory retrieval (right frontal regions). Similar findings using functional magnetic resonance imaging have recently been reported by Kircher et al. (2000).

Perhaps the most dramatic evidence for the independent storage and retrieval of episodic and semantic trait self-knowledge comes from studies of patients with disorders of memory (e.g., Klein, 2001; Klein et al., 1999; Klein & Kihlstrom, 1998; Klein, Loftus, & Kihlstrom, 1996; Klein, Rozendal, & Cosmides, in press; Reinvang & Gjerstad, 1998; Schacter, Wang, Tulving, & Freedman, 1982; Tulving, 1993a). For example, Tulving (1993a), testing a speculation by Klein and Loftus (1993b), found that patient K.C., who permanently lost his entire fund of episodic memory following a motorcycle accident (and underwent a marked personality change), was able to describe his postmorbid personality with considerable accuracy (his mother's ratings served as the criterion). Thus, K.C. was able to acquire semantic knowledge of his new personality without being able to recall information about any of the behavioral episodes through which that knowledge was gained.

More recently, Klein, Loftus, and Kihlstrom (1996) presented the case of patient W.J., who, as a result of a severe head injury, suffered a temporary loss of episodic memory. W.J. was asked both during her amnesia and following its resolution to make trait judgments about herself. Because her responses when she could access behavioral episodes were consistent with her responses when she could not, Klein, Loftus, and Kihlstrom concluded that the loss of trait-relevant episodic memories did not greatly affect the availability of her semantic trait self-knowledge. Results such as these offer strong support for the general proposition that summary knowledge of one's traits is represented and accessed separately from memory of one's trait-relevant behaviors (for related findings, see Klein, Cosmides, et al., in press; Klein et al., 1999).

Why Have Both Inceptive and Derived Memories?

To decide whether a trait describes them, Klein and Loftus's (nonamnesic) subjects retrieved a trait summary. It is important to note, however, that subjects could have decided this question by consulting a library of relevant behavioral episodes; after all, they were able to recall specific episodes in which they manifested the trait at issue. Both types of representation—semantic and episodic—were present. It's just that retrieving a semantic trait summary did not cause subjects to recall episodes faster than otherwise.

But why form trait summaries if one can consult a library of episodes? 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? More generally, are there functional reasons to have databases of both inceptive and derived memories?

Trade-Offs: Memory Versus Computation, Speed Versus Accuracy

Advantageous scheduling of slow computation. Social interaction often requires split-second decisions, some of which are best made by taking into account a diverse array of information. However, it takes time to search for, retrieve, reformat, and integrate the disparate (and sometimes incommensurable) information needed to make certain decisions. For example, behavioral episodes are concrete interactions, each involving different actors and circumstances. They do not come with signs labeling the behavior of an individual actor as “kind,” “honest,” or “cruel”; To make such judgments, one needs to make inferences about the motivations and opportunities of the actors in each episode one has experienced or heard about, and then somehow combine these determinations into a summary judgment. All else being equal, better decisions will be made if this analysis is complete before a decision requiring it is imminent. Performing some (or all) of the necessary computation in advance can, therefore, be advantageous.

Inceptive versus derived memories. It is useful to distinguish inceptive memories from derived memories. Inceptive memories are representations of the world stored more or less in the way they were encoded at their inception (the time at which they were first experienced). Many episodic memories appear to be inceptive: They often contain perceptual detail, proprioceptive information, emotion qualia, a sense of the self as present in the encoding situation, and the sense that the memory is very similar to the original experience. But not all inceptive memories are episodic. It is logically possible to have an inceptive memory about a specific event involving the self without its being episodic in character—it would not be episodic if it lacked the subjective sense of self in space and time that is diagnostic of episodic memory.

Derived memories are so called because they were derived from inceptive memory stores but were computationally transformed to supply information in a form that minimizes the need for further processing by the decision rules that use it. A derived memory system—together with its inferential machinery—prepares answers or preprocesses information into usable and ready forms, to facilitate decisions and judgments of a specific type that would otherwise be detrimentally retarded if they had to be computed from raw data on demand.

The computational operations that generate derived memory are not limited to mere extraction of information from inceptive memory. They may augment it as well. A derived memory can contain information not present in any of the inceptive memories on which it was based, generated by procedures that made inferences from the inceptive data. A derived memory may, in some cases, be the output of the same decision process that uses it at a later time: If one needs to decide at Time 1, on the basis of behavioral episodes, whether Person Y is honest, the conclusion might be stored as a derived memory for use at Time 2. A trait summary is a form of derived memory; a stored episode is an inceptive memory.

Because of their advantages, human memory may be full of independent, derived memory stores linked to various inceptive memory stores (see also Nadel, 1994). The linked memory stores of abstracted personality traits and behavioral episodes that exem-

plify those traits may prove to be an interesting model system that can illuminate properties that are common to many memory stores.

Memory versus computation. The adaptive function of information storage is intrinsically prospective: It is used to support future decisions and judgments, which cannot be known in advance with certainty. To the extent that the character of subsequent decisions and judgments can be predicted, the memory system can be tailored to flag relevant information and precompute variables that are required to make them—that is, it can develop a derived memory store.

However useful it would be to have the answers to all questions that might be asked precomputed and immediately on tap in a derived memory system, this set is indefinitely large and hence precluded by limitations on computational and memory resources. In functional terms, factors that should govern the number and type of derived memory stores include (a) *predictability*: the existence of factors, either ontogenetic or phylogenetic, that allow the successful anticipation of likely varieties of future judgment or decision types; (b) *importance*: the value of those judgments or decisions ought to justify the costs; (c) *urgency*: the requirement that such judgments be made quickly; and (d) *economy*: the number of judgments or decisions supported by a derived memory store should be proportionately large with respect to its size. The use of trait summaries in making decisions about social interactions meets all four criteria.

Storing semantic summaries has both advantages and drawbacks. Direct retrieval of precomputed representations (derived memory) economizes on on-line computation and is fast. But it takes up increasing amounts of memory the larger the set of representations that must be precomputed and on tap, and may provide information that is no longer up to date. In contrast, procedures triggered on line that take inceptive memories as input economize on (derived) memory and provide up-to-date information. But they are expensive computationally and increasingly slow, the more computational steps that are involved. An architecture designed functionally would be shaped by these trade-offs, using different mixes of memory versus computation for different types of adaptive problem (e.g., Pinker, 1997, 1999). Decisions about personality traits provide a case in point.

When Are Trait Summaries Formed?

The evidence on when and how trait summaries are formed, and when they are consulted, is broadly consistent with the functional considerations discussed above. Dynamically, when the data set of relevant events is small, judgments about personality traits appear to be made predominantly on the basis of retrieved behavioral episodes (e.g., Klein et al., 1997; Klein & Loftus, 1990, 1993a, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman & Klein, 1994; but see Klein, Sherman, & Loftus, 1996, for an important qualification). As the behavioral evidence grows, judgments are increasingly made by accessing an emerging trait summary and decreasingly by accessing relevant episodes. When the amount of experience is high, the semantic summary is well developed, and trait judgments are made by accessing this knowledge (e.g., Klein et al., 1997; Klein & Loftus, 1990, 1993a, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992; Schell et al., 1996). This process is faster than retrieving and considering individual instances (e.g., Klein et al., 1989, 1997; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman & Klein, 1994), and it bypasses and

inhibits this other judgment route (e.g., Klein & Loftus, 1993b, 1993c; Klein, Sherman, & Loftus, 1996; Sherman & Klein, 1994).

Evidence suggests that semantic trait summaries are constructed on the basis of the same set of behavioral episodes that are input into the episodic store (e.g., Babey et al., 1998; Sherman, 1996; Sherman & Klein, 1994), although supplementation from other sources of information has not been ruled out (e.g., Klein, Loftus, Trafton, & Fuhrman, 1992; Lord, 1993; Nelson, 1993a). It is not yet clear whether the two stores are constructed in parallel (with summaries updated as new episodes occur) or whether the semantic store is derived from the episodic store after the accumulation of a number of episodes. In any event, neurally impaired individuals who exhibit an inability to retrieve individual events (i.e., access episodic memory) show no apparent impairment in their ability to make accurate trait judgments (e.g., Damasio, Tranel, & Damasio, 1989; M. K. Johnson, Kim, & Risse, 1985; Klein et al., 1999; Klein, Loftus, & Kihlstrom, 1996; Klein, Rozendal, & Cosmides, in press; Schacter et al., 1982; Tranel & Damasio, 1993) and even maintain the ability to revise their trait judgments based on new episodes they cannot recall (e.g., Tulving, 1993a). This suggests that conscious retrieval of behavioral episodes is unnecessary for trait judgments to be made. It further suggests the following: Either updating of semantic information occurs as each episode is experienced (and not later, on the basis of stored episodes) or else the procedures involved in updating do not require conscious access to the episodic store (e.g., Sherman & Klein, 1994; Tulving, 1993a; see also Medin & Bettger, 1994; Riskey, 1979).

Why Store Episodes After a Trait Summary Has Been Formed?

Given the benefits of having a trait summary—speedy access to information, avoiding redundant and costly computation—why have anything else? Once a trait summary is formed—“Tom is rarely cruel,” “I am usually calm”—why not save on memory resources by losing (or overwriting) the episodes on which it was based?

A summary lacks information present in the episodes from which it was derived. Retrieving a trait summary may be faster than constructing a judgment on line from a database of episodes, but it is necessarily less accurate. Maintaining a database of inceptive, episodic memories solves several problems that a trait summary cannot.

Reevaluating Conclusions

New information may cause previous episodes to be reinterpreted, drastically changing one’s judgments of a person or a situation. Fred’s friendly willingness to help you with household repairs may take on different significance if you learn that he is attracted to your wife; Leslie’s willingness to lend you money last year may be reinterpreted when you learn that she had a large and undisclosed financial interest in doing so. If episodes were lost after they had been analyzed to form a summary judgment of Fred’s character, reevaluating his past actions in light of new information about his intentions and values would be impossible.

Keeping a database of episodes is helpful even when a drastic reinterpretation of previous events is not called for. Judgments can be revised in light of new information: If a judgment that “Fred is

usually friendly” was based on 30 episodes, an unfriendly act by Fred should have less impact on the judgment than if it had been based on three episodes (Cosmides & Tooby, 1996; see also Babey et al., 1998; Sherman & Klein, 1994). Without the original database, it is difficult to know whether new, inconsistent information should change one’s summary judgment and, if so, by how much. Moreover, new reference classes can be formed to answer new questions. Suppose you need to decide whether your best friend would make a good employee—something you never considered before. If a database of richly encoded episodes exists, it can be sifted for events relevant to making such a judgment (for discussion, see Cosmides & Tooby, 2000).

Assessing Credal Value

Maintaining source information allows one to evaluate the credal value of people’s statements. Let’s say that Rhonda, who immediately befriended you when you started your new job, told you many terrible things about another coworker, Adam. Much later you find out that she has been stalking him ever since he broke up with her a year ago. As a result, you realize that any unflattering stories you heard from Rhonda may be untrue. In forming an impression of Adam, you integrated information from many sources. One of these sources turned out to be unreliable. How can you update your judgments about Adam? Which of your trait summaries for Adam are still reliable and which are not?

A database of episodic memories would allow you to reevaluate your judgments about Adam. Your “Adam database” would include episodes in which you had interacted with Adam yourself, episodes in which other people told you things about Adam, and episodes in which Rhonda told you Adam stories. A database that preserves these episodes allows you to consider the source: to sort through the database and decide which judgments were based on sound information and which were colored by Rhonda’s lies. Had the episodes on which your judgments of Adam’s character were based been lost, there would be no way to repair the corrupted segments of your semantic store. The ability to judge and reevaluate the credal value of other people’s communications is essential in an organism with language (Cosmides & Tooby, 2000).

Bounding the Scope of Generalizations

An independent store of episodes allows the scope of a summary judgment—the circumstances under which it does, and does not, apply—to be specified. A trait summary such as “He is rarely honest” or “I am usually friendly” gives information about behavior under “average” circumstances, but 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 Demi: “Demi is usually calm.” You are planning what you hope will be a relaxed dinner party with some friends who are political activists of a different stripe than Demi. Access to appropriate episodic memories can delimit the scope of your semantic summary. Recalling that “Demi 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 Demi can be expected to become tense; Wright & Mischel, 1988.)⁶

Episodic Memories As Boundary Conditions on the Scope of Generalizations

This last function leads to a counterintuitive prediction about the relation between semantic and episodic memory when people are called on to make trait judgments. When asked, “Does this describe you: *friendly*?” the decision rule activated should retrieve two kinds of representations: a trait summary from semantic memory and episodes that are inconsistent with the trait one was asked about. For example, “Are you a calm person?” might activate a decision rule that retrieves, “Usually calm (except: the day my tenure was being decided).”

Some traits are invariant over situations: Unless you are Alice in Wonderland, the question, “Does this describe you: 5 ft 9 in. tall?” will have a determinate answer. But personality does not exhibit unwavering cross-situational consistency (e.g., Brody, 1988; Hampson, 1982; Mischel, 1968), and the human cognitive system implicitly recognizes this: People conceptualize personality as conditionally dependent on situations (e.g., Wright & Mischel, 1988).

To rely on semantic memory alone is to sacrifice accuracy for speed: Trait summaries lack the resolution of an episodic memory. Although trait judgments do have some predictive validity, the behavior of an individual manifests a great deal of situation specificity (e.g., Brody, 1988; Funder, 1995; Mischel, 1968). And episodic memories are situation specificity incarnate: They are records of how particular people behaved in specific situations. 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.

The decision context should determine when boundary conditions are needed and which kinds of episodes would be relevant. 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 she or he not?” The decision rule activated is designed to make a judgment about the trait being asked about, whether the question originates with another person or with the individual making the judgment.

In our view, the issue is not, “How are personality traits represented in memory?” because this way of putting the question assumes that the form in which traits are represented is invariant over different decision contexts. We are arguing, instead, that different decision rules are designed to retrieve different kinds of information from memory. The type of decision being made matters, because that determines which search engine gets activated, which will determine which array of memory systems gets searched. When asked to trait-categorize their own behavior, sub-

jects should retrieve a trait summary and episodes inconsistent with the trait they are being asked about. But these representations may be irrelevant in other decision contexts. For example, being asked to define *friendly* does not require a judgment about how often any particular person is friendly. Because no judgment about anyone’s typical behavior needs to be retrieved, there is no need to bound its scope with trait-inconsistent episodes. Therefore, even though defining *friendly* activates representations that share semantic elements with personal trait summaries and personal episodes, one would not expect this decision context to prime either.

In a nutshell, our argument is this:

Every normal adult has a great deal of information about his or her own behavior. We are all “experts” on ourselves and are likely to have formed summary representations for most important traits, which can be retrieved quickly when needed for decision making.

In real life (that is, outside of a psychology laboratory), the question, “Does Trait X describe Person Y?” (where Person Y can be oneself or someone else) is raised when one is considering how best to achieve some outcome, and where the choice of plans depends, in part, on the probability that Person Y will manifest Trait X in Situation Z.

Situations can be categorized in many different ways. A trait summary derived from a large database of episodes (involving many different situations) specifies how frequently Person Y has manifested Trait X across situations. It therefore provides a good default, first-pass answer—one made without prejudging which dimensions of Situation Z might be relevant.

To form a successful plan, one needs to know the scope of this generalization—its boundary conditions. Episodes in which Person Y did not manifest Trait X can provide these boundary conditions. The situations in which Person Y manifested trait-inconsistent behaviors can be compared for similarities with the scenario under consideration—that is, Situation Z.

An “expert” does not need to retrieve supporting evidence for his or her generalizations. The goal is to retrieve accurate information quickly. If one already has a trait summary for Person Y, retrieving episodes in which Person Y manifested Trait X would be redundant, and using them in the decision-making process would delay the decision.

For these reasons, when one is asked to make a categorization judgment about personality traits, the decision rule’s search engine should retrieve the trait summary and trait-inconsistent episodes. It should not retrieve trait-consistent episodes. Trait-consistent episodes are relevant only when one has not yet formed a trait summary for Person Y.

We call this set of propositions the scope hypothesis.

In the experiments discussed below, we used the same priming method as in the previously described experiments by Klein, Loftus, and colleagues. On the basis of the scope hypothesis, we predicted that retrieving a trait summary to answer a describe question would prime trait-inconsistent episodes, but not trait-consistent ones. The first experiment probed knowledge about one’s own traits, because it is fairly certain that the typical adult already has formed a database of trait summaries about himself or herself. The second experiment probed knowledge about another person’s traits—the subject’s mother—to see whether the same results pertain when trait summaries are lacking.

⁶ Implicit in this list of functions is also a set of predictions about which episodes ought to be differentially retained and which might be safely discarded.

Experimental Tests of the Scope Hypothesis

The Self Study

In our first study, we examined whether accessing a trait summary about the self would prime trait-inconsistent episodes but not trait-consistent ones (Klein, Cosmides, Tooby, & Chance, 2001). We used a variant of Klein and Loftus's priming technique (e.g., Klein et al., 1989; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992). Subjects were presented with a list of trait adjectives and asked to perform one of three tasks with them. The describe task asked subjects to decide whether the stimulus trait describes them; the recall task asked subjects to retrieve a specific instance in which they manifested the stimulus trait; the define task asked subjects to generate a definition for the stimulus trait. 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 themselves. These ratings were used to sort subjects' response latencies into three levels of trait self-descriptiveness (high, medium, and low; for details on the sorting procedure, see Klein & Loftus, 1990; Klein, Loftus, Trafton, & Fuhrman, 1992).

If the scope hypothesis is correct, then deciding whether a trait is self-descriptive will activate a summary representation and episodes inconsistent with the trait being asked about. Empirically, then, one should observe two things.

First, for trials on which the initial and target task trait words are antonyms, an initial describe task should be more facilitating than an initial define task to the subsequent performance of a recall task. This is because deciding whether a trait describes oneself activates a trait summary, whereas defining a trait does not (for evidence, see Klein et al., 1997; Klein & Loftus, 1993a, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992). When a trait summary is activated, trait-inconsistent behavioral memories—which are required for performance of the subsequent recall task—should be activated as well. 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 et al., 1989; Klein & Loftus, 1990, 1993b, 1993c; Malt, 1989).

Second, 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 are not activated by a describe judgment (e.g., Babey, Queller, & Klein, 1998; Kihlstrom & Klein, 1994; Klein et al., 1989, 1993, 1997; Klein & Loftus, 1990, 1993a, 1993b, 1993c; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992).

Finally, as noted earlier, work by Klein and Loftus and their colleagues (e.g., Klein et al., 1997; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992; see also Baxter & Goldberg, 1987; Sande, 1990; Sande, Goethals, & Radloff, 1988; Smith, 1990) has shown that trait self-knowledge is represented in summary form across levels of trait descriptiveness. Accordingly, we predicted that the above effects should occur regardless of whether the trait being judged fell in the high, medium, or low self-descriptiveness category.

As Figure 1 shows, the results fully supported these predictions. First, we found strong evidence that trait judgments activated trait-inconsistent episodes: Regardless of level of trait-descriptiveness, subjects were significantly faster to perform the recall task when the initial task was describe than when it was define.

Second, trait judgments did not activate trait-consistent episodes. Replicating findings from previous studies (e.g., Klein et al., 1989, 1997; Klein & Loftus, 1990, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992), the time taken to perform a recall task for trait-consistent episodes was not differentially influenced by the previous performance of a describe task or a define task. Again, this was regardless of whether traits fell in the high, medium, or low descriptiveness categories.

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, subjects took significantly less time to recall a behavior when it exemplified the antonym of the trait judged during an initial describe task than when it exemplified the same trait as that judged during an initial describe task.

In other words, we found that deciding whether a trait describes oneself facilitates retrieval of trait-inconsistent episodes, which suggests that these judgments activate memories of behaviors inconsistent with the trait being judged (Klein et al., 2001). By contrast, we found no evidence that trait self-descriptiveness judgments facilitated retrieval of trait-consistent episodes. From this

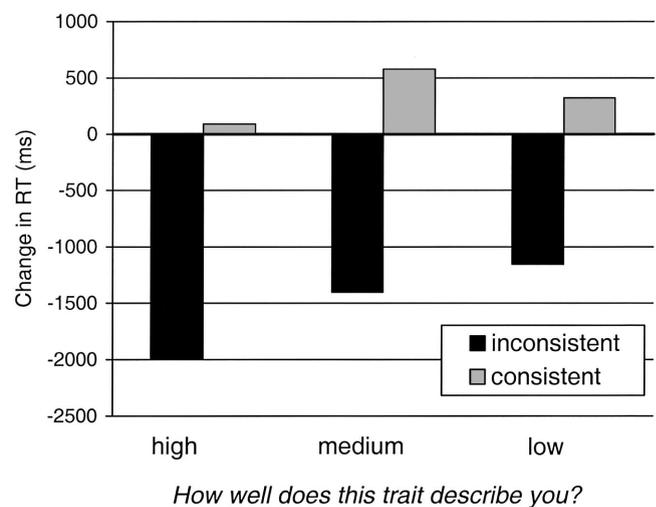


Figure 1. Activating a trait summary about the self can prime episodic memory: self-study. Deciding whether a trait describes oneself activates episodes that are inconsistent with the trait asked about. Trait-consistent episodes are not primed. This holds whether the subject views the trait as highly descriptive or not. Negative numbers indicate that the episode was recalled faster than in a control condition. RT = reaction time.

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.

The Mother Study

Our second study (Klein et al., 2001) 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 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, 1993b; 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 below, 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. Earlier, we described the Klein and Loftus model for the representation of trait knowledge about the self. Recently, Klein, Loftus, and colleagues proposed that this same model can be extended to describe the mental representation of trait knowledge about others (e.g., Babey et al., 1998; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992; 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 inceptive behavioral memories—that is, 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 derived memory representation—a trait summary. Thus, the model proposes that the more 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 similar views, see N. H. Anderson, 1989; Park, 1986).

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, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992, Experiment 1; 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 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, 1993b) 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 derived memory—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 derived memory in the form of 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 subsequent recall task, and one would observe faster retrieval of trait-consistent episodes following a describe task. (The reason for assuming that retrieval will be biased toward trait-consistent behaviors follows from research showing that people rely on "positive-test" strategies in many tasks [e.g., Klayman & Ha, 1987; Wason, 1968], including when answering 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.)

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 prediction can be made for low-descriptive traits; for reasons, see Klein et al., 2001; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992.)

Testing the scope hypothesis when trait summaries are present and absent: Predictions. The results just described show that the mental representation of mother's traits varies with trait descriptiveness: Knowledge of highly descriptive traits consists of derived memories—abstract trait generalizations—whereas knowledge of medium-descriptive traits consists of inceptive memories of specific behavioral episodes. For exactly this reason, however, the scope hypothesis makes different predictions for mother's high- and medium-descriptive traits. When a trait summary exists (as it does for highly descriptive traits), the results from a decision task involving one's mother should exactly parallel those from the self study: Performance of an initial describe task should facilitate the recall of trait-inconsistent behaviors, but not trait-consistent ones.

But when there is no trait summary whose scope needs to be delimited (as for mother's medium-descriptive traits), trait-inconsistent episodes should not be primed (i.e., an initial describe task should not facilitate the recall of trait-inconsistent behaviors).

Instead, one should expect priming of trait-consistent behavioral episodes because these are necessary for making a judgment when no summary exists.

Our data confirmed these predictions (see Figure 2). For traits highly descriptive of one's mother—ones for which a summary should exist—trait-inconsistent episodes were primed, but trait-consistent ones were not. For traits only moderately descriptive of one's mother—ones for which there should be no summary—the opposite pattern was found. Trait-inconsistent episodes were not primed, but trait-consistent ones were.

Because the mother results allowed us to compare priming when trait summaries are present versus absent, they provide strong confirmation of the scope hypothesis. When a decision is made by accessing a trait summary, search engines should simultaneously retrieve inconsistent episodes to provide information about the conditions under which the summary ceases to be useful. Thus, trait-inconsistent behaviors were primed when a decision about one's mother was made by accessing summary representations of her traits (the highly descriptive condition). But when summary representations were absent—as in the medium-descriptive condition—they could play no role in decision making. This explains why trait-inconsistent behaviors were not primed in the medium-descriptive condition. According to the scope hypothesis, when no generalization has been accessed during a decision task, there is no representation whose scope needs to be delimited.

We think the results of the self and mother studies are particularly interesting because they are naturalistic: The subject made trait judgments about real individuals and real experiences. It bears mentioning, however, that priming of trait-inconsistent behavioral episodes following activation of a trait summary also has been found when subjects make judgments about meaningful groups as well as artificial target individuals created in the laboratory (Babey

et al., 1998). Indeed, Babey et al. (1998) found this effect to be highly replicable—it was confirmed in six separate studies and never disconfirmed.

General Discussion

A great deal of memory research is designed to illuminate the design features of memory systems, including their relationships to one another. It has been clear for a long time that processes within a memory system can operate independently of one another: For example, damage to the hippocampus selectively impairs the processes that encode and store new information, leaving retrieval processes intact (e.g., Graham & Hodges, 1997; Squire & Knowlton, 1995; Staubli & Lynch, 1987). It has recently been suggested that relationships between different memory systems are process specific as well: According to Tulving's (1995b) serial-parallel-independence (SPI) model, there is a serial relationship between semantic and episodic memory when it comes to encoding (episodic encoding depends on prior semantic encoding), a parallel relationship when it comes to storage, and functional independence when it comes to retrieval.

We have proposed that this insight about process specificity be carried one step further, to recognize that not all retrieval processes are the same. Each decision rule should have a search engine that is designed to retrieve adaptively appropriate information. Because each decision rule has a different adaptive function, each will have different informational needs: One may require information from semantic memory alone, another from episodic memory alone, and yet another from both. As a result, retrieval from semantic and episodic memory should be independent in some decision contexts but not in others. When it comes to retrieval, there should be no invariant relationship between the semantic and episodic memory systems.

More specifically, we 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. Both speed and accuracy can be engineered into a decision-making system, however, if the system 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 are broad-band encodings of 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, a decision rule needs a search engine that looks for summary information in semantic memory and, on retrieving it, also looks 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.

In support of this argument, we discussed data from experiments on trait judgment. A large body of existing research had already shown that when asked to decide whether a trait describes oneself (e.g., "Does this describe you: *friendly*?"), subjects retrieve a derived memory—a trait summary—from semantic memory. It was also known that, in this decision context, retrieving a trait summary does not activate episodic memories in which the subject manifested behavior that is consistent with the trait under consideration. Therefore, when it comes to behavioral episodes whose

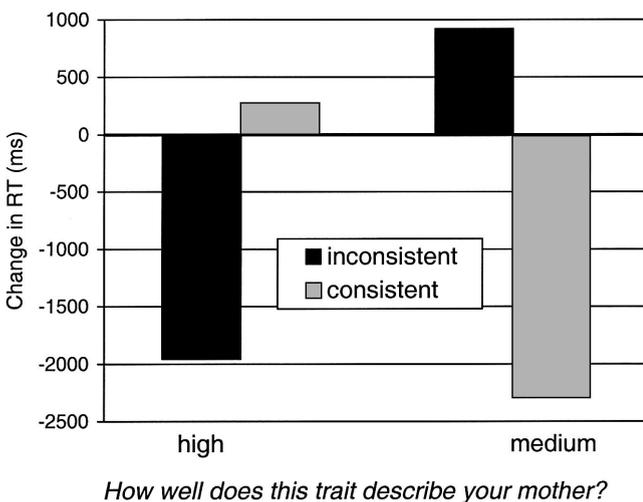


Figure 2. Activating a trait summary about another person primes inconsistent episodes: mother study. Deciding whether a trait describes one's mother activates episodes that are inconsistent with the trait asked about only when the trait is highly descriptive of one's mother. When it is only moderately descriptive, trait-consistent episodes are primed. This is consistent with the notion that subjects have formed summaries only for traits they consider highly descriptive of their mothers. Negative numbers indicate that the episode was recalled faster than in a control condition. RT = reaction time.

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.

We found, however, that this result does not generalize to trait-inconsistent episodes. On the basis of the scope hypothesis, we made the following prediction: When a decision rule designed to make trait judgments about a person succeeds in retrieving a trait summary from semantic memory, it will also retrieve episodic memories of incidents in which that person exhibited behaviors that are inconsistent with that trait. This is because the decision rule is designed to place boundary conditions on the scope of the generalization. We tested this prediction by using a priming paradigm and a decision task in which one is asked whether a trait describes oneself. 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 decision contexts: Asking a subject to simply define a trait—a context that does not cause retrieval of a trait summary (Klein et al., 1989, 1997; Klein & Loftus, 1993b; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992; Schell et al., 1996)—primes neither trait-consistent nor trait-inconsistent episodes.

We demonstrated the generality of this hypothesis by showing that it extends to judgments of others, in this case the subject's mother. Previous research indicates that subjects do have summary representations for traits they consider highly descriptive of their mothers but not of ones they consider only moderately 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. Trait-consistent episodes were primed only for medium-descriptive traits—ones for which the subject appeared 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 inceptive memories, following a positive test strategy (as they do in many domains; Klayman & Ha, 1987).

In short, by considering when a decision 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.

Is Priming Functional?

Tulving (1995b) has argued that perceptual priming is the expression of an adaptation designed to facilitate the identification of “the same or a similar object on a subsequent occasion, in the sense that the identification of the object requires less stimulus information or occurs more quickly than it does in the absence of priming” (p. 841). J. R. Anderson and Milson (1989) also view priming as functional. Just as vision scientists once used the design features of a camera to illuminate the design features of the eye, J. R. Anderson and Milson perceptively used the design features of library systems and computerized databases to illuminate the design features of priming in memory. For example, in libraries and computerized data retrieval systems, the best predictor of whether a

piece of information will be useful in the future is the fact that it was recently used (e.g., books that were recently checked out of the library are more likely to be checked out again than books that have not been checked out for a long time). For this reason, libraries and computer databases are organized such that recently accessed items of information are easier to retrieve than ones that have not been accessed for a long time. J. R. Anderson and Milson proposed that priming—whether perceptual or memorial—is the expression of an adaptation designed to promote the fast retrieval of useful information by exploiting the basic predictive relationship in the world between the frequency and timing of past use and the probability of future use.

We strongly agree that the phenomenon of recent exposure speeding recognition or retrieval is an adaptation that has evolved because of its functional consequences. We began by proposing a functional explanation for the observed phenomenon of priming. According to this view, priming is the result of the designed expediting of information or accessibility to mechanisms on the basis of the probability that the class of information will be useful, and we view prior use as simply one cue of future utility. This raises the possibility that there is a far larger class of computational mechanisms that evolved to cause information expediting than those that operate by prior exposure: Each cue that reliably predicts that a certain kind of information will be useful to a certain kind of decision rule or inference system potentially could serve as the basis for the evolution of a new mechanism for accelerating the accessibility of that category of information to the procedures that require it. Measurement of such facilitation provides the experimental entry point into exploring the potentially heterogeneous, functionally specialized design features of these search engines. In the self and mother experiments we demonstrated the existence of at least one additional variety of priming that operates on an entirely different—yet still functionally advantageous—principle: scope limitation.

Neither of the functions outlined by Tulving, 1995b, or J. R. Anderson and Milson (1989) are adequate for explaining the priming of trait-inconsistent episodes that we found in the experiments discussed herein. Unlike the priming of object representations (Tulving, 1995b), no perceptual processes were involved in our experiments, and our results do not parallel those for perceptual priming (we did not find that “recognizing” oneself as, say, calm speeds recall of instances in which one was calm). The J. R. Anderson and Milson analysis does not apply in this case, either. The finding was not that retrieving a trait summary primes the subsequent retrieval of that same trait summary (although this is, in fact, true; e.g., Klein et al., 1989; Klein, Loftus, Trafton, & Fuhrman, 1992). Nor did we find that retrieving a trait summary primes the subsequent retrieval of trait-consistent episodes—episodes that might well have been recently used in the sense that they contributed to the formation of that summary. The only episodes primed by retrieving a trait summary were ones inconsistent with it. This finding does not fit the J. R. Anderson and Milson profile: It is like finding a library where checking out *The Origins of Virtue* predicts that biographies of Hitler (but not of Mother Theresa!) will be checked out next. In other words, if the priming found in our experiments is the expression of the functional component of an adaptation, we will have to look elsewhere for its explanation.

Is Priming an Adaptation or a By-Product?

These experiments addressed another issue in the priming controversy. Although some researchers regard priming as an evolved adaptation, others regard it as an engineering by-product. For example, the physiological activation of neural tissue could simply lower the threshold for its subsequent activation purely as an accident of cell biology. Indeed, one of the basic principles of adaptationism—the study of the functional design of naturally selected systems—is that not all phenotypic features are functional. Many features of organisms are by-products of functional organization (Gould & Lewontin, 1979; Williams, 1966).

In J. R. Anderson and Milson's (1989) theory, there was a sophisticated matching of data to functional theory. Nevertheless, because by-product explanations are often vague, and can make the same qualitative predictions as the J. R. Anderson and Milson theory, the question persists: Is episode priming a functional component of an adaptation or is it a by-product of one?

If one can demonstrate that priming in another domain takes place according to an entirely different functional logic—one that is inconsistent with any simple physiological or engineering by-product explanation—this would undermine these by-product explanations as counterhypotheses to any functional theory, including those proposed by Tulving (1995b) and by J. R. Anderson and Milson (1989). It would support the hypothesis that all three types of priming are adaptations, and independent ones at that.

Design evidence is criterial. To demonstrate that an aspect of the phenotype is an adaptation, one must first provide design evidence: evidence that the phenotypic feature is well-engineered for solving an adaptive problem. The criteria are similar to those one would use to demonstrate that a human-engineered machine is (e.g.) a TV rather than a stove (a TV has many features arranged in a way that is well suited for transducing electromagnetic radiation into a color bit map—a configuration that is unlikely to have arisen by chance—but few features that would suit it for cooking food). The experiments herein took a first step in that direction. Trait judgment is an adaptive domain in which a generalization can provide fast information to a decision rule, but only by sacrificing information about situation specificity that could improve performance. For adaptive problems of this kind, a well-engineered decision rule—one that can achieve both speed and accuracy—would have search engines that retrieve, along with a generalization, episodes that place boundary conditions on its scope. The experiments we reported demonstrated that trait judgments show just this property: When a trait summary is retrieved, inconsistent episodes, but not consistent ones, are primed. The pattern of results is what one would expect of a priming mechanism designed to delimit the scope of a generalization.

However, providing design evidence is not enough to demonstrate that the machinery involved is an adaptation designed to solve the scope problem. One must also demonstrate that the phenotypic feature in question is not more parsimoniously explained as the by-product of some other adaptation or as a manifestation of incidental features or physical law. To continue the TV analogy, there are no machines designed for sewing, cooking, or any other purpose that just happen, by coincidence, to transduce TV waves into displayed bit maps (leading to the confident rejection that objects that capture and display TV signals do so as a by-product). And TVs—highly entropic collections of matter—do not fly together as the result of Brownian motion, osmosis,

gravity, or any other known physical process or law. Can the same be said for the machinery that caused priming of trait-inconsistent episodes in the self and mother experiments?

Is priming a by-product of engineering constraints? The hypothesis that priming is a by-product of some general engineering constraint of the brain is easily eliminated. Many kinds of stored representations are found in the human brain—from perceptual representations of objects to lexemes to episodic memories—and it is clear that there are many situations in which these are not primed, either by previous exposure or by retrieval (and when priming occurs, one usually sees faster retrieval of the same representation, not of its “opposite”). For example, changing an object's picture-plane orientation impedes priming of perceptual representations of objects, but changing its size does not (Biederman & Cooper, 1992; Cooper, Schacter, & Moore, 1991; Cooper, Schacter, Ballesteros, & Moore, 1992). Perhaps more importantly, the evidence presented herein shows that episode priming after retrieval of a trait summary cannot be a physiologically necessary phenomenon of recent neural firing, because trait-consistent episodes are not primed.

Is priming a by-product of neural activation (spreading activation hypothesis)? Because this is a type of engineering constraint hypothesis, the evidence above speaks against any simple version of a neural activation hypothesis as well.

Could episode priming be a by-product of spreading activation among semantically related items in a neural net? On this view (e.g., J. R. Anderson, 1976, 1983; Collins & Loftus, 1975; McNamara, 1992a, 1992b), retrieving a trait summary would activate other representations with trait-relevant semantic content. Our evidence is not consistent with a spreading activation model. Almost by definition, trait-consistent episodes are semantically closer to a trait summary than inconsistent ones, yet retrieving a trait summary does not prime consistent episodes.

Is priming a by-product of connectionist systems of trait representation? Pattern associator models of memory and cognition have received a lot of attention in the past two decades (McClelland & Rumelhart, 1986; Quinlan, 1991, 1992; Rumelhart & McClelland, 1986). Moreover, they have many appealing features as models for how at least some memory processes might work, including the creation of the derived summary representations explored herein. This raises the question of whether the priming patterns we have reported could be explained as a natural by-product of this kind of architecture.

It is difficult to see how this could be so. In essence, such architectures operate by computing similarities. Although trait-inconsistent episodes share similarities with the summary representation, they obviously share fewer similarities than trait-consistent episodes do, and therefore, trait-consistent episodes should be primed by summary representations to a greater degree than trait-inconsistent episodes. We did not observe this. Moreover, it is difficult to see how the same pattern associator model could produce priming of trait-consistent episodes when subjects are asked about medium-descriptive traits of others, but trait-inconsistent episodes when asked about highly descriptive traits. Although it would be premature, given the great diversity of existing and potential connectionist models, to conclude that it is impossible to contrive connectionist systems that could produce such outcomes, it seems safe to suggest that this pattern of results is unlikely to be a natural by-product of any pattern associator

model not specifically built to incorporate something similar to the functional principles of priming outlined here.

Is priming a by-product of a fan effect? According to fan effect models, a representation must receive a certain amount of “activation” if it is to be consciously retrieved (e.g., J. R. Anderson, 1974, 1976, 1983; Moeser, 1979; Thorndyke & Bower, 1974). A proponent of fan effect models might try to explain the results as follows: Suppose that, for traits that are highly self-descriptive, one’s episodic memory contains more trait-consistent episodes than trait-inconsistent ones. Assume further that when one is asked to retrieve a certain kind of episode, a given amount of “activation” is spread among all of the exemplars. Let’s say the episodic store contains 100 trait-consistent episodes and 10 trait-inconsistent episodes. When asked to retrieve a trait-consistent episode, each incident would acquire 1/100 of the total activation; when asked to retrieve a trait-inconsistent episode, each incident would acquire 1/10 of the total activation. Because each trait-inconsistent episode would have 10 times the activation of each trait-consistent one, trait-inconsistent episodes are more likely to be primed. This would explain why, when subjects are asked to retrieve episodes relevant to traits that are highly self-descriptive, trait-inconsistent episodes, but not trait-consistent ones, are activated. It could also explain why one finds a different priming pattern in the mother experiment. The subject’s memory for episodes involving mother would include many fewer trait-consistent episodes for medium-descriptive traits than for highly descriptive ones. Hence, when the subject is asked about highly mother-descriptive traits, trait-inconsistent episodes should be primed. But trait-consistent episodes should be primed when he or she is asked about medium mother-descriptive traits.

There are a number of problems with this account. First and foremost, there is strong empirical evidence against fan effect models, especially for trait judgment (Babey et al., 1998; Sherman & Klein, 1994). Second, this account depends on the assumption that trait-consistent episodes vastly outnumber trait-inconsistent ones in the episodic memory system. Although this might be true for highly self-descriptive traits, it is unlikely to be true for medium and low descriptive traits. After all, subjects judged that some of the traits were “extremely unlike me” (low descriptive). For traits “extremely unlike me,” the person’s episodic store should contain more trait-inconsistent episodes than trait-consistent ones. Fan effect accounts would have to predict priming of trait-consistent episodes for low descriptive traits, and priming of trait-inconsistent episodes for highly descriptive traits. Such predictions were falsified: Inconsistent episodes were primed, and consistent episodes were not primed, no matter how self-descriptive a person judged a trait term to be.

A third reason to doubt this account is that it does not explain why priming of inconsistent episodes occurs when the subject makes a describe judgment but not when the subject performs other semantically related tasks (such as a define task) or, indeed, no initial task at all (Babey et al., 1998; Klein et al., 2001). To account for the results, the fan effect model must posit that the burst of activation is sent when subjects face the second task, which is the recall task, and not before. This is what ensures that the burst is sent only to (e.g.) friendly episodes, and not to unfriendly ones, when subjects are asked to recall episodes in which they were friendly. (Otherwise, the burst would be spread among all n semantically related episodes, inconsistent or not; each incident would get a charge of $1/n$, and inconsistent episodes

would have no activation advantage.) This means the initial task must be irrelevant. But if the first task is irrelevant, then the fan effect cannot explain why the priming of inconsistent episodes for the self occurs after a describe task, but not after a define task. Indeed, proponents of the fan effect would have to predict that, for highly descriptive traits, inconsistent episodes would be recalled faster than consistent ones even if there were no initial task at all. After all, if there are 10 inconsistent episodes and 100 consistent ones, the charge on each inconsistent one would be 10 times greater than the charge on each consistent one, regardless of any task the subject performs before being asked to recall an episode. Yet inconsistent episodes are not recalled faster than consistent ones in the absence of an initial task (e.g., Babey et al., 1998; Klein et al., 2001).

Lastly, the data on mother judgments are not consistent with the fan effect either. Indeed, it is difficult to reconcile the mother results with the self results given that view. Let’s say, for example, that trait-consistent episodes are primed in the mother experiment for medium-descriptive traits because for these traits, there are many inconsistent episodes and few consistent ones (otherwise the trait would be judged “highly descriptive of mother”). If so, then the same should be true for the medium descriptive traits in the self study, and we would have seen priming of trait-consistent episodes for those traits as well. But this was not the case: Trait-inconsistent episodes were primed for all the self judgments, regardless of how self-descriptive the subject judged the trait queried to be. One could escape this contradiction by assuming that, compared with episodes about the self, the subject has stored many fewer episodes—whether trait-consistent or not—about his or her mother. But if there are a small number of episodes, then each should be able to get a sizable dose of activation. As a result, both trait-consistent and trait-inconsistent episodes would be primed in the medium-descriptive mother condition. Yet trait-inconsistent episodes were not primed at all in that condition.

For these reasons, we conclude that the priming we found in these experiments is the functional product of evolved adaptations designed to expedite information accessibility rather than a by-product of the cognitive architecture. The fact that all of the by-product explanations we considered fail to account for our data has wider implications, however. Existing by-product explanations, by their nature, predict priming in all domains, regardless of content. They contain no mechanism for producing priming of, say, object representations, but not of trait episodes. Therefore, if they fail to account for one domain, they fail as explanations and can be eliminated as viable counterexplanations for the functional theories of priming proposed by J. R. Anderson and Milson (1989), Tulving (1995b), and ourselves.

Conclusions

Toward a General Model of Trait Knowledge

The relation between trait summary knowledge and summary-consistent behavioral memories. The results from the present studies, taken together with the supporting body of preceding research (e.g., Babey, et al., 1998; Craik et al., 1999; Hirshman & Lanning, 1999; Kircher et al., 2000; Klein & Kihlstrom, 1998; Klein & Loftus, 1990, 1993a, 1993b; Klein et al., 1989, 1993, 1997; Klein, Loftus, & Kihlstrom, 1996; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein, Sherman,

& Loftus, 1996; Lord, 1993; Schell et al., 1996; Sherman et al., 1998; Sherman & Klein, 1994; Tulving, 1993a) allow construction of a framework of inferences concerning the relation between summary-consistent behaviors and trait summary knowledge.

First, this body of research supports the view (e.g., Kihlstrom & Klein, 1994; Klein, 1999; Klein & Kihlstrom, 1998; Klein & Loftus, 1993b; Klein, Loftus, & Kihlstrom, 1996) that decision rules designed to render judgments about social entities (self, others, groups) recruit information from two distinct types of memory: an *episodic* store, which represents specific events and behaviors involving the target, and a *semantic* store, which includes summaries of the target's personality traits abstracted from the set of particular events that served as input to the procedures that generated the trait summary (see also Schank, 1982; Tulving, 1993a). Trait knowledge based on individual events and trait knowledge consisting of trait summaries are independently represented and independently used (e.g., Craik et al., 1999; Hirshman & Lanning, 1999; Kihlstrom & Klein, 1994; Klein et al., 1989, 1997; Klein & Loftus, 1990, 1993b; Klein, Loftus, & Kihlstrom, 1996; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992), although under certain circumstances both may be deployed during a given judgment (e.g., Babey et al., 1998; Klein et al., 2001; Klein, Sherman, & Loftus, 1996).

Second, once a summary representation has formed, judgments about a person's traits typically are made on the basis of this abstract knowledge (e.g., Babey et al., 1998; Budesheim & Bonnelle, 1998; Klein et al., 1989, 1993, 1997; Klein & Loftus, 1993a, 1993b; Klein, Loftus, & Plog, 1992; Klein, Loftus, Trafton, & Fuhrman, 1992; Klein, Sherman, & Loftus, 1996; Lord, 1993; Sherman & Klein, 1994). In fact, given the existence of a summary trait representation, specific memories in which the person manifested behavior consistent with a trait summary appear to play little, if any, role in trait judgments (e.g., Buss & Craik, 1984; Chaiken & Baldwin, 1981; Epstein, 1979; Klein et al., 1993; Klein & Loftus, 1993b; Klein, Loftus, Trafton, & Fuhrman, 1992; Schell et al., 1996; Tranel & Damasio, 1993; but see Klein, Sherman, & Loftus, 1996, for an important qualification).

Thus, when it comes to retrieving behavioral episodes whose content is summary consistent, retrieval from the episodic and semantic stores appears to be functionally independent. By this we mean that the operations of the semantic store do not require the operations of the episodic store (e.g., Kihlstrom & Klein, 1994, 1997; Klein et al., 1997; Klein & Loftus, 1993b; Klein, Loftus, & Kihlstrom, 1996; Tulving, 1983).

The relation between summary-inconsistent behavioral memories and trait summary knowledge. The present research indicates that these conclusions do not generalize to summary-inconsistent behavioral episodes. Specifically, we have shown that when a judgment retrieves a trait summary from memory, it also retrieves memories of specific instances in which behavior was inconsistent with that summary (Babey et al., 1998; Klein et al., 2001). Thus, absolute statements about the independence between the episodic and semantic stores may not be possible: One finds independence in some domains but not in others.

Is the Priming of Inconsistent Events Restricted to the Domain of Trait Judgment?

Perhaps because it seems so counterintuitive, the possibility that trait judgments call to mind memories of trait-inconsistent behav-

iors had not been examined empirically before the studies reported herein. However, theoretical developments in three different research domains—script-based learning (e.g., Schank, 1982), schema-driven memory (e.g., Graesser, 1981), and object categorization (e.g., Nosofsky, Palmeri, & McKinley, 1994)—hint strongly at the possibility that similar phenomena will be found in domains other than trait judgment.

The relation between summary knowledge and memory for specific episodes is a central theme of Schank's work on script-based comprehension and learning (e.g., Schank, 1980, 1982; Schank & Abelson, 1977). Schank proposed that when events of a similar class repeat themselves a sufficient number of times, they are assimilated into a knowledge structure he called a *script*. Scripts store generalizations abstracted from repeated experiences of similar events but fail to preserve the specific details of any particular occurrence of those events. They do, however, preserve information about specific events that deviate from the expectations generated by the script.

The special status Schank accorded expectancy-violating events in the representation of script knowledge derived from their role in his theory of failure-driven learning (e.g., Schank, 1980, 1982). According to Schank, the generalizations contained in a script enable a person to know what to expect and how to behave when he or she encounters similar events. However, although these expectations usually are accurate, they occasionally turn out to be wrong. When this happens, it is important that the learner remember the events that led to the failed expectations so better predictions can be made next time the same situation arises. Scripts accomplish this by storing a record of the specific events that caused the failure, which is activated along with general script knowledge when similar situations are encountered. Schank's reasoning was very similar to that which led us to propose the scope hypothesis to explain why trait-inconsistent episodes are primed when a trait summary is retrieved. In our terms, scripts would be a type of derived memory, and the expectancy-violating events would be inceptive memories that are coretrieved to bound the scope of the script's generalizations. When retrieved together, they provide a good combination of speed plus accuracy.

Expectancy-violating events also receive special treatment in several schema models of memory. For example, Bartlett's (1932) "schema with correction" model proposes that the mental representation of a class of similar events consists of a set of generalizations called a *schema* along with which are stored specific deviations from the information contained in the schema (see also Woodworth, 1938). A more recent adaptation of these ideas was Graesser's "schema pointer plus tag" model (e.g., Graesser, 1981; Graesser, Gordon, & Sawyer, 1979; Graesser, Woll, Kowalski, & Smith, 1980; Smith & Graesser, 1981; see also Hudson, 1986). According to this model, a schema contains summary information about the typical attributes and relations that exist among a class of similar events plus a set of tags that index specific episodes inconsistent with, or atypical of, the summary knowledge contained in the schema. These tagged episodes are retrieved when the schema is activated in memory. Again, the schema is a type of derived memory, and the inconsistent episodes are inceptive memories that bound its scope.

Some recent models of category learning also posit a special status for expectancy violations in the representation of category knowledge. For example, the RULEX model proposed by Nosofsky et al. (1994; see also Palmeri & Nosofsky, 1995) claims that

people form abstract, logical rules for categorizing objects through a process of hypothesis testing. However, when an instance occurs that cannot be predicted by those rules, it is stored in memory along with the general rule. People access the rule, along with any exceptions, when making subsequent categorization decisions (for related ideas, see Chumbley, 1986).

Although we outlined the scope problem in the context of trait judgment, it arises in many domains (for discussion, see Cosmides & Tooby, 2000). In the first three cases discussed above, the joint retrieval of a generalization with exceptions to it can be seen as a solution to the scope problem. The generalization provides fast access to information with some predictive validity, and the exceptions enhance accuracy by providing boundary conditions on the scope of the generalization.

One would not expect to see the same pattern across all domains, however. For example, models of how native English speakers learn the use of the past tense (e.g., Brown, 1973; Cazden, 1968; Kuczaj, 1977; Pinker, 1994) bear some interesting similarities to the memory models we have been discussing. According to these models, children hear examples of grammatically correct language (*need-needed, help-helped*), from which they infer a general past-tense rule (i.e., to make past tense, add *ed* to the stem). They next pass through a phase in which they occasionally make errors by overgeneralizing the rule to newly acquired irregular verbs (e.g., *go-goed, catch-catched*). These errors occur when the child fails to retrieve from memory an irregular past tense form (Marcus et al., 1992). In first-language learners, the memory retrieval failure occurs because the child has heard the new word so few times, but such failures can also occur in adults (especially those with organic memory impairments): When they do, the adults over-regularize as well (Ullman et al., 1997).

Appropriate use of the past tense eventually is accomplished when children learn to store in memory both a general past tense rule as well as a record of specific exceptions to that rule (Pinker, 1991, 1999). These exceptions delimit the scope of the past-tense rule. But note that the order of activation goes in the opposite direction from that for trait judgment. Activation of trait summaries activates inconsistent episodes, but activation of the (regular) past-tense rule does not activate irregular exceptions. Instead, activation of an irregular exception blocks the activation of the past-tense rule (Aronoff, 1976). This difference makes sense: In trait judgment, the summary is activated because one is using it to make predictions, but no adaptive function is served by activating a past-tense rule when an irregular past-tense form has already been retrieved.

One should also expect a different pattern for generalizations that are always true and for situations in which exceptions do not illuminate a generalization's boundary conditions. Inconsistent episodes should not be primed by retrieval of generalizations that have always been true over evolutionary time and that have come to be embodied in neural specializations. For example, even at 3 months of age, infants assume that two solid objects cannot pass through one another (Baillargeon, 1986; Spelke, 1990), and sleight of hand is required to make it appear otherwise. But it would be very strange to find that whenever an adult sees two solid objects approach one another—a softball approaching a bat or two cars on a collision course—episodic memories of expectancy-violating magic acts are primed. By the same token, one would not expect retrieval of generalizations to prime inconsistent episodes in domains where, during human evolutionary history, it would not

have helped: Chaotic systems, such as weather (in the absence of modern meteorology), might be an example. Exceptions may lower your confidence that a predicted event will occur, but they do not illuminate the conditions under which the generalization does not hold.⁷

Implications for Studying the Architecture of Memory

Dissociations in retrieval have been used as evidence in support of the existence of multiple memory systems (for reviews, see Foster & Jelicic, 1999; Polster et al., 1991; Schacter & Buckner, 1998; Schacter & Tulving, 1994; Squire, 1992). Do the results of the experiments reported herein mean that we need to abandon the idea of separate memory systems? After all, we have shown that there are decision tasks in which retrieval of episodic knowledge is linked to, rather than independent of, retrieval of semantic knowledge. Does this finding mean that semantic and episodic memories do not reside in two separate memory systems? (For debate, see Squire, 1994; Tulving, 1993b.)

No. In particularly severe cases of the classic amnesic syndrome, patients can still retrieve semantic knowledge, but episodic memories cannot be retrieved no matter what decision task is employed (e.g., Cermak, 1984; Kitchener, Hodges, & McCarthy, 1998; Klein, Loftus, & Kihlstrom, in press; Markowitsch et al., 1993; Milner, Corkin & Teuber, 1968; Tulving, 1989, 1993a; Tulving, Hayman, & MacDonald, 1991; Vargha-Khadem et al., 1997; Wilson & Wearing, 1995; Wood, Brown, & Felton, 1989). That access to episodes can be knocked out across decision contexts, while semantic access is spared, implies that there are, indeed, separate “archives” for episodic and semantic memories.

The results we have presented imply a more complex view. Imagine a library that has a number of different collections, each of which archives a different category of information. It would be surprising if every scholar who uses the library were to access one, and only one, of these collections. This is because each scholar is trying to solve a different problem: The scholar studying sexual mores in Victorian England would be retrieving items from both the special collection on Victorian erotica and general histories of England; the labor economist might only retrieve summary statistics on employment compiled by government sources; the biographer might retrieve only letters from the library's collection of correspondence between John and Abigail Adams. Whether retrieval is from just one archive or from several would vary depending on the purpose of each investigator. The situation in a brain-intact individual is analogous: Whether retrieval is from just the semantic archive, just the episodic archive, or both should vary depending on the adaptive problem that each decision rule that accesses stored information was designed to solve.

This view also carries implications for understanding people who have experienced brain damage. Although some conditions may eradicate access to an entire archive, other conditions might spare the archive while damaging a particular decision rule or

⁷ For example, remembering episodes in which rain did not come, despite the presence of dark clouds, may lower your confidence in dark clouds as a good predictor of rain. But such episodes are unlikely to illuminate the conditions under which the generalization “dark clouds predict rain” fails to hold. (By contrast, if one were trying to decide whether dark clouds are a good predictor of rain, one might want to recall both examples and counterexamples.)

search engine that uses it (Eustache et al., 1999). Returning to the library analogy, instead of losing a particular collection to a fire, one investigator who uses a collection might be harmed, while a different investigator, who uses that same collection, might emerge from the fire unscathed. This raises the possibility of complexly patterned amnesias, where a memory archive can be accessed in some decision contexts but not in others. There do appear to be patterns of brain damage that disrupt a person's ability to make decisions in some domains but not in others (e.g., Cappa, Frugoni, Pasquali, Perani, & Zorat, 1998; Damasio, 1994; McCarthy & Warrington, 1992; Sirigu & Grafman, 1996; Suzuki, Yamadori, & Fugii, 1997; Warrington & Shallice, 1984), lending credence to this possibility. There are also cases in which brain damage creates very content-specific patterns of nonretrieval for lexical knowledge—patients who (for example) can name and describe animals and artifacts but not fruits and vegetables (for review, see Caramazza & Shelton, 1998). Some category-specific deficits might result from the destruction of a content-specific archive (of animal names, for example; see Caramazza, Hillis, Leek, & Miozzo, 1994). But others may involve damage to decision rules that access these archives. The case of M.D., who had difficulty naming fruits and vegetables, but not other categories, is suggestive: Although M.D. could not (for example) name a peach when shown one, when asked, "Show me the peach," he could point to the correct fruit (Hart, Berndt, & Caramazza, 1985). That he could recognize the word indicates that the lexical entry was intact, even though he could not access it in other decision contexts.

A view that emphasizes the adaptive function of decision rules requires that care be taken in analyzing what these functions might be. Restricting our attention once again 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.

There may, in fact, be many social contexts in which a well-engineered response system would have ready access to episodes that contradict—or support—the trait or generalization under discussion. A system that is well engineered for navigating the social world might be designed to prime different kinds of episodes, depending on the type of social negotiation at hand. By analyzing such contexts, one should be able to make principled predictions about the circumstances under which both consistent and inconsistent episodes will be primed.

To do so, however, we must change the way we think about memory in two fundamental ways. First, we must abandon the assumption of function generality that has typified the memory literature. Second, we must remember that decision rules had to coevolve with search engines and memory systems, and that what kind of information they will retrieve will depend on what adaptive

problem they were designed to solve. In short, we must be willing to consider the possibility that natural selection has created machinery that can elegantly solve a far wider array of adaptive problems than we have yet been able to imagine.

References

- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 6, 451–474.
- Anderson, J. R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 22, 261–295.
- Anderson, J. R. (1989). A rational analysis of human memory. In H. Roediger & F. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 195–210). Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, 14, 471–517.
- Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, 96, 703–719.
- 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.
- Aronoff, M. (1976). *Word formation in generative grammar*. Cambridge, MA: MIT Press.
- 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.
- Baillargeon, R. (1986). Representing the existence and the location of hidden objects: Object permanence in 6- and 8-month old infants. *Cognition*, 23, 21–41.
- Barclay, C. R. (1996). Autobiographical remembering: Narrative constraints on objectified selves. In D. C. Rubin (Ed.), *Remembering our past: Studies in autobiographical memory* (pp. 94–125). New York: Cambridge University Press.
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. New York: Macmillan.
- Baxter, T. L., & Goldberg, L. R. (1987). Perceived behavioral consistency underlying trait attributions to oneself and another: An extension of the actor–observer effect. *Personality and Social Psychology Bulletin*, 13, 437–447.
- Beecher, M. (1990). The evolution of parent–offspring recognition in swallows. In D. A. Dewsbury (Ed.), *Contemporary issues in comparative psychology* (pp. 360–380). Sunderland, MA: Sinauer.
- Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6, pp. 1–62). New York: Academic Press.
- Biederman, I., & Cooper, E. E. (1992). Size invariance in visual object priming. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 121–133.
- Bower, G. H. (1972). Stimulus-sampling theory of encoding variability. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory*. New York: Wiley.
- Bower, G. H., & Gilligan, S. G. (1979). Remembering information related to one's self. *Journal of Research in Personality*, 13, 420–432.
- Brewer, L. F. (1986). What is autobiographical memory? In D. C. Rubin (Ed.), *Autobiographical memory* (pp. 25–49). New York: Cambridge University Press.
- Brody, N. (1988). *Personality: In search of individuality*. San Diego, CA: Academic Press.
- Brown, R. (1973). *A first language*. Cambridge MA: Harvard University Press.
- Budeshheim, 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, D. M., & Craik, K. H. (1984). Acts, dispositions, and personality. In B. A. Maher & W. B. Maher (Eds.), *Progress in experimental personality research* (pp. 241–301). San Diego, CA: Academic Press.
- Campbell, R., & Conway, M. A. (Eds.). (1995). *Broken memories: Case studies in memory impairment*. Oxford, England: Blackwell.
- Cappa, S. F., Frugoni, M., Pasquali, P., Perani, D., & Zorat, F. (1998). Category-specific naming impairment for artifacts: A new case. *Neurocase*, 4, 391–397.
- Caramazza, A., Hillis, A., Leek, E., & Miozzo, M. (1994). The organization of lexical knowledge in the brain: Evidence from category- and modality-specific deficits. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the mind: Domain-specificity in cognition and culture* (pp. 68–84). New York: Cambridge University Press.
- Caramazza, A., & Shelton, J. (1998). Domain-specific knowledge systems in the brain: The animate–inanimate distinction. *Journal of Cognitive Neuroscience*, 10, 1–34.
- Cazden, C. B. (1968). The acquisition of noun and verb inflections. *Child Development*, 39, 433–448.
- Cermak, L. S. (1984). The episodic–semantic distinction in amnesia. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory* (pp. 45–54). New York: Guilford Press.
- Cermak, L. S., & O'Connor, M. (1983). The anterograde and retrograde retrieval ability of a patient with amnesia due to encephalitis. *Neuropsychologia*, 21, 213–234.
- 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.
- Chumbley, J. I. (1986). *Concept learning by exemplar memorization*. Unpublished manuscript.
- Cohen, N. J. (1984). Preserved learning capacity in amnesia: Evidence for multiple memory systems. In L. R. Squire & N. Butters (Eds.), *Neuropsychology of memory* (pp. 83–103). New York: Guilford Press.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407–428.
- 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.
- Conway, M. A. (1992). A structural model of autobiographical memory. In M. A. Conway, D. C. Rubin, H. Spinnler, & W. A. Wagenaar (Eds.), *Theoretical perspectives on autobiographical memory* (pp. 167–194). Amsterdam: Kluwer Academic.
- Conway, M. A., & Bekerian, D. A. (1987). Organization in autobiographical memory. *Memory and Cognition*, 15, 119–132.
- Cooper, L. A., Schacter, D. L., Ballesteros, S., & Moore, C. (1992). Priming and recognition of transformed three-dimensional objects: Effects of size and reflection. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 43–57.
- Cooper, L. A., Schacter, D., & Moore, C. (1991, November). *Orientation affects both structural and episodic representations of 3-d objects*. Paper presented at the 32nd annual meeting of the Psychonomic Society, San Francisco, CA.
- Cosmides, L., & Tooby, J. (1987). From evolution to behavior: Evolutionary psychology as the missing link. In J. Dupré (Ed.), *The latest on the best: Essays on evolution and optimality* (pp. 277–306). Cambridge, MA: MIT Press.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163–228). New York: Oxford University Press.
- Cosmides, L., & Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions of the literature on judgment under uncertainty. *Cognition*, 58, 1–73.
- Cosmides, L., & Tooby, J. (1997). Dissecting the computational architecture of social inference mechanisms. In *Ciba Foundation Symposium: Vol. 208. Characterizing human psychological adaptations* (pp. 132–161). Chichester, England: Wiley.
- Cosmides, L., & Tooby, J. (2000). Consider the source: The evolution of adaptations for decoupling and metarepresentation. In D. Sperber (Ed.), *Metarepresentations: A multidisciplinary perspective* (pp. 53–115). New York: 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.
- Daly, M., & Wilson, M. (1995). Discriminative parental solicitude and the relevance of evolutionary models to the analysis of motivational systems. In M. Gazzaniga (Ed.), *The cognitive neurosciences*. (pp. 1269–1286). Cambridge, MA: MIT Press.
- Damasio, A. R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Putnam.
- 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. (1982). *The extended phenotype*. San Francisco: Freeman.
- Dawkins, R. (1986). *The blind watchmaker*. New York: Norton.
- 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.
- de Waal, F. (1982). *Chimpanzee politics: Power and sex among apes*. New York: Harper & Row.
- de Waal, F., & Lanting, F. (1997). *Bonobo: The forgotten ape*. Berkeley, CA: University of California Press.
- 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.
- Eustache, F., Desgranges, B., Laville, P., Guillery, B., Lalevee, C., Schaeffer, S., et al. (1999). Episodic memory in transient global amnesia: Encoding, storage, or retrieval deficit? *Journal of Neurology, Neurosurgery and Psychiatry*, 66, 148–154.
- Evans, J. J., Breen, E. K., Antoun, N., & Hodges, J. R. (1996). Focal retrograde amnesia for autobiographical events following cerebral vasculitis: A connectionist account. *Neurocase*, 2, 1–11.
- Evans, J., Wilson, B., Wraight, E. P., & Hodges, J. R. (1993). Neuropsychological and SPECT scan findings during and after transient global amnesia: Evidence for the differential impairment of remote episodic memory. *Journal of Neurology, Neurosurgery and Psychiatry*, 56, 1227–1230.
- Fong, G. T., & Markus, H. (1982). Self-schemas and judgments about others. *Social Cognition*, 1, 191–204.
- Foster, J. K., & Jelicic, M. (1999). *Memory: Systems, process or function?* New York: Oxford University Press.
- Funder, D. (1995). On the accuracy of personality judgment: A realistic approach. *Psychological Review*, 102, 652–670.
- Funder, D., & Sneed, C. (1993). Behavioral manifestations of personality: An ecological approach to judgmental accuracy. *Journal of Personality and Social Psychology*, 64, 479–490.
- Gallistel, C. R. (1995). The replacement of general-purpose theories with adaptive specializations. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1255–1267). Cambridge, MA: MIT Press.
- Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences*, 20, 1–55.
- Gould, S., & Lewontin, R. (1979). The spandrels of San Marco and the Panglossian program: A critique of the adaptationist program. *Proceedings of the Royal Society of London*, 250, 281–288.
- Graesser, A. C. (1981). *Prose comprehension beyond the word*. New York: Springer.
- Graesser, A. C., Gordon, S. E., & Sawyer, J. D. (1979). Recognition memory for typical and atypical actions in scripted activities: Tests of a

- script pointer + tag hypothesis. *Journal of Verbal Learning and Verbal Behavior*, *18*, 319–332.
- Graesser, A. C., Woll, S. B., Kowalski, D. J., & Smith, D. A. (1980). Memory for typical and atypical actions in scripted activities. *Journal of Experimental Psychology: Human Learning and Memory*, *6*, 503–515.
- Graham, K. S., & Hodges, J. R. (1997). Differentiating the roles of the hippocampal complex and the neocortex in long-term memory storage: Evidence from the study of semantic dementia and Alzheimer's disease. *Neuropsychology*, *11*, 77–89.
- Hampson, S. E. (1982). *The construction of personality: An introduction*. Boston, MA: Routledge & Keegan Paul.
- Hampton, R., Sherry, D., Shettleworth, S., Khurgel, M., Moshe, K., & Ivy, G. (1995). Hippocampal volume and food-storing behavior are related in parids. *Brain, Behavior and Evolution*, *45*, 54–61.
- Hart, J., Berndt, R., & Caramazza, A. (1985). Category-specific naming deficit following cerebral infarction. *Nature*, *316*, 439–440.
- Hirshman, E., & Lanning, K. (1999). Evaluating the role of self-judgments in conscious recollection. *Applied Cognitive Psychology*, *13*, 29–42.
- Hodges, J. R., Spatt, J., & Patterson, K. (1999). "What" and "how": Evidence for the dissociation of object knowledge and mechanical problem-solving skills in the human brain. *Proceedings of the National Academy of Sciences*, *96*, 9444–9448.
- Hudson, J. A. (1986). Memories are made of this: General event knowledge and autobiographical memory. In K. Nelson (Ed.), *Event knowledge: Structure and function in development* (pp. 97–118). Hillsdale, NJ: Erlbaum.
- Jackendoff, R. (1987). *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- Jacobs, W., & Nadel, L. (1985). Stress-induced recovery of fears and phobias. *Psychological Review*, *92*, 512–531.
- 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.
- Johnson, R., Jr., Kreiter, K., Zhu, J., & Russo, B. (1998). A spatio-temporal comparison of semantic and episodic cued recall and recognition using event-related brain potentials. *Cognitive Brain Research*, *7*, 119–136.
- Kendrick, K., Levy, F., & Keever, E. (1992, May 8). Changes in the sensory processing of olfactory signals induced by birth in sheep. *Science*, *256*, 833–836.
- Kihlstrom, J. F., & Cantor, N. (1984). Mental representations of the self. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 17, pp. 1–47). New York: Academic Press.
- Kihlstrom, J. F., Cantor, N., Albright, J. S., Chew, B. R., Klein, S. B., & Niedenthal, P. M. (1988). Information processing and the study of the self. In L. Berkowitz, (Ed.), *Advances in experimental social psychology: Vol. 21. Social psychological studies of the self: Perspectives and programs* (pp. 145–178). San Diego, CA: Academic Press.
- 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* (2nd ed., pp. 153–208). Hillsdale, NJ: Erlbaum.
- Kihlstrom, J. F., & Klein, S. B. (1997). Self-knowledge and self-awareness. In J. G. Snodgrass & R. L. Thompson (Eds.), *The self across psychology: Self-recognition, self-awareness, and the self concept*. *Annals of the New York Academy of Sciences* (Vol. 818, pp. 5–17). New York: New York Academy of Sciences.
- Kircher, T. T. J., Senior, C., Phillips, M. L., Benson, P. J., Bullmore, E. T., Brammer, M., et al. (2000). Towards a functional neuroanatomy of self-processing: Effects of faces and words. *Cognitive Brain Research*, *10*, 133–144.
- Kitchener, E. G., Hodges, J. R., & McCarthy, R. (1998). Acquisition of post-morbid vocabulary and semantic facts in the absence of episodic memory. *Brain*, *121*, 1313–1327.
- 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. 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: Psychology 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., 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., Costabile, K. A., & Mei, L. (in press). Is there something special about the self? A neuropsychological case study. *Journal of Research in Personality*.
- Klein, S. B., Cosmides, L., Tooby, J., & Chance, S. S. (2001). Priming exceptions: A test of the scope hypothesis in naturalistic trait judgments. *Social Cognition*, *19*, 443–468.
- 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. B., & Loftus, J. (1993a). Behavioral experience and trait judgments about the self. *Personality and Social Psychology Bulletin*, *19*, 740–745.
- Klein, S., & Loftus, J. (1993b). The mental representation of trait and autobiographical knowledge about the self. 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. 1–49). Hillsdale, NJ: Erlbaum.
- 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., & Kihlstrom, J. F. (in press). 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*.
- 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., Rozendal, K., & Cosmides, L. (in press). A social-cognitive neuroscience analysis of the self. *Social Cognition*.
- 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.

- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*, 1399–1402.
- Kolar, D., Funder, D., & Colvin, C. (1996). Comparing the accuracy of personality judgments by the self and knowledgeable others. *Journal of Personality*, *64*, 311–337.
- Kopelman, M. D., Wilson, B. A., & Baddeley, A. D. (1989). The autobiographical memory interview: A new assessment of autobiographical and personal semantic memory in amnesic patients. *Journal of Clinical and Experimental Neuropsychology*, *11*, 724–744.
- Kuczaj, S. A. (1977). The acquisition of regular and irregular past tense forms. *Journal of Verbal Learning and Verbal Behavior*, *16*, 589–600.
- 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 knowledge 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.
- Marcus, G., Pinker, S., Ullman, M., Hollander, M., Rosen, T., & Xu, F. (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development*, *57*(4, Serial No. 228).
- Markowitsch, H. J. (1995). Which brain regions are critically involved in the retrieval of old episodic memory? *Brain Research Reviews*, *21*, 117–127.
- Markowitsch, H. J., Calabrese, P., Liess, J., Haupts, M., Durwen, H. F., & Gehlen, W. (1993). Retrograde amnesia after traumatic injury of the fronto-temporal cortex. *Journal of Neurology, Neurosurgery, and Psychiatry*, *56*, 988–992.
- Markowitsch, H. J., Calabrese, P., Neufeld, H., Gehlen, W., & Durwen, H. F. (1999). Retrograde amnesia for world knowledge and preserved memory for autobiographical events: A case report. *Cortex*, *35*, 243–252.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.
- Martin, E. (1971). Verbal learning theory and independent retrieval phenomena. *Psychological Review*, *78*, 314–332.
- Mayr, E. (1983). How to carry out the adaptationist program. *American Naturalist*, *121*, 324–334.
- McCarthy, R. A., & Warrington, E. K. (1990). *Cognitive neuropsychology: A clinical introduction*. San Diego, CA: Academic Press.
- McCarthy, R. A., & Warrington, E. K. (1992). Actors but not scripts: The dissociation of people and events in retrograde amnesia. *Neuropsychologia*, *30*, 633–644.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–437.
- McClelland, J., & Rumelhart, D. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. Psychological and biological models*. Cambridge, MA: MIT Press.
- McCracken, G., & Gustin, M. (1992). Nursing behavior in Mexican free-tailed bat maternity colonies. *Ethology*, *89*, 305–321.
- McNamara, T. P. (1992a). Priming and constraints it places on theories of memory and retrieval. *Psychological Review*, *99*, 650–662.
- McNamara, T. P. (1992b). Theories of priming: I. Associative distance and lag. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 1173–1190.
- Medin, D. L., & Bettger, J. G. (1994). Presentation order and recognition of categorically related examples. *Psychonomic Bulletin & Review*, *1*, 250–254.
- Milner, B., Corkin, S., & Teuber, H. L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H.M. *Neuropsychologia*, *6*, 215–234.
- Mineka, S., & Cook, M. (1988). Social learning and the acquisition of snake fear in monkeys. In T. R. Zentall & B. G. Galef (Eds.), *Social learning: Psychological and biological perspectives* (pp. 51–73). Hillsdale, NJ: Erlbaum.
- Mischel, W. (1968). *Personality and assessment*. New York: Wiley.
- Mooser, S. D. (1979). The role of experimental design in investigations of the fan effect. *Journal of Experimental Psychology: Human Learning and Memory*, *5*, 125–134.
- Nadel, L. (1994). Multiple memory systems: What and why, an update. In D. Schacter & E. Tulving (Eds.), *Memory Systems* (pp. 39–63). Cambridge, MA: MIT Press.
- Nelson, K. (1993a). Developing self-knowledge from autobiographical memory. 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. 111–120). Hillsdale, NJ: Erlbaum.
- Nelson, K. (1993b). The psychological and social origins of autobiographical memory. *Psychological Science*, *4*, 1–8.
- Nelson, K. (1996). *Language in cognitive development: Emergence of the mediated mind*. New York: Cambridge University Press.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, *101*, 53–79.
- Ohman, A., Dimberg, U., & Ost, L. (1985). Biological constraints on the fear response. In S. Reiss & R. Bootzin (Eds.), *Theoretical issues in behavior therapy* (pp. 123–175). New York: Academic Press.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford, England: Clarendon Press.
- Palmeri, T. J., & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 548–568.
- Park, B. (1986). A method for studying the development of impressions of real people. *Journal of Personality and Social Psychology*, *51*, 907–917.
- Perani, D., Bressi, S., Cappa, S. F., Vallar, G., Alberoni, M., Grassi, F., et al. (1993). Evidence of multiple memory systems in the human brain: A [¹⁸F] FDG PET metabolic study. *Brain*, *116*, 903–919.
- Peretz, I. (1996). Can we lose memory for music? A case of music agnosia in a nonmusician. *Journal of Cognitive Neuroscience*, *8*, 481–496.
- Pinker, S. (1979). Formal models of language learning. *Cognition*, *7*, 217–283.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.
- Pinker, S. (1991). Rules of language. *Science*, *253*, 530–535.
- Pinker, S. (1994). *The language instinct*. New York: Morrow.
- Pinker, S. (1997, June 5). Words and rules in the human brain. *Nature*, *387*, 547–548.
- Pinker, S. (1999). *Words and rules: The ingredients of language*. New York: Basic Books.
- Pittman, R., & Orr, S. (1995). Psychophysiology of emotional and memory networks in posttraumatic stress disorder. In J. L. McGaugh, N. Weinberger, & G. Lynch (Eds.), *Brain and memory: Modulation and mediation of neuroplasticity* (pp. 75–83). New York: Oxford University Press.
- Polster, M. R., Nadel, L., & Schacter, D. L. (1991). Cognitive neuroscience analyses of memory: A historical perspective. *Journal of Cognitive Neuroscience*, *3*, 95–116.
- Quinlan, P. (1991). *Connectionism and psychology: A psychological perspective on new connectionist research*. Chicago: University of Chicago Press.
- Quinlan, P. (1992). *An introduction to connectionist modeling*. Hillsdale, NJ: Erlbaum.
- Reber, P. J., Knowlton, B. J., & Squire, L. R. (1996). Dissociable properties of memory systems: Differences in the flexibility of declarative and nondeclarative knowledge. *Behavioral Neuroscience*, *110*, 861–871.
- Reinvang, I., & Gjerstad, L. (1998). Focal retrograde amnesia associated with vascular headache. *Neuropsychologia*, *36*, 1335–1341.
- Riskey, D. R. (1979). Verbal memory processes in impression formation.

- Journal of Experimental Psychology: Human Learning and Memory*, 5, 271–281.
- Roediger, H. L., & Blaxton, T. A. (1987). Retrieval modes produce dissociations in memory for surface information. In D. Gorfein & R. R. Hoffman (Eds.), *Memory and cognitive processes: The Ebbinghaus Centennial Conference* (pp. 349–379). Hillsdale, NJ: Erlbaum.
- Roediger, H. L., & Craik, F. I. M. (Eds.). (1989). *Varieties of memory and consciousness: Essays in honour of Endel Tulving*. Hillsdale, NJ: Erlbaum.
- Roediger, H. L., III, Weldon, M. S., & Challis, B. H. (1989). Explaining dissociations between implicit and explicit measures of retention: A processing account. In H. L. Roediger III & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 3–41). Hillsdale, NJ: Erlbaum.
- Rozin, P. (1976). The evolution of intelligence and access to the cognitive unconscious. In J. M. Sprague & A. N. Epstein (Eds.), *Progress in psychobiology and physiological psychology* (pp. 245–277). New York: Academic Press.
- Rubin, D. C. (Ed.). (1986). *Autobiographical memory*. New York: Cambridge University Press.
- Rumelhart, D., & McClelland, J. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 1. Foundations*. Cambridge, MA: MIT Press.
- 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.
- Schacter, D. L. (1995). Implicit memory: A new frontier for cognitive neuroscience. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 815–824). Cambridge, MA: MIT Press.
- Schacter, D. L., & Buckner, R. L. (1998). On the relations among priming, conscious recollection, and intentional retrieval: Evidence from neuroimaging research. *Neurobiology of Learning and Memory*, 70, 284–303.
- Schacter, D. L., & Tulving, E. (Eds.). (1994). *Memory systems*. Cambridge, MA: MIT Press.
- Schacter, D. L., Wang, P. L., Tulving, E., & Freedman, M. (1982). Functional retrograde amnesia: A quantitative case study. *Neuropsychologia*, 20, 523–532.
- Schank, R. C. (1980). Language and memory. *Cognitive Science*, 4, 243–284.
- Schank, R. C. (1982). *Dynamic memory: A theory of reminding and learning in computers and people*. Cambridge, England: Cambridge University Press.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.
- Schell, T. L., Klein, S. B., & Babey, S. H. (1996). Testing a hierarchical model of self-knowledge. *Psychological Science*, 7, 170–173.
- Schneider, D. J., Hastorf, A. H., & Ellsworth, P. C. (1979). *Person perception* (2nd ed.). Menlo Park, CA: Addison-Wesley.
- Shallice, T. (1979). Neuropsychological research and the fractionalization of memory systems. In L. G. Nilsson (Ed.), *Perspectives on memory research* (pp. 257–277). Hillsdale, NJ: Erlbaum.
- 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.
- Sherry, D. (1997). Cross-species comparisons. In *Ciba Foundation Symposium: Vol. 208. Characterizing human psychological adaptations* (pp. 181–194). Chichester, England: Wiley.
- Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, 94, 439–454.
- Silverman, I., & Eals, M. (1992). Sex differences in spatial abilities: Evolutionary theory and data. In J. H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 533–549). New York: Oxford University Press.
- Singer, J. A., & Salovey, P. (1993). *The remembered self: Emotion and memory in personality*. New York: Free Press.
- Sirigu, A., & Grafman, J. (1996). Selective impairments within episodic memories. *Cortex*, 32, 83–95.
- Smith, E. R. (1990). Content and process specificity in the effects of prior experience. In T. K. Srull & R. S. Wyer Jr. (Eds.), *Advances in social cognition* (Vol. 3, pp. 1–59). Hillsdale, NJ: Erlbaum.
- Smith, D. A., & Graesser, A. C. (1981). Memory for actions in scripted activities as a function of typicality, retention interval, and retrieval task. *Memory and Cognition*, 9, 550–559.
- Snyder, M., & Cantor, N. (1979). Testing hypotheses about other people: The use of historical knowledge. *Journal of Experimental Social Psychology*, 15, 330–342.
- Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, 14, 29–56.
- Squire, L. R. (1987). *Memory and brain*. New York: Oxford University Press.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, 99, 143–145.
- Squire, L. R. (1994). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems* (pp. 203–231). Cambridge, MA: MIT Press.
- Squire, L. R., & Knowlton, B. (1995). Memory, hippocampus, and brain systems. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 825–837). Cambridge, MA: MIT Press.
- Staubli, U., & Lynch, G. (1987). Stable hippocampal long-term potentiation elicited by theta pattern stimulation. *Brain Research*, 435, 227–234.
- Strauman, T. J. (1990). Self-guides and emotionally significant childhood memories: A study of retrieval efficiency and incidental negative emotional content. *Journal of Personality and Social Psychology*, 59, 869–880.
- Suddendorf, T., & Corballis, M. C. (1997). Mental time travel and the evolution of the human mind. *Genetic, Social, and General Psychology Monographs*, 123, 133–167.
- Suzuki, K., Yamadori, A., & Fugii, T. (1997). Category-specific comprehension deficit restricted to body parts. *Neurocase*, 3, 193–200.
- Thorndyke, P. W., & Bower, G. H. (1974). Storage and retrieval processes in sentence memory. *Cognitive Psychology*, 6, 515–543.
- 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. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 381–403). New York: Academic Press.
- Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.
- Tulving, E. (1984). Precis of elements of episodic memory. *Behavioral and Brain Sciences*, 7, 223–268.
- Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie Canadienne*, 26, 1–12.
- Tulving, E. (1987). Multiple memory systems and consciousness. *Human Neurobiology*, 6, 67–80.
- Tulving, E. (1989). Memory: Performance, knowledge, and experience. *European Journal of Cognitive Psychology*, 1, 3–26.
- Tulving, E. (1993a). Self-knowledge of an amnesic individual is repre-

- sented abstractly. In T. K. Srull & R. S. Wyer Jr. (Eds.), *Advances in social cognition: Vol. 5. The mental representation of trait and autobiographical knowledge about the self* (pp. 147–156). Hillsdale, NJ: Erlbaum.
- Tulving, E. (1993b). What is episodic memory? *Current Perspectives in Psychological Science*, 2, 67–70.
- Tulving, E. (1995a). Memory: Introduction. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 751–753). Cambridge, MA: MIT Press.
- Tulving, E. (1995b). Organization of memory: Quo vadis? In M. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 839–847). Cambridge, MA: MIT Press.
- Tulving, E., Hayman, C. A., & MacDonald, C. A. (1991). Long-lasting perceptual priming and semantic learning in amnesia: A case experiment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 595–617.
- Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science*, 247, 301–306.
- Tulving, E., Schacter, D. L., McLachlan, D. R., & Moscovitch, M. (1988). Priming of semantic autobiographical knowledge: A case study of retrograde amnesia. *Brain and Cognition*, 8, 3–20.
- Ullman, M., Corkin, S., Coppola, M., Hickok, G., Growdon, J. H., Korshetz, W. J., & Pinker, S. (1997). A neural dissociation within language: Evidence that the mental dictionary is part of declarative memory, and that grammatical rules are processed by the procedural system. *Journal of Cognitive Neuroscience*, 9, 289–299.
- Van der Linden, M., Bredart, S., Depoorter, N., & Coyette, F. (1996). Semantic memory and amnesia: A case study. *Cognitive Neuropsychology*, 13, 391–413.
- Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Van Paesschen, W., & Mishkin, M. (1997, July 18). Differential effects of early hippocampal pathology on episodic and semantic memory. *Science*, 277, 376–380.
- Warrington, E. K. (1975). The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology*, 27, 635–657.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, 107, 829–853.
- Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, 20, 273–281.
- Weiskrantz, L. (1987). Neuroanatomy of memory and amnesia: A case for multiple memory systems. *Human Neurobiology*, 6, 93–105.
- Weiskrantz, L. (1990). Problems of learning and memory: One or multiple memory systems? *Philosophical transactions of the Royal Society of London B*, 329, 99–108.
- Wexler, K., & Culicover, P. (1980). *Formal principles of language acquisition*. Cambridge, MA: MIT Press.
- Williams, G. C. (1966). *Adaptation and natural selection: A critique of some current evolutionary thought*. Princeton, NJ: Princeton University Press.
- Williams, G. C., & Nesse, R. M. (1991). The dawn of Darwinian medicine. *Quarterly Review of Biology*, 66, 1–22.
- Wilson, B. A., & Wearing, D. (1995). Prisoner of consciousness: A state of just awakening following herpes simplex encephalitis. In R. Campbell & M. A. Conway (Eds.), *Broken memories: Case studies in memory impairment* (pp. 14–30). Cambridge, MA: Blackwell.
- Wood, F. B., Brown, I. S., & Felton, R. H. (1989). Long-term follow-up of a childhood amnesic syndrome. *Brain and Cognition*, 10, 76–86.
- Woodworth, R. S. (1938). *Experimental psychology*. New York: Holt.
- Wright, J. C., & Mischel, W. (1988). Conditional hedges and the intuitive psychology of traits. *Journal of Personality and Social Psychology*, 55, 454–469.
- Wyer, R. S., & Gordon, S. E. (1984). The cognitive representation of social information. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 2, pp. 75–150). Hillsdale, NJ: Erlbaum.

Received April 7, 1999

Revision received October 11, 2000

Accepted December 29, 2000 ■