

Decisions and the Evolution of Memory: Multiple systems, multiple functions

Stanley B. Klein, Leda Cosmides, John Tooby, and Sarah Chance

Psychological Review, in press.

Abstract

We explore the functions of episodic and semantic memory in an unusual way: by considering the computational requirements of the decision rules that access them. Because memory evolved in order to supply useful, timely information to the organism's decision-making systems, we expect that decision rules, multiple memory systems, and the search engines that link them have co-evolved to mesh in a co-adapted, functionally interlocking way. Decision rules designed to use data from memory systems need search engines that find and rapidly deliver adaptively appropriate information, regardless of which memory system(s) the information is stored in. If so, then when it comes to retrieval, dissociations between memory systems should not be absolute: One should find independence for some decision tasks, but not for others. We show that whether retrieval is independent depends in lawful ways on the *content* of the information being retrieved. Using a priming paradigm and a decision task involving person memory, we tested and confirmed a novel prediction that follows from this view: When decision rules designed to make personality judgments retrieve from semantic memory a trait summary of a person, they should also retrieve episodic memories, provided these encode incidents in which that person's behavior was *inconsistent* with that summary. These results 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*). These experiments support the hypothesis that priming is an evolved adaptation designed to expedite information likely to be required by a decision task. We discuss implications of this research for the architecture of

memory systems, the functional significance of priming in different domains, and the interpretation of dissociation data.

Until the paper appears, we can put the abstract, but not the full text, on this webpage. We can, however, send you a preprint copy of this paper.

If you would like a preprint copy via email attachment, please email Stanley B. Klein at klein@psych.ucsb.edu. Please include your mailing address, just in case we need to send a paper copy instead.

Corresponding author:

Stanley B. Klein
Dept. of Psychology
University of California
Santa Barbara, CA 93106

For a summary of *one aspect of the argument in this paper*, see below...

Decisions and the Evolution of Multiple Memory Systems

Using Personality Judgments to Test the Scope Hypothesis

Leda Cosmides, Stanley B. Klein, John Tooby & Sarah Chance

Overheads of talk presented at the Human Behavior and Evolution Society meetings

June 14-17, 2001

University College London.

Using Personality Judgments to Test the Scope Hypothesis

Trait summaries:

Tom: *Rarely Cruel* **Self:** *Usually Friendly*

Why store episodes after a trait summary has been formed?

Benefits of having a trait summary:

- ✓ speedy access to information
- ✓ avoiding redundant and costly computation

Why have anything else? 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.**

To rely on semantic memory alone is to sacrifice accuracy for speed:

- ✓ Trait summaries lack the resolution of an episodic memory
- ✓ Episodic memories are situation-specificity incarnate: They are records of how particular people behaved in specific situations.

Maintaining a database of episodic memories solves several problems that a trait summary cannot.

- ✓ ***Re-evaluating conclusions.*** New information may cause previous episodes to be re-interpreted, drastically changing one's judgments of a person or a situation.
- ✓ ***Assessing credal value.*** Maintaining source information allows one to evaluate the credal value of people's statements.
- ✓ ***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 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.

Demi: *Usually Calm* - EXCEPT [the day her tenure was being decided]

Episodic memories as boundary conditions on the scope of generalizations:

The scope hypothesis.

A system engineered to retrieve, along with a generalization, episodes that place boundary conditions on its scope would achieve the best of both worlds: speed courtesy of semantic memory, accuracy courtesy of episodic memory.

This function leads to a counter-intuitive prediction about the relation between semantic and episodic memory when people are called upon to make trait judgments.

When asked, “Does this describe you: friendly?”, the 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.

SCOPE HYPOTHESIS

in a nutshell...

1. 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.
2. 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.
3. 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.
4. 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 – i.e., situation Z.
5. 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.

6. 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.

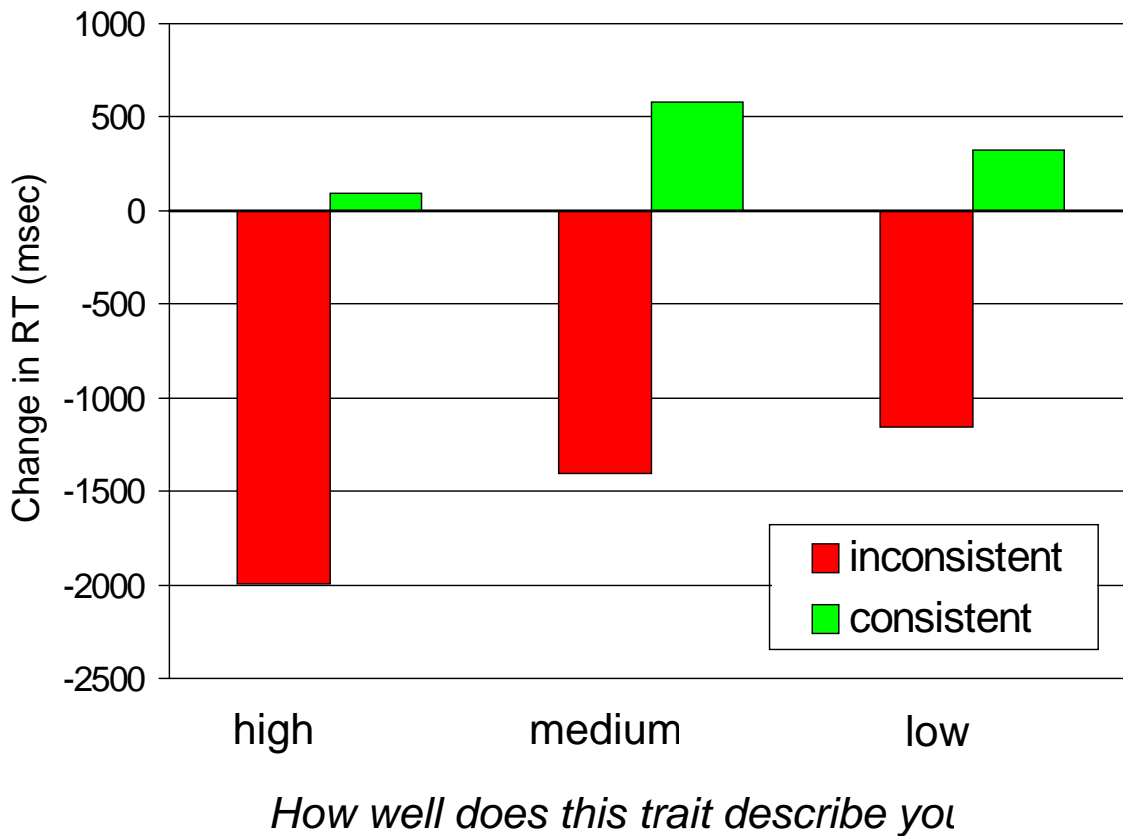
In the experiments, subjects did a series of tasks for many different personality traits. These tasks came in pairs. The initial task of the pair asked: *Does this describe you?: kind* (the trait was varied). The second, target task, was a recall task. The subject was to recall an instance in which s/he manifested the trait flashed on the screen. Sometimes the recall task asked the subject to recall an episode consistent with the trait (e.g., *kind*); sometimes it asked the subject to recall an episode that was inconsistent with the trait (e.g., *unkind*). The time it takes to recall the episode was recorded, and compared to the time for a control task (various types of control tasks were used, all with the same results).

The self study: people are "experts" on themselves – previous results indicate that people have summary representations for all their traits (including ones that they do not feel describe them well (e.g., Self: Rarely Cruel). Because summaries are present for most traits, we expected inconsistent episodes to be primed by the initial "describe" task. (Note: results hold regardless of the trait's social desirability)

The mother study: Previous results indicated that people have summary representations for traits that they consider to be highly descriptive of their mothers, but not for ones that they consider to be only moderately descriptive. Thus we expected that inconsistent episodes would be primed by the initial "describe" task for traits highly descriptive of mother, but not for her moderately descriptive traits. Because people lack summaries for these traits, they will have to retrieve consistent episodes to make a judgment, causing these to be primed on the recall task.

Inconsistent Episodes are Primed


Self Study



Activating a trait summary about the self can prime episodic memory. 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.)

Inconsistent Episodes are Primed Only when a Trait Summary has First been Retrieved

Mother Study



How well does this trait describe your mother?

Activating a trait summary about another person primes inconsistent episodes. 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.)