

Why Cognitive Psychologists Should Know Comparative Psychology; Why Comparative Psychologists Should Know Cognitive Psychology

J. David Smith

State University of New York at Buffalo, U.S.A.

The author contrasts the interpretative perspectives offered by comparative and cognitive psychology. Four strengths of the comparative program are considered in the context of recent research on animals' capacity for uncertainty monitoring or metacognition. However, several historical limitations of the comparative perspective are also discussed—in these areas the cognitive perspective holds the stronger interpretative hand. The author considers the negative consequences that comparative psychology has garnered from the continued premium it has placed on low-level associative explanations of behavioral phenomena, and the constructive synergy that might come from integrating the comparative and cognitive programs.

My assignment for this special issue of the journal was to bridge between the fields of comparative psychology and cognitive psychology. This is a difficult assignment because both fields are large and complex and it is difficult to summarize and integrate their perspectives. Worse, the perspectives typically used to describe animal and human behavior stood at opposite theoretical poles for decades.

Nonetheless, integrating these traditions is important for reasons that will be discussed. Accordingly, this article proceeds as follows. I will describe several aspects of the comparative tradition that have enriched me as a cognitive psychologist. I will consider areas in which comparative psychology has historically been self-limiting and interpretatively cramped. In these areas cognitive psychology has the sharper interpretative tools. I will consider the negative consequences that comparative psychology has garnered from the premium it has placed on low-level associative explanations of behavioral phenomena. I will close by emphasizing the value that a synthesis of the behavioral and cognitive perspectives will have as the field of comparative cognition enters a new century. There is reason to hope that this blended, truly comparative perspective is gaining strength in the field.

The research literature on metacognition provides a good vantage point from which to examine the comparative and cognitive perspectives. This research can serve as a simplifying, bridging, and focusing case study and as a basis for finding constructive common ground between the fields. It can provide a strong contrast between the approaches that theorists in the two fields take to the same paradigms and to the same data patterns produced by human and nonhuman animals (hereafter humans and animals, respectively). One may imagine (rightly) that theoretical sparks will flash from this contrast (Smith, Shields, & Washburn, in press, and accompanying commentaries).

The idea behind research on metacognition is that some minds—human

minds at least—contain a cognitive executive that monitors and, if necessary, controls thought and problem solving. An example of this capacity is that humans feel uncertain—they often know when they do not know—and they often respond intelligently in difficult situations by pausing, reflecting, and seeking help or additional information. These states of (not) knowing and the adaptive responses that accompany them have been the focus of extensive research on metacognition, metamemory, and uncertainty monitoring (Brown, 1991; Brown et al., 1982; Dunlosky & Nelson, 1992; Flavell, 1979; Hart, 1965; Koriati, 1993; Nelson, 1992; Schwartz, 1994; Smith, Brown, & Balfour, 1991).

The monitoring functions of the cognitive executive are studied by having human participants make metacognitive judgments like an ease-of-learning judgment about whether material will be easy or hard to learn, a judgment of learning about how much has been learned, a feeling-of-knowing judgment about whether information is potentially available in memory, and a confidence judgment about a potential answer. The control functions of the cognitive executive are studied by asking whether humans devote disproportionate study time to difficult items, terminate studying when sufficient learning has been achieved, select new retrieval strategies when the present ones are failing, or abandon retrieval efforts if successful retrieval seems unlikely.

Metacognition is taken to be one of humans' most sophisticated cognitive capacities. This capacity shows that there are tiers or hierarchies of oversight and regulation in mind (these are designated the meta and object levels of cognition; Nelson & Narens, 1990). This capacity shows that humans are aware of the processes of mind. Metacognitive states are also a pointer to humans' self-awareness, because uncertainty and doubts, but also knowing and confidence, are personalized (i.e., I am uncertain; I know). Finally, metacognition may also be an indicator of humans' declarative consciousness (Nelson, 1996) because we so easily introspect about those states and speak about them to one another.

For all these reasons, and over all because metacognition is such a high-level, sophisticated cognitive capacity, it is an intriguing question whether animals have analogous or homologous capacities. Accordingly, my colleagues and I have been exploring the possibility of studying animal metacognition (Shields, Smith, & Washburn, 1997; Shields et al., 2003; Smith & Schull, 1989; Smith et al., 1995; 1997; 1998; in press). More recently, other laboratories have taken up the same theoretical and empirical challenge (Hampton, 2001; Inman & Shettleworth, 1999; Son & Kornell, in press). It turns out that the clarifying impact of the comparative perspective is felt immediately as one begins to study the metacognitive capacity of animals. Here now are four strengths the comparative perspective brings to this enterprise.

The Elemental Phenomenon

First, the necessities of animal research require one to distill the theoretical constructs and experimental variables into a form that suits animal participants. Regarding the research area of interest here, the literature on metacognition is filled with difficult and potentially contentious constructs. What is the cognitive executive? What does it monitor and what does monitoring even mean? What is the role (critical or noncritical) of language and verbal self-report in the metacogni-

tion system—or in our theory of the metacognition system? Likewise, what are the roles of consciousness or self-awareness in that system?

Questions like these arise immediately on considering animals' metacognitive capacity because for animals one has to leave aside explicit requests for metacognitive judgments and the verbal self reports featured in so many human metacognitive assessments. Instead the task has to be distilled down to be behavioral, perceptual, and nonverbal.

To illustrate this process and put a relevant phenomenon on the table, let us consider the metamemory paradigm we created for animals. The idea in metamemory is that humans can judge veridically and adaptively whether they remember or not. They can also decline to accept tests of poorly remembered material, as they often do when they forego some questions on essay exams while choosing others. Smith et al. (1998) tried to distill these cognitive capacities and so ask whether animals too can evaluate their state of remembering and respond adaptively when that state does not justify completing a memory test. To do so, we used the fact that some serial positions of a memory list are predictably harder to remember than others. Thus the experimenter can know which memory material will probably be difficult and one can ask whether the animal will selectively decline tests of that memory material. Smith et al. adopted the serial probe recognition (SPR) task that has been so influential in animal memory research (Castro & Larsen, 1992; Roberts & Kraemer, 1981; Sands & Wright, 1980; Wright et al., 1985). Our animals saw a list of four pictures and then were asked whether a 5th picture that was a probe had been in the list or not. Readers can get an idea of what a trial in this task was like by scanning the four shapes in Figure 1 for about 1 s each, preparing for a later possible test of memory. The monkeys had a There response to make if they thought the probe had been in the list, or a Not There response if they thought not. They could also make an Uncertain response if they did not wish to complete the memory test and risk the long timeout that accompanied an error.




Figure 1. An example of the memory lists Smith et al. (1998) presented to monkeys in a serial probe recognition task. The pictures would have been presented to monkeys successively for about 1 s each, by computer and in different colors. Reprinted from Smith (in press) by permission of Oxford University Press.

Figure 2A shows a monkey's performance in this task. On many trials he chose to complete the memory test and, when he did so, he showed the primacy and recency effects (in his percentages correct) that are the hallmark of this task and that confirm that the SPR task was instantiated correctly and probed the monkey's memory appropriately. Most important, though, the animal used the Uncer-

tain response in the mirror image of these serial-position effects. That is, he declined trials selectively when his worst serial positions were probed.

The same was true when humans performed under similar conditions (Figure 2B), though these conditions were not right for letting humans show a strong primacy effect. Humans were instructed explicitly to use the uncertainty response to report on and cope with memory indeterminacy. This makes the similarity of the animal's performance especially intriguing. Humans declined memory tests when they thought they didn't remember. Monkeys behaved like humans.

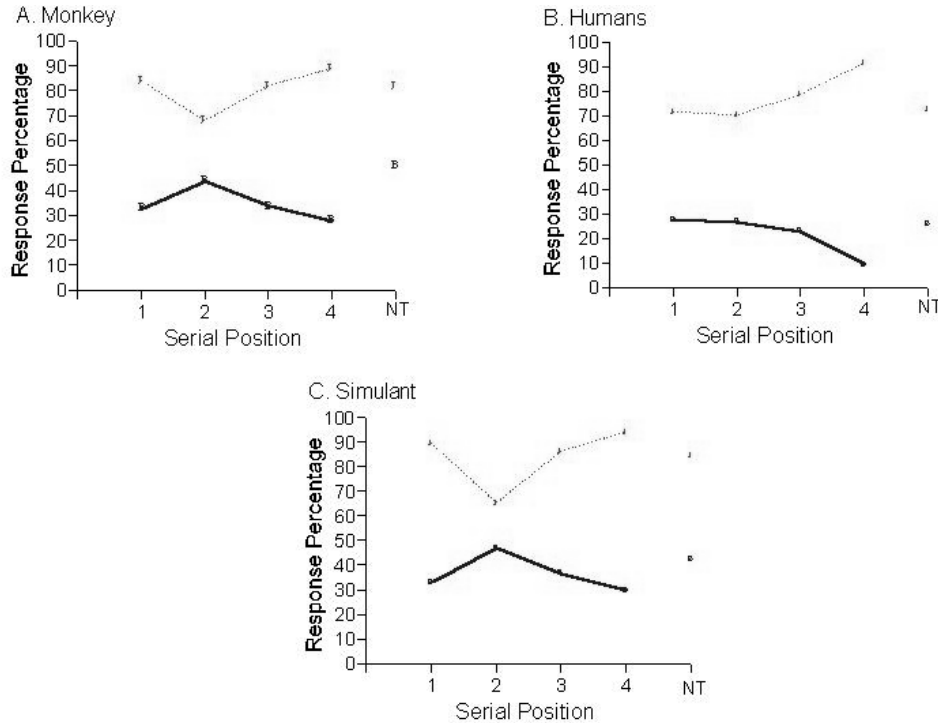


Figure 2. (A) Serial probe recognition (SPR) performance by Monkey Baker in the task of Smith et al. (1998). NT denotes Not There trials. The serial position (1-4) of the probe in the list of pictures is also given along the X-axis for the probes on There trials. The percentage of total trials that received the uncertainty response is shown (bold line). The percentage correct (of the trials on which the memory test was attempted) is also shown (dotted line). (B) Performance by 10 humans in a similar SPR task used by Smith et al. (C) Performance by the simulant that fit best Monkey Baker's performance in the serial probe recognition (SPR) task of Smith et al. (1998). Figure 2A and 2B reprinted from Smith et al. (1998), Copyright © 1998 by the American Psychological Association, adapted with permission. Figure 2C Reprinted from Smith et al. (2003) by permission of Cambridge University Press.

Some view the monkey's performance as just described (and of course the humans' performance) as essentially metacognitive. Some do not. For example, one might think that metacognition has to come with explicit, declarative consciousness or with a self-reflective sense of personhood, and these this experiment cannot show in monkeys. But notice that even if one does decide the question in that way, distilling the performance was still useful theoretically because it shed

light on the component capacities that make up the full-blown human construct and because it lets us sharpen our own theoretical sense of the construct. That is, it let one group of readers realize that metacognition might in some circumstances or in some minds be shorn of declarative consciousness or self-reflective personhood. It let another group of readers realize that in the end the issue isn't mere cognitive monitoring, but that conscious, self-reflective monitoring should be a part of the theoretical construct called metacognition.

Thus, the first moral I took from comparative psychology concerned how clarifying it is that a task is stripped to its essentials for animals. This moral has also been applied productively by cognitive developmental psychologists. In the 1970's, Rochel Gelman (e.g., Gelman & Gallistel, 1978) and others removed the representational obstacles, the information-processing conflicts, and the nonessential sources of difficulty from the then standard (i.e., Piagetian) cognitive-developmental tasks and used the modified tasks to reassess children's cognitive level in domains like conservation, classification, and so forth. Thus they assessed ultimately simply and directly the underlying cognitive competencies. This distillation of the tasks and the issues remade cognitive developmental theory and ultimately replaced Piagetian stage theory.

In fact, it could be a constructive exercise generally for cognitive psychologists to spend a few hours giving themselves and their constructs this Gedanken challenge thrown down by comparative psychology. Accepting that challenge, they could imagine how to tailor their experimental demonstrations for animals, preserving the critical cognitive capacities and sensitivities while omitting the nonessential "human" elements of the procedures. One may find that the comparative analog of a human task clarifies the suite of capacities in the full-blown human performance and clarifies which members of the suite are most theoretically important.

The Relevant Signal or Cue

A second constructive element in comparative psychology's program is the strong premium placed on specifying the cue or signal that the organism senses and responds to during performance. In this area it is possible to evaluate the human literature on metacognition in a critical light, even though that literature is very interesting and theoretically challenging. Yet the literature has been somewhat imprecise on what cue or psychological signal makes humans feel uncertain or what makes them think they know. Part of the reason for the imprecision is that we are humans inquiring about a thing that humans do so easily. We know what humans are doing in the metamemory task just described. They are responding Uncertain when they think they do not remember. We are so comfortable with the state of uncertainty, and with the adaptive reaction to the state, that there does not seem to be anything difficult or problematic to explain. Of course, in reality, to respond Uncertain when you think you do not remember is an incredible psychological feat that needs careful psychological explication.

Comparative psychologists realize clearly this need for explication. Accordingly, when monkeys produce just the same graph as humans do in the SPR task, the explication is required. Partly this requirement arises because comparative psychologists studies species for which phenomena of mind are not accepted so comfortably as we accept them in humans. Partly this requirement arises because

of the inherently skeptical scientific philosophy in comparative psychology. Either way, the requirement can be a constructive thing.

To illustrate why, let us consider what cue or signal monkeys may be sensing and using in the memory task just discussed. The likely answer is a cue of trace activation or trace availability. That is, if a probe picture encountered a very active or available trace in memory, the animal would safely respond There. If a probe picture encountered a very inactive trace, the animal would safely respond Not There. If a probe picture encountered an ambiguously or indeterminately active trace, the animal could respond Uncertain. Figure 3 may let readers feel this cue or signal in action. Consider how strong or how active a trace in memory the two pictures contact. Was the picture on the left There in the list of pictures you saw before, Not There, or would you respond Uncertain? What about the picture on the right? Many will say that they remember more clearly the probe on the left. It was the 4th item in the list of pictures given in Figure 1 and so received a boost of stronger original activation (i.e., a recency effect) from being in that position in the list. The animals were likely to respond There when a probe was presented from this original serial position. The probe on the right may have made a dimmer impression on memory because it was presented in the second serial position that generally made the weakest original memory impression. The animals were likely to respond Uncertain to this kind of probe item.




Figure 3. Two probes that might have followed the memory list given in Figure 1. Both would be there probes. The item on the left reprises the item in the list that occurred at the final, strongest, recency-advantaged serial position. The item on the right reprises the item in the list that occurred at the weakest, second serial position. The monkeys' memory probes were computer-presented in color. Reprinted from Smith (in press) by permission of Oxford University Press.

Closely analyzing the relevant cue or signal for behavior has several benefits. For one thing, it lets one consider the cognitive sophistication of a performance and consider whether it is an associative phenomenon based in stimulus control or whether it relies on more sophisticated cues and signals. Regarding the memory task under discussion, one can understand why low-level interpretations based in stimulus control are insufficient to explain the monkeys' performance. Across trials in this task, all stimuli became targets and foils and were rewarded and nonrewarded following both There and Not There responses. No stimulus cue indicated any response. Only the presence or absence of the probe in the preceding list was ever relevant, and so the relevant cue for behavior had to be something like the subjective strength of the trace the probe contacted. Such a cue is cognitively derived and abstract, and is profoundly different from the signals available in many

traditional operant situations, as when a tone or light stably signals rich or poor reinforcement. Thus pinpointing the signal for behavior in this task shows that animals' performance in it lies far from traditional senses of stimulus control. It is ironic that the standards of comparative psychology require this kind of careful information-processing analysis, even while comparative psychologists are not so comfortable when the conclusion of this analysis is for higher-level or cognitively sophisticated performance. This is exactly the situation in which the marriage of comparative and cognitive psychology can be most fruitful, and I return to this issue below.

Pinpointing the signal for behavior also lets one further refine the theoretical constructs in a field and adjust one's theoretical conclusion. For example, suppose we know that animals do examine their memories, asking whether relevant traces are active/available enough, inactive/unavailable enough, or indeterminate, and respond There, Not There, or Uncertain on this basis. Then, by the current constructs of theory in metacognition, are they showing metacognition or metamemory, or not? One can decide the question either way, but either way it is clear that the care dictated by comparative psychology in establishing the behavioral signals is helpful in making this judgment.

The Formal Description

A third constructive element in the comparative program is the strong emphasis placed on modeling animals' performances formally and mathematically. In the present case of animals' memory-monitoring performances, signal detection theory (SDT) provides a constructive framework for thinking about the psychological organization of animals' performances. Here SDT would assume that the items in the list create subjective memory impressions that lie along a continuum of trace strength (the X-axis in Figure 4). Then the probe picture queries the strength of one trace. Probes on Not There trials will generally point to weak traces, perhaps averaging 0.0 plus or minus the scatter of memory variability (the normal distribution NT—Not There—in the figure). Probes on There trials will point to stronger traces on average though still with memory variability (the four T—There—normal distributions in the figure). The items that occurred early or late in the list will on average point to especially strong traces. These four distributions could be estimated to center at the memory sensitivity (d') appropriate to the performance that the animal showed at each serial position (MacMillan & Creelman, 1991, pp. 209-230). Notice that the overlap between the Not There and There distributions is what makes the SPR task difficult and uncertain because it means that probes on There and Not There trials will often seem equally strongly remembered to the animal.

Facing this difficult memory situation, the animal needs to find a way to divide the memory continuum into three response regions using response criteria, hopefully being able to respond Not There, There, or Uncertain when the probe contacts a memory trace that is weak, strong, or ambiguous/indeterminate. Thus SDT assumes a decision process by which criterion lines are placed along the continuum to define response regions. These criterion lines are also shown in Figure 4. As a probe stimulus contacted a trace that fell to the left of the Not There-Uncertain criterion line, to the right of the Uncertain-There criterion line, or between these two, the participant would make the Not There, There, or Uncertain

response, respectively. Applying this model to the sample probes given in Figure 3, the probe on the left might have contacted a trace in memory above the upper criterion and received a There response, whereas the probe on the right might have contacted a trace whose activity lay between the criterion lines and received an Uncertain response.

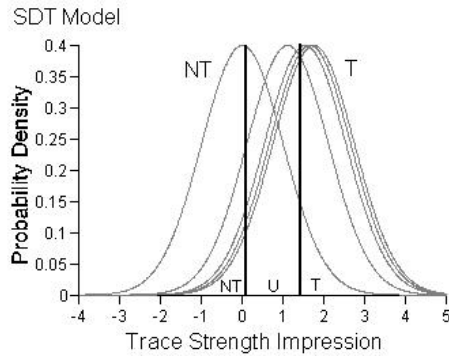


Figure 4. A signal detection theory (SDT) portrayal of Monkey Baker's decision strategy in the serial probe recognition task of Smith et al. (1998). Unit-normal trace-impression distributions are centered at the locations along the trace-strength continuum corresponding to the animal's d' for probes of the four serial positions in the memory lists (T), and at 0.0 for the Not There probes (NT). These normal curves are overlain by the decision criteria that define the animal's three response regions (from left to right, Not There [NT], Uncertain [U], and There [T]). Reprinted from Smith et al. (2003) by permission of Cambridge University Press.

One can conduct simulations to evaluate the SDT model of monkeys' memory-monitoring performances and to find out what decision strategy (i.e., what placement for the two criterion lines) they were probably using. To do so, I evaluated the data patterns produced by many thousands (226,981) of simulated creatures (simulants) who had decision criteria at different points along the trace-strength continuum. Each simulant completed 8,000 trials in a virtual version of the SPR task that the monkeys received. On each trial, the simulant received one of 5 trial types (Not There or a There probe of one of four serial positions), assessed (with memory variability) the trace strength this probe item contacted (following the five probability-density functions shown in Figure 4), and responded according to its criterion placements. The simulant's performance over the 8,000 trials was then summarized and its performance compared mathematically to Monkey Baker's observed performance pattern (Figure 2A). The criterion of best fit was the sum of the squared deviations (SSD) between the observed and simulated percentages. On average, the best-fitting response percentages were within about 3% of their observed targets. Thus the performance of this model was competitive with the performance of other formal models in the experimental literature (Smith & Minda, 1998, 2000).

Figure 2C shows the performance of the simulant that most closely reproduced Baker's performance (compare Figure 2A). Figure 4 shows the criteria for this simulant. One sees that Baker found an adaptive decision strategy. He took the memory test (i.e., he responded Not There or There) on probes that happened to contact quite inactive or quite active traces. He responded Uncertain for indetermi-

nate, ambiguous trace strengths that could easily have been caused by either a There or a Not There probe.

This kind of formal description accomplishes a number of useful ends in the psychological analysis of behavior. First, it supports cross-task and cross-species comparisons among data patterns. In particular, the performance of humans and animals can be summarized in exactly the same formal/mathematical currency and compared directly. Second, this kind of analysis also directly underlies careful studies of the optimality of behavior that figure so prominently in discussions of animal behavior and that are the topic of the next section. Third, this kind of analysis clarifies the formal structure of behavior so that appropriate theoretical perspectives can be brought to bear on it and so that new empirical issues can arise. For example, here the model highlights the fact that animals need to engage in criterion-setting processes or decisional mechanisms that let them organize response regions and choose behaviors. This leads one to ask about the cognitive processes and representations that let the animal choose and adjust response criteria, and about the level of awareness or self-awareness animals and humans have as they carry out these criterion-setting or decisional operations. An additional benefit of the formal model is that it clarifies thinking in an area while still offering a neutral description of performance that is inclusive theoretically because it makes no theoretical commitments toward behaviorism or cognitivism. This is particularly valuable in the area of comparative cognition research where interpretations can be contentious whether they are framed at a high or low cognitive level.

The Adaptiveness of Performance—Optimality

A fourth strength of the comparative program is that it always stays ecologically and situationally grounded. According to the comparative tradition, behaviors do not just happen and cognitive capacities are generally not evolved without a function. Behaviors happen and cognitive capacities exist because they have value for the animal in terms of daily living or value on the longer time scale of evolution for conferring an advantage in reproductive fitness. These kinds of adaptive and evolutionary issues would be constructive for human cognitive psychology to consider more carefully and more frequently. For example, the human metacognitive literature has barely considered these issues. It is mentioned that the capacities for cognitive monitoring and control are good for succeeding on college finals. But when in its evolution did the human species first benefit from the cognitive executive and from the cognitive monitoring and cognitive control it affords? Was this a recent adaptation to living in city states? An earlier adaptation to cooperative hunting? An adaptation common to all social primates? To all foraging generalists? Of course the answer bears directly on whether one thinks that animal species might share the metacognitive capacity with humans. Whatever the answer, that a cognitive system would evolve a conscious, overseeing agent to look in on mental processes to steer and guide them is an extraordinary evolutionary feat and would be an extraordinary evolutionary tale if it could be told.

Short of that grand explanation, the adaptive benefit of different behavioral strategies is often captured in comparative psychology through optimality studies. One can use the formal model just described to illustrate the study of optimality in the research area of metacognition. That is, one can ask how well the animal would do (in rewards per minute, for example) if his response criteria lay at these two

places along the trace-strength continuum, or at some other two places. In effect, one can sample all the possible decision strategies that the animal might use, ask how all of them compare in their reward efficiency, and then consider how relatively well the real monkeys do.

To draw this optimality landscape of the memory monitoring task, I retained the trace-strength continuum shown in Figure 4 and the placement of the five trace-strength distributions along it (for Not There probes and for probes of the four serial positions in the memory list), because trace activation or memory sensitivity is a basic information-processing limit that cannot be increased in the service of greater rewards. Then, I surveyed the reward efficiency of strategies that placed the center of the Uncertain response region at 101 places at each 1% increment along the trace-strength continuum, and, given each center, that widened the Uncertain response region out from having 0 width (zero 1% increments to either side of center) up to 50 width (fifty 1% increments to either side of center). These 5,151 simulants each received 8,000 trials in the virtual SPR task, subject to the trial times, penalty times, and reward structure of the actual task, and responding in accordance with the three response regions in effect.

Figures 5A and 5B show the rewards per minute of simulants that gave the Uncertain response region different centers and widths. B indicates the position in these landscapes of the simulant that performed most similarly to Monkey Baker. Clearly Baker, judging by the performance of his simulated match, centered and widened his Uncertain response region adeptly. He declined those trace strengths that were most indeterminate and that most risked error, and the present formal optimality analysis confirms the near optimality of this decision strategy.

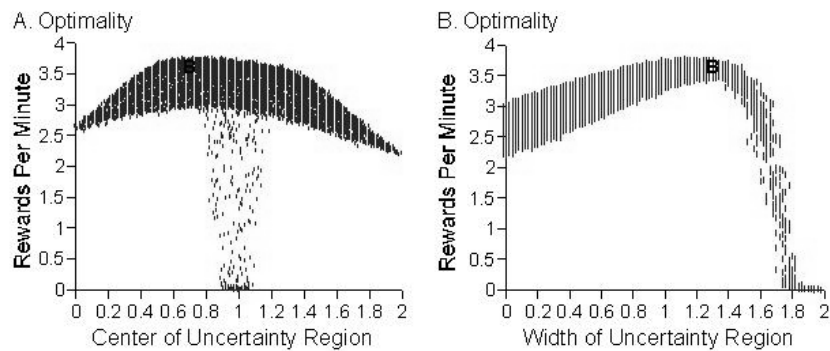


Figure 5. (A) The reward efficiency (in rewards earned per minute) of simulants that centered the Uncertain response region at different places along the trace-strength continuum in a virtual version of Smith et al. (1998) serial probe recognition task. We surveyed the reward efficiency of 5,151 decision strategies when each received 8,000 trials in a simulated version of the task, subject to the trial times, penalty times, and reward structure of the task the monkeys experienced, and using the signal-detection response rule that accorded with the three response regions defined by each simulant's two criterion placements. B represents the position in this optimality space of the simulant that best fit the performance of the real Monkey Baker. (B) The results of the same simulation plotted by the width of the Uncertain response region. Reprinted from Smith et al. (2003) by permission of Cambridge University Press.

Pragmatically, an optimality study like this can facilitate experimental planning in an area. One can preview how different rewards, penalties, and contin-

gencies change the shape of the optimality surface, and one may find experimental parameters that emphasize the value of the metacognitive strategy (for example) over alternatives. This may encourage animal participants to adopt the metacognitive strategy if they can. This is especially constructive when a particular behavioral strategy is cognitively effortful for the animal, as I think the metacognitive strategy is. That is, I think that monkeys gravitate toward an associative, nonmetacognitive performance strategy if they can find an effortless one that earns a decent rate of return. An experiment that maximally rewards the metacognitive approach relative to associative approaches may help persuade animals that the cognitive effort of the former is worthwhile. This same idea extends to the unsuccessful attempts to show a metacognitive capacity in rats and pigeons (Inman & Shettleworth, 1999; Smith & Schull, 1989; Teller, 1989). For these less cognitively sophisticated species, this maximal separation favoring metacognitive monitoring may be even more critical because these species seem to have difficulty expressing the metacognitive capacity at all.

Problems within the Comparative Tradition


The foregoing discussion makes plain the considerable analytic strengths that one encounters within the comparative tradition. There are others as well. Now, though, it is time to consider the other side of the question. Comparative psychology has sometimes been cramped in its interpretations and sometimes deliberately a psychological in its explanations because there has been discomfort with granting animals psychological states and representations that would lie in some unknowable and unstudyable black box in mind. In contrast, cognitive psychology has remained freer to consider these kinds of black-box psychological issues because there have not been such strong theoretical strictures against doing so. In this section of the article, I will note some tendencies that show themselves when the perspectives typically used to explain animal behavior are applied too forcefully or rigidly.

First, there is a risk within the comparative tradition that the optimality of an animal's performance will be taken as a sufficient explanation of the behavior. That is, there is the temptation to say: "Of course animals respond Uncertain for the hardest trials—they avoid timeouts and earn more rewards per minute that way." Now in a sense optimality or reward efficiency is a distal, economic or motivational description of behavior. However, the real goal of cognitive science is to find the proximal, psychological explanation of behavior framed in terms of processes and representations. An optimality analysis does not and, in principle, cannot provide this kind of explanation because it does not explain how the behavior is organized, represented, or understood by the organism. A toaster makes optimal toast. A French chef makes optimal food. But these are completely different acts contrived by completely different systems. Monkeys might perform some optimal responses in a low-level, associative way, but choose other optimal responses in even a consciously deliberate way. The psychological description of these responses would have to be completely different, too.

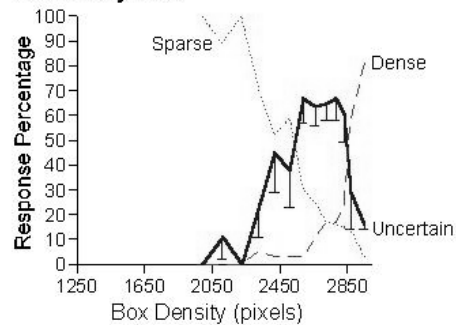
Second, let us consider the SDT model of memory performance in which two criterion lines organize three response regions along a trace-strength continuum. There is also a risk within the comparative tradition that a formal model of performance like this will be taken as a sufficient, low-level explanation of the be-

havior. That is, there is also the temptation to say: "So animals just divide up the continuum using two criterion lines and make three response regions. Easily modeled. Therefore low-level." This is not an appropriate conclusion. A formal model is not a sufficient psychological description of behavior any more than an assessment of optimality is. Moreover, neither easily modeled, nor the use of criterion lines, nor the presence of response regions, has any necessary relation to a low-level explanation of performance. The SDT model does not say anything about the psychological level of the behavior or the animal's understanding of it. The same model would apply to low-level, automatic behavior as to high-level conscious behavior. The model is moot on this fundamental point. The connection that some comparative psychologists would hope to make in this case is wrong and sidesteps important matters of psychological interpretation.


A. Trial Screen



B. Monkey Abel



C. Monkey Baker



D. Humans

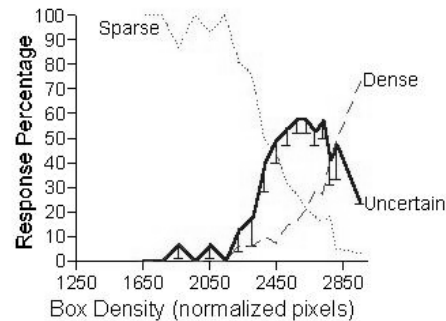


Figure 6. (A) The screen from a trial in the dense-sparse discrimination of Smith et al. (1997). (B) The performance of Monkey Abel in the dense-sparse task. The dense response was correct for boxes with exactly 2950 pixels—these trials are represented by the rightmost data point for each curve. All other boxes deserved the sparse response. The horizontal axis indicates the pixel-density of the box. The solid line represents the percentage of trials receiving the Uncertain response at each density level. The error bars show the lower 95% confidence limits. These were calculated (Hays 1981, pp. 224-26) using the total Uncertain responses as a proportion of total trials at each density level (summing across multiple sessions by the animal). The percentages of trials ending with the dense response (dashed line) or sparse response (dotted line) are also shown. (C) The performance of Monkey Baker in the dense-sparse discrimination depicted in the same way. (D) The performance of seven humans in the dense-sparse discrimination depicted in the same way (here the error bars were calculated using the total Uncertain responses as a proportion of total trials at each density level summing across humans who completed one session each). To equate discrimination performance across participants, the data have been normalized to place each participant's discrimination crossover at a pixel density of about 2700. Figures 6A, 6B and 6C, reprinted from Smith et al. (1997) by permission of Elsevier. Figure 6D reprinted from Smith et al. (2003) by permission of Cambridge University Press.

Third, there is a temptation in the comparative tradition to treat all the diverse cues and signals that can prompt or control behavior as theoretically or psychologically equivalent. It is praiseworthy to insist on isolating the cue that signals the organism to behave in particular ways and patterns. This fosters strong experimental method and can lead to important theoretical insights. However, once that cue is found, there is the temptation to say: "Now we have found the cue that is exerting control over the animal's behavior. So it is just another case of stimulus control over behavior." This statement also skirts important psychological issues. Cues vary widely in their sensory immediacy and in their psychological level. The color red can control a behavior. The belief that red ties make one seem powerful can control the choice of a red tie. The belief that needing a red tie to look powerful makes one seem weak can control the choice of a yellow tie. These are completely different kinds and levels of psychological cues or signals that cannot be given the same kind of psychological explanation and that cannot be lumped together as instances of control by stimulus color. Yet comparative psychologists sometimes collapse, in less extreme settings, different kinds of psychological states, cues, and signals, and try to make them all fit the same low-level cue-control framework that has been dominant for a long time. They may well not all fit that framework. They will need their own, independent psychological descriptions that consider the kinds of intervening representations involved, the cognitive level or sophistication of the relevant processes, and even perhaps the place in the animals' immediate memory or immediate awareness that these representations or processes occupy.

On this important point I would like to give a concrete example from my own research. In an early uncertainty-monitoring experiment, we gave human and monkey participants a visual density-discrimination task. Participants used a joystick to move a cursor to one of three objects on a computer screen (Figure 6A). Moving the cursor to the Box (a Dense response) was correct if the box contained exactly 2950 illuminated pixels. Choosing the S (a Sparse response) was correct if the box contained any fewer pixels. Choosing the Star (an Uncertain response) allowed participants to decline the trial and move into a new, guaranteed-win trial. At mature performance, trial difficulty was adjusted to ensure that animals were experiencing many trials that lay near their perceptual threshold for discriminating Dense from Sparse.

Figures 6B-D shows the performance of two monkeys and a group of seven humans. Dense responses predominated on Dense trials and the most difficult Sparse trials. Sparse responses predominated on the sparser trials. The primary discrimination was performed at chance where these two response curves cross, and the Star was used most in this region of maximum uncertainty. Monkeys and humans assessed accurately when they were liable to make an error in the primary discrimination and they bailed out of those trials selectively and adaptively. In fact, these graphs show one of the strongest performance similarities between humans and animals in the comparative literature. Moreover, humans' descriptions of performance revealed the decisional organization of the task for them. They said that they used the Sparse or Dense responses when they thought they knew what kind of stimulus they were seeing. They said that they reserved the Uncertain response for when they did not know or thought they could not tell the answer on a trial.

Yet one does not have to interpret performance in this way, and theorists of animal behavior would shy away from this kind of interpretation. Instead, the given explanation would emphasize the continuum of stimuli from Sparse to Dense, the possibility that the three behaviors were controlled by three regions along the continuum (sparse, medium, dense), and the conclusion that the monkeys' performance was really just another instance of stimuli controlling behavior.

This emphasis would wrongly trivialize the difficult psychological problem that humans and animals face in this task. The problem is that this perceptual-threshold task deliberately challenges the observer's discrimination ability. Given the variable impressions made by the same objective stimulus event on different occasions, true Dense trials and threshold Sparse trials will often produce exactly the same perceptual impression in the animal's mind. Therefore, the impression of density cannot reliably tell the animal what to do. That is, it cannot control behavior in the normal sense because the perceptual impressions are inherently ambiguous. Shiffrin and Schneider (1977, pp. 167-168) also described how ambiguous mental representations of stimuli would map inconsistently onto behavioral responses, making those representations poor indicators of what the organism should do, and making it impossible to safely associate a response to those stimuli. Instead, in a situation of inconsistent mapping like this, higher levels of cognitive processing (that have come to be called attentional or controlled processes in the cognitive literature) would be needed to resolve the indeterminacy and produce a decision about behavior.

An example will let us distinguish stimulus-controlled and decisionally controlled processes. Stoplights (green-go; red-stop) are a critical real-world conditional discrimination that many humans face every day. The task of the stoplight is consistently mapped. The underlying perceptual representations (green and red) are not confuseable and they dictate absolutely reliably the appropriate behavior. Responses in the stoplight task may be triggered reflexively and automatically just because the task's consistent mapping allows stimulus and response to associate so strongly. In fact, this is the point of stoplights. They need no decision criteria or decision making.

In contrast, imagine if traffic lights gradually morphed between red and green (with red changing to green—pixel by pixel—on the light's surface) so that drivers had to decide whether their light was green enough to go. (Leave out of this example the yellow light that causes us all to accelerate.) The red-to-green situation would require decision criteria and decision making. It would involve controlled processing—slow, attentional, and capacity intensive—about the situation presented by each trial. The situation would also be a nightmare, as orthogonal, hurried travelers applied self-serving criteria in the press of the moment and crashed. This is the kind of judgment humans and animals face in the Dense-Sparse task. They must ask whether the box is dense or sparse enough to try. It is actually an important fact that this cognitive analysis—wherein indeterminate stimulus response mappings encourage controlled decision-making processes—applies no matter the participant species. This realization provides a theoretically principled way to grant animals' uncertainty responses some of the cognitive sophistication they may deserve.

However, the point of this example is that the behavioral perspective that is so important within comparative psychology would tend to coalesce under the

general rubric of stimulus control both a reflexive, stimulus driven response modality borne of consistent mapping and a controlled, decisional, response modality borne of inconsistent mapping. In contrast, cognitive psychology would be friendly to the idea that these response modalities would be different, could occupy different levels in the psychological system, and could even occupy different places and prominences in the organism's immediate memory, executive functions, or awareness. This is an example wherein cognitive psychology would preserve a distinction with a critical difference, whereas comparative psychology would tend to summarize away that distinction following general associative guidelines.

The Philosophical Roots of Comparative Psychology's Self-Imposed Constraints

The theoretical cast of comparative psychology that I have been discussing has a long history and a distinguished pedigree within the field. It goes back to the ideas of behaviorists like Thorndike (1911) and Morgan (1906). The general idea, expressed most famously by Morgan's canon regarding the interpretation of animal behavior, was that one should never interpret an organism's behavior at a higher psychological level if one can interpret that behavior at a lower level. Consequently, given metacognitive-like performances by animals such as those considered in this article, there is a 100-year-old passion to refute them or to dismiss them as associative, conditioned phenomena.

Yet it is critically important to realize that this 100-year-old tendency can be mistakenly applied and can have negative consequences that, despite the craft and elegance of comparative psychology, sometimes have constrained its reach and impeded its progress. I will close by considering this issue.

Morgan hoped by his canon to confer a simplicity or a parsimony on interpretation and explanation in comparative psychology. But one can understand from looking at Figure 6 why the parsimony seemingly embodied in Morgan's canon is false when applied to the data patterns shown there. Humans perform just like monkeys do. Humans are declaratively uncertain as they do so. Humans and monkeys have shared much of their evolutionary histories, especially including the fitness matrices that could have prompted the emergence of an uncertainty-monitoring cognitive system. Humans and monkeys even share homologous brain structures that could provide the neurological substrate for this system.

For these reasons it is unparsimonious to interpret the same graph produced by humans and monkeys in qualitatively different ways—consciously metacognitive vs. low-level associative. It uses two opposed behavioral systems to produce the same phenomenon when one might do. In fact, this duality of interpretation is not even an appropriate scientific stance in a case like this. Imagine if, in any other domain, a researcher showed identical graphs by two populations (e.g., young vs. aged adults, nondepressed vs. depressed individuals) and then nonetheless offered qualitatively different high-level and low-level cognitive interpretations to explain them. The researcher would have no warrant to do so. There is actually no warrant to do so in the case of human and animal performance, either. To the contrary, correct scientific inference in such a case would require that the two graphs be provisionally interpreted as instances of the same phenomenon—a phenomenon that humans describe very clearly in metacognitive terms—unless and

until other evidence required the monkey graph to be given a qualitatively different and lower interpretation. Indeed, in a case like this, the burden of proof would actually fall on the behaviorist theorist who would try to interpret the monkeys' performance in a low-level, associative manner.

It is an extraordinary thing that this duality of high-level and low-level interpretations is given momentary credence when the two populations are humans and animals. However, this is an accident of the history of comparative psychology that serves the field poorly—from the perspectives of parsimony and of scientific inference—when humans and animals perform identically in paradigms such as those under consideration here.

Meanwhile, the preconceived downgrading of animals' performance has consequences for the character of interpretation in comparative psychology and for the conduct of science within it. This preconceived idea helps encourage us to suffice with formal/mathematical, distal explanations of behavior because we are less comfortable with engaging fully the problems of psychological process and representation that lie behind the phenomena we model. For the same reasons the preconceptions of comparative make distal descriptions of behavior in terms of optimality or fitness seem more attractive and sufficient. The same stance leads us to coalesce diverse phenomena downwards toward the lowest level associative denominator, even when, as we saw in the contrast between automatic and controlled processes, the psychological structure of the diverse phenomena are different and have different psychological organizations behind them.

Moreover, by defending Morgan's brand of parsimony, many studies of animal behavior have *de facto* hindered the development of a true comparative psychology, because they have emphasized just the kinds of explanations that apply least to humans. From the perspective of the sociology of science, this downward interpretative bias, and its inherent species exclusivity, has aided comparative psychology's downward spiral in the academy. The courses are fewer. The textbooks on learning are aging and growing "classic." From a scientific standpoint, one doesn't have to care about this, but one will if one believes that the issues of comparative psychology are lasting and foundational. The loss of interest in this field arises because so much of the theory applied to animal behavior has been so deliberately kept at theoretical arms length from the kinds of cognitive, decisional, and metacognitive issues that would apply to the human cognitive system and that would interest broadly human students of animal cognitive systems.

There is an old idea from clinical psychology that defense mechanisms start out serving a healthy purpose for the organism, but that gradually they become overused, controlling, and harmful. Comparative psychology illustrates this principle well from the perspective of the philosophy of science. It was probably critical in the early history of comparative psychology—that is, in a climate of anecdotal accounts of intelligent animal behavior and generous inferences about animal "mind" (e.g., Romanes, 1894; reacted strongly against by Thorndike, 1911)—to enforce a high threshold for accepting high-level, deliberate, or decisional explanations of animal behavior. This kept anecdotes and just-so stories out of the literature. It gave methodology and inference a sharp edge. It kept a theoretical lid on things.

But this early and valuable protective stance gradually became a systematic bias in studies of animal behavior. This cannot have helped but exclude the

consideration of some phenomena, color the interpretations of phenomena that are studied, affect the framing of research questions, and partially determine the character and the content of the articles that appear in the journals. It may be a good time, as comparative psychology enters its second century, for comparative psychologists to reevaluate their position on their discipline's historically preconceived notions.

As one route to this reevaluation, comparative psychology might benefit from becoming deliberately more comparative. It is striking how many comparative articles are noncomparative in the sense of being about only one species. The single-species character of much of comparative psychology has encouraged a focus on low-level, associative kinds of parsimony because that kind of explanation will often suffice for one kind of animal participant doing one kind of discrimination task. However, a truly comparative empirical picture (for example, Figures 2 and 6 in this article) makes different explanatory demands on the scientist. In these figures one sees human and monkey participants performing uncertainty-monitoring tasks in the domains of perception and memory. The phylogenetic and information-processing breadth of this data pattern is important. It makes clear that the kind of integrative parsimony and simplicity of explanation one seeks when explaining the performance of several species in several tasks is different from the parsimony one seeks when explaining the performance of a single species in a single task. More multiple-species studies might foster interest in more integrative kinds of parsimony that could make important contributions to theory in the comparative literature in coming years.

Synthesizing the Comparative and Cognitive Traditions

Cognitive psychology does not completely escape criticism in these areas. The problem that our comparisons of human and animal performance are often polarized is a two-way street. We usually describe humans' conscious, verbal metacognition in a way that definitely excludes animals, even as behaviorist theorists try to describe animals' similar performances using low-level, descriptions that do not fit the human case. The same polarization colors discussions of tool use, language, and so forth. There is real language as opposed to complex sequential discriminations. There is real tool use vs. local or situational enhancement and trial/error learning. Indeed, sometimes when animals seem to meet some competence criterion both theoretical camps quietly raise the standard animals must meet in order to show the capacity. This has the effect (ironically) of keeping comparative phenomena incomparable.

Yet in cases like this the issue does not have to reduce to either elevating humans or dismissing animals. Rather, given truly comparative data like that shown in Figures 2 and 6, and clearly equivalent performance by humans and animals, it could be possible to strive for a common psychological interpretation that could acknowledge both the strong processing similarities across species and seek an information-processing middle ground that lies between merely associative and fully conscious interpretations.

The uncertainty-monitoring tasks considered in this article illustrate how this integrative approach might begin. One would acknowledge that the psychophysical procedures and uncertainty-monitoring tasks ensure indeterminate stimulus-response mappings and encourage controlled decision-making processes—for

both humans and animals. One would construe the uncertainty response as one manifestation of the controlled indeterminacy-resolution processes that are necessary—for both species—near threshold. There is a lot of common cognitive ground here that could prompt further research and careful information-processing analyses. In the end, it might still prove necessary to conclude for both common processing principles and experiential differences in some way that was dictated by the whole empirical picture. But this conclusion would be science- or data-driven, not driven by either cognitive psychology's or comparative psychology's prior assumptions.

In moving toward this integrative comparative psychology, it is fair to say that cognitive psychology is ahead and better prepared. Cognitive psychology allows for cues and signals of different character and sophistication. Cognitive psychology allows for qualitatively different kinds of processes in mind, that we sometimes call automatic and effortful, or automatic and controlled. Cognitive psychology allows for processes occurring at lower, object levels in mind, and at higher, meta levels in mind. Cognitive psychology allows for the possibility of consciousness and self-awareness. Because cognitive psychology grants humans a wide variety of perceptual, memory, and cognitive processes, it is flexibly prepared to ask which of these capacities members of other species have and which not.

But this preparedness by cognitive psychology does not mean that comparative psychology brings little to the table. To the contrary, most of this article was devoted to showing that comparative brings elegant tools of careful science, crisp operationalization, clear inference, and cautious interpretation. These aspects of comparative psychology have deepened enormously my understanding of the metacognitive phenomenon I study—including that phenomenon as it is instantiated by the human mind.

Moreover, one can see that comparative psychology is presently moving in positive directions. This movement can be seen in the innovative research of Reiss and Marino (2001) on the mirror dye-test applied to dolphins, in the elegant writings of Rumbaugh and his colleagues on an emergent level of cognition in the primate mind (Rumbaugh, Savage-Rumbaugh, & Washburn, 1996), in the studies of Povinelli and his colleagues on theory of mind (e.g., Povinelli, Parks, & Novak, 1991), in the studies by Boysen and Washburn and their colleagues on numerical and symbolic functioning in primates (Boysen, Berntson, Hannan, & Cacioppo, 1996; Washburn, 1994), and in the theoretically challenging research on episodic memory in animals (Menzel, 1999; Schwartz et al., 2002). One could cite other examples as well. One can also see this positive movement in the 21 commentaries written in response to our BBS article on the comparative study of metacognition (Smith et al., in press). These commentaries represent a serious, careful dialog that is truly comparative in spirit and cognitive in its interpretative framework. Only occasionally in that dialog did the quiet historical bias toward behaviorism still show. Thus I am confident that in coming years many laboratories will continue to integrate cognitive psychology's flexible information-processing vision with comparative psychology's great skill and inferential sharpness. And I think this will lead to an exciting time and to a synergy that will show comparative cognition research at its best.

References

- Boysen, S. T., Berntson, G. G., Hannan, M. B., & Cacioppo, J. T. (1996). Quantity-based interference and symbolic representations in chimpanzees (*Pan troglodytes*). *Journal of Experimental Psychology: Animal Behavior Processes*, **22**, 76-86.
- Brown, A. S. (1991). A review of the tip-of-the-tongue experience. *Psychological Bulletin*, **109**, 204-223.
- Brown, A. L., Bransford, J. D., Ferrara, R. A., & Campione, J. C. (1982). Learning, remembering, and understanding. In J. H. Flavell & E. M. Markman (Eds.), *Handbook of child psychology, Vol. 3* (pp. 77-164). New York: Wiley.
- Castro, C. A., & Larsen, T. (1992). Primacy and recency effects in nonhuman primates. *Journal of Experimental Psychology: Animal Behavior Processes*, **18**, 335-340.
- Dunlosky, J., & Nelson, T. O. (1992). Importance of the kind of cue for judgments of learning (JOL) and the delayed JOL effect. *Memory and Cognition*, **20**, 374-380.
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, **34**, 906-911.
- Gelman, R., & Gallistel, C. R. (1978). *The child's understanding of number*. Cambridge, MA: Harvard University Press.
- Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, **98**, 5359-5362.
- Hart, J. T. (1965). Memory and the feeling-of-knowing experiments. *Journal of Educational Psychology*, **57**, 347-349.
- Hays, W. L. (1981). *Statistics*. New York: CBS College.
- Inman, A., & Shettleworth, S. J. (1999). Detecting metamemory in nonverbal subjects: A test with pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **25**, 389-95.
- Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, **100**, 609-639.
- MacMillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide*. Cambridge, UK: Cambridge University Press.
- Menzel, C.R. (1999) Unprompted recall and reporting of hidden objects by a chimpanzee (*Pan troglodytes*) after extended delays. *Journal of Comparative Psychology*, **113**, 426-34.
- Morgan, C. L. (1906). *An introduction to comparative psychology*. London, UK: Walter Scott.
- Nelson, T. O. (Ed.). (1992). *Metacognition: Core readings*. Toronto, Canada: Allyn and Bacon.
- Nelson, T. O. (1996). Consciousness and metacognition. *American Psychologist*, **51**, 102-116.
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation*, **26**, 125-141.
- Povinelli, D. J., Parks, K. A., & Novak, M. A. (1991). Do rhesus monkeys (*Macaca mulatta*) attribute knowledge and ignorance to others? *Journal of Comparative Psychology*, **27**, 318-325.
- Reiss, D., & Marino, L. (2001). Self-recognition in the bottlenose dolphin: A case of cognitive convergence. *Proceedings of the National Academy of Sciences U.S.A.*, **98**, 5937-5942.
- Roberts, W. A., & Kraemer, P. J. (1981). Recognition memory for lists of visual stimuli in monkeys and humans. *Animal Learning and Behavior*, **9**, 587-594.
- Romanes, G. J. (1884). *Mental evolution in animals*. New York: Appleton.
- Rumbaugh, D. M., Savage-Rumbaugh, E. S., & Washburn, D. A. (1996). Toward a new outlook on primate learning and behavior: Complex learning and emergent processes in comparative perspective. *Japanese Psychological Research*, **38**, 113-125.
- Sands, S. F., & Wright, A. A. (1980). Primate memory: Retention of serial list items by a rhesus monkey. *Science*, **209**, 938-939.
- Schwartz, B. L. (1994). Sources of information in metamemory: judgments of learning and feelings of knowing. *Psychonomic Bulletin and Review*, **1**, 357-375.
- Schwartz, B. L., Colon, M. R., Sanchez, I. C., Rodriguez, I. A. & Evans, S. (2002). Single-trial learning of "what" and "who" information in a gorilla (*Gorilla gorilla gorilla*): Implications for episodic memory. *Animal Cognition*, **5**, 85-90.
- Shields, W. E., Smith, J. D., & Washburn, D. A. (1997). Uncertain responses by humans and rhesus monkeys (*Macaca mulatta*) in a psychophysical same-different task. *Journal of Experimental Psychology: General*, **126**, 147-164.

- Shields, W. E., Smith, J. D., Guttmanova, K., & Washburn, D. A. (2003). *Confidence judgments by humans and rhesus monkeys*. Paper submitted for consideration.
- Shiffrin, R.M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, **84**, 127-190.
- Smith, J. D. (in press). Studies of uncertainty monitoring and metacognition in animals and humans. To appear in H. S. Terrace & J. Metcalfe (Eds.). *The missing link in cognition: Origins of self-knowing consciousness*. Oxford, UK: Oxford University Press.
- Smith, S. M., Brown, J. M., & Balfour, S. P. (1991). TOTimals: A controlled experimental method for studying tip-of-the-tongue states. *Bulletin of the Psychonomic Society*, **29**, 445-447.
- Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: the early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **24**, 1411-1436.
- Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **26**, 3-27.
- Smith, J. D., & Schull, J. (1989). *A failure of uncertainty monitoring in the rat*. Unpublished data.
- Smith, J. D., Schull, J., Strote, J., McGee, K., Egnor, R., & Erb, L. (1995). The uncertain response in the bottlenosed dolphin (*Tursiops truncatus*). *Journal of Experimental Psychology: General*, **124**, 391-408.
- Smith, J. D., Shields, W. E., Allendoerfer, K. R., & Washburn, D. A. (1998). Memory monitoring by animals and humans. *Journal of Experimental Psychology: General*, **127**, 227-250.
- Smith, J. D., Shields, W. E., Schull, J., & Washburn, D. A. (1997). The uncertain response in humans and animals. *Cognition*, **62**, 75-97.
- Smith, J. D., Shields, W. E., & Washburn, D. A. (in press). The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences*.
- Son, L. K., & Kornell, N. (in press). Confidence judgments in rhesus macaques: Does self-awareness follow? In H. S. Terrace & J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-knowing consciousness*. Oxford, UK: Oxford University Press.
- Teller, S. A. (1989) *Metamemory in the pigeon: Prediction of performance on a delayed matching to sample task*. Unpublished undergraduate thesis, Reed College, Portland, Oregon.
- Thorndike, E. L. (1911). *Animal intelligence*. New York: Hafner Publishing Company.
- Washburn, D. A. (1994). Stroop-like effects for monkeys and humans: Processing speed or strength of association? *Psychological Science*, **5**, 375-379.
- Wright, A. A., Santiago, H. C., Sands, S. F., Kendrick, D. F., & Cook, R. G. (1985). Memory processing of serial lists by pigeons, monkeys, and people. *Science*, **229**, 287-289.

Received January 10, 2003.

Revision received April 28, 2003.

Accepted May 18, 2003.