

COMPARATIVE COGNITION & BEHAVIOR REVIEWS

2009

Volume 4, pp 1 -16

Metacognition in animals

Jonathon D. Crystal & Allison L. Foote
University of Georgia

Metacognition is thinking about thinking. There is considerable interest in developing animal models of metacognition to provide insight about the evolution of mind and a basis for investigating neurobiological mechanisms of cognitive impairments in people. Formal modeling of low-level (i.e., alternative) mechanisms has recently demonstrated that prevailing standards for documenting metacognition are inadequate. Indeed, low-level mechanisms are sufficient to explain data from existing methods. Consequently, an assessment of what is 'lost' (in terms of existing methods and data) necessitates the development of new, innovative methods for metacognition. Development of new methods may prompt the establishment of new standards for documenting metacognition.

Keywords: Metacognition, comparative metacognition, uncertainty monitoring, metamemory, quantitative modeling.

A defining feature of human existence is the ability to reflect on one's own mental processes, termed metacognition (Descartes, 1637; Metcalfe & Kober, 2005). Consequently, a fundamental question in comparative cognition is whether nonhuman animals (henceforth animals) have knowledge of their own cognitive states (Smith, Shields, & Washburn, 2003). Answering this question not only provides critical information about the evolution of mind (Emery & Clayton, 2001), but also provides a potential framework for investigating the neurobiological basis of cognitive impairments in people (Hoerold et al., 2008; Robinson, Hertzog, & Dunlosky, 2006; Shimamura & Metcalfe, 1994).

The presentation of a stimulus gives rise to an internal representation of that stimulus (which is referred to as the primary representation). Primary representations are the basis for many behaviors. For example, when presented with an item on a memory test, it is possible to evaluate familiarity with the item to render a judgment that the item is new or old. Metacognition involves a secondary representation

which operates on a primary representation. For example, a person might know that he does not know the answer to a question, in which case appropriate actions might be taken (such as deferring until additional information is available). To document metacognition, we need a method that assigns performance to the secondary representation (i.e., we need to be certain that performance is not based on the primary representation).

Carruthers (2008) distinguishes between first-order explanations and metacognition. First-order explanations are "world-directed" rather than "self-directed" according to Carruthers. According to this view, first-order explanations are representations about stimuli in the world (i.e., beliefs about the world), whereas metacognition involves representations about beliefs (i.e., knowing that you hold a particular belief). Note that according to the definition provided above, metacognition involves knowledge about one's cognitive state. Thus, a variety of other conditional arrangements would not constitute metacognition. For example, discriminating an internal, physiological state would not constitute metacognition. Similarly, discriminating hierarchical relations between a variety of stimuli and responses (e.g., occasion setting) would not constitute metacognition. Carruthers argues that putative metacognitive phenomena in animals may be explained in first-order terms¹.

With human participants, we can ask people to report about their subjective experiences using language. Self

Jonathon D. Crystal and Allison L. Foote, Department of Psychology, University of Georgia, Athens GA 30602-3013

This work was supported by National Institute of Mental Health grant R01MH080052 to JDC. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of Mental Health or the National Institutes of Health. Correspondence concerning this article should be addressed to JDC at jcrystal@uga.edu.

reports of subjective experiences play a prominent role in investigations of metacognition in people (Nelson, 1996). Although these reports may not be perfect, they provide a source of information that is not available from nonverbal animals. Consequently, the difficult problem of assessing metacognition in animals requires the development of behavioral techniques from which we may infer the existence of metacognition. A frequent approach is to investigate the possibility that an animal knows when it does not know the answer to a question; in such a situation, an animal with metacognition would be expected to decline to take a test, particularly if some alternative, desirable outcome is available. Importantly, it is necessary to rule out simpler, alternative explanations. In particular, we need to determine that the putative case of metacognition is based on a secondary representation rather than on a primary representation. For example, if principles of associative learning or habit formation operating on a primary representation may account for putative metacognition data, then it would be inappropriate to explain such data based on metacognition (i.e., based on a secondary representation); the burden of proof favors primary representations, by application of Morgan's canon (Morgan, 1906). We shall refer to explanations that apply primary representations without appeal to secondary representations as simpler or low-level alternative hypotheses to metacognition. Such considerations raise the question of the standards by which putative metacognition data are to be judged. A standard specifies criteria that must be met to infer metacognition using methods that cannot be explained by simpler, alternative hypotheses. We recognize that the details of an alternative hypothesis need to be specific (and specification is provided below), but it is worth recognizing that alternatives to metacognition are simpler (i.e., only primary representations are required). We also note that use of a complex experimental task does not imply that data from such a task require a complex explanation (e.g., a secondary representation). From our perspective, the main issue is the appropriateness of appealing to a complex proposal. Thus, the purpose of testing a less complex proposal is to determine if the output of the low-level model can account for the data. If the output of the model accounts for the data, then it is not appropriate to select the more complex proposal to explain the data (absent an independent line of evidence that cannot be explained by the low-level model). Thus, it is ill-advised to choose to not apply the low-level model because of claims that the primary task is sophisticated, especially if the low-level model can produce the observed pattern of data.

It has long been recognized that an animal might learn to decline difficult tests by discriminating the external stimuli that are associated with such tests (Inman & Shettleworth, 1999); we refer to this class of explanations as a stimulus-

response hypothesis (i.e., in the presence of a particular stimulus, do a specific response). Consequently, an important standard by which to judge putative metacognition data emerged (Inman & Shettleworth, 1999), according to which task accuracy provided an independent line of evidence for metacognition.

The goal of this article is to apply low-level explanations of putative metacognition data to a broad series of experiments in this domain. We find that existing experiments on uncertainty monitoring can be explained by low-level explanations without assuming metacognition.

Predictions about task accuracy

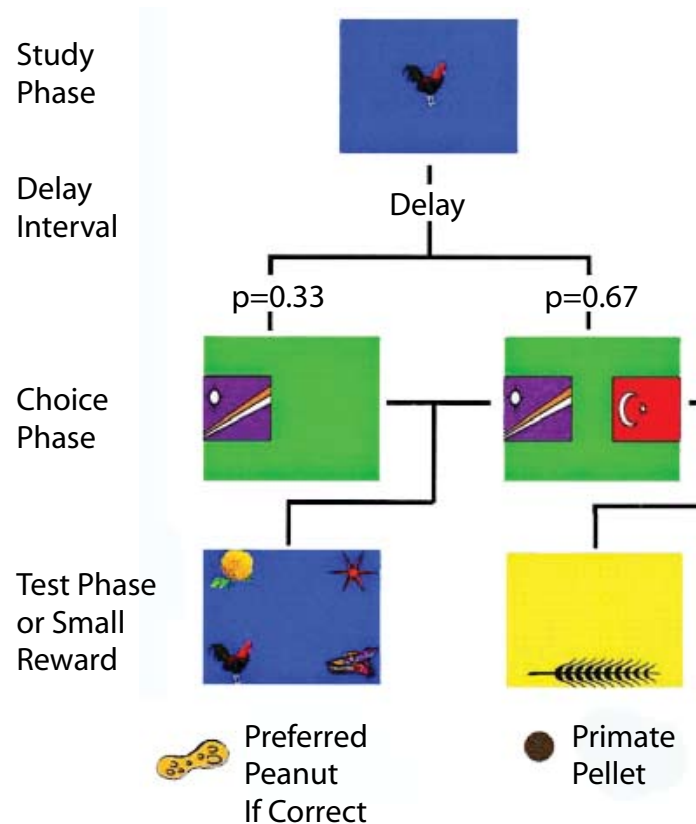
An influential article by Inman and Shettleworth (1999) introduced the idea that it is critical to assess accuracy with and without the opportunity to decline difficult tests. They argued that an animal without metacognition would have the same level of accuracy when tested with and without the opportunity to decline tests (we note that this hypothesis has recently been challenged by quantitative modeling (Smith, Beran, Couchman, & Coutinho, 2008), as discussed below). Inman and Shettleworth hypothesized that an animal with metacognition should have higher accuracy when it chooses to take a test compared with accuracy when it is forced to take the test. The rationale for this hypothesis follows: If the animal 'knows that it does not know' the correct response, then it will decline the test; moreover, being forced to take a test is likely to degrade performance because forced tests include trials that would have been declined had that option been available.

Thus, the prevailing standard since Inman and Shettleworth (1999) includes two criteria: (1) the frequency of declining a test should increase with the difficulty of the task and (2) accuracy should be higher on trials in which a subject chooses to take the test compared with forced tests, and this accuracy difference should increase as task difficulty increases (we refer to this latter pattern as the Chosen-Forced performance advantage). Inman and Shettleworth also emphasized that it is necessary to impose the choice to take or decline the test before being presented with the test.

Representative data

We show two examples of data that meet the prevailing standard, from a rhesus monkey (*Macaca mulatta*) (Hampton, 2001) and rats (*Rattus norvegicus*) (Foote & Crystal, 2007). Hampton (2001) used daily sets of four clip-art images in a matching to sample procedure (i.e., reward was contingent on selecting the most recently seen image from a set of distracter images). The procedure is outlined in Figure

A



B




Figure 1. Schematic representation of design of study and data. Procedure for monkeys (left panel; Hampton, 2001): After presentation of a clip-art image to study and a retention-interval delay, a choice phase provided an opportunity for taking or declining a memory test; declining the test produced a guaranteed but less preferred reward than was earned if the test was selected and answered correctly (test phase); no food was presented when a distracter image was selected in the memory test. Items were selected by contacting a touch-sensitive computer monitor. Data (right side; Hampton, 2001): Performance from a monkey that both used the decline response to avoid difficult problems (i.e., relatively long retention intervals) and had a Chosen-Forced performance advantage that emerged as a function of task difficulty (i.e., accuracy was higher on trials in which the monkey chose to take the test compared with forced tests, particularly for difficult tests). Filled squares represent the proportion of trials declined, and filled and unfilled circles represent proportion correct on forced and chosen trials, respectively. Error bars represent standard errors. (Adapted from Hampton, R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences of the United States of America*, 98, 5359-5362. © 2001 The National Academy of Sciences. Reprinted with permission.)

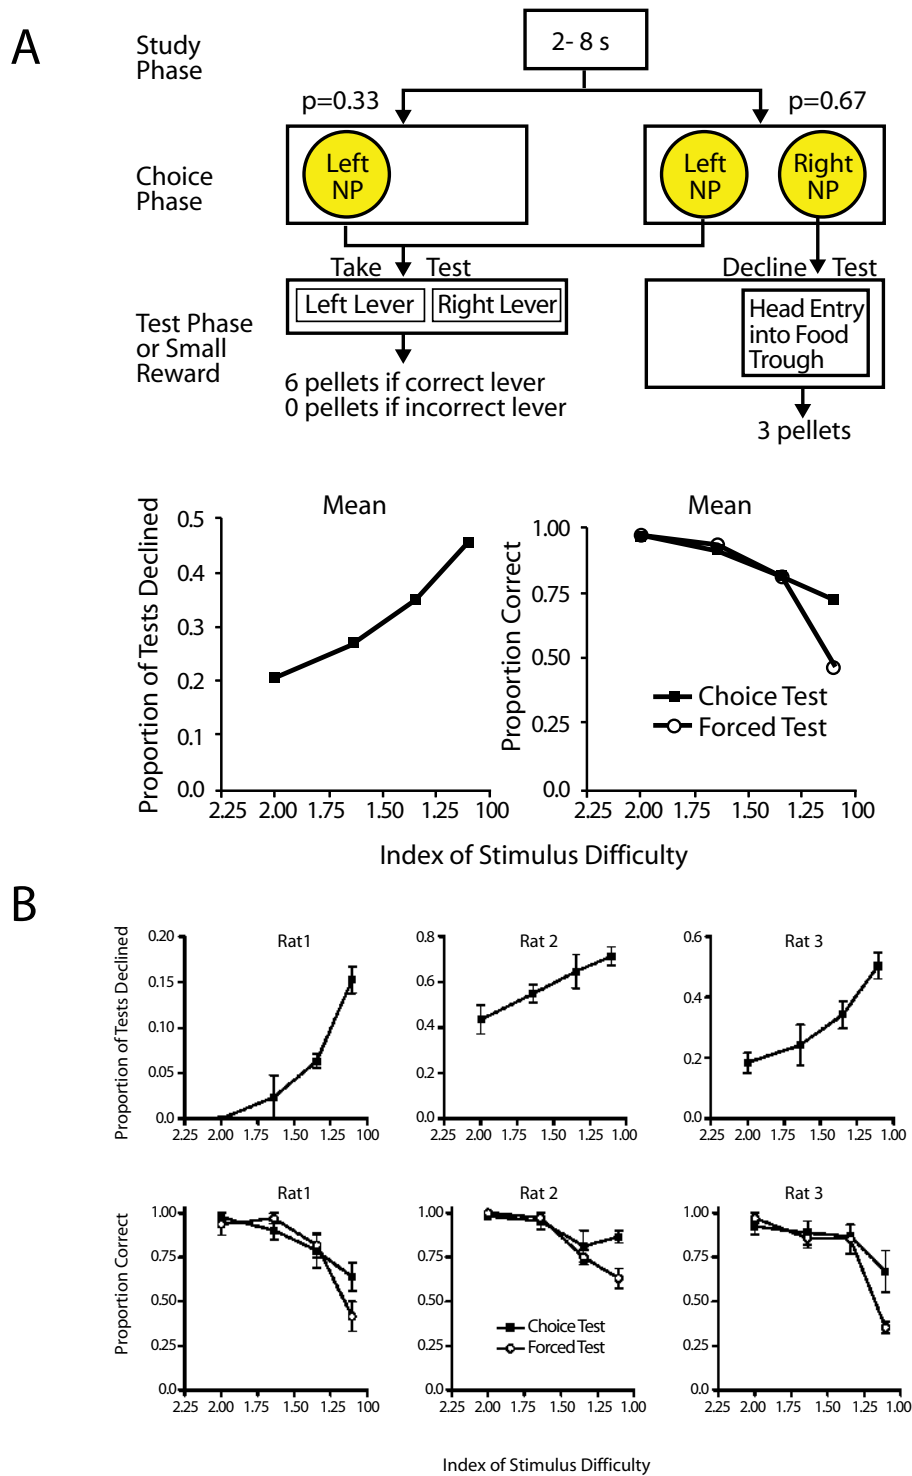


Figure 2. Procedure for rats (top left panel; Foote & Crystal, 2007): After presentation of a brief noise (2-8 s; study phase), a choice phase provided an opportunity for taking or declining a duration test; declining the test produced a guaranteed but smaller reward than was earned if the test was selected and answered correctly (test phase). The yellow shading indicates an illuminated nose-poke (NP) aperture, used to decline or accept the test. Data (Foote & Crystal, 2007): Performance from three rats (bottom panels) and the mean across rats (top-middle and top-right panels). Difficult tests were declined more frequently than easy tests; difficulty was defined by proximity of the stimulus duration to the subjective middle of the shortest and longest durations). The decline in accuracy as a function of stimulus difficulty was more pronounced when tests could not be declined (forced test) compared to tests that could have been declined (choice test). Error bars represent standard errors. (Adapted from Foote, A. L., & Crystal, J. D. (2007). Metacognition in the rat. *Current Biology*, 17, 551-555. © 2007 by Elsevier Ltd.)

1. Foote and Crystal (2007) presented a noise from a set of eight durations, which was to be categorized as short or long (i.e., reward was contingent on judging the four shortest and four longest durations as short and long, respectively). The procedure is outlined in Figure 2. The two experiments had the following common features. Before taking the test, the animals were given the opportunity to decline it. On other trials, the animals were not given the option to decline the test. Accurate performance on the test yielded a valuable reward, whereas inaccurate performance resulted in no reward. Declining a test yielded a less valuable but guaranteed reward. The decline rate increased as a function of difficulty (longer retention intervals for the monkey or proximity to the subjective middle of short and long durations for the rats) and accuracy was lowest on difficult tests that could not be declined. Note that the data in Figures 1 and 2 meet the prevailing standard: not only did the animals appear to use the decline response to avoid difficult problems, but the Chosen-Forced performance advantage emerged as a function of task difficulty.

Quantitative modeling

Recent quantitative modeling by Smith and colleagues (2008) shows that low-level (i.e., alternative) mechanisms can produce both apparently functional use of the decline response and the Chosen-Forced performance advantage; these alternatives are low level in the sense that they use primary representations without application of secondary representations. Consequently, the formal modeling suggests that the prevailing standard is inadequate to document metacognition.

Smith and colleagues used basic associative and habit formation principles in their quantitative model. They proposed that direct reward of the decline response produces a low-frequency tendency to select that response independent of the stimulus in the primary discrimination². We note that Smith et al proposed that the decline response has a constant attractiveness across the stimulus continuum; constant attractiveness means that the tendency to produce the response is constant across stimulus conditions. We refer to this class of threshold explanations as a stimulus-independent hypothesis to contrast it with the stimulus-response hypothesis outlined above (the explanatory power of a stimulus-independent hypothesis will be evaluated below). For the primary discrimination, Smith et al. used standard assumptions about exponential decay of a stimulus (i.e., generalization decrements for an anchor stimulus in a trained discrimination); exponential decay is also commonly used to model a fading memory trace (Anderson, 2001; Killeen, 2001; Sargisson & White, 2001, 2003, 2007; Shepard, 1961; Sikström, 1999; White, 2001, 2002; Wixted, 2004). Such exponential decay functions have extensive empirical and theoretical support

(Shepard, 1961, 1987; White, 2002). Thus, the primary discrimination and the decline option give rise to competing response-strength tendencies, and Smith et al. proposed a winner-take-all response rule (i.e., the behavioral response on a given trial is the one with the highest response strength). A schematic of the formal model appears in Figure 3a. Simulations with this quantitative model document that it can produce both aspects of the prevailing standard (Figure 3b): the decline response effectively avoids difficult problems, and the Chosen-Forced performance advantage emerges as a function of task difficulty. Note that both empirical aspects of putative metacognition data are produced by the simulation (Figure 3b) without the need to propose that the animal ‘knows when it does not know’ or any other metacognitive process.

Our goal is to assess the impact of applying Smith and colleagues’ (2008) model depicted in Figure 3a to a broad set of experiments on metacognition. In our view, it is important to note that the modeling generates predictions that are stimulus independent in contrast to the traditional stimulus-response hypothesis. According to a stimulus-response hypothesis, an animal is assumed to learn to do a particular response in the presence of a particular stimulus. For example, with a stimulus response mechanism, an animal can learn to do a particular response (such as a decline response) in particular stimulus conditions at a higher rate than in other stimulus conditions; such a stimulus-response hypothesis would take the form of an inverted U-shaped function in Figure 3a for the decline response (i.e., replacing the constant attractiveness proposed by Smith et al.’s threshold in Figure 3a). By contrast, according to a stimulus-independent hypothesis, previous reinforcement with a particular response is sufficient to produce that response in the future at a relatively low frequency. Note that the response has a constant attractiveness independent of stimulus context. Our application of the ideas described above suggests that many studies in metacognition are well equipped to test stimulus-response hypotheses, but are not adequate to test the stimulus-independent hypothesis (the details to support this conclusion appear below).

We emphasize at the outset that our view is not that stimulus-response learning is absent in these types of experiments; indeed, stimulus-response learning is likely at work, meaning that principles of generalization of training stimuli to new stimulus conditions are at work too, in addition to stimulus-independent factors. It is also worth emphasizing at the outset, that Smith et al. (2008) did not propose that animals would learn a stimulus-independent use of a decline response in all experiments. However, they did emphasize that a history of reinforcement is sufficient to establish a low-frequency threshold that is independent of stimulus conditions. Thus, our goal is to evaluate the implications of

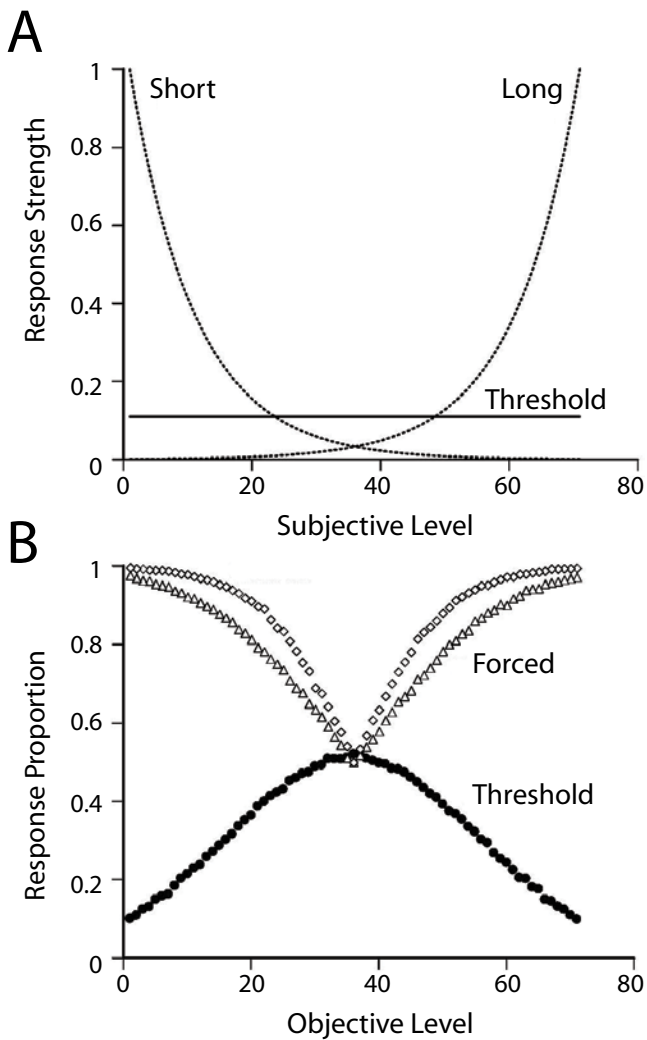


Figure 3. Schematic of low-level, response-strength model and simulation. (a) Presentation of a stimulus gives rise to a subjective level or impression of that stimulus. For any given subjective level, each response has a hypothetical response strength. The schematic outlines response strengths for two primary responses in a two-alternative forced-choice procedure and for a third (i.e., decline or uncertainty) response (labeled threshold). Note that response strength is constant for the third response (i.e., it is stimulus independent). By contrast, response strength is highest for the easiest problems (i.e., the extreme subjective levels). Note also that for the most difficult problems (i.e., middle subjective levels) the decline-response strength is higher than the other response strengths. Reproduced from Smith et al. (2008). (b) Simulation of schematic shown in (a). Simulation of a response-strength model with a flat threshold produces apparently functional use of the decline response (i.e., intermediate, difficult stimuli are declined more frequently than easier stimuli). The Choice-Forced performance advantage emerges as a function of stimulus difficulty. Reproduced from Smith et al. (2008). (From Smith, J. D., Beran, M. J., Couchman, J. J., & Coutinho, M. V. C. (2008). The comparative study of

the stimulus-independent hypothesis and to determine how much of the existing data can be explained by applying the stimulus-independent hypothesis to established methods of assessing metacognition in animals.

Assessment of what is 'lost' (methods and corresponding data)

Formal modeling of low-level alternative mechanisms that account for putative metacognition data necessitates an assessment of what is 'lost' in terms of existing methods and data. In our assessment, data from existing methods do not withstand the scrutiny of the formal modeling. We believe that recognizing that existing methods are inadequate represents progress. One potential benefit of such an assessment is that it may lead to the development of new methods to examine metacognition that are not subject to low-level explanations.

The methods and data shown in Figures 1 and 2 are subject to modeling by low-level explanations. Indeed, the schematic in Figure 3 was designed to specifically explain the rat data in Figure 2 (Foote & Crystal, 2007). We agree that the formal model of Smith et al. (2008) applies to our data from rats (throughout the remainder of this article we focus primarily on data from rhesus monkeys because these are the most extensive and thorough tests of comparative metacognition). However, in our view, the formal situation is essentially the same for the monkey data in Figure 1 (Hampton, 2001). We outline the formal situation in Figure 4. In the case of delayed-matching to sample, we may assume that the presentation of a sample stimulus gives rise to a fading memory trace after stimulus termination (Anderson, 2001; Killeen, 2001; Sargisson & White, 2001, 2003, 2007; Shepard, 1961; Sikström, 1999; White, 2001, 2002; Wixted, 2004). Thus, the horizontal axis in Figure 3 could be represented as trace decay, which grows as a function of retention interval (i.e., the independent variable in Hampton's study). A low-frequency threshold is used for the decline response. Smith and colleagues' simulations suggest that this type of model can produce both apparently functional use of the decline response to avoid difficult problems and a Chosen-Forced performance advantage that emerges as a function of task difficulty.

Metamemory

The case of memory is less perceptually grounded than temporal discrimination, and it may be argued that memory is more abstract or sophisticated. However, we believe that it is possible to apply the formal model by using a trace-decay continuum for a fading stimulus trace. Thus the model may be sufficient to explain the decline rate and the Chosen-

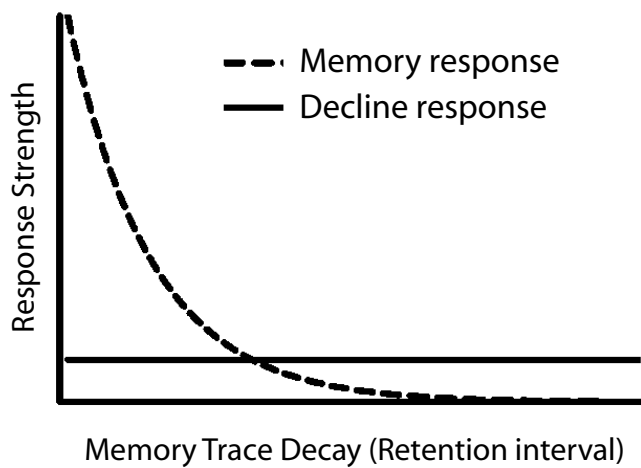


Figure 4. Schematic of low-level, response-strength model of metamemory. Presentation of a stimulus gives rise to a fading memory trace after stimulus termination. Trace decay (which is shown on the horizontal axis) grows as a function of retention interval. A low-frequency threshold is used for the decline response. Note that response strength is constant for the decline response. By contrast, memory response strength is highest for the shortest retention intervals. Note that for the most difficult problems (i.e., long retention intervals) the decline response strength is higher than the memory response strength. Also note that the horizontal axis may be viewed as a primary representation (see text for details).

Forced performance advantage. In particular, it is possible that the monkey's performance depicted in Figure 1 could be based on a primary representation of trace strength. According to this view, use of the decline response is based on a fading memory trace just as the old-new responses from the primary task are based on a fading memory trace. Because the same fading memory trace (i.e., the same primary representation) is used for both the primary memory task and the decline response, it is not clear that a secondary representation is needed to explain the data. The use of two different responses (decline and matching responses) does not, in itself, indicate that the two responses are based on different types of representations, as outlined next.

The interpretive problem here is how to determine if the monkey is responding on the basis of a primary representation (i.e., the strength of stimulus representation is very weak) or on the basis of a secondary representation (i.e., the monkey knows that it does not know the correct answer). It is not sufficient to claim that any paradigm that uses memory as the primary task will, by definition, result in secondary representations about memory (and thus, by definition, constitute evidence for metamemory). What data specifically implicates the use of a secondary representation? Before Smith and colleagues documented putative metacognitive

data patterns based on low-level mechanisms, the answer to this question was that the Chosen-Forced performance advantage could not be explained without appeal to metacognition. However, this pattern of data is not as informative as previously supposed. The burden of proof, in this situation, is on providing evidence that implicates a secondary representation, and until such evidence is provided the cautious interpretation is to claim that a primary representation is sufficient to explain the data. We also note that the observation that the memory trace is an internal representation is not adequate to answer the question posed above. Indeed, all representations are internal. If all that is needed is an internal representation (what we have been referring to as a primary representation), then what is to prevent the assertion that performance on matching to sample is based on metacognition (i.e., a secondary representation)? There are multiple responses in these types of experiment (i.e., the decline response and the primary response of choosing a correct/incorrect choice in matching to sample). Thus, it may be argued that the decline response is dedicated to reporting about a secondary representation, whereas the other responses are dedicated to reporting about the primary representation. But how do we know if this is the case? Clearly, what is needed is an independent line of evidence. In any case, Smith et al.'s (2008) model deals very nicely with competition between responses (i.e., the model successfully picks responses based on low-level mechanisms without application of a secondary representation).

Hampton's (2001) study had several elegant features, and several of these features were included in a series of recent elegant tests with pigeons (Sutton & Shettleworth, 2008). For example, after training with one retention interval, Hampton's monkeys received a novel probe in which no sample stimulus was presented. This is a direct manipulation of memory, and it is intuitive that an animal with metacognition would respond adaptively by declining the test (this is what the monkeys did). However, this could also be based on the primary representation. Indeed, if a sample is omitted on a probe trial, the trace strength from the most recently presented sample (i.e., the one presented on the trial that preceded the probe) would have had a very long time to decay. In such a situation, the trace strength from the primary representation would be very low, in which case it would likely be lower than the threshold for declining the test. Thus, a decline response would be expected based on the primary representation. Again, the interpretive problem here is how to know if the monkey is responding on the basis of a primary representation (i.e., the strength of stimulus representation is very weak) or on the basis of a secondary representation (i.e., the monkey knows that it does not know the correct answer – in this case because there is no correct answer).

Pervasiveness of reinforcement of uncertainty responses

Some early experiments on uncertainty used direct reinforcement variables to influence the behavior of the monkeys; for example, an uncertain response sometimes produced a hint or identification of the currently correct response (e.g., Smith, Shields, Allendoerfer, & Washburn, 1998), a guaranteed-win trial (e.g., Shields, Smith, & Washburn, 1997; Smith, Shields, Schull, & Washburn, 1997), a time-out delay for over-use of the uncertainty response (e.g., Shields et al., 1997; Smith et al., 1998; Smith et al., 1997), or food (Hampton, 2001).

Smith and colleagues (Smith et al., 2008) suggested that a history of reinforcement associated with the decline response is responsible for the deployment of low-level alternative explanations. This observation has led to some creative attempts to circumvent the role of reinforcement (Beran, Smith, Redford, & Washburn, 2006; Smith, Beran, Redford, & Washburn, 2006). However, the functional use of a decline or uncertainty response may be due to the existence of residual reinforcement variables (see below). Consequently, the low level threshold from Figure 3 may apply, meaning that the animal's uncertainty behavior could be explained by low-level mechanisms.

Pure uncertainty response. Smith and colleagues suggest that reinforcement of the uncertainty response is responsible for the deployment of low-level alternative explanations. Indeed, this observation represents a significant hurdle for any method that employs an uncertainty response. For example, Beran and colleagues (2006) sought to develop a "pure" uncertainty response that would not be contaminated by reinforcement.

Beran and colleagues (2006) trained monkeys in a numerosity discrimination. Between one and nine circles were presented on a computer screen. When the display had less or more than a designated center value, the monkeys were rewarded for using a joystick to move a cursor to an "L" (less) or "M" (more) on the computer screen, respectively. A wide range of center values was systematically explored, many configurations of dots were used across trials, and brightness was controlled. The uncertainty response was a "?" at the bottom-center of the screen. Moving the joystick to this position ended the trial and initiated the next trial. Importantly, the authors emphasize that this method represents a pure uncertainty response in the sense that the uncertainty response was not reinforced by food, information about the correct answer, or the presentation of an easy next trial. Thus, they conclude that this was the purest trial-decline response possible.

However, even this valuable attempt to curtail reinforcement may leave some residual reinforcement in place; thus,

Smith et al.'s low-level explanation may apply in this experiment (using the primary representation of numerosity). The two rhesus monkeys in this study had previous experience with an uncertainty response from an earlier study that did not deploy the purest trial-decline procedure (Shields, Smith, Guttmanova, & Washburn, 2005; Smith et al., 2006); we note that the problem is not that the monkeys had successfully used an uncertainty response in a previous task, but rather that they had received reinforcement in the past. Specifically, the monkeys had been rewarded in the past for moving the joystick down (which was the response in their study). Training the monkeys to use the joystick involved requiring the monkeys to learn to (1) approach a perch to view the video display, (2) reach through the cage mesh to manipulate the joystick below the monitor, (3) move the joystick so that the cursor on the screen contacted computer-generated stimuli; the joystick response was rewarded with food (Rumbaugh, Richardson, Washburn, Savage-Rumbaugh, & Hopkins, 1989; Shields et al., 2005; Washburn & Rumbaugh, 1992). A history of reinforcement associated with moving the joystick down would presumably be sufficient to generate a low-frequency tendency to select this response. Smith and colleagues' model would then apply (a winner-take-all decision between a low-frequency tendency to move the joystick down to the "?" vs. selection of the "L" or "M" responses). The model predicts that "L" or "M" would win when the number of stimuli is far from the center value, whereas the uncertain response would win near the center value; thus, the model predicts an increase in the uncertainty response for the difficult central numbers even if they are not explicitly reinforced for making these choices in the current experiment. Moreover, these monkeys are relatively task savvy, given that they have a long history of participating in laboratory tasks with joysticks and moving icons to target locations in addition to other laboratory tasks. The efficiency benefits of using task-savvy subjects means that these subjects generalize from earlier experiments to the current experiment. This generalization is not surprising given the similarity between earlier experiments and a current experiment (e.g., sitting at the experimental perch, observing the computer display, reaching an arm through the mesh cage, contacting and moving the joystick, receiving reinforcement for joystick movements, etc.); all of these factors promote the use of responses from within their experimental repertoire in new experiments, thereby allowing the experimenters to forgo the extensive training experience that would otherwise be required if new subjects were tested in each experiment.

Other valuable features of the task (re-training with new central values, many different configurations of the circles on the computer screen, etc.) do not mitigate against that application of the low-level explanation. In this regard, it

is worth noting that the ability to perform the numerosity discrimination is presumably based on a primary representation, and for easy numerical discriminations the response strengths for “L” or “M” would be higher than the hypothesized low-level threshold for responding down to the “?”.

Moreover, there may have been concurrent reinforcement because the uncertainty option reduced the delay to reinforcement in subsequent trials. Reducing delay to reward is a reinforcement variable (Carlson, 1970; Kaufman & Baron, 1968; Richardson & Baron, 2008), which could maintain the low-frequency flat threshold for the uncertainty response. To examine the role of delay to reinforcement in these types of experiments we conducted a simulation of reinforcement rate. For the simulation, we used the exact feedback described by Beran et al. (2006) for their purest trial-decline response. On the primary task, a correct response produced 1 food pellet, and an incorrect response did not produce any food pellets. Critically, in their procedure, an incorrect response produced a time out of 20 sec. An uncertainty response did not produce food and did not produce a time out. We used a flat uncertainty threshold, as proposed by Smith et al. (2008). In the simulations, we varied the response strength for the uncertainty response from 0 to 1 using many intervening values and held all other aspects of the simulation constant³. If delay to reinforcement is not a reward variable in these studies, then the amount of food per unit time will be constant as a function of the threshold values in the simulations. By contrast, if delay to reinforcement functions as a reward variable, then there will be some threshold parameter for the uncertainty response that maximizes food per unit time.

Figure 5 shows the results of the simulation. Note that there is a peak in food per unit time. Thus, it is possible that a subject in these types of experiments could adjust its threshold level to maximize food per unit time, and this adjustment of the “non-reinforced” uncertainty response is reinforced by reduced delay to reinforcement in the overall procedure. This simulation shows that despite the lack of direct reward for use of the uncertainty response, there are residual reinforcement variables at work in these types of experiments. Thus, the uncertainty response was indirectly reinforced by increased food rate; application of the Smith et al. (2008) model would predict use of the uncertainty response for the intermediate stimuli. Our simulation is consistent with the hypothesis that there are negative affective consequences of time outs, which has been verified through independent approaches (Richardson & Baron, 2008).

In summary, it is important to note that although delay to reinforcement could occur on a trial with any numerosity display, it is not necessary to assume uncertainty monitoring in order to produce apparently functional use of the uncer-




Figure 5. Results of a simulation of reinforcement density as a function of variation in threshold for the uncertainty response. The simulation used the generalization and constant-threshold concepts from Smith et al. (2008). Reinforcement and delays were based on Beran et al. (2006). Although no food was delivered upon selecting the uncertainty response, the simulation shows that the value of the threshold for selecting the uncertainty response influences the amount of food obtained per unit time in the primary discrimination. Thus, the uncertainty response was indirectly reinforced despite efforts to eliminate reinforcement.

tainty response. First, a previous history of reinforcement of joystick responses is sufficient to establish a low-frequency tendency to select the “?”. The “?” response loses to “L” or “M” because “L” or “M” are high based on training in the numerosity task, except for the most difficult trials in which the response strength of “L” and “M” are lower than that of “?” (which produces a preference for “?” at the most difficult trials). Second, concurrent reinforcement may maintain the tendency to select “?” at a low frequency, as suggested by the simulation described above. Third, training with the primary task (i.e., changing the central value across phases of training) is responsible for changing the response gradients associated with the “L” and “M” responses. Certainly, these gradients would be modeled as primary representations. Thus, the movement of the expected high point in the use of “?” is based on the change in the shape of the response gradients for “L” and “M” as judged against a relatively constant low-frequency threshold for “?”.

Any of the above sources of reinforcement may be sufficient to apply the low-level threshold from Figure 3, meaning that the monkey’s uncertainty behavior could be explained by reinforcement and application of the Smith et al. (2008) model. We also note that the complexity of the primary task does not play a role in the application of Smith et al.’s model. The main issue is the appropriateness of ap-

peeling to a complex proposal (i.e., metacognition). Thus, the purpose of testing a less complex proposal (i.e., Smith et al.'s response strength model) is to determine if the output of the model can account for the data. If the output of the model accounts for the data, then it is not appropriate to select the more complex proposal to explain the data (absent an independent line of evidence that cannot be explained by the low-level model). It is ill-advised to choose to not apply the low-level model because of claims that the primary task is sophisticated, especially if the low-level model can produce the observed pattern of data.

Trial-by-trial feedback. Another example of the difficulty encountered in overcoming the low-level explanations comes from a recent study by Smith and colleagues (Smith et al., 2006). In this uncertainty monitoring study, trial-by-trial feedback was delayed and uncoupled from the responses that earned the feedback. In particular, the monkeys were presented with a computer display that had a variable number of randomly placed pixels. Some displays were sparse and others were dense. The monkeys were required to use a joystick to move a cursor to an "S" (sparse) or "D" (dense) on the screen. A "?" appeared at the bottom-center, to be selected for the uncertainty response. As the trials progressed, the monkeys earned food rewards or 20-s penalties (i.e., a time-out with a buzzer sound) based on correct or incorrect responses, respectively. However, the earned rewards or penalties were not delivered at the end of each trial. Instead, these consequences were delayed until the completion of a block of four trials. To further uncouple consequences from the responses that earned them, the feedback was not presented in the order in which they were earned. Instead, when the block ended, all rewards were presented first, followed by all time outs. The proportion of "S" and "D" responses tracked the density of the stimuli and declined toward the central value. The use of the "?" response peaked near the central value for one of the monkeys.

This is a highly innovative method to uncouple feedback in the density discrimination from the specific stimuli that were present when the feedback was earned, which likely has many applications. Indeed, it is very impressive that monkeys learned the task contingencies under these circumstances. Moreover, there were many other admirable features of the design (e.g., the monkeys initially received blocks of one trial – meaning transparent feedback – but the critical data was subsequently collected using different density ranges). The central question for our purpose is what can be predicted from Smith et al.'s (2008) model.

The study employed monkeys with previous reinforcement of the joystick response, so their history of reinforcement could contribute to a low-threshold tendency to select that option. A history of reinforcement with moving the joy-

stick down may be sufficient to generate a low-frequency tendency to do this response in the future as discussed above. The rest of the work is done by the response strengths for the "S" and "D" options. Because the proportion of "S" and "D" responses tracked the density of the stimuli, we may conclude that the animals had lower response tendencies for "S" and "D" near the central value based on learning the density discrimination. We note that they had these response tendencies despite the lack of transparent feedback, but density-discrimination performance is presumably based on a primary representation. The remaining question is: Do we need to hypothesize a secondary representation to explain the use of the uncertainty response? Because response strengths for "S" and "D" are expected to decline toward the most difficult density discriminations, a flat low-frequency threshold for the "?" would selectively produce higher response tendencies for "?" at the difficult discriminations. By contrast, as the density discrimination becomes easier for sparse and dense problems, "S" and "D" response strengths would progressively begin to exceed the response strength for "?". In addition to the history of reinforcement of the joystick response described above, there also may have been a residual source of concurrent reinforcement that would maintain the tendency to select "?" at a relatively low frequency. The selection of the uncertainty response would reduce delay to reinforcement in the next block of trials as outlined in the simulation above. Consequently, reinforcements per unit time would be higher when the monkey selected the uncertain option (which would be primarily restricted to difficult discriminations based on a comparison of response strengths as outlined above) compared to the scenario of not using the uncertain option. It is worth noting that the above analysis does not require that we assume uncertainty monitoring of a secondary representation. All that is required is a comparison of response strength of "?" (which is relatively low and flat) with the response strength of "S" and "D", which decline for difficult discriminations. Previous reinforcement of the joystick is sufficient to produce a constant attractiveness of this option. Moreover, the level of the low threshold for the uncertainty response could be maintained by sensitivity to food per unit time as suggested by the simulation above. Therefore, the analysis from Figure 3 may apply even to this study.

It is worth noting that although this is an impressive procedure that made significant progress in making feedback on the primary task opaque, these features of the experiment do not eliminate response strengths for the primary task. The monkeys in Smith and colleagues' (2006) study responded with higher accuracy on easy problems near the end of the stimulus continuum compared to difficult problems near the middle of the stimulus continuum. Thus, we could trace out a psychophysical function for the sparse-dense continuum.

This function is consistent with high response tendencies to respond “sparse” for the least-dense stimuli and to respond “dense” for the most-dense stimuli. Thus, response strengths appear to be much like what would be observed if feedback was transparent. To claim otherwise amounts to claiming that the animals do not have a customary psychophysical function. The above discussion suggests that it is reasonable to assume that the animals have response strengths for the primary task that appear to be similar to the response strength functions used in Smith et al.’s (2008) model. Although it is quite impressive that the animals learn the sparse-dense discrimination despite delayed and re-ordered feedback, once such learning has been documented, we can infer the existence of response strength tendencies and apply Smith et al.’s (2008) model.

It is also worth noting that asymmetries in the primary task and in the use of the uncertainty response do not provide definitive evidence for metacognition. In particular, Smith et al.’s (2006) monkey had a leftward shift in its use of sparse, dense, and uncertainty responses. If there is no response bias, then a subject would be expected to have response tendencies that are symmetrically distributed across the sparse-dense continuum. For example, in a discrimination with 41 stimulus levels, the frequency of sparse and dense responses would be expected to cross over at the middle stimulus level (i.e., 21) assuming a linear scale. Instead, the cross-over point occurred at approximately stimulus 16. The highest response proportion also occurred for stimulus level 16, but the distribution of all uncertain responses appears to be shifted even further to the left. Smith et al. (2008) argue that the leftward shift in the uncertainty-response distribution provides strong evidence for uncertainty monitoring. We offer an alternative explanation. The monkey had a relatively strong bias to judge displays as dense; the monkey was virtually perfect on all dense displays. Accuracy variation across stimuli was mainly restricted to the sparse response, with virtually perfect performance at the eight sparsest stimulus levels. If a psychophysical function plotted the probability of a sparse response as a function of stimulus level, the function would be shifted toward the left (i.e., the point of subjective equality [$p(\text{sparse response}) = .5$] was below stimulus level 21). The standard way to model this type of data pattern in a primary sparse-dense task is to propose that the psychophysical function for density is biased toward the left (Blough, 1998, 2000). Similarly, the uncertainty response is biased to the left, although the magnitude of the bias is slightly large in the case of the uncertainty response. A bias parameter is needed for both distributions, and the only anomaly is that the bias is slightly larger in one of the cases. From another perspective, it is worth noting that a low-level model predicts the use of the uncertainty response for difficult stimuli in the experiment reported by Smith et al. (2006), although

without an extra parameter, it cannot account for the exact location of the peak. Although these small modifications to the generalization account increase its complexity, the use of bias parameters to explain a small difference in an individual subject’s data is not unprecedented. Moreover, a slightly more complex generalization model would remain less complex than the proposal that animals exhibit metacognition.

Summary. Despite creative attempts to curtail reinforcement, the functional use of an uncertainty response may be due to the existence of residual reward variables. Indeed, if an uncertainty response was never reinforced, it seems unlikely that it would be produced by the subject, and it seems more unlikely that it would be used functionally to express uncertainty or escape a difficult trial⁴. We note that this pessimistic assessment is not meant to restrict inquiry. Instead, it is our hope that recognizing the limits of existing methods may help foster the development of new methods to assess the use of a secondary representation.

Transfer Tests

Given the pervasiveness of reinforcement variables in the training of decline or uncertainty responses, despite careful attempts to curtail reinforcement, it is important to examine alternative techniques that do not try to eliminate reward. The transfer test is the major technique that is used to test a stimulus-response hypothesis. In a transfer test, stimuli from training are replaced with new stimuli in test. Although there is an intuitive appeal to the transfer test methodology, the formal situation in a transfer test is much the same as discussed above. Thus, the animal’s uncertainty behavior in a new stimulus context may be explained by low-level mechanisms (as outlined below).

Formal modeling of transfer test methodology. If uncertainty responding is conditioned on occurrence of a specific stimulus, then transferring to a novel stimulus context is sufficient to prevent application of the stimulus-response mechanism. Thus, it is intuitive that the functional use of the uncertainty response in a novel stimulus context would strengthen the claim that generalized uncertainty – rather than stimulus context – controls uncertainty responding (i.e., metacognition). Although a transfer test is a powerful technique to assess representations that govern performance in many domains (Cook & Wasserman, 2007; Heyes, 1993; Reid & Spetch, 1998; Wright & Katz, 2007), an analysis of low-level mechanisms suggests that this intuition does not apply in the case of metacognition (details to support this conclusion appear below).

Typically, experiments with monkeys use relatively task-savvy subjects with a long history of participating in laboratory tasks. Thus, these subjects bring to each new experiment

an extensive repertoire of experiment-related behaviors. Indeed, the ability to readily draw on this behavioral repertoire facilitates using such experienced subjects. For example, the monkeys have extensive experience with joysticks and moving a cursor to target locations. Consequently, these subjects may generalize from earlier experiments to the current experiment. As described above, this generalization is not surprising given the similarity between earlier experiments and a current experiment (e.g., the monkeys approached the perch to view the video display, reached through the cage mesh to manipulate the joystick below the monitor, moved the joystick so that the cursor on the screen contacted computer-generated stimuli, etc.); all of these factors promote the use of responses from their experimental repertoire in new experiments, thereby allowing the experimenters to forgo the extensive training that would otherwise be required if new subjects were tested in each experiment.

In our analysis of transfer tests, uncertainty responding is modeled by a flat, low-frequency threshold in a transfer test because the uncertainty response is stimulus independent. In each case, response strengths determine selection of the uncertainty or transfer response. It is worth noting that a stimulus independent mechanism can operate in parallel with a stimulus-response mechanism. Consider the situation in which an animal learned a stimulus-response rule for a set of specifically trained stimuli. In a transfer test, the previously trained stimuli are not presented, and so the stimulus-response function is not available to guide the use of a decline response. However, if the animal also has a low-frequency tendency to select the response in a stimulus independent fashion, then what remains in the transfer condition is the low-frequency threshold. In such a situation, a generalization decrement would be expected for the stimulus-response mechanism but not for the stimulus-independent mechanism. In any case, a low-frequency tendency to select a decline response may exist once it has a history of reinforcement. Our goal is to evaluate the ability of this simple mechanism when the animal is subjected to a transfer test. We consider three types of transfer tests.

Case 1: Training on the primary task withholds a subset of stimulus conditions to be used for a future transfer test. Figure 6 shows the formal situation using a memory task at multiple retention intervals. The animal has been trained on the primary task using a variety of stimulus conditions (i.e., retention intervals), but a subset of stimulus conditions have been withheld (i.e., these stimuli have been reserved for use in a future test). When the retention interval is unusually short, trace strength is unusually high (and higher than the uncertainty response). For longer retention intervals, it is increasingly likely that the uncertainty option will be selected. As discussed above (see section on metamemory), the memory trace decay continuum is a primary representation

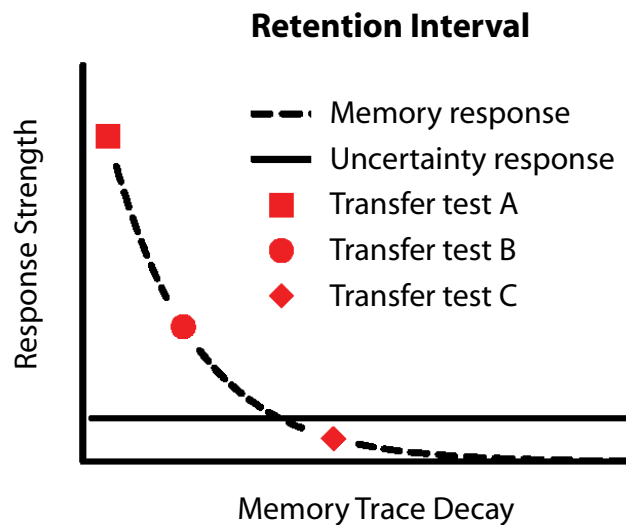


Figure 6. Schematic of response-strength model for a memory task using transfer to novel retention intervals. The primary task is well trained, and use of the uncertainty response has occurred along the trained stimulus continuum. Square, circle, and diamond represent response strengths for novel stimuli (i.e., novel retention intervals).

absent specific data that would pinpoint the use of a secondary representation. Note that a graded level of transfer is expected in Case 1 (with more responding to transfer test A than to transfer test B, which is higher than transfer test C).

Case 2: Training has occurred on two primary tasks, but the uncertainty response has not been presented in the transfer task prior to collection of the critical transfer data. Figure 7 shows the formal situation using two two-alternative forced choice tasks. Training on one discrimination (i.e., stimulus dimension 1, responses 1 and 2 in the figure) included the uncertainty response, whereas training on the other discrimination (i.e., stimulus dimension 2, responses 3 and 4) did not include the uncertainty option; the presentation of the uncertainty option together with stimulus dimension 2 has been reserved for use in a future transfer test. When the transfer task is difficult, the uncertainty response has the highest response strength. When the uncertainty option is presented with the transfer task for the first time, the response strength gradients are available for the transfer task based on previous training. Note that Responses 3 and 4 are expected to occur at the extreme ends of stimulus dimension 2 in the transfer test in Case 2, whereas the uncertainty response is expected to occur for intermediate stimuli based on the primary representations.

Case 3: Training has occurred with a primary task and the uncertainty response, and transfer assesses an untrained task. Thus, the uncertainty response and the primary task, but not the transfer task, are well trained. Figure 8 shows the

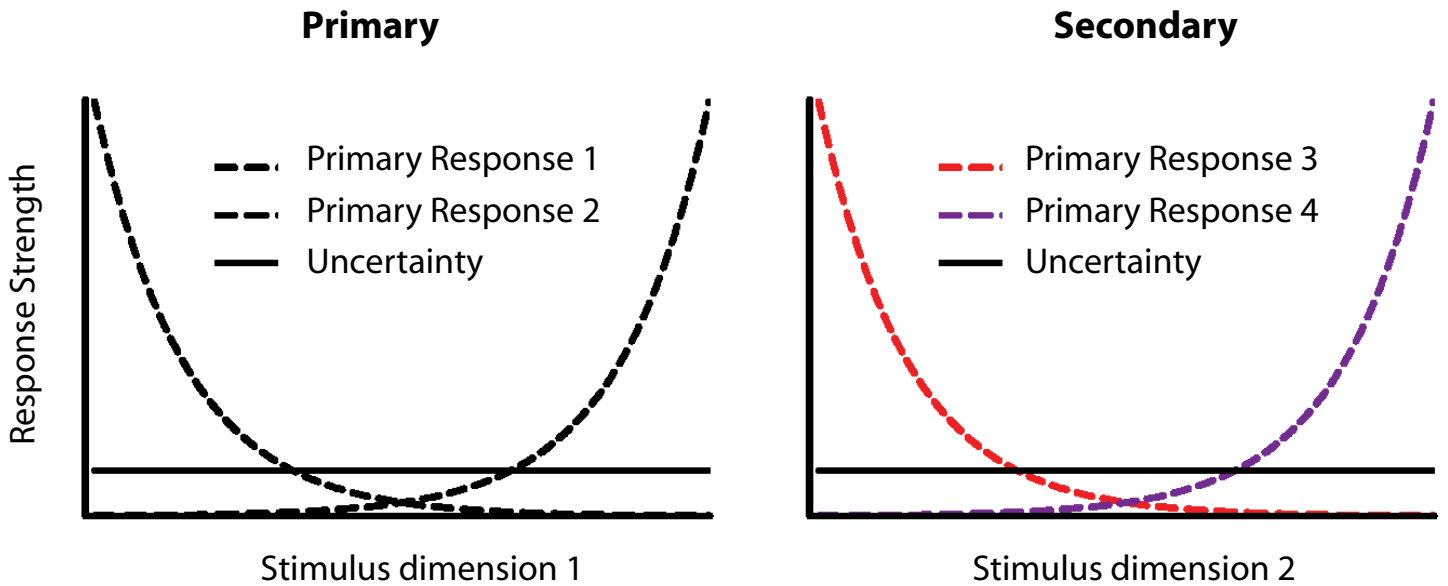


Figure 7. Schematic of response-strength model for two well-trained two-alternative forced-choice procedures. When the uncertainty option is presented with the transfer task for the first time, the response strength gradients are available for the transfer task (based on previous training).

formal situation. Unlike Case 2 above, training has occurred only on stimulus dimension 1. Thus, response-strength gradients are absent in the transfer test because the transfer task is not yet trained. This is the most difficult transfer case to model because assumptions need to be made about the response strengths in the untrained primary task. We assume that response strengths would be low for the untrained, primary-task response options, in which case the uncertainty option likely has the highest response strength, on average; the same situation occurs when a to-be-remembered item is

omitted as a transfer test. Note that in Case 3, the uncertainty response is expected to win compared to the untrained task (transfer stimuli A, B, and C).

The analysis of the three cases above focused on evaluating the implications of a flat low-frequency threshold for the uncertainty response. However, the above analysis also applies to the case of a stimulus-response hypothesis in the following way. Suppose that an animal learns a specific stimulus-response rule. To be concrete, let's consider Case 2 above, in




Figure 8. Schematic of response-strength model for one well-trained two-alternative forced-choice procedure. Training has occurred with a primary task and the uncertainty response (together), and transfer assesses an untrained task. Response-strength gradients are absent in the transfer test because the transfer task is not yet trained. Square, circle, and diamond represent response strengths for novel stimuli.

which the decline response is trained on stimulus dimension 1 and examined in a transfer test using stimulus dimension 2. If the animals have learned a stimulus-response rule in Case 2, then an inverted U-shaped function would appear in the left panel of Figure 7 (in place of the flat response threshold in the left panel of Figure 7). A transfer test is well suited to test this stimulus-response hypothesis because the transfer test uses stimuli from stimulus dimension 2, in which case none of the stimuli in the original stimulus-response rule are available in the transfer test. This is the conventional rationale for a transfer test. However, we need to evaluate the implications of an ineffective stimulus-response rule in the transfer condition. What is the attractiveness of the decline response in the transfer test (i.e., in stimulus dimension 2)? Because a stimulus-response rule has not been learned for any of the stimuli in dimension 2, the stimuli in dimension 2 would support only a very low-level of attractiveness, and this attractiveness would be constant across stimuli in dimension 2. Thus, the next step in our analysis is to evaluate the predictions of a low-frequency threshold for the uncertainty response. Of course, this is what is displayed in the right panel of Figure 7, and Smith et al.'s (2008) model can predict apparently functional use of the uncertainty response and a performance difference based on a winner-take all rule applied to response strengths. Thus, we conclude that the status of stimulus-response learning in the original task is not critical to our analysis of the transfer test.

In each case, putative transfer of metacognition can be based on low-level, response-strength mechanisms. Although transfer tests are a powerful technique to evaluate stimulus-response hypotheses, they are of limited utility here because the low-level mechanism may be stimulus independent (i.e., a low-frequency, flat threshold).

Parsimony and metacognition

We have argued that explanations based on primary representations should be tentatively accepted before asserting explanations based on secondary representations. In this respect, Smith et al.'s (2008) response-strength model is less complex than a metacognition model. In our review, we found that the existing data on uncertainty responses can be explained by a low-level model without appealing to metacognition. It is important to note that the situation would change dramatically if other data were to emerge that could not be explained by low-level models but could be explained by metacognition. In that situation, we would agree that animals are capable of metacognition. Moreover, if new data emerged that required a metacognitive explanation, then we would encourage a re-evaluation of older data. If animals were shown to be metacognitive in some tasks, then it would not be simpler to invoke alternative explanations for the re-

maining cases. For example, consider a case in which putative metacognition tasks 1-5 were adequately explained by low-level mechanisms but new tasks 6-7 emerged that could only be explained by metacognition. In this case, we would support the reinterpretation of all of the data as metacognition given the converging lines of evidence. Thus, there is a great need to explore new methods that can be explained by metacognition but cannot be explained by a low-level hypothesis.

Conclusions

We believe that existing methods have approached the difficult problem of metacognition in innovative and creative ways. A large array of techniques are being used (examination of accuracy predications, attempts to curtail reinforcement of uncertainty responses, uncoupling of feedback from the responses that earned the feedback). Although our assessment is that existing methods do not pinpoint the use of a secondary representation, we are optimistic that new methods can be developed. The objective of developing new methods would be to make predictions that cannot be explained based on primary representations alone.

Smith and colleagues' (2008) use of a flat threshold in their models has far-reaching implications. It suggests that stimulus-response learning about an uncertainty response (i.e., a curved or inverted U-shaped function) is not required for apparently functional use of the uncertainty response. Moreover, an independent line of evidence from accuracy data (i.e., a Chosen-Forced performance advantage) is not available from established methods. In addition, it is desirable to simulate low-level mechanisms before conducting new experiments. The advantage of such an approach is that it requires precise specifications of the method and the model. Such an approach may be helpful in identifying new, innovative methods that can pinpoint the use of a secondary representation (i.e., a method in which some pattern of data is not predicted by application of a primary representation alone).

We believe that an assessment of what is 'lost' in terms of existing metacognition methods and data may prompt the development of new methods to examine metacognition that are not susceptible to low-level explanations. Such an assessment is critical for the development of new standards by which to evaluate tests of metacognition. A periodic re-evaluation of standards will facilitate progress in our understanding of metacognition in animals.

References

- Anderson, R. B. (2001). The power law as an emergent property. *Memory and Cognition*, **29**, 1061-1068.
- Beran, M. J., Smith, J. D., Redford, J. S., & Washburn, D.

- A. (2006). Rhesus macaques (*Macaca mulatta*) monitor uncertainty during numerosity judgments. *Journal of Experimental Psychology: Animal Behavior Processes*, **32**, 111-119.
- Blough, D. S. (1998). Context reinforcement degrades discriminative control: A memory approach. *Journal of Experimental Psychology: Animal Behavior Processes*, **24**, 185-199.
- Blough, D. S. (2000). Effects of priming, discriminability, and reinforcement on reaction-time components of pigeon visual search. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**, 50-63.
- Carlson, J. G. (1970). Delay of primary reinforcement in effects of two forms of response-contingent time-out. *Journal of Comparative and Physiological Psychology*, **70**, 148-153.
- Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind & Language*, **23**, 58-89.
- Cook, R. G., & Wasserman, E. A. (2007). Learning and transfer of relational matching-to-sample by pigeons. *Psychonomic Bulletin & Review*, **14**, 1107-1114.
- De Lorge, J. O., & Clark, F. C. (1971). Observing behavior in squirrel monkeys under a multiple schedule of reinforcement availability. *Journal of the Experimental Analysis of Behavior*, **16**, 167-175.
- Descartes, R. (1637). *Discourse on method*.
- Emery, N. J., & Clayton, N. S. (2001). Effects of experience and social context on prospective caching strategies by scrub jays. *Nature*, **414**, 443-446.
- Footo, A. L., & Crystal, J. D. (2007). Metacognition in the rat. *Current Biology*, **17**, 551-555.
- Hampton, R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences of the United States of America*, **98**, 5359-5362.
- Heyes, C. M. (1993). Anecdotes, training, trapping and triangulating: Do animals attribute mental states? *Animal Behaviour*, **46**, 177-188.
- Hoerold, D., Dockree, P. M., O'Keefe, F. M., Bates, H., Pertl, M., & Robertson, I. H. (2008). Neuropsychology of self-awareness in young adults. *Experimental Brain Research*, **186**, 509-515.
- Inman, A., & Shettleworth, S. J. (1999). Detecting metamemory in nonverbal subjects: A test with pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **25**, 389-395.
- Jozefowicz, J., Staddon, J. E. R., & Cerutti, D., T. (submitted). *The behavioral economics of choice and interval timing*.
- Kaufman, A., & Baron, A. (1968). Suppression of behavior by timeout punishment when suppression results in loss of positive reinforcement. *Journal of the Experimental Analysis of Behavior*, **11**, 595-607.
- Killeen, P. R. (2001). Writing and overwriting short-term memory. *Psychonomic Bulletin & Review*, **8**, 18-43.
- Metcalfe, J., & Kober, H. (2005). Self-reflective consciousness and the projectable self. In H. Terrace & J. Metcalfe (Eds.), *The Missing Link in Cognition: Origins of Self-Reflective Consciousness* (pp. 57-83). New York: Oxford University Press.
- Morgan, C. L. (1906). *An introduction to comparative psychology*. London: W. Scott.
- Nelson, T. O. (1996). Consciousness and metacognition. *American Psychologist*, **51**, 102-116.
- Reid, S. L., & Spetch, M. L. (1998). Perception of pictorial depth cues by pigeons. *Psychonomic Bulletin & Review*, **5**, 698-704.
- Richardson, J. V., & Baron, A. (2008). Avoidance of timeout from response-independent food: Effects of delivery rate and quality. *Journal of the Experimental Analysis of Behavior*, **89**, 169-181.
- Robinson, A. E., Hertzog, C., & Dunlosky, J. (2006). Aging, encoding fluency, and metacognitive monitoring. *Aging, Neuropsychology, and Cognition*, **13**, 458-478.
- Rumbaugh, D. M., Richardson, W. K., Washburn, D. A., Savage-Rumbaugh, E. S., & Hopkins, W. D. (1989). Rhesus monkeys (*Macaca mulatta*), video tasks, and implications for stimulus-response spatial contiguity. *Journal of Comparative Psychology*, **103**, 32-38.
- Sargisson, R. J., & White, K. G. (2001). Generalization of delayed matching-to-sample following training at different delays. *Journal of the Experimental Analysis of Behavior*, **75**, 1-14.
- Sargisson, R. J., & White, K. G. (2003). The effect of reinforcer delays on the form of the forgetting function. *Journal of the Experimental Analysis of Behavior*, **80**, 77-94.
- Sargisson, R. J., & White, K. G. (2007). Remembering as discrimination in delayed matching to sample: Discriminability and bias. *Learning & Behavior*, **35**, 177-183.
- Shahan, T. A., & Podlesnik, C. A. (2008). Quantitative analyses of observing and attending. *Behavioural Processes*, **78**, 145-157.
- Shepard, R. N. (1961). Application of a trace model to the retention of information in a recognition task. *Psychometrika*, **26**, 185-203.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, **237**, 1317-1323.
- Shields, W. E., Smith, J. D., Guttmanova, K., & Washburn, D. A. (2005). Confidence judgments by humans and rhesus monkeys. *Journal of General Psychology*, **132**, 165-186.
- 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.

- Shimamura, A. P., & Metcalfe, J. (1994). The neuropsychology of metacognition. In *Metacognition: Knowing about knowing*. (pp. 253-276). Cambridge, MA, US: The MIT Press.
- Sikström, S. (1999). Power function forgetting curves as an emergent property of biologically plausible neural network models. *International Journal of Psychology*, **34**, 460-464.
- Smith, J. D., Beran, M. J., Couchman, J. J., & Coutinho, M. V. C. (2008). The comparative study of metacognition: Sharper paradigms, safer inferences. *Psychonomic Bulletin & Review*, **15**, 679-691.
- Smith, J. D., Beran, M. J., Redford, J. S., & Washburn, D. A. (2006). Dissociating uncertainty responses and reinforcement signals in the comparative study of uncertainty monitoring. *Journal of Experimental Psychology: General*, **135**, 282-297.
- 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. (2003). The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences*, **26**, 317-373.
- Staddon, J. E. R., Jozefowicz, J., & Cerutti, D. (2007). Metacognition: A problem not a process. *PsyCrit*, April **13**, 2007.
- Steiner, J. (1970). Observing responses and uncertainty reduction: II. The effect of varying the probability of reinforcement. *The Quarterly Journal of Experimental Psychology*, **22**, 592-599.
- Sutton, J. E., & Shettleworth, S. J. (2008). Memory without awareness: Pigeons do not show metamemory in delayed matching to sample. *Journal of Experimental Psychology: Animal Behavior Processes*, **34**, 266-282.
- Washburn, D. A., & Rumbaugh, D. M. (1992). Testing primates with joystick-based automated apparatus: Lessons from the Language Research Center's Computerized Test System. *Behavior Research Methods, Instruments & Computers*, **24**, 157-164.
- White, K. G. (2001). Forgetting functions. *Animal Learning & Behavior*, **29**, 193-207.
- White, K. G. (2002). Psychophysics of remembering: The discrimination hypothesis. *Current Directions in Psychological Science*, **11**, 141-145.
- Wixted, J. T. (2004). On common ground: Jost's (1897) law of forgetting and Ribot's (1881) law of retrograde amnesia. *Psychological Review*, **111**, 864-879.
- Wright, A. A., & Katz, J. S. (2007). Generalization hypothesis of abstract-concept learning: Learning strategies and related issues in *Macaca mulatta*, *Cebus apella*, and *Columba livia*. *Journal of Comparative Psychology*, **121**, 387-397.
- Zentall, T. R., Clement, T. S., & Kaiser, D. H. (1998). Delayed matching in pigeons: can apparent memory loss be attributed to the delay of reinforcement of sample-orienting behavior? *Behavioural Processes*, **43**, 1-10.
- Zentall, T. R., Hogan, D. E., Howard, M. M., & Moore, B. S. (1978). Delayed matching in the pigeon: Effect on performance of sample-specific observing responses and differential delay behavior. *Learning and Motivation*, **9**, 202-218.

Footnotes

- ¹ We note that Carruthers's first-order explanation proposes that a "gate-keeping" mechanism is used to select from differing goals, each with a different degree of strength, thereby determining a behavioral outcome. In our analysis that follows, it is not necessary to propose a gate-keeping mechanism. Rather a winner-take all response rule is used to select a response with the highest response strength.
- ² Smith et al. (2008) described two alternative proposals and Staddon and colleagues (Jozefowicz, Staddon, & Cerutti, submitted; Staddon, Jozefowicz, & Cerutti, 2007) have described an additional alternative. Each proposal has a similar function form for the decision-making process. Thus, we examine in detail one of Smith et al.'s proposals here.
- ³ We used 0.7 for the "sens" parameter in Smith et al. (2008), and for variability in mapping physical stimuli into subjective representations, we used a normal distribution with a mean of 0 and a standard deviation of 1.5. The delays to reinforcement were based on information from Beran et al. 2006 (20-s time out for incorrect responses, 1-s inter-trial interval), and we estimated the amount of time for viewing the stimuli and producing the choice response at 2 s. The center value for the primary discrimination was 5.
- ⁴ A subject may engage in observing responses to test the functional role played by different response options. Although observing responses may have a basis in reinforcement (De Lorge & Clark, 1971; Shahan & Podlesnik, 2008; Steiner, 1970; Zentall, Clement, & Kaiser, 1998; Zentall, Hogan, Howard, & Moore, 1978), even if they did not have a basis in reinforcement, the Smith et al. (2008) low-level model could be applied to this situation because observing responses may occur in a stimulus independent fashion.