“PLAY IT AGAIN”: A NEW METHOD FOR TESTING METACOGNITION IN ANIMALS

by

ALLISON L. FOOTE

(Under the Direction of Jonathon D. Crystal)

ABSTRACT

Putative metacognition data in animals may be explained by non-metacognition models (e.g., stimulus generalization; Smith et al., 2008). The objective of the present study was to develop a new method for testing metacognition in animals that cannot be explained by non-metacognition models. Animals are sometimes in a high or low state of performance. Metacognition is the hypothesis that animals know that they are in a high or low state. On a difficult problem we assume that animals choose to repeat a stimulus if they are in a low state of performance. Rats were first presented with a brief noise duration which they would subsequently classify as short or long. Rats were sometimes forced to take an immediate duration test, forced to repeat the same duration, or had the choice to take the test or repeat the duration. Metacognition, but not alternative non-metacognition models, predicts that accuracy on difficult durations is higher when subjects are forced to repeat the stimulus compared to trials in which the subject chose to repeat the stimulus, a pattern observed in our data. Simulation of a non-metacognition model supports the conclusion that our data document metacognition in rats.

INDEX WORDS: Metacognition, Metacognitive Control, Methods, Simulations, Rats, Animals

“PLAY IT AGAIN”: A NEW METHOD FOR TESTING METACOGNITION IN ANIMALS

by

ALLISON L. FOOTE

B.A., Auburn University, 2003

M.S., The University of Georgia, 2007

A Dissertation Submitted to the Graduate Faculty of The University of Georgia in Partial Fulfillment of the Requirements for the Degree

DOCTOR OF PHILOSOPHY

ATHENS, GEORGIA

2010

“PLAY IT AGAIN”: A NEW METHOD FOR TESTING METACOGNITION IN ANIMALS

by

ALLISON L. FOOTE

Major Professor:	Jonathon D. Crystal
Committee:	Dorothy M. Fragaszy Adam Goodie Philip V. Holmes Richard Marsh

Electronic Version Approved:

Maureen Grasso
Dean of the Graduate School
The University of Georgia
May 2010

ACKNOWLEDGEMENTS

I would like to thank my advisor Dr. Jonathon D. Crystal for his guidance, support, expertise, and most of all his patience in the completion of my doctoral degree. I would also like to thank my committee members Drs. Dorothy M. Fragaszy, Adam Goodie, Philip V. Holmes, Richard Marsh for their advice and suggestions regarding my dissertation. Special thanks to the following individuals for listening to my dissertation woes: Jessi Crast, Allison Eury, Freya Liu, Desiree Sharpe, Brian Stone, George Wilson, and Wenyi Zhou. I also want to thank Tony Snodgrass for his help with writing the programs needed for my simulations. Most importantly, I would like to thank my rats (KK1-KK8) that worked their little tails off to provide the important data discussed in my dissertation. The work for my dissertation was supported by National Institute of Mental Health Grants R01MH64799 and R01MH080052.

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

Page

CHAPTER 1

INTRODUCTION

Metacognition has been defined as the ability to reflect upon one's own internal cognitive states (Metcalfe & Kober, 2005). Although humans regularly engage in metacognition in daily life, it may not be immediately apparent how humans use metacognition. Imagine that you and a friend are asking each other trivia questions and your friend asks you who the last person to win 14 major championships in men's tennis was before the 2009 French Open. There may be many answers to this question or just one specific answer. One strategy would be to list each tennis player that enters your mind to your friend and, assuming your friend knows the answer, see if you are correct. Or, you could use metacognition. Metacognition would allow you to quickly and efficiently determine your knowledge of how many men have won 14 major championships in men's tennis and this would narrow down your list of potential players to just one, Pete Sampras. In the above scenario, metacognition has allowed you to quickly and efficiently evaluate your memory instead of being subjected to the more clumsy process of trial-and-error learning (Hampton, 2001).

Metacognition has been a topic of interest for at least the past 3 decades in the study of human cognition (see Nelson & Narens, 1980; Nelson, 1984). Nelson and Narens (1990) proposed the idea that a cognitive executive regulates the flow of information between an object-level and a meta-level by using control and monitoring processes (see Figure 1.1 below). The object-level is a reservoir for an individual's cognitions, behaviors, memories, and descriptors of a current situation and the meta-level monitors and controls the object-level (Son & Kornell, 2005). Control and monitoring processes are two mechanisms used by the central executive to allow for communication between meta- and object- levels. Generally, monitoring processes consist of such phenomena as confidence judgments, feeling-of-knowing judgments, ease-of-learning judgments, and judgments-of-learning (Smith, 2005). On the other hand, control processes are composed of phenomena that determine the selection and kind of processing, selection of a search strategy, termination of study and search, and the allotment of time for study (Smith, 2005).

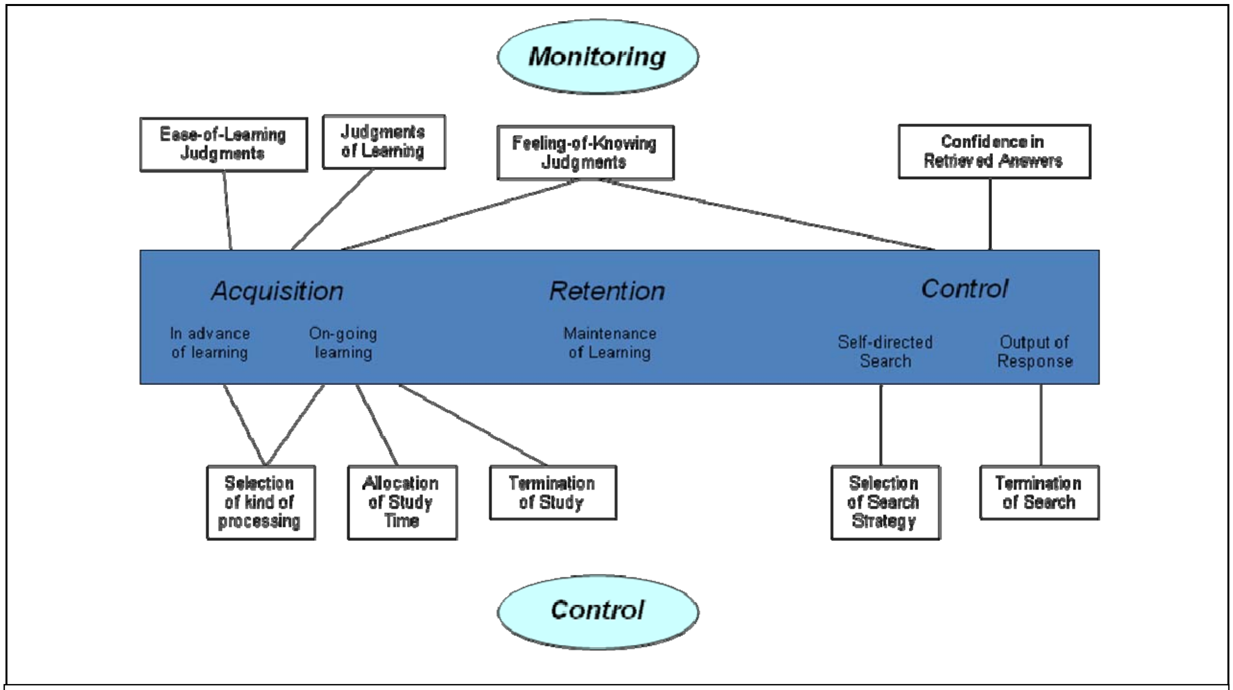

Figure 1.1. Nelson and Narens' (1990) model of metamemory in humans. Only the meta-level is depicted in the figure. The cognitive executive allows for communication between meta-level and the object-level via monitoring and control processes such as feeling-of-knowing judgments and the selection of search strategies, respectively.

CHAPTER 2

METACOGNITION IN ANIMALS

Perhaps one of the most fascinating reasons for studying metacognition in animals is discovering how metacognition evolved. Metacognition, like other cognitive, physiological, and morphological characteristics, was likely shaped by the unique ecological challenges in an animal's environment. Such challenges would include the pressures of uncertainty about finding food, determining the location of predators, finding shelter safe from weather and predators, and locating and securing a mate (Kornell, 2009). Additionally, individual and species survival would be dependent upon an animal's ability to control cognitive decisions. Cognitive control would be important for dividing waking hours into proportionate units of time to spend foraging, finding shelter, and selecting a mate, as well as determining priority for daily activities. In order to survive and successfully pass on these adaptive traits, animals would certainly benefit from cognitive abilities such as metacognition or other alternative strategies.

Another important reason for studying metacognition in animals is the development of animal models. Neurobiological research often involves the use of animal models. Furthermore, most neurobiological research is conducted using a rodent model because the neuroanatomy is well understood. Thus, a rodent model of a cognitive process would be beneficial for future studies on the underlying anatomy and physiology involved in metacognitive processes. In addition, the most important reason for having a reliable and valid animal model for metacognition is the impact it can have on research involving pathological brain diseases in humans. Specifically, an animal model would allow for an in-depth exploration and understanding of what changes occur in the brains of humans who are afflicted with cognitive diseases or impairments (e.g. understanding the failure of the distinctiveness heuristic in Alzheimer's disease, understanding how changes in brain morphology affect rumination in patients with depression; Budson, Dodson, Daffner, and Schacter, 2005; Roelofs, Papageorgiou, Gerber, Huibers, Peeters et al., 2007). In addition, an animal model of metacognition would be greatly beneficial for the development and refinement of treatments that could mitigate the physical and emotional challenges faced by cognitively impaired individuals and their families (e.g., alleviating the effects of nicotine withdrawal on memory and metacognition; Kelemen and Fulton, 2008). Moreover, findings of metacognition in animals would provide important insight into the evolution of the mind (i.e., by discovering which species have metacognition), but need not necessarily imply that animals have some form of consciousness (Clayton and Dickinson, 1998; Clayton and Dickinson, 1999a; Clayton and Dickinson 1999b; Clayton, Yu, and Dickinson, 2001; Emery and Clayton, 2001; Hampton, 2001, Kornell, 2009).

Hampton's (2001) study on metacognition in rhesus monkeys is one example of a successful demonstration of metacognition in animals. Hampton used delayed-matching-to-sample where images (i.e., icons) were presented on a touch screen. Hampton predicted that monkeys that have forgotten an answer to a memory test and who are able to discriminate between the presence and absence of a memory should be more accurate on tests in which they have the option to decline it. Additionally, Hampton predicted that monkeys should decline memory tests more often when memory is experimentally manipulated (e.g., giving monkeys probe trials which do not include a stimulus and manipulating the delay interval after the presentation of a sample stimulus).

Hampton's method included four different phases: a study phase, a delay interval, a choice phase, and a test/small reward phase (see Figure 2.1). In the first phase, the study phase, a monkey was presented with a sample image that was followed by the second phase containing the delay period. Afterwards, during the third phase, the choice phase, the monkey had the option to either take or decline a memory test. If the monkey opted to take a memory test, it selected an icon that let it proceed directly to the test phase where the matching stimulus and three distracter images were presented. If the monkey chose the matching stimulus it received a highly desirable peanut reward. On the other hand, if the monkey had the option to decline a memory test, it selected the respective icon and then directly proceeded to the test. Hampton also conducted a second experiment which included no-sample test trials to determine if monkeys were choosing to decline memory tests because their memories for the stimulus were weak (or in this case non-existent). Hampton found that monkeys chose to decline the trials that contained no sample more frequently than normal trials and that the monkeys did so from the first session. Lastly, in a third experiment, Hampton manipulated the retention interval and found that the monkeys' accuracy was higher on trials with short retention intervals compared to trials with a long retention interval.

Whereas some studies of animal metacognition have used dolphins and pigeons, the majority of studies have used non-human primates as subjects. Foote and Crystal (2007) expanded the search for metacognition in other species by adapting Hampton's (2001) method for use with rats. Instead of using a visual, delayed-matching-to-sample task, Foote and Crystal used an auditory time discrimination task. Furthermore, the authors made similar predictions to Hampton's about behavior and performance. First, if rats had knowledge of their own cognitive states then they would be expected to decline difficult tests more often than easy tests. Second, accuracy should be worse for difficult tests in which rats did not have the option to decline taking the test.

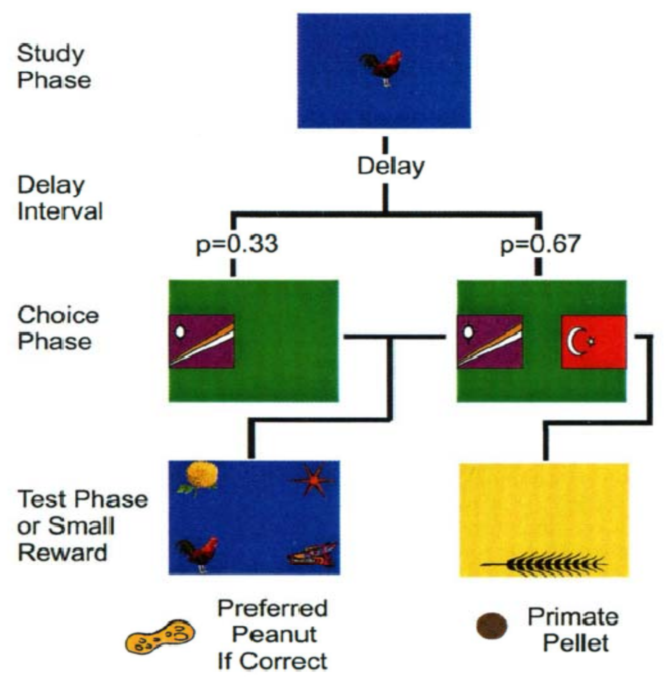

Figure 2.1 Hampton's (2001) method for testing metacognition in rhesus monkeys.

Foote and Crystal's (2007) method was similar to Hampton's method in that it contained a study phase, a choice phase, and a test/small reward phase (see Figure 2.2). Eight logarithmically spaced, white-noise durations ranging from 2-8 seconds were used for the auditory stimuli. Rats were required to classify stimulus durations as either short or long, depending upon the length of each stimulus, in order to receive a food reward. During the study phase, rats were presented with a white-noise stimulus duration that was followed by an approximately 8-minute intertrial interval. In the choice phase, the rat had the option to either take or decline a memory test. If the rat opted to take a memory test, it selected the take-test response by using the respective nose-poke aperture which proceeded directly to the test phase. During the test phase the rat could classify the stimulus duration as short or long by pressing one lever for short or the other lever for long. If the rat was correct it received a large reward of six food pellets; no pellets were delivered if an incorrect duration classification occurred. On the other hand, if the rat selected the option to decline a memory test, by using the other nose-poke aperture, it proceeded directly to the small reward phase. Instead of taking a memory test during the test phase the rat was given a less desirable reward of one food pellet. However, if the rat was forced to take a memory test, its only option was to select the nose-poke aperture associated with taking the memory test. Foote and Crystal found that rats declined difficult tests more often than easy tests and that accuracy was worse on difficult tests when rats did not have the option to decline.

Challenges of Studying Metacognition in Animals

Recently, Smith, Beran, Couchman, and Coutinho (2008) constructed a model (referred to hereafter as the response-strength model) that has revealed three previously undetected methodological problems with existing work on the metacognition. One problem with existing research is reinforcement of the uncertainty response (a methodological norm in metacognition experiments). Smith and colleagues have argued that reinforcement of the uncertainty response could create an independent response strength associated with the uncertainty response which could be the primary reason that animals use it (Smith et al., 2008). A second methodological problem emerges from providing transparent feedback. Smith et al. have argued that giving feedback about a specific response on every trial (also a methodological norm) makes reinforcement predictable because the consequence of each behavior can be associated with the specific stimulus-response pair that produces the positive or negative outcome (Smith et al., 2008). Lastly, Smith et al. argue that the manner in which metacognitive data patterns are interpreted may result in incorrect conclusions. In other words, it is neither sufficient nor safe to assume that the existing methodology for studying metacognition provides conclusive results without having first comprehensively modeled the assumptions on which the paradigm rests. In order to fully appreciate how Smith et al. arrived at the arguments discussed above it is necessary to discuss the various components used in the model.

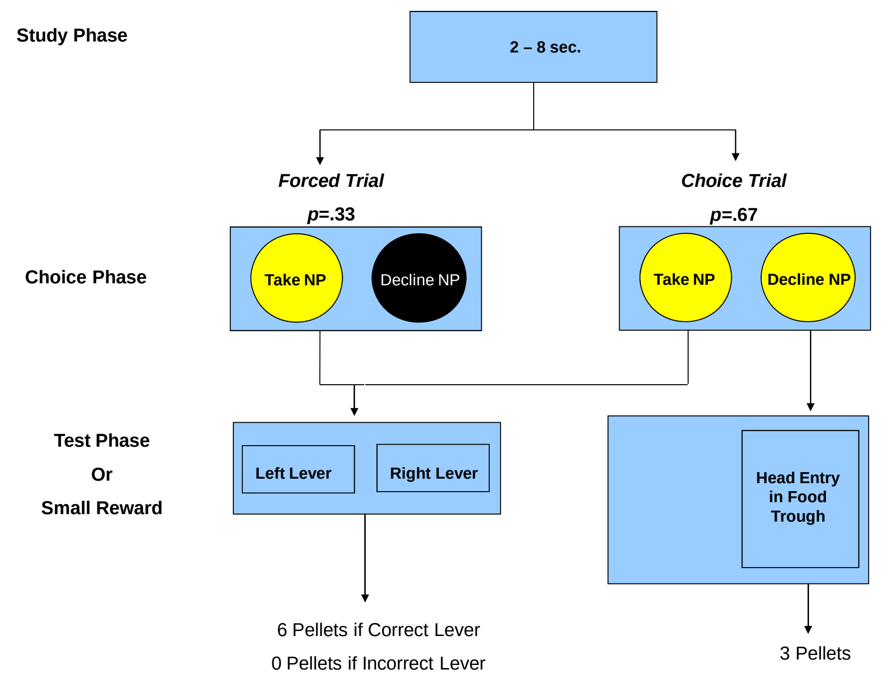

Figure 2.2. Schematic representation of procedure from Foote & Crystal (2007).

Traditionally, the difference between accuracy on forced and choice trials in metacognition paradigms has been the putative evidence needed for a conclusion of animal metacognition. The important feature of Smith et al.'s (2008) response-strength model is that it can produce a large difference in accuracy between choice trials and forced trials by using simpler alternative explanations (see Figure 2.3). Specifically, the response-strength model uses associative and habit-formation principles to explain that observed accuracy and choice data are the result of the response strength (i.e., associative strength) associated with rewarding the decline (or uncertainty) response across a continuous range of stimuli (i.e., the stimulus continuum). Smith et al. argue that rewarding the decline response results in the production of a low-frequency tendency to select the decline response which is independent of stimulus properties (e.g., duration, density, brightness) and is constant across the stimulus continuum. In other words, the associative strength of the reward, not metacognition, is responsible for the production of the decline response. Additionally, Smith et al. argue that generalizations of anchor stimuli (e.g., *Stimulus 2* and *Stimulus 8* in Figure 2.3) used to train animals on a discrimination task decrease exponentially away from the anchor stimuli and have equivalently low response strengths at the point where the generalization gradients cross. Thus, generalizations of the anchor stimuli are sufficient for explaining the observed decrease in accuracy across the stimulus continuum. Perhaps the most interesting aspect of the response-strength model is the notion that the response strengths associated with the anchor stimuli and the decline response follow a winner-take-all response rule. For example, if on a given trial the response strength associated with *Stimulus 2* is stronger than the response strength associated with the *decline response* then the animal will respond *Stimulus 2*. On the other hand, if the response strength associated with *Stimulus 2* is less than that of the *decline response* the animal will respond *decline*. Importantly, Smith et al. used simulations to show that this non-metacognition model produces two key pieces of data. First, the decline response increases as a function of increasing difficulty, and second,

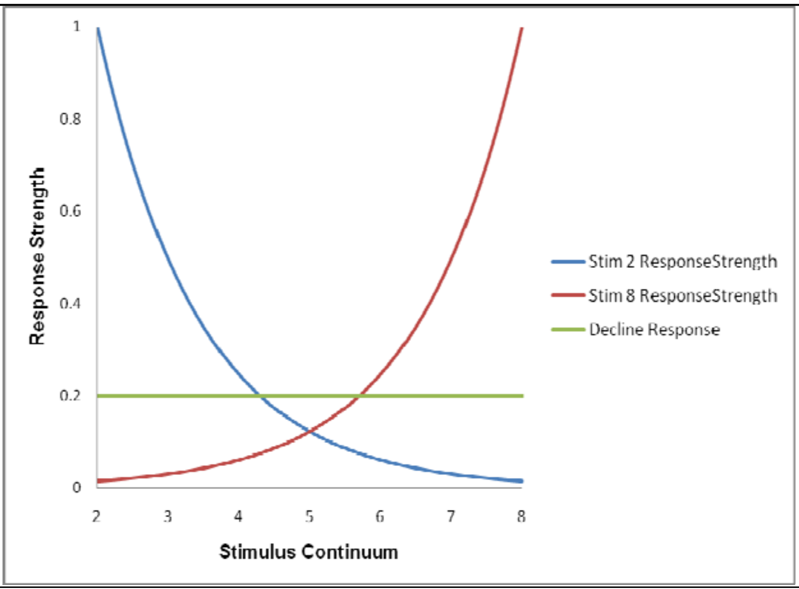

Figure 2.3. Smith, Beran, Couchman & Coutinho's (2008) associative model of stimulus-generalization/response strength. Smith, Beran, Couchman & Coutinho's (2008) associative model of stimulus-generalization/response strength. Smith and colleagues have developed a model of the putative metacognitive data pattern by using exponential curves, response strength, and a reward threshold as a function of the subjective impression of the stimulus. One exponential curve represents the response strength of a short temporal duration while the other exponential curve represents the response strength of a long temporal duration. Both exponential curves cross one another as response strength decreases. Additionally, there is a relatively low and flat threshold for the escape response. Because the reward threshold is above the point where the two exponential curves cross, the escape response is selected for trials that would otherwise generate low accuracy as the response strength for the temporal stimulus begins to weaken. Importantly, this model predicts the emergence of an accuracy divergence because poor-performing trials are selectively removed when tests are declined on difficult trials.

superior performance on trials in which a test is chosen (rather than forced) increases as a function of increasing stimulus difficulty (i.e., an accuracy divergence).

Future Directions

As a result of the response-strength model, current methods for experiments in comparative metacognition have been called into question. The implications from response-strength model pose a challenge for future research on comparative metacognition. Furthermore, an additional problem arises for the existing methodology of metacognition experiments because the traditional method for obtaining accuracy is no longer valid (i.e., the accuracy advantage of choice trials over forced trials). The simple associative mechanisms used in the response-strength model are able to produce a simulated accuracy function that is similar to observed accuracy functions in the literature. Therefore, if it is possible for putative metacognition performance to be produced using simple associative mechanisms, then it is quite feasible that observed accuracy functions in the existing literature could be the result of these same mechanisms. However, one advantage of identifying alternative mechanisms is the potential for refinement of existing experimental paradigms and the development of new experimental paradigms. Comparative psychology must now improve existing methods or develop new ones that are able to completely rule out associative mechanisms. Essentially, what is needed to distinguish metacognition from other simpler processes is an independent line of evidence that cannot be explained by associative mechanisms.

CHAPTER 3

“PLAY IT AGAIN” EXPERIMENT

The goal of the current experiment was to develop a new method for testing metacognition in animals that would not be subject to alternative explanations (e.g., associative or habit formation explanations). Rats were first presented with a brief duration of noise that they would need to classify as either short or long in a later duration test (Figure 3.1). On some trials rats were forced to immediately take a duration test (i.e., a forced-test trial). On other trials, rats were sometimes required to repeat the same brief noise duration (i.e., repeat the stimulus) that they had just heard (i.e., a forced-repeat trial). On a third type of trial, rats had the option to take a duration test immediately or repeat the noise duration (i.e., choice-take trials and choice-repeat trials, respectively). The intent of forcing rats to repeat the brief noise duration was to establish a baseline for comparing performance on forced-repeat tests to performance on tests where rats had the option to repeat the noise duration.

We assume that animals are sometimes in a low internal state of performance and sometimes in a high internal state of performance. Metacognition is the hypothesis that animals can discriminate internal, low versus high states of performance. We assume that animals choose to repeat the stimulus if they are in a low state of performance on difficult trials. Although a second presentation of the stimulus is expected to increase accuracy, it is assumed that the initial low state of performance continues throughout the trial in which the animals choose to repeat. Trials in which animals choose to repeat the stimulus function to isolate low states of performance. By contrast, trials in which animals are forced to repeat the stimulus have a combination of low and high states of performance. Therefore, if rats have knowledge about their own cognitive states, rats should be less accurate when they have the option to repeat difficult stimulus duration tests than when they are forced to repeat difficult stimulus duration tests. This prediction is unique to difficult stimulus durations because low states of performance rarely occur on easy stimulus durations. In contrast, alternative (i.e., non-metacognitive) proposals predict equal performance on easy and difficult trials when rats have the choice to repeat the stimulus and when they are forced to repeat the stimulus.

General Method

Subjects

Eight male Sprague-Dawley rats (*Rattus norvegicus*; Harlan, Madison, WI; 85 days old) were individually housed in a colony room with a reversed 12-12 light-dark schedule (light offsetat 07:00, onset at 19:00). Testing began when the rats were approximately 131 days old and weighed an average of 277 g. During pre-training and testing sessions rats received 45-mg pellets (F0165, Bio-Serv, Frenchtown, NJ) and later received a supplemental ration of 5001-Rodent-Diet (Lab Diet, Brentwood, MO) for a total daily ration of 15-20 g. Water was available continuously. All procedures were approved by the institutional animal care and use committee and followed the guidelines set forth by the National Research Council *Guide for the Care and Use of Laboratory Animals*.

Pilot Study. Rats were initially trained to discriminate short and long stimulus durations (see duration discrimination training under preliminary training below) and to use nose-pokes (see addition of take and repeat responses under preliminary training below). Terminal performance from these data were originally used to identify the difficult stimulus for each rat.

Next, rats received forced-repeat, forced-test, and choice trials, as outlined in Figure 3.1 and described in greater detail below. However, one of the limitations of this approach in the pilot study was that the identification of a difficult stimulus was fixed and was based on increasingly old baseline data. Therefore, the approach developed in the pilot was refined to retrain the rats on the duration discrimination in the first half of the daily session, followed by the repeat-the-stimulus task for the remainder of the session (as described below). One improvement over the

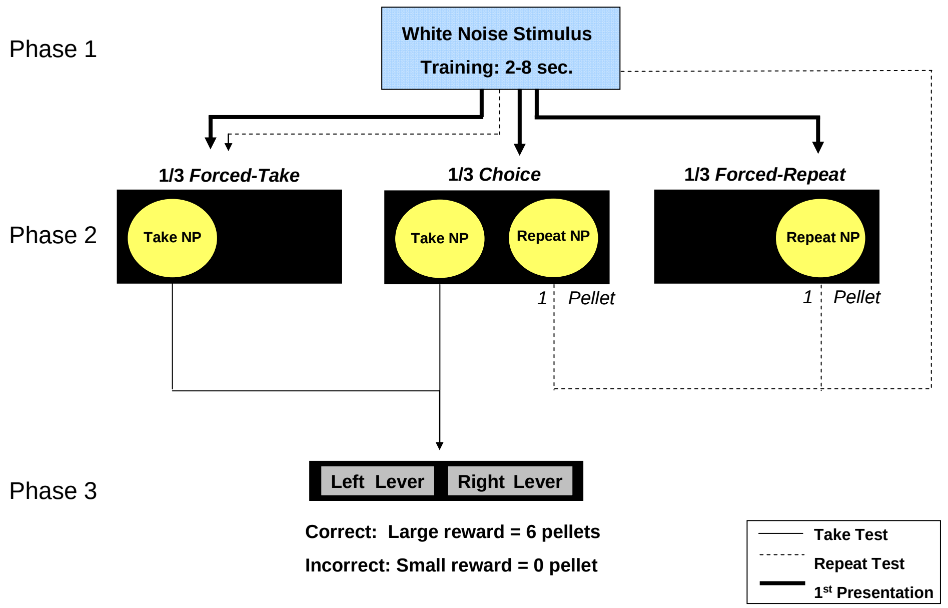

Figure 3.1. Schematic representation of the "Play it Again" method. Nose-poke is abbreviated as NP. Each trial consisted of three phases. In phase 1 rats were presented with a white noise stimulus duration. In phase 2 rats were forced to either take a duration test, repeat the stimulus from phase 1, or were given the choice to take a test or repeat the stimulus. In phase 3 rats used levers to classify the stimulus from phase 1 as short and long. The stimulus duration and trial type were randomly selected before the beginning of each trial. Bold lines indicate one of three ways in which a trial could proceed. Thin solid lines indicate how a forced-take (left side of diagram) or a choice-take trial (center of diagram) would proceed. Dashed lines indicate the two ways a trial could be repeated, a choice-repeat trial (center of diagram) or a forced-repeat trial (right side of diagram).

pilot is the ability to obtain a daily estimate of the difficult stimulus. The total number of sessions for the pilot study varied (due to individual differences in learning) according to subject as follows: KK1 and KK7 completed 11 sessions; KK2, KK3, KK4, KK6, and KK8 completed 21 sessions; and KK5 completed 14 sessions.

Apparatus

Eight identical operant chambers (30 cm × 28 cm × 23 cm, width × height × depth; Med Associates ENV-007, Georgia, VT), each located within a ventilated sound-attenuation cubicle (ENV-016M, 66 cm × 56 cm ×36 cm, W×H×D), were used for the experiment. Each operant chamber contained a recessed food trough (ENV-200R2M, 5 cm × 5 cm) equipped with photobeams (used to detect head entries; ENV254, 1 cm in from food trough, 1.5 cm from bottom of food trough) that was centered horizontally (63 cm above the floor) between two retractable levers (ENV-112CMX) on one wall of the chamber. A 45-mg pellet dispenser (ENV-203-45IRX) was located on the outside wall of the chamber and was attached to the food trough. A photobeam located on the feeder detected successful pellet delivery. A pellet dispenser would make up to four additional attempts to dispense pellets if a failure was detected. A water bottle with an attached sipper tube was placed on the outside wall opposite of the food trough. The sipper tube was inserted into the chamber via a 1 cm × 1.5 cm opening in the wall. A nose-poke aperture was located to either the left or the right of the sipper tube and contained a photobeam that detected individual entries. The nose-poke apertures were small recessed openings (2.5 cm diameter) located on one wall of the operant chamber. A retractable automated guillotine door (ENV-210M) was used to give/restrict access to each nose-poke opening. The floor of the chamber was made of 19 stainless steel rods (4 mm diameter, 15.5 mm spacing) and a stainless steel waste tray was located below the chamber floor. Other equipment included a clicker (ENV-

135M), lights (ENV-215M and ENV-227M), and speaker (ENV-225SM), a photobeam lickometer (ENV-251L), and four equally spaced photobeams that were 4 cm above the floor. A computer with a Celeron processor (850 MHz) running Med-PC (version 4.0) was located in a nearby room and controlled experimental events and recorded the time at which each event occurred (10-ms accuracy).

Procedure

Preliminary training

Pre-training. Rats were given feeder training that consisted of one food pellet being delivered per minute, accompanied by a click before pellet delivery, for one 30-minute session. Next, rats underwent 3 daily sessions of lever training. Lever training occurred on alternate levers (10 trials on the left lever followed by 10 trials on the right lever for a total of 30 trials for each lever), for 60 minutes or until 60 pellets were earned for each session. The delivery of one pellet was contingent upon a single lever press. Feeder and lever training were followed by 4 daily sessions of nose-poke training. Nose-poke training occurred on alternate nose-pokes (10 trials on the left nose-poke followed by 10 trials on the right nose-poke for a total of 30 trials for each nose-poke), for 60 minutes or until a total of 60 pellets had been earned for each session. Initially, one guillotine door retracted to allow access to the nose-poke opening. The delivery of one pellet was contingent upon a rat inserting its snout into the accessible nose-poke opening. Immediately after 10 pellets were earned the guillotine door closed and the other nose-poke became accessible.

Duration-discrimination training. The stimuli that were used for the duration-discrimination training were eight logarithmically spaced white noise stimuli: 2.00, 2.44, 2.97, 3.62, 4.42, 5.38, 6.56 and 8.00 seconds. Stimuli were chosen by independent random selection before the start of each trial. Rats were trained to discriminate short and long noise durations. Short durations ranged from 2.00-3.62 seconds and long durations ranged from 4.42-8.00 seconds. Duration discrimination became more difficult as a stimulus approached 4.00 seconds (i.e., the easiest durations to discriminate were 2.00, 2.44, 2.97, 5.38, 6.56, and 8.00 seconds while this most difficult durations to discriminate were 3.62 and 4.42 seconds). Rats were trained with all 8 stimulus durations. The inter-trial interval (ITI) was 8-10 minutes for each 9-hour daily session. A trial began with the presentation of a 70-dB white noise stimulus duration that rats had to classify as either short or long. Rats indicated their choice by pressing one lever for short and one lever for long (lever assignment was counterbalanced across subjects prior to the beginning of the experiment). Rats received a large reward of 6 pellets for correctly discriminating a stimulus and received no reward for incorrectly discriminating a stimulus. Duration-discrimination training continued until each subject achieved an average accuracy score of at least 75% across all eight stimulus durations. The number of sessions for duration discrimination training varied by subject (due to individual differences in learning) as follows: KK1, KK5, and KK7 completed 15 sessions, KK2, KK4, KK6, and KK8 completed 10 sessions.

Addition of “take” and “repeat” responses. The “take the duration test” and “repeat the stimulus” responses were added next. Left and right nose-pokes were randomly assigned as the “repeat the stimulus” and “take the test” nose-pokes; assignment of nose-pokes was counterbalanced across subjects and conditions (i.e., taking the duration test or repeating the stimulus). Trial type (i.e., “take the duration test” or “repeat the stimulus”) was chosen by independent random selection prior to the beginning of each trial. Rats were allowed to choose to “take the duration test” and “repeat the stimulus” but were only allowed to repeat the stimulus once per trial. Rats were trained until they achieved an accuracy score of at least 75% on the four easiest short and long stimulus durations (easy short = 2.00 and 2.44; easy long = 6.56 and 8.00) before moving to testing. The number of sessions needed to learn the "take" and "repeat" responses varied by subject (due to individual differences in learning) as follows: KK 1 and KK 7 completed 11 sessions; KK2, KK3, KK4, KK6, and KK8 completed 21 sessions; and KK5 completed 15 sessions.

Metacognition Testing

Repeat the Stimulus Procedure

Metacognition was assessed by using the repeat-the-stimulus procedure that was developed as a refinement of the procedure used in the pilot experiment. The repeat-the-stimulus procedure established a daily estimate of the most difficult stimulus for each subject. Each daily session was comprised of two parts: 1.) In the first part of each daily session rats received retraining on the duration discriminations, 2.) In the second part of the session rats proceeded through the repeat-the-stimulus task (see Figure 3.1). The transition between the two parts occurred approximately halfway through the procedure.

Each repeat-the-stimulus trial consisted of three phases. During Phase 1 a white-noise stimulus duration was presented to the rat and was the stimulus the rat was later tested on in Phase 3 (see below for a detailed description of each condition). During Phase 2, the rat could either be forced to take a duration test (1/3 of trials), forced to repeat the stimulus (1/3 of trials), or have the choice to take a duration test or repeat stimulus depending upon trial type (1/3 of trials). In Phase 3, rats pressed one lever to identify the stimulus as short or pressed the other lever to identify the stimulus as long. If the rat correctly identified the stimulus duration it received a large reward of 6 pellets. However, if it identified the stimulus duration incorrectly it received no reward (i.e., 0 pellets). Rats received an additional small, one pellet reward immediately upon choosing or being forced to choose to repeat the stimulus. Trial type was randomly selected prior to the start of each trial. The stimulus duration was chosen by independent random selection before the start of each trial. Rats were tested using three different trial types (outlined below) until their average accuracy level was greater than or equal to 75% for the four easiest short and long stimulus durations (easy short = 2.00 and 2.44; easy long = 6.56 and 8.00).

Conditions

Forced-test trials. On forced-test trials rats were forced to take a stimulus duration test. Forced-test trials began with the presentation of a white-noise stimulus duration (i.e., Phase 1, see left side Figure 3.1). In Phase 2 the guillotine door covering the “take the test” nose-poke retracted and allowed access to *only* the “take the test” nose-poke (e.g., the left nose-poke in Figure 3.1). After the guillotine door retracted, rats were required to break the photobeam in the “take the test” nose-poke aperture (e.g., the left nose-poke in Figure 3.1) in order to move to Phase 3. As soon as the rat broke the photobeam in the nose-poke the guillotine door closed. In Phase 3 levers inserted into the chamber and, rats were required to press one lever. If the rat classified the stimulus duration correctly it received a reward of 6 pellets, and if it was incorrect it received no pellets.

Forced-repeat trials. On forced-repeat trials rats were forced to repeat a stimulus (i.e., the same stimulus duration was presented again) which was later followed by a forced-test. Forced-repeat trials began with the presentation of a white-noise stimulus duration (i.e., Phase 1, see right side of Figure 3.1). In Phase 2, rats were forced to hear a re-presentation of the stimulus duration presented during Phase 1. Phase 2 began with the retraction of the guillotine door covering the “repeat the stimulus” nose-poke (e.g., the right nose-poke in Figure 3.1). Only the “repeat the stimulus” nose-poke was accessible and the guillotine door on the other nose-poke remained closed. Rats were required to break a photobeam in the “repeat the stimulus” nose-poke in order to hear a re-presentation of the stimulus duration presented in Phase 1 and receive a 1 pellet reward. After rats heard the stimulus duration for a second time the remainder of the trial proceeded in the same manner as a forced-take trial.

Choice trials. In choice trials rats had the opportunity to choose to take the duration test (i.e., choice-take) or to hear a re-presentation of the stimulus (i.e., choice-repeat). Choice trials began with the presentation of a white-noise stimulus duration (i.e., Phase 1, center of Figure 3.1). Afterwards, in Phase 2, both guillotine doors covering the “take the test” and the “repeat the stimulus” nose-pokes retracted which allowed the rats to access both nose-pokes. If rats chose to take a duration test, the remainder of the trial proceeded as in a forced-take trial. On the other hand, if rats chose to repeat the stimulus (i.e., hear a re-presentation) the remainder of the trial proceeded as in a forced-repeat trial.

CHAPTER 4

EXPERIMENTAL RESULTS

The total number of sessions for the metacognition testing procedure differed for individual rats. Fifteen sessions were omitted for all subjects due to an equipment problem. Subject KK1 exhibited a response bias on the duration discrimination task and did not progress to subsequent testing procedures because its performance never exceeded 75% correct for the four easiest short and long stimulus durations (easy short = 2.00 and 2.44; easy long = 6.56 and 8.00). As a result, subject KK1 never received the metacognition testing procedure. Additionally, seven of KK5's test sessions were omitted due to a brief illness. Table 4.1 contains the total number of metacognition test sessions for each subject.

The proportion of correct answers (i.e., accuracy) and the frequency for each trial type were recorded for all eight stimulus durations during each daily session for each subject. Data from each daily session was divided into two parts. Data from the first part of each daily session were used to estimate accuracy for the difficult stimuli and for the easy stimuli. The first part of each daily session identified the stimulus that each subject found most difficult by determining whether the accuracy for each stimulus duration was less than 75% correct. To estimate accuracy for the easy stimuli, the following stimulus durations were used: 2.00 s, 2.44 s, 6.56 s, and 8.00 s. To estimate accuracy for the difficult stimuli, the following durations were used: 3.62 s and 4.42 s. These two stimulus durations were selected based upon what subjects found to be the most difficult stimulus on a given day. Data from the second part of each daily session were then examined for easy and difficult conditions using the stimuli identified for each subject.

Table 4.1

Total Number of Metacognition Test Sessions

Subject	Total Sessions
KK-2	70
KK-3	150
KK-4	130
KK-5	140
KK-6	130
KK-7	110
KK-8	160

Note. Totals reflect the final number of sessions that were included in the data analysis. Totals exclude omitted sessions.

Afterwards, proportion correct was calculated by dividing the total number of correct trials into the total number of trials for each stimulus duration separately for each trial type (i.e., choice-repeat or forced-repeat). The analyses were performed on the 70 terminal sessions for all subjects (70 sessions was the minimum number of sessions completed by all subjects).

If rats have knowledge about their own cognitive states, then rats should be more accurate when forced to repeat a difficult stimulus duration test than when they have the option to do so on difficult trials. Figure 4.1 shows accuracy as a function of difficulty for both choice-repeat and forced-repeat trials. As expected, accuracy is lower on difficult trials compared to easy trials. Importantly, choice-repeat accuracy is lower than forced-repeat accuracy on difficult trials. A two factor analysis of variance (ANOVA) was performed on accuracy for trial type (i.e., choice-repeat trials vs. forced-repeat trials) and stimulus difficulty (i.e., easy vs. difficult). The ANOVA revealed a significant main effect of stimulus difficulty, $F(1, 6) = 133.05$, $p <.001$, partial $\eta^2 = .957$, and a significant main effect of trial type, $F(1, 6) = 13.68$, $p <.01$, partial $\eta^2 = .695$. Additionally, the ANOVA revealed a significant interaction between stimulus difficulty and trial type, $F(1, 6) = 6.98$, $p <.038$, partial $\eta^2 = .538$. A planned comparison of accuracy on difficult stimuli for forced-repeat trials and choice-repeat trials was analyzed using a paired samples *t*-test (*N* = 7, *M* = - .051, *SEM* = .012). As predicted, rats were more accurate on trials in which they were forced to take a difficult stimulus duration test than on trials where they chose to take a difficult test, *t* (6) = - 4.39, *p* = .005 (two-tailed), *d* = 0.81, 95 % CI [- .079, - .02].

On choice trials, rats would be expected to repeat the stimulus more often on difficult stimulus durations than on easy stimulus durations according to both metacognition and non-metacognition proposals. The observed rate of choosing to repeat the stimulus was 0.638 and 0.671 for easy and difficult conditions, respectively. A paired samples *t*-test (*N* = 7, *M* = .033, *SEM* = .022) did not reveal a significant difference in the frequency of choosing to repeat the stimulus more often on difficult stimulus durations than on easy stimulus durations, *t* (6) = 1.49, *p* = .187 (two-tailed), *d* = 0.09, 95 % CI [-.021, .088]. The data are consistent with the predicted direction for both metacognition and non-metacognition proposals (i.e., rats had a very small tendency to repeat the stimulus more often on difficult stimulus durations) although the difference was not statistically significant.

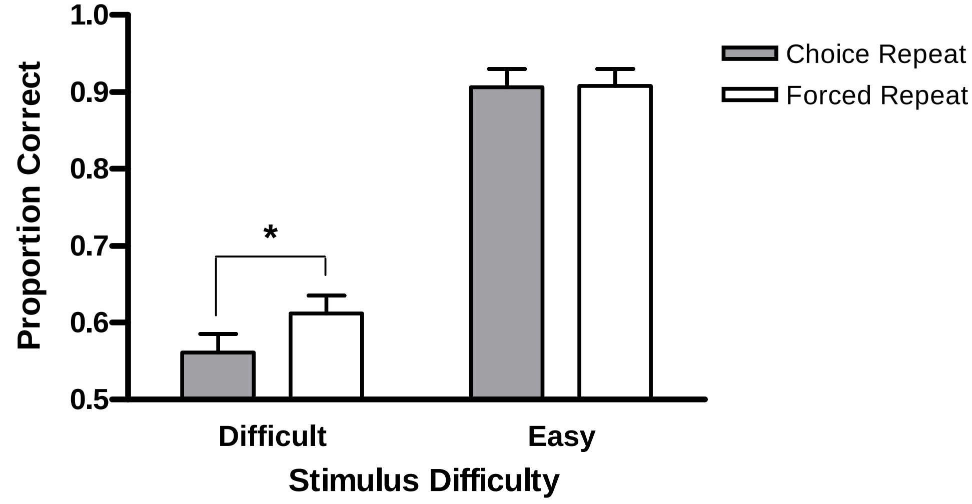

Figure 4.1. Accuracy on each trial type as a function of stimulus difficulty. Difficult stimuli included stimulus durations 3.63s and 4.42 s. Easy stimuli included stimulus durations 2.00 s, 2.44 s, 6.56 s and 8.00 s. The asterisk indicates a significant difference at an alpha-level of $p = .005$. Error bars represent standard errors.

CHAPTER 5

SIMULATION OF REPSONSE STRENGTH MODEL

Existing methodology in comparative metacognition lacks validation because putative metacognition data patterns may be explained by a low-level response-strength model (i.e., without metacognition). Although the empirical data demonstrate what we believe to be results of metacognition, it is imperative that the critical metacognition prediction made by the newly developed “Play it Again” method be thoroughly evaluated. This objective was accomplished through simulations. We assume that animals are sometimes in a high or low state of performance. Metacognition is the hypothesis that animals know when they are in a high or low internal state of performance. We also assume that when animals are in a low state of performance, and they know that they are, they will choose to repeat the stimulus. In the “Play it Again” procedure, choice-repeat trials functioned to isolate low states of performance while forced-repeat trials contained a mixture of high and low states of performance. Accuracy measured from choice-repeat trials and forced-repeat trials is the most important and critical comparison made in the procedure. If animals have metacognition, we would expect to see an accuracy difference between choice-repeat and forced repeat trials when tested with a difficult problem. Specifically, we would expect to see lower accuracy on choice-repeat trials than on forced-repeat trials. In contrast, the response strength model predicts no accuracy difference on forced-repeat and choice-repeat trials. Simulations were conducted to determine whether the response strength model could fit the accuracy difference produced in our data.

Method

The simulation began with an exhaustive search of the parameter space in order to identify the least-squares best fitting parameters (i.e., the parameters that minimized the sum of squared deviations between the empirical data and the simulated values). A minimum, maximum, and step-size was used for each parameter. The parameters that minimized the difference between the data and the simulation were identified.

The simulation closely followed the “Play it Again” procedure and therefore, the process of identifying parameters for the simulation is henceforth described in procedural terms. The range of stimulus durations (2-8 seconds) was expressed as values within the range of 1-71 (following Smith et al.’s simulations). An objective physical stimulus is perceived with variability. This concept was modeled by sampling from a normal distribution with a mean that corresponds to the objective physical stimulus (Stimulus Mean) with a parameter for the standard deviation of the distribution (Stimulus SD). Therefore, a subjective duration was determined on each simulated trial by a random number, a mean, and a standard deviation. The response strength to judge the subjective duration as short or long was determined by an exponential curve (see Figure 2.3), a sensitivity parameter in the exponent (sens), and the subjective duration described above (errs; i.e., the distance between the subjective and physical durations). The exponential curve was calculated using Smith et al.’s (2008) equation $e^{\text{-sens}\times\text{errs}}$ (p. 691). The decision to repeat the stimulus or take an immediate test was modeled by a flat response threshold (i.e., a constant level of attractiveness, independent of the magnitude of the objective physical stimulus). The flat response threshold was modeled by sampling from a normal distribution with a mean (Threshold Mean) and standard deviation (Threshold SD) as parameters.

To simulate a repeat-the-stimulus condition, each stimulus presentation was modeled with an independent random sample using the same parameters. The subjective duration after two stimulus presentations was modeled by a weighted average of the two independent stimulus presentations (Weighted Average). A winner-take-all response rule was applied to response strengths for the short, long, and repeat-the-stimulus responses. The duration classification response was based on a winner-take-all response rule for short and long. Impossible values (e.g., durations below zero) were discarded and re-sampled. Accuracy is based on averaging (i.e., the relative frequency of) outcomes for incorrect (represented as 0) and correct (represented as 1) outcomes.

Each simulation consisted of 10,000 trials. The parameter set for an individual simulation was selected from the minimum, maximum, and step size values shown in Table 5.1. An exhaustive search of the 282, 240 sets of parameters was conducted for easy (corresponding to durations 2.00 and 2.44 seconds; stimulus categories 1 and 11 in the simulation) and difficult (corresponding to 3.62 seconds; stimulus category 31 in the simulation) conditions. Because difficulty in the duration task is indexed by the distance between stimulus duration and the point of subjective equality of 4.0 seconds, we simulated conditions below 4.0 seconds. Duplication of equivalent levels of difficulty above 4.0 seconds would have increased the number of simulations to over 2.5 million simulations, which was deemed prohibitive with respect to the computer resources. The simulation outcomes for the easy conditions were averaged (as was done for the data) at each parameter set.

Table 5.1

Minimum, Maximum, and Step Size Values

Parameter	Minimum Value	Maximum Value	Step Size
Threshold Mean	0.1	1.0	0.1
Threshold SD	0.0	0.0	0
SenseS Mean	0.005	0.305	0.05
SenseS SD	0.005	0.305	0.1
SenseL Mean	0.005	0.305	0.05
SenseL SD	0.005	0.305	0.1
Stimulus SD	0.0	25.0	5.0
Weighted Average	0.0	1.0	0.2

Note. The stimulus mean was 1, 11, and 31 in separate simulations. The stimulus range was 1 to 71, which corresponds to 2.0 and 8.0 seconds. Easy conditions correspond to durations 2.00 and 2.44 seconds and stimulus categories 1 and 11 in the simulation. Difficult conditions correspond to the duration of 3.62 seconds and stimulus category 31 in the simulation.

In each simulation, six values were estimated: the proportion of trials in which a choice to repeat the stimulus occurred in easy and difficult conditions; proportion correct on forced repeat trials in easy and difficult conditions; and proportion correct on choice repeat trials in easy and difficult conditions. The sum of squared differences between observed (data) and expected (simulation) values was calculated for each set of parameters using the six proportions listed above. The set of parameters that minimized the sum of squared deviations is presented in Table 5.2. The set of parameters identified by the least-squares method described above was examined further as follows. Seven new simulations (each consisting of 10,000 trials) were performed to estimate variability across seven simulated subjects, which corresponds to the sample size in the data.

Table 5.2

Least-Squares Best Fit Parameter Values

Parameter	Best fit
Threshold Mean	0.2
Threshold SD	0.0
SenseS Mean	0.255
SenseS SD	0.005
SenseL Mean	0.055
SenseL SD	0.205
Stimulus SD	20.0
Weighted Average	0.0

Simulation Results

The goal of the simulation was to determine whether the response-strength model could fit the observed accuracy difference in the experimental data. The exhaustive search of the parameter space identified a set of parameters that minimized the error between the empirical data (see Figure 4.1) and simulated values. The results of the simulation are shown in Figure 5.1. Unlike the experimental data, the simulated data do not show a difference in accuracy between choice repeat trials and forced repeat trials for difficult stimuli, t (6) = 0.85, $p = 0.4$ ($N = 7$, $M = 0.0016$, $SEM = 0.0018$). According to the results of the simulation, we would expect no significant difference in accuracy on choice repeat trials and forced repeat trials for difficult stimulus durations, which is contrary to the data shown in Figure 4.1.

Importantly, the empirical data produced a statistically significant difference between forced repeat and choice repeat accuracy in the difficult condition. Consequently, we wanted to determine whether the magnitude of this difference was statistically different from the expected difference according to the response strength simulation. An independent samples t-test ($M = 0.0493$, $SEM = 0.0117$) was conducted on the empirical data and the simulated data. The magnitude of the accuracy difference in forced-repeat and choice-repeat conditions was larger in the empirical data than in the simulation, *t* (12) = -4.21, *p* = 0.001 (two-tailed). Hence, the empirical data showed an accuracy difference that could not be explained by the response strength model.

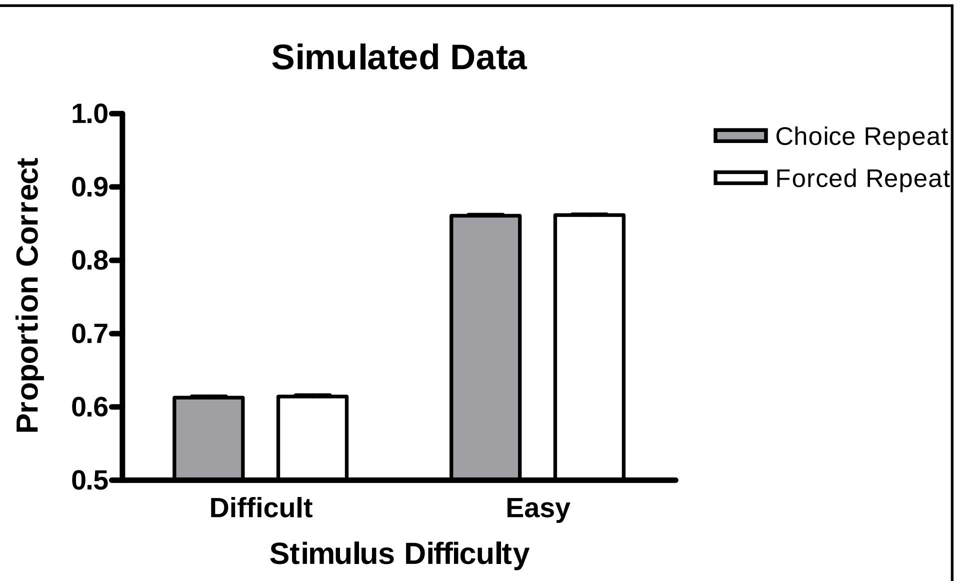

Figure 5.1. Simulated data. Unlike the empirical data shown in Figure 4.1, the simulations show no difference in accuracy for difficult choice repeat trials and difficult forced repeat trials. Accuracy on each trial type as a function of stimulus difficulty for simulated data. Difficult stimuli included stimulus Level 31 (i.e., stimulus duration 3.62 s). Easy stimuli included stimulus Level 1 and 11 (i.e., stimulus durations 2.00 s, 2.44 s). The simulated rate of choosing to repeat the stimulus was 0.578 and 0.730 for easy and difficult conditions, respectively. Error bars represent 1 SEM.

CHAPTER 6

DISCUSSION

The objective of the experiment and simulation was to develop a method for testing metacognition in animals that could not be interpreted by a response-strength model. In the “Play it Again” procedure, rats first heard a brief noise duration that they had to classify as either short or long in a subsequent duration test (Figure 3.1). In some trials, rats were forced to repeat the stimulus duration that they had just heard (forced-repeat trials), forced to immediately take a duration test (forced-take trials), or were required to choose to take a duration test immediately or repeat the same noise duration (choice-repeat trials). Metacognition, but not the response strength model, predicts that performance on choice-repeat trials would be lower than on forced-repeat trials.

Metacognition is the hypothesis that rats (or other animals) can discriminate between low and high internal states of performance. We assume that rats are sometimes in low or high internal states of performance. We assume that they choose to repeat a difficult noise duration if they are in a low state of performance. While a second presentation of the noise duration may increase performance, we assume that the initial low state of performance continues throughout the trial in which a rat chooses to repeat the noise duration. Consequently, choice-repeat trials served to isolate low states of performance. By contrast, forced-repeat trials contained a combination of low and high states of performance. Therefore if rats had metacognition, we would expect to see an accuracy difference between choice-repeat and forced repeat trials. Specifically, we would expect to see lower accuracy on choice-repeat trials than on forced-repeat trials. Our data showed the predicted critical difference in accuracy on choice-repeat and forced-repeat trials. In contrast, the response-strength model does not predict an accuracy difference on choice-repeat and forced-repeat trials.

While our primary interest lay in detecting the accuracy difference predicted by metacognition, we were also interested in determining the frequency with which rats chose to repeat difficult stimulus discriminations. Both metacognition and the response-strength model predict that rats should choose to repeat difficult stimulus discriminations more often than easy stimulus discriminations. Our results suggested that rats had a small tendency for choosing to repeat the stimulus more often for difficult stimulus durations. Although this finding is consistent with the direction of our prediction for both metacognition and the response-strength model, it was not a statistically significant difference. Reinforcement of the repeat nose-poke response (see Figure 3.1) may have masked the expected difference in take versus choice nose-poke responses. To encourage the rats to sample the repeat nose-poke response, a pellet reward was delivered contingent on the repeat nose-poke response; this reinforcement was needed to offset the more attractive delay to reinforcement available when the rat chose to take an immediate test. Therefore, we suspect that the design of the procedure allowed for greater sensitivity for detecting differences in duration-discrimination accuracy than for choice of differentially reinforced nose-poke responses. Another possible explanation for why rats did not exhibit a preference for choosing to repeat difficult stimuli could be that the effect size of the difference was too small to be detected as significant. The small effect size could be the result of lack of power due to a small sample size of seven rats. For future studies, this problem could be remedied by performing an a priori power analysis to determine the appropriate number of subjects to use with this procedure.

The results from the simulations also provide evidence for the conclusion that rats have metacognition. We performed simulations of the “Play it Again” procedure to determine whether the response-strength model could fit the accuracy difference between forced-repeat and choice-repeat trials. Specifically, we conducted an exhaustive search of parameters and found that the simulated data predicted equal performance on difficult stimulus discriminations for both forced-repeat and choice repeat trials. In contrast, our experimental data showed a significant accuracy difference between choice-repeat and forced-repeat trials as predicted by metacognition. The simulation demonstrated that the response-strength model did not fit our experimental data, which validated the “Play it Again” procedure.

Conclusions

The goal of the “Play it Again” experiment was to develop a new method for testing metacognition in animals that would not be subject to associative or habit formation explanations. We believe that the findings from the “Play it Again” experiment and the simulations document a valid method for testing metacognition in animals. The simulations suggest that the response strength model does not predict an accuracy difference between choice-repeat trials and forced-repeat trials in our data. This prediction is contrary to the prediction made by metacognition and to the accuracy difference observed in our data. Consequently, our findings suggest that rats have metacognition.

Like the present study, Foote and Crystal’s (2007) study examined metacognition in rats. The response-strength model (Smith et al., 2008) represented an attractive alternative explanation to Foote and Crystal’s metacognition data prior to the collection of new data using the “Play it Again” procedure. However, the response-strength model cannot explain the new data. Because it is necessary to propose metacognition in rats to explain the new data, it is most parsimonious to also explain Foote and Crystal’s data by appeal to metacognition. Thus, the new and old data may be viewed as providing converging lines of evidence for metacognition in rats.

Our findings demonstrate that valid methods for testing metacognition in animals can be developed. Importantly, performing simulations required us to precisely develop our new method and specify the predictions made by non-metacognitive explanations. Although our method was designed to test metacognition in rats, an important next step is to use similar methods (i.e., a combination of novel procedures and simulations) with other animals. Using similar methods with other animals would provide a means of validation for our metacognition method and would also provide converging lines of evidence for metacognition in other animals. We believe the “Play it Again” method and our simulations have the potential to resolve controversies about the existence of metacognition in non-human animals.

REFERENCES

Basile, B.M., Hampton, R.R., Suomi, S.J., & Murray, E.A. (2009). An assessment of memory awareness in tufted capuchin monkeys (*Cebus apella*). *Animal Cognition, 12,* 1169-180.

Budson, A.E., Dodson, C.S., Daffner, K.R., and Schacter, D.L. (2005). Metacognition and false recognition in Alzheimer's disease: Further exploration of the distinctiveness heuristic. *Neuropsychology, 19*(2), 253-258.

Call, J., & Carpenter, M. (2001). Do apes and children know what they have seen? *Animal Cognition, 4*, 207-220.

Clayton, N.S. & Dickinson, A.D. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature, 395*, 272-278.

Clayton, N.S. & Dickinson, A.D. (1999a). Scrub jays (*Aphelocoma coerulescens*) remember the relative time of caching as well as the location and content of their caches. *Journal of Comparative Psychology, 113*, 403-416.

Clayton, N.S. & Dickinson, A.D. (1999b). Memory for the contents of caches by scrub jays. *JEP: ABP, 25*, 82-91.

Clayton, N.S., Yu, K. & Dickinson, A.D. (2001), Scrub jays (*Aphelocoma coerulescens*) can form integrated memory for multiple features of caching episodes. *JEP: ABP, 27*, 17-29.

Crystal, J.D. & Foote, A.L. (2009). Metacognition in animals. *Comparative Cognition & Behavior Reviews, 4,* 1-16.

Emery, N. J. & Clayton, N.S. (2001). Effects of experience and social context on prospective caching strategies by scrub jays. *Nature, 414*, 443-446.

Foote, A. L., & Crystal, J. D. (2007). Metacognition in the rat. *Current Biology, 17*, 551-555.

Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences of the United States of America, 98*(9), 5359-5362.

Hampton, R. R., Zivin, A., & Murray, E. A. (2004). Rhesus monkeys *(Macaca mulatta)* discriminate between knowing and not knowing and collect information as needed before acting. *Animal Cognition, 7,* 239-246.

Hampton, R. R. (2009). Multiple demonstrations of metacognition in nonhumans: Converging evidence or multiple mechanisms? *Comparative Cognition & Behavior Reviews, 4,* 17-28.

Kelemen, W.L., & Fulton, E.K. (2008). Cigarette abstinence impairs memory and metacognition despite administration of 2 mg nicotine gum. *Experimental and Clinical Psychopharmacology, 16*(6), 521-531.

Kornell, N., Son, L., & Terrace, H. (2007). Transfer of metacognitive skills and hint seeking in monkeys. *Psychological Science, 18*(1), 64-71.

Kornell, N. (2009). Metacognition in humans and animals. *Current Directions in Psychological Science, 18*(1), 11-15.

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.

Nelson, T.O. & Narens, L. (1980). Norms of 300 general-information questions: Accuracy of recall, latency of recall, and feeling-of-knowing ratings. *Journal of Verbal Learning and Verbal Behavior, 19*, 338-368.

Nelson, T. O. (1984). A comparison of current measures of the accuracy of feeling-of-knowing predictions. *Psychological Bulletin, 95*(1), 109-133.

Nelson, T. O., & Narens, L. (1990). Metamemory: a theoretical framework and new findings. *The Psychology of Learning and Motivation, 26*, 125-173.

Roelofs, J., Papageorgiou, C., Gerber, R.D., Huibers, M., Peeters, F., & Arntz, A. (2007). On the links between self-discrepancies, rumination, metacognitions, and symptoms of depression in undergraduates. *Behaviour Research and Therapy, 45*, 1295-1305.

Smith, J.D. (2005). Studies of uncertainty monitoring and metacognition in animals and humans. In H. S. Terrace, & J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 242-271). New York: Oxford University Press.

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*(4), 679-691.

Smith, J.D., Beran, M. J., Couchman, J. J., Coutinho, M. V. C., Boomer, J. B. (2009). Animal metacognition: Problems and prospects. *Comparative Cognition & Behavior Reviews, 4,* 40-53.

Son, L. K., & Kornell, N. (2005). Metaconfidence judgments in rhesus macaques: explicit versus implicit mechanisms. In H. S. Terrace, & J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 296-320). New York: Oxford University Press.

APPENDIX

LITERATURE REVIEW

Call and Carpenter (2001) performed three experiments with chimpanzees, orangutans, and children to determine whether these participants could remember perceived information about what they had or had not seen. More specifically, Call and Carpenter wanted to know if apes and children were able to access and use information about their own visual perceptions to maximize reward. Call and Carpenter were interested in the "looking behavior" of their participants and made two predictions about their participants' "looking behavior". First, if participants continued to search for a reward in spite of already having found the reward this would suggest an inability to access perceptions about visual information. Second, if participants were able to use different strategies then this would suggest access to, and flexibility in, participants' visual perceptions.

In their first experiment, Call and Carpenter baited one of two square tubes either directly in front of chimps and orangutans or behind a screen that obstructed the subjects' view of the baiting process. Call and Carpenter hypothesized that when baiting was in full view of chimps and orangutans that the animals would be less likely to look into the tubes, and that the animals would perform better than when baiting was obstructed by a screen placed between the tubes and the animals. In addition to manipulating visual information, Call and Carpenter also manipulated the delay (no delay or a 5 sec delay) between baiting the tubes and presenting the tubes to the apes. Call and Carpenter hypothesized that when presentation of the baited tubes was delayed, apes and orangutans would look into the tubes more often and be more successful than when presentation of the baited tubes was not delayed. In order to test their predictions, Call and Carpenter measured both looking and choosing behavior; looking was defined as head or body movement which allowed an animal to gaze into a tube, whereas choosing behavior was defined as touching one of the tubes. Additionally, each subject's search pattern or strategy was measured in an effort to better understand what type of reasoning subjects might be using to solve the task. Results for Experiment 1 indicated that chimpanzees and orangutans looked more often when they had not seen the baiting process, looked more often when the presentation of the tubes was delayed, performed better when they looked into the tubes, and obtained more food in all experimental conditions than would be expected by chance. As a result, Call and Carpenter concluded that apes were able to gather information about the location of the food reward when they did not have the relevant information, and that apes stopped gathering information when they found the food reward. However, limitations from Experiment 1 lead Call and Carpenter to perform Experiment 2.

In Experiment 2, where chimpanzees were the only species examined (Call and Carpenter wanted a more accurate measure for one species), the delay between baiting and tube presentation was removed, and sample size was increased from 3 to 11 chimps. Importantly, 3 of the 11 chimps had participated in Experiment 1, which could have biased results in favor of supporting the proposed hypotheses. Call and Carpenter added a third tube to give their participants an opportunity to use multiple search strategies. The third tube also made looking into the tubes slightly more expensive. Results from Experiment 2 indicated that when chimps had not seen the tubes directly being baited they looked into the tubes more often and chimps performed better when they had looked into the tubes. There were two additional findings from Experiment 2. The first was that some subjects never looked into the tubes, and second there was excessive searching of the tubes after having found the food reward. Call and Carpenter suggested that the aforementioned effects could have been due to the addition of a third tube. For their third, and final, experiment Call and Carpenter implemented a slightly different version of Experiment 2 (they added the delay back to the procedure) for 12 children that were approximately 2 years old. Call and Carpenter hypothesized that children would perform better in trials where they had seen the baiting process and in trials with a delay. Results from the children were found to be very similar to the results of the apes. Children chose the baited tube more often than would be expected by chance for seen, unseen immediate, and unseen delayed conditions. Children also looked more and performed better by looking in the tubes on trials in which they had not seen the baiting process and looking behavior increased slightly when a delay was added after baiting. The main differences between the results of the children and the chimps or orangutans were that the children stopped searching after finding a reward and children were less likely to stop searching upon finding empty tubes. Based upon their findings, Call and Carpenter concluded that their participants were able to seek out additional information when they knew they did not know where a reward was located.

In a similar study to Call and Carpenter's (2001), Hampton, Zivin, and Murray (2004) investigated whether rhesus monkeys could discriminate between a state of knowing and not knowing and, in the case of not knowing, if it was possible for monkeys to collect the necessary information before solving a task. Hampton et al. used a set of opaque tubes where food rewards could be hidden to test monkeys' ability to collect more information. Specifically, Hampton et al. hypothesized that if a monkey is aware of its own memory then it should be able to discriminate between whether it knows (or does not know) a food reward's location. Additionally, if monkeys are aware of the location of the food then they should select the tube that contains the food without looking. Hampton et al. tested nine rhesus monkeys by using a tray that had four tubes attached at one end that allowed monkeys to lift each tube in order to look for the food reward hidden inside. Another feature of the tray was that it could be adjusted to five different heights so that looking could be made easier or more difficult for each monkey. Hampton et al. first familiarized the monkeys with the task by using clear tubes so that monkeys could see that a food reward had been “hidden” in one of them. After monkeys reached a performance criterion of one hundred percent they moved to a similar phase which used opaque tubes. Monkeys also received training in which they learned how to look into tubes in order to find the hidden food. After completing training, the monkeys progressed to the main task in which they observed a mixture of seen trials (where a transparent screen was raised and they could see the tubes being baited) and unseen trials (where a translucent screen was raised to block their view of the baiting process). Next, the screen was raised, the monkey was allowed to select a tube, and the selection was recorded by live coding.

Results from the familiarization process demonstrated that monkeys learned how to choose and pull tubes, learned to attend to the baiting process, successfully transitioned from clear to opaque tubes, and learned to look into the tubes with the tray positioned at different heights. Analysis of the main task revealed that seven of nine monkeys looked into the tubes significantly more on trials in which they had not seen the baiting process. All of the monkeys, except for one, selected a tube in each trial. Monkeys accurately selected the baited tube on trials in which they had seen the tubes being baited. Monkeys were also found to perform significantly better than chance if they first looked into the tubes (on trials in which the baiting process had been obscured) and then made their selection. Hampton et al.’s main findings revealed that monkeys were able to discriminate between knowing and not knowing the location of the food and that the majority of the monkeys were able to look into the tubes when they did not know the location of the food (i.e. seeking out more information).

Basile, Hampton, Suomi, and Murray (2009) have recently performed a series of similar experiments on memory awareness in tufted capuchin monkeys. The goal of their study was to determine whether a species of New World monkey would show evidence of metamemory. A successful demonstration of metamemory would require that monkeys meet three behavioral criteria: looking more often on unseen trials than on seen trials, whether accuracy improved as a function of looking on unseen trials, and whether monkeys immediately stopped searching for food after finding it. Basile et al. used the aforementioned criteria to test capuchins' memory awareness in three successive experiments by using a method comparable to Hampton et al.'s (2004). The same apparatus described in Hampton et al.'s (2004) study was used in all of the following experiments. In Experiment 1, five capuchins were allowed to adapt to the testing procedures in a series of three phases where they learned how to select and pull transparent tubes, made the transition to opaque tubes, and finally learned how to search for hidden food. During the main task, Basile et al. titrated the height of the tray containing the tubes, by using the total number of looks, in order to control for potential ceiling and floor effects. The monkeys were exposed to a pseudorandom mixture of seen and unseen trials in which they had to choose a tube to obtain food. Despite the incremental procedure, capuchin monkeys failed to meet all three behavioral criteria in Experiment 1 due to a failure to attend to baiting. Therefore, in Experiment 2 Basile et al. encouraged monkeys to attend to baiting by mixing sessions in which a short visual barrier was used to block the monkeys' view of baiting with test trials. Basile et al. hypothesized that mixing sessions would increase the number of first-looks into the tubes and decrease the number of searches on seen trials. In addition to the visual barrier, the tray containing the tubes remained in the lowest position throughout the second experiment forcing the monkeys to attend to tube being baited. Three of five capuchin monkeys met Basile et al.'s proposed criteria for memory awareness, which suggests that capuchins can adaptively seek information when they do not know the answer to a problem. In Experiment 3, Basile et al. increased the cost of looking behavior by adding hinged metal flaps to the front of each tube. The intent of adding the flaps was to further decrease the number of unnecessary looks and to obtain a larger difference between the number of looks on seen and unseen trials (i.e., looking was expected to increase on unseen trials after adding the flaps). Results from the third experiment indicated that adding the flaps did decrease the number of unnecessary looks however it also decreased the number of necessary searches. Additionally, none of the monkeys searched more on unseen trials than on seen trials. Basile et al. concluded that they indeed found evidence for memory awareness in three capuchins based upon their three behavioral criteria: seeking out more information when needed, an increase in successful performance as a result of searching, and termination of searching after finding food.

Kornell, Son, and Terrace (2007) emphasized the importance of transfer tests for both metacognitive monitoring and control processes. Kornell et al. performed two experiments with transfer tests to determine the influence of potential associative mechanisms. The goal of the first experiment was to determine whether a monkey was able to accurately monitor its performance from one perceptual domain to another. Both the primary and the transfer tasks utilized a token economy and consisted of the presentation of a perceptual task that was immediately followed by a confidence-rating task. For the primary task, monkeys were trained using a circle-size discrimination task in which they were required to select the circle with either the larger or smaller diameter. The monkeys were required to give their answer to the perceptual task before proceeding to the confidence-rating task, where they could then rate their confidence as either *high* or *low*. Monkeys were trained on a serial working memory task for the transfer test, and were presented with a specific sequence of trial-unique photographs they would later need to place in correct order. Results for the primary (circle-size discrimination) task indicated that monkeys responded metacognitively, meaning that they rated their confidence as higher after correct responses and lower after incorrect responses. Additionally, the same pattern of responding was found for the serial working memory transfer task. Reaction time duration for correct and incorrect responses on both tasks was determined to have no influence upon metacognitive responding. Experiment 1 indicated, for the first time, that monkeys can transfer their metacognitive abilities from a perceptual task to a working memory task.

In their second experiment, Kornell et al. (2007) were interested in the control component of metacognition, specifically whether monkeys had control over their own knowledge. For this experiment, monkeys were required to learn “lists” of four novel photographs in a simultaneous-chaining paradigm (all items are displayed simultaneously throughout each trial) and were given the opportunity to request “hints” that would reveal the next photograph in the sequence. Essentially, the authors wanted to determine whether monkeys would learn to request hints when learning new lists and if hint-seeking behavior decreased as accuracy for the new lists increased. Results revealed that hint-seeking behavior for new item lists decreased across sessions and, that requests for hints for familiar lists were less frequent when compared to the number of hint requests for new lists. Based upon the aforementioned findings, it appears that monkeys do have some control over their own knowledge, meaning they seek out hints when they lack the necessary information to complete a task.

Current Opinions in Comparative Metacognition

Presently, there is a debate about the methodology used in studies on comparative metacognition and the manner in which metacognitive evidence has been interpreted. Essentially, there are two different positions: proponents who believe that metacognition data have been misinterpreted and proponents who believe that metacognitive data should meet a specified criteria, or standard, for metacognition. Smith, Beran, Couchman, Coutinho, and Boomer (2009) have recently written a review in which they discuss this methodological debate and in particular, questions surrounding the roles of associative mechanisms (e.g., stimulus cues and reinforcement contingencies). In Smith et al.'s opinion, the debate is not centered on associative explanations but rather on how comparative psychologists have misinterpreted findings from certain metacognition paradigms. For example, Smith et al. believe paradigms that use more abstract stimuli (e.g., psychophysical Same-Different tasks) are less grounded in associative explanations than those that focus on stimulus-based qualities (e.g., pitch discrimination, noise-duration discriminations). Furthermore, Smith et al. argue that abstract judgments about the relationship between two stimuli in a Same-Different task are more cognitively sophisticated than judgments about the absolute properties inherent to task stimuli. Lastly, Smith et al. claim that the uncertainty response cannot be due to associative mechanisms in paradigms that use abstract stimuli because they are relational in nature and, the uncertainty response can only be due to the indeterminacy of a relationship between stimuli.

Alternatively, Crystal and Foote (2009) believe that while studies of metacognition are well suited for testing stimulus-response hypotheses they are insufficient for testing the stimulus independent-hypothesis described by the response threshold in Smith et al.'s model. Stimulus-response hypotheses have historically been the alternative explanations proposed for metacognitive data. For example, a stimulus-response hypothesis would explain metacognitive data by stating that an animal has learned to perform a specific response in the presence of a specific stimulus (e.g., pressing the left lever to indicate short stimuli in the presence of a short 2 sec stimulus). However, a new challenge has arisen from Smith et al.'s recent response threshold model (although Smith et al. did not specify that their response threshold was stimulus-independent), the stimulus-independent hypothesis. A stimulus-independent hypothesis, for example, would explain metacognition data by stating that the reinforcement of a specific response is adequate for producing the same response at a low frequency in the future (i.e., an operant rate). Perhaps the most important feature of the stimulus-independent hypothesis is that a response that has been reinforced in the past carries with it a constant attractiveness that is not dependent upon the stimulus context in which it was presented.

Transfer tests are a simple and effective means for evaluating the impact of stimulus-response hypotheses. Traditionally, transfer tests have been used to evaluate stimulus-response explanations for metacognition data. Furthermore, transfer tests have been easy to implement because they are procedurally the same as the original task but they use novel stimuli. However, the problem with using a transfer test in comparative metacognition is that it cannot rule out alternative explanations that are stimulus-independent. This is because a low tendency to use a specific response (that has previously established by reinforcement history) carries forward to the transfer task even though novel stimuli are used. In other words, the animal is able to use its reinforcement history when responding to novel stimuli. Therefore, the stimulus-independent nature of the response threshold in Smith et al.'s model suggest that transfer tests are no longer adequate for successful demonstrations of metacognition.

Hampton (2009) has also contributed unique ideas regarding the current methodological debate. Hampton has proposed that there are private mechanisms and public mechanisms that are responsible for current data on animal metacognition. Private mechanisms are defined as introspective (i.e., having access to private cognitive states) while public mechanisms are defined as being more dependent on associative mechanisms about observable stimuli (e.g., perceived difficulty, reinforcement history). In addition, Hampton argues that the observed data from experiments on animal metacognition can be mostly, if not fully, explained by public mechanisms. Four classes of stimulus control mechanisms best exemplify public mechanisms: environmental cue associations, behavioral cue associations, response competition, and introspection. Environmental cue associations are discriminative cues that are associated with test difficulty. Behavioral cue associations are cues that control the metacognition response by stimulus-response behaviors that are associated with discriminative stimuli. Response competition is the competition between two behavioral responses presented at the same time (e.g., such as the option to take or decline a test). Introspection is the private assessment of information that is solely possessed by a learner or performer. Specifically, it is the failure to eliminate environmental cue associations, behavioral cue associations, and response competition from introspection that has led to the debate about methodology. Hampton argues that the only way to be sure that animals are indeed using introspection as a metacognitive control mechanism is to eliminate the other three sources of stimulus control.