

Dissociating Uncertainty Responses and Reinforcement Signals in the Comparative
Study of Uncertainty Monitoring

J. David Smith

University at Buffalo, State University of New York

Michael J. Beran

Language Research Center, Georgia State University

Joshua S. Redford

University at Buffalo, State University of New York

David A. Washburn

Language Research Center, Georgia State University

July 20, 2005

Running head: DISSOCIATING UNCERTAINTY AND REINFORCEMENT

DISSOCIATING UNCERTAINTY AND REINFORCEMENT 2

Abstract

Though researchers are exploring animals' capacity for monitoring their states of uncertainty, some current paradigms allow the criticism that animals map avoidance responses to error-causing stimuli not because of uncertainty monitored but because of feedback signals and stimulus aversion. We addressed this criticism with an uncertainty-monitoring task in which participants completed blocks of trials with feedback deferred so that they could not associate reinforcement signals to particular stimuli or stimulus-response pairs. Humans and one of two monkeys were able to make cognitive, decisional uncertainty responses that were independent of feedback signals or reinforcement histories. This finding unifies the comparative literature on uncertainty monitoring. The dissociation of performance from reinforcement has theoretical implications, and the deferred-feedback technique has many applications.

*Dissociating Uncertainty Responses from Reinforcement Signals
in the Comparative Study of Uncertainty Monitoring*

Humans have feelings of doubt and confidence, of (not) knowing, of (not) remembering. They often respond appropriately to these feelings by reflecting, rethinking, and seeking help or information. These responses are the referent phenomena for the literatures on uncertainty monitoring and metacognition (Brown, 1991; Brown, Bransford, Ferrara, & Campione, 1983; Dunlosky & Nelson, 1992; Flavell, 1979; Hart, 1965; Koriat, 1993; Metcalfe & Shimamura, 1994; Nelson, 1992; Schwartz, 1994).

Human metacognition is defined as thinking about thinking. The idea in this research area is that the human mind (perhaps uniquely or perhaps along with animal minds) contains a cognitive executive that monitors perception, memory, or problem solving to see how it is faring and how it could be facilitated. This happens, for example, when we realize we have not understood a two-year-old's utterance and try to find its probable meaning. These monitoring functions are assessed in the laboratory by collecting metacognitive judgments (confidence ratings, feelings of knowing, etc.).

Researchers take humans' metacognitive behaviors to indicate important aspects of mind. Metacognition is interpreted to show hierarchical layers of cognitive control—the metacognitive processes monitor and direct lower-level cognitive processes. It is taken to show humans' awareness of their thought processes. Researchers also link metacognitive states to self-awareness (because uncertainty and doubts are subjective and personal) and to declarative consciousness (because humans easily introspect and communicate these states—see Nelson, 1996). Thus, metacognition is believed to be

one of humans' most sophisticated cognitive capacities and one of those most likely to be uniquely human. For this reason, it is an important question whether nonhuman animals (hereafter animals) have similar cognitive capacities. Indeed, the sophistication of the metacognitive capacity could let it rise to the level of tool use and language in its potential to highlight both continuities and discontinuities between human and animal minds.

Accordingly, researchers have inaugurated a new area of comparative inquiry that considers the extent to which animals have a capacity for cognitive monitoring (Hampton, 2001; Inman & Shettleworth, 1999; Shields, Smith, & Washburn, 1997; Shields, Smith, Guttmanova, & Washburn, 2005; Smith & Washburn, 2005; Smith, Schull, et al., 1995; Smith, Shields, Schull, & Washburn, 1997; Smith, Shields, Allendoerfer, & Washburn, 1998; Smith, Shields, & Washburn, 2003a,b; Son & Kornell, 2005). In these studies, researchers used tasks that had a mix of easy and difficult trials. They gave animals an additional response—beyond the primary discrimination responses—that let them decline to complete any trials of their choosing. If animals can monitor cognition accurately, they should recognize difficult trials as problematic and error-risking and choose to decline these trials selectively. In fact, animals do so, producing data patterns in some cognitive-monitoring tasks that are strikingly like those of humans (Shields et al., 1997; Smith et al., 1997; Smith et al., 1998). This additional response has come to be called the uncertainty response and this article concerns its appropriate psychological interpretation.

If animals' uncertainty responses are indicative of uncertainty monitoring and metacognition, then of course they are important behavioral ambassadors. However,

caution is needed in considering this possibility. For one thing, human metacognition presents a high standard of reflective and conscious behavior—animals might have only part of its full suite of capacities. For another thing, there is an established tradition in comparative psychology of interpreting an organism's behavior at the lowest possible psychological level. This tradition defends the principle of parsimony embodied in Morgan's Canon (1906. p. 53). Thus, even given performances by animals that might be about metacognition and uncertainty monitoring, one must consider carefully the alternative possibility that these performances might be explained using lower-level, associative mechanisms.

We can illustrate this tension between theoretical perspectives using the paradigm that initiated the comparative study of uncertainty monitoring (e.g., Smith et al., 1995; Smith et al., 1997). This paradigm combined a psychophysical discrimination (in which one presents a range of trial difficulties, including difficult trials near participants' perceptual thresholds) with an uncertainty response (with which participants can decline to make one of the primary responses). Smith et al. (1997) instantiated this paradigm using a Sparse-Dense discrimination (Figure 1A). Choosing the Box as a response was correct if the Box contained exactly 2,950 illuminated pixels. Choosing the S was correct if the Box contained any fewer pixels. Choosing the Star allowed participants to decline the trial and move into a new, guaranteed-win trial. Difficulty was varied so as to map carefully animals' response patterns in the region surrounding their threshold for distinguishing Sparse from Dense. A variant of this Sparse-Dense task is the focus of the present research.

Figures 1B,C show the monkeys' performance (Smith et al., 1997). Box or Dense responses predominated on Dense trials and the most difficult Sparse trials. S or Sparse responses predominated on sparser trials. The primary discrimination was performed at chance where these two response curves cross, and the Star or uncertainty response was used most in this region of maximum uncertainty. Both monkeys assessed accurately when they were liable to make an error in the primary discrimination and they declined those trials selectively and adaptively. Humans did so, too (Figure 1D), and attributed their uncertainty responses to uncertainty monitored.

However, there is a possible criticism of this study as it bears on showing a capacity for uncertainty monitoring in animals. When animals made one of the two primary discrimination responses—Box or S—on near-threshold trials, they were often wrong and timeouts attended these errors. In contrast, when animals made uncertainty responses on these difficult trials, the consequence could have felt more felicitous to them. Uncertainty responses brought a sure-win trial that was rewarded when completed. So, the uncertainty response delayed reward for a few seconds, but reward did come when the next trial was responded to and there were never timeouts. One could say, then, that the problematic stimuli and the use of the primary responses in those stimulus contexts came to feel slightly aversive to the animals. One could say that the animals were conditioned to use the uncertainty response in the context of difficult, near-threshold stimuli. If so, they would have been using the uncertainty response to avoid aversive stimuli, not to express uncertainty about difficult trials. These are not the same interpretation. These are different interpretations that inhere in different psychological levels.

Moreover, one can see in the description of the study the seed of the confound between aversion-avoidance and uncertainty monitoring. Smith et al. (1997) gave the animals feedback on every trial. So, every consequence was felt by the animal and could be directly credited to the stimulus or to the stimulus-response combination that had caused the error or reward. Smith et al. created a situation in which reinforcement was transparent and in which reinforcement histories and response tendencies could be associatively constructed from the animal's trial-by-trial experience.

Fortunately, one can see in the description of the confound a route to removing it. What if one could deny animals the possibility of directly associating outcomes with particular stimuli or stimulus-response pairs? What if one could make the reinforcement situation opaque to animals, so that they could not construct reinforcement histories or response tendencies from their trial experience? Then, if they still made adaptive uncertainty responses on difficult trials, despite never having experienced a direct outcome regarding those stimuli, one would know they were responding to some higher-level cognitive state because the uncertainty response would have been dissociated and insulated from the processes of reinforcement, association, and conditioning. Creating this dissociation, and asking whether animals can still make adaptive uncertainty responses, was the purpose of the present research.

This approach might contribute to other lines of comparative research. Historically, the contingencies regarding reward and punishment for responses made in different stimulus contexts were made transparent to animals through direct and immediate feedback signals. This methodological approach produced beautiful research, but it betrayed psychologists' originally skeptical view of animal minds, it gave animals an

associative way to perform many tasks, and it supported an associative theoretical framework that in a sense confirmed the original skepticism. As we all know, this approach tended to create interpretative distance between human and animal minds. Accordingly, the broader goal of this article is to explore with humans and animals the effect of taking the opposite methodological tack. By removing immediate reinforcement from the experimental context, we hope to force animals to reckon with more cognitive task approaches, and ourselves to reckon with higher-level descriptive frameworks, and by doing so join other current comparative research programs in reconsidering the interpretative distance between human and animal minds.

Experiment 1: Humans

Method

Participants. Sixty undergraduates from the University at Buffalo, State University of New York (UB) participated in a 55-minute session to fulfill a course requirement. Thirty-three additional participants failed to complete all the session's tasks, either because they had longer response latencies or because they made more errors and endured more penalty timeouts. They produced exactly the same data pattern as that reported, but they were not included in the main statistical analyses for the sake of having a balanced within-participant design.

Apparatus. A computer generated stimuli on a monitor. Participants seated in front of the monitor used keyboard keys to indicate their responses as described below.

Psychophysical tasks. Each participant completed six psychophysical tasks with the following character. The stimulus to be judged was a 200 X 100 unframed box presented in the top center of the computer screen that contained a variable number of

randomly placed lit pixels within its invisible borders. The participant was to discriminate Sparse boxes that had fewer than the task's midpoint of lit pixels from Dense boxes that had more. Tasks were constructed from a base number of pixels called Level 0. For Tasks 1 to 6, respectively, these bases were 900, 1800, 600, 2400, 3000, and 300. The tasks were then run along a 21-step continuum from Level 1 to Level 21, with each level's pixel count given by the formula $\text{Pixels} = \text{round}(\text{Level 0 Pixels} \times 1.018^{\text{Level}})$. So, for example, the density continuum underlying Task 1 (base pixels = 900) ran from 916 pixels (Level 1) to 1309 pixels (Level 21). Ten density levels between 916 and 1076 pixels (Levels 1-10) were designated as Sparse stimuli, and 10 density levels between 1115 and 1309 pixels (Levels 12-21) were designated as Dense stimuli. To be fair, the box holding 1095 pixels (Level 11) was never presented. The same rules and designations applied to the other tasks built on different pixel bases.

Responses. Along with the stimulus box on each trial, participants saw a large S to the bottom left of the box, a large D to the bottom right, and a large ? at the bottom center. These three responses were selected by pressing 3 labeled keyboard keys arranged to duplicate the spatial layout of the response icons on the screen.

Delayed feedback. Participants received no direct feedback for responses in the psychophysical tasks. Instead, a blue bar—shown momentarily after each trial—gradually lengthened in successive eighths of the screen while a block of trials continued. This let participants know where they were in the block and when feedback would come. Then, with the 8-trial block completed, participants received feedback on all the block's trials but *not* in the order in which the trials were completed. First, participants heard a reward whoop for each correct response they had made during the

block. For each of these, two points were added to their total score. Then, for each incorrect response, participants received a 4s timeout accompanied by a low, computer-generated buzzing sound and lost two points from their total score. For both kinds of responses, the participant's total score in points after the outcome was flashed on the screen so they could track their progress. The delayed feedback after 8 trials, rearranged by outcome, instantiated the critical feature of the experiment that participants could not tell which trials they had gotten correct or incorrect. They could not assign credit for different outcomes to different stimulus-response combinations, or build up reinforcement histories of any kind based on trial-to-trial experience in the task.

The uncertainty response did not earn any reward and could not cause any timeout penalty. The use of the uncertainty response was also only indirectly reflected in the feedback that occurred after each trial block. For example, if participants responded Uncertain on 3 of 8 trials, their feedback summary would have 5 events in it, some combination of whoops and buzzes based on the trials they attempted. To give the uncertainty response some attractiveness to humans, and to foster its use a little, we arranged it so that the uncertainty response produced a momentary indication as to the correct response on the declined trial. To this end, the correct response on the declined trial (S or D) went green for a brief time to indicate to the participant what the correct response would have been. Thus, for humans, the uncertainty response gave no reinforcement, but some knowledge. Experiment 3 examines the effect of removing this informational contingency. Following the rewards and punishments of the feedback cycle, the screen cleared and a new block of non-reinforced trials began.

Instructions. Entering the first psychophysical task, participants read the following: “You will see boxes that are SPARSE or DENSE with dots. If the Box is Sparse, type the key labeled S (Sparse). If the box is Dense, type the key labeled D (Dense). When you don't know what kind of box you are seeing, type the key labeled ? to say you are uncertain and to decline to answer the trial. You will receive feedback on your performance every eight trials. You will hear a whoop sound and GAIN 2 POINTS for each correct response. You will hear a 4-second buzz and LOSE TWO POINTS for each incorrect response. Using the ? to say you are uncertain neither gains nor loses points and lets you avoid the long buzzes. After you use the ?, you will also see the answer that would have been correct on that trial flash green. This week, a \$5 prize goes to the four participants who earn points most efficiently. Remember: errors will cost you points. Remember: the 4-second buzz will cost you time to earn points. So, to earn points efficiently, be sure to respond ? when you think you don't know, and respond S or D when you think you do know.” Entering subsequent psychophysical tasks, participants were also told that the task would change because all the boxes would get sparser or denser. They were told to still distinguish relatively sparse from relatively dense boxes or to respond ? if they were uncertain. Experiment 3 considers humans' performance when minimal instructions are given instead.

Training. Training began with only the presentation of Level 1 and Level 21 stimuli intermixed randomly and occurring equally often with the provision that runs of more than four of one trial type were disallowed. After each block of 8 trials was completed, the program assessed the participant's recent percent accuracy in the primary discrimination. If performance was better than 70% correct, the initial trial distribution

was changed. The new distribution invoked a variable called Instep and presented Sparse trials across the range of 1+random (Instep) and Dense trials across the range of 21-random (Instep). If Instep was 2, Trials at Levels 1-2 or 21-20 were possible. If Instep was 10, trials at Levels 1-10 or Levels 21-12 (i.e., all trial levels) were possible. Thus, the Instep variable let us gradually move in toward the difficult center of the continuum. Once this new trial distribution was instituted, the program consulted recent percent correct every two trials, and increased or decreased the value of Instep by 1 when performance was greater than or less than 70% correct, respectively. In this way, participants were eased toward the mature phase of data collection. When Instep reached 10, with performance still beyond 70% correct, a final dual-rectangular trial distribution was instituted. Now, on half of the trials, the density level was sampled randomly from the whole range of the task—that is, Levels 1-21 (but never Level 11). On half the trials, the density level was sampled randomly from the twelve most difficult density levels—Levels 5-10 and 12-17. This dual distribution let us increase the task's difficulty and emphasize uncertainty responding by testing especially intensely the region of the continuum near the discrimination's breakpoint. Only the trials in the final phase of testing were analyzed (at least 100 trials, and on average 104 trials, per participant per task).

Results

Overall performance. The 60 participants completed 37,559 mature trials with an overall Sparse/Dense discrimination accuracy of .76. This left room for higher and lower performance, respectively, on the easier and more difficult stimulus levels. Participants responded uncertain to .26 of all trials across tasks and stimulus levels, leaving room for

lower (more confident) and higher (less confident) levels of uncertainty responding, respectively, on the easy and difficult stimulus levels. Participants showed the phenomenon—observed for 100 years—that humans disprefer to use the uncertainty response. Perhaps a longer timeout would have produced more uncertainty responding. But we needed to have participants complete 6 psychophysical tasks (about 1,000 trials) in an hour. Our 4 s error timeout was a necessary methodological decision, and the .26 level of declined trials its result.

Density bins. Given the structure of the dual-rectangular distribution, we knew that Levels 1-4, 5-10, 12-17, and 18-21, respectively, would receive on average about 2.5, 6.7, 6.7, and 2.5 trials per level per participant per task. Accordingly, we binned the data to roughly equalize trial counts and the stability of the performance estimates within each bin. Levels 1-2 and Levels 3-4, respectively, became Bins 1 and 2. Levels 5 to 10 became Bins 3-8; Levels 12-17 became Bins 9-14. Levels 18-19 and Levels 20-21, respectively, became Bins 15 and 16.

Density-bin effects. We analyzed the data by entering the dependent measures (the proportion correct and the level of uncertainty responding each participant showed in each bin of each task) into the General Linear Model (GLM) procedure of SAS 9.0. Bin and Task were within-subject factors. This sophisticated multivariate procedure let us accommodate the lack of cross-cell balance in this analysis and the missing cells that resulted when participants happened to never see any stimuli in some density bin. There was a significant effect of Bin on proportion correct in the analysis, $F(15, 885) = 60.83$, $p < .05$, $MSE = 0.076$; $\eta^2 = .171$. Figure 2A shows this classic psychophysical data pattern, with excellent performance for the easy density bins at the ends of the

stimulus continua, but with poor performance for the density bins near the discriminations' breakpoints between Bins 8 and 9. This data pattern is the backdrop against which one may see complementary changes in uncertainty responding.

The analysis of uncertainty responding revealed these complementary changes, $F(15, 885) = 19.70$, $p < .05$, $MSE = 0.031$, $\eta^2 = .021$. Figure 2A shows the level of uncertainty responding for each bin. The level of declined trials reflected the primary performance achieved when participants attempted the same kinds of trials at other times. Participants declined easier trials less often and difficult trials more often. Even denied direct signals of reinforcement and assignable feedback, participants still used the uncertainty response adaptively and appropriately. For them, its use was not dependent on trial-by-trial feedback to condition or to occasion particular responses given particular trial contexts. In fact, it even appears that humans were gradually overcoming their overconfidence in the tasks as the session continued and were coming to have a greater appreciation of the uncertainty response's function and value. Their levels of uncertainty responding were .17, .18, .29, .27, .25, and .39 for the six tasks taken in the order in the session in which participants received them, $F(5, 295) = 29.78$, $p < .05$, $MSE = 0.207$, $\eta^2 = .070$.

Experiment 2: Humans

Experiment 2 replicated Experiment 1 with three changes. We eliminated the two sparsest tasks that produced the least stable performance. We challenged participants' ability to transfer the uncertainty response from task to task by having the stimulus continua overlap, so that one task's easiest trials became another task's most uncertain trials. And we presented the six task densities in different orders to different participants

so that changing levels of uncertainty responding through the session could be evaluated with task density controlled.

Method

Participants. Sixty UB undergraduates participated in a 55-minute session to fulfill a course requirement.

Apparatus. The apparatus was identical to that in Experiment 1.

Psychophysical tasks. Participants completed 6 tasks with base pixels (Level 0) of 900, 1095, 1333, 1622, 1973, and 2401. The structure of these tasks was described in Experiment 1. We eliminated the 300- and 600-pixel tasks to create 6 tasks of more homogeneous difficulty and stable character. We also restricted the density range that the 6 tasks spanned from 300-3,000 in Experiment 1 to 900-2401 in Experiment 2. This restriction meant that the base pixel count for each next denser task lay at Level 11 of its sparser neighbor. So, for example, Level 11 (1095 pixels) of the 900-pixel task was Level 0 for the next denser task, and that task's Level 11 (1333) was Level 0 for the next denser task. This overlap challenged participants' ability to reconstrue flexibly specific stimuli as easy or uncertain because these designations changed from task to task. Therefore, it evaluated whether participants carry forward residual associative tendencies from task to task, or whether they can treat each task as decisionally new. All other aspects of the procedure were like those described for Experiment 1.

Results

Overall performance. The 60 participants completed 37,558 mature trials with an overall discrimination accuracy of .78 and an overall rate of uncertainty responding of

.19. Compared to Experiment 1, they performed a little better and declined trials a little less, probably because we had replaced the two hardest and most uncertain tasks.

Density-bin effects. The data were binned and analyzed as before. The analysis revealed a significant effect of Density Bin for both dependent measures, $F(15, 885) = 71.57, p < .05, MSE = 0.072, \eta^2 = .217$ for proportion correct; $F(15, 885) = 19.49, p < .05, MSE = 0.033, \eta^2 = .029$ for uncertainty responding. Figure 2B shows the proportion correct and the proportion of trials declined for each density bin. Again, the idealized psychophysical data pattern in primary performance was answered by complementary changes in uncertainty responding. Once again the adaptive use of the uncertainty response was not dependent on trial-by-trial feedback to condition or to occasion particular responses given particular trial contexts. Here, too, humans gradually came to have a greater appreciation of the uncertainty response's function and value. Their levels of uncertainty responding were .12, .15, .19, .20, .24, and .25 for the six tasks taken in the order in the session in which participants received them, $F(5, 295) = 15.70, p < .05, MSE = 0.153, \eta^2 = .037$.

Individual differences in uncertainty responding. Smith et al. (1997) found strong individual differences among humans in the level of uncertainty responding. These individual differences have empirical importance given the monkeys' performance as described below. They also have theoretical importance because they constrain the appropriate psychological interpretation of uncertainty responses (for humans, and perhaps for monkeys). Indeed, as discussed and referenced below, the lability of the uncertainty response is a 100-year-old theoretical issue in psychophysics. Accordingly, Figure 3 confirms the expected strong individual differences in uncertainty responding in

the present data. It shows the extremes of uncertainty responding across Experiments 1 and 2. One participant (Panel A) used the uncertainty response frequently and appropriately to decline trials surrounding the breakpoint of the Sparse-Dense discrimination. One participant (Panel B) did not, and essentially did not use the uncertainty response at all.

Experiment 3: Humans

There are different approaches to collecting human uncertainty-monitoring data for comparison to animal data. First, one can tell humans to monitor their cognition, to use one response to report uncertainty, and so forth. This provides a face-valid performance standard to which animals' performance can be compared. Experiments 1 and 2 took this approach. Another approach is to give humans minimal task instructions and context so that they face a similar problem to monkeys who can be given no instructions. Experiment 3 took this approach.

Method

Participants. Twenty UB undergraduates participated in a 52-minute session to fulfill a course requirement. The earlier experiments had made it clear that a smaller sample would still provide enough power to show the crucial results.

Apparatus. The apparatus was identical to that in Experiments 1 and 2.

Psychophysical tasks. Each participant completed two psychophysical tasks (acknowledging that humans would need longer learning periods given minimal task orientation). These tasks were designed as follows to mirror closely the tasks that monkeys received in Experiment 4. First, the base pixel counts were 753 and 3012. Second, the tasks operated on a 41-step density continuum. Third, the blue bar

indicating the progress of a block was removed. Fourth, humans performed 4-trial blocks before feedback. Fifth, uncertainty responses produced no information about the correct answer for trials declined—their information value was eliminated. Sixth, the error penalty for humans was doubled to 8 s. Seventh, the humans received virtually no instructions (see below). Eighth, the S, D, and ? icons (that might have stood in humans' minds for Sparse, Dense, and Uncertainty) were replaced with X, O, and Star icons that should have had no extrinsic meaning. Ninth, the humans experienced a transition from single trials with immediate feedback to deferred, blocked feedback that was like the transition the monkeys made during pilot testing (see Experiment 4).

Instructions. Entering the first psychophysical task, participants received these instructions: “Use the keys labeled X, O, or Star to respond to the boxes you will see. Try to use your time and your responses to earn points as fast as you can. Your point score will be shown at the top-left of the screen. Later in the experiment your points will be updated every four trials. This week, a \$5 prize goes to the four people who earn the most points. The experimenter will help you begin the experiment now.”

Procedure. Training in the first task began with the presentation of Level 1 and Level 41 stimuli intermixed randomly. When performance surpassed 85% correct, the dual-rectangular trial distribution was instituted, with 67% of trials drawn from the entire 41-step continuum and 33% of trials drawn from the most difficult stimulus levels (Levels 15-20 and 22-27 on the 41-Step continuum). When performance surpassed 85% correct with this distribution, participants began to receive feedback every 4 trials. Following at least 160 trials of 4-trial blocks (ending with performance still beyond 85%), the base pixel level was switched to 3012 and the presentation of only Level 1 and

Level 41 stimuli was re-instituted. During this change, blocked feedback continued and there was no cue that the base pixel level had changed other than the appearance of the boxes. Again, when performance surpassed 85% correct, the dual-rectangular trial distribution was instituted. Participants completed trials in the mature phase of the second task until their session ended. Only the trials in this phase were analyzed.

Results

Overall performance. The 20 participants completed 11,105 mature trials in the second task with a discrimination accuracy of .83 and a .25 level of uncertainty responding overall.

Density-bin effects. Levels 1-14 became Bins 1-7 (2 levels per bin); Levels 15-20 became Bins 8-13; Levels 22-27 became Bins 14-19; and Levels 28-41 became Bins 20-26 (2 levels per bin). The data were analyzed as before. The analysis revealed a significant effect of Density Bin for both dependent measures, $F(25, 459) = 13.94, p < .05, MSE = 0.034, \eta^2 = .401$ for proportion correct; $F(25, 472) = 7.03, p < .05, MSE = 0.024, \eta^2 = .097$ for the level of uncertainty responding. Figure 2C shows that changes in discrimination accuracy were still answered by complementary changes in uncertainty responding. Humans produced the same results as in Experiments 1 and 2 even when they had to construe the tasks and responses for themselves.

Experiment 4: Monkeys

The monkey participants had experimental histories comprising hundreds of thousands of trials, almost all of which led to immediate reinforcement if correctly completed. A principle methodological problem we faced was to move them from this place to one in which they would complete a trial block absent any feedback, receiving

only reordered feedback following each block. This required training and new procedures because models for this training were lacking. We describe our procedures to make our approach transparent and replicable, to point out helpful methods that did move animals ahead, and to point out methods that failed in our view. In this developing area, describing both kinds of methods is constructive.

Training and Pilot Testing with Monkeys

Participants. Murph and Lou—two 10-year-old male rhesus monkeys (*Macaca mulatta*)—were tested. They had been trained, using procedures described elsewhere (Rumbaugh, Richardson, Washburn, Savage-Rumbaugh, & Hopkins, 1989; Washburn & Rumbaugh, 1992), to respond to computer-graphic stimuli by manipulating a joystick. They had also been tested in prior studies on a variety of computer tasks, including a related Sparse-Dense discrimination in which they rated their low or high “confidence” in their discrimination response before receiving feedback (Shields et al., 2005). (In this task, monkeys completed a Sparse-Dense discrimination, and then made a high-stakes bet to earn a double food reward while risking a double timeout or made a low-stakes bet to earn a single food reward while risking a single timeout.) The monkeys were tested in their home cages with *ad lib* access to the test apparatus, working or resting as they chose during long sessions. They were neither food deprived nor weight reduced for the purposes of testing and they had continuous access to water.

Apparatus. The monkeys were tested using the Language Research Center’s Computerized Test System—LRC-CTS (described in Rumbaugh et al., 1989; Washburn & Rumbaugh, 1992)—comprising a Compaq DeskPro computer, a digital joystick, a color monitor, and a pellet dispenser. Monkeys could manipulate the joystick through

the mesh of their home cages, producing isomorphic movements of a computer-graphic cursor on the screen. Contacting appropriate computer-generated stimuli with the cursor resulted in the delivery of a 94-mg fruit-flavored chow pellet (Bioserve, Frenchtown, NJ) using a Gerbrands 5120 dispenser interfaced to the computer using a relay box and output board (PIO-12 and ERA-01; Keithley Instruments, Cleveland, OH).

Psychophysical tasks. The monkeys were trained with the 200 X 100 unframed pixel box and the S, D, and ? response icons. Potential tasks were constructed from base pixel numbers of 753, 3012, 1509, 1066, and 2155. Each task occupied a 41-step continuum from Level 1 to Level 41 with the pixel counts determined by the formula already given. The wider continua let us create a stronger contrast between easy and difficult/uncertain trials. To challenge associative response tendencies and the appropriate transfer of the uncertainty response, the tasks had overlapped stimulus ranges so that one task's Sparse, Uncertain, or Dense trials could be another task's Dense, Sparse, or Uncertain trials. The Base-1066 task had as its easiest Sparse and Dense trials the most uncertain trials from the Base-753 and Base-1509 tasks, respectively. The Base-2155 task had as its easiest Dense and Sparse trials the most uncertain trials from the Base-3012 and Base-1509 tasks, respectively. If animals did bring their associative history with them into these tasks, they might well transfer wrongly their uncertainty-response tendencies.

Responses. The monkeys used the joystick to move their cursor to touch the S, D, and ? icons placed on the screen as already described.

Within-task training. As with humans, training began with only the most discriminable stimuli presented (Levels 1 and 41) randomly intermixed and equally often

except that runs of 4 or more of one trial type were disallowed. Next, the trial distribution gradually expanded inward toward the midpoint of the discrimination so that difficulty gradually ramped up as already described. In some cases as noted below, the phase of gradually increasing difficulty was omitted and the monkey moved directly into the dual-rectangular phase of mature testing. Finally, a dual-rectangular distribution was used to evaluate participants' mature performance as already described. Only the trials in this last phase of each task were analyzed. We used a criterion of 80% correct as the performance level needed to advance into a new distribution.

Training toward delayed feedback. Our goal was to teach the monkeys to perform a block of trials, only then receiving deferred feedback rearranged by trial outcome. Thus, monkeys were to hear a reward whoop and receive a food pellet for each correct response and a penalty timeout (about 20s) accompanied by a low, computer-generated buzzing sound for each error. As with humans, this delayed feedback was reorganized by outcome so that first all rewards were given successively (about 1 s apart) followed by all the penalty timeouts incurred during the block (about .5 s apart but lasting about 20 s each). This manipulation ensured that monkeys could not tell which trials they had gotten correct or incorrect because the sequences of stimuli and responses during the trial block and the sequences of outcomes during feedback did not correspond either temporally or ordinally. Therefore, monkeys could not attribute particular outcomes to particular stimulus-response combinations or build up reinforcement histories based on their trial-to-trial experience.

The monkeys were trained to tolerate delayed feedback as follows. They first experienced extensive training with every trial given feedback. During this period, the

uncertainty response was introduced and stabilized within the response repertoire. Note that during this period, monkeys might have been able to associate certain stimuli with certain reinforcement contingencies (i.e., to learn associatively). Thus, the crux of the experiment came later when, having trained them to tolerate deferred feedback, we moved them to new psychophysical tasks with different density ranges.

Next we introduced 2-trial blocks with deferred reinforcement. This caused the problem that one animal began to use the uncertainty response somewhat rhythmically, using it appropriately but more often on Trial 1 than Trial 2 of the block (i.e., Lou declined 54.4% and 11.8% of trials in the first and second position, respectively). Given this problem, we instituted instead a mixed-block format that included random mixtures of 1-, 2-, 3-, and 4-trial blocks. We could control the exact frequencies of each block length so that we could gradually emphasize longer block lengths and gradually eliminate shorter block lengths. This approach was more effective, perhaps because it broke up the task's rhythm. In the end, we weaned both animals to the exclusive use of 4-trial blocks.

From this point on, animals made a voluntary trial-initiation response (moving the cursor into a 3 cm circle) to begin each block of trials. We thought that this might create conditions for them to complete blocks when they were most ready to do so. As with humans in Experiment 3, the progression through a trial block was not indicated by any lengthening blue bar. We thought the bar might be distracting to animals and add complexity to a task that already required a sophisticated performance strategy. We kept this possibility in reserve in case the animals needed more explicit bridging toward eventual reinforcement at the end of the block, but it turned out that they did not.

The uncertainty response: training contingencies and consequences. This project's most difficult methodological problem was to get monkeys to sustain the uncertainty response within their trial repertoires even across 4-trial blocks during which its function would be deferred, indirect, and subtle. In fact, the uncertainty response did occasionally threaten to drop from animals' response repertoires. At different times during training, we applied different supports to animals' use of the uncertainty response, but in the later, transfer phases of pilot testing we removed all of these supports.

First, we sometimes forced animals to choose the uncertainty response on a few randomly chosen trials. These forced trials may have let animals experience that response's function and helped them sustain it within the response repertoire.

Second, we sometimes gave animals a deferred reward for uncertainty responses. That is, we made the expected reward value of every block of 4 trials 4 pellets, with only discrimination errors subtracting from this total. This effectively rewarded uncertainty responses. However, by only instituting this contingency after we were at 100% 4-trial blocks, we hoped to soften and veil this reward property so that the uncertainty response remained attractive and stable in the response repertoire without dominating it. This approach seems to have had very little impact on performance.

Third, we sometimes let uncertainty responses dictate that the next "random" trial chosen by the computer would be a Level 1 or Level 41 trial. We thought that this release into an easy subsequent trial might give the uncertainty response a generally positive feel that also sustained it within the response repertoire. However, we believe

this approach did harm. It seemed to tempt both animals to make serial uncertainty responses, so that many Level 1 and Level 41 stimuli were declined.

Finally, as in our previous research (Shields et al., 1997; Smith et al., 1997; Smith et al., 1998), we sometimes mildly discouraged the overuse of the uncertainty response as follows. The program constantly calculated the quantity (Number of Recent Uncertainty Responses - Number of Recent Primary Responses). When that quantity was greater than or equal to 2 (or sometimes 3), the use of the ? froze the screen for $(\text{Recent Uncertainty Responses} - \text{Recent Primaries})^2 \times 100 \text{ ms}$ (or sometimes 50 ms). Thus, overuse of the uncertainty response could slow the task proportionally to the overuse, but the animals had control over their level of uncertainty responding and the amount of slowing they would tolerate. In our view, this is a preferred, gentle approach toward helping animals regulate for themselves their overall use of the uncertainty response. This contingency might govern the overall use of the uncertainty response, but it should have no impact on the distribution of uncertainty responses across the stimulus continuum. Given the levels of uncertainty responding animals showed in these tasks, this contingency was felt by them only 5-10% of the time. This slowing contingency was also removed later without affecting performance.

Like humans in Experiment 3, the monkeys were never given any information about the answer to a trial that they had declined.

Results: Experimental Testing with Monkey Murph

In experimental testing, Murph was presented with three tasks (hereafter called Tasks 1, 2, and 3). For Tasks 1, 2, and 3, respectively, he received the dual-rectangular distribution for 9,085, 4,251 trials, and 6,599 trials. In all three tasks, feedback was

delivered every 4 trials blocked by outcome and the animal performed with no supports to the uncertainty response. The uncertainty response removed that trial's outcome from the summary feedback following the block. It is worth noting that in Task 2 Murph needed only 24 trials of the Level 1-Level 41 trial distribution before reaching criterion and moving directly into the dual-rectangular distribution. In Task 3, he needed only 84 trials before making this transition. These rapid transitions to mature testing are a distinctive feature of this animal's performance and show that, once he knew the task's grammar, he learned new Sparse and Dense perceptual anchors almost instantly and was ready for mature performance.

Performance bins. Levels 1-14, Levels 15-20, Levels 22-27, and Levels 28-41, respectively, became Bins 1-7 (2 levels per bin), Bins 8-13; Bins 14-19; and Bins 20-26 (2 levels per bin).

Murph Task 1. Figure 4A shows Murph's performance over 9,085 trials. Though forced to learn the discrimination with no trial-by-trial feedback, he did so. (Shortly we explain why he showed below-chance performance in some bins.) Moreover, he basically understood where the difficult region of the stimulus continuum lay and he organized an adaptive region of uncertainty responding there. This is a crucial result. Even denied direct signals of reinforcement and assignable feedback, this monkey (just as humans) used the uncertainty response adaptively and appropriately. For him, the use of that response was not dependent on trial-by-trial feedback to condition or to occasion particular responses given particular trial contexts.

Murph Task 2. Figure 4B shows Murph's performance over 4,251 trials. Once again he learned the basic discrimination without trial-by-trial feedback and organized an

adaptive region of uncertainty responding spanning the difficult region of the stimulus continuum. In fact, Figure 4C shows that the animal developed a plausible task construal within the first 400 trials. (This graph is noisier because—given only 400 trials distributed across 26 performance bins—the estimates of the 2 response percentages for each bin are less stable.) Thus, in Task 2 also, even denied assignable feedback, this monkey (like humans) still made uncertainty responses appropriately, showing that those responses can be produced at a cognitive or decisional level.

Murph Task 3. Figure 5A shows Murph's performance over 6,599 trials. All aspects of the data patterns seen in Tasks 1 and 2 were seen here, strengthening the conclusion that the animal was making cognitive or decisional uncertainty responses.

Figure 5A shows another crucial result. At Bins 5-6, the animal was 95% correct and appropriately declined trials at a low rate. But at Bins 12-13, the animal performed far worse (24% correct) and even far below chance and yet declined trials at the same rate. If he had tracked the associative history of these trials, and declined trials on this basis, these trials surely would have been aversive and he would have tried to decline them all. He did not, confirming that he lacked access to those associative cues and was not responding to them. Making this important point more generally, Figure 5B plots the proportion of trials declined for each density bin against the proportion of trials correct for that bin. Across nearly all the range of percent correct in the task there was no relationship between trial failure and trials declined as the associative hypothesis would predict. The animal's uncertainty responses did not follow the task's reinforcement and associative patterning.

In contrast to an associative strategy, Figure 5C lets one see Murph's decisional strategy. In his task construal, the Sparse-Dense crossover lay at Bin 9, not at Bin 13-14 as it actually does. At Bin 9, his use of the two discrimination responses equalized and his uncertainty response peaked (declining symmetrically to each side). This task construal explains why Murph performed below chance on Bins 10-13. He thought those trials were Dense, even though by the task's objective structure and reinforcement patterning they were Sparse. This task construal also explains why Murph felt equally confident about Bin 5 and Bin 13 stimuli. They were equally far from his decisional breakpoint, and equally clear as Sparses and Denses, respectively, even though the Bin 13 stimuli were a reinforcement nightmare for him and the Bin 5 and Bin 13 stimuli were associative opposites. Making this important point more generally, Figure 5D shows that there was a strong relationship between the proportion of trials declined in a bin and the decisional distance of the bin from the animal's decisional breakpoint (Bin 9 = 0; Bins 8 and 10 = 1; Bins 7 and 11 = 2, etc.). This correlation is $r = .92$, with the relation accounting for 85% of the variance in Murph's uncertainty responding.

Thus one sees that the delayed and blocked feedback fundamentally denied the animal access to the reinforcement framework of the task. Absent that, he built his own decisional framework for it. He did so well but imperfectly. By partially succeeding, he showed that an animal can organize a decisional framework for a task—with two response criteria defining three response regions—absent trial-by-trial signals of reinforcement. By partially failing, he showed something more important. For by doing so, his decisional framework for the task became dissociated from the task's

reinforcement patterning, and he showed clearly that he was performing according to the former. This may be the first dissociation of its kind reported in the literature.

We confirmed our understanding of Murph's performance through formal modeling grounded in Signal Detection Theory (SDT—e.g., MacMillan & Creelman, 1991). This class of model assumes that objective stimuli will be perceived with perceptual error, so that the same stimulus level on different occasions will be perceived differently. Thus, one parameter of the SDT model was the size of this perceptual error in steps along the 41-step continuum. The SDT model then also assumes that the animal's decisional framework for the task consists of two decision criteria, one separating the Sparse and Uncertain response regions and one separating the Uncertainty and Dense response regions. The placements of these two criterion lines along the 41-step continuum were also parameters in the SDT model. The SDT model captured Murph's performance well, with a fit index that is competitive with other formal cognitive models. Figure 6 shows the performance pattern of the best-fitting parameter configuration (compare Figure 5C). With best-fitting criteria of Step 13 (Bin 7) for the Sparse-Uncertain decision and Step 19 (Bin 12) for the Uncertain-Dense decision, the modeling confirmed that Murph's decisional framework for the task centered at Step 16 (Bin 9) as already described. Some details of these modeling and fitting procedures are described in the Appendix.

With his performance in these three psychophysical tasks, Murph met every cognitive requirement of the experiment. He organized adaptive uncertainty regions along the stimulus continua even though denied trial-by-trial feedback and a way to build up reinforcement histories or associative tendencies regarding stimuli or stimulus-response pairs. His uncertainty responding followed the patterning indicated by his

decisional construal of the task and not the associative patterning of the task that was opaque to him. Thus, the experiment showed for the first time a dissociation between the cues of reinforcement/association and the judgments from decision making under uncertainty.

Results: Experimental Testing with Monkey Lou

Lou progressed less far into the experiment, and his data provide another kind of information--about the psychological and behavioral fragility of the uncertainty response. In experimental testing, Lou received two tasks (hereafter called Tasks 1 and 2). For Tasks 1 and 2, respectively, he received the dual-rectangular distribution for 5,160 and 7,711 trials. In both tasks, feedback was delivered every 4 trials blocked by outcome and the animal performed without any supports to the uncertainty response.

Lou Task 1. Figure 7A shows Lou's performance on 5,160 trials of the dual-rectangular distribution. Though Lou was forced to learn the discrimination with no trial-by-trial feedback, he did so. However, Lou basically did not respond Uncertain. Any chance that he might have reserved the uncertainty response more for the most difficult trials—he declined 2.9% of trials in his easiest bins, and 4.0% of trials in his difficult bins—was overwhelmed by his tendency not to use the uncertainty response.

Lou Task 2. Figure 7B shows Lou's performance over the 7,711 trials in the rectangular distribution. Though forced to learn the discrimination with no trial-by-trial feedback, he did so. However, again Lou basically did not decline trials.

We applied the SDT modeling framework to Lou's performance in order to ask how much his behavioral strategy of not declining trials cost him in his rate of reinforcement. For this purpose, we created a more detailed virtual task environment that incorporated

all of the task's response times and penalty times. In this way, we could estimate total rewards, total punishment times, total time on task, and so forth, and thus finally the rewards per minute that would be earned by a particular decisional strategy in the task. This optimality analysis showed that by barely responding uncertain, Lou reduced his rate of reinforcement by 24% compared to the optimal strategy that used the uncertainty response more to decline adaptively the most difficult trials. Some details of these modeling and fitting procedures are described in the Appendix.

Thus Lou, even denied trial-by-trial feedback, successfully learned new psychophysical tasks and organized appropriate Sparse and Dense response regions. However, his uncertainty response in the deferred-reinforcement context proved fragile, and he did not sustain its use as Murph did. Because individual differences in uncertainty responding have long raised theoretical issues within perceptual psychology and psychophysics, we return to this aspect of Lou's performance below.

General Discussion

Humans and monkeys learned new psychophysical tasks without the trial-by-trial feedback that could support associative learning about stimulus-response pairs. Humans and one monkey constructed adaptive uncertainty-response regions that let them decline trials in difficult regions of new stimulus continua. They did this even though the patterning of reinforcement by density level could not be tracked because reinforcement was deferred and rearranged by outcome. They did this even though the function of the uncertainty response was more subtle than in any previous comparative study of uncertainty monitoring. All in all, the monkey (Murph) showed a distinctive decision-making performance by an animal because he defined for himself the decision

criteria to separate Sparse, Uncertain, and Dense response regions. One understands how reinforcement signals could entrain criterion lines. But criterion setting absent reinforcement is a more sophisticated process that must be guided by the animal's own judgments about which trials can or cannot safely be called Sparse or Dense. The present results make a strong contribution to the comparative decision-making literature because they may show the first example of criterion setting by an animal who was denied direct reinforcement.

More generally, we believe that the dissociative technique used here—to divorce learning and performance from reinforcement and associative cues—could contribute to other areas of comparative research. We provided an existence proof here that animals will learn tasks even denied any trial-by-trial feedback. This approach potentially opens to any research area a new window on the cognitive construal that animals make about tasks independent of the tasks' contingencies. Theoretical progress could result from sometimes understanding animals' performances in this more cognitive/decisional, and less associative/conditioned, way.

In fact, we were able to demonstrate here, because of imperfections in Murph's cognitive construal, that his responses followed his decisional framework and definitely not the reinforcement structure of the task. This dissociation contributes to the comparative literature on uncertainty monitoring because it shows that animals' uncertainty responses can be generated at a decisional level that is dissociable from the associative level. These results clarify in several ways the theoretical situation regarding the psychological organization of uncertainty responses.

First, that monkeys can make cognitive, decisional uncertainty responses fills an empirical gap. Some studies of monkeys' uncertainty monitoring had been criticized for their possible associative basis (i.e., animals might have been escaping aversive middle stimuli with known reinforcement histories that included errors and timeouts). But the present findings show that this interpretation is not sufficient to describe what animals can do in these tasks. A higher-level interpretation is required that allows for decision-making processes that are independent from the task's reinforcement structure.

Second, our findings narrow a task disconnect that had arisen in this field. That is, some had attributed the psychophysical (e.g., Sparse-Dense) uncertainty-monitoring findings to associative phenomena though other findings indicated that this description was insufficient. For example, Hampton (2001) and Smith et al. (1998) showed that monkeys could recruit adaptive uncertainty responses when facing difficult probes of their memory. Here, low-level interpretations cannot be defended—monkeys had to be responding to a more cognitively derived cue like trace familiarity. The present results are helpful because they eliminate the need to divide tasks along associative or cognitive-monitoring lines. Instead, all the referent tasks can be given the same interpretation in terms of decision making under uncertainty, uncertainty-response regions as part of an overall task construal, and so forth. That a common information-processing framework encompasses the existing tasks supports theoretical progress.

Third, our findings narrow a species disconnect in this field. Humans attribute their uncertainty responses in psychophysical tasks to uncertainty monitoring. Researchers have comfortably applied these descriptors to their performance (see below). However, some animal researchers have pointed out that the identical graphs shown in Figure 1

could represent analogous performances achieved by metacognitive mechanisms (humans) and associative mechanisms (animals). But the present results, by showing that the associative mechanism is insufficient, and by definitively moving the uncertainty-monitoring performances of animals onto the cognitive or decisional plane, show that this interpretative species divide may not be necessary. Uncertainty-monitoring performances by humans and animals are probably partially homologous in their mechanism, though, as we discuss presently, not perfectly so.

Thus, this article closes important circles in comparative research on uncertainty monitoring, simplifies the findings theoretically, and grants the field a constructive interpretative parsimony. This parsimony does not require the assumption that animals or humans show full-blown metacognition or consciousness in these tasks. But it means that one can credit the performance similarities across tasks and species, and accept that there may be information-processing homologies across these. On doing so, one sees that the state of this field is intriguing. Research shows that animals' uncertainty responses are not low-level reactions to specific stimuli (Hampton, 2001; Smith et al., 1998). The present findings show that animals' uncertainty responses need not be reinforcement based. What theoretical framework is appropriate when animals decline difficult trials for cognitive and decisional reasons that are dissociable from the traditional associative dimensions of stimuli and reinforcement? The theoretical framework of uncertainty monitoring fits onto the existing evidence well.

Earlier we discussed the difficult criterion-setting problem that an animal faces when denied any trial-by-trial reinforcement. Notice that an uncertainty-monitoring interpretation of the uncertainty response provides a natural solution to this problem.

The animal would only have to respond Sparse or Dense when he thought he knew the answer, and respond Uncertain when he thought the trial was too doubtful. Three adaptive response regions would follow in just the pattern Murph showed. Clearly animals experience perceptual close calls of this kind in nature—Griffin (2003) gave lovely examples (also Smith et al., 2003b, p. 367). So it is plausible that animals would have benefited from having a generalized uncertainty-monitoring capacity that let them cope with difficult perceptual problems.

Lou's performance here raises complementary issues. This animal barely responded Uncertain despite serving 2,492 penalty timeouts—lasting 49,840 seconds or almost 14 hours—that reduced his rewards per minute by 24%. Many of these timeouts were avoidable using the uncertainty response. If this animal could have kept that response in his repertory—even by tracking the reinforcement history of the density levels and escaping the middle, aversive stimuli—he should have done so. In previous research, we have shown that animals can and often do use the uncertainty response to optimize their performance (e.g., Smith et al., 1998, p. 236, Figure 5). That Lou did not do that here suggests again that the associative structure of the task was invisible to him. We point out that in Shields et al. (2005), Lou (but not Murph) showed a related poor ability to use low-and high-risk responses that for humans amount to low- and high-confidence ratings. Yet in Beran, Smith, Redford, & Washburn (2005), Lou was able to use an uncertainty response appropriately to decline difficult numerical discriminations, though he did so less often than Murph. We are interested in Lou's occasional reluctance to use confidence-based or uncertainty-based responses, but we do not have a complete explanation.

Notice that humans and both animals used the primary responses (Dense and Sparse) similarly. The uncertainty response alone was behaviorally fragile, so that it fell out of the response repertoire of one animal. Smith et al. (1997, Figure 4, p. 87) observed this same fragility—only one of two monkeys used the uncertainty response appropriately to decline the difficult trials in a related Sparse-Dense discrimination. Smith et al. observed the same individual difference across humans in their studies (Figure 3, p. 86). The same individual difference among humans was present here (Figure 3). Hampton (2001) also found in his Experiment 3 that only one of two monkeys showed the uncertainty-monitoring data pattern.

For a century, individual differences like these (among humans) suggested to psychophysicists that the uncertainty response is behaviorally fragile because it is psychologically different from the primary discrimination responses. In fact, some wanted to disallow uncertainty responses from psychophysical tasks because of this special psychological status (Fernberger, 1914; Woodworth, 1938, pp. 419-427). That response—more than the primary discrimination responses—seemed to be affected by instructional set (Brown, 1910; Fernberger, 1914, 1930; Woodworth, 1938), linked to participants' temperamental tentativeness (Angell, 1907; Fernberger, 1930; Thomson, 1920), less perceptual (Boring, 1920; George, 1917), and more reflective and cognitive (Angell, 1907; Fernberger, 1930; George, 1917). So the view arose that uncertainty responses were meta- to the primary discrimination and were a comment on the participant's failure to assign a stimulus to one of the primary input classes. This explains the behavioral lability of the uncertainty response for humans—some would naturally use it more or less, whereas all would use the Sparse and Dense responses

similarly because without them the task is incoherent. This interpretation makes the behavioral fragility of monkeys' uncertainty response interesting—it might be meta- to the primary discrimination for them as well. Notice that monkeys' (and humans') uncertainty responses could be meta- to their discriminations in some ways without being meta- in every possible way.

Our ongoing research may help explain the behavioral fragility that attends uncertainty responses and explain their different psychological role. Redford, Chapman, & Smith (2005) added to a Sparse-Dense discrimination a concurrent Stroop task that occupied humans' working memory during the Sparse-Uncertain-Dense judgment. They found that uncertainty responding only, not Sparse or Dense responding, was felled by the cognitive load. Then they added the Stroop task to a Sparse-Middle-Dense discrimination. Cognitive load had no impact on Middle responses, even though these responses were made to the same trial levels that formerly earned uncertainty responses. These experiments support the idea that humans' uncertainty response is less grounded perceptually and more executive in its psychological character. One could ask the same question about monkeys' uncertainty response.

This idea, and the individual difference between the monkeys, raise the possibility that the monitoring capacity required in the present tasks is near the limit of what monkeys can achieve. And for this reason, too, the present research could foster empirical and theoretical progress in this area of comparative research. For if animals are on the edge of not succeeding in our tasks, it places us in a strong position to ask about the similarities and differences between human and animal performance, and about aspects of the whole suite of metacognitive abilities that animals may have or

may lack. Of course animals perform these tasks less verbally than humans. Perhaps animals perform them less consciously or more implicitly or with less self-awareness than humans. For a variety of reasons that are grounded in both cognitive psychology and neurobiology, monkeys may have greater difficulty recruiting or sustaining the cognitive effort to make sophisticated decisional or executive responses like uncertainty responses (Smith et al., 2004). This is why we and many commentators on Smith et al. (2003—see commentaries by Call, Carruthers, Flavell, King, Metcalfe, Proust, and others) are so interested in the possibility that animals could have some phylogenetic precursors of human metacognition or some of the suite of metacognitive capacities but not all. We have even called what monkeys might have “junior metacognition”—granting that metacognition is a multifaceted thing and that having it is not an all-or-none thing. But whatever the final set of answers about what aspects of metacognition monkeys do or do not have, the best place from which to look for similarities, differences, gradations, and facets will be the level interpretative playing field and the common information-processing description for monkeys and humans that the present research recommends and demonstrates is justified.

References

- Angell, F. (1907). On judgments of "like" in discrimination experiments. *American Journal of Psychology*, 18, 253.
- Beran, M. J., Smith, J. D., Redford, J. S., & Washburn, D. A. (2005). Rhesus macaques (*Macaca mulatta*) monitor uncertainty during numerosity judgments. Manuscript under review.
- Boring, E. G. (1920). The control of attitude in psychophysical experiments. *Psychological Review*, 27, 440-452.
- Brown, A. S. (1991). A review of the tip-of-the-tongue experience. *Psychological Bulletin*, 109, 204-223.
- Brown, A. L., Bransford, J. D., Ferrara, R. A., & Campione, J. C. (1983). Learning, remembering, and understanding. In J. H. Flavell and E. M. Markman (Eds.), *Handbook of child psychology* (Vol. 3, pp. 77-164). New York: Wiley.
- Brown, W. (1910). The judgment of difference. *University of California Publications in Psychology*, 1, 1-71. Berkeley CA: The University Press.
- Call, J. (2003). On linking comparative metacognition and theory of mind. *The Behavioral and Brain Sciences*, 26, 341-342.
- Carruthers, P. (2003). Monitoring without metacognition. *The Behavioral and Brain Sciences*, 26, 342-343.
- Dunlosky, J., & Nelson, T. O. (1992). Importance of the kind of cue for judgments of learning (JOL) and the delayed JOL effect. *Memory & Cognition*, 20, 374-380.
- Fernberger, S. W. (1914). The effect of the attitude of the subject upon the measure of sensitivity. *American Journal of Psychology*, 25, 538-543.

- Fernberger, S. W. (1930). The use of equality judgments in psychophysical procedures. *Psychological Review*, 37, 107-112.
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, 34, 906-911.
- Flavell, J. H. (2003). Varieties of uncertainty monitoring. *The Behavioral and Brain Sciences*, 26, 344.
- George, S. S. (1917). Attitude in relation to the psychophysical judgment. *American Journal of Psychology*, 28, 1-38.
- Griffin, D. R. (2003). Significant uncertainty is common in nature. *The Behavioral and Brain Sciences*, 26, 346.
- Hampton, R. R. (2001). Rhesus monkeys know when they remember, *Proceedings of the National Academy of Sciences*, 98, 9, 5359-5362.
- Hart, J. T. (1965). Memory and the feeling-of-knowing experiments. *Journal of Educational Psychology*, 57, 347-349.
- 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.
- King, J. E. (2003). Parsimonious explanations and wider evolutionary consequences. *The Behavioral and Brain Sciences*, 26, 347-348.
- Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, 100, 609-639.
- MacMillan, N. A., & Creelman, C. D. (1991). *Detection theory: a user's guide*. Cambridge, UK: Cambridge University Press.

- Metcalfe, J. (2003). Drawing the line on metacognition. *The Behavioral and Brain Sciences*, 26, 350-351.
- Metcalfe, J., & Shimamura, A. (1994). *Metacognition: Knowing about knowing*. Cambridge, MA: Bradford Books.
- Morgan, C. L. (1906). *An introduction to comparative psychology*. London: Walter Scott.
- Nelson, T. O. (Ed.) (1992). *Metacognition: Core readings*. Toronto: Allyn and Bacon.
- Nelson, T. O. (1996). Consciousness and metacognition. *American Psychologist*, 51, 102-116.
- Proust, J. (2003). Does metacognition necessarily involve metarepresentation? *The Behavioral and Brain Sciences*, 26, 352.
- Redford, J., Chapman, P., & Smith, J. D. (2005). Reducing uncertainty about the cognitive processes associated with uncertainty monitoring. Manuscript in preparation.
- 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.
- Schwartz, B. L. (1994). Sources of information in metamemory: Judgments of learning and feelings of knowing. *Psychonomic Bulletin and Review*, 1, 357-375.
- Shields, W. E., Smith, J. D., & Washburn, D. A. (1997). Uncertain responses by humans and rhesus monkeys (*Macaca mulatta*) in a psychophysical same-different task. *Journal of Experimental Psychology: General*, 126, 147-164.

- Shields, W. E., Smith, J. D., Guttmanova, K., & Washburn, D. A. (2005). Confidence judgments by humans and rhesus monkeys. *Journal of General Psychology, 132*, 165-186.
- Smith, J. D., Minda, J. P., & Washburn, D. A. (2004). Category learning in rhesus monkeys: A study of the Shepard, Hovland, and Jenkins tasks. *Journal of Experimental Psychology: General, 133*, 398-414.
- Smith, J. D., Schull, J., Strote, J., McGee, K., Egnor, R., & Erb, L. (1995). The uncertain response in the bottlenosed dolphin (*Tursiops truncatus*). *Journal of Experimental Psychology: General, 124*, 391-408.
- Smith, J. D., Shields, W. E., Schull, J., & Washburn, D. A. (1997). The uncertain response in humans and animals. *Cognition, 62*, 75-97.
- Smith, J. D., Shields, W. E., Allendoerfer, K. R., and Washburn, W. A. (1998). Memory monitoring by animals and humans. *Journal of Experimental Psychology: General, 127*, 227-250.
- Smith, J. D., Shields, W. E., & Washburn, D. A. (2003). The comparative psychology of uncertainty monitoring and metacognition. *The Behavioral and Brain Sciences, 26*, 317-339. (a)
- Smith, J. D., Shields, W. E., & Washburn, D. A. (2003). Inaugurating a new area of comparative cognition research. *The Behavioral and Brain Sciences, 26*, 358-373. (b)
- Smith, J. D., & Washburn, D. A. (2005). Uncertainty monitoring and metacognition by animals. *Current Directions in Psychological Science, 14*, 19-24.

- Son, L. K., & Kornell, N. (2005). Metaconfidence judgments in rhesus macaques: explicit vs. implicit mechanisms. In H. S. Terrace and J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 296-320). New York: Oxford University Press.
- Thomson, G. H. (1920). A new point of view in the interpretation of threshold measurements in psychophysics. *Psychological Review*, 27, 300-307.
- 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, and Computers*, 24, 157-164.
- Woodworth, R. S. (1938). *Experimental psychology*. New York: Holt.

Appendix

Murph. To fit the SDT model of performance, we built a virtual version of the Sparse-Dense task as the animals experienced it, and then placed in that task environment simulated observers with differently sized perceptual errors and differently placed criterion points. To model Murph's performance in his Task 3, we evaluated 7,560 simulated observers whose perceptual errors varied from 1 to 12 steps along the stimulus continuum, whose Sparse-Uncertain criterion varied from Step 1 to Step 20, and whose Uncertain-Dense criterion varied from each value of the Sparse-Uncertain criterion up to Step 41. Each simulated observer was given 6,600 trials in the task to match Murph's trial count, misperceiving each objective stimulus (with its perceptual error the standard deviation of a Gaussian distribution around the objective stimulus level) into a subjective stimulus impression. Then the simulated observer responded Sparse, Dense, or Uncertain as the impression fell below the Sparse-Uncertain criterion, above the Uncertain-Dense criterion, or between these two. For each simulated observer, the overall performance pattern (i.e., the use of the Sparse, Dense, and Uncertain responses at each of 26 bin levels) was compared to the observed performance mathematically, yielding a fit index that we minimized to find the decisional framework that fit most closely that of the animal. This fit measure was the sum of the squared deviations between the 78 observed and predicted data pairs. For Murph in his Task 3, the best-fitting parameters were 7 steps (along the 41-step continuum) for the Gaussian perceptual error, Step 13 (Bin 7) for the Sparse-Uncertain criterion point, and Step 19 (Bin 12) for the Uncertain-Dense criterion point.

Lou. The SDT modeling procedures described for Murph were applied to Lou's Task 2 performance as well. His perceptual error was also found to be 7 steps (along the 41-step continuum). Both of his criterion points were estimated to be at Step 24 (Bin 16). The 0-width of his uncertainty region (from Step 24 to Step 24) is consistent with his non-use of the uncertainty response. The model fit his observed data pattern very well. Next we considered the optimality of Lou's performance in comparison to a family of simulated observers who all had perceptual errors of 7 (because this is a real perceptual limit) but who had differing criterion placements. These simulated observers had Sparse-Uncertain criterion points that varied from Step 1 to Step 24, and Uncertain-Dense criterion points that varied from each value of the Sparse-Uncertain criterion up to Step 41. Each of 708 simulated observers was given 7,711 trials in the task to match Lou's trial count, misperceiving each objective stimulus (with 7 the standard deviation of a Gaussian distribution around the objective stimulus level) into a subjective stimulus impression. Then the simulated observer responded Sparse, Dense, or Uncertain as the impression fell below the Sparse-Uncertain criterion, above the Uncertain-Dense criterion, or between these two. For each simulated observer, given all the times involved on every trial, we calculated the rewards per minute that that decisional strategy would have earned and so gauged the relative adaptiveness of the strategy Lou chose.

Author Notes

J. David Smith, Department of Psychology and Center for Cognitive Science,
University at Buffalo, State University of New York.

Michael J. Beran, Language Research Center, Georgia State University.

Joshua S. Redford, Department of Psychology, University at Buffalo, State
University of New York.

David A. Washburn, Department of Psychology and Language Research Center,
Georgia State University.

Correspondence concerning this article should be addressed to J. David Smith, 346
Park Hall, SUNY Buffalo, Buffalo, NY 14260, or by email to psysmith@buffalo.edu.

The preparation of this article was supported by the National Institute of Child
Health and Human Development (HD-38051).

This article honors the research careers of Thomas O. Nelson and Donald R. Griffin
who pioneered the study of human metacognition and animal awareness.

Figure Captions

Figure 1. *A.* The screen from a trial in the Sparse-Dense discrimination of Smith et al. (1997). *B.* The performance of Monkey Abel in the Sparse-Dense task. The Dense response was correct for boxes with exactly 2950 pixels—these trials are represented by the rightmost data point for each curve. All other boxes deserved the Sparse response. The horizontal axis indicates the pixel-density of the box. The solid line represents the percentage of trials receiving the uncertainty response at each density level. The error bars show the lower 95% confidence limits. The percentages of trials ending with the Dense response (dashed line) or Sparse response (dotted line) are also shown. *C.* The performance of Monkey Baker in the Sparse-Dense discrimination, depicted in the same way. *D.* The performance of seven humans in the Sparse-Dense discrimination, depicted in the same way. To equate discrimination performance across participants, the data were normalized to place each participant's discrimination crossover at a pixel density of about 2700. The horizontal axis indicates the normalized pixel-density of the box.

Figure 2. *A.* The performance of humans in the deferred-reinforcement uncertainty-monitoring tasks of Experiment 1. The horizontal axis indicates the density bin of the box. The Sparse and Dense responses, respectively, were correct for boxes at Density Bins 1-8 and 9-16. The open squares show the proportion of trials attempted that were answered correctly. The dark circles show the proportion of trials receiving the uncertainty response at each density bin. *B.* The performance of humans in Experiment 2, depicted in the same way. *C.* The performance of humans in Experiment 3, depicted

in the same way. In this case, the Sparse and Dense responses, respectively, were correct for boxes at Density Bins 1-13 and 14-26

Figure 3. A,B. The performance of two humans in the Sparse-Dense discriminations of Experiments 1 and 2, depicted as in Figure 2.

Figure 4. A. The performance of Murph in the deferred-reinforcement uncertainty-monitoring task of Experiment 4 (Task 1). The horizontal axis indicates the density bin of the box. The Sparse and Dense responses, respectively, were correct for boxes at Density Bins 1-13 and 14-26. The open squares show the proportion of trials attempted that were answered correctly. The dark circles show the proportion of trials receiving the uncertainty response at each density bin. Representative 95% confidence intervals are shown for the peak of uncertainty responding near the task's breakpoint and for the first bins to the right and left of that peak in which uncertainty responding was significantly reduced relative to the peak. *B.* Murph's overall performance in his Task 2, depicted in the same way. *C.* Murph's performance in the first 400 trials of his Task 2, depicted in the same way (but with error bars omitted because this is a small sample of the animal's behavior).


Figure 5. A. Murph's performance in his Task 3, depicted as in Figure 4. *B.* Murph's performance in his Task 3, with the proportion of trials declined in each density bin plotted against the proportion correct for that bin. *C.* Murph's performance in his Task 3, depicted showing separately his use of the Sparse response (open circles), Dense response (open triangles), and Uncertainty response (dark circles). *D.* Murph's performance in his Task 3, with his proportion of trials declined in each density bin

plotted against the decisional distance of the bin from his decisional breakpoint (Bin 9 = 0; Bins 8 and 10 = 1; Bins 7 and 11 = 2, etc.).

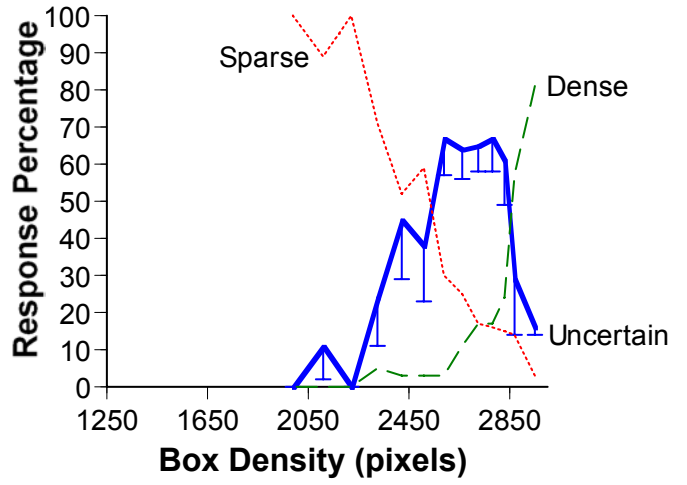
Figure 6. The best-fitting prediction of Murph's Task 3 performance as made by a formal SDT model.

Figure 7. A. The performance of Lou in his Task 1, depicted as in Figure 4. *B.* Lou's performance in Task 2, depicted in the same way.


A.



B. Monkey



C. Monkey



D. Humans

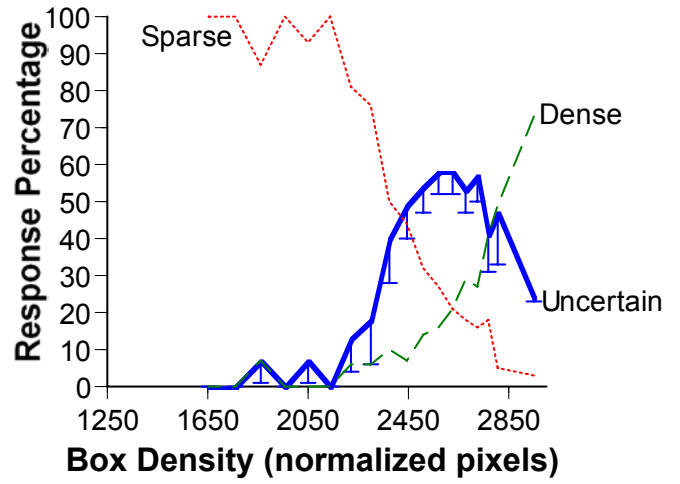


Figure 1.






Figure 2.

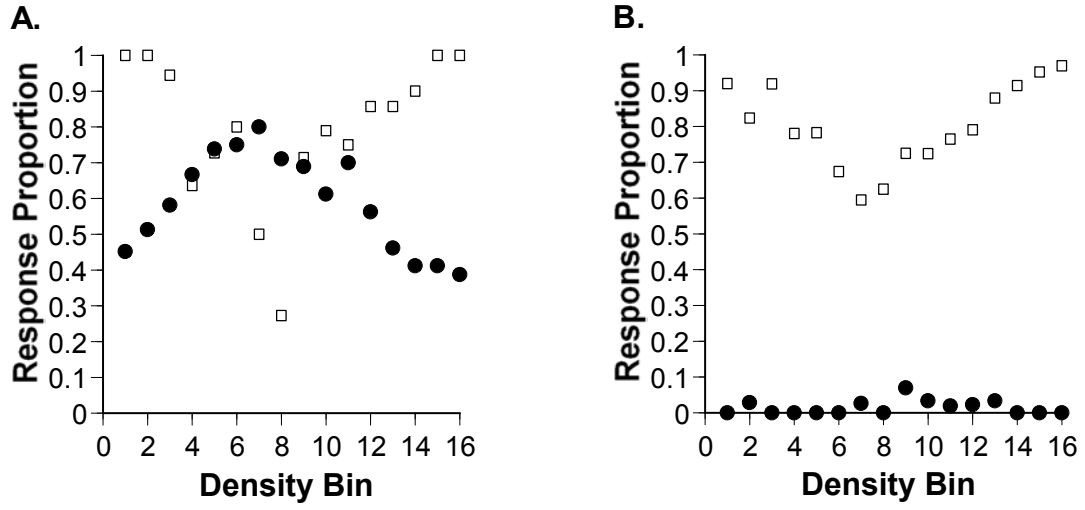


Figure 3.




Figure 4.

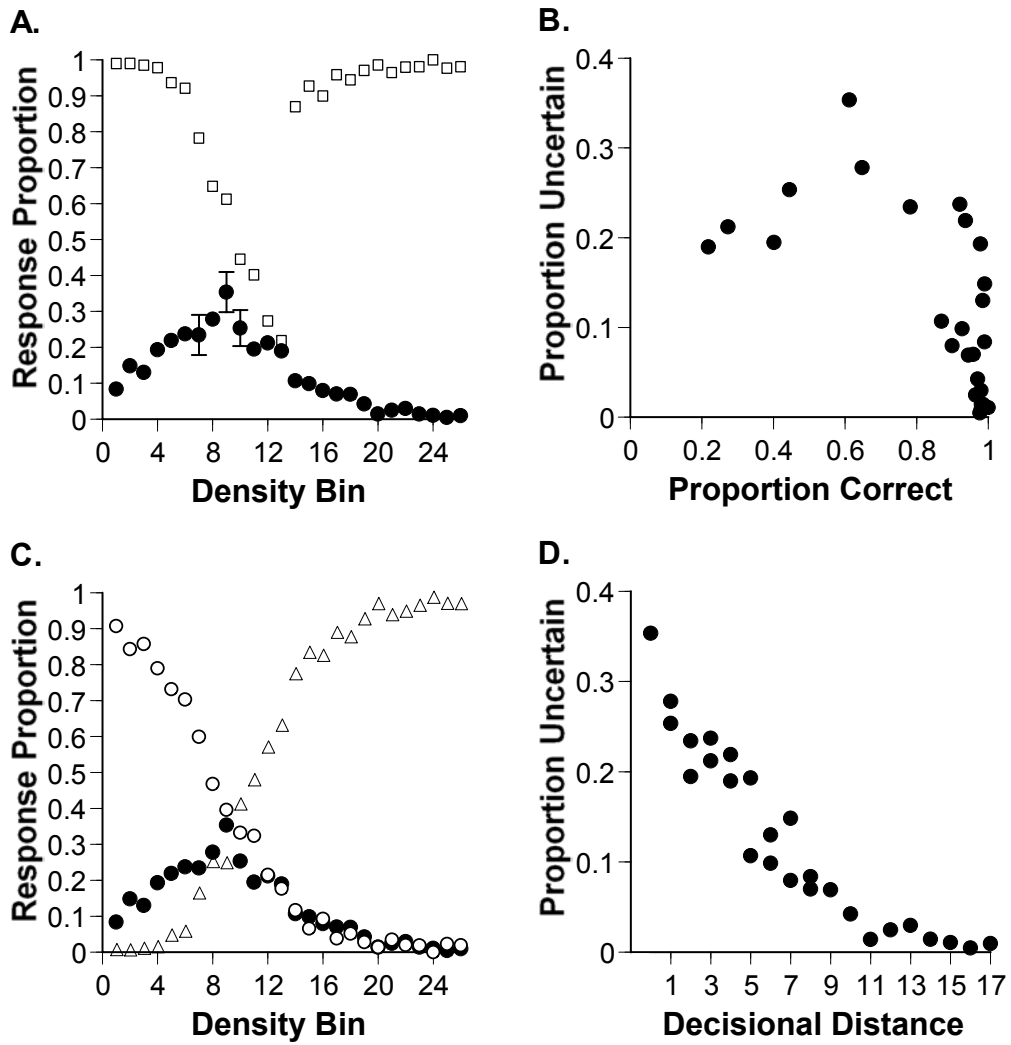


Figure 5.

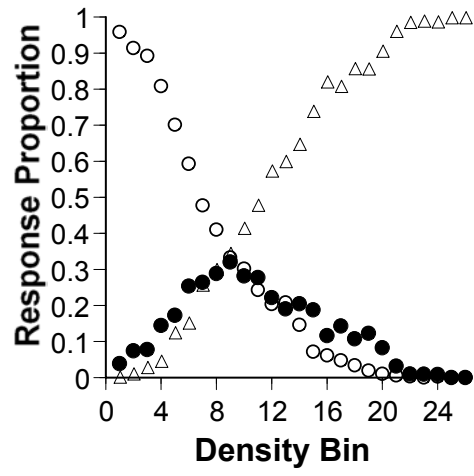


Figure 6.

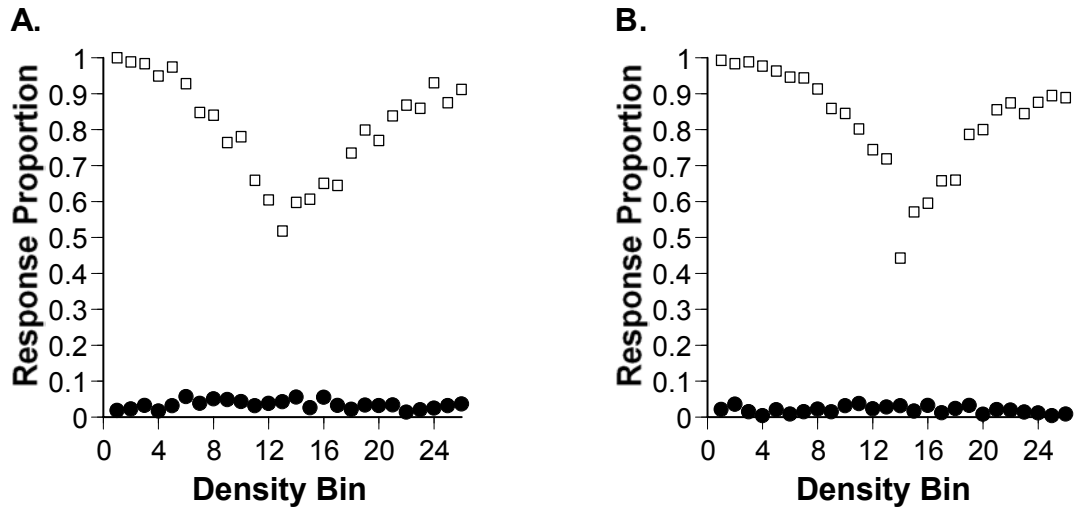


Figure 7.