



The Cognitive Architecture of Uncertainty

J. David Smith, Brooke N. Jackson, and Barbara A. Church*

Department of Psychology, Georgia State University

*Corresponding author (Email: bchurch@gsu.edu)

Citation – Smith, J. D., Jackson, B. N., & Church, B. A. (2019). The cognitive architecture of uncertainty. *Animal Behavior and Cognition*, 6(4), 236-246. <https://doi.org/10.26451/abc.06.04.03.2019>

Abstract – The authors consider theory in the animal-metacognition literature. Theoretical interpretation was long dominated by associative interpretations, a conservative approach well illustrated in the 2009 special issue of animal metacognition in *Comparative Cognition and Behavioral Reviews*. We suggest, though, that this approach risks a self-limiting understanding of animal minds, and an imprecise understanding of the cognitive requirements inherent in metacognition tasks. In fact, some tasks self-entail the need for higher-level decision-making processes, processes that—in humans—we would call explicit, declarative, and conscious. These points are illustrated using the inaugural study on dolphin metacognition. We urge researchers to turn more toward illuminating the cognitive architecture of capacities like metacognition, including illuminating the depth, and structure, the learning/memory systems, the cognitive levels, and the declarative awareness possibly present in animals' minds. The empirical development of this literature demonstrates that researchers are now prepared to do so. This study can produce strong synergies across the allied fields of biopsychology, comparative and cognitive psychology, and neuroscience.

Keywords – Metacognition, Uncertainty monitoring, Comparative psychology, Cognitive psychology, Cognitive architecture

Metacognition—the monitoring and controlling of one's own cognition—is a significant research topic in cognitive psychology, cognitive neuroscience, and educational psychology (Dunlosky & Bjork, 2008; Flavell, 1979; Fleming & Frith, 2014; Nelson & Narens, 1990). Metacognition is a sophisticated capacity that may be linked to self-awareness and consciousness (Koriat, 2007; Nelson, 1996). It deserves complete scientific understanding, including understanding its evolutionary origin. This understanding might reveal the roots of reflective mind in humans and its emergence during cognitive development. It could provide animal models for studying cognitive awareness and its neuroscientific organization. It might even give researchers leverage on the problem of animal consciousness. In this study also lies the potential for cross talk among comparative psychologists, cognitive psychologists, and neuroscientists—cross talk still far too infrequent given the powerful synergies possible. For these reasons, in 2019 as in 2009 for the *Comparative Cognition and Behavioral Reviews (CCBR)* special issue, metacognition is still a worthy focus of cross-species research.

The Dolphin's Tale

Smith et al.'s (1995) first application of the uncertainty paradigm—25 years ago—tested a bottle-nosed dolphin (*Tursiops truncatus*) in a psychophysical threshold task. Natua (Figure 1) discriminated High tones (half the trials, always 2,100 Hz, deserving the “High” response) from Low tones (half the trials, 1,200-2,099 Hz, deserving the “Low” Response). The dolphin could also make a

third response that has become known as the uncertainty response (UR). That is, he could choose to decline to answer any trials he chose. In Smith et al.'s original paradigm, the UR received no primary reinforcer. Instead, it advanced the task into a following trial that was chosen to be easy for the dolphin. Figure 1A shows that the dolphin responded "Low" and "High," respectively, for many Low and High tones. His extraordinary auditory performance produced a Low-High psychophysical threshold at about 2,080 Hz. But the crucial result was that Natua judged correctly when he was in perceptual difficulty in the discrimination and likely to err. On those threshold trials, he made URs selectively. The fact that the UR was used most at the dolphin's true psychophysical limit must shape one's theoretical interpretation of that response. We will come to the matter of that theoretical interpretation.

Natua produced an additional crucial result. On some trials, he swam slowly and hesitantly, or waggled his head, or worked his jaw. We conducted a factor-analytic study of these behaviors, asking observers—who were blind to the trials' pitch levels—to provide ratings of these incidental behaviors. A strong general factor of hesitation and wavering emerged. Figure 1B shows that scores on this factor varied systematically with pitch level, reproducing the curve of the animal's URs.

Thus, Natua not only used the experimenter-provided UR to manage difficulty and uncertainty. He also expressed symptoms of something like his own uncertainty, through dolphin-specific behaviors. These facts must also help shape one's interpretation of Smith et al.'s (1995) experiment.

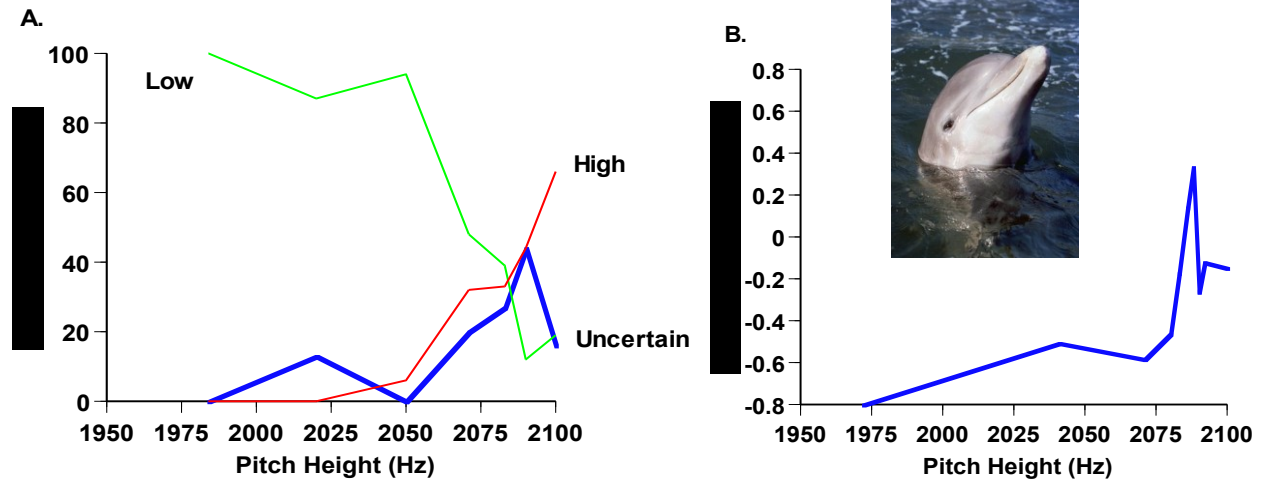


Figure 1. **A.** Natua's performance in a threshold-uncertainty paradigm. He touched a High response when a 2,100 Hz tone sounded. He touched a Low response when any lower tone sounded. The uncertainty response let him decline the trial and move instead into an easy next trial. The pitch of the low tones was varied to find his auditory psychophysical threshold and to study his uncertainty responses in this threshold region of difficult trials. The horizontal axis indicates the frequency of the tone. The green and red lines show the percentage of Low and High responses, respectively, at each pitch level. The blue line shows the percentage of trials receiving the uncertainty response at each pitch level. **B.** The dolphin's weighted overall Factor 1 behavior (hesitancy, slowing, wavering) for different pitch levels. From 'The Uncertain Response in the Bottlenosed Dolphin (*Tursiops truncatus*),' by J. D. Smith, J. Schull, J. Strode, K. McGee, R. Egnor, and L. Erb, 1995, *Journal of Experimental Psychology: General*, 124, 399, 402. Copyright 1995 by the American Psychological Association. Reprinted with permission.

The Associative Debate

As (human) cognitive researchers, these results raised many questions for us. We wanted to understand the cognitive architecture of uncertainty, and whether nonhumans shared that architecture. We wanted to know at what cognitive level humans and animals feel doubt, and whether it is the same level. We wondered about the role of working memory and working consciousness in monitoring uncertainty, and whether animals, as humans, do engage in such monitoring with some conscious awareness. We thought that metacognition experiments might give scientists leverage on the problem of animal consciousness, an exciting possibility.

But the theoretical trajectory of the literature did not go in those directions. Instead, for years thereafter, the dolphin experiment, along with companion experiments with rhesus monkeys (*Macaca mulatta*), received heavy incoming constructive criticism (e.g., Carruthers, 2008; Hampton, 2009; Staddon, Jozefowicz, & Cerutti, 2007). Many colleagues chose to emphasize the associative-learning processes they thought could produce results like the dolphin's over the alternative explanation using processes like metacognition or cognitive monitoring. For example, in some studies, even in the dolphin study, researchers gave animals indeterminate stimuli they hoped would stir up animal uncertainty. But those stimuli also had "associative" effects—they caused errors and earned scant rewards. Associative-learning systems could sense the relationship between such stimuli and sparse rewards. Avoidance responses could entrain to problematic stimuli, fending off errors. Thus the "metacognitive" uncertainty response could really be a low-level avoidance response.

Other colleagues incorporated these theoretical ideas into formal-mathematical models (Jozefowicz, Staddon, & Cerutti, 2009; Le Pelley, 2012). The goal of these models was to show that mathematical gradients could reproduce the observed findings. These "associative" models included abstract representational registers arranged as a series of perceptual or memory psychological states (e.g., from "Low" to "High" tones in the dolphin's case). These registers contained response-strength levels, so that some register positions produced one response and other positions produced another response. The registers were updated based on trial outcomes, so that the registers' response strengths changed to reflect the reinforcement history in the task. For example, a threshold-trial register level would contain many errors from primary discrimination responses ("Low," "High"), producing weak response strengths that the "uncertainty" response would out-compete. But an easy-trial register level would contain a rich reinforcement history of wins and rewards, dictating a primary discrimination response.

In fact, examining the *CCBR* issue from 2009, one sees that it focused sharply on this associative issue. Hampton (2009) concluded that "many of the demonstrations of metacognition in nonhumans can be explained in terms of associative learning." Jozefowicz et al. (2009) attributed animals' so-called metacognitive performances to "familiar learning principles, and thus lacking any metacognition construct." Crystal and Foote (2009) concluded that "low-level mechanisms are sufficient to explain data from existing methods."

Shortly thereafter, Le Pelley (2014) wrote his self-titled *Primate Polemic*—criticizing the attribution of metacognition to animals and attributing their "metacognitive" performances to associative learning. Readers will understand that it was not actually the primates' polemic. Le Pelley's (2012) article title made the same point: *Metacognitive monkeys or associative animals? Simple reinforcement learning explains uncertainty in nonhuman animals*.

In fact, one can summarize simply the second epoch of the animal-metacognition literature. It was dominated by an associative debate that overshadowed the field's theoretical development (e.g., Carruthers, 2008; Hampton, 2009; Jozefowicz et al., 2009; Le Pelley, 2012, 2014; Smith, 2009; Smith, Couchman, & Beran, 2014b; Smith, Zakrzewski, & Church, 2016; Staddon et al., 2007; review in Smith, Church, Beran, & Washburn, 2018).

Therefore, researchers turned to the problem of eliminating associative concerns. There was not a choice in this. The associative debate was front and center. It had to be adjudicated in each study and article. During this time, human cognitive researchers asked us to please stop harping on associationism. (We could not do so.) They said it was obvious what was going on—animal metacognition was going on. (No comment.) The gulf between human-cognitive interpretations and animal-associative interpretations still must be taken seriously today, to understand what it means for collaboration, cross talk, and synergies across the disciplines of bio-behavioral research.

Yet it has been a decade of positive development since 2009. There have been elegant studies on animal metamemory (e.g., Morgan, Kornell, Kornblum, & Terrace, 2014; Templer & Hampton, 2012; Templer, Lee, & Preston, 2017). Various associative explanations of metacognitive performances have been carefully and systematically ruled out (e.g., Basile, Schroeder, Brown, Templer, & Hampton, 2015).

For example, consider Hampton's (2001) influential metamemory study. His macaques performed a delayed matching-to-sample task. As they decided whether to complete each memory test or

to respond Uncertain instead, they had *no associative stimuli available* that could trigger one behavior instead of another. Instead, they had to monitor their memory to see if it warranted confidence or doubt. This is metamemory, not associative responding. Hampton showed that even the forgetting delay was not an associative cue. The monkeys responded Uncertain far more on blank trials (no sample item given) than otherwise, even with forgetting time equated. This is because they rightly had no memory of the sample on the blank trials. They knew they did not know. Finally, sometimes the monkeys chose to complete the memory test even after quite long forgetting delays. They performed those trials accurately. This is because they were monitoring some internal memory signal. They knew they knew.

Many similar demonstrations have been forthcoming. URs have been shown to be recruited effectively even when immediate reinforcement is denied the animal, so that the UR cannot be conditioned to “difficult” stimuli (Smith, Beran, Redford, & Washburn, 2006). URs have been shown to be based on executive cognitive processes that are especially working-memory intensive (Smith, Coutinho, Church, & Beran, 2013; see below). Beautiful comparisons across animal species have been arranged that are beginning to draw the map of the evolutionary emergence of the metacognitive capacity (Beran, Smith, Coutinho, Couchman, & Boomer, 2009; Fujita, 2009; Paukner, Anderson, & Fujita, 2006; Roberts et al., 2009; Sutton & Shettleworth, 2008). This map would not look as it does if all these species shared a basic capacity to make “associative” uncertainty processes.

The modeling approach to asserting associationism also failed—for a simple reason. The models depended on an abstract mathematical framework. They were unable to incorporate the organism’s true processes, representations, introspections, feelings, and awareness. For example, even the obviously metacognitive performance of an obviously conscious human could be laid onto, and fit by, the associative models. Thus, the clearest instance of metacognition would be declared associative. The associative models were never associative (Smith et al., 2014b; additional commentary in Basile & Hampton, 2014; Carruthers, 2014; Le Pelley, 2014; Smith, Couchman, & Beran, 2014a). They were not anything. They were psychologically empty. They never could bear on the crucial question: the correct psychological interpretation of animals’ metacognitive responses in many tasks.

The associative debate has largely evaporated, replaced by a strong theoretical consensus that some animals share some aspects of humans’ metacognitive capacity, though not necessarily all aspects of that capacity. The field has arrived at a strong theoretical belay stance, based on a decade of productive research.

Two Possible Lessons from the Associative Decade

Nonetheless, the associative decade in animal metacognition may offer lessons for comparative psychology as it enters another productive decade. First, associative descriptions of animals’ performance generally draw a flat, dimensionless theoretical picture of animal mind. There is not a focus on the mind’s cognitive architecture, but on the stimulus-response bonds. There are not cognitive levels, just the associative level. There are not different systems guiding learning and behavior, just the associative system. There are not processes especially executive, or especially working-memory intensive—these ideas clash with the associative construct. There are not declarative or conscious cognitive processes in the associative framework, for the same reason. Consequently, empirical investigations into these kinds of issues are not directly encouraged by theory, and they may be resisted by some.

Possibly, though, animal minds may be trying to tell us that they have depth, and levels, and awarenesses, some consciousness, structure, and different systems guiding learning and behavior. If so, and if these kinds of empirical explorations could be more proactively fostered by the field, then the field’s theoretical development would benefit.

Second, we believe that the associative attribution is sometimes applied too easily, too glibly. The result is that we may take our theoretical eyes off the specifics of the cognitive problem the animal faces, the exact mental representations that need to be generated and sustained during performance, and the necessary structure of a cognitive system that can achieve the given performance. But our tasks may be trying to tell us that they require a mind that has depth, and levels, and systems, and so forth. If so, then

there is great benefit in moving beyond the surface associative description to focus sharply on the task's cognitive demands. The stimulus-response bond is really only the hitch joining the cart to the tractor. We need also to ask what is in the cart, where is the tractor going, and why.

The Inescapable Logic of Doubt

The dolphin experiment illustrates these problems. In fact, we believe the focus on associative explanations of the dolphin study was a theoretical mistake that arose from imprecise and misapplied constructs. It was not understood a) that threshold stimuli are profoundly special and are not even associable stimuli; b) that uncertainty cannot control behavior as supposed by associative-learning theory; c) that doubt is definitionally not an associative construct; and d) that uncertainty self-entails the higher-level cognitive psychology that must manage it. Here is how these problems can be seen reflected in Smith et al. (1995).

Figure 1A showed that the dolphin's URs peaked near 2,085 Hz. In essence, with the 2,100 Hz "High" tone as the musical pitch C, the dolphin felt maximally uncertain to musical pitch B#####: that is, $1/9^{\text{th}}$ of a half step away. It is highly unlikely that there was a discrete, intermediate pitch class there that could have prompted some conditioned response—we believe that no known biological auditory system would achieve this feat of resolution (especially given successive presentations of tones a minute apart, separated by trainer interaction including auditory stimulation and a reward protocol also including auditory stimulation). The threshold stimulus was not a stimulus in this discrete, intermediate sense.

In fact, the threshold stimulus is not a stimulus at all. Imagine yourself in the shower, when you suddenly wonder if the phone is ringing. No one in this moment has ever said: There is one of those it-might-be-a-phone-ringing-or-it-might-not-be-type stimuli. No, they say: I don't know if the phone is ringing. It is interpretatively important that there is even an I, a me, there—a selfness. The psychology is self-evidently no longer about the world of stimuli and responses.

Likewise, it is possible that the dolphin at threshold was expressing that he did not know whether the High "phone" was ringing. Note carefully that this is not some speculative, anthropomorphic analogy. This is a purely logical point about the psychological nature of the "threshold state," sharply contrasted to the nature of an associable stimulus that can bond to behavior.

Next consider that at threshold the relevant, competing response strengths are at equipoise. So, the threshold state cannot condition, or exert stimulus control over, either primary response ("High," "Low"). Nor can equipoise power the uncertainty response. In many experiments, the UR earns no reward, no differential consequence, no special next trial, nothing. It is not rewarding in the required sense to exert behavioral control at all. And yet some species use it to manage difficulty. Why? The threshold state cannot be about stimulus or behavioral control as associative-learning theory understands these constructs. Indeed, even committed associative theorists have acknowledged that the threshold state changes profoundly the nature of stimulus control and associative behavior (Boneau & Cole, 1967; Commons, Nevin, & Davison, 1991; Davison, McCarthy, & Jensen, 1985; Miller, Saunders, & Bourland, 1980; Terman & Terman, 1972).

To illustrate this point, we will lay down a challenge here that we know we will regret. It is this: *No-one will ever be able to condition any classical reflex (not salivation, not eye blinks, not fear suppression) using a threshold feeling-state of uncertainty as the CS.*

On several occasions, we tried to remind associative theorists that threshold uncertainty tasks were inconsistently mapped in the manner discussed by Shiffrin and Schneider (1977; also Schneider & Shiffrin, 1977). At threshold, the animal genuinely cannot discern what pitch (High? Low?) it is hearing. The perceptual system is variable and prone to some error. Either stimulus could be present and either response could be warranted. Therefore, making either reflexive response would be unwise. In that case, the cognitive system must resort to some kind of cognitive referee to break the associative tie, to set response criteria, to remember what has worked out well in this situation in the past. Shiffrin and Schneider called these higher-level processes controlled cognitive processes.

Probably Natua brought an analog of these controlled processes to bear on the threshold task that always demands them. If the task does not demand them, then one has not found the animal's threshold yet.

Another consideration goes to the dolphin's own uncertainty behaviors—the slowing, the head wagging, and so forth. Why was his behavior so distinctively different on the threshold trials? It would not be if the UR were just the 3rd reaction to the 3rd stimulus class. Yet these behaviors found important historical resonance in the work of Tolman. He observed similar behaviors in rats that he called “lookings and runnings back and forth” and that he thought might become “A Behaviorist's Definition of Consciousness” (Tolman, 1927, 1938) because they made visible mental turmoil in the animal and changing cognitive determinations in mind. Here, Tolman raised the striking possibility that the higher-level, cognitive-resolution processes at threshold might even be conscious.

If the dolphin's conflict behaviors were expressing his conscious decisional turmoil on difficult trials, it is an extraordinary demonstration. We do not assert that. However, the possibility is a matter of fact, a matter of theoretical consideration, and a matter of empirical investigation that must be clearly acknowledged.

In a sense, therefore, threshold uncertainty is the associative theorist's nightmare. The usual associative constructs like stimulus control, stimulus-response bonds, and so forth explode in the tight quarters of a psychophysical threshold task. But some minds rally at threshold, bringing alternative criterial and decision-making and conflict-resolving processes to bear. The essential requirement of a threshold task is for a higher-level, controlled process that can adjudicate uncertainty. Animals time and again have shown that their reflective minds have the depth and structure to manage effectively this adjudication. For these reasons, threshold uncertainty is a goldmine for the cognitive or metacognitive psychologist, or the behavioral neuroscientist.

Confirming the Distinctive Cognitive Architecture of Uncertainty

It is twenty-five years on from Smith et al. (1995). As a practical matter, it has emerged that nonhuman primates (enjoying tang-flavored Cap'n Crunch food rewards) can perform 10 times the daily trials that a dolphin (enjoying whole-herring rewards) can. Nonetheless, the lessons discussed here from the earliest experiment—that the threshold state requires a qualitatively distinctive psychological explanation, and that the UR cannot be understood associatively—still continue to be confirmed. In closing, we will illustrate this with two recent studies, with monkeys and with the human primate.

First, Smith et al. (2013) gave macaques a concurrent memory load while they performed the Sparse-Uncertain-Dense discrimination task shown in Figure 2. Macaques saw an unframed box filled to one of 60 densities with lit pixels. Each level was 1.8% more pixel dense. They touched the S or D icons with their computer cursor to respond that a box was Sparse or Dense (Stimulus Levels 1-30 and 31-60, respectively). Trial difficulty centered and focused on the trial levels near the Sparse-Dense breakpoint (Levels 30-31). Monkeys touched the ? icon to decline the trial, transitioning out of any difficult trials they did not wish to complete.




Figure 2. Illustrating a trial in a Sparse-Uncertain-Dense task

Smith et al.'s (2013) first hypothesis was that Sparse and Dense responses would be grounded in associative learning as traditionally understood, making few working-memory demands so that they would be unaffected by the concurrent memory load. Their second hypothesis was that URs would be an expression of higher-level, controlled decision making. Thus, URs might be strongly affected by the load—uncertainty responding might evaporate under load conditions. Figure 3 shows that URs were selectively disrupted by the concurrent memory task. Sparse and Dense responses were not. The observation that the metacognitive responses of macaques can be working-memory intensive fits perfectly with the idea that uncertainty self-entails the intervention of more executive processes.

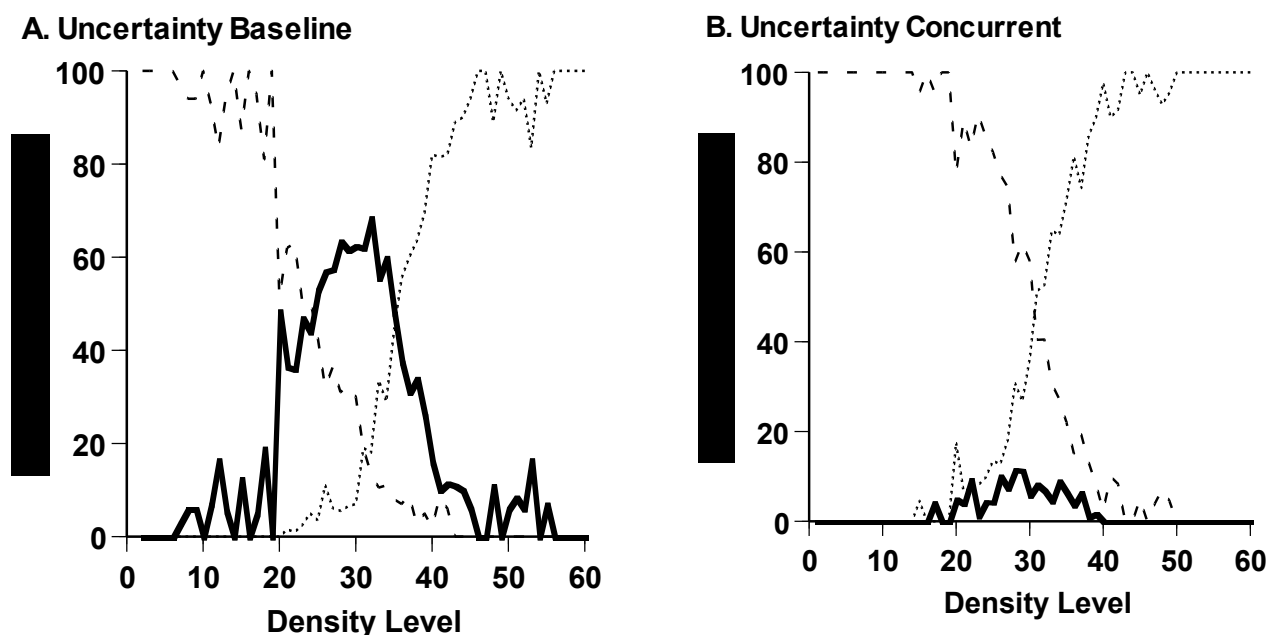


Figure 3. Percentage of uncertainty responses (solid line), sparse responses (dashed line), and dense responses (dotted line) made by macaques Murph and Lou in their baseline performance (A) and in their first phase of concurrent-load testing (B). From "Executive-Attentional Uncertainty Responses by Rhesus Macaques (*Macaca mulatta*)," by J. D. Smith, M. V. C. Coutinho, B. A. Church, & M. J. Beran, *Journal of Experimental Psychology: General*, 142, 472. Copyright 2013 by the American Psychological Association. Reprinted with permission.

Second, Paul et al. (2015) placed humans in a similar Sparse-Uncertain-Dense task. Then, using rapid event-related fMRI, they showed that the neural circuit recruited during humans' URs is distinctively different from that recruited during humans' Sparse and Dense responses. Figure 4 (top) shows the simple activation pattern when participants made correct Sparse or Dense responses in lieu of URs. There was selective activation in the occipital lobe bilaterally and in the caudate nucleus and nucleus accumbens. This supports the interpretation that Sparse and Dense responses are first-order associative responses in which discrete stimuli elicit well-conditioned responses. Figure 4 (bottom) shows the complex activation pattern when participants made URs in lieu of Sparse/Dense responses. URs reflected the activation of a broad-based network including prefrontal cortex, anterior and posterior cingulate cortex, anterior insula, and posterior parietal areas. Thus, as discussed above, uncertainty monitoring apparently benefits from, and may inherently demand, a larger-scale cognitive-control network.




Figure 4. Top. Whole-brain results from the contrast Correct Sparse-Dense Response > Uncertainty Response projected on an inflated lateral and medial cortical surface. Table 2 in Paul et al. (2015) listed the coordinates and number of voxels for every significant cluster. Images were cluster thresholded (correcting for multiple comparisons) at $z > 3.54$, $p < .01$. *Bottom.* Whole-brain results from the contrast Uncertainty Response > Correct Sparse-Dense Response projected on an inflated lateral and medial cortical surface. Table 3 in Paul et al. (2015) listed the coordinates and number of voxels for every significant cluster. Images were cluster thresholded (correcting for multiple comparisons) at $z > 3.54$, $p < .01$. From “Neural Networks of the Psychophysical Uncertainty Response,” by E. J. Paul, J. D. Smith, V. Valentin, B. O. Turner, A. K. Barbey, & F. G. Ashby, *Cortex*, 71, 316, 317. Copyright 2015 Elsevier Ltd. Reprinted with permission.

To illustrate this point, we will lay down another challenge. It is this: When well-acclimated macaques eventually participate in the equivalent fMRI assessment to that in Paul et al. (2015), they will confirm that the uncertainty-monitoring circuit is broad-based and complex, even in a species that thinks without any language or declarative strategies. We think this must be so, because the threshold state of doubt inevitably self-entails a higher-level cognitive capacity in mind—for humans, for macaques, and for that awesome dolphin Natua.

We hope that a special issue organized in 2029 will reflect that comparative psychology turned more toward exploring the cognitive architecture of important capacities like uncertainty monitoring and metacognition, including considering how animal minds may show depth, and structure, and systems, and cognitive levels, and awareness, and declarative capability. Recent research in the field of animal metacognition makes it clear that researchers are now prepared to do so. We note that this approach will simultaneously be the study of humans’ cognitive emergence during evolution, and so it will knit together more closely and constructively the allied fields of biopsychology, comparative and cognitive psychology, and neuroscience.

Acknowledgments

The preparation of this article was supported by Grants HD-060563 and HD-093690 from the NICHD. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

- Basile, B. M., & Hampton, R. R. (2014). Metacognition as discrimination: Commentary on Smith et al. (2014). *Journal of Comparative Psychology, 128*, 135–142. <https://dx.doi.org/10.1037/a0034412>
- Basile, B. M., Schroeder, G. R., Brown, E. K., Templer, V. L., & Hampton, R. R. (2015). Evaluation of seven hypotheses for metamemory performance in rhesus monkeys. *Journal of Experimental Psychology: General, 144*, 85–102. <https://dx.doi.org/10.1037/xge0000031>
- Beran, M. J., Smith, J. D., Coutinho, M. V. C., & Couchman, J. J., & Boomer, J. (2009). The psychological organization of “uncertainty” responses and “middle” responses: A dissociation in capuchin monkeys (*Cebus apella*). *Journal of Experimental Psychology: Animal Behavior Processes, 35*, 371–381. <https://dx.doi.org/10.1037/a0014626>
- Boneau, C. A., & Cole, J. L. (1967). Decision theory, the pigeon, and the psychophysical function. *Psychological Review, 74*, 123–135. <https://dx.doi.org/10.1037/h0024287>
- Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind & Language, 23*, 58–89. <http://dx.doi.org/10.1111/j.1468-0017.2007.00329.x>
- Carruthers, P. (2014). Two concepts of metacognition. *Journal of Comparative Psychology, 128*, 138–139. <https://dx.doi.org/10.1037/a0033877>
- Commons, M. L., Nevin, J. A., & Davison, M. C. (1991). *Signal detection: Mechanisms, models, and applications*. Hillsdale, NJ: Erlbaum.
- Crystal, J. D., & Foote, A. L. (2009). Metacognition in animals. *Comparative Cognition & Behavior Reviews, 4*, 1–16. <https://dx.doi.org/10.3819/ccbr.2009.40001>
- Davison, M., McCarthy, D., & Jensen, C. (1985). Component probability and component reinforcer rate as biases of free-operant detection. *Journal of the Experimental Analysis of Behavior, 44*, 103–120. <https://dx.doi.org/10.1901/jeab.1985.44-103>
- Dunlosky, J., & Bjork, R. A. (2008). *A handbook of metamemory and memory*. Hillsdale, NJ: Psychology Press.
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. *American Psychologist, 34*, 906–911. <https://dx.doi.org/10.1037/0003-066X.34.10.906>
- Fleming, S. M., & Frith, C. D. (2014). Metacognitive neuroscience: An introduction. In S. M. Fleming & C. D. Frith (Eds.), *The cognitive neuroscience of metacognition* (pp. 1–6). Berlin Heidelberg: Springer-Verlag. https://dx.doi.org/10.1007/978-3-642-45190-4_1
- Fujita, K. (2009). Metamemory in tufted capuchin monkeys (*Cebus apella*). *Animal Cognition, 12*, 575–585. <http://dx.doi.org/10.1007/s10071-009-0217-0>
- Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences, 98*, 5359–5362. <https://dx.doi.org/10.1073/pnas.071600998>
- Hampton, R. R. (2009). Multiple demonstrations of metacognition in nonhumans: Converging evidence or multiple mechanisms? *Comparative Cognition & Behavior Reviews, 4*, 17–28. <https://dx.doi.org/10.3819/ccbr.2009.40002>
- Jozefowicz, J., Staddon, J. E. R., & Cerutti, D. T. (2009). Metacognition in animals: How do we know that they know? *Comparative Brain & Behavior Reviews, 4*, 29–39. <http://dx.doi.org/10.3819/ccbr.2009.40003>
- Koriat, A. (2007). Metacognition and consciousness. In P. D. Zelazo, M. Moscovitch & E. Thompson (Eds.), *The Cambridge handbook of consciousness* (pp. 289–326). Cambridge, UK: Cambridge University Press. <https://dx.doi.org/10.1017/CBO9780511816789.012>
- Le Pelley, M. E. (2012). Metacognitive monkeys or associative animals? Simple reinforcement learning explains uncertainty in nonhuman animals. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38*, 686–708. <http://dx.doi.org/10.1037/a0026478>
- Le Pelley, M. E. (2014). Primate polemic: Commentary on Smith, Couchman, and Beran (2014). *Journal of Comparative Psychology, 128*, 132–134. <http://dx.doi.org/10.1037/a0034227>

- Miller, J. T., Saunders, S. S., & Bourland, G. (1980). The role of stimulus disparity in concurrently available reinforcement schedules. *Animal Learning & Behavior*, *8*, 635–641. <https://dx.doi.org/10.3758/BF03197780>
- Morgan, G., Kornell, N., Kornblum, T., & Terrace, H. S. (2014). Retrospective and prospective metacognitive judgments in rhesus macaques (*Macaca mulatta*). *Animal Cognition*, *17*, 249–257. <https://dx.doi.org/10.1007/s10071-013-0657-4>
- Nelson, T. O. (1996). Consciousness and metacognition. *American Psychologist*, *51*, 102–116. <https://dx.doi.org/10.1037/0003-066X.51.2.102>
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 125–173). New York: Academic Press. [https://dx.doi.org/10.1016/S0079-7421\(08\)60053-5](https://dx.doi.org/10.1016/S0079-7421(08)60053-5)
- Paukner, A., Anderson, J. R., & Fujita, K. (2006). Redundant food searches by capuchin monkeys (*Cebus apella*): A failure of metacognition? *Animal Cognition*, *9*, 110–117. <http://dx.doi.org/10.1007/s10071-005-0007-2>
- Paul, E. J., Smith, J. D., Valentin, V. V., Turner, B. O., Barbey, A. K., & Ashby, F. G. (2015). Neural networks underlying the metacognitive uncertainty response. *Cortex*, *71*, 306–322. <https://dx.doi.org/10.1016/j.cortex.2015.07.028>
- Roberts, W. A., Feeney, M. C., McMillan, N., MacPherson, K., Musolino, E., & Petter, M. (2009). Do pigeons (*Columba livia*) study for a test? *Journal of Experimental Psychology: Animal Behavior Processes*, *35*, 129–142. <http://dx.doi.org/10.1037/a0013722>
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I Detection, search, and attention. *Psychological Review*, *84*, 1–66. <https://dx.doi.org/10.1037/0033-295X.84.1.1>
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II Perceptual learning, automatic attending and a general theory. *Psychological Review*, *84*, 127–190. <https://dx.doi.org/10.1037/0033-295X.84.2.127>
- Smith, J. D. (2009). The study of animal metacognition. *Trends in Cognitive Sciences*, *13*, 389–396. <https://dx.doi.org/10.1016/j.tics.2009.06.009>
- Smith, J. D., Beran, M. J., Redford, J. S., & Washburn, D. A. (2006). Dissociating uncertainty responses and reinforcement signals in the comparative study of uncertainty monitoring. *Journal of Experimental Psychology: General*, *135*, 282–297. <https://dx.doi.org/10.1037/0096-3445.135.2.282>
- Smith, J. D., Church, B. A., Beran, M. J., & Washburn, D. A. (2018). Meta-cognition. In J. Vonk & J. T. Shackelford (Eds.), *Encyclopedia of animal cognition and behavior*. New York: Springer. https://dx.doi.org/10.1007/978-3-319-47829-6_1822-1
- Smith, J. D., Couchman, J. J., & Beran, M. J. (2014a). A tale of two comparative psychologies: Reply to commentaries. *Journal of Comparative Psychology*, *128*, 140–142. <https://dx.doi.org/10.1037/a0034784>
- Smith, J. D., Couchman, J. J., & Beran, M. J. (2014b). Animal metacognition: A tale of two comparative psychologies. *Journal of Comparative Psychology*, *128*, 115–131. <http://dx.doi.org/10.1037/a0033105>
- Smith, J. D., Coutinho, M. V. C., Church, B. A., & Beran, M. J. (2013). Executive-attentional uncertainty responses by rhesus macaques (*Macaca mulatta*). *Journal of Experimental Psychology: General*, *142*, 458–475. <http://dx.doi.org/10.1037/a0029601>
- 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. <https://dx.doi.org/10.1037/0096-3445.124.4.391>
- Smith, J. D., Zakrzewski, A. C., & Church, B. A. (2016). Formal models in animal-metacognition research: The problem of interpreting animals' behavior. *Psychonomic Bulletin & Review*, *23*, 1341–1353. <https://dx.doi.org/10.3758/s13423-015-0985-2>
- Staddon, J. E. R., Jozefowicz, J., Cerutti, D. (2007). Metacognition: A problem not a process. *PsyCrit*, 1–5.
- Sutton, J. E., & Shettleworth, S. J. (2008). Memory without awareness: Pigeons do not show metamemory in delayed matching to sample. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*, 266–282. <http://dx.doi.org/10.1037/0097-7403.34.2.266>
- Templer, V. L., & Hampton, R. R. (2012). Rhesus monkeys (*Macaca mulatta*) show robust evidence for memory awareness across multiple generalization tests. *Animal Cognition*, *15*, 409–419. <https://dx.doi.org/10.1007/s10071-011-0468-4>
- Templer, V. L., Lee, K. A., & Preston, A. J. (2017). Rats know when they remember: Transfer of metacognitive responding across odor-based delayed match-to-sample tests. *Animal Cognition*, *20*, 891–906. <https://dx.doi.org/10.1007/s10071-017-1109-3>

- Terman, M., & Terman, J. S. (1972). Concurrent variation of response bias and sensitivity in an operant-psychophysical test. *Perception & Psychophysics*, *2*, 428–432. <https://dx.doi.org/10.3758/BF03206285>
- Tolman, E. C. (1927). A behaviorist's definition of consciousness. *Psychological Review*, *34*, 433–439. <https://dx.doi.org/10.1037/h0072254>
- Tolman, E. C. (1938). The determiners of behavior at a choice point. *Psychological Review*, *45*, 1–41. <https://dx.doi.org/10.1037/h0062733>