

ABSTRACT

The ability to reflect on one's own mental processes, termed metacognition, is a defining feature of human existence. Consequently, a fundamental question in comparative cognition is whether nonhuman animals have knowledge of their own cognitive states. Recent evidence suggests that people and nonhuman primates but not less "cognitively sophisticated" species are capable of metacognition. Here, we demonstrate for the first time that rats are capable of metacognition — i.e., they know when they do not know the answer in a duration-discrimination test. Before taking the duration test, rats were given the opportunity to decline the test. On other trials, they were not given the option to decline the test. Accurate performance on the duration test yielded a large reward, whereas inaccurate performance resulted in no reward. Declining a test yielded a small but guaranteed reward. If rats possess knowledge regarding whether they know the answer to the test, they would be expected to decline most frequently on difficult tests and show lowest accuracy on difficult tests that cannot be declined. Our data provide evidence for both predictions and suggest that a nonprimate has knowledge of its own cognitive state.

Metacognition: A Problem not a Process

"Metacognition" in animals can be explained by familiar learning principles

PsyCrit, April 13, 2007

Metacognition is defined in Foote and Crystal's paper by reference to phenomenal experience: "People are sometimes aware of their own cognitive processes. For example, a college student entering a classroom to take a test will often have some knowledge about how she will perform on the test. It is noteworthy that this knowledge (whether accurate or inaccurate) is available *before* the student actually responds to the test questions and obtains feedback about performance."




Figure 1. Schematic Representation of Experimental Design Showing Trial Sequence in Study, Choice, and Test Phases. After presentation of a white noise (2–8 s; study phase), a choice phase provided an opportunity for taking or declining a duration test; declining the test produced a guaranteed but smaller reward than was earned if the test was selected and answered correctly (test phase). To the extent that a rat is uncertain about how to classify the interval, it will (1) decline to take the test and (2) show lower accuracy when forced to take the test relative to trials on which it chose to take the test. The gray shading indicates an illuminated nose poke (NP) aperture, but smaller reward than was earned if the test was selected and answered correctly (test phase). To the extent that a rat is uncertain about how to classify the interval, it will (1) decline to take the test and (2) show lower accuracy when forced to take the test relative to trials on which it chose to take the test. The gray shading indicates an illuminated nosepoke (NP) aperture.

before the student actually responds to the test questions and obtains feedback about performance."

Animals cannot be queried in this way. They cannot be asked a question *for the first time*. We can train them on a task, but then they will have a past history with the task which they can react to in ways not available to the human student confronted with the task for the first time. Animal subjects need not examine their inner selves to judge their ability to answer the question. They know, from their past experience with the very same question. This procedural limitation makes it very difficult to be sure what we are testing for when we look for metacognition in animals, as this study shows.

The Experiment

In the Foote and Crystal experiment, rats were trained to discriminate between two classes of noise duration: short (4 stimuli, 2–3.62 s) – press the left (say) lever, or long (4 stimuli, 4.42–8 s) – press the right lever (this account is slightly simplified: Figure 1,

taken from the Foote & Crystal paper, gives more details). Correct responses got them 6 pellets.

On test trials the animals were allowed to choose between a third option, a weakly rewarded (3 pellets) sure thing, vs. one of the two test levers (a possible 6 pellets).

Obviously in training the rats did less well when responding to stimuli near the middle of the range. Perhaps unsurprisingly, therefore, they tended to choose the weakly rewarded alternative more frequently when confronted with such “difficult” choices.

The Explanation

The simplest explanation here is that because of the animals’ history with this procedure, the sample stimuli are associated with different payoff probabilities: choice in the presence of difficult stimuli (in the middle of the range) is associated with a low payoff, close to .5 (expected value $.5*6 = 3$ pellets), in the presence of easy stimuli, the expected value approaches $6*1 = 6$ (see Foote & Crystal, Figure 2). Thus, the animals should prefer the discrimination task on easy trials but approach indifference between the discrimination task and the sure thing on difficult ones.

There are two apparent obstacles to this simple account. First, given that the expected value of the discrimination task is no less than 3 pellets ($.5*6$), whereas the fixed alternative is always exactly 3 pellets, surely the rational rat should always pick the discrimination task? Yet some rats some of time chose the sure thing. Second: “Accuracy declined as the difficulty of the discrimination increased, *but this decline was greater when the rats were forced to take the test* compared to trials on which the rats chose to take the test. [our italics]” This finding – that animals do worse on forced trials than voluntary ones – is critical for the metacognition idea because it seems to require some knowledge of the animal’s own mental state.

The answer to the first obstacle is simply *risk aversion* with respect to reward amount. Pigeons do not seem to show it (Staddon & Innis, 1966) but rats often do (Real & Caraco, 1986; Kagel et al, 1983; Roche et al., 1997), which means that in a discrete-trial situations like the Foote & Crystal experiment many rats may be expected to prefer a certain alternative to an equal-expected-value risky one.

The second obstacle is more difficult. A behavioristic “out” is that the animals are responding not to their own mental state but just to their own behavior: “If forced choices had been surprising, we would expect longer latencies to respond in forced tests relative to choice tests. However, there was no evidence for a difference in the latency to respond with either a nose-poke or lever-press response.” This failure to find a latency difference on forced vs. free tests seems to rebut the idea that the animals were somehow reacting to their own (latency) behavior.

A Psychophysical Model

A simple psychophysical model overcomes the second obstacle and explains the other results of the Foote & Crystal experiment. The experiment involves two choice tasks, a *discrimination* task where the animal chooses based on whether it judges a given sample time as “long” or “short,” and *simple choice* task, where it chooses between the discrimination task and an alternative that yields a fixed, smaller alternative reward.

First, the discrimination task: The 8 blue diamonds in the middle of Figure 2 are the equally-spaced training times (2–3.62 s and 4.42–8 s, plotted on a log scale) and the associated reward size (6 pellets) for correct responses. The four on the left require one response (LEFT, say), the four on the right another (RIGHT). Because time perception is not perfect, each one of these 8 times is associated with its own Weber-law-obeying error distribution – the 8 thin, magenta normal distributions at the bottom of the figure. The sum of the 4 distributions for which a LEFT choice is appropriate is shown by the heavy red line; the sum of the 4 for which RIGHT is appropriate by the heavy green line. There is some overlap between the red and blue distributions, in the middle of the range. Thus, some errors are unavoidable.

Clearly, if we consider this task on its own, animals should choose LEFT for time-perception values to the left of the black vertical line, and RIGHT for values to the right. We

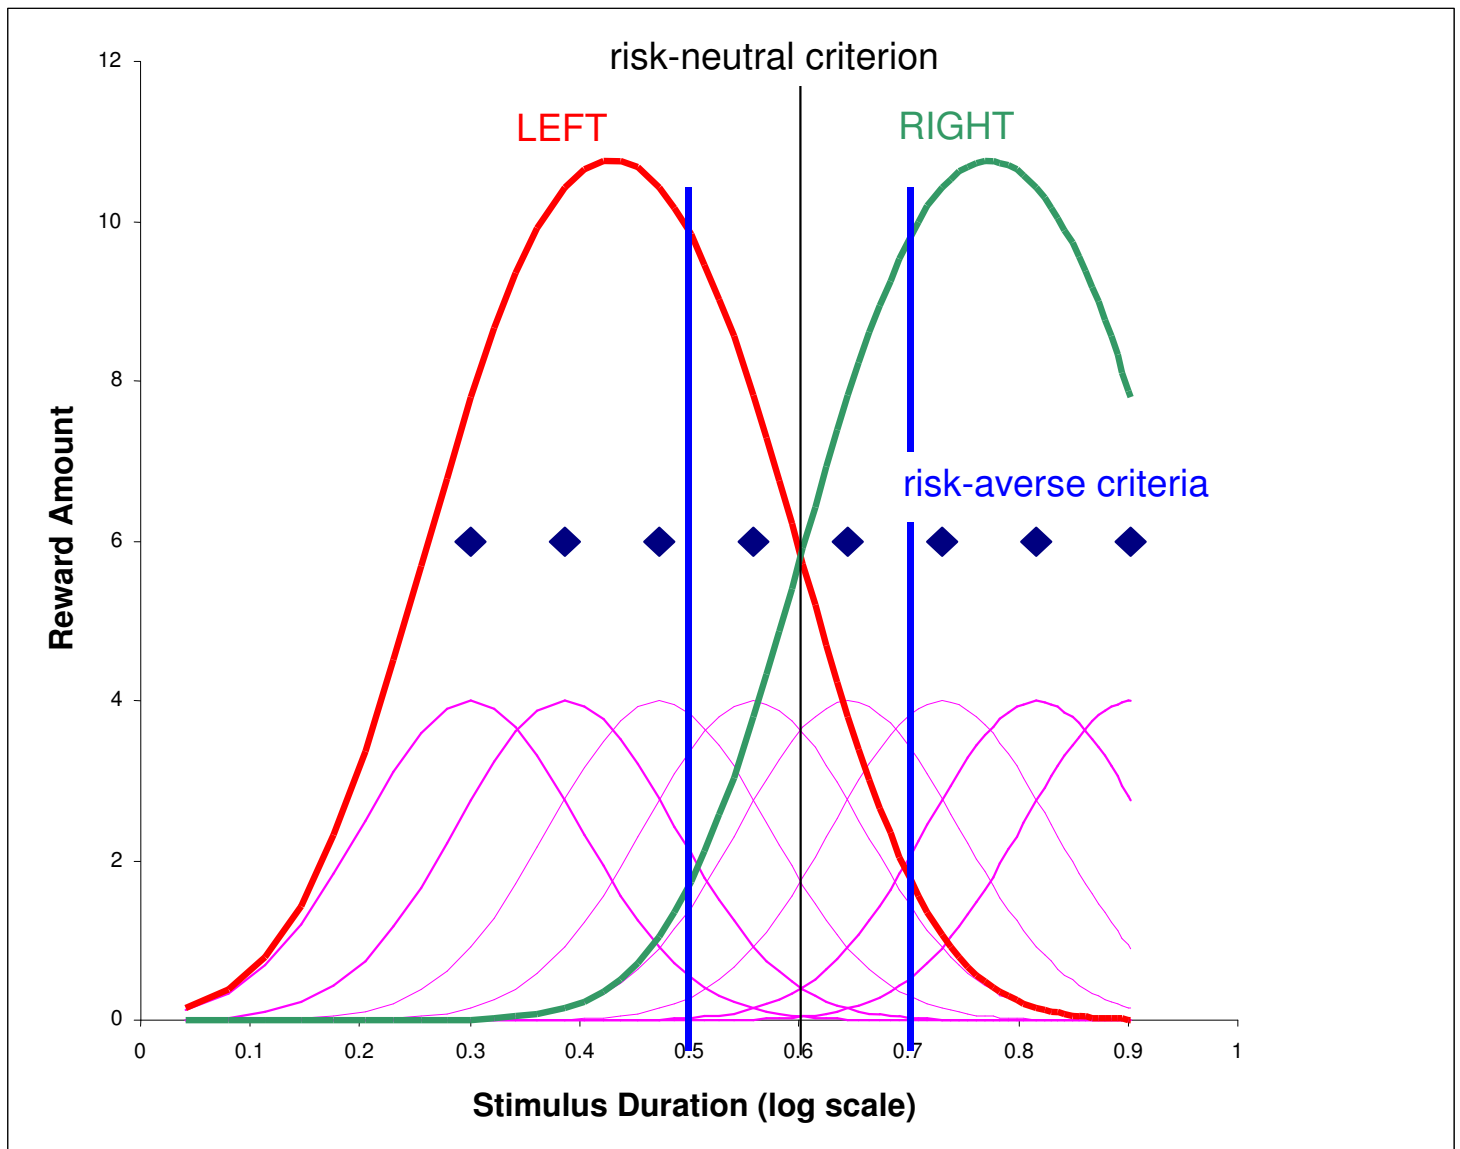


Figure 2 The discrimination task. A psychophysical analysis of the “metacognition” task. The 8 diamonds are the 8 noise durations, equally spaced on a log scale. Associated with each duration is a Weber-law-compliant error distribution (at the bottom of the figure). Correct response for the four stimuli on the left is LEFT and the sum of those four error distributions is the heavy red distribution – and similarly for the four on the right (heavy green distribution). Left and right choices in the absence of an alternative, or in the presence of the sure-thing 3-pellet choice under risk-neutrality, is guided by the black vertical line – the risk-neutral criterion i.e., always choose the discrimination task. But a risk-averse animal, offered the sure-thing alternative, may choose it when presented with stimuli in the middle of the range, i.e., when the perceived duration fall in between the two blue risk-averse criteria.

can call this line the *risk-neutral* choice criterion. It is all that is available to the animal in the absence of an alternative to the LEFT and RIGHT choices. The payoff probability under these conditions is shown in Figure 3.

Now for the simple choice task, which provides an alternative to the discrimination task. If the animal is risk averse, it may choose the smaller fixed alternative for time-perception regions where discrimination is difficult so that payoff is uncertain – in the middle, where the red and green distributions overlap (the dip in payoff probability in the middle region on the discrimination task is shown in Figure 3). By setting its criteria for switching to the certain option

in between the two blue risk-averse criterion lines in the figure, the animal will be almost 100% correct when it chooses the discrimination task. It will thus be more accurate when it chooses the discrimination task freely than when its choice is forced, just as Foote and Crystal, and other authors who have used similar tasks, report – and it will also only make this choice when presented with stimuli in the middle of the range: “choice accuracy was indeed less when the test was forced (no escape trials) compared to when it was declinable. *However, this effect was only apparent for one of the trial types (when the stimuli to be discriminated between were most similar) and not for the remaining three (easier) trial types.* [our italics]”

Our conclusion is that the animal in this task may be reacting simply to a learned association between a set of stimulus values and the probabilities of payoff. For stimuli 1-3 (on the left in Figure 2) and 6-8 (on the right), payoff probability is close to 100%, so the animal chooses the discrimination task and responds correctly most of the time. But a risk-averse animal, given a sure-thing option will choose it when presented with the two difficult middle stimuli, 4 and 5.

Tests

Metacognition not a process, not an explanation for anything, but just a name for a set of experimental results which need to be explained by some process model such as this. Thus, the most productive way to proceed is to test models for these data. We have proposed a simple one that seems to work well.

Apparently some animals (e.g., pigeons: Sutton & Shettleworth, 2007) cannot do the metacognition task (do no better in free- vs. forced-choice tests). But pigeons rarely show risk aversion for reward amount, so would not in fact be expected to perform like Foote and Crystal’s rats. Moreover, there should be a correlation in individuals rats between those that show risk-aversion in an independent test and those that show “metacognition” in the Foote-Crystal procedure.

Risk-sensitivity for amount can probably be manipulated by varying the range of outcomes: an animal that prefers a certain outcome of 3 vs. 50:50 outcomes of 0 and 6 may not be risk averse if the risky alternative is 50:50 2 vs. 4, for example. Such an animal should show “metacognition” in the Foote-Crystal procedure with the 6 vs. 0 discrimination choice but not with the 4 vs. 2.

Another way to test the plausibility of this account would be to train animals on its elements. Thus, a simple experiment could involve just two choices and a set of temporal stimuli, 2 – 6 s, say, associated with different payoff probabilities from 0 – 1. Responses on the right would pay off proportional to the duration of the temporal sample stimulus: in the presence of the 2-s stimulus, the payoff on the right would be 0, in the presence of 6-s, 1, and so on for intermediate values. On the left lever the payoff would always be .5. Would animals adjust their




Figure 3: Payoff probability as a function of stimulus duration for responses based on the risk-neutral criterion (black line). Payoff under the risk-averse criteria (blue lines) excludes the region in between the two lines where the sure-thing is chosen.

choices according to the duration of the temporal stimulus, responding on the left after a 2-s stimulus right after a 6-s? Would their choices change if the probability function is changed? If so, one element of the psychophysical model would be strengthened. Experiments with non-monotonic probability functions on the right lever could then be done to see if animals can set up two or more criteria as the present model requires.

Finally, this same account may apply with little change to familiar procedures such as temporal bisection (e.g., Church & Deluty, 1977; Machado & Arantes, 2006; Bizo & White, 1997).

J. E. R. Staddon
Staddon@psych.duke.edu

Jeremie Jozefowicz
jeremie@psyc.duke.edu

Dan Cerutti
Cerutti@psych.duke.edu

Duke University

(Research supported by grants from the National Institute of Mental Health to Duke University, JERS Principal Investigator.)

REFERENCES

- Bizo, L. A. & White, K. G. (1997). Training with controlled reinforcer density: Implications for models of timing. *Journal of Experimental Psychology: Animal Behavior Processes*, 23, 44-55.
- Kagel, J. H., Macdonald, D. N., Battalio R. C., White, S. & Green L. (1986) Risk aversion in rats (*Rattus norvegicus*) under varying levels of resource availability. *Journal of Comparative Psychology*, 100, n°2, 95-100.
- Machado, A. & Arantes, J. (2006) Further tests of the scalar expectancy theory (SET) and the learning to time (LET) model in a temporal bisection task. *Behavioural Processes*, 72, 195-206.
- Real, L. & Caraco, T. (1986) Risk and Foraging in Stochastic Environments. *Annual Review of Ecology and Systematics*, 7, pp. 371-39.
- Roche, J. P., Timberlake, W. & McCloud, C. (1997) Sensitivity to variability in food amount: Risk aversion is seen in discrete-choice, but not in free-choice, trials. *Behaviour* 134 (15-16), 1259-1272.
- Staddon, J. E. R., & Innis, N. K. (1966) Preference for fixed vs. variable amounts of reward. *Psychonomic Science*, 4, 193-194.
- Sutton, J. E. & Shettleworth, S. J. (2007) Memory without awareness: Pigeons do not show metamemory in delayed matching-to-sample. Submitted.