

COMPARATIVE COGNITION & BEHAVIOR REVIEWS

2009

Volume 4, pp 29 -39

Metacognition in animals: how do we know that they know?

J. Jozefowicz
Universidade do Minho

J. E. R. Staddon
Duke University

D. T. Cerutti
California State University-East Bay

Research on animal metacognition has typically used choice discriminations whose difficulty can be varied. Animals are given some opportunity to escape the discrimination task by emitting a so-called uncertain response. The usual claim is that an animal possesses metacognition if (a) the probability of picking the uncertain response increases with task difficulty, and (b) animals are more accurate on “free-choice” trials —i.e., trials where the uncertain response was available but was not chosen— than on “forced-choice” trials, where the *uncertain response* is unavailable. We describe a simple behavioral economic model (BEM), based on familiar learning principles, and thus lacking any metacognition construct, which is able to meet both criteria in most of these tasks. We conclude that rather than designing ever more complex experiments to identify “metacognition,” a necessarily ill-defined concept, knowledge might better be advanced not by further refining behavioral criteria for the concept, but by the development and testing of theoretical models for the clever behavior that many animals show in these experiments.

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

Can you recall what you did yesterday around 2 PM? Probably yes. Can you recall what you did 10 years ago around 2 PM? Probably no. You did not actually need to retrieve any information to answer these questions: you knew immediately that the first question could be answered but not the second. This is an example of what has come to be called

metacognition: the ability to judge one’s chances of success or failure at a cognitive task before actually carrying it out. More abstractly, metacognition is the ability to perceive one’s own mental states and cognitive processes (Metcalf & Kober, 2005; Metcalfe, in press). Some (e.g. Nelson & Narens, 1990), consider metacognition to be a higher level of cognitive functioning, monitoring and regulating lower-level cognitive processes via self-awareness and consciousness (Nelson, 1996).

Metacognition is a concept that arose from contemplation of our own subjective, phenomenal experience but can it be found in other species? Because the verbal methods used to investigate metacognition in humans cannot be used with animals, researchers have come up with behavioral criteria for metacognition. Smith, Shields & Washburn (20083, for example (see also Sutton & Shettleworth, 2008), claim that given a discrimination task whose difficulty can be controlled by the experimenter, the animal is considered to show

J. Jozefowicz, Instituto de Educação e Psicologia, Universidade do Minho, Braga, Portugal; J. E. R. Staddon, Department of Psychology and Neuroscience, Duke University, Durham, North Carolina 27708; D.T. Cerutti, Department of Psychology, California State University - East Bay, Hayward, California. Part of this work was done while Jeremie Jozefowicz was a post-doctoral fellow in the laboratory of Dr. Ralph Miller at SUNY-Binghamton and, as such, supported by a grant from National Institute of Mental Health to Binghamton University (MH 033881). Correspondence concerning this article should be addressed to Jeremie Jozefowicz, Instituto de Educação e Psicologia, Universidade do Minho, 4710 Braga, Portugal. E-mail: jeremie@iep.uminho.pt.

metacognition if, (a) it is more likely to avoid the task (i.e., of emitting a so-called “uncertain” response) on trials where the discrimination is difficult; and (b) is more accurate on trials where the “uncertain” response is available and the discrimination task can be avoided (i.e., when allowed to choose between the uncertain response and the discrimination task) than on “forced-choice” trials (i.e., where the uncertain response is not available). Using these criteria, researchers have concluded that rhesus monkeys (e.g. Hampton, 2001; Shields, Smith, & Washburn, 1997; Smith, Shields, Allendoerfer, & Washburn, 1998; Beran, Redford, & Washburn, 2006; Washburn, Smith, & Shields, 2006), dolphins (Smith et al., 1995), rats (Foote & Crystal, 2007) but not, apparently, pigeons (Sutton & Shettleworth, 2008; Sole, Shettleworth,

& Bennett, 2003 but see Inman & Shettleworth, 1999) have metacognition (see a review in Smith, et al., 2003). But are these two criteria indeed sufficient to conclude that an animal possesses metacognition, i.e., to rule out simpler explanations, based on familiar learning principles?

A behavioral economic model (BEM) of choice

We begin by presenting briefly a simple discrimination model based on some basic learning principles. This model, nicknamed BEM (for Behavioral Economic Model), is shown in Figure 1. It just assumes that when confronted with a stimulus, the subject emits the behavior which is associated with the higher payoff. The only other assumption is that perception of the stimulus is noisy.

A full description of BEM is presented in Jozefowicz, Staddon, and Cerutti (in press). We now illustrate how it works in a simple discrimination task, where response R1 is reinforced after stimulus S1 while response R2 is reinforced after stimulus S2. We will denote by I1 and I2 the intensity of respectively stimuli S1 and S2 and we will assume that both responses are reinforced with the same amount, A, of reinforcer.

How could the animal solve this task? If it had a noiseless perception of the objective intensity of the stimuli, it would learn that the payoff for emitting R1 when the stimulus intensity is I1 is A units of reinforcer while it is 0 when the stimulus intensity is I2. It would learn that the payoff function for R2 is exactly the reverse. It would then be able to derive from this knowledge of the payoff functions an optimal policy: basically, since R1 has a higher payoff than R2 when the stimulus intensity is I1 and vice versa when the stimulus intensity is I2, emit R1 when the stimulus intensity is I1, R2 otherwise. This is basically BEM except that we add the additional assumption that animals’ perception of stimulus intensity is not noiseless, but on the contrary, noisy.

Stimulus S1 has an objective intensity of I1, but because of noise in the sensory system, its subjective intensity is a random variable following a Gaussian distribution with mean $\ln I1$ and constant standard deviation σ , respecting the Weber-Fechner law. The model was developed initially to account for experiments on interval timing, hence its emphasis on Weber’s law. Yet, it is not fundamental to our account of experiments on animal metacognition, at least as far as qualitative predictions are concerned: other random distributions could also be used.

Suppose that a stimulus is presented to the subject. If the stimulus is S1, it should emit R1. If it is S2, it should emit R2. But, the subject has no way of knowing for sure what

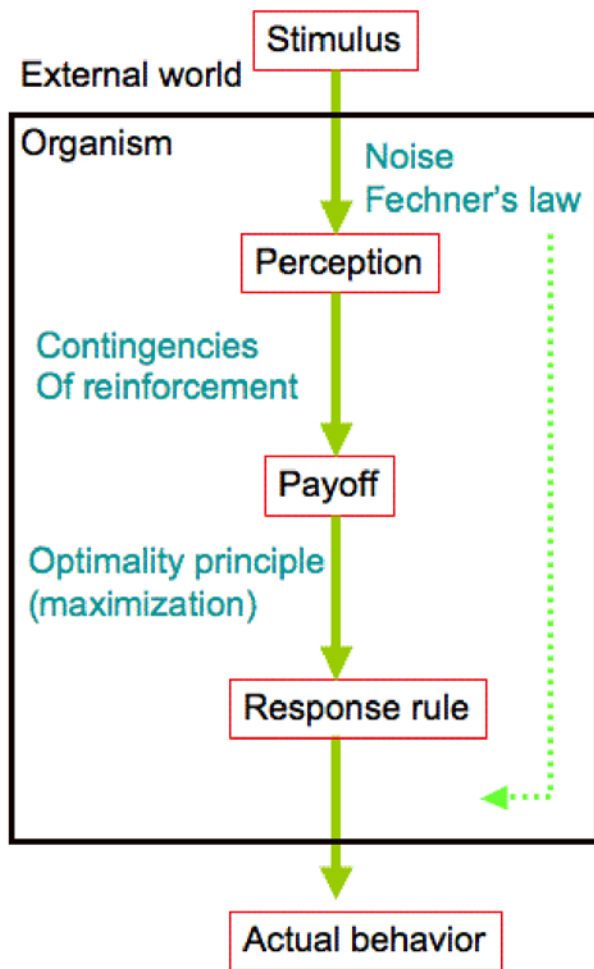


Figure 1. BEM at a glance: (a) based on its perception of the stimulus, the animal emits the behavior which leads to the higher payoff; (b) but its perception is noisy, following the Weber-Fechner law. At objective time t , the representation of the test stimulus (e.g., a duration) is a random variable drawn from a Gaussian distribution with a mean equal to $\ln t$ and a constant standard deviation. From *The Behavioral Economics of Choice and Interval Timing*, by J. Jozefowicz, J.E.R. Staddon, and D.T. Cerutti, in press, *Psychological Review*. Reprinted with permission.

stimulus has actually been presented: it has access only to its subjective perception of the stimulus intensity, which we denote by x . The expected payoff for emitting R1 when the subjective intensity of the stimulus is x , $Q1(x)$, is therefore

$$Q1(x) = P(S1|x)P1A \quad (1)$$

where $P(S1|x)$ is the probability that stimulus S1 has been presented given that the subjective intensity of the stimulus is x , which is given by Bayes' theorem

$$P(S1|x) = \frac{P(S1)P(x|S1)}{P(x)} \quad (2)$$

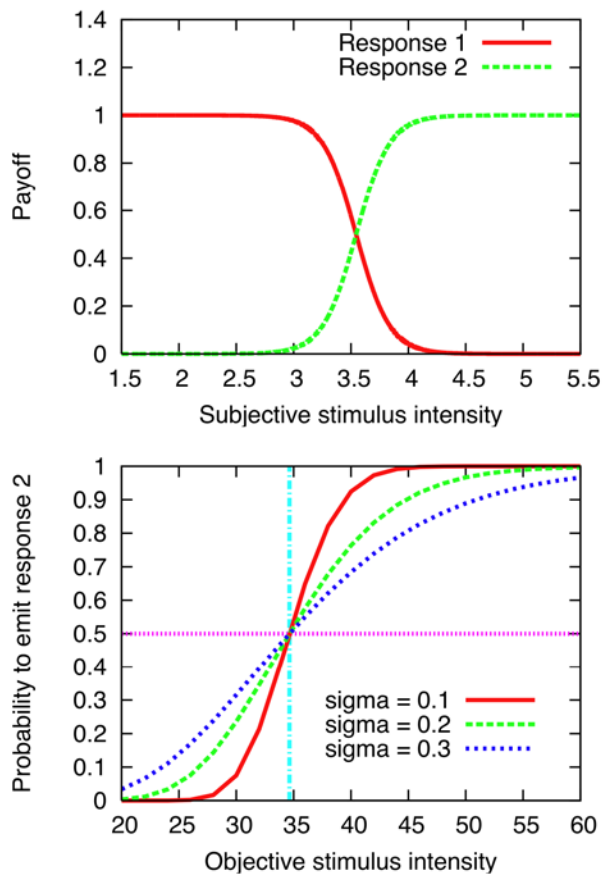


Figure 2. Simulation of a discrimination task by BEM. Response 1 is reinforced after stimulus S1 (intensity = 20); response 2 after stimulus S2 (intensity = 60). Top panel: Payoff function for each response as a function of the subjective stimulus intensity. Bottom panel: Probability of emitting response 2 as a function of the objective stimulus intensity for various values of standard deviation (σ) d . From *The Behavioral Economics of Choice and Interval Timing*, by J. Jozefowicz, J.E.R. Staddon, and D.T.Cerutti, in press, *Psychological Review*. Reprinted with permission.

$P(S1)$, the probability that stimulus S1 is presented on a trial, is a variable under the control of the experimenter. $P(x) = P(S1)P(x|S1) + P(S2)P(x|S2)$ while, in this case $F(x, \ln I, \sigma)$ ($F(x, m, d)$ being the density function for a Gaussian distribution with mean m and standard deviation d) can be substituted for $P(x|Si)$ in equation (2) (see Jozefowicz et al., in press, for a more rigorous treatment). The equation for $Q2(x)$, the payoff for emitting R2 when the perceived stimulus intensity is x , can be deduced from the above equations.

The top panel of Figure 2 shows the payoff functions for both responses R1 and R2. BEM assumes that the subject follows a simple maximization response rule based on these functions: emit R1 if the payoff for that response is higher than the payoff for R2, otherwise emit R2. This policy maps subjective stimulus intensity on to response probabilities. To predict behavior, we need a policy that maps objective stimulus intensity on to response probabilities — which means taking into account the random nature of the relation between objective and subjective stimulus intensity. Suppose a stimulus S with intensity I is shown to the subject. It could be S1, S2 or a new stimulus the subject has never encountered before. The subjective intensity of that stimulus will be a random variable with mean $\ln I$ and standard deviation σ . Let $p2(x)$ be the probability of emitting response R2 when the subjective stimulus intensity is equal to x . Then $P2(I)$, the probability of emitting response R2 when the objective stimulus intensity is I is equal to

$$P2(I) = \int_{-\infty}^{+\infty} p2(x)P(x|I)dx \quad (3)$$

The bottom panel of Figure 2 shows what $P2(I)$ looks like in this example.

A behavioral account of animal metacognition

BEM is obviously a very general model, which can be applied to many experimental situations. Indeed, we initially developed it to account for data showing an interaction between interval timing and reinforcement (Jozefowicz et al., in press). It assumes nothing beyond basic discrimination processes. Since it uses a strictly deterministic response rule, there is no room for uncertainty in the model: A stimulus is always categorized as belonging to one category or the other. The question of interest here is whether it can account for data on animal metacognition (see Staddon, Jozefowicz, & Cerutti, 2007, for an earlier treatment).

The first type of task used to demonstrate metacognition in animals is categorization. Categorization tasks use stimuli varying along a stimulus continuum: all stimuli below a critical value are associated with one response, all stimuli above that value are associated with the other response. The task is obviously harder with stimuli close to the critical value.

Indeed, in dolphins (with tones Smith et al., 1995), rhesus monkeys (with both stimuli differing in terms of pixel density, Shields et al., 1997), or number of elements (Washburn et al., 2006), rats (with stimuli differing in terms of their duration, Foote & Crystal, 2007) and pigeons (with stimuli differing in terms of pixel density, Sole et al., 2003), the probability of picking the uncertain response increases the closer the stimulus is to the critical value. With pigeons the sole exception, all species so far tested are more accurate in their categorization on free-choice trials than on forced-choice ones.

Can BEM account for these results? We applied it to the task used by Foote and Crystal (2007) since the amounts

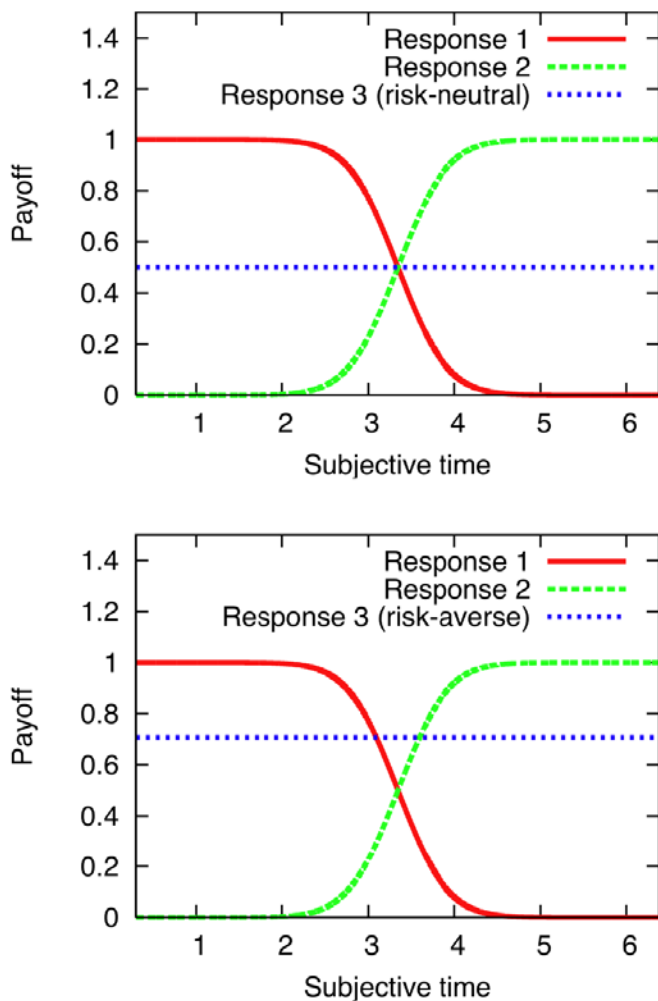


Figure 3. Payoff functions in a simulation of Foote and Crystal (2007). Top panel: the subjective reward magnitude is equal to the objective reward magnitude. Bottom panel: the subjective reward magnitude is a power function of the objective reward magnitude with exponent, c , smaller than 1 to account for risk aversion. From *The Behavioral Economics of Choice and Interval Timing*, by J. Jozefowicz, J.E.R. Staddon, and D.T. Cerutti, in press, *Psychological Review*. Reprinted with permission.

of reinforcer for each response, including the uncertain response, are clearly identified in that study. Foote and Crystal (2007) showed rats duration stimuli, evenly spaced (on a log scale) between 2 and 8 s. If the stimulus duration was less than 4 s, one response (R1) was reinforced while if more than 4 s another response (R2) was. On some trials (free-choice trials), the animals had the opportunity of picking a third response (the *uncertain* response) which was reinforced no matter the stimulus duration, but with only half the amount of reinforcer that the animal could obtain in case of an accurate categorization response. On the other hand, it would get zero reward in case of a wrong categorization.

The top panel of Figure 3 shows the payoff functions for R1, R2 and R3. As you can see, the payoff functions for R1 and R2 are always above the ones for R3. Hence, according to the model, the animals should never pick the uncertain response. But this is because we have assumed that the objective amount of reinforcer an animal receives and the subjective amount it experiences are the same. This is not a valid assumption as it does not take into account the well-established fact of risk sensitivity: when given the choice between an alternative delivering a fixed amount of reinforcer (say, 2 units) and one delivering a variable amount (say, either 1 unit or 3) which, on average, is equal to the amount delivered by the fixed alternative, many animals prefer the fixed alternative (Bateson & Kacelnik, 1995; Kacelnik & Bateson, 1996; Roche, Timberlake, & McCloud, 1997; Staddon & Innis, 1966). This is *risk-aversion*. The reverse pattern is called *risk-proneness*¹.

The usual way to explain risk aversion is to assume that the subjective reward function, which maps objective amount of reward collected on to subjective amount experienced, is negatively accelerated, following the principle of diminishing marginal value. To incorporate this in BEM, we need, in equation (1), to substitute A , the objective amount of reward collected with A^c , the subjective amount of reward experienced. c is a free-parameter representing risk-sensitivity: if $c = 1$, the animal is risk-neutral; if $c < 1$, the animal is risk-averse; if $c > 1$, the animal is risk-prone.

The bottom panel of Figure 3 shows the payoff functions for the three responses once risk sensitivity is taken into account. As can be seen, there are now some subjective stimulus intensity values for which the payoff function for the “uncertain” (certain-outcome) response exceeds the (*uncertain*) choice.

Figure 4 shows a quantitative fit of the model to the data of Foote and Crystal (2007). To obtain those fits, we first adjusted σ so as to predict accuracy on the forced-choice trials as well as possible by minimizing square error between the model and the data. The predictions of the model in the

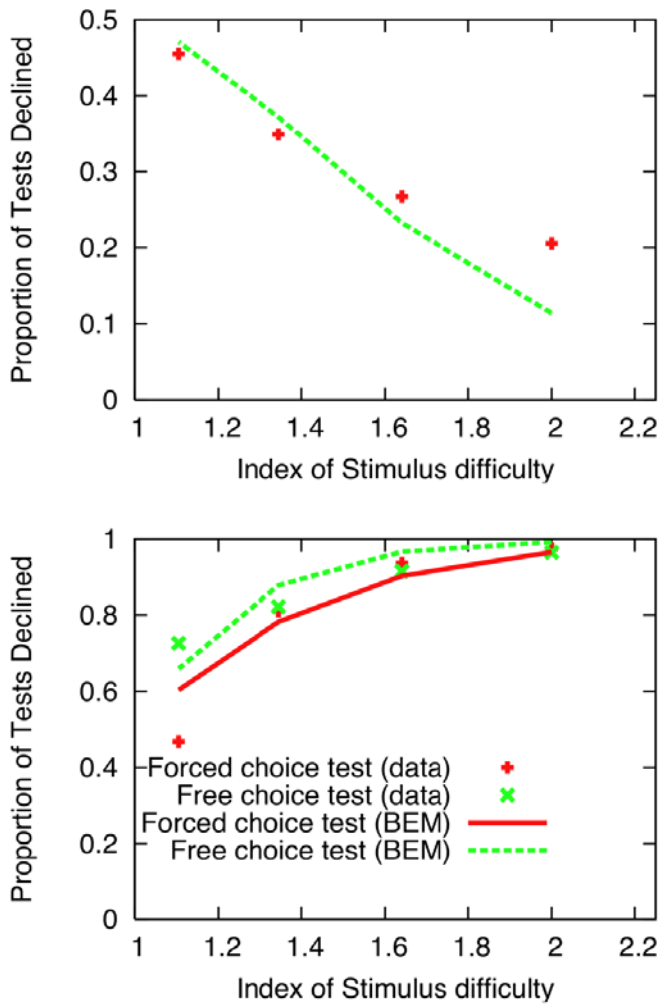


Figure 4. Top panel: Probability of declining a test (that is to say, of selecting the uncertain response) as a function of the index of stimulus difficulty used by Foote and Crystal (2007). (Since this index represents the distance to the boundary between the stimulus classes, stimuli on the fringe of the stimulus range, hence easier to discriminate, have a higher index of stimulus difficulty.) The points are data from Foote and Crystal (2007) while the line is the prediction from BEM. Bottom panel: accuracy in the forced choice (2 responses available) and free-choice (3 responses available) trials as a function of the index of stimulus difficulty. The points are the data from Foote and Crystal (2007), the lines are the predictions from BEM. $d = 0.38$, $c = 0.46$. From *The Behavioral Economics of Choice and Interval Timing*, by J. Jozefowicz, J.E.R. Staddon, and D.T. Cerutti, in press, *Psychological Review*. Reprinted with permission.

force-choice trials are not affected by the risk-sensitivity parameter c (below). Then, we adjusted c so as to predict as well as possible (minimization of the square error between the data and the model's prediction) the proportion of tests declined by the rats—that is to say, the proportion of free-choice trials where the rats chose the uncertain response. The fit to “tests declined” is adequate (although BEM un-

derestimates the proportion of trials where the rats should decline the test for the stimuli at the ends of the stimulus range). But, the model in this case predicts that accuracy should have been higher for all stimulus durations when the weakly reinforced sure-thing response was available, while this was observed only for the most difficult test in the Foote and Crystal (2007)'s data. BEM shows this limitation in some of the simulations we ran but not with the parameters necessary to best fit Foote and Crystal's (2007) data.

But this is irrelevant to our main point. Beyond the quantitative fit, it is important to note that BEM, which lacks any metacognitive ability—only basic discrimination processes—satisfies the two generally accepted criteria for metacognition: that the probability of picking the uncertain response increases with the difficulty of the task (Figure 4, top panel); and that the subject is more accurate on free-choice trials than on forced-choice trials (Figure 4, bottom panel). (Indeed, paradoxically, the model shows rather better “metacognition” than the rats, since it is less accurate on all forced trials, not just those in the middle of the range.)

This account of animal metacognition experiments is similar to the one recently proposed by Smith, Beran, Couchman, and Coutinho (2008). Those authors proposed that animals map subjective stimulus intensity on to response strength. For a stimulus of intensity I_1 , the response strength for the response reinforced in presence of that stimulus is maximum at the point on the subjective stimulus intensity continuum corresponding to I_1 and follows an exponentially decaying (as opposed to Gaussian) generalization gradient to the left and right of that value. On the other hand, since the uncertain response is reinforced no matter what the stimulus intensity, its response strength is assumed to be constant across the stimulus continuum. When a given stimulus is presented, the response with the higher response strength is emitted.

The Smith et al. (2008) model parallels BEM. They both lead to a very similar conceptualization of the problem but BEM derives it from a basic optimization analysis of operant behavior while several assumptions in the Smith et al. (2008) model are more specific to the metacognition paradigm. For instance, while BEM uses risk sensitivity to explain why the payoff function of the uncertain response is sometimes below the payoff functions of the two other responses, no such reason is given as for why the response strength of the uncertain response is below the response strength of the two other responses in the Smith et al. (2008)'s model. BEM is also much easier to use as it does not require the extensive simulation work Smith et al. (2008) had to run in order to get predictions from their model. But, overall, the general philosophy behind the two models is the same.

BEM is also able to account for data showing that animals

are able to generalize the use of the uncertain response to new tasks (e.g., Washburn et al., 2006). As long as the task requires a discrimination along the same stimulus dimension as the one used in the task where the uncertain response was initially trained, the subject should still be able to compare the payoff for the uncertain response to the payoff for the responses in the new task. Depending on the amount of generalization between stimulus dimensions, the model would also be able to account for generalization of the use of the uncertain response to tasks employing a stimulus dimension different from the one used initially to train the uncertain response.

Metacognition has also been investigated in animals using a delayed-matching-to-sample (DMTS) procedure. In this procedure, the animal is shown a sample stimulus. After a subsequent retention interval during which the sample is absent, the animal is asked to choose between two responses R1 or R2. Which is reinforced depends on the identity of the sample. An “uncertain” response is available on some choice trials, on others a retention choice is forced.

The longer the retention interval, the less accurate the animal becomes, presumably because of the decay of the short-term memory (STM) trace of the sample stimulus. In such a task, the probability of picking the uncertain response increases with the retention interval in both rhesus monkeys (Hampton, 2001) and pigeons (Inman & Shettleworth, 1999; Sutton & Shettleworth, 2008). But, as in other studies, pigeons’ accuracy is no better on free-choice trials as compared to forced-choice ones.

Although BEM is principally a model of discrimination, it can be extended to DMTS by borrowing from White and Wixted (1999) the idea that forgetting in STM is more a discrimination problem than a memory one. According to this view, the value of a stimulus is represented in STM by a Gaussian distribution whose standard deviation increases with the time since the stimulus presentation ($\sigma = d_0 + kD$, where d_0 is the standard deviation of the Gaussian distribution for a retention interval of 0, D is the retention interval and k is a free parameter). As retention interval increases, the distribution widens, making it more difficult for the animal to correctly assign a choice response to the presented stimulus. Hence, in this view, the difference between the categorization tasks discussed previously (where the animal is asked to make a choice when the sample is still present in the environment) and the DMTS (where the animal is asked to make a choice when the sample is no longer present in the environment) is purely in the eye of the beholder: the underlying processes are exactly the same. White and Wixted (1999) have shown that this approach can account for such forgetting phenomena as the power forgetting function and the fact that animals actually perform worse if they are tested

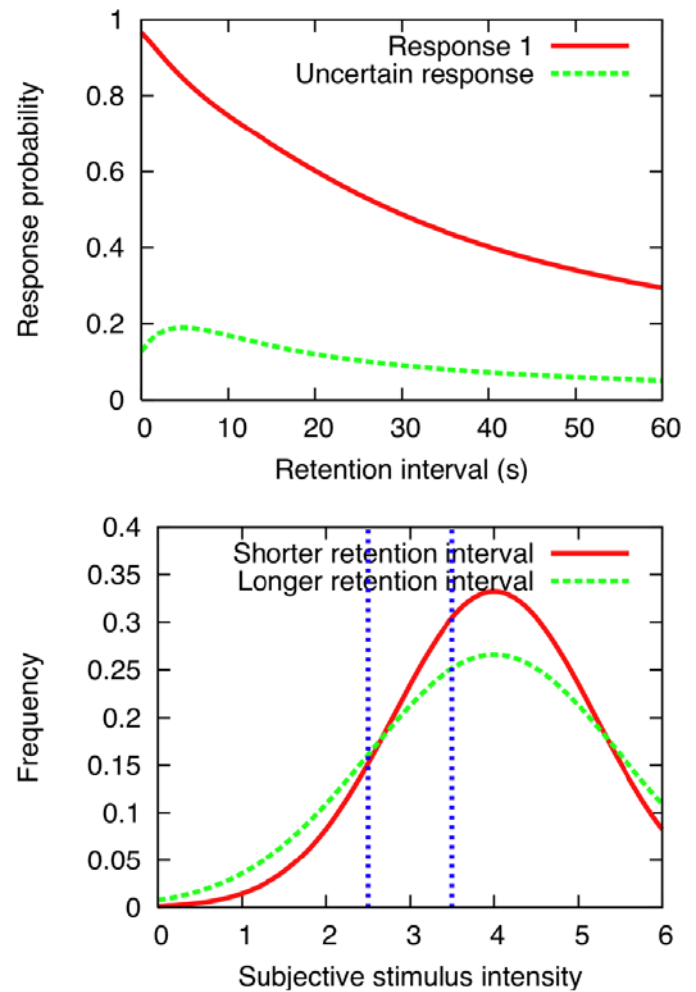


Figure 5. Top panel: Simulation of a delayed matching-to-sample task by BEM. Response 1 is reinforced after stimulus S1 while response 2 is reinforced after stimulus S2. The uncertain response is reinforced after both stimuli but with only half the amount of reinforcer that could be obtained by emitting response 1 or 2. The graph shows the probability on forced-choice trials (the uncertain response is not available) of picking response 1 after stimulus S1 has been presented as a function of the retention interval as well as the probability of pickings the uncertain response on free-choice trials, also as a function of the retention interval. Bottom panel: Gaussian distribution of the subjective stimulus intensity for two retention intervals. The probability of emitting the uncertain response corresponds to the area under the curve between the two horizontal lines.

with a retention interval shorter than the one for which they have been trained.

We simulated a simple DMTS task with BEM: response R1 is reinforced after stimulus S1 while response R2 is reinforced after stimulus S2; the uncertain response is reinforced after both stimuli but with only half the amount of reinforcer that could be obtained by emitting R1 or R2; the retention

interval during training is 0 s.

The results of the simulations are shown in Figure 5, which plots the probability of picking R1 on forced-choice trials after S1 has been presented as a function of the retention interval and the probability of picking the uncertain response after S1 has been presented, also as a function of the retention interval. As can be seen, the forgetting curve for R1 is the familiar power function that has been shown experimentally in many studies (Staddon, 2001) but the predictions of the model about the probability of picking the uncertain response are at odds with the data: instead of increasing with the retention interval, as in the data by Hampton (2001), the probability function is non-monotonic, first increasing, then decreasing with the retention interval.

The reasons for this are shown in the bottom panel of Figure 5. The payoff functions divide the subjective stimulus intensity into 3 areas. The uncertain response is emitted when the subjective stimulus intensity falls in the center area (between the two vertical lines shown in Figure 5) where the payoff for the uncertain response is higher than the payoff for either response R1 or R2. When the stimulus is shown, its subjective value is drawn randomly from a Gaussian distribution whose standard deviation increases with the retention interval. The probability of emitting the uncertain response corresponds to the area of that Gaussian function falling into the uncertain response decision area. The bottom panel of Figure 5 shows the Gaussian distribution for 2 retention intervals. Because when the standard deviation is increased, the peak of the distribution decreases, the area under the curve in the uncertain-response decision area is obviously smaller for the longer retention interval than for the shorter one. The model does not fare any better if it is given explicit training with all the retention intervals (actually, since in some cases, when the retention intervals are spread far apart, the payoff function becomes non-monotonic, the problem gets worse).

This failure is partly due to the fact that we allowed the model to make its decision based only on the subjective stimulus intensity, so that the criteria (vertical lines) are the same for all delays. But it is likely that animals in a DMTS also use other kinds of information, notably, given the ubiquity of interval timing in appetitive learning procedures (e.g., Wynne & Staddon, 1988), the duration of the retention interval itself. Indeed, the fact that, in a DMTS task, pigeons (Sutton & Shettleworth, 2008) increase their probability of picking the uncertain response without increasing their accuracy is usually explained by the fact that their behavior is under the control of the duration of the retention interval since taking the test after a long retention interval is usually correlated with a low payoff (the same pattern is observed in one of Hampton, 2001's monkeys).

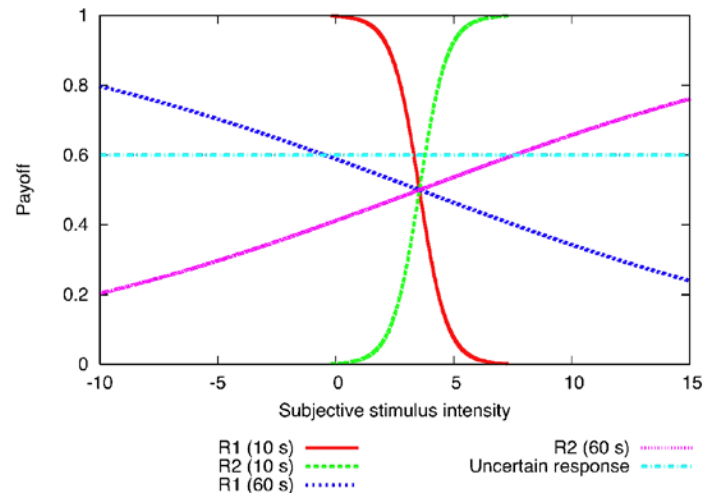


Figure 6. Payoff function in a delayed matching-to-sample if the subject is allowed to make its decision based on both subjective stimulus intensity and subjective time. In this case, the payoff functions for responses R1 and R2 are not the same after a 10-s retention interval then after a 60-s retention interval while the payoff function for the uncertain response remains identical.

Hence, a more realistic model would have the payoff function mapping both subjective stimulus intensity and subjective time since the sample offset onto reward expectation instead of simply subjective stimulus intensity on reward expectation. Also, to stay coherent, subjective time should be a random variable following Weber's law. This later constraint makes the model a little bit more computationally intensive on the simulation side and so, in this paper, we used instead a simpler version which assumes that the animal has a noiseless linear representation of time. Although the representation of time is definitely noisy and, in our opinion logarithmic, this simple model is sufficient to make our point here. We will postpone the exploration of the more complete model (which should allow predictions about the effect of the spacing of the retention intervals) to a future paper.

Figure 6 illustrates how the model works with this additional assumption according to which the decision of the animal is controlled not only by the subjective stimulus intensity but also by the retention interval durations. Because it is able to discriminate between retention intervals, the animal is able to have different payoff functions for each retention interval. As Figure 6 shows, because the payoff function for the uncertain response is constant across subjective stimulus intensity and subjective time, the region of subjective stimulus intensity for which the uncertain response has a higher payoff than the two responses increases with the retention interval. As the top panel of Figure 7 shows, this leads to an increased choice of the uncertain response as retention interval increases. Moreover, the lower panel of Figure 7 shows that the model is more accurate on free-choice trials than

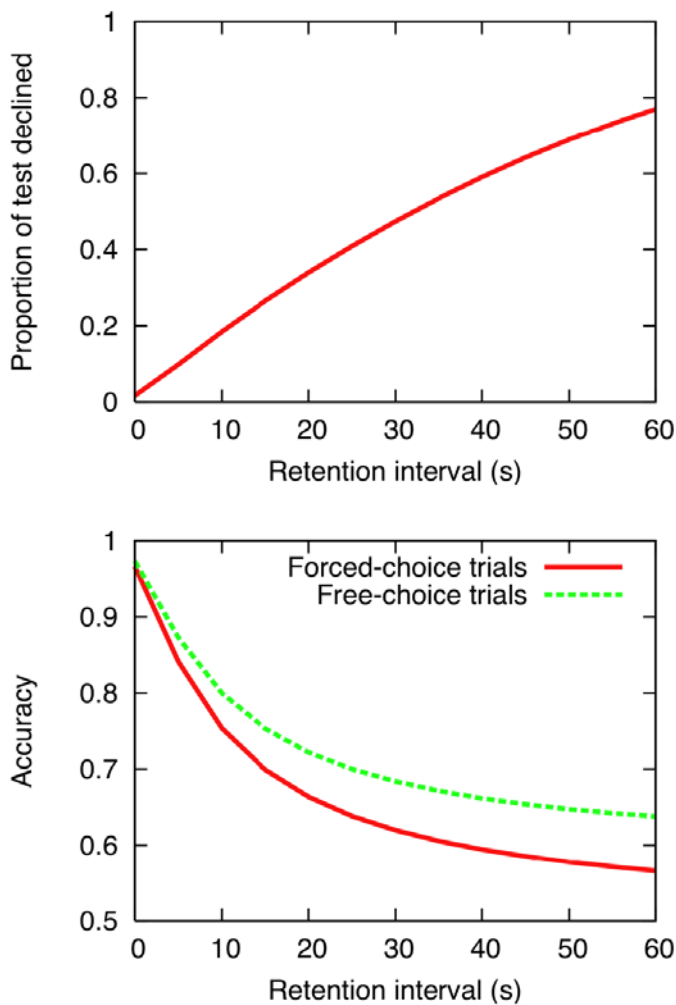


Figure 7. *Top panel: Probability of picking the uncertain response in a delayed matching-to-sample if the subject is allowed to make its decision based on both subjective stimulus intensity and subjective time since sample offset. Bottom panel: Accuracy in forced and free-choice trials in a delayed matching-to-sample if the subject is allowed to make its decision based on both subjective stimulus intensity and subjective time since sample offset.*

on forced-choice ones. Once again, despite its lack of any construct for metacognition, the predictions of BEM satisfy both the conventional behavioral criteria for metacognition.

One could object that this explanation cannot account for Hampton (2001)'s data because the monkeys received only 100 trials with several retention intervals, not enough to allow them to learn new criteria for each of the various retention intervals. But both monkeys had received previously extensive training with two retention intervals. Even if 100 trials was not enough for them to directly learn new criteria, simple generalization could explain how the monkeys would have been able, based on this training, to extrapolate new criterion at new retention intervals. Moreover, the fact that

100 trials is not enough for the monkeys to learn new criteria is not clear. Indeed, one of the monkey showed no improvement in its accuracy in free-choice trials, indicating that its behavior was only under the control of the duration of the retention interval. This proves that temporal learning, eventually boosted up by generalization, proceeded fast enough for our explanation of Hampton (2001)'s data to be plausible².

Finally, note that the view that choice is based on several stimulus dimensions allows us to explain why some animals (like pigeons) do not meet the criteria for metacognition. Criteria for metacognition will be met only if the behavior is at least partially under the control of a stimulus dimension correlated with the subject's chance of success on the task (as it is the case of the subjective stimulus dimension in BEM). If that dimension is overshadowed by other more salient ones (as the temporal dimension certainly is for pigeons), the criteria for metacognition will not be met. This seems more satisfying than saying that those animals do not "have metacognition" which fails to account for the choice pattern observed. After all, if the animals did not have metacognition, they should never pick the uncertain response or pick it with a low probability that would remain constant no matter the difficulty of the task, not increase their probability of picking the uncertain response as the difficulty of the task is increased. Nor does "metacognition" account for the fact that, as one monkey in Hampton (2001)'s study showed, some animals meet the criteria in one task but then suddenly fail to meet them in another.

Conclusion

We have described a simple behavioral model of choice (BEM) and showed that it is able to account for data supposed to demonstrate animal metacognition in categorization tasks and delayed matching-to-sample. If BEM, a model lacking any such construct, predicts that the probability of picking the uncertain response increases with the difficulty of the task, and that the subject is more accurate on free-choice trials than on forced-choice trials, then maybe animals also lack the faculty of metacognition. Indeed, Occam's razor almost compels that conclusion.

Another possible conclusion, for believers in metacognition, is that the usual behavioral criteria for metacognition are inadequate. Should we craft new behavioral criteria, extending the list that an animal must fulfill before we can declare it has metacognition? Recognizing the limits of the current criteria, several researchers have proposed just that. Metcalfe (in press) has proposed that, besides the two criteria described in this paper, the animal must based its decision on an internal representation (as in a DMTS task) instead of on its perception of a stimulus present in the environment (as in Foote & Crystal, 2007). But, in a framework like BEM,

there is no real difference between these two kinds of task. Smith et al. (2008) have added that the uncertainty response must not be explicitly reinforced. This is rather puzzling: the controversy between models like BEM, which does not need the concept of metacognition, and models that require it is about the information they use to make their decisions. A point which should not be controversial is that, in the end, the reinforcement contingencies explain why animals make these decisions at all. As a thought-experiment, imagine a hypothetical organism having real, genuine metacognition. Even such an organism would never pick the uncertain response if the payoff for doing so is less than what it would get by responding randomly to the test. In other words, if you are given the choice between taking a test, and winning either 100 dollars if you are correct or 0 dollar if you are not, or declining the test and not getting anything, then you have nothing to lose in taking the test, even if you don't know the answer and if, having metacognition, you know that you don't. Hence, not only is reinforcing the uncertain response necessary if we want the animal to pick it, adding reinforcement does not compromise the procedure as far as metacognition is concerned.

Some researchers have started to develop new procedures altogether, allowing the animal subject to ask for further information (e.g. Hampton, Zivin, & Murray, 2004) or to make confidence judgements (e.g. Shields, Smith, Guttmanova, & Washburn, 2005; Son & Kornell, 2005; Kornell, Son, & Terrace, 2007). We have not yet applied BEM to those tasks and it is very possible that it might fail. For instance, as we said, BEM can account for generalization of the use of the uncertain response if, without further training, there is some ground for stimulus generalization between the task where the uncertain response was initially trained and the task where it is introduced. BEM may therefore have difficulty in cases where stimulus generalization between tasks is not plausible, as for example, in recent research by Son and Kornell (2005) and Kornell et al. (2007) where monkeys trained to emit confidence judgements in a perceptual task are then able to generalize their use to a new short-term memory task (BEM has no problem with the confidence judgment task itself as, despite the labels used to describe it, it is basically a DMTS task not very different in principle from the one used by Hampton, 2001).

Does it mean that those tasks would have demonstrated genuine animal metacognition? Obviously, no. The failure of BEM has implications for BEM only. It says nothing about whether or not metacognition is required by these new data. Maybe another model, based on different assumptions, will be able to perfectly account for them, also without using the concept of metacognition.

The core problem with metacognition is that it is not a

scientific explanation in the usual sense. No process is proposed, no computable theory offered. Rather, like "theory of mind" or "insight," it is a faculty, a hypothetical new capacity supposedly inexplicable by existing theory, especially behavioristic theory. It is defined by analogy to our own subjective experience—and by exclusion: an aspect of behavior not attributable to simple reinforcement principles or to any kind of associative learning. And this is simply not satisfying. We cannot accept metacognition as a "true" concept simply because existing models cannot account for some data—no more than we can accept that some peculiarity of molecular biology which currently eludes evolutionary theory is proof of "intelligent design." That would be making an argument from ignorance. Even if we do not have a model now, we might have in the future. This is especially true now when rather few attempts have been made in studies of animal metacognition to evaluate the ability of simple, behavioral models to account for the data.

Hence, we do not think the metacognition problem can be solved by adding new behavioral criteria that animals must display in order to qualify. Whether or not animals display metacognition is not an empirical problem. It is a theoretical one. Researchers must develop models, mathematical or computational, describing how metacognition works in these tasks. The proof that animal has metacognition will come, if it ever does, not from the inability of models without metacognition to account for the data but from the ability of a model with metacognition to (better) account for them. Since no such model exists at this time, proof that animals display metacognition is still lacking.

Perhaps that doesn't matter. The tasks and results we reviewed in this article are intrinsically interesting. Hence, the issue might not so much "do animals display metacognition in those tasks?" but simply "how do they perform those tasks?" If the best model is one requiring metacognition, then we will have discovered that animals indeed have metacognition. Otherwise, we might still learn something valuable about the mechanisms of behavior and how they contribute to the genesis of complex activities.

Progress on these lines might well lead to a reevaluation of the usefulness of the concept of metacognition in human beings. After all, had many of the tasks described in this paper been performed by humans, metacognition would have been invoked. Studies (i.e., Smith et al., 1998) which have compared human with animal performance and have found them to be extremely similar, suggesting similar underlying mechanisms. Thus, if a metacognition-free model such as BEM can account for the animal data, one may wonder if it can account for the human data as well. Indeed, experimental work on metacognition has shown that feeling-of-knowing judgments in humans are based on the ability of the sub-

ject to retrieve cues associated with the retrieval of the target item (see Metcalfe, 2008 for reviews). This is what would have been expected if simple associative learning mechanisms were underlying feeling-of-knowing judgments. The issue here is not whether or not animals or humans can judge the difficulty of the task. This is an empirical question and the answer is obviously yes. It is how to explain this ability: do we need to invoke a mysterious new faculty, metacognition, that would eventually be uniquely human, or can basic associative learning process account for it? Again, this issue is theoretical, not empirical. It will not be decided by finding facts alone but by testing models to explain them.

References

- Bateson, M., & Kacelnik, K. (1995). Preference for fixed and variable food sources: variability in amount and delay. *Journal of the Experimental Analysis of Behavior*, **63**, 313-329. doi:10.1901/jeab.1995.63-313
- Beran, M. J., Redford, J. S., & Washburn, D. A. (2006). Rhesus monkeys (*Macaca mulatta*) monitor uncertainty during numerosity judgements. *Journal of Experimental Psychology: Animal Behavior Processes*, **32**, 111-119. doi:10.1037/0097-7403.32.2.111
- Footo, A., & Crystal, J. (2007). Metacognition in rats. *Current Biology*, **17**, 551-555. doi:10.1016/j.cub.2007.01.061
- Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, **98**, 5359-5362. doi:10.1073/pnas.071600998
- Hampton, R. R., Zivin, A., & Murray, E. A. (2004). Rhesus monkeys (*Macaca mulatta*) discriminate between knowing and not knowing and collect information as needed before acting. *Animal cognition*, **7**, 239-246. doi:10.1007/s10071-004-0215-1
- 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. doi:10.1037/0097-7403.25.3.389
- Jozefowicz, J., Staddon, J. E. R., & Cerutti, D. T. (in press). The behavioral economics of choice and interval timing. *Psychological Review*.
- Kacelnik, A., & Bateson, M. (1996). Risky theories: the effects of variance on foraging decisions. *American Zoologist*, **36**, 293-311.
- Kornell, N., Son, L. K., & Terrace, H. S. (2007). Transfer of metacognitive skills and hint seeking in monkeys. *Psychological Science*, **18**, 64-71. doi:10.1111/j.1467-9280.2007.01850.x
- Metcalfe, J. (2008). Metamemory. In Roediger, H. & Byrne, J. H. (Eds), *Learning and memory: a comprehensive reference. volume 2: Cognitive psychology of memory* (pp. 349-362). Oxford: Elsevier.
- Metcalfe, J. (in press). Evolution of metacognition. In J. Dunlosky & R. Bjork (Eds.), *Handbook of metacognition and learning*. Lawrence Erlbaum.
- Metcalfe, J., & Kober, H. (2005). Self-reflective consciousness and the projectable self. In Terrace, H. & Metcalfe, J. (Eds), *The missing link in cognition: origins of self-reflective consciousness* (pp. 57-83). New York: Oxford University Press.
- Nelson, T. O. (1996). Metacognition and consciousness. *American Psychologist*, **51**, 102-116. doi:10.1037/0003-066X.51.2.102
- Nelson, T. O., & Narens, L. (1990). Metamemory: a theoretical framework and new findings. In Bower, G. H. (Ed), *The psychology of learning and motivation* (pp. 125-173). New York: Academic press.
- Roche, J. P., Timberlake, W., & McCloud, C. (1997). Sensitivity to variability in food amount: risk aversion is seen in discrete-choice, but not free choice trials. *Behaviour*, **134**, 1259-1272. doi:10.1163/156853997X00142
- Shields, W. E., Smith, J. D., Guttmanova, K., & Washburn, D. A. (2005). Confidence judgments by humans and rhesus monkeys. *Journal of General Psychology*, **132**, 165-186.
- Shields, W. E., Smith, J. D., & Washburn, D. A. (1997). Uncertain responses by humans and rhesus monkeys (*Macaca mulatta*) in psychophysical same-different task. *Journal of Experimental Psychology: General*, **126**, 147-164. doi:10.1037/0096-3445.126.2.147
- Smith, J. D., Beran, M. J., Couchman, J. J., & Coutinho, M. V. C. (2008). The comparative study of metacognition: sharper paradigms, safer inferences. *Psychonomic Bulletin and Review*, **15**, 679-691. doi:10.3758/PBR.15.4.679
- Smith, J. D., Schull, J., Strote, J., McGee, J., Egnor, R., & Erb, L. (1995). The uncertain response in the bottlenosed dolphin (*Tursiops truncatus*). *Journal of Experimental Psychology: General*, **124**, 391-408. doi:10.1037/0096-3445.124.4.391
- Smith, J. D., Shields, W. E., Allendoerfer, K. R., & Washburn, D. A. (1998). Memory monitoring by animals and humans. *Journal of Experimental Psychology: General*, **127**, 227-250. doi:10.1037/0096-3445.127.3.227
- Smith, J. D., Shields, W. E., & Washburn, D. A. (2003). The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences*, **26**, 317-373. doi:10.1017/S0140525X03000086
- Sole, L. M., Shettleworth, S. J., & Bennett, P. J. (2003). Uncertainty in pigeons. *Psychonomic Bulletin and Review*, **10**, 738-745.
- Son, L. K., & Kornell, N. (2005). Metacognitive judgments in rhesus macaques: Explicit versus implicit mechanisms. In Terrace, H. & Metcalfe, J. (Eds), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 296-320). New York: Oxford University Press.
- Staddon, J. E. R. (2001). *Adaptive dynamics: The theoretical analysis of behavior*. Cambridge, MA: MIT Press.

- Staddon, J. E. R., & Innis, N. K. (1966). Preference for fixed vs. variable amounts of reward. *Psychonomic Science*, **4**, 193-194.
- Staddon, J. E. R., Jozefowicz, J., & Cerutti, D. T. (2007). Metacognition: a problem, not a process. *Psycrit*, April. (www.psycrit.com).
- 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. doi:10.1037/0097-7403.34.2.266
- Washburn, D. A., Smith, J. D., & Shields, W. E. (2006). Rhesus monkeys (*Macaca mulatta*) immediately generalize the uncertain response. *Journal of Experimental Psychology: Animal Behavior Processes*, **32**, 185-189. doi:10.1037/0097-7403.32.2.185
- White, K. G., & Wixted, J. T. (1999). Psychophysics of remembering. *Journal of the Experimental Analysis of Behavior*, **71**, 91-113. doi:10.1901/jeab.1999.71-91
- Wynne, C. D. L., & Staddon, J. E. R. (1988). Typical delay determines waiting time on periodic-food schedules: static and dynamic tests. *Journal of the Experimental Analysis of Behavior*, **50**, 197-210. doi:10.1901/jeab.1988.50-197

Footnotes

- ¹ Risk-sensitivity is necessary to explain Foote and Crystal's (2007) data, even if the rats had metacognition. On a trial where they would not have known the answer, the animal has to choose between the uncertain response, delivering 3 pellets for sure, and the other two responses, which will provide on average 3 pellets of food. Hence, in the absence of risk-sensitivity, the animal would not be biased toward the uncertain response but indifferent between this response and the discrimination task.
- ² Here is a more traditional alternative. Let's assume that, when presented with the sample, the animal forms a STM trace which can be in two states: present or absent. The probability that the trace moves from one (present) to the other (absent) increases with the retention interval. In this case, simple reinforcement processes would allow the animal to learn to pick the uncertain response when the memory is absent, to decline it if the memory is present. This does not require metacognition: it requires the animal to remember or not, not to know that it remembers or not. Again, whether or not DMTS task involves metacognition depends on how the forgetting problem is theorized.