

Chapter 11

Do animals know what they know?

Sara J. Shettleworth and Jennifer E. Sutton

Abstract: Using well-established paradigms for studying animal perception and memory, researchers have begun to ask whether animals can monitor the status of their knowledge in a behavioural task - whether they know what they know. Generally, such metacognitive ability is tested by giving animals the opportunity to avoid (or 'escape') a test of memory or perceptual discrimination. The pattern of escapes can then be analyzed in a number of ways, including whether the subject escapes more often from difficult tests, where a correct answer is less likely, than from easy tests. A number of non-metacognitive strategies can be used by animals in these experiments, however, and it is important to control carefully for alternative explanations. Moreover, only rigorous, controlled tests will determine whether current suggestions of species differences in metacognitive abilities are correct.

Some of the most interesting and controversial work in animal behaviour involves the search for relatively complex mental processes in nonhuman animals. Investigations of episodic memory (Clayton et al, this volume), mind reading (Tomasello & Call; Povinelli & Vonk, this volume), intentional deception (Kummer et al 1996), and metacognition (Hampton, 2001; Inman & Shettleworth, 1999; Smith, Shields, & Washburn, 2003; Sole,

Shettleworth, & Bennett, 2003) have established experimental situations in which animals may be tested for behaviour functionally similar to behaviour that is accompanied by distinctive mental states in humans.

The term *functional similarity* (Hampton, 2001) captures the idea that the best we can do in such investigations is to define rigorously the behaviour accompanied by a given mental process and see if the animals show it. If the process is one usually assessed by verbal report, we will never be able to have the same kind of evidence for it in other species as in humans. For example, episodic memory in adult humans is an integrated memory for the what, where, and when of an event that comes with an awareness of re-experiencing that event (Tulving, 1972). As discussed by Clayton et al. (this volume), it may be possible to devise ways to probe an animal's memory to see if it remembers the 'what', 'where' and 'when' of an event in an integrated way, but this is not the same as asking about its private experiences while recalling the event. People may never agree on the degree to which similar nonverbal behaviour implies similar mental events in very different species (c.f. Griffin, 2001), but when the phenomena under study are well enough defined, agreement on functional similarity may be possible.⁹⁷ The fact that researchers interested in episodic memory in nonhuman animals refer to what they study as 'episodic-like' memory (Griffiths et al 1999; Morris, 2001) acknowledges this basic methodological and philosophical truth.

⁹⁷ Editors' note: see and cf. the methodological points about distinguishing rational and associative processes made by Heyes & Papineau, this volume.

In this chapter we discuss whether animals know what they know. That is, do any animals have metacognitive abilities? There are well-established paradigms for assessing cognitive processes such as memory, perceptual processing, categorization, timing, and numerical competence in nonhuman animals (henceforth simply 'animals'), and we know a great deal about how they perform in these tasks (Roberts, 1998; Shettleworth, 1998). But in addition to knowing which stimulus was presented most recently or how long it lasted, do animals know *that* they know such things? That is to say, in addition to reporting what it had just seen or heard, could any animal additionally report on the status of its knowledge? Adult humans can. In a variety of laboratory paradigms, reports of a 'feeling of knowing' in memory tests or perceptual certainty in psychophysical tasks are positively associated with accurate performance on these tasks (Nelson, 1996). Thus people not only feel that they have metacognitive abilities, but metacognition is generally accurate.

In everyday human life, metacognition clearly contributes to rational, efficient, behaviour. For example, before setting out on a shopping expedition, a person will be aware of whether or not he needs to consult a map. Similarly, one usually knows whether or not it's necessary to look in the phone book before making a call. We also behave as if we assume accurate metacognition in others. For example, lost in a new city, one approaches a stranger not by asking 'Where is the train station?' but more likely by saying, 'Can you tell me the way to the train station?' If the stranger replies 'Sorry, I can't', we ask someone else. It is precisely the dissociation between knowing and awareness

of knowing in some brain-damaged patients that makes them so fascinating (e.g. Weiskrantz, 1986), but on the whole we can safely treat others as if they can tell us whether or not they know something, or how sure they are that they know it. Can members of any other species do the same thing, or are they like patients who have knowledge without awareness?

11.1. How can we test animals' metacognition?

As the preceding examples indicate, the challenge to researchers is simultaneously to test animals' ability to perform a cognitive task (to answer a direct question like 'What is the way to the station?') and to test their ability to report on the state of their knowledge (to answer a question like 'Do you know the way to the station?'). Consequently, metacognition experiments include both a test of memory or perceptual discrimination that can vary in difficulty, and an additional response that allows the animal to accept or decline these tests. To encourage the animals to use it appropriately, this 'don't know' or 'uncertain' response receives a small or delayed reward, more than the non-reward typically given for an incorrect response in the cognitive test but less than could be obtained by responding correctly in that test. Thus this opportunity to 'opt out' or escape from an impending test should be taken only if the animal knows it is likely to fail the test.

Any method for studying the functional similarity of animal behaviour and complex human mental processes must carefully control for alternative interpretations. Therefore, one further step is needed to be sure such choices are based on metacognition rather than on learning about distinctive external

stimuli that predict the animal's performance (and therefore reward rate) in the primary cognitive task. For example, suppose we are using a test of memory for shapes in which we vary memory strength by varying the interval between showing the animal the shape and testing its memory, as illustrated in Figure 11.1. The longer the interval between this stimulus (the 'sample') and the test, the worse performance will be. Our subject therefore can maximize its reward rate by choosing to complete the memory test after a short interval, when it will likely receive a large or immediate reward, and by opting out after a long interval, when the relatively small reward for escaping exceeds the reward expected from performing at chance in the memory test, that is, without the aid of memory. That is to say, in this case behaviour functionally similar in one respect to behaviour consistent with metamemory could result from discriminating a long from a short retention interval and choosing accordingly. We know that animals are very good at discriminating among time intervals (Shettleworth, 1998), so this is a very real possibility. Obviously, then, functional similarity must be defined more richly, so that it specifies a constellation of behaviours that together are uniquely consistent with metacognition.

There are two ways in which this has been done, and in a handful of experiments both have been used. One is to include randomly presented trials without the escape option, the 'forced trials' in Figure 11.1. On such trials, the animal is forced to perform the primary memory task whether it knows the answer or not. An animal that is escaping selectively on trials when it 'knows it does not know' will do worse on these forced trials than on trials that it

freely chose when it could have escaped. Notice, however, that because many animals tend to do much worse than normal when conditions change, performance may also be impaired if trials without the escape option are relatively novel. Therefore, it is important to mix in such trials throughout the experiment rather than presenting them only occasionally.

The second, and very important, way to address the possibility that subjects simply learn to escape in the presence of specific external stimuli is to include transfer tests with new stimuli. This is the well-established policy in psychology of using convergent operations or 'triangulation' (Heyes, 1993) to test for an inferred cognitive process in multiple ways. In our example, after testing memory with varying retention intervals and finding behaviour consistent with metamemory, we could now keep the retention interval constant at some intermediate value and vary the length of time for which the shapes are presented to vary memory strength in a new way. Longer presentations should improve memory, shorter presentations weaken it, and the animals' use of the escape option should change accordingly. It should also change immediately upon introduction of this new test. Other ways of varying memory strength can be imagined, such as varying the inter-trial interval or the salience of the stimuli. In one clever test with monkeys (Hampton, 2001), the to-be-remembered stimuli were occasionally omitted altogether. An animal using metamemory must report 'I don't remember' on such trials, which is what the monkeys did.

Not all tests of metacognitive abilities in animals have provided such clear cut results as did the one by Hampton just mentioned. In the rest of this

chapter, we briefly review the research in this area, identify some gaps in knowledge, and sketch possible future directions. Smith, Shields, and Washburn (2003) have recently reviewed research on animal metacognition in more detail. We emphasize more than they do the methodological issues sketched above, and we are accordingly less inclined to interpret the results as evidence for human-like awareness or metacognition.

11.2. Uncertainty in perceptual discriminations

The first published studies of metacognition in animals asked whether animals could report on their level of perceptual certainty when performing a difficult discrimination (Smith et al., 1995; Shields, Smith, & Washburn, 1997; Smith, Shields, Schull, & Washburn, 1997; Sole et al., 2003). In the first such study, Smith et al. (1995) tested a dolphin in an auditory psychophysical task. The dolphin was presented with a single tone from an underwater speaker on each trial. Two underwater response paddles were also available, with a response to one resulting in reinforcement after a high-pitched tone and the other resulting in reinforcement after a low-pitched tone. After a series of training phases, the final testing phase contained probe trials where the pitch of the relatively lower tone was raised, making it more similar to the high tone. Along with this increase in difficulty, the dolphin also had a third response paddle available, which allowed it to escape from categorizing the tone and endure a delay followed by an easily discriminated tone instead. The important prediction was that the dolphin should preferentially choose

the escape option and the less-preferred delayed reward on difficult trials but attempt to categorize the tone on easier trials for the immediate payoff.

At first, the dolphin chose the escape option whenever it was presented, so an additional cost was added by increasing the delay to the easy trial if the escape option had been chosen frequently in the previous trials. Eventually, the dolphin escaped more when the tone was hard to classify than when the tone was easily classified as high or low. In a comparable experiment (Smith et al., 1997), monkeys were trained to classify the density of illuminated pixels on a video monitor as either low or high. Like the dolphin, the monkeys escaped most on difficult trials, with stimuli in the middle of the continuum.

A nice feature of both of these studies is that people were tested on almost identical tasks so their performance could be directly compared to the animals'. It was indeed similar, and the data from humans, monkeys and dolphin could all be accounted for by a signal detection model in which responses are based on dividing the stimulus continuum into 'low', 'uncertain', and 'high' regions. The human participants reported that they chose the escape option when they experienced a feeling of uncertainty about how to classify the stimulus. Explaining human and animal data using a single model makes it tempting to conclude the same underlying, conscious process is employed by all the participants. However, although signal detection models refer to setting criteria and making decisions, there is no necessary implication that these processes are conscious (Macmillan & Creelman, 1991). Moreover, because the studies with the monkeys and the

dolphin did not include regular, randomly presented, tests of the animals' discrimination performance in the absence of the escape option, it is unclear whether or not they were actually using the escape response selectively when they 'knew they did not know' how to classify the stimulus. In addition, because no transfer tests were included, it seems possible that the animals had simply learned that for stimuli in a certain part of the tone frequency or pixel density continuum, the escape response had the shortest average delay to reward. Thus, although these studies are ground-breaking and the results are consistent with those of some of the more stringent tests reviewed below, they lacked all the controls required for a conclusive demonstration of metacognition.

Pigeons have been widely tested in psychophysical procedures and the results successfully accounted for by signal detection models (Blough & Blough, 1977; Commons, Nevin, & Davison, 1991). This suggests that pigeons should behave similarly to the subjects in the studies of Smith et al. if they were given a similar test. This was the idea behind a study by Sole, Shettleworth, and Bennett (2003). Like the monkeys in Smith et al. (1997), the pigeons were required to classify the number of pixels on a computer monitor as sparse or dense. They were also provided with an escape option that could be chosen instead of classifying the image. However, this experiment also included two control procedures lacking in the earlier studies. First, rather than a delay to an easy trial, the reward for escaping was three pigeon pellets, a mediocre reward compared to the six pellets that could be gained if the stimulus was correctly classified. This reward was constant for each pigeon

rather than changing dynamically throughout the experiment, which helped to ensure that the birds knew the outcome of each available option and facilitated modelling the results. Second, after experience with the escape option, the birds received a transfer phase where some stimuli were red instead of green. They needed time to acquire the density discrimination with these new stimuli and performed poorly at first. Therefore, in the first session of transfer when performance was poor, the birds should choose to escape if uncertainty was truly governing the use of the escape option. Finally, classification accuracy was compared on forced and freely chosen trials throughout the experiment.

Like the monkeys and the dolphin in the experiments by Smith and colleagues, the pigeons escaped more difficult trials than easy trials. However, accuracy on forced versus chosen tests did not differ. Moreover, when presented with the new transfer stimuli that were difficult to classify, the birds performed poorly but did not opt for the escape option any more than they did once they were classifying the new stimuli accurately. Thus the pigeons did not selectively choose to take the test when they 'knew that they knew' how the pixels should be classified. However, their performance could be accounted for very well by a signal detection model that assumed they were maximizing perceived reward for correct responses in a consistent way across all the experimental conditions. That is to say, they behaved as if they had learned which was the most profitable response to each region of the continuum of internal responses evoked by the displays. [Fitting the model

did not require postulating some extra reward such as 'reduction of uncertainty' on the difficult trials.

Is it the case then, that pigeons do not monitor uncertainty while monkeys, dolphins, and humans do? Because there were procedural differences between experiments, it is still difficult to form a firm conclusion. The tests with pigeons may be viewed as more rigorous given the multiple predictions afforded by the design, but it is still not known how monkeys, dolphins, or people would behave with identical tests.

11.3. Monitoring the status of a memory

Like psychophysical investigations, the study of memory processes in animals is well established (Roberts, 1998; Shettleworth, 1998) and provides a strong basis for tests of metacognition. A substantial series of experiments on monkeys and one on pigeons has sought to determine whether animals can assess the status of a memory. These tests of metamemory have employed two different paradigms: serial probe recognition and matching-to-sample.

Smith, Shields, Allendorfer, & Washburn (1998) used the well-established serial probe recognition task with monkeys. In a serial probe recognition task (Wright, 1989), a series of images is displayed (a 'list' of to-be-remembered items), and then a probe image is presented. The subject's task is to determine whether the probe image was or was not displayed in the preceding list. Under the conditions used by Smith et al. (1998), monkeys typically show a recency effect, performing best when the probe item was presented late in the list. Longer lists are also more difficult than short ones.

In Smith et al's (1998) study, images in the list were chosen randomly from a bank of over 100 digital files, enabling a different combination of images on each trial (so-called trial-unique items).

After the list of two to six images was presented, a probe image and three stimuli indicating different response options appeared on the screen. Responses to one stimulus were reinforced if the probe image was presented in the previous list, one was reinforced if the image was not presented, and one stimulus served as the 'escape' option and was followed by a delay and then the correct answer. Only one item from each list was tested, and the position of the tested item was varied randomly. For example, on one trial the probe might be the first item from the most recent list. Then a new list would be presented and the last item from that list presented as the probe, and so on. Importantly, the use of this probe procedure together with trial-unique stimuli eliminated the possibility that the monkeys could learn specific stimulus-outcome pairings. Instead, choosing the escape option on difficult trials could be confidently attributed to monitoring the status of memory for the probed item. The monkeys chose the escape option more often on tests of earlier items that were more difficult to remember. That is, accuracy showed a significant recency effect, and escaping mirrored this effect, with fewer escapes for most recently presented items. Generally, accuracy was higher on trials when the monkeys chose to take the test than when the escape option was not available and they were 'forced' to answer, although they did not receive both types of trial in a single session.

Hampton (2001) also tested rhesus monkeys' metamemory, but he used a delayed matching-to-sample procedure similar to that depicted in Figure 11.1. On each trial, the monkeys were shown an image to study (the 'sample'), followed by a delay with no stimulus present. After this retention interval, the monkeys chose between two stimuli that led to different conclusions of the trial. A response to one stimulus resulted in a memory test, where the sample was presented along with three distractor stimuli, and a correct match was rewarded with a highly-valued peanut. Responses to the other stimulus led to one final stimulus instead of a memory test, and a primate pellet reward, a treat considered rather mediocre by the monkeys. Tests at the end of each day confirmed that the monkeys always preferred peanuts to pellets.

Hampton's (2001) study was a powerful test for a number of reasons. First, the monkeys were required to make a decision about the status of their memories *before* the test was presented. That way, retrieval cues from the presence of the correct stimulus or interference from the incorrect stimuli could not influence the memory judgment. This procedure is demanding because the animal must access the strength of its memory for the most recent sample stimulus in the absence of that sample. Second, the use of two different kinds of rewards produced a natural cost for opting out of the memory test and eliminated the need to punish the monkeys for overusing that option. Finally, Hampton manipulated the strength of monkeys' memories for the sample in more than one way, providing an opportunity for the analysis of transfer performance. Two monkeys were first trained on the

task just described with a single retention interval, chosen so they made errors on about 20% of trials. Consistent with the use of metamemory, both animals performed better on freely chosen than on forced memory tests. To test whether use of the escape option was equivalent to reporting 'I don't remember', they were then given probe trials where no sample was presented. Finally, their memory strength was manipulated by increasing the retention interval.

The results of all three phases of the study were consistent with memory monitoring. The monkeys were more accurate when they chose to take a test than when they were forced to, they declined memory tests much more often when no sample had been presented, and they performed more poorly and declined more tests after longer delays. In addition, performance on freely chosen memory tests declined less at longer delays than did performance on forced tests, consistent with the monkeys selectively choosing the memory test when they knew they remembered the sample. Taken together, the serial probe recognition test of Smith et al. (1997) and the matching-to-sample experiments by Hampton (2001) provide strong evidence that monkeys can monitor the status of a memory.

In the first published investigation of any metacognitive process in pigeons, Inman and Shettleworth (1999) used a procedure nearly identical to Hampton's (2001) and found that procedural factors were crucial to determining whether pigeons' behaviour was consistent with memory monitoring. They used a matching-to-sample design with several retention intervals, and as in Sole et al.'s (2003) study, escaping the test resulted in

fewer pellets than a correct answer. As in Hampton's (2001) transfer test, at the longer delays memory should be worse, escape should be chosen more often, and there should be a bigger difference between forced and chosen memory performance. Unlike the case in Hampton's experiment, however, but as in the study of Shields et al., in Inman and Shettleworth's Experiment 1 the memory test and the escape option were presented at the same time on choice trials, rather than each one following an explicit choice as depicted in Figure 11.1. In this experiment, three of the four birds chose the escape option more at the longer delays, where memory should be poor, than at the short delays. The birds were also less accurate on forced than on freely chosen memory tests.

The pattern of results was different in a second experiment with the same birds. Here, the escape option was presented before the opportunity to complete the memory test, as in Hampton's experiments and as shown in Figure 11.1. On those trials, delay length had no effect on the probability of escaping the test, and the birds were no more accurate on chosen than forced tests. Thus, when the birds had to consult their memory strength before choosing whether or not to have their memory tested, they failed the test of metamemory, whereas if they were confronted with the items in the memory test and the escape option at the same time, they passed. One might want to conclude that the pigeons can pass an easy but not a difficult test of metamemory. However, one reason why this interpretation might not be warranted is that in Inman and Shettleworth's second experiment, most birds were performing so well in the primary test of memory that there was little to

be gained by escaping from it at even the longest retention interval. Another potential problem is the lack of evidence that the pigeons actually had any preference for the mixture of six and zero pellets determined by their proportion of correct responses over the three sure pellets from the escape option. Thus, Inman and Shettleworth had to conclude that while their results were not incompatible with pigeons being able to use metamemory, they were not conclusive one way or the other.

Given that Hampton's (2001) study using very similar procedures to Inman and Shettleworth's subsequently provided strong evidence for metamemory in monkeys, we (Sutton & Shettleworth, in preparation) took up the study of pigeon metamemory again in an experiment very similar to Hampton's. One important change from the earlier studies with pigeons is that we verified pigeons' preferences among the rewards we used: all birds in an independent study preferred a constant twelve pellets every time to a constant five pellets, but they preferred a constant five pellets to receiving twelve pellets on 25% of trials and zero pellets otherwise. In the main experiment, memory tests had four alternatives with twelve pellets' reward for a correct choice, and escaping was rewarded with five pellets. Pigeons should therefore escape the memory tests when they know that they do not know and therefore can expect only a 25% chance of getting 12 pellets.

Like Hampton's monkeys, the birds were first trained with a single retention interval in a procedure like that depicted in Figure 11.1. Unlike the monkeys, at this stage they performed no better on freely chosen than on forced tests of memory. This same pattern of results was maintained when

testing continued with a mixture of a long and a short retention interval. However, most birds apparently learned that the value of taking the memory test declined relative to the value of escaping at the longer delay, since they escaped more at this delay. Finally, the birds did not consistently reject the test of memory after no sample had occurred. Like Inman and Shettleworth's Experiment 2, these results indicate that pigeons cannot report on their metacognitions under conditions where monkeys can.

11.4. Summary and future directions

Consistent with the literature on human subjects (e.g. Nelson, 1996), we and other researchers using animals have tended to lump together tests of perceptual certainty with tests of metamemory as if they assay a single cognitive process. This assumption may be encouraged by the fact that nearly all the tests of either one to date consist of offering the animal the option of escaping from the test of cognition (the uncertain response, Smith et al. 1997). But perceptual certainty ('Do I know what I am seeing?') and memory monitoring ('Do I remember what I saw?') are not necessarily the same process, even though people as well as animals might report on it by 'saying' they are uncertain. Some species might show one and not the other. For example, whereas pigeons passed one easy test of metamemory (Inman & Shettleworth, 1999, Experiment 1), they failed a parallel test of perceptual certainty (Sole et al, 2003). In the absence of further studies with pigeons, the existing data are not conclusive, but they do suggest that future studies might more carefully distinguish between perceptual certainty and memory

monitoring abilities. Consistent with this suggestion, in humans 'blindsight', or vision without awareness (Weiskrantz, 1986) and specific loss of episodic but not semantic memory are associated with damage to different brain areas.⁹⁸

A further important distinction is that between showing the animal the item(s) to be classified or recognized at the same time as offering it an alternative to completing the primary task and offering this choice before presenting the primary task (as in Figure 11.1). To our knowledge, the latter procedure has not been used in a perceptual task with animals. In a memory task it is surely the more difficult procedure because it requires the animal to call to mind an absent stimulus. Intuitively, displaying the test stimuli and asking 'Do you recognize one of these or would you rather not answer?' tests memory strength in a more direct way than asking 'Do you think you will recognize the sample when you see it?' As pointed out in the introduction to this chapter, tests, which do the latter, are most closely analogous to tests of human metamemory. So far, however, they have been used only in one of the studies with monkeys (Hampton, 2001) and two with pigeons (Inman & Shettleworth, 1999, Experiment 2; Sutton & Shettleworth, in preparation).

All the studies we have reviewed so far involve training animals very extensively to give an explicit report on memory strength or perceptual

⁹⁸ Editors' note: It would be interesting to relate the paradigms described here for studying metacognition in animals to the paradigms used to study blindsight in animals: would blindsighted monkeys opt to escape no-stimulus classification tasks? See Stoerig & Cowey 1999, Stoerig et al 2002.

certainty. In contrast, the examples from everyday life sketched at the beginning of this chapter illustrate how people make spontaneous implicit use of metacognition. Might these suggest new ways to look for metacognitive abilities in animals, related to ways they might spontaneously use these abilities in biologically relevant contexts? Call and Carpenter (2001) have recently devised such a test (see also Call, this volume). Chimpanzees, orangutans, and children saw food (for the apes) or a sticker (for the children) hidden in one of two or three horizontal tubes and were allowed to retrieve it. Sometimes it was difficult to know which tube was baited either because the subject did not see the baiting or because there was a delay before a choice was permitted. These conditions increased the chance that subjects would bend down and peer into the tubes before choosing, as if they knew they did not know and were looking for more information. Not all the apes' data were consistent with looking into tubes being information-seeking, however. For instance, subjects often continued looking after they had found the food, and they did not always look when the baiting had been concealed. Hampton, Zivin, and Murray (2004) adapted this task for rhesus macaques and most monkeys looked down the tubes more often when they were ignorant of which one held the bait. This finding is particularly important because it means that memory awareness in this species is supported by convergent evidence from two quite different kinds of task.

So far, tests of whether animals know what they know suggest that despite the fact that both monkeys and pigeons can perform difficult tests of perceptual discrimination or memory, only the monkeys and perhaps other

primates can report on their cognitive states while they are doing these tasks. Thus only primates and perhaps a dolphin have been proven rational in a strong sense that requires having access to the reasons for their behaviour. But because their choices in similar tasks optimize reward under the constraint of imperfect discrimination, pigeons are at least biologically rational (see Kacelnik, this volume). One challenge for the future might be to seek naturalistic situations in which a biologically rational animal must be rational in the stronger sense as well.

References

- Blough, D., and Blough, P. (1977). 'Animal psychophysics', in W. K. Honig and J. E. R. Staddon (eds.), *Handbook of Operant Behavior*. Englewood Cliffs, NJ: Prentice-Hall, 514-539.
- Call, J., and Carpenter, M. (2001). 'Do apes and children know what they have seen?' *Animal Cognition*, 4: 207-220.
- Commons, M. L., Nevin, J. A., and Davison, M. C. (eds.). (1991). *Signal Detection: Mechanisms, Models, and Applications*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Griffin, D. R. (2001). 'Animals know more than we used to think'. *Proceedings of the National Academy of Sciences (USA)*, 98: 4833-4834.
- Griffiths, D., Dickinson, A., and Clayton, N. S. (1999). 'Episodic memory: what can animals remember about their past?' *Trends in Cognitive Science*, 3: 74-80.
- Hampton, R. R. (2001). 'Rhesus monkeys know when they remember'. *Proceedings of the National Academy of Sciences (USA)*, 98: 5359-5362.
- Hampton, R. R., Zivin, A., and Murray, E. A. (2004). 'Rhesus monkeys (*Macaca mulatta*) discriminate between knowing and not knowing and collect information as needed before acting'. *Animal Cognition*, 7. Published online 23 April, 2004. DOI: 10.1007/s10071-004-0215-1.
- Heyes, C. M. (1993). 'Anecdotes, training, trapping and triangulating: do animals attribute mental states?' *Animal Behaviour*, 46: 177-188.

- Inman, A., and Shettleworth, S. J. (1999). 'Detecting metamemory in nonverbal subjects: A test with pigeons'. *Journal of Experimental Psychology: Animal Behavior Processes*, 25: 389-395.
- Kummer, H., Anzenberger, G., and Hemelrijk, C. K. (1996). 'Hiding and perspective taking in long-tailed macaques (*Macaca fascicularis*)'. *Journal of Comparative Psychology*, 110: 97-102.
- Macmillan, N. A., and Creelman, C. G. (1991). *Detection theory: A user's guide*. Cambridge: Cambridge University Press.
- Morris, R. G. M. (2001). 'Episodic-like memory in animals: psychological criteria, neural mechanisms and the value of episodic-like tasks to investigate animal models on neurodegenerative disease'. *Philosophical Transactions of the Royal Society (London) B*, 356: 1453-1465.
- Nelson, T. O. (1996). 'Metacognition and consciousness'. *American Psychologist*, 51: 102-116.
- Roberts, W. A. (1998). *Principles of Animal Cognition*. New York: McGraw Hill.
- Shettleworth, S. J. (1998). *Cognition, Evolution, and Behavior*. Oxford: Oxford University Press.
- Shields, W. E., Smith, J. D., and Washburn, D. A. (1997). 'Uncertain responses by humans and rhesus monkeys (*Macaca mulatta*) in a psychophysical same-different task'. *Journal of Experimental Psychology: General*, 126: 147-164.
- Smith, J. D., Schull, J., Strote, J., McGee, K., Egnor, R., and Erb, L. (1995). 'The uncertain response in the bottlenosed dolphin (*Tursiops truncatus*)'. *Journal of Experimental Psychology: General*, 124: 391-408.

- Smith, J. D., Shields, W. E., Allendoerfer, K. R., and Washburn, D. A. (1998). 'Memory monitoring by animals and humans'. *Journal of Experimental Psychology: General*, 127: 227-250.
- Smith, J. D., Shields, W. E., Schull, J., and Washburn, D. A. (1997). 'The uncertain response in humans and animals'. *Cognition*, 62: 75-97.
- Smith, J. D., Shields, W. E., and Washburn, D. A. (2003). 'The comparative psychology of uncertainty monitoring and metacognition'. *Behavioral and Brain Sciences*, 26: 317-339.
- Sole, L. M., Shettleworth, S. J., and Bennett, P. J. (2003). 'Uncertainty in pigeons'. *Psychonomic Bulletin and Review*, 10: 738-745.
- Stoerig, P. & Cowey, A. (1997). 'Blindsight in man and monkey.' *Brain* 129: 535-559.
- Stoerig, P., Zontanou, A., & Cowey, A. (2002). 'Aware or unaware: assessment of cortical blindness in four men and monkeys.' *Cerebral Cortex* 12: 565-574.
- Sutton, J. E., and Shettleworth, S. J. (Sutton, J. E., and Shettleworth, S. J. (in preparation). 'Further tests of metamemory in the pigeon'.
- Tulving, E. (1972). 'Episodic and semantic memory', in E. Tulving and W. Donaldson (eds.), *Organization of Memory*. New York: Academic Press, 381-403.
- Weiskrantz, L. (1986). *Blindsight*. Oxford: Clarendon Press.
- Wright, A. A. (1989). 'Memory processing by pigeons, monkeys, and people'. *The Psychology of Learning and Motivation*, Vol. 24. Academic Press, 25-70.