



People's study time allocation and its relation to animal foraging

Janet Metcalfe*, W. Jake Jacobs

Columbia University, University of Arizona, United States

ARTICLE INFO

Article history:

Received 6 October 2009

Received in revised form

10 December 2009

Accepted 11 December 2009

Keywords:

Metacognition

Foraging

ABSTRACT

In this article we suggest a relation between people's metacognitively guided study time allocation strategies and animal foraging. These two domains are similar insofar as people use specific metacognitive cues to assist their study time allocation just as other species use cues, such as scent marking. People decline to study items that they know they already know, just as other species use a win-shift strategy – avoiding already visited and depleted patches – in foraging. People selectively study the easiest as-yet-unlearned items first, before turning to more difficult items just as other species take the 'just right' size and challenge of prey—the so-called Goldilocks principle. People use a stop rule by which they give up on one item and turn to another when the returns diminish just as others species use a stop rule that guides shifting from one patch to another. The value that each item is assigned on the criterion test, if known during study, influences which items people choose to study and how long they study them just as knowledge of the nutritional or energy value of the food influences choices and perseverance in foraging. Finally, study time allocation strategies can differ in their effectiveness depending upon the expertise of the student just as some species forage close to optimally while others do not.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

The analogy we explore in this article (following in the tradition of Hills, 2006, Kamil, 1983, Nairne et al., 2009, and Pirolli and Card, 1999) is between people's study time allocation, in the service of learning, and animals' foraging. Much research on human metacognition – their knowledge of what they know and do not know – has been directed at its potential usefulness in terms of people's optimization of study time allocation. The assumption here is that people *do* try to optimize their time and effort in a study situation, as other species try to optimize their efficiency in a foraging situation, that people's metacognitions play an active role in that process, and that their success as learners depends on the effectiveness of these strategies, just as the animals' evolutionary fitness depends on their foraging effectiveness. Metacognition is a cue or marker that students can use to make study more efficient, and, as such, metacognition may be analogous to some of the cues that other species use in foraging. Presumably, it is only if people know what they know, that they will be able to effectively direct their own study time and efforts to learn new things in a manner that is efficient and not entirely driven by the environment. Despite an apparently vast difference in levels of cognition, the parallels between what people do when strategically using their metacog-

nitions to enhance learning, and what other species do when they are foraging for food are striking, as we outline in this essay.

We do not claim that animals' foraging behaviour is metacognitively driven. Indeed, it is only recently that rigorous experimentation has been directed at whether any non-human primates have metacognitions under any circumstances (see, Hampton, 2001, 2005; Kornell et al., 2007; Smith et al., 2003; Terrace and Metcalfe, 2005; Terrace and Son, 2009), and the results remain controversial. Although there is no doubt that people have a metacognitive capability, it is not so clear that animals are capable of metacognition. To assess the possibility, though – when there is a question – the definition of metacognition itself needs to be strict: it involves making a judgment, commentary or reflection about a mental representation. There is little doubt that many animals, including, notably, primates, are able to mentally represent objects and events. But that is not enough for metacognition. There is also little doubt that animals are able to make judgments or decisions, even non-binary judgments, about objects, events and perceptual stimuli, in the world. Again, that is not enough. For a judgment to be a metacognitive judgment, it must be made about an internal representation. If the stimulus that the animal is making a judgment about is present in the environment (or even is a well-learned discriminatory or conditioned response) then that is a simple judgment but not a metacognitive judgment. Under this strict definition, metacognition is routinely found in humans, and almost never found in any other animal.

The extent of the favorable evidence for any metacognition meeting this definition in non-humans is based on one out of two

* Corresponding author at: Department of Psychology, Columbia University, 406 Schermerhorn Hall, 1190 Amsterdam Avenue, New York, NY 10027, United States.
E-mail address: jm348@columbia.edu (J. Metcalfe).

rhesus monkeys in a study by Hampton (2001), and two out of two monkeys in Kornell et al.'s (2007) study. After intensive training these three monkeys demonstrated metacognitive abilities. In Hampton's study, the monkeys were trained to do a delayed-match-to-sample task, where the to-be-matched stimulus changed on every trial. They had to remember this current sample for a few moments, and during those few moments the stimulus was not present in the environment, so the memory for the sample (marginally) qualified as being a mental representation. Hampton also trained the monkeys to do a choice task, in which they could press one icon if they wanted to take a test, and another if they did not. If they took the test, and were correct they got a favored treat, but if they were wrong they got nothing. If they chose to not take the test, they got a less favored treat regardless of whether they were right or wrong. After training on both the delayed-match-to-sample task and the choice task, these two were combined, so that the choice to take or not take the test was made about the sample during the delay interval—a choice about a representation in working memory. The logic is that if the monkeys knew that they remembered the sample, they should choose to take the test (and get the good treat), whereas if they did not remember the sample they should choose to not take the test (and be sure to get at least some treat). This meets the definition of metacognition (though minimally, since some theorists might deny that working memory traces are true mental representations). Of course, the monkeys might just choose randomly (as pigeons do, in this same task) and such a result would not indicate that they have metacognition. To ascertain whether the monkeys were, indeed, choosing appropriately, Hampton included some trials on which the monkeys were forced to take the test. He then compared performance on those trials on which they chose to take the test to those trials on which they were forced to take the test. If the monkey was correctly assessing his (or her) own memory of the sample when s/he made the choice, then the probability of being correct on the match-to-sample test should have been higher when the monkey chose to take the test than when forced to take the test. And, for one of the two monkeys it was! Kornell et al. (2007) used a similar logic, but with the metacognitive assessment being retrospective rather than prospective. They, also, found evidence for metacognition in two rhesus monkeys (whose names, appropriately enough, were Lashley and Ebbinghaus). In both Hampton's experiment and that of Kornell et al. (2007), training in the procedures took many months and thousands of trials. There is no evidence that any such capabilities are manifested in any behaviour exhibited by these primates, or, indeed, any other animals (except humans), in the wild.

Hence, given the difficulty in training any non-human animal to exhibit even the slightest evidence of metacognition, it is unlikely that metacognitions direct any of their behaviour. Hence the analogy between metacognitive control of study in humans and the cues that animals use to control foraging behaviour is not dependent upon identical mechanisms being used in the two situations. But the mechanisms may, nevertheless, be functionally analogous. What causes the similarities in these two situations? Did people exapt an already established strategy, designed for an entirely different purpose, to enhance their cognitive capabilities (see Hills, 2006; Pirolli and Card, 1999 for similar suggestions)? We will look, first, at what people do when they allocate their study time to a to-be-learned set of materials. With these processes and strategies in hand, we will explore parallels between this behaviour and what other species do when they forage for food.

1.1. Metacognition as a control cue in study time allocation

The link between metacognition and control of study in humans seems obvious, conceptually, at least. Many researchers have argued that if people do not know what they know or do not know,

then they cannot take appropriate measures to improve their learning or knowledge base. They will not know what to study. Although this may seem self-evident, only a few recent experiments have explicitly tested whether or not people, in fact, use metacognition to control study (see Finn, 2008; Metcalfe and Finn, 2008a; Thiede et al., 2003). These three studies, all conducted on adults, indicate that, indeed, adults do use their metaknowledge to determine what to study. In Thiede et al.'s (2003) study, for example, a method was used whereby people's metacognitions about their comprehension of text passages could be improved (or left unimproved). Normally, metacomprehension accuracy is rather poor, and people often do not know very well what they have understood and what they have not understood. To improve metacomprehension accuracy, Thiede et al. (2003) used a technique called the delayed keyword method, whereby, after reading text passages, and then waiting some time (i.e., delaying) people are asked to provide 5 keywords related to each passage. This effort to find keywords results in people subsequently making more accurate metacomprehension judgments about those passages, presumably because when it is difficult to generate the keywords people correctly infer that they have not understood the passage very well, whereas when it is easy to generate the keywords they know that they have understood. When people used this helpful delayed keyword method, they made more accurate metacomprehension judgments, and they also chose which texts to restudy more wisely—putting the results of this improved metacognition to effective use. Being allowed to study their own choices for restudy, resulted, in this condition, in superior eventual performance on a later test on the text passages.

Finn (2008) and Metcalfe and Finn (2008a) showed that when people's metacognitions were manipulated such that they were more confident, they studied less than when they were manipulated to be less confident. Finn (2008), for example, induced people into a less confident or a more confident metacognitive state by asking them to think about whether they would forget or remember the target items, respectively. Although actual learning was the same in the two conditions, when asked to think about whether they would forget the items, the participants were less overconfident than when asked whether they would remember the items. Critically, when they were less confident, the participants chose to study more than when they were more confident—indicating that the metacognitions themselves were responsible for the study choices or lack thereof. Thus, adults not only make metacognitive judgments with facility, but they use them – at least some of the time – to determine what and how much to study.

Metacognitive control is not universal even in humans, however, and can easily be undermined. For example, there appear to be implementation deficits in some populations, notably, young children. They may know what they do and do not know, if queried appropriately, and yet do not use this knowledge to control their study choices (Metcalfe and Finn, 2009). An implementation deficit may occur in other populations, but those populations have not yet been tested, although the popular refrain of the failing student who appears at the professor's office claiming to have studied so hard and to know the materials so well, suggests that many students may not know how or what to study. Thus, while there does appear to be a link between adult humans' metacognition and their control of study, and hence of learning, it may not be a necessary or an automatic connection.

1.2. Other cues in study time allocation

Although people can use their metacognitions to control study, these cues may be overridden rather easily. Sometimes discounting metacognitive cues may be rational and enhance the person's study effectiveness, but sometimes not. For instance, Ariel et al. (2009)

demonstrated that, when the experimenters assigned and clearly marked items as either having a high or low probability of showing up on a test, students selectively allocated their study time to the high rather than the low probability items. The students did this despite metacognitive cues about how easy or difficult it would be to learn the items. The students were not, however, entirely rational about doing so. They chose to probability match, that is to make choices that match the ratio of the low probabilities and high probabilities, rather than to use the more optimal strategy of devoting study time exclusively to the high probability items and ignoring the low probability items. By probability matching they omit study of many high probability items, in favor of studying items that are very unlikely to be on the test. But humans often choose the non-optimal strategy of probability matching rather than optimizing study (see, Vulkan, 2000, for review, c.f., Gaissmaier and Schooler, 2008). We will detail later that animals, in foraging situations, also probability match (Herrnstein, 1961). Thus, this irrational behaviour in humans may be a co-opted strategy – though not necessarily a good one – that other animals have used in food seeking (see, however, Bradshaw et al., 1976; Horne and Lowe, 1993).

The strategy of taking into account (however inefficiently) expectations about the test, even at the expense of one's metacognitions, is well known. Virtually all students who have succeeded in getting into college use it routinely. If the professor indicates that certain pages in a textbook will not be on the test, few students focus on those pages the night before the test. This external cue, then, results in people disregarding their metacognitions.

Sometimes, people even override metacognitive cues in favor of external (or internal) cues with no obvious validity for enhancing study time allocation. For example, social psychologists delight in showing us that people's behaviour is biased by the position of items on the page—even when that position has no predictive validity. It can also be biased by the weather, by whether people's emotions were suppressed during a movie, by whether people have just engaged in the act of proofreading, or even by whether they have just eaten radishes instead of chocolates (Bargh and Chartrand, 1999; Baumeister et al., 1998).

1.3. Distortions and individual differences impacting the goodness of metacognitive cues

Furthermore, even when people use their metacognitive cues, these cues, themselves, are not always good indications of what to study: it is rather easy to distort people's metacognitions. For example, if people are given both the cue and the target (or, in common parlance, the question and the answer), and asked to estimate the difficulty of the problem or how easy it will be to remember, their metacognitions are systematically distorted in the direction of believing that the difficulty is much less than it really is (Bjork, 1994). Given that metacognitions can be distorted, people's study time allocation based on the metacognitions, might also be expected to be suboptimal. Although metacognitive control of learning, even in humans, is something that may be used to enhance learning, it does not always do so. When metacognitions are accurate, and when they are used strategically, as will be detailed next, the learning that is the goal of such self-control strategies, can show enhancement. People can and, in some situations, do, escape from being stimulus bound, and take control over their own learning and cognition.

Under the best of conditions, there is a method – studied first by Nelson and Dunlosky (1991) – that produces highly accurate metacognitions (see, e.g., Thiede and Dunlosky, 1994). The method is called the 'cue only delayed judgment of learning' method. To illustrate, suppose a person learns a series of vocabulary items, where the cue is in regular font and the target, that is, what

the participant is being asked to learn, is in uppercase font, such as follows: Given to complaining, fretful, whining: QUERULOUS; Praising, laudatory, eulogistic: ENCOMIASTIC; Partial shadow (in an eclipse): PENUMBRA. To make cue only delayed judgments of learning, the person first needs to wait some time after the original learning episode, allowing at least enough time to ensure that the target is no longer in working memory. Then, s/he would be given the cues alone and asked to assess the probability that, later, s/he will remember the target. So, s/he might be given 'Partial shadow (in an eclipse)' and be asked to give a rating between, say, 0 and 1.00, of the estimated probability of correctly recalling the target an hour later. Typically, the participant then receives the test an hour later. The individual ratings given while making the judgments are then correlated with whether particular items were correct or incorrect on the test, resulting in a gamma correlation for each participant that indicates the extent to which they knew what they would know. Using this method, the correlations between the judgments and later test performance are nearly always extremely high.

What people appear to do (see Son and Metcalfe, 2005 and Metcalfe and Finn, 2008b) to make these ratings is first determine whether they recognize the cue. If they do not recognize it, they stop and very quickly give that item a very low value. If they recognize it, then they go on to try to retrieve the target. If they do retrieve the target, they give the item a high judgment of learning, but the value of the judgment decreases depending on the length of time it takes them to do the retrieval—presumably because the longer retrieval time, or the lack of fluency, indicates that the target is less well learned than are those retrieved quickly. This method of making the judgments produces extremely accurate assessments of what has and has not been learned.

Note that the passage of time, and the forgetting that ensues with it will automatically be registered in metacognitive judgments made this way. If a judgment is made soon after study, and little forgetting has occurred, the person will recognize the cue easily, retrieve the target quickly, and give a high judgment of learning—indicating that immediate study will have little benefit. If, however, a considerable amount of time has passed, and much forgetting has occurred, the person may be less likely to recognize the cue, and if s/he does, may have difficulty retrieving the target. The resultant lowered judgment of learning will indicate to the learner that restudy of that item is in order. Thus, forgetting due to the passage of time or to interference, will be well assessed by the delayed cue only method of making the judgments of learning.

Using this method of assessing their learning, one might expect people to be in a good position to use their metacognitions effectively. So, the next questions are, how do people choose the items to study and how long do they persist, once studying an item?

1.4. Study choice and persistence: declining study of already mastered items

All theories (e.g., Dunlosky and Hertzog, 1998; Metcalfe, 2009; Metcalfe and Kornell, 2003, 2005; Thiede and Dunlosky, 1999) agree that metacognitively informed learners, should, and do, decline to study items they have truly mastered. This tendency alone accounts for the nearly ubiquitous finding (Son and Metcalfe, 2000) of a negative correlation between judgments of learning and study choice. When an individual already knows a subset of the to-be-remembered items, those items will be given high judgments of learning. Because the items are already learned, adults rarely choose to study them. For example, Son (2004) asked people to make judgments of learning, and then gave them three choices: They could decline study entirely, study again immediately, or study again, but later. These people declined to study only the items with very high judgments of learning. (It is interesting that a substantial number of items with even the very highest judgments

of learning were not declined – suggesting that people may not believe they absolutely ‘know’ an item, when they give it a 100% judgment of learning. They may simply mean – as the above explanation of how they make judgments of learning suggests – that they can retrieve the item quickly at that time.) Nevertheless, the tendency to decline study was relatively high for the highest judgments of learning and very low for items with even slightly lower judgments of learning. There appears to be some marker – and it appears to be a metacognitive marker – indicating that immediate study will have little benefit. We will suggest, shortly, that this marker is the metacognitive ‘scent’ of the item that allows such a decision, and that if an item has such a scent, declining immediate study, or deferring further study until later, is wise.

1.5. Study the items nearest to being mastered before turning to more difficult items

A number of experiments point to the idea that people select, for study, items in their own perceived Region of Proximal Learning (e.g., Kornell and Metcalfe, 2006; Metcalfe, 2002; Metcalfe and Kornell, 2007). This region is one in which the items are not already well mastered, nor are they virtually impossible for that particular person to learn. Rather they are items that are just beyond their current mastery. . . a kind of Goldilocks principle: not too easy, not too difficult, but just right. A person can assimilate items that are in the Region of Proximal Learning into their knowledge store quickly and with relatively little effort, but without that bit of effort being exerted, the item may be incorrectly recalled, or missed on the test. The person will recall items that are too easy, that is, already mastered, with or without study; investing time and effort in them is unproductive. Difficult items require a great deal of time and effort to master. It is not cost effective to seek out these most difficult items while easier items with equal payoff remain unstudied: the time and effort spent on the very difficult items may produce substantial opportunity costs (i.e., many relatively easy items being missed that could otherwise be remembered).

The idea of a range of difficulty just beyond what the person has already mastered – a range that is particularly open to learning – is not new (see, Atkinson, 1972; Piaget, 1952; Vygotsky, 1987). These may be the materials and items that are particularly open to scaffolding, hold the greatest promise for learning, and hold the greatest interest for the student (see, Berlyne, 1978). All else being equal, people should choose to study the easiest as-yet-unlearned items, not only because these items have the greatest interest for them, but because they will also give the maximum payoff in terms of percent correct on the test, with the least effort, and in the minimal time.

This scheme depends, of course, on all of the items being equally weighted on the test. If some more difficult items are accorded much more value on a criterion test (either in terms of the probability that they will appear on the test or in terms of the point value they will be assigned on the test), then it may be optimal to attend to the more difficult items and accord them study time, rather than focusing exclusively on easier items that are credited by many fewer points. Thus, as Ariel et al. (2009) astutely pointed out, people’s preferences, and the optimal predictions about their choices, need to be weighted by the value of the materials in the situation, and not only by their perceived degree of learning. Furthermore, as people master items, the items change status; they convert from being the highest return items to being already mastered items, on which immediate study time would be wasted.

1.6. Stop studying when the learning returns are negligible

Once an item has been chosen, the person then goes about the actual business of studying that item – doing what they can to learn

it. But, at some point they stop studying a given item and turn to another item. Clearly, some kind of monitoring is recruited to make this decision. The situation is one that resonates with the problem in optimal foraging theory of how the animal decides to leave a patch and turn to another. There have been two stop rules proposed in the literature on human study time allocation: an absolute criterion rule and a rate rule.

The absolute criterion rule says that people continue studying an item until they have learned it to criterion – the absolute degree of learning must be monitored (Thiede and Dunlosky, 1999). In contrast, the rate rule is more relative and is based on the derivative of the rate of learning function rather than on an absolute criterion. According to this second rule, people study an item as long as they perceive themselves to be learning at a sufficiently fast pace, or as long as their perceived information uptake function has not yet reached some, presumably low, rate. To use this rule, people need to be able to monitor the derivative of their perceived learning function, not their absolute level of learning. Although there is little research on this fascinating issue, the extant evidence favors the second rule: people stop studying when their rate of learning is low (see Metcalfe and Kornell, 2005). The rate can be low for two reasons: (1) because the person has already learned the item and hence further study is not increasing learning, or (2) because the item is so difficult the person makes no progress in learning it. This first case produces the same kind of empirical results as does the absolute rule: people do not study once an item is learned. But it produces the same result for a different reason – not because learning reaches an absolute criterion but rather because perceived learning plateaus. The second reason a person may stop studying an item, according to this rate rule is more interesting, and it distinguishes between the two rules. According to the rate rule, a person may stop studying because the item is intractable and efforts at learning do not produce tangible results. The absolute rule predicts that a person will persevere in studying extremely difficult items, once chosen, because such items are far from the absolute criterion of being learned. The rate rule predicts that, when the going gets too tough, people will give up and turn to other easier items.

Studies have shown, not only that people seem to make their choices as outlined above – declining study of items that are already well known, choosing the easiest as-yet-unlearned items first, and stopping study of one item to turn to another when the rate of perceived information uptake of the first item becomes too low – but also that so doing is effective. Kornell and Metcalfe (2006) conducted a number of studies in which people made their own choices for further study of half of the items in a to-be-tested set. The researchers then either honored those choices or gave people the other non-chosen half of the items to study. The experiments consistently showed, with college-student adults, at least, that honoring people’s own choices for study resulted in superior later test performance over dishonoring their choices. Thus, although we cannot determine anything like true ‘optimality’ in this situation, we can, at least measure effectiveness. These data indicate that the metacognitively guided choices of adult college students were effective.

2. The analogy

The following characteristics of human study time allocation appear to be related to foraging strategies that animals use. First, people appear to use specific metacognitive cues to assist their study time allocation, and the simplest way in which they use them is to determine which items they will decline to study – they decline items because they already know them. These markers and their use relate to how animals detect cues – the simplest being

scent markings – and use those cues to allow them to implement win-shift strategies. Second, people pluck off the easiest as-yet-unlearned items first for learning, before they turn to more difficult items – just as other species use the Goldilocks principle in foraging. Third, people have a stop rule by which they give up on one to-be-learned item and turn to another, just as animals obey a stop rule when shifting from one patch to another. Fourth, the value that each item is assigned on the final test, or the probability that it has of showing up on the test, influences which items people choose to study and how long they study them – just as the learned nutritional or energy value of the food influences choices and perseverance in foraging. And, finally, individual differences are apparent: study time allocation strategies differ depending up on the expertise of the student just as some animals forage close to optimally while others do not. We discuss the similarities between the two domains, as well as possible alignments in mechanism that give rise to these similarities, below.

2.1. *Specific cues are used to decline study of already known items: the win-switch strategy*

As described above, people decline study of items they have just studied, or that they know they already know. This phenomenon seems similar to the fact that animals and insects do not return – at least immediately – to a food source that has just been depleted. One of the most prominent examples of the effect of a depleted food source is in the honeybee (*apis mellifera*). Careful studies have shown that once a bee takes the nectar from a flower, it moves on to new flowers. Furthermore, once it has gone to a new flower, the bee does not revisit the first, already foraged flower, even when it is closer than the flowers to which it turns. One might wonder whether the bee has memory for previously foraged flowers, such that it knows that the nectar is now gone – perhaps a kind of primordial metacognition.

Furthermore, some, but not all species of bee, have a fairly tight temporal calibration between the replenishment rate of a flower and the time they return to that flower. They do not return immediately, but they do return later, and some do so just as the flower replenishes its nectar. It is notable that adults in Son's (2004, 2005) study choice experiments also showed good calibration of the temporal components of restudy: her participants knew that they should return to certain already studied items later, but not immediately – for spaced (but not massed) practice.

Although people enact these strategies by using their metacognitions, no evidence suggests that bees do. Instead, the bee does it by scent marking. Each time it enters a flower, the bee rubs a bit of a pheromone onto the flower (Yokoi and Fujisaki, 2007). They can later detect this scent, which is the marker that tells them that they have already visited said flower. Some bees are sensitive to the scent markings of other bees from their hive (Stout et al., 1998), and some even to the scent of other species – an adaptation that saves them from wasting time foraging on what has already been foraged by themselves or others, and which will not pay off in a nectar haul.

The scent left on the flower by some, but not all bee species (Yokoi and Fujisaki, 2007) wears off at just the rate that the flower replenishes nectar. This is an adaptation allowing the bee to time return visits to just the moment at which they will obtain profitable nectar. This is akin to a highly tuned adaptation noted by Son (2005) whereby people know when to revisit to-be-remembered items, and do so at the proper temporal spacing. Interestingly, although Son found that adults are quite accurate in this timing, children are not (Son, 2005) – just as certain species of bees are well tuned to temporal replenishment parameters while others are not. In some bee species, the timing of the dissipation of their own scent is tied to the rate of replenishment rates of the flowers on which they forage;

in other species scent depletion is not. Those that specialize on a particular kind of flower, which replenishes at a particular rate, are more likely to produce scents timed to the replenishment rates of that flower, while those that forage more broadly usually do not have such finely tuned scent marking.

Thus, in bees, a mechanism for declining to return to a just-visited source is overt and physical, rather than depending upon memory for the source. The mechanism is neither cognitive nor metacognitive. Other animal species, also, do not return to sources just visited. In some, the mechanisms are thought to be memory based. For example, there are many studies on foraging rats that show that they, too, avoid just-visited sites, choosing a 'win-shift' strategy. Timberlake and Hoffman (2002) argued against the scent-marking hypothesis as the mechanism that allows the shift, however, observing that the rats in their experiments, though displaying the shift behaviour, were never seen to body mark or leave deposits. They also noted that in the wild, in the kind of sandy substrate that is the natural home of these rats, winds would likely obliterate scent marks. Hence, scent would not be a reliable cue. Instead, some kind of memory marking appears to be implicated.

Similarly, although bats, too, use win-shift strategies, and soon learn to not return to an already foraged site, Winter and Stich (2005) have argued that they do so, not by scent marking, but instead through a memory-based strategy. As such, these and other species may provide a bridge toward a purely metacognitive strategy that people use in information foraging.

A win-switch strategy, in some species, is odd, from a certain perspective. It depends on the animals acting against the principles of reinforcement. Presumably, the rat receives reinforcement by obtaining food in a particular location. Should that positive reinforcement not strengthen the probability that the organism return to just that place? Instead though, the animal does just the opposite. It does not immediately revisit the location at which it received a reinforcer.¹ A similar problem arose in strength-based theories of rehearsal in human information processing, such as that of Metcalfe and Murdock (1981). In that model, a simple rule concerning what to rehearse next, that was later abandoned in the model for a win-shift rule, was that the next item that would enter the model's consciousness buffer and be rehearsed would simply be the strongest item. But this rule of rehearsing the strongest item – while simple and even compelling – immediately encountered the problem that the strongest item is virtually always the last item in consciousness, that is, the last item that was rehearsed. Hence, a model using this rule gets stuck, and always rehearses what was just rehearsed, never turning to weaker items. Yet, the data on human rehearsal patterns (see, Murdock and Metcalfe, 1978) indicated that, aside from the patterns of people who have the odd syndrome called echolalia, or who have Korsakoff amnesia, (see Cermak et al., 1976), people virtually never get stuck in a loop in which they repeat the same thing over and over. Instead, that strongest item appeared to be suppressed or inhibited: a win-shift strategy was engaged. (The rule in the model was eventually changed to allow it to do the same thing that people do.)

¹ There may be several behavioral reasons for this, however. For some organisms, a re-visit might occur after depletion—which produces extinction. As importantly, the 'value' of particular foodstuffs changes as satiation occurs permitting at least two phenomena to occur: contrast effects and a dramatic diminution in the impact of the item as a reinforcer. For example, something that is reinforcing under one condition when presented alone (e.g., fried liver) may not be reinforcing in the presence of other foodstuffs (e.g., filet mignon); equally what is reinforcing at one state of deprivation (e.g., liver to a starving person) may not be reinforcing under another state (e.g., fried liver to a person who has just finished a gourmet meal). Although some might think it forced, a reasonably sensitive behavioral model may provide predictions superior to the simple-minded view of reinforcement outlined above.

2.2. The Goldilocks principle

As noted above, adults attempt to focus their study in their own Region of Proximal Learning, preferring items that are not so easy that they have already been learned, and not so difficult that they are impossible to learn, but just right. There is considerable evidence that other species do the same when foraging. For example, certain seabirds, particularly, oystercatchers, selectively choose bivalves that are smaller than optimal as estimated by the maximal dietary return per oyster, but they also spurn very small oysters (Rutten et al., 2006). The driving factor that determines this 'just right' size appears to be the maximum shell size they can pry open without damaging their beaks. How they know – without having first broken their beak – what size shell to choose may be partly genetically determined. Studies have shown that oystercatchers who do not do damage to their beaks by attempting too large oysters, have much higher survival and reproductive rates. Those birds that prefer oysters of the just right size might pass on this preference to their offspring, and hence assure a survival advantage. Nevertheless, the fact that some birds break their beaks at a high rate indicates that such a selection effect is not pervasive. Additionally, there might be some cues – other than innate preferences for a certain sized oyster – such as pain – that might foreshadow bad things to come from a potentially broken beak, which might dissuade the oystercatchers from prying outside the range of their beaks.

Oystercatchers are not the only species to choose the just right materials. Indeed, this idea is a mainstay of optimal foraging theory (Pyke et al., 1977). Usually species take as prey other species that are smaller than they are. They also take into account their own energy expenditure, and seek near sources and low hanging fruit, rather than far sources.

The idea that foraging occurs within the 'just right' range is so widely held that a recent paper in PNAS (Berlow et al., 2008) – commenting on a longer but similarly oriented article by Petchey et al. (2008) – was entitled "The 'Goldilocks factor' in food webs". The Petchey et al. (2008) model of optimal foraging, which emphasizes body size, has successfully predicted predacious and herbivorous feeding interactions. It has less success in predicting parasitic and pathogenic interactions, however.

2.3. The stop rule or when to leave

As noted above, two stop rules have been proposed in the human study time allocation literature. The first, absolute rule, proposed within the discrepancy reduction framework (Dunlosky and Hertzog, 1998), said that, once they have started, people continue to study an item until its degree of learning meets an internal criterion (of being learned or as being as much learned as the person thinks is necessary), and then, and only then, do they stop and turn to other items. Absolute criterion rules resemble the Fixed-number (or Fixed-amount) rule taken from the foraging literature – if one thinks of achieving an absolute threshold as a person continuing to study an item until acquiring some constant, X , (a number of bits, or amount) of information (see, e.g., Stephens and Krebs, 1986; Wilke et al., 2009). Some animals appear to do this, but only in very select situations, consistent with the notion that it is only used by people in pathological cases: the individual who uses it consistently may become trapped in the untenable situation of studying an impossible-to-learn item forever. There is no guarantee that very difficult items would ever reach an absolute criterion, and, there is no opt out (stop) procedure.

The second stop rule (Metcalfe and Kornell, 2005) – the rate rule – while primitive by comparison to the Marginal Value Theorem used in optimal foraging theory, is closer in spirit. The rate rule says, essentially: continue to study as long as the rate of perceived

information uptake continues apace, but stop if learning is too slow. This rule bears a non-trivial resemblance to Charnov's classic 1976 Marginal Value Theorem which predicts that an organism will leave a patch (cease foraging in the patch) when its instantaneous rate of return from the current patch falls below the average return rate from the environment. Similar to the marginal value theorem, Waage's (e.g., 1979) threshold rule predicts that indicators of the size of the patch and host density determine the tendency to stay in a patch. When the tendency to stay in a patch drops below a certain threshold, the organism leaves the patch. Using these strategies (Charnov's rule, Waage's rule, or the rate rule) appears to produce optimal fitness when foraging for food or for information. The Marginal Value Theorem is more complex than the Rate Rule in study persistence. It takes into account not only the rate of return of energy from the patch being foraged (which is analogous to the perceived learning) but also the expectation of the return rate from competing patches that could be foraged instead. Furthermore, it also factors in the switch costs – the energy expenditure involved in switching from the current patch, and turning to, including traveling to, another patch. Thus, the wealth, and hence, seductiveness, of the entire foraging landscape is included in the foraging stop rule, making the study time allocation stop rule seem primitive by comparison. While the rate rule in the human study time allocation literature constitutes an advance over the absolute criterion rule, human theory could benefit from study of the more complex, but also more realistic, persistence equations used in optimal foraging theory.

2.4. Value influences choices and perseverance

Humans use the probability that an item will be on the test as well as the perceived difficulty of an item, to determine whether or not to choose to study an item. They also respond to factors such as position of items, the salience or importance of items, and the perceptual fluency of items, in determining their choices. These latter kinds of predispositions may or may not be helpful to study behaviour. For example, Shah and Oppenheimer (2007) showed that people weighted information in a clear font more heavily than information in a less clear font; in focus information was also weighted more heavily than out of focus information, and they accord more credibility to financial information from a brokerage firm whose name was easy to pronounce than the same information from a hard-to-pronounce firm. Indeed, Rhodes and Castel (2008) showed that font size, itself, directly influenced people's metacognitive judgments, presumably because the apparent fluency that accompanied large fonts translated into a feeling that the item had been learned.

Other species, too, respond differentially based on perceptual factors and ease of processing. Spaethe et al. (2001) showed that the sensory-perceptual processes that constrain the search for food, and not just prey density and other factors normally considered in optimal foraging theory, are crucial. They evaluated the flight of bumblebees searching for artificial flowers of various sizes and colors. Perceptual factors were highly significant in determining these pathways. When flowers were large, search times correlated well with the color contrast of the targets with their background, and the bees' used their UV, blue, and green receptors. Targets that made poor color contrast with their backdrop, took longest to detect, even though brightness contrast with the background was pronounced. When searching for small targets, the bees changed their strategy, flying slower and closer to the ground, and, they used only the green receptor for detection. The authors concluded that foraging speed was not limited exclusively by factors such as prey density, or flight energetics, but that factors such as the bees visual information processing also made a significant contribution.

In addition, other species – even insects – can learn to make particular choices – rather than simply being hard-wired to use a particular foraging strategy. Things like the color of the flower that has a high return can be learned. Goulson and Cory (1993), for example, showed learning in foraging preferences of the green-veined white butterfly. When nectar was freely available from flowers, independent of color, equal numbers of butterflies visited each of two colors, but individual butterflies exhibited flower color preferences, foraging particularly on one color or the other. When the experimenters manipulated the availability of nectar to flowers of one color only, the butterflies developed a preference for this color. This preference persisted when both flower colors were refilled but could subsequently be reversed to the other flower color following further experience with the color-nectar relationships. These learned preferences have obvious selective benefits in the field, enabling butterflies to prefer flower species which experience has shown are a rich source of nectar, and to adapt to temporal and spatial changes in nectar availability. People make tradeoffs in variables like learnability and probability on test, as Ariel et al. (2009) have shown.

This kind of learning may be theoretically explained by a kind of optimal foraging behaviour called area-restricted search. Species ranging from nematodes through humans tend use behavioural strategies (e.g., slowing movement and remaining longer in a vicinity) that produce foraging in locations or environments near where they, in the past, have found resources. The resemblance between the learning strategies shown in Ariel et al.'s (2009) experiments and area restricted search may be more than merely superficial, for, according to Hills (2006), neural processes that support goal-directed area restricted search (primarily dopaminergic systems) have been exapted (shifted in function during evolution) in the service of learning, memory retrieval, and other cognitive processes (see Hills, 2006 for a beautiful review of this literature; in addition, see Kareiva and Odell, 1987; Grunbaum, 1998).

Other species make foraging tradeoffs as do humans in allocating study time. Arcis and Desor (2003) showed that, although rats prefer to forage in an area that offers cover when all else is equal, and they prefer to forage in areas where the food density is greater, again, all else being equal, when these two factors are crossed, the rats take both into account. The authors suggested that foraging decisions relied mostly, but not exclusively, on safety needs. They also suggested that food deprivation might change that priority.

2.5. Individual differences in the quest for optimality

Although some bees produce scents that monitor the replenishment of the flowers they forage on, not all do. Furthermore, some species pick up on conspecific scents, and some do not. The analogy suggests that we might look for individual differences in perception and use of metacognitive cues. Cosentino et al. (2007) have found that some people do not use some cues. For example, not everyone uses the memory for past test heuristic, a cue that lowers metacognitive judgments if one remembers that one was previously wrong, regardless of one's current state. Elders who have a memory deficit and who are aware of having such a deficit use the memory for past test heuristic, that is, they underweight the learning of items on which they were incorrect on a past test. But elders with a memory deficit but who are unaware of having the deficit do not use this metacognitive cue. They think that they know these items well—too well. Finn and Metcalfe (2007, 2008) have shown that normal college students are nearly always sensitive to this cue, so the finding of a population that does not use it is of interest. Furthermore, Dunlosky (personal communication) has found that children in kindergarten do not use the memory for past test heuristic either.

Bjork (1994) noted that college students often do not use the strategy of spacing practice, while Son (2008) showed that at least some college students do use this strategy, but that children do not (Son, in press). These and other individual differences (Griffin et al., 2008) in humans' metacognition may be of great interest, but have been under explored, in part, perhaps, because it is often assumed that human cognitive capabilities are universal.

In contrast, species and individual differences are the mainstay of evolutionary theory. The organism that uses a particular strategy more effectively than others has an adaptive and hence survival advantage. Hence, these kinds of differences have received much more attention in the animal foraging literature. For example, Ritchie (1990) determined the optimal diets for individual Columbian ground squirrels in northwestern Montana. Body mass, daily activity time, and vegetation consumption rates for individuals were measured in the field, along with the average water content of vegetation at each ground squirrel colony along with literature estimates of average daily energy requirements as a function of body mass and digestible energy content of vegetation. He then determined the consequences of selecting an optimal diet for energy intake and fitness. Growth rate, yearly survivorship, and litter size increased with energy intake, and the optimal foragers had six times the reproductive success of deviators by age three.

Rutten et al. (2006) showed that oystercatchers that did not choose food in the optimal range, and, as a result, broke their beaks, showed a 23% reduction in food intake and an associated reduction in body mass. Because lower body mass is associated with higher mortality probability, there were long-term survival costs associated with an inappropriate feeding strategy.

Finally, there may be more to human study time behaviour than just optimizing study and subsequent performance on the test itself. Other considerations weigh in on animal foraging behaviour, as noted by Green et al. (1985) who observed the foraging behaviour of male cunner fish. Individual territorial male cunners were observed in the morning and afternoon during pre-spawning, spawning, and post-spawning seasons in Conception Bay, Newfoundland. The cunners fed at a constant rate throughout the day in all seasons, with one exception. When they were spawning they shifted their foraging to just the morning. This shift ensured that their foraging did not conflict with afternoon spawning – a strategy shift that the authors viewed as a true optimal foraging strategy, where fitness rather than energy intake was optimized. It appears there is more to passing the test than just the test.

3. Conclusion

The analogy between human study time allocation and animal foraging may allow insight into weaknesses in current theory of human study time allocation and push the further development of those theories. For example, a number of variables extensively explored in animal foraging remain unexplored in human study time allocation, but they may be ripe for exploration. The stop rule in the study time allocation domain should probably take into account all of the variables that the analogous stop rule does in foraging theory. At the moment, even the rate rule does not factor into the equation the influence of the alternatives that could be but are not being studied at any given time. But as in foraging theory, this factor must surely matter. The change in the theory that could result by scrutinizing considerations such as this one that already are well studied in foraging theory could advance current theory of study time allocation. The kinds of individual differences that exist across people as well as their consequences deserve consideration. Following the lead of foraging theory in scrutiny and evaluation of different strategies could advance the field. Finally, the notion of trying to come up with a formula that would optimize study is itself a productive notion. Whether we can actualize such a rule

remains to be seen. To date, unlike in foraging theory, we have no definition even of the currency of learning that could be used to ground such optimization equations. But serious consideration of the problem might well lead to advances. More thorough investigation of the strategies of study time allocation, and their relation to similar strategies in food foraging, might not only allow us to advance theory in the former domain, but in addition, it might put us in a position to better teach them and make learning itself more efficient and enjoyable.

Acknowledgements

This research was supported by National Institute of Mental Health grant RO1-MH60637. The authors are solely responsible for the views expressed herein. We thank Lisa K. Son for her help and comments.

References

- Arcis, V., Desor, D., 2003. Influence of environmental structure and food availability on the foraging behaviour of the laboratory rat. *Behavioural Processes* 60, 191–198.
- Ariel, R., Dunlosky, J., Bailey, H., 2009. Agenda-based regulation of study-time allocation: when agendas override item-based monitoring. *Journal of Experimental Psychology: General* 138, 432–447.
- Atkinson, R.C., 1972. Optimizing the learning of a second language vocabulary. *Journal of Experimental Psychology* 96, 124–129.
- Bargh, J.A., Chartrand, T.L., 1999. The unbearable automaticity of being. *American Psychologist* 54, 462–479.
- Baumeister, R.F., Bratslavsky, E., Muraven, M., Tice, D.M., 1998. Ego depletion: is the active self a limited resource? *Journal of Personality and Social Psychology* 74, 1252–1265.
- Berlow, E.L., Brose, U., Martinez, N.D., 2008. The “Goldilocks factor” in food webs. *Proceedings of the National Academy of Sciences of the United States of America* 105, 4079–4080.
- Berlyne, D.E., 1978. Curiosity and learning. *Motivation and Emotion* 2, 97–175.
- Bjork, R.A., 1994. Memory and metamemory considerations in the training of human beings. In: Metcalfe, J., Shimamura, A.P. (Eds.), *Metacognition: Knowing about Knowing*. MIT Press, Cambridge, MA, pp. 185–206.
- Bradshaw, C.M., Szabadi, E., Bevan, P., 1976. Behaviour of humans in variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behaviour* 26, 135–141.
- Cermak, L.S., Naus, M.J., Reale, L., 1976. Rehearsal strategies of alcoholic Korsakoff patients. *Brain and Language* 3, 375–385.
- Charnov, E.L., 1976. Optimal foraging, the Marginal Value Theorem. *Theoretical Population Biology* 9, 129–136.
- Cosentino, S.A., Metcalfe, J., Butterfield, B., Stern, Y., 2007. Objective metamemory testing captures awareness of deficit in Alzheimer's disease. *Cortex* 43, 1004–1019.
- Dunlosky, J., Hertzog, C., 1998. Training programs to improve learning in later adulthood: helping older adults educate themselves. In: Hacker, D.J., Dunlosky, J., Graesser, A.C. (Eds.), *Metacognition in Educational Theory and Practice*. Erlbaum, Mahwah, NJ, pp. 249–275.
- Finn, B., 2008. Framing effects on metacognitive monitoring and control. *Memory & Cognition* 36, 815–821.
- Finn, B., Metcalfe, J., 2007. The role of memory for past test in the underconfidence with practice effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 33, 238–244.
- Finn, B., Metcalfe, J., 2008. Judgments of learning are influenced by memory for past test. *Journal of Memory and Language* 58, 19–34.
- Gaissmaier, W., Schooler, L.J., 2008. The smart potential behind probability matching. *Cognition* 109, 416–422.
- Goulson, D., Cory, J.S., 1993. Flower constancy and learning in foraging preferences of the green-veined white butterfly *Pieris napi*. *Ecological Entomology* 18, 315–320.
- Green, G.M., Martel, G., Kingsland, E.A., 1985. Foraging time allocation in a territorial fish: influence of reproductive activities. *Marine Ecology Progress Series* 24, 23–26.
- Griffin, T.D., Wiley, J., Thiede, K.W., 2008. Individual differences, rereading, and self-explanation: concurrent processing and cue validity as constraints on meta-comprehension accuracy. *Memory & Cognition* 36, 93–103.
- Grunbaum, D., 1998. Using spatially explicit models to characterize foraging performance in heterogeneous landscapes. *The American Naturalist* 151, 97–115.
- Hampton, R.R., 2001. Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences of the United States of America* 98, 5359–5362.
- Hampton, R.R., 2005. Can monkeys discriminate between remembering and forgetting? In: Metcalfe, J., Terrace, H. (Eds.), *The Missing Link in Cognition: Origins of Self-Reflective Consciousness*. Oxford University Press, New York, NY, pp. 272–295.
- Herrnstein, R.J., 1961. Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behaviour* 4, 267–272.
- Hills, T.T., 2006. Animal foraging and the evolution of goal-directed cognition. *Cognitive Science* 30, 3–41.
- Horne, P.J., Lowe, C.F., 1993. Determinants of human performance on concurrent schedules. *Journal of the Experimental Analysis of Behaviour* 59, 29–60.
- Kamil, A.C., 1983. Optimal foraging theory and the psychology of learning. *American Zoologist* 23, 291–302.
- Kareiva, P., Odell, G., 1987. Swarms of predators exhibit “prey taxis” if individual predators use area-restricted search. *American Naturalist* 130, 233–270.
- Kornell, N., Metcalfe, J., 2006. Study efficacy and the region of proximal learning framework. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32, 609–622.
- Kornell, N., Son, L.K., Terrace, H., 2007. Transfer of metacognitive skills and hint seeking in monkeys. *Psychological Science* 18, 64–71.
- Metcalfe, J., 2002. Is study time allocated selectively to a region of proximal learning? *Journal of Experimental Psychology: General* 131, 349–363.
- Metcalfe, J., 2009. Metacognitive judgments and control of study. *Current Directions in Psychological Sciences* 18, 159–163.
- Metcalfe, J., Finn, B., 2008a. Evidence that judgments of learning are causally related to study choice. *Psychonomic Bulletin & Review* 15, 174–179.
- Metcalfe, J., Finn, B., 2008b. Familiarity and retrieval processes in delayed judgments of learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 34, 1084–1097.
- Metcalfe, J., Finn, B., 2009. Metacognition and control of study in grade 3–5 children. In: O'Donnell, C.L. (Chair), Albro, E.R. (Chair), *Examining Self-regulated Learning and Metacognition in Authentic Kindergarten Through College Classrooms*. Symposium conducted at the Annual Meeting of the Association for Psychological Sciences, San Francisco.
- Metcalfe, J., Kornell, N., 2003. The dynamics of learning and allocation of study time to a region of proximal learning. *Journal of Experimental Psychology: General* 132, 530–542.
- Metcalfe, J., Kornell, N., 2005. A region of proximal learning model of study time allocation. *Journal of Memory and Language* 52, 463–477.
- Metcalfe, J., Kornell, N., 2007. Principles of cognitive science in education: the effects of generation, errors and feedback. *Psychonomic Bulletin and Review* 14, 225–229.
- Metcalfe, J., Murdock, B.B., 1981. An encoding and retrieval model of single-trial free recall. *Journal of Verbal Learning and Verbal Behaviour* 20, 161–189.
- Murdock, B.B., Metcalfe, J., 1978. Controlled rehearsal in single trial free recall. *Journal of Verbal Learning and Verbal Behaviour* 17, 309–324.
- Nairne, J.S., Pandeirada, J.N.S., Gregory, K.J., Van Arsdall, J.E., 2009. Adaptive memory: fitness-relevance and the hunter-gatherer mind. *Psychological Science* 20, 740–746.
- Nelson, T.O., Dunlosky, J., 1991. When people's judgments of learning (JOLs) are extremely accurate at predicting subsequent recall: the “delayed-JOL effect”. *Psychological Science* 2, 267–270.
- Petchey, O.L., Beckerman, A.P., Riede, J.O., Warren, P.H., 2008. Size, foraging, and food web structure. *Proceedings of the National Academy of Sciences of the United States of America* 105, 4191–4196.
- Piaget, J., 1952. *The Origins of Intelligence in Children*. International Universities Press, New York.
- Pirolli, P.L., Card, S.K., 1999. Information foraging. *Psychological Review* 106, 643–675.
- Pyke, G.H., Pulliam, H.R., Charnov, E.L., 1977. Optimal foraging: a selective review of theory and tests. *The Quarterly Review of Biology* 52, 137–154.
- Rhodes, M.G., Castel, A.D., 2008. Memory predictions are influenced by perceptual information: evidence for metacognitive illusions. *Journal of Experimental Psychology: General* 137, 615–625.
- Ritchie, M.E., 1990. Optimal foraging and fitness in Columbian ground squirrels. *Oecologia* 82, 56–67.
- Rutten, A.L., Oosterbeek, K., Ens, B.J., Verhulst, S., 2006. Optimal foraging on perilous prey: risk of bill damage reduces optimal prey size in oystercatchers. *Behavioural Ecology* 17, 297–302.
- Shah, A.K., Oppenheimer, D.M., 2007. Easy does it: the role of fluency in cue weighting. *Judgment and Decision Making* 2, 371–379.
- Smith, J.D., Shields, W.E., Washburn, D.A., 2003. The comparative psychology of uncertainty monitoring and metacognition. *Behavioural and Brain Sciences* 26, 317–339.
- Son, L.K., 2004. Spacing one's study: evidence for a metacognitive control strategy. *Journal of Experimental Psychology: Learning, Memory and Cognition* 30, 601–604.
- Son, L.K., 2005. Metacognitive control: children's short-term versus long-term study strategies. *The Journal of General Psychology* 132, 347–363.
- Son, L.K., 2008. The spacing effect, metacognitively speaking. In: Paper presented at the 1st Annual Meeting of the Subway Summit. Fordham University, New York, NY.
- Son, L.K. Metacognitive control and the spacing effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, in press.
- Son, L.K., Metcalfe, J., 2000. Metacognitive and control strategies in study-time allocation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26, 204–221.
- Son, L.K., Metcalfe, J., 2005. Judgments of learning: evidence for a two-stage model. *Memory & Cognition* 33, 1116–1129.

- Spaethe, J., Tautz, J., Chittka, L., 2001. Visual constraints in foraging bumblebees: flower size and color affect search time and flight behaviour. *Proceedings of the National Academy of Sciences of the United States of America* 98, 3898–3903.
- Stephens, D.W., Krebs, J.R., 1986. *Foraging Theory*. Princeton University Press, Princeton, New Jersey.
- Stout, J.C., Goulson, D., Allen, J.A., 1998. Repellent scent marking of flowers by a guild of foraging bumblebees (*Bombus spp.*). *Behavioural Ecology and Sociobiology* 43, 317–326.
- Terrace, H.S., Metcalfe, J., 2005. *The Missing Link in Cognition: Origins of Self-reflective Consciousness*. Oxford University Press, New York, NY.
- Terrace, H.S., Son, L.K., 2009. Comparative metacognition. *Current Opinion in Neurobiology* 19, 67–74.
- Thiede, K.W., Anderson, M.C.M., Theriault, D., 2003. Accuracy of metacognitive monitoring affects learning of texts. *Journal of Educational Psychology* 95, 66–75.
- Thiede, K.W., Dunlosky, J., 1994. Delaying students' metacognitive monitoring improves their accuracy at predicting their recognition performance. *Journal of Educational Psychology* 2, 290–302.
- Thiede, K.W., Dunlosky, J., 1999. Toward a general model of self-regulated study: an analysis of selection of items for study and self-paced study time. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25, 1024–1037.
- Timberlake, W., Hoffman, C., 2002. How does the ecological foraging behaviour of desert kangaroo rats (*Dipodomys deserti*) relate to their behaviour on radial mazes? *Learning & Behaviour* 30, 342–354.
- Vulkan, N., 2000. An economist's perspective on probability matching. *Journal of Economic Surveys* 14, 101–118.
- Vygotsky, L.S., 1987. In: Rieber, R.W., Carton, A.S. (Eds.), *The Collected Works of L.S. Vygotsky, vol. 1: Problems of General Psychology Including the Volume Thinking and Speech*. Plenum Press, New York.
- Waage, J.K., 1979. Foraging for patchily-distributed hosts by the parasitoid, *Nemeritis canescens*. *Journal of Animal Ecology* 48, 353–371.
- Wilke, A., Hutchinson, J.M.C., Todd, P.M., Czienskowski, U., 2009. Fishing for the right words: decision rules for human foraging behaviour in internal search tasks. *Cognitive Science* 33, 497–529.
- Winter, Y., Stich, K.P., 2005. Foraging in a complex naturalistic environment: capacity of spatial working memory in flower bats. *The Journal of Experimental Biology* 208, 539–548.
- Yokoi, T., Fujisaki, K., 2007. Repellent scent-marking behaviour of the sweat bee *Halictus (Seladonia) aerarius* during flower foraging. *Apidologie* 38, 482–504.