

Decision ecology: Foraging and the ecology of animal decision making

DAVID W. STEPHENS

University of Minnesota, Saint Paul, Minnesota

In this article, I review the approach taken by behavioral ecologists to the study of animal foraging behavior and explore connections with general analyses of decision making. I use the example of patch exploitation decisions in this article in order to develop several key points about the properties of naturally occurring foraging decisions. First, I argue that experimental preparations based on binary, mutually exclusive choice are not good models of foraging decisions. Instead, foraging choices have a sequential foreground-background structure, in which one option is in the background of all other options. Second, behavioral ecologists view foraging as a hierarchy of decisions that range from habitat selection to food choice. Finally, data suggest that foraging animals are sensitive to several important trade-offs. These trade-offs include the effects of competitors and group mates, as well as the problem of predator avoidance.

Foraging in Patches

In an alpine meadow, a worker bumblebee plows through the air. Its large body and implausibly small wings remind one of an enormous but absurdly miniaturized transport plane. Like their better known relatives the honeybees, bumblebees live in colonies and collect pollen and nectar to feed their developing brood. Foraging bumblebees visit flowers, of course, flying to them and crawling into them to find pools of nectar. A natural meadow is not a uniform floral carpet. Instead, the bumblebee finds flowers in clumps. Most plants, for example, present a small cluster of flowers together on a single stalk called an *inflorescence*. What sorts of decisions does the bumblebee need to make as it moves through this world of flowers and flower clumps? Of course, it must somehow choose which inflorescence to visit and which flower to visit on the inflorescence. A somewhat less obvious problem is the necessity of deciding whether to stay and exploit another flower on this inflorescence or to leave and find a new inflorescence to exploit.

Many—probably most—animals face leave-versus-stay decisions like this. To understand the advantages or disadvantages of leaving or staying, we need to know how the bee accrues food as it spends time exploiting the inflorescence. Typically, the bee will acquire food fairly quickly at first, so that we would see a roughly linear relationship between nectar gained and the time spent exploiting the inflorescence. This cannot last, because the bee's exploitation eventually depletes the inflorescence, so our plot of gains versus time will start to bend downward (Figure 1) and must ultimately asymptote to some maximum that is set by the limited resources of the inflorescence. So patches exhibit diminishing returns; fresh, unexploited

patches yield food quickly, but food gains per unit time inevitably decline. Consider two possible strategies: A patch-exploiting bumblebee could make short visits that "skim the cream," moving quickly on to fresher patches elsewhere. At the other extreme, our bumblebee could exploit each inflorescence thoroughly, working the inflorescence to extract every last drop of nectar. Which should the bumblebee do?

To answer this question economically, we want to compare the value of staying to the value of leaving. We can use the relationship between exploitation time and food gains discussed earlier to find the value of staying. Obviously enough, if this relationship shows diminishing returns, the value of staying for one more time interval must also decline steadily. To find the value of leaving, however, we need to consider things beyond the current patch. We need to know the value of the bee's options elsewhere: What can it gain if it leaves the patch? This depends on the overall richness of the bee's habitat. In a rich habitat with lots of inflorescences dripping with nectar, the value of leaving is obviously higher than in a poor habitat, where low-quality inflorescences are few and far between. With these elements in mind, the reader can probably piece the whole story together. One would expect a forager's patch-leaving behavior to reflect a balance between the diminishing returns of patch gain and the value of the forager's options elsewhere. This covers a lot of territory, of course, because many things can influence the forager's patch-gain function, as well as the properties of the options available elsewhere. We can reduce this complexity somewhat by considering a simplified habitat in which all the patches are the same—that is, all patches exhibit the same relationship between exploitation time and amount

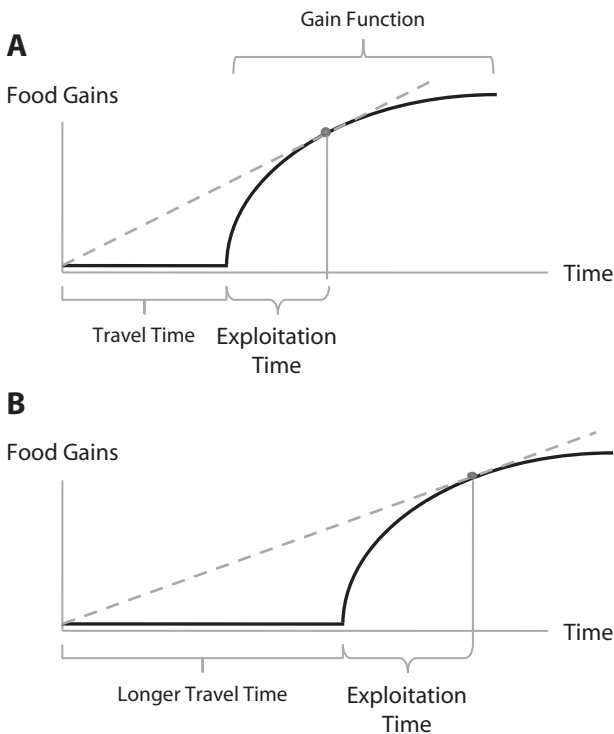


Figure 1. Graphical model of patch exploitation. Each plot shows the relationship between time and the amount of energy extracted from a patch. The forager gains no energy while it is traveling to the patch, so we show a flat line of height 0 during the travel phase of a patch exploitation cycle. The forager begins to obtain resources when it enters a patch, and we characterize this gain process by a decelerating but increasing curve. This gain function bends down because the resources in the patch are finite, so the rate of extraction must ultimately slow down. To find the patch exploitation time that gives the highest intake rate, we construct a straight line from the origin to the gain curve, so that this line is just tangential to the gain curve. The slope of this line is the intake rate (gains/total time), and the rate-maximizing patch exploitation time corresponds to the point of tangency (where the dashed line touches the curve). Panel A shows a short travel time (rich habitat), and panel B shows a long travel time (poor habitat). One can see that this tangent construction model predicts longer exploitation times in poorer habitats.

gained, which we call the *gain function*. This simplified situation lets us separate the values of staying and leaving quite cleanly, because if all patches are the same, richer habitats simply have more of them; to represent this, we focus on the time required to travel between patches. This travel time will be short in rich habitats and long in poor habitats. Therefore, long travel times should promote staying, and short travel times should promote leaving. Figure 1, which will be familiar to some readers, shows this logic graphically. Of course, this leads to a simple behavioral prediction: Animals should spend more time in and extract more resources from patches when travel times are long. The driving force behind this prediction is opportunity cost. In a rich environment, a forager loses future opportunities by overstaying, whereas in a poor habitat, an understaying forager loses opportunities associated with exploiting the current patch. Readers may recognize the ideas developed here as the so-called *marginal value theo-*

rem originally developed by Charnov (1976) and Parker (1978; see Stephens & Krebs, 1986, for a comprehensive review of this classic idea).

Do foraging animals follow this prediction? The answer is a resounding and quite impressive *yes!* We have evidence from over 60 different species ranging from invertebrates, such as spiders and wasps, to vertebrates, such as birds and mammals (Nonacs, 2001; Stephens & Krebs, 1986). In a well-known study, Richard Cowie (1977) created artificial patches for an insectivorous bird (*Parus major*) by hiding mealworms in sawdust-filled yogurt cups and found a very orderly relationship between travel time and exploitation time.

An especially elegant setup uses animals that forage from a central place. Chipmunks studied by Giraldeau and Kramer (1982) made forays from their burrows collecting seeds and returned, adding those seeds to larders in their burrows. In situations such as this one, one can manipulate travel time experimentally by simply placing piles of seeds different distances from the home burrow. As with Cowie's (1977) birds, one finds an orderly relationship between travel time and the time chipmunks spend collecting seeds from a patch. It would be misleading to claim that the simplest patch models (as outlined here) predict behavior perfectly—far from it. Yet, they clearly predict this basic and widely observed qualitative relationship between travel time and patch exploitation time. Finally, I comment that food patches are quite literally everywhere; even things that do not seem like patches have patch-like properties. Think again about our foraging bumblebee. We developed the idea of patches by focusing on clumps of flowers, but we could have focused on a single flower: As the bee laps the nectar from a flower, it experiences diminishing returns, just as it does at the level of the whole inflorescence.

Self-control versus patch exploitation. In this article, I review the behavioral ecology of animal foraging, highlighting connections with psychology and neurobiology. The introduction described the simple natural problem of patch exploitation. This section begins to make connections with psychology and neuroscience by exploring the connection between patch exploitation and impulsive choice. As many readers will know, animals show strong preferences for immediate food delivery (see, e.g., Ainslie, 1974; Bateson & Kacelnik, 1996; Killeen, 1968; Mazur, 1987; Rachlin & Green, 1972). This behavior travels under many names: impulsivity, failure to delay gratification, delay discounting, and a few others.

Consider a simple impulsivity study. The investigator arranges a situation in which the animal waits for a fixed time, τ . When this intertrial interval expires, the apparatus offers the animal a choice between a smaller-sooner option and a larger-later option (Figure 2A). Obviously, the animal must do this repeatedly, and it will learn through experience that, for example, choosing a red stimulus leads to the larger-later option, whereas choosing a green stimulus leads to a smaller-sooner consequence. We have a large body of results from this simple experimental paradigm, which is usually called the *self-control preparation* (because an animal choosing the larger-later option is said to demonstrate self-control). The basic observa-




Figure 2. A comparison of the self-control and patch choice preparations. Panel A shows a single trial in a conventional choice study. The animal waits for time τ ; the apparatus then offers a choice between a smaller-sooner and a larger-later option. The cross shows the point at which the animal makes a choice. Panel B shows an economically analogous patch situation. In the patch preparation situation, the animal waits (travels) for time θ , then encounters a single stimulus (a patch); after entering the patch, the animal waits d_1 sec for a delivery of amount B_1 . At this point, the animal can choose to stay in the patch and obtain additional amount B_2 after delay d_2 , or it can leave and start a new trial. Subjects must choose between a smaller amount in a smaller total trial time and a larger amount in a larger trial time in both situations.

tion, however, is that animals have strong preferences for smaller-sooner options, even though simple economic considerations suggest that the larger-later option is better. The economic considerations in question derive from comparisons of long food-intake rates. So when we say that animals prefer smaller-sooner options, even though they should not, we mean that they continue to prefer smaller-sooner options, even when they could obtain more food in the long run by waiting for the larger-later option. Note that this smaller-sooner versus larger-later problem has much in common with patch exploitation: Should I spend less time extracting less or spend more time to extract more? We have clear-cut evidence showing that animals in poor habitats spend more time in patches extracting more food, effectively choosing larger-later consequences (Nonacs, 2001). Similarly, animals in rich habitats spend less time in patches and extract less food, effectively preferring smaller-sooner consequences. To be sure, exploiting a natural food patch differs in many ways from choosing between red and green stimuli, but the problems are quite similar economically. Given this economic similarity, we would expect that creating a poor habitat by experimentally creating a long intertrial interval (the economic analogue of travel time) should shift preference toward the larger-later option. On the other hand,

creating a rich habitat by creating short intertrial intervals should shift preference toward the smaller-sooner option. Unfortunately, the data clearly show that changes in intertrial interval do not influence choice in this preparation (see, e.g., Bateson & Kacelnik, 1996; Mazur, 2001).

Why is the self-control preparation so different from patch exploitation, even though the same lost-opportunity effects occur in both? Over the past 10 years, my laboratory has developed a choice preparation that captures some basic properties of naturally occurring patch exploitation, and we can directly compare data gathered from this preparation with data from self-control studies (Stephens & Anderson, 2001; Stephens, Kerr, & Fernández-Juricic, 2004; Stephens & McLinn, 2003). Figure 2B shows the basic idea. The subject (in our experiments, captive blue jays) waits for a fixed time between presentations. This waiting time is economically analogous to the travel time in patch exploitation and intertrial interval in the self-control preparation. When this “travel time” expires, the apparatus presents a single stimulus that indicates a patch encounter. The animal makes some response that serves as the experimental analogue of entering the patch (in our preparation, it hops from the “travel perch” at the rear to the “exploitation perch” at the front of the apparatus). After an initial delay, the apparatus delivers a small amount of food. At this point, the animal faces the patch exploitation decision that interests us. The animal can stay where it is and wait for an additional delivery of food from the apparatus. This mimics natural staying decisions that generate more food but take more time. Alternatively, the subject can leave the front perch and return to the travel perch at the rear. If the subject chooses to leave at any time, this initiates a new travel-encounter-enter cycle.

The downside of the patch choice preparation is, of course, that it is quite simple in comparison with natural patches. It has two key virtues, however. Its simplicity makes it tractable and easy to implement, and we can easily engineer situations where our patch preparation creates exactly the same economic situation as a self-control situation. One can do this fairly simply (Figure 2) by arranging things so that (1) a leave decision yields the same amount of food and takes the same amount of time as choosing smaller-sooner in self-control, and (2) a stay decision yields the same total amount of food and takes the same amount of time as choosing larger-later in self-control. I hasten to add that this arrangement only creates situations that are economically equivalent. The two situations still differ in many ways, as will be discussed below. For example, in the patch situation, staying produces two food deliveries (within a choice-to-choice trial), whereas the economically equivalent larger-later option in self-control produces only one.

In our first experimental study comparing the patch and self-control preparations (Stephens & Anderson, 2001), we created 12 pairs of equivalent patch and self-control situations by systematically varying the temporal parameters of the situation—that is, the intertrial interval, the delay to small, and the delay to large. In comparing these 12 economically equivalent situations, we found that our subjects made many impulsive errors (i.e., they chose

smaller—sooner in situations where larger—later would have produced a larger intake rate) in the self-control situation, but they achieved a much higher level of performance in the patch situation. Specifically, in cases where birds in the self-control preparation “mistakenly” chose smaller—sooner, the same birds chose to stay, achieving relatively high long-term gains in patch. In another study (Stephens & McLinn, 2003), we showed that changes in intertrial interval (i.e., travel time) influenced choice behavior in the patch situation, as it does in naturally occurring patch exploitation. Indeed, this study showed that animals in the patch situation treated the intertrial interval (θ , in Figure 2) and initial delay (d_1) as a single combined delay, a result that contrasts sharply with the self-control preparation, where subjects show exquisite sensitivity to delay, but ignore the intertrial interval.

In addition to these empirical results, our group has observed that impulsive decision rules—that is, rules that evaluate options shortsightedly—could, in theory, explain these results (Stephens & Anderson, 2001; Stephens et al., 2004). Specifically, an animal using a shortsighted rule could achieve a high long-term intake rate in the patch preparation, even though the same rule leads to impulsive mistakes in the self-control situation. This observation led us to offer a novel explanation for the evolution of animal impulsivity; we call this the *ecological rationality hypothesis*. According to our hypothesis, selection has favored impulsive choice mechanisms, because natural choice situations are frequently like our experimental patch situation—which, it turns out, favors shortsightedness. In other words, animal impulsivity is rational when viewed in the correct ecological context. Although this article is not, primarily, a review of our work on ecological rationality, this approach does illustrate why it is important to consider the natural context of the choice mechanisms we all study. And, of course, this article argues that foraging is an important and tractable “natural context” in which to study decision making.

What is the difference between patch exploitation decisions and conventional approaches to choice? There are, in fact, several critical differences:

1. Foraging choices are repeated, not single-shot, choices. The single-shot choice (a situation in which a decision maker must choose between alternatives once and only once and be committed in perpetuity thereafter to whatever the decision maker selects) is a useful theoretical construct because it isolates *this* decision from the remaining complications of the world. Of course, the real world does not work this way—certainly not the world in which animals acquire food. Acquiring food is a lifetime endeavor for most animals, and decisions made today have implications for foraging tomorrow.

2. Natural foraging decisions are seldom choices between mutually exclusive options. When an animal chooses to stay in a patch, it is, in one sense, choosing not to leave; but choosing to stay does not prevent the animal from eventually leaving. Yet, most choice experiments enforce mutual exclusivity; choosing option A means that option B will not be available until the next trial, when another A versus B choice is offered.

Combining these two points, we can see that natural foraging (and patch exploitation decision, specifically) has a special structure that we call *foreground–background choice*. In foreground–background choice, the forager faces a series of “act or continue”-style choices (stay vs. leave, attack this prey item vs. continue searching for others, etc.). Within this framework, we imagine that a decision maker follows a default series of actions (e.g., search for more prey or patches) that represents a background exploitation strategy (see Stephens et al., 2004, for a complete development). This background strategy represents the animal’s default actions, and it generates an ongoing stream of food gains.

Now assume that a forager encounters a potential food item that it can attack or ignore. If it attacks, this means that it puts its background strategy on hold for the time required to exploit this foreground resource, but that it will return to its background tactic when it is finished. If it ignores the foreground item, it simply continues using its background tactic. In patch exploitation, therefore, choosing to stay for an additional time unit is a foreground option, but choosing to leave and search for further patches is always in the background. Foraging decisions do not, then, necessarily fit the decision-making model that experiments using binary, mutually exclusive options seem to imply. Using this new terminology, then, the ecological rationality hypothesis holds that impulsive choice mechanisms evolved because they perform well in foreground–background choice, even though they produce relatively poor performance in the laboratory-enforced regime of binary, mutually exclusive choices.

Of course, this critique of binary choice is not unique. In Hastie’s (2001) thoughtful review of problems in the study of decision making, he ridicules “the image of a decision-maker standing at a choice point like a fork in a road” (p. 665). Yet, recognizing this problem within natural foraging gives us an important natural context in which to investigate nonbinary choice; it also suggests alternative formulations, such as our foreground–background distinction.

I have taken some time to develop the idea of patch exploitation, in part, to illustrate the important properties of foraging decisions outlined above. However, I also seek to remind readers of this important and basic decision problem: Typical animals will make thousands of patch exploitation decisions during their lives, and my group’s research suggests that they may be especially good at making these sorts of decisions. In contrast to the natural simplicity of patch exploitation, the contrived situations that some workers describe as foraging are, to me, simply surprising. For example, I find very little similarity between natural foraging choice and the concurrent variable interval schedules used in the study of matching. Yet, it is not uncommon to read psychological articles that offer matching as an analogue to foraging (e.g., Gallistel, 1994).

Foraging: Models and General Principles

The patch exploitation situation offers us a glimpse into the problems that foragers face, but there is, of course, a great deal more to natural foraging. The remainder of this

article offers an overview of foraging problems that is intended for students of decision making with backgrounds in psychology and neurobiology. Before I launch into this review, however, I consider a more basic question: Why should students of decision making care about natural foraging? I offer three points in answer to this question: (1) As scientists, we all recognize that simplified and even contrived situations (such as the self-control preparation discussed above) can offer important insights in underlying decision processes. Yet, our interpretation of results can depend on an understanding of how the laboratory situation relates to situations that our subjects normally face. In my group's studies of animal impulsivity, for example, interpreting self-control results as patch exploitation rules used out of context gives us a dramatically different interpretation of animal impulsivity (see Stephens et al., 2004). (2) An ecological approach to decision making can reveal new patterns and generate new hypotheses. For example, patch exploitation models led investigators to document the nearly universal relationship between travel time and patch exploitation time. (3) Foraging is special. The animal way of life is, quite literally, defined by an imperative to acquire food, and this imperative has surely played a central role in shaping the neural mechanisms that control movement and, ultimately, decision making.

Behavioral ecology is a subdiscipline of biology that seeks to understand behavior in an ecological context—especially, how ecological forces, via natural selection, have shaped behavior. Foraging behavior has been a major topic within behavioral ecology since the 1970s. In this section, I review the behavioral ecologist's approach to foraging; but before doing so, let us pause to discuss the nature of models in behavioral ecology.

The foraging models outlined here follow a long tradition of evolutionary economics within the field of behavioral ecology. Natural selection, the argument goes, acts like an optimizing agent, because it selects alternatives on the basis of fitness (e.g., reproductive success). Although the mathematics of these models are nearly identical to those used in economic models, the two types of models are not the same. Economic studies of choice assume that some system of value (e.g., utility) exists that guides the decision maker's behavior, and the economic models are not especially fussy about where this system of value comes from; it is often fitted from observed behavior via the so-called *revealed preference approach*, or it could be derived from some a priori model. Evolutionary models, in contrast, have one and only one system of value: evolutionary fitness. In practice, however, this crisp logical distinction can become fuzzy, because the connection between fitness and behavior is often purely hypothetical.

Reconsider the patch models outlined earlier. These models predict the behavior of an animal exploiting patchy resources, following the premise that natural selection has favored increases in the long-term rate of food intake. This model, correct or not, represents an evolutionary hypothesis. If selection has acted according to assumptions X, Y, and Z, then behavior should have properties A and B. In our example, the patch model successfully predicts a relationship between travel time and exploitation time.

Now consider an alternative model that considers the underlying neural mechanisms of patch departure. Although I have no idea what form this model would actually take, such a model might identify neural activity associated with patch departure and somehow specify a connection between the properties of this activity and the locomotion that precedes patch encounter.

A mechanistic model like this could surely explain the travel time–exploitation time relationship. Are the two models in conflict? Are they alternatives? Is one about true causation and the other something less? Surely not. Animals can exploit patches more thoroughly in poor habitats because selection has favored actions that maximize intake rate *and* because they have neural mechanisms that control this behavior. Behavioral ecologists call this the *proximate–ultimate distinction* (Alcock, 2005; Alcock & Sherman, 1994). In this terminology, ultimate models deal with the evolutionary causes of behavior, and proximate models deal with the underlying mechanisms. Some readers will recognize the similar distinction made by Tinbergen (1963) in his famous “four questions.”

In practice, many investigators are perfectly happy to ask questions about only one of these two causes (mechanistic or evolutionary) of behavior. Most students of behavior would agree, however, that these two accounts of behavior should ultimately be drawn together into an integrated whole. Unfortunately, the current state of the art falls short of this integration. Instead, we simply hope that each approach informs the other. For example, studies of mechanism reveal constraints that evolutionary modelers can incorporate into their models, whereas evolutionary economics often reveal new behavioral phenomena (e.g., the travel time–exploitation relationship, or the input-matching behavior of group-feeding animals) for students of mechanism to explore physiologically.

A decision hierarchy. The process of foraging involves many decisions, and behavioral ecologists find it useful to think of an interrelated hierarchy of decisions (Stephens & Krebs, 1986). At the highest level, foraging animals must choose where to forage: Should our bumblebee forage in a clump of *Delphinium* on a slope, or from the spikes of *Aconitum* along a neighboring stream? We call these *habitat choice* problems, and the defining feature of these choices is that foragers' behavior has little influence on the benefits obtained from a habitat. In the simplest habitat choice model, therefore, we would characterize each option by rate of gain (e.g., 8 J/min in the *Delphinium* habitat and 11 J/min in the *Aconitum* habitat) and predict that the animal should prefer the habitat with the higher food gain. Habitat choice is probably the closest that natural foraging comes to binary mutually exclusive choice, although this binary choice happens on fairly large spatial and temporal scales, and, of course, habitat-exploiting animals are free to switch to another habitat at any time. Below, I will describe two important complications—the effects of predators and those of competitors—in habitat choice.

Within a habitat, we imagine that the animal experiences a temporally structured encounter process (encountering prey items or patches) as it moves through the

habitat. At its simplest, this encounter process could be something like *encounter a flower every 25 sec*. In a more complex encounter process, we would specify the average rates at which the forager encounters different types of potential resources. (Although investigators seldom discuss it, a forager can make several decisions that influence the encounter process, such as movement speed.) Given a well-defined encounter process, the forager must make a sequence of accept- or reject-style decisions. We often call these *prey choice* decisions, even though the selected entities sometimes do not seem much like prey. A cow, for example, may accept or reject a clump of grass. At the lowest level, we are back to the patch exploitation problem discussed above. Given that the cow accepts the clump of grass, how extensively should it exploit it? Clearly, prey choice and patch exploitation decisions have a foreground–background structure that differs strikingly from the binary, mutually exclusive choice paradigm favored by experimentalists and decision modelers. Finally, I comment that, whereas behavioral ecologists find this three-level hierarchy useful, one can certainly expand and embellish this simple hierarchy to address special situations. We could, for example, push things down one more level by considering digestive and processing decisions (e.g., how long to chew).

Rate-maximizing models. Behavioral ecologists have expended considerable effort in exploring this hierarchy of decisions, both theoretically and empirically. The classic models of these problems adopt a rate-maximizing premise. Our basic model of patch exploitation discussed above, for example, solves for the patch exploitation tactic that maximizes long-term food intake (usually, but not necessarily, measured in joules). These models provide a rich set of predictions too extensive to review here (Stephens & Krebs, 1986, provide the classic discussion of these models). It is useful, however, to recognize the principle of lost opportunity that is the common thread connecting this family of basic models.

For example, a forager that stays in a patch too long pays an opportunity cost because it wastes time exploiting a depleted patch when fresher patches remain unexploited. Similarly, in rate-maximizing models of prey choice, choosing a mediocre prey item with a long handling time carries an opportunity cost, because the forager forgoes the opportunity to search for a superior prey item while processing the mediocre item. One can certainly quibble over the details of this simple family of models, but the principle of lost opportunity does seem to offer important insight into the economic determinants of forager behavior.

Predators and competitors in habitat choice. Economically, the problem of habitat choice seems to be the least interesting problem in our hierarchy. If one habitat offers food at rate r_1 , and the other offers food at rate r_2 , then we should surely expect a forager to exploit the habitat that provides the higher intake rate. Of course, it may have to learn this, but this does not seem to be a terribly challenging problem. Perhaps due to the logical simplicity of habitat choice problems, studies of habitat choice have given us a foothold on two more complex problems: the

effects of competitors on foraging behavior and those of predators on foraging behavior.

To see the effects of competitors on habitat choice, consider two habitats that deliver food at rates r_1 and r_2 , respectively. As explained above, we expect a solitary forager to occupy the habitat with the highest intake rate. To include the effects of competitors, assume that there are N total foragers who must somehow arrange themselves so that n_1 individuals occupy Habitat 1 and n_2 individuals occupy Habitat 2 (where we obviously require that $n_1 + n_2 = N$). Suppose that these habitats are ideal in the sense that the n_1 individuals in Habitat 1 share the resources equally, so that each obtains an intake rate of r_1/n_1 , and it follows that the individuals in Habitat 2 obtain an intake rate of r_2/n_2 . Now suppose that each of N individuals is free to move from one habitat to the next. If individuals can move and evaluate their gains in the different habitats, we should expect a net movement to Habitat 1 any time r_1/n_1 exceeds r_2/n_2 . Indeed, the only situation in which there would be no net movement—anthropomorphically, we might say that all the foragers are happy where they are—is when $r_1/n_1 = r_2/n_2$, or, equivalently, when $r_1/r_2 = n_1/n_2$. This is called the *input matching rule*, because it predicts that relative numbers of competitors in each habitat should match the relative food inputs. As readers will observe, this is an incredibly simple model, so it is quite surprising to find that it is very successful experimentally. Abrahams (1989) showed, for example, that guppies foraging in an aquarium followed the predictions of the input matching rule. Simple studies like this confirm that groups tend to distribute themselves in a way that matches inputs; if we deliver food at twice the rate to the left, we will typically find twice as many fish there. Behavioral ecologists call this the *ideal-free distribution* because it assumes that individuals choose freely and share food according to an ideal assumption of equal sharing (Fretwell & Lucas, 1970). Of course, real animals are seldom ideal or free. In particular, some individuals compete more effectively than others, getting a larger share of the food, or even defending a site against intruders. Yet, even in the face of these obvious failures of our assumptions, we still frequently observe input matching among groups of foragers (see, e.g., Tregenza, 1995). Due to its robustness and its logical simplicity, the ideal-free distribution provides the foundation of modern studies of social foraging. These models use a leaving–joining framework that clearly derives from the ideal-free model (Waite & Field, 2007). This input match framework resembles Herrnstein’s well-known matching law (Herrnstein, Rachlin, & Laibson, 1997; see also Baum, 1974), which predicts the allocation of behavior by individual animals in operant choice situations. Interestingly, the only study I know of having individual animals in a habitat choice situation is from my own laboratory (Stephens & Stevens, 2001). This study showed that single individuals do not match behaviorally in habitat choice situations. Instead, they use the best habitat exclusively, or—if the two food sources are close together—they position themselves to use both, rather than allocating their time to each op-

tion in portion to its richness, as the input matching law would predict. This is hardly surprising, notwithstanding the algebraic similarity of the two “matching” formalisms, since the two approaches actually deal with strikingly different scenarios and dependent measures (e.g., allocation of choice behavior vs. allocation of individuals between habitats).

As these ideas suggest, different habitats offer different levels of food reward. Of course, they also vary in other ways. One key dimension is in the degree of risk from predators; indeed, many of the best habitats are also the most dangerous. Foragers clearly face a critical trade-off between acquiring food efficiently and avoiding death. Here again, habitat choice provides a useful framework. For example, Abrahams and Dill (1989) engineered a simple extension of conventional “fish in a tank” studies by cleverly positioning a predator at one habitat, but not at the other (using a screen that guppies, but not predators, could pass through). They found that the presence of a predator changed the distribution of fish quite dramatically; but by systematically increasing the food input rate in the predator habitat, they were able to create a 50:50 distribution of guppies in the risky and safe habitats. This suggests that foraging animals somehow evaluate combinations of danger and food reward. Behavioral ecologists now have many results like this (probably hundreds of examples) that suggest that animals make sophisticated trade-offs between food and safety. For example, grasshoppers adjust their position on grass stems in response to predators. When they primarily suffer bird predators, they forage near the ground; but when predatory lizards are present, they shift upward (Pitt, 1999). In another instance, small birds have two key defenses against predators: vigilance (looking up) and foraging in groups. They look up more frequently in small groups, suggesting that they are sensitive to a trade-off between these two predator avoidance techniques (see, e.g., Fernández-Juricic, Kerr, Bednekoff, & Stephens, 2004). In addition, the relationship between group size and vigilance depends on factors that influence the predator risk (e.g., a recently sighted predator, tall grass that reduces the value of vigilance, etc.).

These observations suggest some interesting problems for psychologists and neurobiologists interested in decision making. The ideal-free distribution, for example, represents a relatively simple (perhaps fundamental) problem in social behavior that psychologists and neurobiologists could profitably explore. Several recent articles suggest that the neural underpinnings of social interaction could be experimentally accessible (Flombaum & Santos, 2005; Sanfey, 2007), but these studies tend to focus on complex animals in complex situations (e.g., primate strategies in social games). It could be quite revealing to apply similar techniques to these basic problems of competition and resource use. To an outside observer, at least, it seems that the neural bases of reward and fear are studied as two distinct phenomena (see Davis, 2006), but this mass of evidence from behavioral ecology suggests not only that animals can integrate decisions based on reward and fear, but also that this represents an important natural problem that animals have experienced through many generations

of selection (see Mobbs et al., 2007, for a study that shows how these analyses might proceed).

Why Foraging Matters

In this final section, I discuss connections between foraging and other approaches to decision making. In doing so, I hope to make the point that an understanding of animal foraging behavior can enrich our analyses of decision making. Of course, the basic point here is that understanding the problems that the decision-making apparatus has been designed to solve offers a template we can use to make stronger inferences about underlying mechanisms. I offer some specific examples below to explain how thinking about foraging can inform psychological and neurobiological approaches to decision making.

Psychology and foraging. The interplay between foraging and animal psychology is rich and increasingly well developed (three recent articles review this area: Adams-Hunt & Jacobs, 2007; Sherry & Mitchell, 2007; Stephens, 2007). Studies of foraging may, for example, rely on psychological results about discrimination in order to discern how animals discriminate between palatable and unpalatable prey items. One of the most basic connections focuses on foraging and learning. Many learning studies—starting with Pavlov—focus on learning about food and feeding resources. Foraging animals clearly learn many features of their foraging environments (e.g., properties of prey, locations of resources, etc.), yet they respond with simple fixed preferences in other feeding situations. In order to understand this, behavioral ecologists have studied the benefits associated with learning (using experience to track the properties of the environment), both theoretically and empirically (Stephens, 2007). The basic premise of this approach is that statistical properties of the environment, such as predictability and change, determine the value of attending to experience.

Consider, for example, another choice problem that our foraging bumblebee might face. Assume that our bumblebee finds itself foraging in a clump of flowers. A given percentage of all flowers provide rich nectar rewards (proportion $p = \text{good}$), whereas the remaining flowers (proportion $1 - p$) provide almost nothing. The variable p tells us something about the underlying uncertainty about flower quality. For example, if p is near 1, all the flowers are good; our bumblebee can simply visit every flower it encounters. Now imagine that some of the flowers are blue and some are yellow, and there is a correlation between flower color and nectar rewards, so that blue flowers tend to be richer. Although this color-quality correlation could take many forms, let us use a simple reliability parameter to represent the strength of the correlation; we will assume that blue tends to signal a good flower, and yellow tends to signal a bad flower. We can use a conditional probability q to represent this reliability. That is, we let q be the probability that a blue flower is good and the probability that a yellow flower is bad. A q value near 1 means that color reliably indicates quality; conversely, a q value of .5 means that color provides no new information. When should the bee use its experience of flower color to change its behavior, as opposed to simply ignoring its experience

and attacking or ignoring everything it encounters? A moment's reflection will show that reliability and uncertainty are both important here. If uncertainty is low (e.g., most flowers are good, or most flowers are bad), the bee already knows what to do, and flower color has little importance, regardless of its reliability. On the other hand, if flower quality is uncertain, we still need a reliable correlation between experience (color, in our example) and consequence to make attending to experience worthwhile. Even though this example is oversimplified (and even contrived), we clearly expect that an interaction between uncertainty and reliability should determine the value of attending to experience. Do real animals respond to these variables as we predict? Using our colony of captive blue jays, my research group created an experimental situation that parallels our flower color example (McLinn & Stephens, 2006; see Stephens, 2007, for a recent review). We factorially manipulated uncertainty and reliability, and we found that these variables predicted whether our subjects responded to a color stimulus roughly as hypothesized. In particular, we found that our jays came to ignore colored stimuli in conditions of low uncertainty and that the same animals changed their behavior in response to colored stimuli in conditions of high uncertainty and high reliability. This is, of course, only one of many approaches within foraging and behavioral ecology to learning and learning-related phenomena (see Stephens, 2007, for a review). Focusing on whether experience should shape behavior, I remark that the phenomenon discussed here actually cuts a somewhat broader swath than learning. This is a defining property of learning, and it also admits phenomena such as stimulus–response relationships, which we would not call learning.

Although there is an increasingly well-developed interaction between behavioral ecologists interested in foraging and animal psychologists, these two fields do, of course, offer quite different perspectives. As the example above suggests, behavioral ecologists emphasize adaptive differences in cognitive abilities. These adaptive differences may be short-term differences that reflect recent changes in an animal's habitat or longer term differences (between species and populations) that reflect differences in evolutionary history (which we might call *cognitive specialization*, following Shettleworth, 1998). In contrast, the psychological tradition emphasizes general processes of behavioral control. This, for example, leads psychologists to assume that principles of behavioral control revealed by studies with rats will easily generalize to small birds or slugs. Behavioral ecologists often find this a difficult argument to accept. Overall, these differing perspectives provide a healthy dose of tension, which forces workers in both fields to enlarge and defend their paradigms. This difference can get in the way of productive communication, however, because the researchers sometimes fail to recognize their different but complementary goals.

Foraging and neuroeconomics. Neural studies of decision making, and the field of neuroeconomics, specifically, are relatively young, so rather than outline how ideas about foraging have contributed to these fields, I will offer

a prospectus about how this interaction might develop. Although students of neuroeconomics show a growing sensitivity to the natural context of decision making (see, e.g., Platt, 2002), their studies still largely use traditional decision-making paradigms, such as binary, mutually exclusive choice. One reason for this, of course, is that one can implement these conventional tests easily. Of course, I advocate that neuroeconomists consider choice paradigms (such as patch exploitation) that have a more realistic structure. I do not necessarily argue that binary choice situations are irrelevant, but that understanding their relationship with the types of problems animals face in nature will provide important comparative data and possibly reveal that different neural mechanisms are in play.

I also think that the potential to incorporate naturally important trade-offs in neural analyses of decision making is quite exciting. As explained above, evidence from behavioral ecology suggests that animals can make rather sophisticated decisions that balance predation risk and food intake. We have, of course, a growing base of knowledge about the neural basis of fear (e.g., Davis, 2006); this has not yet been integrated into the neuroeconomic approach. Evidence from behavioral ecology suggests that this is a compelling and naturally significant problem: How is information about danger and food gain combined, and how is the balance struck?

Caching and memory. The phenomenon of food caching represents one area where a reasonably well-developed interaction between foraging and neurobiology already exists. Clark's nutcrackers (*Nucifraga columbiana*) live at high elevations in western North America, where ripening pine seeds are superabundant in the fall. Nutcrackers can store up to 30,000 seeds for later use, using them to rear their offspring the following spring. Many foraging animals collect food and store it for later use. This simple aspect of natural history illustrates the power and importance of understanding the problems animals face in nature, because caching has proven to be an extremely rich area for the study of memory and spatial abilities. For example, many studies have demonstrated that caching animals achieve higher levels of performance in spatial memory tasks than their noncaching counterparts, and this enhanced ability seems to be supported by an enlarged hippocampus. Sherry (2006) reviews this important interaction between foraging and neural structures, and outlines a review program he calls *neuroecology*.

The status of foraging models, as well as results. As I have outlined results from behavioral ecology, I may have inadvertently painted a picture of foraging as a solved and completely understood problem. Of course, this is emphatically not the case; the models and data outlined above draw a fairly crude sketch of real animal foraging, and they are very much works-in-progress (see Stephens, Brown, & Ydenberg, 2007, for a review of current problems).

Several investigators focus on violations of these basic claims about foraging choice. Collectively, these studies fall under the heading of *irrationality* or *bounded rationality*. We have several examples of apparent irrationality in feeding animals. These include examples of intransitive

preferences (Shafir, 1994; Shafir, Waite, & Smith, 2002); examples of decoy effects, in which the presence of a mediocre third option changes preference for two better alternatives (see, e.g., Schuck-Paim, Pompilio, & Kacelnik, 2004); and so-called *work effects*, where subjects prefer options previously associated with hard work (Pompilio, Kacelnik, & Behmer, 2006). These are all intriguing and active topics of study. For behavioral ecologists, the working hypothesis for these observations is the ecological rationality hypothesis. As explained above, this view holds that observed irrational choice is an emergent property of mechanisms that work well (i.e., typically make rational choices) in natural choice situations (see Todd & Gigerenzer, 2000). This basic idea stretches back to the early days of behavioral ecology, when it was known as the *rules-of-thumb problem* (see, e.g., Cheverton, Kacelnik, & Krebs, 1985). After a brief flurry of interest, behavioral ecologists seemed to abandon the rules-of-thumb problem until its reemergence under the heading of *ecology rationality*. Like its predecessor, the ecology rationality approach faces some difficulties. For example, it sometimes seems too easy. Virtually any deviation from expectation could easily be dismissed as being ecologically rational. It is clear that we still need to develop some standards for how this approach can be rigorously applied.

Summary

Among living systems, animals are the ultimate decision-making machines. The animal way of life has the business of acquiring resources at its core. It follows, I argue, that those who wish to understand decision making need to appreciate the decision-making problems that foraging animals have faced throughout their evolutionary history. Foraging animals face a hierarchy of decisions in which they must choose among feeding habitats, select prey items and patches within these habitats, and allocate their foraging efforts in the face of conflicting demands, such as that between efficient feeding and avoiding predators. I have argued that binary, mutually exclusive choice offers a very poor model of naturally occurring foraging decisions. In contrast, foraging choices are fundamentally sequential and commonly have a foreground-background structure, in which the background alternative is a component of all possible alternatives. Behavioral ecologists have many lines of evidence suggesting that animal decisions strike a sophisticated balance between food gains and predation risk. Finally, many other lines of evidence suggest that foraging animals systematically make errors or irrational choices in certain experimental situations. Modern thinking about these results appeals to the idea of ecological rationality, which hypothesizes that animals make errors because natural selection has shaped their decision-making apparatus to achieve high levels of benefit in specific natural contexts.

AUTHOR NOTE

I am grateful for the support of the U.S. National Science Foundation, which supported my research via Grants IOB-0235261 and IOS-0727221. Phillipe Tobler, Tobias Kalenscher, Alex Kacelnik, and two anonymous referees provided comments that greatly improved this

article. I appreciate their assistance. Correspondence concerning this article should be addressed to D. W. Stephens, Department of Ecology, Evolution, and Behavior, University of Minnesota, 1987 Upper Buford Circle, Saint Paul, MN 55108 (e-mail: dws@umn.edu).

REFERENCES

- ABRAHAM, M. V. (1989). Foraging guppies and the ideal free distribution: The influence of information on patch choice. *Ethology*, **82**, 116-126.
- ABRAHAM, M. V., & DILL, L. M. (1989). A determination of the energetic equivalence of the risk of predation. *Ecology*, **70**, 999-1007.
- ADAMS-HUNT, M. M., & JACOBS, L. F. (2007). Cognition for foraging. In D. W. Stephens, J. S. Brown, & R. C. Ydenberg (Eds.), *Foraging: Behavior and ecology* (pp. 105-138). Chicago: University of Chicago Press.
- AINSLIE, G. W. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*, **21**, 485-489.
- ALCOCK, J. (2005). *Animal behavior: An evolutionary approach* (8th ed.). Sunderland, MA: Sinauer.
- ALCOCK, J., & SHERMAN, P. (1994). The utility of the proximate-ultimate dichotomy in ethology. *Ethology*, **96**, 58-62.
- BATESON, M., & KACELNIK, A. (1996). Rate currencies and the foraging starling: The fallacy of the averages revisited. *Behavioral Ecology*, **7**, 341-352.
- BAUM, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, **22**, 231-242.
- CHARNOV, E. L. (1976). Optimal foraging: The marginal value theorem. *Theoretical Population Biology*, **9**, 129-136.
- CHEVERTON, J., KACELNIK, A., & KREBS, J. R. (1985). Optimal foraging: Constraints and currencies. In B. Hölldobler & M. Lindauer (Eds.), *Experimental behavioral ecology and sociobiology* (pp. 109-126). Sunderland, MA: Sinauer.
- COWIE, R. J. (1977). Optimal foraging in great tits (*Parus major*). *Nature*, **268**, 137-139.
- DAVIS, M. (2006). Neural systems involved in fear and anxiety measured with fear-potentiated startle. *American Psychologist*, **61**, 741-756.
- FERNÁNDEZ-JURICIC, E., KERR, B., BEDNEKOFF, P. A., & STEPHENS, D. W. (2004). When are two heads better than one? Visual perception and information transfer affect vigilance coordination in foraging groups. *Behavioral Ecology*, **15**, 898-906.
- FLOMBAUM, J. I., & SANTOS, L. R. (2005). Rhesus monkeys attribute perceptions to others. *Current Biology*, **15**, 447-452.
- FRETWELL, S. D., & LUCAS, H. L. (1970). On territorial behavior and other factors influencing habitat distribution in birds. *Acta Biotheoretica*, **19**, 16-36.
- GALLISTEL, C. R. (1994). Foraging for brain stimulation: Toward a neurobiology of computation. *Cognition*, **50**, 151-170.
- GIRALDEAU, L.-A., & KRAMER, D. L. (1982). The marginal value theorem: A quantitative test using load size variation in a central place forager, the eastern chipmunk, *Tamias striatus*. *Animal Behaviour*, **30**, 1036-1042.
- HASTIE, R. (2001). Problems for judgment and decision making. *Annual Review of Psychology*, **52**, 653-683.
- HERRNSTEIN, R. J., RACHLIN, H., & LAIBSON, D. I. (Eds.) (1997). *The matching law: Papers in psychology and economics*. Cambridge, MA: Harvard University Press.
- KILLEEN, P. R. (1968). On the measurement of reinforcement frequency in the study of preference. *Journal of the Experimental Analysis of Behavior*, **11**, 263-269.
- MAZUR, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior: The effect of delay and of intervening events on reinforcement value* (pp. 55-73). Hillsdale, NJ: Erlbaum.
- MAZUR, J. E. (2001). Hyperbolic value addition and general models of animal choice. *Psychological Review*, **108**, 96-112.
- MCLINN, C. M., & STEPHENS, D. W. (2006). What makes information valuable? Signal reliability and environmental uncertainty. *Animal Behaviour*, **71**, 1119-1129.
- MOBBS, D., PETROVIC, P., MARCHANT, J. L., HASSABIS, D., WEISKOPF, N., SEYMOUR, B., ET AL. (2007). When fear is near: Threat imminence

- elicits prefrontal–periaqueductal gray shifts in humans. *Science*, **317**, 1079-1083.
- NONACS, P. (2001). State dependent behavior and the marginal value theorem. *Behavioral Ecology*, **12**, 71-83.
- PARKER, G. A. (1978). Searching for mates. In J. R. Krebs & N. B. Davies (Eds.), *Behavioural ecology: An evolutionary approach* (pp. 214-244). Sunderland, MA: Sinauer.
- PITT, W. C. (1999). Effects of multiple vertebrate predators on grasshopper habitat selection: Trade-offs due to predation risk, foraging, and thermoregulation. *Evolutionary Ecology*, **13**, 499-516.
- PLATT, M. L. (2002). Neural correlates of decisions. *Current Opinion in Neurobiology*, **12**, 141-148.
- POMPILO, L., KACELNIK, A., & BEHMER, S. T. (2006). State-dependent learned valuation drives choice in an invertebrate. *Science*, **311**, 1613-1615.
- RACHLIN, H., & GREEN, L. (1972). Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior*, **17**, 15-22.
- SANFEY, A. G. (2007). Social decision-making: Insights from game theory and neuroscience. *Science*, **318**, 598-602.
- SCHUCK-PAIM, C., POMPILO, L., & KACELNIK, A. (2004). State-dependent decisions cause apparent violations of rationality in animal choice. *PLoS Biology*, **2**, 2305-2315.
- SHAFIR, S. (1994). Intransitivity of preferences in honey bees: Support for comparative evaluation of foraging options. *Animal Behaviour*, **48**, 55-67.
- SHAFIR, S., WAITE, T. A., & SMITH, B. H. (2002). Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioral Ecology & Sociobiology*, **51**, 180-187.
- SHERRY, D. F. (2006). Neuroecology. *Annual Review of Psychology*, **57**, 167-197.
- SHERRY, D. F., & MITCHELL, J. B. (2007). Neuroethology of foraging. In D. W. Stephens, J. S. Brown, & R. C. Ydenberg (Eds.), *Foraging: Behavior and ecology* (pp. 61-102). Chicago: University of Chicago Press.
- SHETLEWORTH, S. J. (1998). *Cognition, evolution, and behavior*. New York: Oxford University Press.
- STEPHENS, D. W. (2007). Models of information use. In D. W. Stephens, J. S. Brown, & R. C. Ydenberg (Eds.), *Foraging: Behavior and ecology* (pp. 31-58). Chicago: University of Chicago Press.
- STEPHENS, D. W., & ANDERSON, D. (2001). The adaptive value of preference for immediacy: When shortsighted rules have farsighted consequences. *Behavioral Ecology*, **12**, 330-339.
- STEPHENS, D. W., BROWN, J. S., & YDENBERG, R. C. (Eds.) (2007). *Foraging: Behavior and ecology*. Chicago: University of Chicago Press.
- STEPHENS, D. W., KERR, B., & FERNÁNDEZ-JURICIC, E. (2004). Impulsiveness without discounting: The ecological rationality hypothesis. *Proceedings of the Royal Society B*, **271**, 2459-2465.
- STEPHENS, D. W., & KREBS, J. R. (1986). *Foraging theory*. Princeton: Princeton University Press.
- STEPHENS, D. W., & MCLINN, C. M. (2003). Choice and context: Testing a simple short-term choice rule. *Animal Behaviour*, **66**, 59-70.
- STEPHENS, D. W., & STEVENS, J. R. (2001). A simple spatially explicit ideal-free distribution: A model and an experiment. *Behavioral Ecology & Sociobiology*, **49**, 220-234.
- TINBERGEN, N. (1963). On aims and methods of ethology. *Zeitschrift für Tierpsychologie*, **20**, 410-433.
- TODD, P. M., & GIGERENZER, G. (2000). Precis of simple heuristics that make us smart. *Behavioral & Brain Sciences*, **23**, 727-780.
- TREGENZA, T. (1995). Building on the ideal free distribution. *Advances in Ecological Research*, **26**, 253-307.
- WAITE, T. A., & FIELD, K. L. (2007). Foraging with others: Games social foragers play. In D. W. Stephens, J. S. Brown, & R. C. Ydenberg (Eds.), *Foraging: Behavior and ecology* (pp. 331-362). Chicago: University of Chicago Press.

(Manuscript received March 24, 2008;
revision accepted for publication August 7, 2008.)