Reprint:
G. J. Madden & W. K. Bickel (Eds.). (2009). Impulsivity. *The Behavioral and Neurological Science of Discounting* (pp. 361–387). Washington, DC: APA.

13

THE ADAPTIVE NATURE OF IMPULSIVITY

JEFFREY R. STEVENS AND DAVID W. STEPHENS

An old joke circulates among animal behavior instructors. One can, the joke goes, divide the topics of animal behavior into four *F*s: fighting, fleeing, feeding, and reproduction. This somewhat tired joke carries considerable truth. Animals behaving in nature surely must make decisions about conflicts, predator avoidance, feeding, and mating. Male crickets, for example, are notoriously combative. Studies have shown, however, that they escalate fights in some situations and retreat to fight another day in others (Beaugrand, 1997; Parker, 1974; Parker & Rubenstein, 1981). Squirrels, like many small animals, respond to the presence or absence of protective cover; for example, they will carry large food items into the safety of the bushes to consume them but eat small items immediately (Lima, Valone, & Caraco, 1985). Female widow birds prefer males with long tails, and evolutionary theorists have argued that tail length correlates with male quality (Andersson, 1982, 1994). So a female confronted with a short-tailed male faces a dilemma: mate now or keep looking. Notice that in all of these choice situations, time complicates the animal's problem: Risk injury by fighting now or retreat to fight later; stay exposed to possible predation or invest time in moving to a safer place; settle for the short-tailed male or keep looking. Each of these situations, and indeed virtually any naturally occurring choice situation one can imagine, is

an intertemporal choice problem. We define these as choice situations in which an animal's alternatives vary in the time at which the animal realizes consequences and in the quality of those consequences once the animal secures them.

Although intertemporal choice applies to many domains (and all four *F*s), we need to focus on a specific situation to make scientific headway, and for virtually all behavioral ecologists interested in intertemporal choice that focal situation is foraging. We can observe animal foraging choices easily (e.g., animals eat more often than they reproduce), and we can manipulate the time and magnitude of foraging options much more easily than we can manipulate mate quality or predation risk. Moreover, we have a large base of theoretical and empirical results that help us frame the intertemporal choice problem in the context of animal foraging behavior. Foraging is not only a convenient topic but also a fundamentally important one; actively seeking food is a basic part of animal existence that deserves our attention. In the first part of this chapter, we focus on adaptive aspects of intertemporal choice in animal foraging behavior, and especially on the problem of impulsivity, which we see as a central problem in intertemporal choice. In the second part of this chapter, we take a broader perspective, including domains other than food and extending beyond impulsivity to a more encompassing view of intertemporal choice. Within this general view, we explore the adaptive nature of impulsivity.

FORAGING AND INTERTEMPORAL CHOICE

A pileated woodpecker (*Dryocopus pileatus*)—a crow-sized North American woodpecker—works its way through the trees along a suburban lake. It typically lands on a tree trunk and moves up the trunk, making distinctive hopping motions as it goes. The pileated woodpecker feeds on wood-boring insects, and it uses its beak to chisel its prey from their galleries below the bark. What sorts of choices must a foraging woodpecker make? It will surely make choices about where to search (along the lake shore or along a ridge), which behavioral ecologists call *habitat choice* decisions. It will also make choices about how to search (how fast to fly, where to land on a tree, which parts of the tree to focus on). It will make choices about what to eat; in lean periods, it will attack small prey that it might pass by in better times.

Although the woodpecker must make many decisions as it forages, we focus for the moment on the woodpecker's problem of deciding how thoroughly to exploit each tree before flying to the next one. Should it make a few quick probes and move on to the next tree, or should it exhaustively check every crevice and abnormality in the bark? The reader will immediately rec-

ognize costs and benefits of both strategies. The quick, "cream-skimming" visits will lead to much time spent traveling, and the departing woodpecker may leave many good food items behind, whereas the "bowl-licking" strategy will produce many prey per tree, but the woodpecker may waste time extracting the last dregs instead of moving to a fresh tree. This observation should lead the reader to think that foragers should be sensitive to environmental richness. In a rich environment, the next tree may offer a feast, so we would expect rich environments to favor cream skimming, whereas lean environments should favor thorough exploitation.

Behavioral ecologists call this the problem of *patch exploitation*, and it plays a central role in our thinking about animal foraging decisions. Moreover, we have both a well-developed body of theory about this problem and extensive experimental and observational data (Stephens, Brown, & Ydenberg, 2007; Stephens & Krebs, 1986). Formal models of how animals exploit patches of food begin by considering the relationship between the time spent exploiting a patch and the amount of food the animal extracts. We call this relationship the *gain function*. Although the gain function can take many forms, the most plausible and common form is a negatively accelerated shape, as Figure 13.1 shows. The amount of food in the patch sets an upper limit, so the gain function will asymptote to this maximum. The bending of the gain function captures an important property of natural resource exploitation: Resources deplete, and finding the next unit of food from a clump is typically harder and more time consuming than finding the previous unit of food. During a foraging bout, the animal spends its time doing two things: exploiting

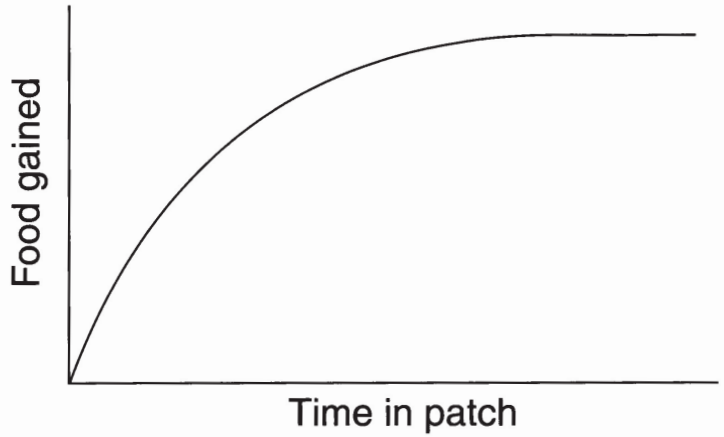

Figure 13.1. Hypothetical gain function. The gain function associated with extracting food from a patch is likely not linear. Instead, there are diminishing returns: As a patch depletes, it takes longer to find food.

patches or traveling to new patches. If τ represents the average time required to travel from one patch to the next, then the overall rate of food intake is

$$\frac{g(t)}{\tau + t},$$

where $g(t)$ is the patch gain function, and t is the time spent exploiting each patch. (Notice that for simplicity, we assume a situation with only one type of patch; we can easily generalize this approach to more types.) Mathematically, we can ask what value of t (the time spent exploiting patches, or patch-residence time) maximizes this intake rate, and we find that the best patch-residence time (t) is large when travel times between patches are large (i.e., in poor environments) and small when travel times are small (i.e., in rich environments; see Stephens & Krebs, 1986).

This model makes the qualitative prediction that animals should skim the cream (spend a short time, extracting relatively little) when the travel time is short (rich environments) but that they should adopt a thorough exploitation strategy (spend more time, extracting more) when travel time is long. Do the data support this claim? Yes, they do; we have very strong evidence from laboratory experiments, field experiments, and simple observations across many species and habitats that travel time affects patch-exploitation behavior, with increases in travel time leading to increases in patch-exploitation time (Stephens & Krebs, 1986). We develop a simple analysis of patch exploitation here to highlight an important naturally occurring problem in intertemporal choice and its basic structure. Two key features of the patch-exploitation problem stand out. First, the patch-exploitation problem is fundamentally sequential. Foragers make a sequence of patch visits such that the quality of the next patch influences the relative value of the present patch. Second, notice that although a forager exploiting a patch must choose between staying and leaving, choosing to stay does not prevent the animal from eventually leaving. This contrasts sharply with many experimental studies of choice behavior in which the investigators set up a situation in which subjects must make mutually exclusive choices: Choosing Option A irrevocably slams the door on Option B.

IMPULSIVITY AND SELF-CONTROL

Traditional models of foraging, like the patch-exploitation model, are derived from the premise of maximizing intake rate over a sequence of foraging choices, an assumption we call *long-term rate maximization*. Although one can raise many possible objections to this premise, it has served students of foraging quite well overall, as the success in predicting the relationship between

travel time and patch-residence time shows. The long-term rate approach has, however, repeatedly failed to predict choice in laboratory studies using the self-control (or delay-discounting) paradigm. In self-control studies, the investigator trains subjects (typically pigeons [*Columba livia*], or rats [*Rattus norvegicus*]) to choose between a small reward it can obtain quickly (usually called smaller–sooner), and a larger reward it must wait a bit longer to obtain (larger–later; see Figure 13.2, Panel (a)). Using this scheme, the investigator can explore the effects of delay and amount on preference. The reader may notice that this procedure crudely resembles patch exploitation. The smaller–sooner option resembles a short patch stay yielding a smaller amount; conversely, the larger–later option is like staying longer and obtaining more. We might expect,

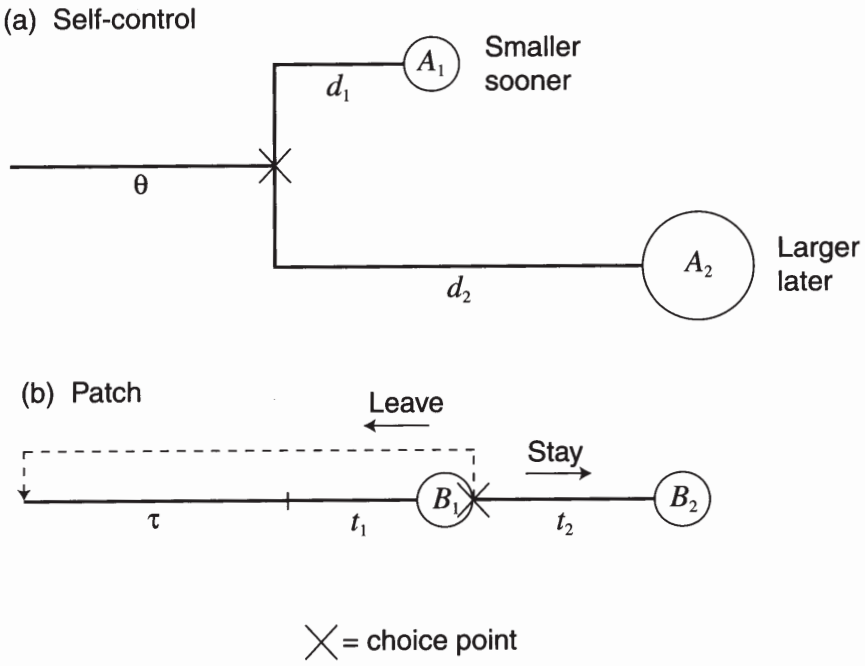

Figure 13.2. Self-control and patch experimental designs. (a): In the self-control condition, subjects begin with an intertrial interval θ before facing a simultaneous choice between a smaller–sooner reward (A_1 = small amount, d_1 = short delay) and a larger–later reward (A_2 = large amount, d_2 = long delay). After experiencing the delay and consuming the reward, subjects begin another intertrial interval. (b): In the patch condition, subjects face a sequential choice between staying in the patch or leaving to find a new patch. Thus, all subjects wait for an intertrial interval τ, experience a short delay t_1 and receive a small reward B_1. Then, they choose between staying in the patch an additional delay t_2 and receiving the additional reward B_2 or leaving the patch to start another intertrial interval. Stephens and Anderson (2001) offered these conditions to blue jays and set the parameters such that choices in the self-control and patch treatments were economically equivalent.

therefore, some correspondence between observed choice in the two situations. Yet, results from the two approaches diverge dramatically. Animals in self-control situations commonly prefer the smaller–sooner option even when the larger–later option leads to a higher long-term rate of intake (Bateson & Kacelnik, 1996; Mazur, 1987; McDiarmid & Rilling, 1965). In addition, the time between choice presentations (intertrial interval—the experimental analogue of travel time) has virtually no effect on choice in the self-control preparation (Mazur & Romano, 1992), even though patch studies have nearly universally shown that travel time affects patch exploitation.

We describe this pattern of choice in self-control situations as impulsive. We define *impulsivity* as choosing a smaller–sooner option when a larger–later option produces a better outcome. In broad strokes, this definition agrees with the day-to-day meaning that impulsive decisions lead to error. "I bought the iPhone impulsively" means that had I stopped to think through the long-term consequences, I would not have made this purchase. Of course, to identify impulsivity in nonhumans, we need to say what we mean by a "better outcome." Motivated by foraging models and the tools they offer for calculating rate, here we operationally define impulsivity as choosing smaller–sooner alternatives when the larger–later option yields a higher long-term rate of intake. Clearly, observed behavior in the self-control situation satisfies this definition. In these experiments, subjects often obtain less food than they could, creating an evolutionary puzzle. Why should natural selection favor choice mechanisms that produce less food? It would seem to be relatively simple to "engineer" a decision-making system that does better. Why has natural selection not done this?

EVOLUTIONARY APPROACHES TO IMPULSIVITY

As we have outlined here, the data suggest that animals consistently favor immediate rewards even though it seems that they could achieve higher fitness gains by choosing more delayed options. For behavioral ecologists, the natural first response to this puzzle is to speculate that immediacy is valuable in some way that our traditional rate-based models fail to capture. Specifically, some have suggested that increasing delay reduces, or "discounts," the value of delayed benefits. We remark that the phrase *delay discounting* is closely linked to the study of impulsive choice, so much so that delay discounting and impulsivity are identical in the minds of some authors. From an evolutionary perspective, however, we see the puzzle of impulsivity as a description of the phenomenon of interest and delay discounting as an explanatory hypothesis derived from economic principles.

How can delay reduce value? There are two possibilities: costs due to collection risk and lost investment opportunity. According to the collection-risk hypothesis, the animal has a better chance of collecting (or realizing) more

immediate options. Behavioral ecologists often call this the "discounting-by-interruptions" hypothesis. In this hypothesis, we suppose that interruptions prevent an animal from collecting any delayed benefit. The model requires that interruptions occur in some time-dependent way, so that they occur more frequently in long intervals, thereby enhancing the relative value of short delays. In nature, these interruptions could come from groupmates, from predators, from weather, or even from prey escaping, and we have clear evidence that interruptions occur. For example, groupmates often steal food from each other, so that a social forager who has located food surely experiences some risk of losing it, and presumably this risk increases with the time spent handling the food. Unfortunately, although we have many casual observations like this, we have virtually no rigorous information about the critical statistical properties of these interruption processes. To account for observed levels of impulsivity, we would need fairly high interruption rates, likely much higher than exist in nature.

For lost investment opportunities, consider a human investor who must choose between $100 now and $100 delayed for 1 year. Clearly, the $100 now is more valuable because if the investor waits, she or he pays an opportunity cost by forgoing the benefits of investing the $100 for 1 year. We can imagine situations in which this may be relevant to nonhuman animals. For example, if the next unit of food allows a forager to improve its breeding status or dominance, this could produce investment opportunity costs if benefits begin to accrue as soon as the animal gains its new status; so the sooner you obtain your new status, the better. Behavioral ecologists have not paid much attention to this possibility, perhaps because this sort of lost investment seems remote from the day-to-day world of animal behavior.

DELAY-DISCOUNTING MODELS OF IMPULSIVITY

We can explore more rigorously the delay-discounting idea that delay reduces value by using mathematical models. Students of choice have focused on two delay-discounting models, the exponential and hyperbolic models, which we discuss in the following sections (see also chaps. 1 and 3, this volume).

Exponential Model

Imagine that the value of a delayed resource decays at a constant rate as delay increases. This could happen because interruptions occur at constant rate or because investment opportunity costs accrue at a constant rate, or some combination of both. Restricting our attention to the interruptions idea simplifies our development. If interruptions occur at constant rate λ, then the expected value of a food reward of size A (for amount) delayed by d seconds is simply $V_d = Ae^{-\lambda d}$. If we knew or could estimate the interruption rate, λ, then we could, in theory, predict the outcome of a self-control experiment by comparing $A_S e^{-\lambda d_S}$ to $A_L e^{-\lambda d_L}$, where A_S, d_S, A_L, and d_L represent the amounts and delays associated with the smaller–sooner and larger–later alternatives. This otherwise appealing model fails to predict the phenomenon of preference reversal, which is a key observation in the self-control preparation (Ainslie, 1975; Ainslie & Herrnstein, 1981; Green, Fisher, Perlow, & Sherman, 1981). Preference reversal occurs when a subject's preference switches from smaller–sooner to larger–later as the delays associated with both options increase by the same amount. The exponential model cannot predict this because it predicts that adding a constant delay should modify the discounted value of both options by the same factor and therefore should not influence choice. Yet, animals are sensitive to this sort of general increase in delay; increasing both delays shifts preference toward the larger–later alternative.

Hyperbolic Model

At least partially in response to the failings of the exponential model, Mazur (1987) offered an algebraically simple delay-discounting model termed *hyperbolic discounting* (Ainslie, 1975), and Mazur's elegant empirical work has lead to broad acceptance of this approach. According the hyperbolic model, the expression

$$V_d = \frac{A}{1 + kd}$$

describes the decline of value with delay. The hyperbolic model can easily accommodate preference reversal. To predict the outcome of a self-control test we would compare $\frac{A_S}{1 + kd_S}$ and $\frac{A_L}{1 + kd_L}$. We call k the hyperbolic discount factor, and its meaning crudely parallels the meaning of λ in the exponential model: High k means that value declines more steeply with increasing delay—more discounting. Notice, however, that we can, in principle, calculate the exponential model's λ parameter a priori, for example, by observing or manipulating the interruption rate. In contrast, we must estimate the hyperbolic model's k parameter from observed preference. That is, k is a fitted parameter. From the perspective of our attempts to explain how natural selection influences patterns of intertemporal choice, the hyperbolic model offers a description of observed choice rather than an explanation.

Although these standard delay-discounting models play important roles in the fields of animal behavior and psychology, they inevitably leave those interested in the evolution of intertemporal choice unsatisfied. The exponential model offers the promise of explanation from first principles but ultimately fails empirically. The hyperbolic model describes many data but has little

explanatory power. In addition, the models suffer from a seldom-recognized disconnect with naturally occurring decision making in foraging animals. Both focus on a single decision abstracted from the remainder of the animal's activities (Kacelnik, 2003). Although one can see that this "single-shot" approach may be an informative theoretical device—like a physicist's assumption of frictionless world—it is difficult to imagine animals facing single-shot decisions in the real world. Experiments with nonhumans require repeated trials for subjects to learn the payoffs and delays. Data from natural foraging (like patch exploitation) have suggested sensitivity to future consequences; indeed, even within the laboratory paradigm of self-control testing, we have evidence of sensitivity to consequences beyond the current choice (Mazur, 1994). In other words, it is clear that animal choice typically occurs in a sequential context, so that any model purporting to account for the economic consequences of choice needs to consider this basic reality. Yet, the empirical success of the single-shot hyperbolic model shows that we can often describe choice without references to this sequence. In our view, this is another perspective on the puzzle of impulsivity: Animals live in a world of sequential decision making, yet they often behave as if only the next choice matters to them.

Short-Term Rate Model

For some purposes, investigators prefer a simplified version of the hyperbolic model called the *short-term rate model*. Although not technically a discounting model, it can predict impulsive choice. To apply this model to the self-control situation, we would compare $\frac{A_S}{d_S}$ and $\frac{A_L}{d_L}$; that is, a rate comparison without the intertrial interval that we need to calculate the long-term rate over a sequence of trials. In effect, this model hypothesizes that the subject considers the rate from the choice point to the food delivery but nothing else. In practice, the hyperbolic and short-term rate models make similar predictions except when delays are very small, but the short-term rate model is convenient conceptually because we can compare it so easily to our standard of "economically sound" choice, the long-term rate model. Empirical studies that have tested a modified version of this model have shown that it can account for data in European starlings (*Sturnus vulgaris*) and cotton-top tamarins (*Saguinus oedipus*), even without the fitted k parameter of the hyperbolic model (Bateson & Kacelnik, 1996; Stevens, Hallinan, & Hauser, 2005).

Discounted-Sequence Model

Nothing about the concept of discounting necessitates single-shot decision making. We can, for example, build delay-discounting models that consider sequences of gains. Stephens (2002) offered a simple discounted-sequence model by considering a sequence of exponentially discounted gains, which he called the *exponentially discounted-sequence model*. This logically simple approach yields a hybrid of foraging theory's long-term rate model and the single-shot exponential model. Like the single-shot exponential model, it includes an a priori discount rate parameter—so, for example, we know how variations in interruption rate should affect the value of a sequence. It can, in principle, accommodate preference reversal (as any rate-based model can). Unfortunately, like most a priori models of discounted value, it fails to account for the self-control data quantitatively. For example, we need very high discount rates to explain the strength of observed preferences for immediacy, yet at these high discount rates the model loses its ability to predict preference reversal.

The failure of the exponentially discounted-sequence model illustrates the problems associated with finding a successful economic account of impulsive choice behavior. We can capture ratelike properties such as preference reversal easily enough, but the strength of preferences for immediacy suggests very high discount rates, and it is hard to imagine natural processes (interruption or opportunity cost) that could create such high discount rates. In addition, the discounting-by-interruption hypothesis that behavioral ecologists have long favored seems increasingly strained. Experiments on birds and humans (Henly et al., 2008; King & Logue, 1992) failed to find the predicted increase in preference for smaller–sooner options with experimentally created interruptions in a self-control situation. Suffice it to say, then, efforts to understand observed choice in terms of a priori models of discounted value leave something to be desired. Even though the economic forces that reduce the value of delayed benefits (opportunity costs and collection risk) are surely real, they seem too weak to explain the strong preferences for immediacy we see in the self-control situation.

Ecological Rationality of Impulsivity

As explained earlier, animals exploiting patches change their behavior when travel time increases, spending more time in patches and extracting more food; in contrast, animals tested in the self-control preparation seem insensitive to changes in the experimental analogue of travel time, the intertrial interval. Intrigued by this difference, Stephens and Anderson (2001) created an experimental comparison of the patch and self-control situations. To achieve this, Stephens and Anderson needed an experimentally tractable patch-choice situation that they could compare with self-control treatments. Figure 13.2 shows their approach. In the self-control situations (Figure 13.2, Panel A), the subject waits for a fixed intertrial interval θ and chooses between amount A_1 after delay d_1 and amount A_2 after delay d_2. In the patch situation (Figure 13.2, Panel B), the subject waits for a fixed intertrial interval–travel time, τ (note that $\theta = \tau$, but we use different notation to distinguish

between self-control and patch variables). When this time expires, the apparatus presents a single stimulus; in the jargon of foraging theory, this single stimulus indicates a patch encounter. The subject "enters the patch" (in Stephens and Anderson's preparation, a blue jay [*Cyanocitta cristata*] hops forward to a perch just below the patch stimulus), and then the subject waits again: waiting time t_1 for a food delivery of B_1 (two pellets). After this initial food delivery, the subject must make a choice: look for a new patch (i.e., start a fresh cycle of wait–encounter–wait–get food), or stay in the patch, waiting a bit longer (additional waiting time t_2) for an additional two food pellets (amount B_2). Although natural patches can, in principle, contain many items, these experimental patches never produce more than two food deliveries, so the trial effectively ends after this second food delivery and the subject must start a new cycle of wait–encounter and so on.

By choosing the amounts and delays carefully, Stephens and Anderson (2001) created economically comparable patch and self-control situations, that is, situations in which the "leave and start over" choice produced the same amount of food ($B_1 = A_1$ = two pellets) and took the same time ($\tau + t_1 = \theta + d_1$) as the smaller–sooner option and the "stay" choice produced the same amount of food and took the same time as the larger–later option ($B_1 + B_2 = A_2$ = four pellets and $\tau + t_1 + t_2 = \theta + d_2$). Even if this experimental patch situation is hopelessly simple minded as compared with the complexity of naturally occurring patch exploitation, it frames the intertemporal choice problem in terms of leave versus stay rather than the more conventional A-versus-B–style choice.

Stephens and Anderson (2001) created a wide range of treatments systematically varying delays, amounts, and intertrial intervals but always ensuring comparable patch and self-control versions of each condition. Figure 13.3 shows conventional dose–response style plots that assess the subjects' sensitivity to long-term rate. As the figure shows, the jays followed an orderly sigmoid response in the patch situation, roughly following the predictions of long-term rate: choosing to stay when this gave the highest long-term rate and choosing to leave when the opposite was true. Moreover, with extreme differences in rate, the jays strongly preferred the "correct" option. In contrast, the self-control data show a very weak response to long-term rate; although they tend to increase their preference for larger–later when this gives the higher long-term rate, they still take the smaller–sooner option quite frequently. Indeed, one can readily see the notorious impulsiveness of self-control treatments by observing that subjects chose the smaller–sooner option more frequently under almost all conditions.

WHY BE IMPULSIVE?

Animals perform poorly in the self-control situation, achieving lower rates of intake and acquiring less food than a hypothetical "more patient" ani-

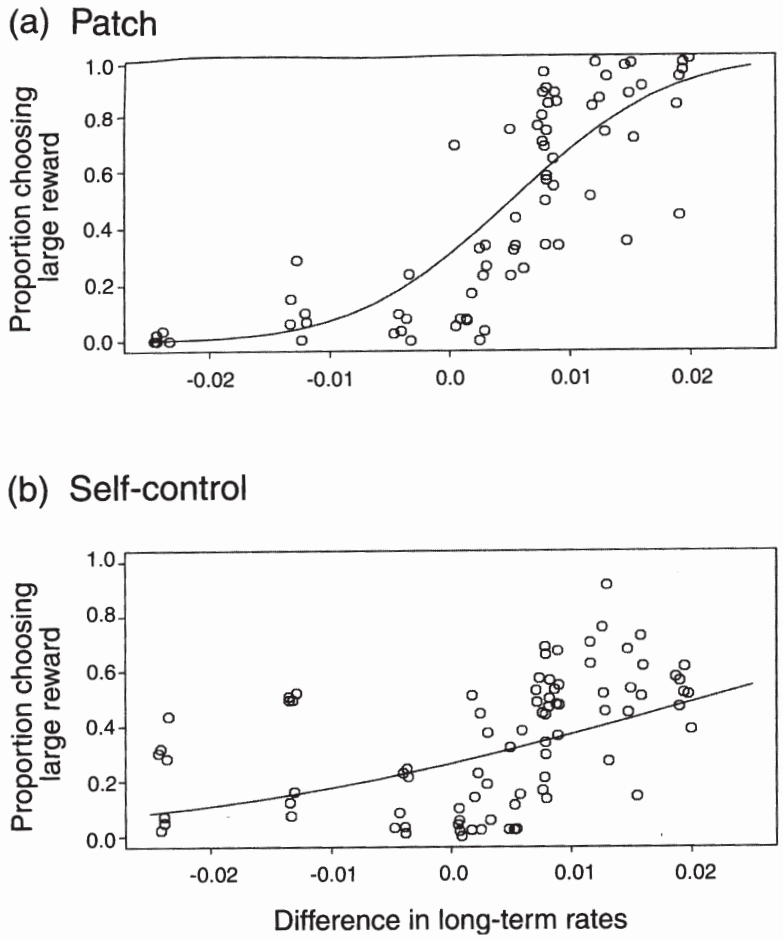

Figure 13.3. Influence of self-control and patch treatments on blue jay temporal preferences. The figure shows proportional choice of the option yielding the larger amount as a function of the difference in long-term rates associated with two options (Stephens & Anderson 2001). For example, zero means that large and small options produced the same long-term rate; a negative value means that the smaller option produced the higher long-term rate; and a positive value means the larger option yielded the higher long-term rate. Blue jays experienced treatments that varied in the intertrial interval, the delay to the smaller option, and the delay to the larger option. Each treatment pair created the same difference in long-term rates in the self-control and patch situations. The data points represent the performance of different individuals in each of the 12 treatments. The solid curves show the best least-square fit to a conventional dose–response relationship. As the data illustrate, the subject's behavior seems to track the difference in long-term rate in the patch situation but not in self-control. From "The Adaptive Value of Preference for Immediacy: When Shortsighted Rules Have Farsighted Consequences," by D. W. Stephens and D. Anderson, 2001, *Behavioral Ecology, 12,* p. 337. Copyright 2001 by Oxford University Press. Adapted with permission.

mal could. Yet, Stephens and Anderson's (2001) blue jays did quite well in the economically analogous patch situation. What explains this? There are two possibilities. Subjects could use different decision rules in the two situations; alternatively, the same mechanisms (e.g., decision rules such as "choose the option with the highest short-term rate") may simply work better in the patch situation.

Short-Term Rule, Long-Term Benefits

Stephens and Anderson (2001) argued for the second possibility. To see their argument, consider the long-term rates associated with the two options in patch choice. If the subject repeatedly leaves, it obtains a long-term rate of

$$\frac{B_1}{\tau + t_1}.$$

If the subject repeatedly stays, however, it obtains a long-term rate of

$$\frac{B_1 + B_2}{\tau + t_1 + t_2}.$$

Obviously, staying is better than leaving in the long run if

$$\frac{B_1 + B_2}{\tau + t_1 + t_2} > \frac{B_1}{\tau + t_1}.$$

This is a classical comparison in foraging theory that asks the following question: Will the forager benefit from adding an activity that yields amount B_2 in time t_2 to the things it is already doing (which, on average, yields amount B_1 every $\tau + t_1$ time units)? We can easily show that the answer is yes if and only if

$$\frac{B_2}{t_2} > \frac{B_1}{\tau + t_1}.$$

Although it may not be immediately obvious, this comparison is precisely the impulsive short-term rate rule we discussed earlier. In the patch situation, then, a subject who makes a short-term comparison (leaving, yielding B_1 in $\tau + t_1$ seconds, vs. staying, yielding B_2 in t_2 seconds) will coincidentally also be choosing the option that produces highest long-term rate. Yet, a subject who made the same short-term comparison in the self-control situation would compare B_1 in t_1 seconds and B_2 in t_2 seconds. By ignoring τ, subjects may be more likely to choose the smaller–sooner reward even though the larger–later reward offers a better long-term rate.

In light of this argument, the *ecological rationality* hypothesis holds that natural selection has favored short-sighted rules of choice because these rules fare well (achieve high long-term intake rates) in naturally occurring choice situations that have a structure similar to patch exploitation. However, these same rules fare poorly in tests of simultaneous, mutually exclusive choice, which are likely rare in nature. Notice that according to this hypothesis, selection could favor choice mechanisms that produce impulsive choice without discounting, even though the ecological rationality approach is not necessarily incompatible with discounting.

This application of the short-term rule to the patch situation hypothesizes that the subject views the intertrial interval (τ) and the first time to food (t_1) as a single combined delay. This is rather a bold claim given that self-control results suggest a strong asymmetry (delay has a powerful effect on choice, but the intertrial interval has little effect). To test this claim, Stephens and McLinn (2003) provided blue jays with a range of conditions using the same total delay (i.e., $\tau + t_1 = \text{constant}$) but different mixes of intertrial interval (τ) and initial delay (t_1). They showed that increasing the total delay ($\tau + t_1$) shifted preference toward staying as predicted in the patch situation, but the jays did not respond to the different mixes of intertrial intervals and initial delay.

Foreground/Background

Our description of the ecological rationality hypothesis describes the choice structure that favors impulsiveness rather vaguely as "similar to patch exploitation." What does that mean? Stephens, Kerr, and Fernandez-Juricic (2004) have developed a more explicit description of situations that favor impulsive decision making that they call *foreground/background* structure. As a forager moves through its habitat, it encounters many types of resources and presumably makes many choices about how to exploit them. At any moment, we suppose that the forager follows what we call a background strategy that represents its default or standard behavior. Obviously, the background strategy yields a stream of gains, and we could characterize the benefits associated with this background strategy using any model we favor (e.g., long-term rate). Occasionally, the forager encounters a new resource that represents an opportunity to deviate from the background strategy (Figure 13.4). We refer to this as a "foreground" option, and the forager must now choose between continuing with the background strategy or deviating to exploit the foreground option and then returning to the background strategy. Notice that the critical feature of this form is that the background strategy is part of both options, and this fact greatly reduces the costs of short-sighted decision making. Many naturally occurring choices follow this pattern, including but by no means limited to patch exploitation. We can view diet choice, mate choice, and many other resource exploitation decisions in this way.

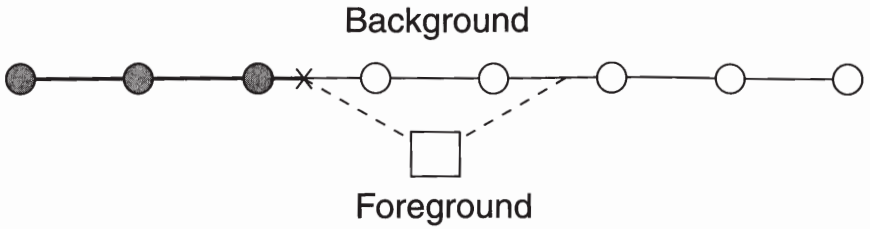

Figure 13.4. Foreground–background comparison. Animals often face choices between a "default" series of options (background) and the occasional chance to temporarily deviate from the default (foreground). In this example, a hypothetical forager has consumed three background options (filled circles) and must now choose between continuing with the background options (open circles) or opting for the foreground option (open square). It is important to note that choosing the foreground option forgoes the opportunity for some background options but eventually results in returning to the background strategy.

Discrimination Advantage

As we explained earlier, in the patch situation an animal using the short-term rule can achieve the same long-term intake rate as an animal that directly compares long-term rates. Some recent analyses have suggested an even stronger result (Stephens, 2002; Stephens et al., 2004). If judgments about differences in time follow Weber's law (Gibbon, 1977; Gibbon, Church, Fairhurst, & Kacelnik, 1988), then a short-term comparison can, in principle, be superior to a long-term comparison in the self-control situation. An animal that compares long-term differences will necessarily compare larger quantities because the intertrial interval is included, whereas an animal comparing short-term differences will compare smaller quantities. Because perceptual mechanisms usually generate more error when they compare larger quantities (e.g., 3 seconds vs. 5 seconds is easier to discriminate than 11 seconds vs. 13 seconds; Figure 13.5), an animal making long-term comparisons would discriminate between the two delays less accurately, resulting in occasional choices for a lower intake rate. This leads to the paradoxical conclusion that an animal making short-term comparisons could, in principle, achieve a higher long-term intake rate than an animal making long-term comparisons.

Although one can use discounting functions to describe animal impulsivity in the self-control paradigm, finding explanatory discounting models has proved quite challenging. Current evidence has suggested, for example, that interruptions do not create a bias favoring immediacy as delay-discounting models have long hypothesized (Henly et al., 2008; King & Logue, 1992). In addition, although we find explanatory models that predict the qualitative properties (e.g., preference reversal) correctly, the same models fail to capture the quantitative properties of observed choice. In the face of these challenges,

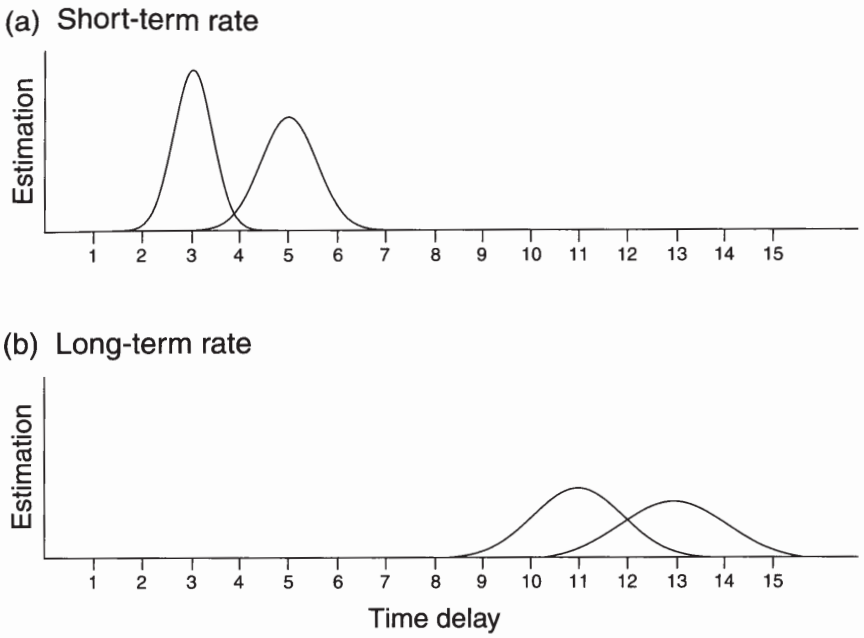

Figure 13.5. Discrimination advantage of short-term rates. Scalar expectancy theory states that the variance around estimates of temporal duration increases with the magnitude of the duration (Gibbon, 1977). Shorter durations, therefore, are easier to discriminate than longer durations. For instance, there may be little variance for estimations of 3- and 5-second durations, so discriminating between them is easy. When adding a constant time (e.g., the intertrial interval) of 8 seconds to these durations, however, the variance dramatically increases for 11 and 13 seconds, and estimates greatly overlap for these durations, making them difficult to distinguish.

the ecological rationality hypothesis offers an exciting and important alternative that can explain impulsive choice without discounting, although we could incorporate discounting into ecological rationality models if this proves to be important. We are just beginning to explore the ecological rationality idea; further experimentation and theorizing will undoubtedly refine this approach.

Adaptive Hypotheses for Intertemporal Choice

Up to now, our discussion has focused on the evolutionary puzzle of impulsivity, and we have argued that ecological rationality offers a new and compelling explanation of impulsivity in self-control experiments. In this final section, we take a larger view in two senses. First, we focus on the more general problem of intertemporal choice rather than the narrower topic of impulsivity. That is, we consider a broader view of intertemporal choice that

includes the vast range of situations in which animals must choose between options that vary in both time and magnitude. Second, we consider how intertemporal choice problems arise in nature, and how natural selection, acting in real-world ecological situations, may have shaped animal mechanisms of intertemporal choice. Drawing from the models of impulsivity reviewed earlier, we recognize three explanatory principles. First, we can probably explain many aspects of naturally occurring intertemporal choice using simple rate models, as traditional models of foraging behavior advocate. Second, the economic forces associated with delay discounting (collection risk and the opportunity costs of investments) may be important for many species. Finally, approaches based on the premise of ecological rationality emphasize the mismatch between rules that work well in natural situations and the rules required to make economically perfect decisions. We recognize these as distinct approaches mostly because they have been handed down to us via different research traditions, yet the boundaries between them are often fuzzy. Nevertheless, we find it helpful to use these three approaches to organize this final section on the adaptive nature of intertemporal choice.

RATE EFFECTS ON INTERTEMPORAL CHOICE

Theoretical behavioral ecology offers an extensive menu of rate-based models, covering many different domains. We have rate models that deal with patch exploitation, diet choice, habitat choice, joining social groups, and so on (reviewed in Giraldeau & Caraco, 2000; Stephens et al., 2007; Stephens & Krebs, 1986). Yet, students of intertemporal choice often pass by rate-based explanations in the search for more elaborate approaches. Indeed, our discussion of impulsivity earlier in the chapter explicitly defined rate out of the problem. Although there certainly are situations in which we need to look beyond rate models to explain intertemporal choice, we should not throw the baby out with the bath water. Rate models describe the basic economics of repeated choice in a simple and powerful way. At worst, rate models provide baseline expectations (as in the study of impulsivity) and at best they make well-supported predictions (as in patch exploitation).

So what can rate models tell us about intertemporal choice in nature? Consider again the basic patch-exploitation problem described earlier. At one extreme, we can imagine an "impatient" cream-skimming tactic that grabs the good stuff in a patch and quickly moves on, and at the other extreme we might have a "patient" bowl-licking strategy extracting even the last dregs before moving to the next patch. Rate-based models tell us that an animal's options elsewhere should set the balance between these two approaches. In a rich environment, the animal has many good alternatives beyond the current patch, so we predict something like cream skimming; however, in a poor environment

the current patch looks pretty good compared with the animal's options elsewhere, and we predict a strategy approximating bowl licking. Crudely speaking then, we expect that species adapted to rich environments should be less patient. Experimental data on chimpanzees (*Pan troglodytes*) and bonobos (*Pan paniscus*) in captivity support this prediction (Rosati, Stevens, Hare, & Hauser, 2007). In the wild, bonobos live in a richer environment, feeding more heavily on abundant herbaceous vegetation. Chimpanzees, in contrast, forage on fruit more than herbaceous vegetation, which is less abundant in their habitats (Malenky & Wrangham, 1994). When tested in repeated self-control tasks, bonobos seemed less patient, opting for the cream-skimming technique.

Our hypothetical cream skimmer leaves food patches early because staying longer would reduce its chance to acquire food at a higher rate elsewhere. This, of course, follows from the assumption that an animal spending time in the current patch cannot simultaneously spend its time looking for new patches. This is the bedrock assumption of rate-based models. They owe their successes (and their limitations) to the simple way in which they caricature this basic trade-off. Yet, this is not always true. Some animals can search for new resources while they exploit others. A web-building spider can, for example, multitask in this way because it can extract nutrients from one prey item while its web works to capture a second. Some species, then, experience a sort of release from the conventional exploit versus search trade-off. Animals that experience no exploit–search trade-off should not really care about exploitation delays—after all, it costs them nothing.

Discounting Effects on Intertemporal Choice

Delay discounting refers to a reduction in a reward's value caused by delay. We are cautious with applying discounting as an explanation of choice in the self-control situation for two reasons. First, we now have several hypotheses that can, in principle, account for specific laboratory results without discounting (Bateson & Kacelnik, 1996; Kacelnik, 2003; Stephens & Anderson, 2001). Second, direct manipulation of delay-discounting variables did not have the predicted effect on choice (Henly et al., 2008; King & Logue, 1992). Despite these setbacks in the laboratory, we have good reasons to believe that delay-discounting effects—both collection risk and opportunity costs—have shaped animal temporal preferences in some important situations. The laboratory tasks may not capture the ecologically relevant risks and opportunity costs that animals face in their natural environments.

Collection Risk

As we explained previously, animals may discount delayed rewards because the future is uncertain; some event—a so-called interruption—may

prevent the animal from collecting a delayed reward. The discounting-by-interruptions hypothesis seems compelling to many students of animal behavior because natural interruptions can come in many forms, such as social interruptions, interruptions by predators, and so on. Indeed, some evidence has suggested that interruptions and cues predicting interruptions can influence intertemporal choice. This section reviews naturally occurring interruptions and the evidence that these interruptions affect intertemporal choice.

Life Expectancy

Death ultimately interrupts us all, humans and nonhumans alike. An evolutionary approach predicts differences in intertemporal choice for species that vary in life span. Mathematically, we can substitute the mortality rate (a classic parameter from demography) into the exponential discounting model to motivate a life-expectancy–based model of delay discounting. Clearly, adult insects with very short life spans would not do well to delay payoffs for long periods, when they are unlikely to survive long enough to reap the rewards. We intuitively expect, therefore, that intertemporal choice should scale with expected life span. To explore this, we can evaluate the comparative data shown in Figure 13.6. Here, we see that Old World monkeys and apes wait much longer than birds, rodents, and other primates. Although at first glance, life span appears to correlate with these comparative data, on closer inspection this is not the case. On the one hand, pigeons, cotton-top tamarins, common marmosets (*Callithrix jacchus*), and capuchin monkeys (*Cebus apella*) live much longer than rats but share similar temporal preferences. On the other hand, chimpanzees and bonobos have similar life spans but quite different temporal preferences. Life span and mortality-rate effects probably only influence intertemporal choice at the extremes—for example, explaining differences between very short-lived and very long-lived species.

Although the effect of life span on intertemporal choice may be swamped by other ecological factors at the species level, life expectancy may have a strong effect at the individual level. An adaptive approach predicts that temporal preferences should change over the life span of organisms (Daly & Wilson, 2005; Sozou & Seymour, 2003). In fact, as organisms age or perceive cues associated with a short life expectancy, they often prefer immediate payoffs rather than waiting for delayed payoffs. For instance, when parasitoid wasps (*Leptopilina heterotoma*) detect cues associated with a short life expectancy, they lay more eggs in lower quality hosts than in the absence of the cues (Roitberg et al., 1992). The wasps accept a lower reproductive output when the future appears uncertain. Thus, individuals can flexibly respond to a shortened temporal horizon, a finding also reported in older human adults (Carstensen, 2006).

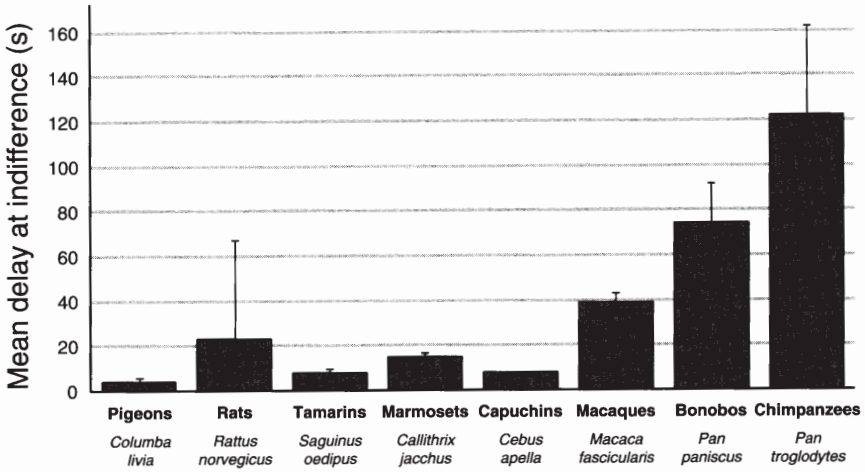

Figure 13.6. Comparison of animal species tested in self-control paradigm. Researchers have tested pigeons, rats, tamarins, marmosets, capuchin monkeys, macaques, bonobos, and chimpanzees in the self-control paradigm. In this comparison, individuals chose between two immediate rewards and six delayed rewards. In most cases, the bars represent the mean delay to receiving the large reward at which subjects were indifferent between the smaller–sooner and larger–later reward. Error bars represent the range of minimum and maximum indifference points for subjects. The capuchin data represent an interpolation of approximate indifference points with twofold and fourfold differences in food amounts. Data sources: pigeons and rats, Green, Myerson, Holt, Slevin, and Estle (2004); tamarins and marmosets, Stevens et al. (2005); capuchins, Ramseyer, Pele, Dufour, Chauvin, and Thierry (2006); macaques, Tobin, Logue, Chelonis, and Ackerman (1996); and bonobos and chimpanzees, Rosati et al. (2007).

Competition

For most species, competition with groupmates or other species poses an enormous interruption risk. Waiting to obtain a food item or mating opportunity gives others the chance to grab it in the meantime. For instance, bypassing a smaller, closer food item in favor of a larger, more distant one means that a social forager may well end up with nothing because groupmates have arrived first. We would expect, therefore, that species living in larger, competitive groups should exhibit stronger preferences for immediacy than solitary species. At the moment, we have not tested enough species to evaluate this comparative hypothesis.

At the individual level, however, we have clear evidence of competitive interference effects on intertemporal choice. In parallel with the wasp example mentioned previously, some species are quite sensitive to cues associated with social interruptions, especially caching species. *Caching* (or storing) food offers one of the most striking instances of delayed rewards for

nonhuman animals because caching animals actively choose to delay consumption, often for months. This long-term storage makes cached food vulnerable to pilferage from competitors (Vander Wall, 1990). Both natural observations and laboratory experiments have shown that individuals adaptively respond to this social risk by eating rather than caching when in the presence of potential pilferers (Carrascal & Moreno, 1993; Emery, Dally, & Clayton, 2004). The threat of competitive interruption triggers a preference for immediate payoffs.

Opportunity Costs

In addition to increasing collection risk, waiting imposes opportunity costs because an animal must wait to put delayed rewards to use. In one of the first attempts to address species differences in intertemporal preference, Tobin and Logue (1994) argued that species with high metabolic rates simply cannot afford to wait long delays to receive food and will prefer immediate rewards more than will species with lower metabolic rates. Thus, for animals with high metabolic rates, waiting imposes high investment opportunity costs because they can put smaller–sooner rewards to use sooner to meet their high metabolic demands. Tobin and Logue used an analysis of previous experiments on pigeons, rats, and humans to support this hypothesis. Given their much smaller body size, pigeons and rats have higher specific metabolic rates than humans, and metabolic rates correlated with impulsive choice in self-control procedures. Although consistent with the differential metabolic rate hypothesis, these data do not demonstrate a definitive link between metabolic rate and intertemporal choice because these three species vary in many ways other than their metabolic rates. Nevertheless, metabolic costs—and investment opportunity costs more generally—may play a key role in determining temporal preferences, and an analysis with a wider range of species would certainly be illuminating.

Ecological Rationality of Intertemporal Choice

Todd and Gigerenzer (2000) defined ecological rationality as "adaptive behavior resulting from the fit between the mind's mechanisms and the structure of the environment in which it operates" (p. 728). That is, the appropriateness of a decision mechanism depends on the decision environment. In laboratory settings, ecological rationality may result in a mismatch between the domain of selection and the domain of testing. The Stephens and Anderson (2001) result clearly exemplifies this. According this hypothesis, natural selection has favored mechanisms that perform well in solving patch-exploitation problems (the domain of selection), yet these same mechanisms lead to errors in conventional laboratory tests of choice (the domain of testing).

The key point analytically is that identifying and understanding the domain of selection is critical to understanding the behavior revealed in a range of testing domains.

Clearly, this mismatch problem applies quite broadly, given that testing situations almost never match the domain of selection. For example, Stevens et al. (2005) tested temporal preferences in two closely related species of monkeys, cotton-top tamarins and common marmosets. Although similar in many ways, these two species differ in one crucial aspect: their diets. Tamarins eat primarily insects—which require quick action to capture—whereas marmosets feed primarily on sap and gum exuding from trees (Snowdon & Soini, 1988; Stevenson & Rylands, 1988). Marmosets chew on tree bark and wait for seconds, minutes, or hours for the sap to flow. Thus, the marmosets are adapted to waiting for food. In a self-control paradigm, marmosets waited almost twice as long as tamarins for the same amount of food (Stevens et al., 2005). Recall that the tamarins were one of the species whose preferences matched the predictions of the short-term rate model. Marmosets, therefore, exhibited more patience than this model predicts. The bonobo–chimpanzee difference we discussed previously also shows this pattern. The bonobo preferences matched predictions of the rate model (this time, the long-term rate model), whereas the chimpanzees were more patient than expected (Rosati et al., 2007). In both of these cases, the species that deals with longer delays in the natural foraging environment also waits longer than expected by a rate model. These may offer cases of the decision mechanisms being tuned to the delayed nature of their natural foraging environment. Under the artificial conditions of the self-control paradigm, the marmosets and chimpanzees seem overly patient when using ecologically rational decision mechanisms.

Uniquely Human Patience?

A glance at Figure 13.6 shows something striking. The animals tested so far in the self-control situation do not wait more than a few seconds or minutes for a threefold increase in food amount. Although we offer good explanations for why this is the case, these preferences, nevertheless, pale in comparison to the temporal preferences documented in humans. Economists and psychologists have repeatedly shown that humans are willing to wait months or years for delayed payoffs (reviewed in Frederick, Loewenstein, & O'Donoghue, 2002), leading some to propose patience as a unique capacity in humans (e.g., McClure, Laibson, Loewenstein, & Cohen, 2004).

Readers should evaluate these claims cautiously, however, because the methodologies used in human and nonhuman testing often differ dramatically. These different decision-making environments could tap different decision mechanisms, resulting in divergent but ecologically rational preferences. As we have already mentioned, in self-control experiments animals always

experience repeated choices between food rewards with no other way to acquire food and few, if any, alternative activities available (for an alternative paradigm, see Beran, Savage-Rumbaugh, Pate, & Rumbaugh, 1999; Evans & Beran, 2007). In contrast, investigators usually ask humans about their temporal preferences for money (but for primary rewards such as food and juice, see Lagorio & Madden, 2005; McClure, Ericson, Laibson, Loewenstein, & Cohen, 2007). Often, both the monetary rewards and the delay periods are hypothetical, but even if real money is offered, participants can obviously leave the experiment and go about their daily lives while waiting for the payoffs. Thus, important differences exist in the reward types, number of exposures to options, sampling methods required (descriptions vs. experiences with options), and costs associated with waiting between human and nonhuman experiments.

Rosati et al. (2007) explored whether humans do indeed exhibit more patience than other animals by testing humans and chimpanzees in a comparable self-control experiment. Both species experienced repeated choices between a smaller food reward available immediately and a larger food reward available after 2 minutes. Humans rarely waited for the large rewards (even less than the chimpanzees did), demonstrating that (a) humans are not always patient and (b) comparing the existing human and animal data is not a valid comparison. When tested with similar rewards using similar procedures, humans look much more like other animals, showing a strong preference for immediate rewards. Thus, the intertemporal choice decision mechanisms are tuned to specific decision environments, supporting the ecological rationality hypothesis.

SUMMARY

The puzzle of impulsivity has been documented repeatedly in self-control experiments in pigeons, blue jays, and starlings. In each of these species, individuals prefer a smaller–sooner reward more often than expected by a long-term rate model. Much like the paradox of altruism, we believe that impulsivity is in the eye of the beholder—there are a number of reasonable explanations for it. The long-term rate-based approach of classic foraging theory fails to account for the data in this situation, although short-term models do quite well. The alternative, delay-discounting approach nicely captures quantitative aspects of the data, but offers no explanation of observed preferences, only a description. When viewed through the lens of an evolutionary approach, a preference for immediate rewards appears not impulsive but adaptive in a naturally occurring behavioral situation. A decision mechanism adapted to a common foraging problem may not work as well in an artificial situation contrived in the laboratory. This is a specific example of the more general phenomenon of ecological rationality—the adaptive match between decision mechanisms and the decision environment. We argue that this ecological rationality approach can be very informative to the study of impulsivity and to the study of intertemporal choice more broadly. In fact, the ecological rationality approach is broad enough to include both the rate and the delay-discounting approaches. It offers predictive models of intertemporal choice and emphasizes the general nature of trading off time delays and reward amounts, an important and ubiquitous class of decisions that all organisms face.

REFERENCES

Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin, 82,* 463–496.

Ainslie, G., & Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning & Behavior, 9,* 476–482.

Andersson, M. (1982, October 28). Female choice selects for extreme tail length in a widowbird. *Nature, 299,* 818–820.

Andersson, M. (1994). *Sexual selection.* Princeton, NJ: Princeton University Press.

Bateson, M., & Kacelnik, A. (1996). Rate currencies and the foraging starling: The fallacy of the averages revisited. *Behavioral Ecology, 7,* 341–352.

Beaugrand, J. P. (1997). Resolution of agonistic conflicts in dyads of acquainted green swordtails (*Xiphophorus helleri*): A game with perfect information. *Behavioural Processes, 41,* 79–96.

Beran, M. J., Savage-Rumbaugh, E., Pate, J. L., & Rumbaugh, D. M. (1999). Delay of gratification in chimpanzees (*Pan troglodytes*). *Developmental Psychobiology, 34,* 119–127.

Carrascal, L. M., & Moreno, E. (1993). Food caching versus immediate consumption in the nuthatch: The effect of social context. *Ardea, 81,* 135–141.

Carstensen, L. L. (2006, June 30). The influence of a sense of time on human development. *Science, 312,* 1913–1915.

Daly, M., & Wilson, M. (2005). Carpe diem: Adaptation and devaluing the future. *Quarterly Review of Biology, 80,* 55–60.

Emery, N. J., Dally, J. M., & Clayton, N. S. (2004). Western scrub-jays (*Aphelocoma californica*) use cognitive strategies to protect their caches from thieving conspecifics. *Animal Cognition, 7,* 37–43.

Evans, T. A., & Beran, M. J. (2007). Chimpanzees use self-distraction to cope with impulsivity. *Biology Letters, 3,* 599–602.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40,* 351–401.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84,* 279–325.

Gibbon, J., Church, R. M., Fairhurst, S., & Kacelnik, A. (1988). Scalar expectancy theory and choice between delayed rewards. *Psychological Review, 95,* 102–114.

Giraldeau, L.-A., & Caraco, T. (2000). *Social foraging theory*. Princeton, NJ: Princeton University Press.

Green, L., Fisher, E. B., Perlow, S., & Sherman, L. (1981). Preference reversal and self-control: Choice as a function of reward amount and delay. *Behaviour Analysis Letters, 1,* 43–51.

Green, L., Myerson, J., Holt, D. D., Slevin, J. R., & Estle, S. J. (2004). Discounting of delayed food rewards in pigeons and rats: Is there a magnitude effect? *Journal of the Experimental Analysis of Behavior, 81,* 39–50.

Henly, S., Ostdiek, A., Blackwell, E., Knutie, S., Dunlap, A. S., & Stephens, D. W. (2008). The discounting-by-interruptions hypothesis: Model and experiment. *Behavioral Ecology, 19,* 154–162.

Kacelnik, A. (2003). The evolution of patience. In G. Loewenstein, D. Read, & R. F. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 115–138). New York: Russell Sage Foundation.

King, G. R., & Logue, A. W. (1992). Choice in a self-control paradigm: Effects of uncertainty. *Behavioural Processes, 26,* 143–153.

Lagorio, C. H., & Madden, G. J. (2005). Delay discounting of real and hypothetical rewards III: Steady-state assessments, forced-choice trials, and all real rewards. *Behavioural Processes, 69,* 173–187.

Lima, S. L., Valone, T. J., & Caraco, T. (1985). Foraging efficiency–predation risk trade-off in the grey squirrel. *Animal Behaviour, 33,* 155–165.

Malenky, R. K., & Wrangham, R. W. (1994). A quantitative comparison of terrestrial herbaceous food consumption by *Pan paniscus* in the Lomako Forest, Zaire, and *Pan troglodytes* in the Kibale Forest, Uganda. *American Journal of Primatology, 32,* 1–12.

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* (Vol. 5, pp. 55–73). Hillsdale, NJ: Erlbaum.

Mazur, J. E. (1994). Effects of intertrial reinforcers on self-control choice. *Journal of the Experimental Analysis of Behavior, 61,* 83–96.

Mazur, J. E., & Romano, A. (1992). Choice with delayed and probabilistic reinforcers: Effects of variability, time between trials, and conditioned reinforcers. *Journal of the Experimental Analysis of Behavior, 58,* 513–525.

McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience, 27,* 5796–5804.

McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004, October 15). Separate neural systems value immediate and delayed monetary rewards. *Science, 306,* 503–507.

McDiarmid, C. G., & Rilling, M. E. (1965). Reinforcement delay and reinforcement rate as determinants of schedule preference. *Psychonomic Science, 2,* 195–196.

Parker, G. A. (1974). Assessment strategy and the evolution of fighting behaviour. *Journal of Theoretical Biology, 47,* 223–243.

Parker, G. A., & Rubenstein, D. I. (1981). Role assessment, reserve strategy, and acquisition of information in asymmetric animal conflicts. *Animal Behaviour, 29,* 221–240.

Ramseyer, A., Pele, M., Dufour, V., Chauvin, C., & Thierry, B. (2006). Accepting loss: The temporal limits of reciprocity in brown capuchin monkeys. *Proceedings of the Royal Society of London, Series B, 273,* 179–184.

Roitberg, B. D., Mangel, M., Lalonde, R. G., Roitberg, C. A., van Alphen, J. J. M., & Vet, L. (1992). Seasonal dynamic shifts in patch exploitation by parasitic wasps. *Behavioral Ecology, 3,* 156–165.

Rosati, A. G., Stevens, J. R., Hare, B., & Hauser, M. D. (2007). The evolutionary origins of human patience: Temporal preferences in chimpanzees, bonobos, and adult humans. *Current Biology, 17,* 1663–1668.

Snowdon, C. T., & Soini, P. (1988). The tamarins, genus *Saguinus*. In R. A. Mittermeier, A. B. Rylands, A. F. Coimbra-Filho, & G. A. B. Fonseca (Eds.), *Ecology and behavior of neotropical primates* (Vol. 2, pp. 223–298). Washington, DC: World Wildlife Fund.

Sozou, P. D., & Seymour, R. M. (2003). Augmented discounting: Interaction between ageing and time-preference behaviour. *Proceedings of the Royal Society of London, Series B, 270,* 147–153.

Stephens, D. W. (2002). Discrimination, discounting and impulsivity: A role for an informational constraint. *Philosophical Transactions of the Royal Society of London, Series B, 357,* 1527–1537.

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. (2007). *Foraging: Behavior and ecology*. Chicago: University of Chicago Press.

Stephens, D. W., Kerr, B., & Fernandez-Juricic, E. (2004). Impulsiveness without discounting: The ecological rationality hypothesis. *Proceedings of the Royal Society of London, Series B, 271,* 2459–2465.

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: 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.

Stevens, J. R., Hallinan, E. V., & Hauser, M. D. (2005). The ecology and evolution of patience in two New World monkeys. *Biology Letters, 1,* 223–226.

Stevenson, M. F., & Rylands, A. B. (1988). The marmosets, genus *Callithrix*. In R. A. Mittermeier, A. B. Rylands, A. F. Coimbra-Filho, & G. A. B. Fonseca (Eds.), *Ecology and behavior of neotropical primates* (Vol. 2, pp. 131–222). Washington, DC: World Wildlife Fund.

Tobin, H., & Logue, A. W. (1994). Self-control across species (*Columba livia*, *Homo sapiens*, and *Rattus norvegicus*). *Journal of Comparative Psychology*, *108*, 126–133.

Tobin, H., Logue, A. W., Chelonis, J. J., & Ackerman, K. T. (1996). Self-control in the monkey *Macaca fascicularis*. *Animal Learning & Behavior*, *24*, 168–174.

Todd, P. M., & Gigerenzer, G. (2000). Précis of *Simple heuristics that make us smart*. *Behavioral and Brain Sciences*, *23*, 727–741.

Vander Wall, S. B. (1990). *Food hoarding in animals*. Chicago: University of Chicago Press.