

When rate maximization is impulsive

Theodore P. Pavlic · Kevin M. Passino

Received: 3 September 2009 / Revised: 17 February 2010 / Accepted: 19 February 2010 / Published online: 11 March 2010
© Springer-Verlag 2010

Abstract Although optimal foraging theory predicts that natural selection should favor animal behaviors that maximize long-term rate of gain, behaviors observed in the laboratory tend to be impulsive. In binary-choice experiments, despite the long-term gain of each alternative, animals favor short handling times. Most explanations of this behavior suggest that there is hidden rationality in impulsiveness. Instead, we suggest that simultaneous and mutually exclusive binary-choice encounters are often unnatural and thus immune to the effects of natural selection. Using a simulation of an imperfect forager, we show how a simple strategy (i.e., an intuitive model of animal behavior) that maximizes long-term rate of gain under natural conditions appears to be impulsive under operant laboratory conditions. We then show how the accuracy of this model can be verified in the laboratory by biasing subjects with a

short pre-experiment ad libitum high-quality feeding period. We also show a similar behavioral mechanism results in diet preferences that are qualitatively consistent with the digestive rate model of foraging (i.e., foraging under digestive rate constraints).

Keywords Impulsiveness · Impulsivity · Rationality · Self-control · Optimal foraging · Simultaneous encounters

Introduction

As discussed by Schoener (1971) and Pyke et al. (1977), an animal must make choices that put upward pressure on both the number of prey encountered in its finite lifetime and the gain (e.g., calories or another surrogate for Darwinian fitness) returned from each prey. To attain maximum gain in minimum time, optimal foraging theory (OFT) predicts that natural selection will favor behaviors that maximize the long-term rate of gain (Stephens and Krebs 1986). However, in operant binary-choice experiments in the laboratory (e.g., Ainslie 1974; Bateson and Kacelnik 1996; Bradshaw and Szabadi 1992; Green et al. 1981; McDiarmid and Rilling 1965; Rachlin and Green 1972; Stephens and Anderson 2001; Snyderman 1983; Siegel and Rachlin 1995), animals will often make impulsive choices that favor short handling times regardless of foraging gain. Conversely, as reviewed by Giraldeau and Caraco (2000, pp. 155–167), in experiments that do not force animals to make binary-choice decisions, animals not only maximize their rate of gain but also respond to changes in the environment by dynamically adjusting their behavior to maintain maximal long-term rate of

Communicated by J. Lindström

Electronic supplementary material The online version of this article (doi:10.1007/s00265-010-0940-1) contains supplementary material, which is available to authorized users.

T. P. Pavlic (✉) · K. M. Passino
Department of Electrical and Computer Engineering,
Ohio State University, 2015 Neil Avenue,
Columbus, OH 43210, USA
e-mail: pavlic.3@osu.edu

K. M. Passino
Department of Evolution, Ecology, and Organismal
Biology, Ohio State University, 318 W. 12th Avenue,
Columbus, OH 43210, USA
e-mail: passino@ece.osu.edu

gain. That is, natural selection has gone so far as to bestow on animals the ability to calculate rate-maximizing behaviors in real time. In this paper, we propose a simple behavioral strategy consistent with real-time rate maximization in nature and show how it appears to be impulsive in laboratory binary-choice experiments. Additionally, we show how rate maximization in the laboratory can be restored with appropriate pre-experiment treatment. If our hypothesis is shown to be empirically robust, then impulsiveness may be a behavioral nuance that is immune to the effects of natural selection.

Impulsiveness without discounting

As reviewed by Ainslie (1975), early theories of impulsiveness are based on the assumption of temporal discounting. Animals are assumed to discount future rewards by some decreasing function of the time until the reward. As long as the discounting function is sufficiently concave, a smaller-sooner reward can have greater value than a larger-later reward. However, Stephens (2002) argues that realistic discount rates will be too shallow to impact laboratory experiments. In particular, animals with relatively long lifetimes may value rewards today more than rewards tomorrow, but a difference in delay of less than a minute should not be enough to cause a preference reversal. Additionally, Henly et al. (2008) show that likelihood of interruptions, which is often used to justify the discounting hypothesis, has negligible impact on impulsiveness observed in the laboratory.

To explain impulsiveness without discounting, recent attention focuses on how laboratory methods may artificially bias subjects toward impulsive behaviors. Arguing that impulsiveness is the result of an informational constraint, Stephens and Anderson (2001) and Stephens et al. (2004) show how a simple rule leads to impulsive decisions in the typical binary-choice schedule and rate-maximizing decisions in a sequential-choice schedule, but evidence from Stephens and McLinn (2003) suggests that animals do not use the same choice rule in both contexts. In a different experiment, pigeons whose attention is trained on a larger-later option lose their impulsive tendencies (Siegel and Rachlin 1995), which leads Monterosso and Ainslie (1999) to suggest that impulsiveness may be the result of attention and not deliberate choice. To investigate state-dependent effects, Houston and McNamara (1985) use a dynamic programming model to show that an impulsive strategy minimizes probability of starvation when an animal is in a low-energy state. This result agrees with studies (e.g., Bradshaw and Szabadi 1992; Snyderman 1983) that show that animals

are more likely to prefer smaller-sooner alternatives when they are subjected to greater food deprivation. Because food deprivation accompanies conventional operant methodology, this result may explain all observed laboratory impulsiveness; however, the model of Houston and McNamara is based on a sequential-choice assumption that is rarely met in the mutually exclusive binary-choice experiments used in most tests of impulsiveness. Here, we describe a simple behavioral mechanism that leads to impulsiveness under typical operant conditioning and rate maximization otherwise. As this single mechanism is highly influenced by both the subject's attention and energetic state, it synthesizes ideas from the three different approaches to explaining impulsiveness without discounting.

A graph of the prey model

A detailed discussion of classical OFT is given by Stephens and Krebs (1986). Here, we summarize a central OFT result and recreate a useful graphical interpretation first presented by Charnov (1976). Consider a forager that searches through its environment for $n \in \mathbb{N}$ types of prey. Encounters with prey of type $i \in \{1, 2, \dots, n\}$ come from an independent Poisson process with rate $\lambda_i > 0$ and those type- i prey that are chosen for processing return an average g_i gain (e.g., calories) to the forager after an average h_i handling time. When not handling prey, the forager pays a cost c per unit search time (i.e., c is a cost rate). Assuming that natural selection favors foraging behaviors that maximize lifetime gain, the optimal behavior should trade large reward per encounter for increased number of encounters (i.e., less time per encounter), and so an optimal behavior maximizes the long-term average rate

$$R \stackrel{\text{def}}{=} \frac{-c + \sum_{i=1}^n \lambda_i p_i g_i}{1 + \sum_{i=1}^n \lambda_i p_i h_i} \quad (1)$$

where $p_i \in [0, 1]$ represents the probability that the forager processes an encounter with a prey of type i . Each behavior is represented by a particular collection (p_1, p_2, \dots, p_n) representing the diet of the forager.

The central result of this so-called contingency model (Pulliam 1974; Charnov 1976) or prey model (Stephens and Krebs 1986) is that if prey types are ordered so that

$$\frac{g_1}{h_1} > \frac{g_2}{h_2} > \dots > \frac{g_{n-1}}{h_{n-1}} > \frac{g_n}{h_n},$$

then an optimal behavior is to process all encounters with prey of type $m \leq k^*$ and to ignore all encounters

with prey of type $m > k^*$ where $k^* \in \{0, 1, 2, \dots, n\}$ is the smallest k that satisfies

$$\frac{-c + \sum_{i=1}^k \lambda_i g_i}{1 + \sum_{i=1}^k \lambda_i h_i} > \frac{g_{k+1}}{h_{k+1}}$$

Consequently,

$$\underbrace{\frac{g_1}{h_1} > \frac{g_2}{h_2} > \dots > \frac{g_{k^*}}{h_{k^*}}}_{\text{Processed types}} > \underbrace{\left(R^* \stackrel{\text{Optimal rate}}{\text{def}} \frac{-c + \sum_{i=1}^{k^*} \lambda_i g_i}{1 + \sum_{i=1}^{k^*} \lambda_i h_i} \right)}_{\text{Optimal rate}} > \underbrace{\frac{g_{k^*+1}}{h_{k^*+1}} > \dots > \frac{g_{n-1}}{h_{n-1}} > \frac{g_n}{h_n}}_{\text{Ignored types}} \quad (2)$$

where the optimal rate R^* partitions the list of profitabilities (i.e., the ratio g_i/h_i for each type i) into an acceptable set and an unacceptable set. This result, which is known as the zero–one rule (Stephens and Krebs 1986), is shown graphically in Fig. 1. For each

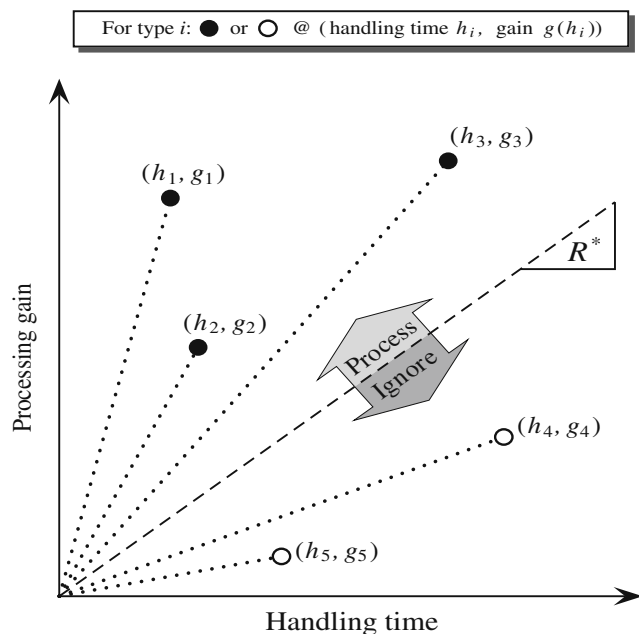


Fig. 1 Graphical summary of prey model result. For a prey type $i \in \{1, 2, 3, 4, 5\}$, the average handling time h_i and average gain g_i is plotted as a dot. The maximum long-term rate of gain R^* is the slope of the dashed line which separates the processed types, 1, 2, and 3, from the ignored types, 4 and 5. The profitability of each type is the slope of the dotted line connecting the origin to its (gain, time) coordinate

type, a solid dot is plotted with the type’s average handling time as its abscissa and the type’s average processing gain as its ordinate, and so the slope of each dotted line corresponds to the profitability of the corresponding type. If the optimal long-term rate of gain is R^* , then the dashed line with slope R^* divides the plane into an acceptable upper region and an unacceptable lower region. In this example, the three steep lines corresponding to types 1, 2, and 3 are contained in the acceptable region, and the shallow lines corresponding to types 4 and 5 are contained in the unacceptable region and so $k^* = 3$ in Eq. 2. As encounter rates of acceptable types increase or decrease, the dashed line rotates to exclude or include more types.

Justification for adaptive rate-maximization model

Optimal foraging theory describes the ultimate behaviors favored by natural selection (Stephens and Krebs 1986). In the example in Fig. 1, a heritable preference for only the three highest-profitability types will readily spread through future generations of the population. If an environmental disaster causes encounter rates to fall sharply, then more-inclusive heritable preferences present in the background population will successfully invade and dominate future generations of the population. So optimal foraging theory posits that rate-maximizing behaviors will be the adaptive outcome of the gradual process of natural selection. However, Giraldeau and Livoreil (1998) show that birds in the laboratory respond to a sharp change in the environment with a new behavior that matches the long-term rate-maximizing behavior predicted by OFT. This response is consistent with real-time calculation of the optimal rate-maximizing behavior. Moreover, in a survey by Sih and Christensen (2001) of 74 recent foraging studies, 22 studies are given that show animals flexibly adapting to their experimental scenarios in a way that is at least qualitatively consistent with optimal foraging theory. Thus, OFT may also be viewed as describing the proximate outcome of dynamic behaviors that adapt to changing environments.

Other ostensible violations of rate maximization

Impulsiveness describes only a subset of behaviors that appear to violate rate maximization. In the recent survey by Sih and Christensen (2001) of 60 foraging studies originally surveyed by Stephens and Krebs (1986) as well as 74 more recent studies, optimal foraging theory is shown to have mixed success in describing animal behavior. Sih and Christensen argue that the prey model is a poor descriptor of the foraging behaviors of

animals with mobile prey. However, even in herbivores (Hirakawa 1997a; Verlinden and Wiley 1989) and molluscivores (van Gils et al. 2005; Quaintenne et al. 2010) that have essentially immobile prey, optimal foraging theory fails to predict the observed diet preferences. In these cases, the digestive rate model (DRM; Verlinden and Wiley 1989; Hirakawa 1995, 1997b) is shown to better predict animal diet preference, but the DRM itself is the result of generalizing the prey model (i.e., rate maximization) to include the effect of digestive constraints.

The DRM of Hirakawa (1995) is a correction of related foraging constraint models from Pulliam (1974), Stephens and Krebs (1986), and Verlinden and Wiley (1989). It maximizes the same rate expression from Eq. 1; however, it also introduces the material intake rate

$$X \stackrel{\text{def}}{=} \frac{\sum_{i=1}^n \lambda_i p_i k_i}{1 + \sum_{i=1}^n \lambda_i p_i h_i} \quad (3)$$

where k_i is the average quantity (e.g., mass) of ingested material from prey of type i . Depending on whether the material is a nutrient or a toxin, the intake rate X will have a lower or upper bound, respectively. In the molluscivore shorebird examples described by van Gils et al. (2005) and Quaintenne et al. (2010), X represents the intake rate of ballast material (e.g., shells of bivalves and snails). When the gizzard of one of these cropping birds is full, the bird must pause its foraging behavior until its gizzard has emptied enough for it to continue. Thus, these shorebirds have a maximum intake rate of ballast material that is not negligible. A key result of the DRM is that prey types should not be ordered by profitability as in Eq. 2; instead, they should be ordered by so-called (digestive) quality which is defined by g_i/k_i for each type i . The DRM gives an algorithm to find the prey type whose quality partitions the group into take-always and take-never groups. However, only a fraction of encounters with prey of the partitioning type should be ingested. Hence, even though the DRM is the result of rate maximization under a constraint, its result both re-orders prey preferences and appears to violate the zero-one rule. Although our main focus is to present a simple model of foraging that explains impulsiveness in the laboratory, we shall also show how our simplified model has similar qualitative features as the DRM. In particular, our DRM-consistent behavioral model reconciles the differences between rate maximization and DRM using an internal handling time approach, which is similar to the ecological-physiological hybrid method described by Whelan and Brown (2005).

Model

State-based real-time adaptive rate-maximization model

Using a version of the prey model of Stephens and Krebs (1986) put into an engineering context, Andrews et al. (2004) and Quijano et al. (2006) implement a heuristic algorithm that achieves real-time rate maximization through computation of an estimated optimal behavior. In particular, on encounter with a prey of type i , the algorithm estimates the encounter rate λ_i and then updates its exclusion rule based on Eq. 2 assuming that the present λ_i estimate is correct. That is, the algorithm calculates the optimal rate R^* based on the estimated rates and then updates its prey model exclusion policy based on that estimated optimal rate.

Instead of estimating each of the n encounter rates to find the critical k^* , our algorithm keeps a running estimate of the long-term rate of gain R by dividing total accumulated gain by total time. Then, on each encounter, the forager compares the encountered prey's profitability with the current R estimate and rejects the prey if its profitability is lower than R . By accepting prey with super- R profitability and rejecting those with sub- R profitability, processed prey only cause the estimated R to rise. Between processed encounters, the R estimate falls due to the effect of increased time without increased gain. Over time, the oscillating R estimate converges on the optimal R^* ; hence, the process-ignore policy converges to the prey model described by Eq. 2. This method is essentially an application of the patch model of Stephens and Krebs (1986) applied to a forager whose processing-length decision is limited to leaving immediately or staying for a patch's entirety. Not only does this strategy have less computational complexity than those based on encounter-rate estimation, but it has an intuitive state-based interpretation; when the forager's present energy stores are low (i.e., when R is low due to low accumulated gain or high total running time), the forager becomes more inclusive. We caution that this is a loaded interpretation. Among other things, it neglects the effects of non-foraging activities (e.g., reproduction) that both depress present energy stores and take time. However, although a forager's instantaneous energy stores may be a poor estimator of R , low-frequency trends in the energy store signal should correlate well with R . For example, it may be maladaptive for momentary hunger to drive foraging decisions, but its persistent hunger is an important signal that the forager should generalize more. Furthermore, we show later that a small adjustment to this model accounts for optimal

foraging under digestive rate constraints. In that case, a stronger connection exists between energy state and R (e.g., high R may indicate that less bulky foods should be preferred).

Our algorithm is depicted in Fig. 2a. The solid line represents the trajectory of the gain of the forager over time. Each dotted line has a slope matching one of the n profitabilities. The estimated rate R at each encounter is the slope of an imaginary line drawn from the origin to the gain trajectory at that time. If the profitability of an encountered prey is steeper than the present R estimate, then the prey is processed, and the current R estimate increases; otherwise, the prey is ignored. Encounters are shown in this depiction as points on the gain trajectory. Those encounters that are processed are shown as circles, and encounters that are ignored are shown as diamonds. So an encounter is processed if its graph location, which corresponds to the estimated R at the time of the encounter, is below the dotted profitability line of its task type. For this example trajectory, the first four encounters are chosen for processing because they fall below their corresponding profitability lines. The fifth encounter is with a prey of type 3, and it is ignored because it falls above the third profitability line. Even though the sixth encounter is also with a prey of type 3, it is chosen for

processing because enough time has passed to depress the estimated rate R . As time continues, the estimated rate R eventually converges to the maximum long-term rate of gain R^* , which is shown as the slope of the dashed line.

State-based real-time adaptive model consistent with DRM

With only a small modification, the adaptive rate-maximization model can be made to be consistent with the qualitative predictions of the DRM. In particular, for each type i , let ballast time κ_i be a parameter with units of time that is proportional to the quantity k_i . For example, κ_i may be viewed as the average amount of time prey of type i spend in the gizzard of a molluscivore shorebird. If the profitability g_i/h_i for each type i used by the forager to determine whether a particular encounter should be accepted or rejected is replaced by the $g_i/(h_i + \kappa_i)$ or simply g_i/κ_i , then the gain trajectory will follow (i.e., continually switch across) one of the profitability lines. The prey type corresponding to the switching line will also represent the prey that is only partially preferred. The other prey types will be preferred via the zero–one rule. This modified process is depicted in Fig. 2b. Hence, the introduction of internal

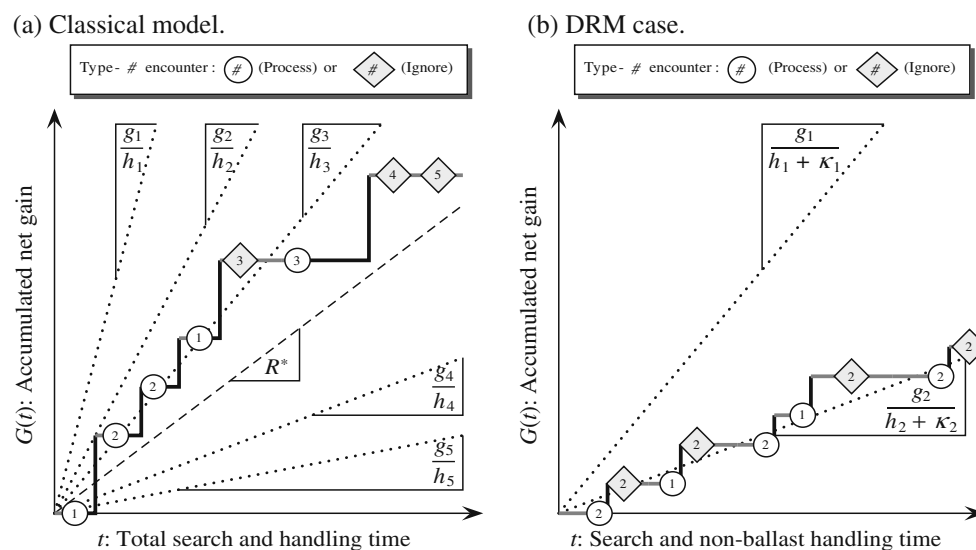


Fig. 2 Adaptive long-term maximizer trajectory. The *solid line* represents a forager's accumulated gain over time. Each encounter with a new prey is shown with as a *circle* or a *diamond* surrounding the prey's type. Encounters shown with a *circle* are chosen for processing and encounters shown with a *diamond* are ignored. Hence, jumps in the gain trajectory occur at circled encounters. Encounters that occur below the prey type's profitability slope are processed; otherwise, encounters are

ignored. In **a**, profitability is defined in the classical way as g_i/h_i for each type i , and the gain trajectory converges on a line with slope R^* , which is the maximum long-term rate of gain. In **b**, profitability is defined as $g_i/(h_i + \kappa_i)$ (or, alternatively, g_i/κ_i with similar results) where κ_i has units of time and is proportional to k_i , and the gain trajectory converges on the profitability line of the partially preferred prey type

handling time κ_i mitigates the need for a hard digestive rate constraint, which is similar to the ecological–physiological hybrid model described by Whelan and Brown (2005) that also models internal handling time. Coopting their example, we notice that in our model, as R increases (e.g., stomach fills), calorie-dense foods (e.g., cookies) may be preferable over high-bulk foods (e.g., salad) even though profitabilities may be similar, which matches conventional human experience.

Although this algorithm may not accurately describe the proximate mechanisms for decision making in foragers with digestive constraints, it may provide helpful intuition for visualizing DRM results. Additionally, if the ballast-time proportionality constant (i.e., κ_i/k_i for each type i) is identical across all types and digestive profitability is defined as $g_i/(h_i + \kappa_i)$ for each type i , then this model behaves as a classical rate maximizer when h_i is much larger than κ_i and as a digestively constrained forager when κ_i is much larger than h_i . Hence, the previous model is a limiting case of this model; as ballast times become vanishingly small, the acceptance proportion of the partially accepted prey type approaches unity or nullity. Thus, this simple behavioral model reconciles differences between classical optimal foraging theory models and digestive constraint models. Moreover, this model justifies the otherwise questionable use of energy availability (e.g., hunger) as an estimator of R . In the case of pure rate maximization, an animal should only use its energy state as an estimator of a low R if its energy stores are persistently low. However, when using digestive profitability (i.e., when including ballast effects), using energy stores as an estimator of R implicitly accounts for digestive constraints; less bulky foods should be preferred when R is high.

Results

The most prevalent foraging models used in OFT make the simplifying assumption that encounters with more than one prey at a time occur with zero probability (Charnov 1973). This assumption does not lack realism. A forager that searches for sparse stationary morsels of food may rarely encounter items simultaneously, and even when it does, its choices need not be mutually exclusive. If such an environment is the source of natural selection, then it makes sense that the animal's behavior may be poorly adapted for the mutually exclusive, simultaneous encounter, binary-choice experiments of the conventional operant laboratory. Here, we show how our real-time algorithm is impulsive under traditional laboratory conditions and maximizes

long-term rate of gain under natural conditions. We also show how our DRM-inspired real-time algorithm results in foraging preferences consistent with results from the DRM. All simulations were implemented using MATLAB ((R14SP3) The MathWorks, Natick, MA, USA), and all MATLAB source code is available as an electronic supplement (see S1 Appendix).

Simulation: simultaneous encounters lead to suboptimality

A simulated trajectory for our algorithm over 100-s is shown in Fig. 3a where $n = 2$, $(\lambda_1, g_1, h_1) = (2\text{-s}^{-1}, 5, 2\text{-s})$, $(\lambda_2, g_2, h_2) = (5\text{-s}^{-1}, 1, 1\text{-s})$, and $c = 0.1\text{-s}^{-1}$. So type 1 has the highest profitability and type 2 has the lowest average handling time. For each type, the gain and handling time have nonzero variance and positive correlation. Additionally, on an encounter with a single prey, there is a 5% probability that the forager will make an algorithmic mistake. In this simulation, encounters with each type come from separate and independent Poisson processes, and so simultaneous encounters occur with zero probability. However, in the rare case of a simultaneous encounter, this forager arbitrarily attends to the prey with the lower average handling time first with a 95% probability. This assumption of biased attention is motivated by the training experiment of Siegel and Rachlin (1995) and the suggestion by Monterosso and Ainslie (1999) that impulsiveness may be the result of attention span and not deliberate choice. In this example, the steady-state foraging behavior accepts all high-profitability encounters and rejects all low-profitability encounters, and the gain trajectory behaves like a line with R^* slope. Because this behavior eventually achieves the maximum long-term rate of gain, it is eligible for proliferation by natural selection.

In Fig. 3b, a simulated trajectory for the same behavior that maximizes its long-term rate in Fig. 3a is shown when facing simultaneous and mutually exclusive encounters at rate $\lambda = 2\text{-s}^{-1}$. Because it attends to the prey with the lower average handling time first and the experimenter removes the other prey on each encounter, the trajectory achieves a suboptimal long-term rate $R_{\text{type-2-only}}^*$, and the experimenter concludes that the forager is not a rate maximizer. However, the rate $R_{\text{type-2-only}}^*$ is the maximum rate in an environment with only the low-profitability type. Under these experimental conditions, the estimated rate is initially very low, and so the low-profitability prey are initially acceptable for processing. Because the forager attends to these prey first and the experimenter immediately removes the other prey, the new estimated rate R

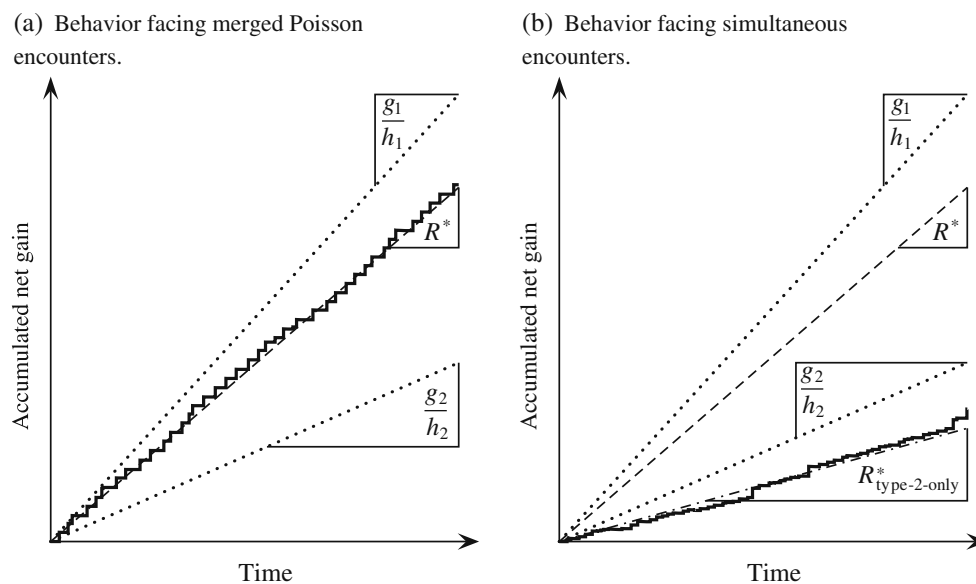


Fig. 3 Simulation of an adaptive and impulsive foraging behavior facing different encounter processes. In the 100-s foraging bout, the forager encounters high-profitability prey of type 1 ($g_1 = 5$, $h_1 = 2$ -s) and low-handling-time prey of type 2 ($g_2 = 1$, $h_2 = 1$ -s). During search, it pays cost $c = 0.1$ -s $^{-1}$. On an encounter with a prey, it is chosen for processing if its profitability is greater than the quotient of the accumulated gain and total lifetime, and the forager makes a mistake in this calculation with 5%

probability. On simultaneous encounters, the forager attends to the prey with lowest handling time first with 95% probability (i.e., its attention is impulsive). In **a**, encounters come from a merged Poisson process with rates $\lambda_1 = 2$ -s $^{-1}$ and $\lambda_2 = 5$ -s $^{-1}$, and the behavior is optimal. In **b**, encounters are simultaneous and mutually exclusive with rate $\lambda = 2$ -s $^{-1}$, and the same behavior is suboptimal. In particular, the performance in **b** is optimal in an environment of only type 2 prey

remains too low to exclude the next encountered low-profitability type. Hence, the forager never experiences high-profitability prey gains and is destined for poor performance. If these operant laboratory conditions (i.e., starvation followed by sequences of mutually exclusive binary-choice trials) are unlike those in nature, this short-time attentive foraging strategy may still maximize its long-term rate of gain in nature and thus will still be present in the animal population.

Simulation: pre-experiment feeding restores optimality

This hypothesis can be tested by preloading the forager's gain trajectory immediately before the experiment (e.g., to alleviate any persistent hunger signals that drive the forager to be maximally inclusive). The trajectory in Fig. 4a is the result of a second simulation with merged Poisson encounters; however, during the first 2-s of the experiment, the forager faces very frequent encounters with prey of a very high profitability g_p/h_p . As a consequence, the trajectory rises very quickly and then plateaus off until the estimated rate of gain falls below the profitability of type 1. After that time, the forager follows a typical rate-maximizing trajectory. Similarly, the trajectory in Fig. 4b is the result of a second simulation with simultaneous and

mutually exclusive encounters, but the forager is given the same initial 2-s feeding period in order to bias its behavior into a rate-maximizing mode. As with the previous case, the gain rises steeply and then plateaus until the point where the estimated gain falls below the profitability of type 1. After that time, the previously impulsive behavior follows a true rate-maximizing trajectory. That is, because the estimated rate is so high, the forager never attends to the low-profitability type; it ceases to appear impulsive. This restoration of rationality will rarely occur under conventional operant methods because those subjects are typically intentionally deprived of food outside of the experimental apparatus (e.g., Bradshaw and Szabadi 1992; Green et al. 1981; Stephens and Anderson 2001; Stephens and McLinn 2003; Snyderman 1983; Siegel and Rachlin 1995).

Simulation: equal-opportunity foragers and simultaneous encounters

Another set of simulated simultaneous-encounter trajectories of our algorithm is shown in Fig. 5. However, in these examples, the forager has unbiased attentiveness. That is, in the case of a simultaneous encounter with two prey, the forager attends first to one of them at random following a uniform distribution.

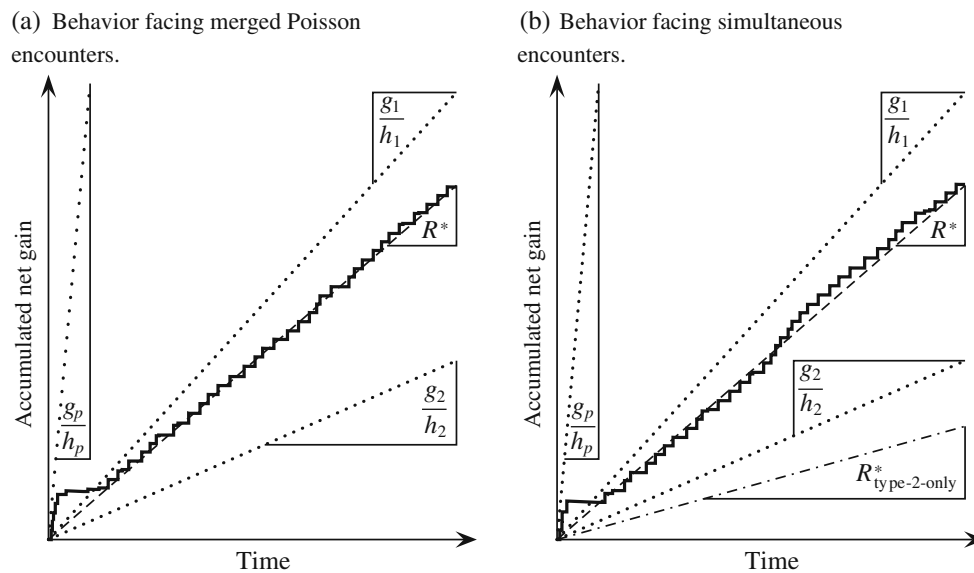


Fig. 4 Simulation of an impulsive foraging behavior after initial ad libitum feeding period. In most of the 100-s foraging bout, the forager encounters high-profitability prey of type 1 ($g_1 = 5$, $h_1 = 2$ -s) and low-handling-time prey of type 2 ($g_2 = 1$, $h_2 = 1$ -s). However, during the first 2-s of the bout, the forager encounters type- p prey of very high profitability and very high rate ($\lambda_p = 10\text{-s}^{-1}$, $g_p = 5$, $h_p = 0.2$). During search, it pays cost $c = 0.1\text{-s}^{-1}$. On an encounter with a prey, it is chosen for processing if its profitability is greater than the quotient of the accumulated

gain and total lifetime, and the forager makes a mistake in this calculation with 5% probability. On simultaneous encounters, the forager attends to the prey with lowest handling time first with 95% probability (i.e., its attention is impulsive). In **a**, encounters come from a merged Poisson process with rates $\lambda_1 = 2\text{-s}^{-1}$ and $\lambda_2 = 5\text{-s}^{-1}$, and in **b**, encounters are simultaneous and mutually exclusive with rate $\lambda = 2\text{-s}^{-1}$. In both cases, the behavior is optimal

Under merged Poisson encounters, this behavior will still maximize the forager's long-term rate of gain because simultaneous encounters are so rare. However, when encounters are artificially made by an experimenter to be simultaneous and mutually exclusive, the foraging behavior is not guaranteed to maximize the long-term rate. In particular, runs can be put into two groups characterized by Fig. 5a, b. In the former case of Fig. 5a, the gain accumulates initially fast enough to exclude the low-profitability type from consideration by the algorithm, and the foraging behavior achieves the optimal long-term rate of gain R^* . In the latter case of Fig. 5b, the equal-opportunity foraging behavior achieves a long-term rate of gain $R^*_{\text{depressed}}$ equal to the optimal long-term rate of gain when encounter rates are halved. Hence, forcing animals to make mutually exclusive decisions about repeated simultaneous encounters has the effect of scaling encounter rates by the animal's attention pattern. However, as with the impulsive behavior, a rate-maximizing result like the one shown in Fig. 5a can be forced by preloading the forager with an ad libitum feeding period to ensure the initial behavior is as exclusive as possible (i.e., ensuring that the initial rate R estimate is relatively high). Again, treating animals with pre-experiment feeding is un-

common in conventional operant laboratory methods because animals are starved to enforce compliance with the experimental apparatus. Unfortunately, it is possible that starved preferences are overly inclusive; a starved individual may eat foods that it would rarely choose willingly in its natural environment.

Simulation: DRM-inspired rule has DRM-like preferences

In Fig. 6a, the gain trajectory for a forager using the DRM-style decision algorithm is shown. The gain trajectory switches across the profitability line for type 2. Moreover, as shown in Fig. 6b, the prey types are partitioned into a set of always-accepted types, always-rejected types, and a single partially accepted type (i.e., type 2 in this example). Similar results can be found by defining digestive profitability as g_i/κ_i for each type i , which guarantees the same type ordering as the DRM model provided that the proportionality constant (i.e., κ_i/k_i for each type i) is identical across all types. However, when digestive profitability is defined as $g_i/(h_i + \kappa_i)$ for each type i , then this forager behaves as a traditional rate maximizer when h_i is significantly larger than κ_i .

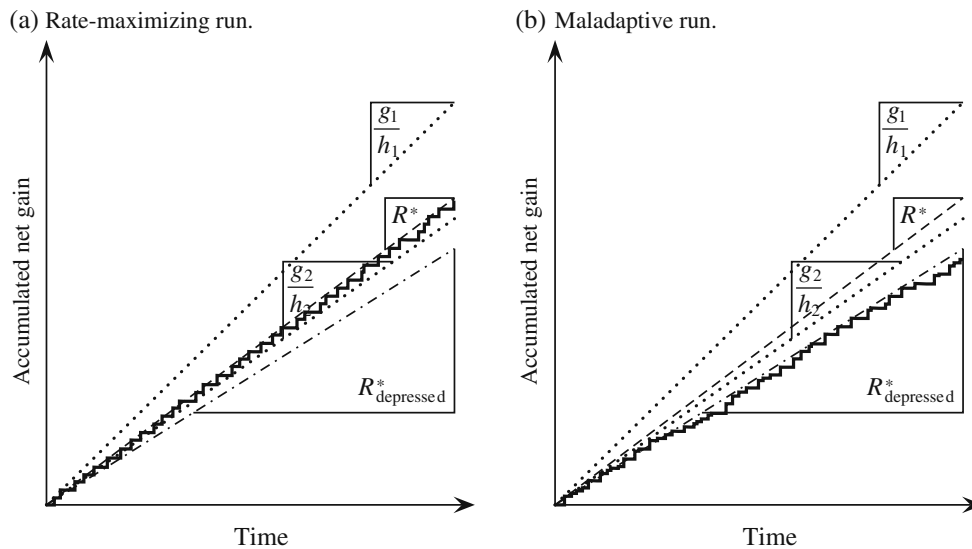


Fig. 5 Equal-opportunity behavior facing simultaneous encounters. During the 100-s foraging bout, the forager encounters high-profitability prey of type 1 ($g_1 = 4.5$, $h_1 = 2$ -s) and low-handling-time prey of type 2 ($g_2 = 1.6$, $h_2 = 1$ -s). All encounters are simultaneous and are generated by a Poisson process with a 1.7-s^{-1} encounter rate. During search, the forager pays cost $c = 0.1\text{-s}^{-1}$. On an encounter with a prey, it is chosen for processing if its profitability is greater than the quotient of the accumulated

gain and total lifetime, and the forager makes a mistake in this calculation with 5% probability. On simultaneous encounters, the forager randomly attends to one of the encountered prey. In the typical run shown in **a**, the gain trajectory converges to the rate-maximizing optimal result. In another typical run shown in **b**, the gain trajectory converges to a result congruent with rate maximization if encounter rates were halved

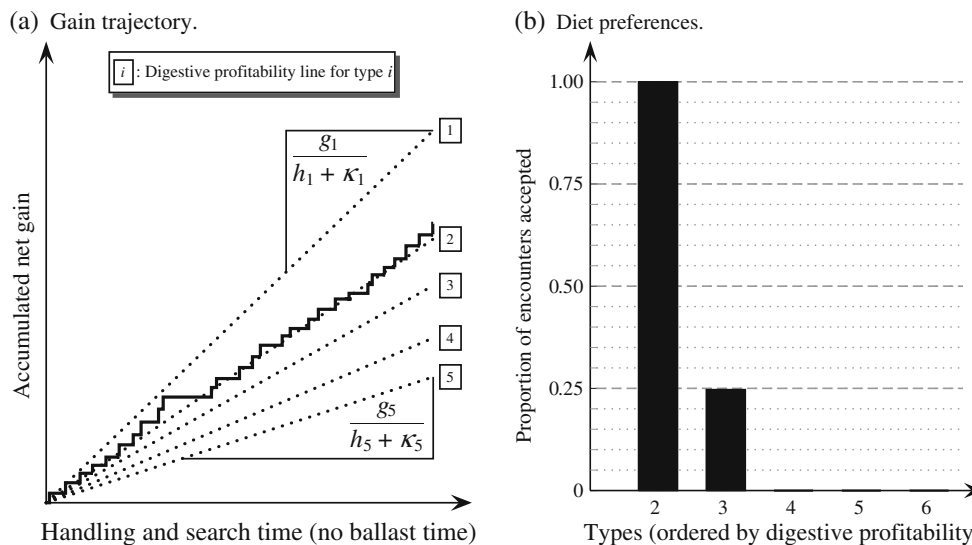


Fig. 6 Simulation of a DRM-inspired foraging behavior that has DRM-like results. During the 100-s foraging bout, the forager encounters prey of types 1 ($\lambda_1 = 1\text{-s}^{-1}$, $g_1 = 3$, $h_1 = 0.5$ -s, $\kappa_1 = 5$ -s), 2 ($\lambda_2 = 0.75\text{-s}^{-1}$, $g_2 = 5$, $h_2 = 2$ -s, $\kappa_2 = 5$ -s), 3 ($\lambda_3 = 0.75\text{-s}^{-1}$, $g_3 = 4$, $h_3 = 0.25$ -s, $\kappa_3 = 4$ -s), 4 ($\lambda_4 = 1\text{-s}^{-1}$, $g_4 = 4$, $h_4 = 0.5$ -s, $\kappa_4 = 3$ -s), and 5 ($\lambda_5 = 0.1\text{-s}^{-1}$, $g_5 = 5$, $h_5 = 0.1$ -s, $\kappa_5 = 3$ -s) according to a five-way merged Poisson process. The prey types are ordered by digestive profitability, which is defined as $g_i/(h_i + \kappa_i)$ for each type i . In this example scenario, there is no

search cost (i.e., $c = 0\text{-s}^{-1}$). On simultaneous encounters (which occur with probability 0 due to the merged Poisson process), the forager attends to the prey with lowest handling time first (i.e., its attention is impulsive). The gain trajectory is shown in **a**, and the proportion of encounters that are accepted is shown in **b**. As shown, prey of type 1 are always accepted, prey of types 3, 4, or 5 are always rejected, and prey of type 2 are accepted on 25% of encounters with them

Discussion

Binary-choice impulsiveness can be sequentially optimal in nature

In the laboratory, animals forced to make a choice between two mutually exclusive alternatives show an impulsive tendency for short handling time even when that behavior fails to maximize long-term rate of gain. We have shown that animals that are impulsive in binary-choice experiments can still be rate maximizers in nature. Our hypothesis is that on an encounter with a prey, the animal chooses to process the prey based on whether the prey's gain–time ratio is greater than the animal's present accumulated gain–total time ratio. In the rare case when multiple prey are encountered at the same time, the animal arbitrarily attends to the one with shortest handling time first. Because these cases are so rare, the animal's attention pattern has little impact on its lifetime success. However, because laboratory encounters are always simultaneous and mutually exclusive, the food-deprived animal fails to make good decisions.

To test this hypothesis, animals can be given a short ad libitum feeding period directly before testing. This feeding period should raise their gain–time ratios so that the purely impulsive choice is not a viable alternative. If this hypothesis is robust, then impulsiveness may not be a curiously irrational behavior so much as a behavioral polymorphism that is normally masked to the effects of natural selection. Additionally, this simple behavioral algorithm may provide intuition for understanding complex state-based animal behavior. An important future direction is the testing of this theory in the laboratory; however, it is consistent with previous experiments (e.g., Bradshaw and Szabadi 1992; Snyderman 1983) that show that impulsiveness is reduced when animals have more access to food. Additionally, the R process is a discrete-time denumerable Markov chain, and it is likely that the convergence we have demonstrated in simulation can be shown analytically via mathematical proof. To establish confidence that this algorithm is a robust rate maximizer, the stochastic convergence of R should be analytically characterized.

A notable weakness of the central hypothesis of this paper is that it takes for granted that the animal attends first to prey with short handling times when facing simultaneous encounters. The salient feature of this assumption is that the animal's attention modulates the perceived encounter rates of encountered prey, which is consistent with the attention-training experiment of Siegel and Rachlin (1995) and the conclusion of

Monterosso and Ainslie (1999) that impulsiveness may be linked to attention and not choice. If the forager only sees simultaneous encounters with a low-profitability short-time prey and a high-profitability long-time prey and always attends to the short-time prey first, it is as if the high-profitability prey has a null encounter rate. Moreover, an equal-opportunity forager facing all simultaneous encounters may also appear to be overly inclusive to short-time prey. During the period of time where its initial rate estimate is low enough to consider the low-profitability short-time prey, the forager forced to make mutually exclusive choices on each simultaneous encounter will behave as if it were facing encounter rates that are halved and thus will continue to be overly inclusive. It is true that a forager that attends first to high-profitability prey on simultaneous encounters will outperform others in binary-choice experiments. However, if simultaneous encounters are rare in nature, it is unlikely that natural selection will have shaped these attention mechanisms, and a wide variation of attention schemes may be present in a background population. Furthermore, the assumption that animals exhibit their natural preferences when starved seems flawed regardless of the underlying behavioral mechanism. This paper serves as an example of how starved preferences can be unique.

A mechanism consistent with the digestive rate model

We have also shown how a simple modification to the our rate-maximizing behavior leads to adaptive diet choices that are qualitatively consistent with the digestive rate model of digestively constrained rate maximizers. In particular, the mechanism automatically orders types by digestive profitability, which is a quantity similar to digestive quality in the DRM, and results in partial preference for a single prey type that partitions the other prey types into a set of always-accepted and always-ignored types. To construct digestive profitability, we defined an internal ballast time that is likely proportional to the material quantity used in DRM constraints. Hence, digestive profitability for each type is the ratio of its average gain to the sum of its handling time and ballast time. Like the approach of Whelan and Brown (2005), this model achieves qualitatively similar results as DRM without a hard digestive rate constraint. Like the short-time attentive rate-maximizing mechanism we introduced, this DRM-consistent mechanism may also be prone to maladaptive impulsiveness in the case of simultaneous encounters. However, our objective in presenting this simple DRM-consistent behavioral model is to suggest intuition for visualizing DRM results and

reconciling differences between ecological and physiological schools of foraging thought. For example, our model shows clearly how a forager with high energy stores may prefer the less bulky of two items with equal profitability while a forager with low energy stores will accept both. An important future direction is to precisely define what a ballast time is in nature.

Acknowledgements We thank Thomas A. Waite for his helpful insights and instruction and Ian M. Hamilton for his comments on and suggestions for this paper. We are also grateful for the comments and suggestions of two anonymous referees. An anonymous reviewer on a related submission also provided helpful comments that influenced the presentation of this work.

References

- Ainslie GW (1974) Impulse control in pigeons. *J Exp Anal Behav* 21(3):485–489
- Ainslie GW (1975) Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychol Bull* 82(4):463–496
- Andrews BW, Passino KM, Waite TA (2004) Foraging theory for decision-making system design: task-type choice. In: Proceedings of the 43rd IEEE conference on decision and control, Nassau, Bahamas, vol 5, pp 4740–4745. doi:10.1109/CDC.2004.1429539
- Bateson M, Kacelnik A (1996) Rate currencies and the foraging starling: the fallacy of the averages revisited. *Behav Ecol* 7(3):341–352. doi:10.1093/beheco/7.3.341
- Bradshaw CM, Szabadi E (1992) Choice between delayed reinforcers in a discrete-trials schedule: the effect of deprivation level. *Q J Exp Psychol* 44(1):1–16
- Charnov EL (1973) Optimal foraging: some theoretical explorations. Dissertation, University of Washington
- Charnov EL (1976) Optimal foraging: attack strategy of a mantid. *Am Nat* 110(971):141–151
- Giraldeau L-A, Livoreil B (1998) Game theory and social foraging. In: Dugatkin LA, Reeve HK (eds) *Game theory and animal behavior*. Oxford University Press, New York, pp 16–37
- Giraldeau L-A, Caraco T (2000) *Social foraging theory*. Princeton University Press, Princeton
- Green L, Fisher Jr EB, Perlow S, Sherman L (1981) Preference reversal and self-control: choice as a function of reward amount and delay. *Behav Anal Lett* 1:43–51
- Henly SE, Ostdiek A, Blackwell E, Knutie S, Dunlap AS, Stephens DW (2008) The discounting-by-interruptions hypothesis: model and experiment. *Behav Ecol* 19(1):154–162. doi:10.1093/beheco/arm110
- Hirakawa H (1995) Diet optimization with a nutrient or toxin constraint. *Theor Popul Biol* 47(3):331–346
- Hirakawa H (1997a) Digestion-constrained optimal foraging in generalist mammalian herbivores. *Oikos* 78(1):37–47
- Hirakawa H (1997b) How important is digestive quality? A correction of Verlinden and Wiley's digestive rate model. *Evol Ecol* 11(2):249–251
- Houston AI, McNamara JM (1985) The choice of two prey types that minimizes the probability of starvation. *Behav Ecol Sociobiol* 17(2):135–141. doi:10.1007/BF00299245
- McDiarmid CF, Rilling ME (1965) Reinforcement delay and reinforcement rate as determinants of schedule preference. *Psychon Sci* 2:195–196
- Monterosso J, Ainslie GW (1999) Beyond discounting: possible experimental models of impulse control. *Psychopharmacology* 146:339–347
- Pulliam HR (1974) On the theory of optimal diets. *Am Nat* 108(959):59–74
- Pyke GH, Pulliam HR, Charnov EL (1977) Optimal foraging: a selective review of theory and tests. *Q Rev Biol* 52(2):137–154
- Quaintance G, van Gils JA, Bocher P, Dekinga A, Piersma T (2010) Diet selection in a molluscivore shorebird across Western Europe: does it show short- or long-term intake rate-maximization? *J Anim Ecol* 79(1):53–62. doi:10.1111/j.1365-2656.2009.01608.x
- Quijano N, Andrews BW, Passino KM (2006) Foraging theory for multizone temperature control. *IEEE Comput Intell Mag* 1(4):18–27. doi:10.1109/MCI.2006.329704
- Rachlin H, Green L (1972) Commitment, choice and self-control. *J Exp Anal Behav* 17(1):15–22
- Schoener TW (1971) Theory of feeding strategies. *Annu Rev Ecol Syst* 2:369–404
- Siegel E, Rachlin H (1995) Soft commitment: self-control achieved by response persistence. *J Exp Anal Behav* 64(2):117–128
- Sih A, Christensen B (2001) Optimal diet theory: when does it work, and when and why does it fail? *Anim Behav* 61(2):379–390. doi:10.1006/anbe.2000.1592
- Snyderman M (1983) Optimal prey selection: the effects of food deprivation. *Behav Anal Lett* 3:359–369
- Stephens DW (2002) Discrimination, discounting and impulsivity: a role for an informational constraint. *Philos Trans R Soc B* 357(1427):1527–1537. doi:10.1098/rstb.2002.1062
- Stephens DW, Krebs JR (1986) *Foraging theory*. Princeton University Press, Princeton
- Stephens DW, Anderson D (2001) The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behav Ecol* 12(3):330–339
- Stephens DW, McLinn CM (2003) Choice and context: test a simple short-term choice rule. *Anim Behav* 66(1):59–70. doi:10.1006/anbe.2003.2177
- Stephens DW, Kerr B, Fernández-Juricic E (2004) Impulsiveness without discounting: the ecological rationality hypothesis. *Proc R Soc B* 271(1556):2459–2465. doi:10.1098/rspb.2004.2871
- van Gils JA, de Rooij SR, van Belle J, van der Meer J, Dekinga A, Piersma T, Drent R (2005) Digestive bottleneck affects foraging decisions in red knots *Calidris canutus*. I. Prey choice. *J Anim Ecol* 74(1):105–119. doi:10.1111/j.1365-2656.2004.00903.x
- Verlinden C, Wiley RH (1989) The constraints of digestive rate: an alternative model of diet selection. *Evol Ecol* 3(3):264–272. doi:10.1007/BF02270727
- Whelan CJ, Brown JS (2005) Optimal foraging and gut constraints: reconciling two schools of thought. *Oikos* 110(3):481–496. doi:10.1111/j.0030-1299.2005.13387.x