

The Sunk-cost Effect as an Optimal Rate-maximizing Behavior

Theodore P. Pavlic · Kevin M. Passino

Received: 7 March 2009 / Accepted: 30 June 2010 / Published online: 24 July 2010
Springer Science+Business Media B.V. 2010

Abstract Optimal foraging theory has been criticized for underestimating patch exploitation time. However, proper modeling of costs not only answers these criticisms, but it also explains apparently irrational behaviors like the sunk-cost effect. When a forager is sure to experience high initial costs repeatedly, the forager should devote more time to exploitation than searching in order to minimize the accumulation of said costs. Thus, increased recognition or reconnaissance costs lead to increased exploitation times in order to reduce the frequency of future costs, and this result can be used to explain paradoxical human preference for higher costs. In fact, this result also provides an explanation for how continuing a very costly task indefinitely provides the optimal long-term rate of gain; the entry cost of each new task is so great that the forager avoids ever returning to search. In general, apparently irrational decisions may be optimal when considering the lifetime of a forager within a larger system.

Keywords Solitary animal behavior · Patch residence time · Rationality · ← Concorde fallacy · Escalation error · Optimal foraging theory

T. P. Pavlic (✉) · K. M. Passino

Department of Electrical and Computer Engineering, The Ohio State University, 205 Dreese Labs,
2015 Neil Avenue, Columbus, OH 43210, USA
e-mail: pavlic.3@osu.edu

K. M. Passino
e-mail: passino@ece.ohio-state.edu

K. M. Passino

Department of Evolution, Ecology, and Organismal Biology, The Ohio State University,
318 W. 12th Avenue, Columbus, OH 43210, USA

1 Introduction

Influential work by Schoener (1971) postulates that natural selection favors behaviors on a continuum from foraging time minimization to net energetic gain maximization. Techniques developed by Pyke et al. (1977) and Charnov (1973) use gain-to-time rate maximization to quantitatively predict which behavior from this continuum will be favored in particular cases. These techniques have been popularized by Stephens and Krebs (1986) as the *prey* and *patch* models of classical optimal foraging theory (OFT). They respectively describe which prey foragers should include in their diet and how long foragers should exploit a patch of prey, which are the two central questions of solitary foraging theory.

The rate-maximizing prey and patch models have mixed success in explaining behavioral observations in the field. The prey model accurately describes patterns of preference, and the patch model makes accurate predictions about how foraging durations should change when background parameters change, but the patch model has poor success in many cases when predicting actual foraging durations (Stephens and Krebs 1986; Nonacs 2001; Sih and Christensen 2001). In a review by Nonacs (2001), several examples are picked that show foraging durations tend to be longer than expected by the classical patch model. Nonacs concludes that OFT is incomplete and that optimal behavior is not described by rate maximization. However, recent studies show how observed behaviors that are inconsistent with classical OFT are indeed rate maximizing under an adjusted foraging model. For example, shorebirds that appear to violate classical OFT are shown to be rate maximizers when explicitly modeling digestive constraints (van Gils et al. 2005) and the value of information discovery (van Gils et al. 2003).

Here, we show how explicitly modeling foraging costs can answer Nonacs' criticisms of the classical OFT's underpredictions of foraging duration. Additionally, we show how the same modifications lead to foraging theoretic explanations of the *sunk-cost effect* (Arkes and Blumer 1985; Arkes and Ayton 1999; Kanodia et al. 1989; Staw 1981), which describes behaviors that invest more time in the more costly of two otherwise equivalent resources. This effect is also known as the *Concorde fallacy* (Dawkins and Carlisle 1976) because it describes an apparent logical fallacy analogous to the continued development of a supersonic jet that never returns a profit. However, by giving a foraging theoretic explanation, we show that the fallacy is actually an optimal behavior. This explanation is also consistent with observations of animals that commit longer feeding times after moving into areas where prey requires greater energy to acquire (Nolet et al. 2001). Additionally, if a forager's uncertainty in estimating patch quality is modeled as an initial average cost associated with each patch type, the extended foraging durations predicted by our augmented OFT is consistent with the "wait for good news" behaviors predicted and observed in supposed Bayesian foragers (Olsson and Holmgren 1998; van Gils et al. 2003; Olsson and Brown 2006).

This paper is organized as follows. First, in Sect. 2, we summarize the principal results from classical OFT. Next, in Sect. 3, we present a common empirical criticism of OFT and provide a mathematical response supporting OFT. In Sect. 4,

we augment the traditional graphical analysis method from OFT to include effects of search and in-patch costs. This analysis shows how explicitly including costs leads to longer patch residence times that are more consistent with observed behavior. We use similar graphical methods in Sect. 5 to provide a detailed explanation of the sunk-cost effect as an adaptive rate-maximizing behavior. Finally, in Sect. 6, we give some concluding remarks.

2 Classical Optimal Foraging Theory

Optimal foraging theory models a forager that faces $n \in \mathbb{N}$ types of prey. The forager encounters prey of type $i \in \{1, 2, \dots, n\}$ according to an independent Poisson counting process with rate $\lambda_i > 0$. Each encountered prey item of type i is picked for handling with probability $p_i \in [0,1]$, and those chosen items are exploited for $\tau_i \geq 0$ time on average. The forager pays a cost (e.g., energy) of c^s per unit search time, but it receives an average reward of $g_i(\tau_i)$ after handling an item of type i . Assuming that natural selection favors foragers that maximize their lifetime gain, behaviors should trade increased reward per encounter for increased number of encounters in a lifetime. So the optimal behavior will maximize the rate

$$\frac{-c^s + \sum_{i=1}^n \lambda_i p_i g_i(\tau_i)}{1 + \sum_{i=1}^n \lambda_i p_i \tau_i}. \quad (1) \leftarrow$$

The *prey model* treats the case where τ_i is fixed for each type i and the forager must make a binary choice to handle or ignore each encountered prey. The *patch model* treats the opposite case where $p_1 = p_2 = \dots = p_n = 1$ (i.e., all encountered prey are exploited) and the forager must choose when to stop exploiting each prey. A central result of the patch model is the *marginal value theorem* (MVT) (Charnov 1973, 1976), which states that nonzero exploitation times $\tau_1^*, \tau_2^*, \dots, \tau_n^*$ that make

$$g'_i(\tau_i^*) = \frac{-c^s + \sum_{i=1}^n \lambda_i g_i(\tau_i^*)}{1 + \sum_{i=1}^n \lambda_i \tau_i^*} \quad \text{and} \quad g''_i(\tau_i^*) < 0 \quad (2)$$

for each type i will correspond to an optimal patch foraging behavior. That is, behaviors that equate per-prey rate of gain with per-lifetime rate of gain will appropriately balance present reward with future opportunity cost so as to be optimal over a lifetime.

3 OFT Criticism and Explicit Processing Costs

Nonacs (2001) bases his criticisms on the observation that Eq. 2 implies that $g'_1(\tau_1^*) = g'_2(\tau_2^*) = \dots = g'_n(\tau^*)$. That is, if behaviors in nature could be described by the MVT, then every per-type rate of average gain (i.e., “speed” of gain) should be equal to some global *giving-up density* (GUD). In the studies that Nonacs reviews, foragers stop exploiting different prey types at different speeds, which leads him to the conclusion that the MVT does not hold in reality.

As discussed by Stephens and Krebs (1986) the function g_i is not meant to be a gross observable reward to the forager. Instead, it models the energetic reward to the forager minus the internal handling cost, which is usually not externally observable. To make this distinction clearer, we explicitly introduce an average handling cost $c_i(\tau_i)$ for exploiting prey of type i for an average of τ_i time, and we let $g_i(\tau_i)$ be the corresponding average observable reward. Thus, the forager should maximize

$$\frac{-c^s + \sum_{i=1}^n \lambda_i p_i (g_i(\tau_i) - c_i(\tau_i))}{1 + \sum_{i=1}^n \lambda_i p_i \tau_i}. \quad (3)$$

In this framework, the MVT states that times $\tau_1^*, \dots, \tau_n^*$ that make

$$g'_i(\tau_i^*) - c'_i(\tau_i^*) = \frac{-c^s + \sum_{i=1}^n \lambda_i (g_i(\tau_i^*) - c_i(\tau_i^*))}{1 + \sum_{i=1}^n \lambda_i \tau_i^*} \quad \text{and} \quad g''_i(\tau_i^*) - c''_i(\tau_i^*) < 0 \quad (4)$$

will correspond to an optimal patch behavior. From Eq. 4, it must be that $g'_1(\tau_1^*) - c'_1(\tau_1^*) = g'_2(\tau_2^*) - c'_2(\tau_2^*) = \dots = g'_n(\tau_n^*) - c'_n(\tau_n^*)$. The rare case when the observable rewards are such that $g'_1(\tau_1^*) = g'_2(\tau_2^*) = \dots = g'_n(\tau_n^*)$ is when the internal handling cost rates are also all equal, and so a global GUD is exceptional. Differences in terminal handling densities across types reflect differences in the handling burden of those types. For example, for types with the same reward–time curve and linear handling cost functions (i.e., $c_i(\tau_i) \triangleq c_i \tau_i$ for type i and constant $c_i > 0$), smaller marginal handling costs (i.e., $c'_i(\tau_i) = c_i$ for type i) should lead to longer optimal exploitation times—foragers spend relatively less time in patches with steep handling-cost functions.

4 Graphical Optimization and Long Residence Times

The intuitive graphical analysis approach frequently used by Stephens and Krebs (1986) for the single-type case can be extended to our model, which makes all costs explicit and uses multiple types. To do so, we define $\lambda \triangleq \lambda_1 + \lambda_2 + \dots + \lambda_n$ to be the average Poisson encounter rate for all prey types combined. Then $1/\lambda$ is the average search time between encounters, and

$$\bar{g} \triangleq \sum_{i=1}^n \frac{\lambda_i}{\lambda} p_i g_i(\tau_i), \quad \bar{c} \triangleq \sum_{i=1}^n \frac{\lambda_i}{\lambda} p_i c_i(\tau_i), \quad \text{and} \quad \bar{\tau} \triangleq \sum_{i=1}^n \frac{\lambda_i}{\lambda} p_i \tau_i$$

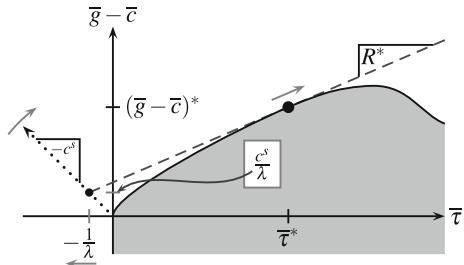
are the average per-encounter handling gain, cost, and time. Equation 3 then becomes

$$\frac{\bar{g} - \bar{c} - \frac{c^s}{\lambda}}{\frac{1}{\lambda} + \bar{\tau}}. \quad (5)$$

Graphical optimization of Eq. 5 is shown in Fig. 1.

Each point in the shaded area corresponds to a particular choice of preferences, p_1, p_2, \dots, p_n , and exploitation times, τ_1, \dots, τ_n . The upper boundary of the shaded area is the optimal frontier on which all optimal behaviors are found. Optimization consists of rotating a ray originating from $(-1/\lambda, c^s/\lambda)$ counter-clockwise from a

Fig. 1 Effect of searching on optimal patch residence times. As search cost c^s or average interarrival time $1/\lambda$ increases, the slope R^* of the optimal ray will decrease, and the optimal average exploitation time $\bar{\tau}^*$ will increase



–90° angle to the least upper bound of the angles less than 90° where the ray intersects the shaded region. The slope of the optimal ray R^* is the optimal value of Eq. 5, and the corresponding values $(\bar{g} - \bar{c})^*$ and $\bar{\tau}^*$ are the optimal per-encounter average net handling gain and exploitation time. Hence, this graphical interpretation of rate maximization allows for predictions of how perturbing parameters of the model will impact the average handling time $\bar{\tau}$. An increase in the average handling time may be due to an increase in the optimal exploitation time for each patch (i.e., changes in τ_i^* for each type i) or due to additional patch types being added to the optimal diet (i.e., changes in p_i^* for each type i) or a mixture of the two effects. This graphical approach does not suggest how to predict each individual τ_i^* or p_i^* ; it is only meant to investigate how certain parameter perturbations will affect general foraging trends. Specific optimal diet content and patch exploitation times can be found using algorithms from classical OFT (Charnov 1973, 1976; Stephens and Krebs 1986; Pavlic 2007).

When Stephens and Krebs (1986) consider graphical solutions to the single-type case, they frequently discuss how changes in encounter rate λ_1 lead to changes in optimal behavior. For multiple types, the overall rate λ can be changed while holding the density λ_i/λ constant for each type i . The graphical example in Fig. 1 shows that the optimal average exploitation time $\bar{\tau}^*$ will increase if either the overall encounter rate λ decreases or the search cost rate c^s increases. In both cases, to maximize lifetime reward, the forager must increase present gains to compensate for future search losses—the *opportunity cost* of more exploitation decreases when search time or cost increase. Likewise, if the average per-encounter handling cost \bar{c} is increased, the $\bar{g} - \bar{c}$ curve will be shifted down the vertical axis and the optimal average exploitation time $\bar{\tau}^*$ will increase. This prediction suggests an explanation for the ostensibly anomalous observation by Nolet et al. (2001) that tundra swans spend more time feeding in areas of deep water where more energy is required to acquire similar prey as in shallow water. The increased costs necessarily decrease the maximal long-term rate of gain, which is the opportunity cost of increased exploitation, and so patches are exploited longer. This effect is investigated in more detail in Sect. 5.

5 The Sunk-Cost Effect

Graphical optimization of Eq. 3 helps explain apparently irrational behaviors discussed in the fields of economics, psychology, and biology in terms of lifetime

gain maximization. In economics, the propensity to continue a costly task after paying some initial cost is sometimes called an *escalation commitment*, *escalation behavior*, or *escalation error*, but it is more commonly known as the *sunk-cost effect* (Kanodia et al. 1989; Staw 1981). This nomenclature is also used by psychologists (e.g., Arkes and Blumer 1985). In fact, psychologists Arkes and Ayton (1999) note that the sunk-cost effect is equivalent to the *Concorde fallacy* first described by biologists Dawkins and Carlisle (1976). Although none of these terms are used, the same phenomena is also observed by Nolet et al. (2001). In particular, tundra swans must expend more energy to “up-end” to feed on deep-water tuber patches than they do to “head-dip” to feed on shallow-water patches; however, contrary to the expectations of Nolet et al., the swans feed for a longer time on each high-cost deep-water patch. In every context, the observation of the sunk-cost effect is an enigma because intuition suggests that this behavior is suboptimal. Here, we show how optimization of Eq. 3 predicts the sunk-cost effect for certain scenarios; a common element of every case is a large initial cost.

5.1 Initial Costs: Recognition, Acquisition, Reconnaissance

For simplicity, we make a graphical argument with the assumptions that $n = 1$, $c^s = 0$, and $p_1 = 1$ (i.e., the patch case). We also revert to interpreting g_1 as the average net handling gain (i.e., the sum of an observable gain and an internal cost). Under these assumptions,

$$\frac{g_1(\tau_1) \leftarrow}{\frac{1}{\lambda_1} + \tau_1} \quad (6)$$

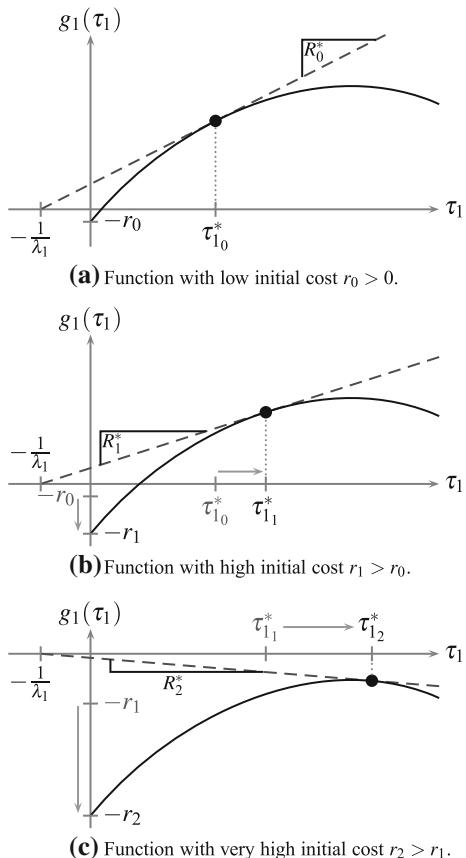
is the objective function for optimization. Similar arguments can be made about the general multivariate form in Eq. 3.

Consider the case when gain function g_1 is initially:

- (1) negative (i.e., $g_1(0) < 0$),
- (2) increasing (i.e., $g'_1(0) > 0$), and
- (3) decelerating (i.e., $g''_1(0) < 0$)

Functions meeting items (2) and (3) are treated by Stephens and Krebs (1986). However, these functions are all initially zero. Here, we let the average handling gain $g_1(\tau_1)$ be initially negative to reflect some initial cost. For example, rather than treating recognition cost as a separate quantity (e.g., Stephens and Krebs 1986, pp. 79–81), we consider them to be a characteristic of the gain returned to the forager from handling each item (e.g., an initial energy expenditure required simply for access to or identification of a prey). Alternatively, as in the case of the tundra swans observed by Nolet et al. (2001), the initial cost may model the extra energy required to acquire prey (e.g., by “up-ending” instead of simply “head-dipping”). Furthermore, just as Olsson and Brown (2006) show how Bayesian foragers receive a relative “foraging benefit of information” for spending extra time in a patch, the initial cost here may be viewed as an average penalty due to uncertainty about the patch gain function.

Fig. 2 Different recognition costs $r_2 > r_1 > r_0$ have different optimal exploitation times $\tau_{l_2}^* > \tau_{l_1}^* > \tau_{l_0}^*$ and optimal rates $R_2^* < R_1^* < R_0^*$. For two equally shaped gain functions, the one with the higher initial cost will also have the higher optimal residence time



The graphical optimization procedure for such a initially negative gain function is shown in Fig. 2. The same gain function is shown in Figs. 2a–c except that initial recognition cost is low, moderate, and high, respectively. When the initial cost increases from r_0 to r_1 , the optimal exploitation time increases from $\tau_{l_0}^{*\leftarrow}$ to $\tau_{l_1}^{*\leftarrow}$. Even when the initial cost increases from r_1 to r_2 and makes the gain function strictly negative, the exploitation time still increases from $\tau_{l_1}^{*\leftarrow}$ to $\tau_{l_2}^{*\leftarrow}$. This propensity to exploit items longer when the initial cost is increased is exactly the sunk-cost effect (Arkes and Ayton 1999; Arkes and Blumer 1985). However, it also maximizes the long-term net rate of gain, and so it is the rational behavior.

This result can be explained using an *opportunity cost* interpretation (Houston and McNamara 1999; Stephens and Krebs 1986). The optimal rate R^* is an opportunity cost for spending additional time exploiting an item. It represents the maximum expected gain possible per unit time, and so it is costly to exploit items of a particular type for so long that the type's rate of average gain falls below R^* . When the initial cost of handling increases on all encounters, the total gain decreases, and so the opportunity cost for additional exploitation decreases (e.g., $R_1^* < R_0^*$). Thus, the marginal handling gain must decrease further in order for marginal costs and

marginal benefits to equalize. This effect manifests itself through the increase in optimal exploitation time.

By increasing exploitation time when initial cost increases, the forager reduces the amount of time it spends paying high costs. Leaving patches at an earlier time is equivalent to volunteering for paying recognition costs more frequently. It is not the increased initial cost on a single encounter that is important; it is the increased initial cost on all encounters that causes the decrease in opportunity cost for additional exploitation.

5.2 Human and Nonhuman Examples

5.2.1 Feeding Time Increases in Areas of High-Cost Prey

In a study by Nolet et al. (2001), the foraging behavior of tundra swans in shallow water is compared to the behavior in deep water. The swans forage on belowground tubers, and thus the swans need only to “head-dip” in shallow water whereas they need to “up-end” in deep water to find and retrieve the tubers. The model that the authors use to explain the tundra swan behavior predicts that there will be a decrease in feeding time in areas where feeding has larger power requirements. However, in deep water where the power requirements to up-end are apparently larger than the requirements to head-dip, the observed tundra swans spent longer times feeding on tuber patches. Although this result is a contradiction of the model used by Nolet et al., it follows directly from the discussion in Sect. 5.1. In particular, let $n \geq 1$ (i.e., multiple tuber patch types are available) and $p_1 = \dots = p_n = 1$ (i.e., encounters are not ignored). For each patch type i , the average net gain function g_i is assumed to be initially negative in order to model the energetic burden of acquiring the belowground tubers. Hence, we use the classical marginal value theorem result in Eq. 2 applied to initially negative gain functions.

The curve in Fig. 3a models the average value g_i of a tuber patch of type i after exploiting the patch for time τ_i . The average cost c_i of a patch entry reduces the maximum possible long-term rate of gain R^* from all patches. The MVT predicts that the optimal exploitation time τ_i^* occurs when $g'_i(\tau_i^*) = R^*$, which is depicted in Fig. 3b for three different R^* maximum rates corresponding to three different c_i entry costs. As the average entry cost of a patch type increases, the maximum rate decreases, and so average exploitation time increases. When R^* is very high either due to low cost c_i or high return rates from other patches, the optimal commitment time drops to zero because the shallow value function is dominated by steep returns from other options.

5.2.2 Human Preference for High Price

One recorded example of the sunk-cost effect in humans comes from studying behavior at the cinema. In a controlled study, Arkes and Blumer (1985) charge movie patrons different prices for tickets to see a particular movie. They show that the likelihood that people attend the movie is positively correlated with the cost of the ticket, and they conclude that this behavior is an irrational mistake. However,

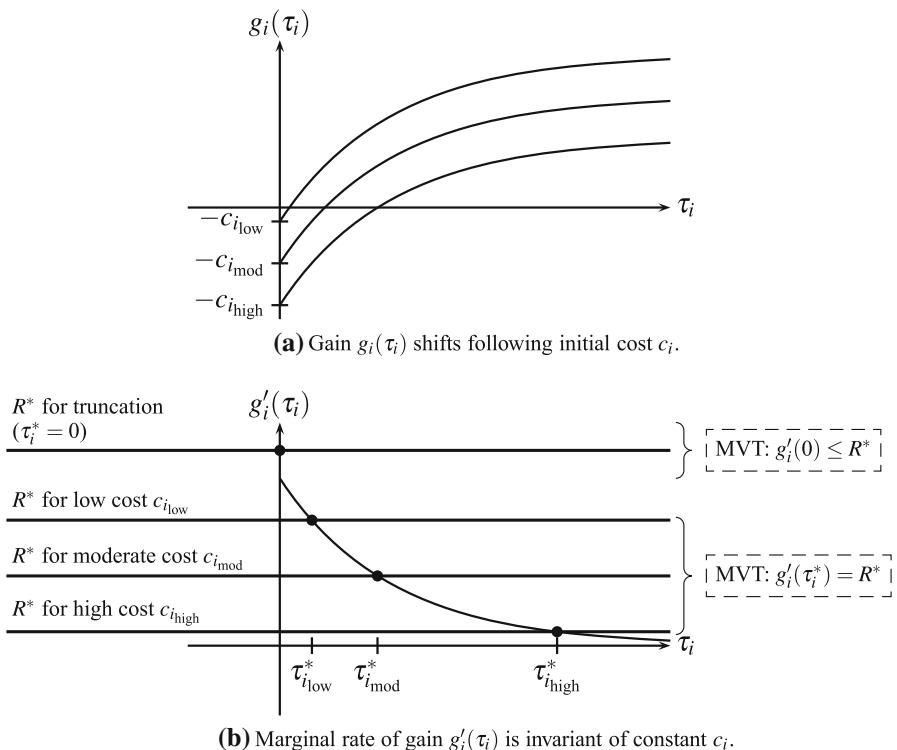


Fig. 3 Optimal exploitation time for patch type i with acquisition cost c_i . Increased costs lead to decreased overall return rate R^* , which leads to greater exploitation time in each costly patch. That is, when most encountered patches are costly to enter, exploitation time in each of these high-cost patches will be long in order to reduce their frequency. However, when enough other low-cost patches are available to raise the overall return rate sufficiently high, exploitation time to patches with relatively shallow marginal returns may decrease to zero. That is, when there are enough high-quality patches, time is better spent searching for them

this result can also be predicted by the classical marginal value theorem result in Eq. (2) applied to initially negative gain functions.¹ Again, consider the case where $n \geq 1$ (i.e., multiple types are available) and $p_1 = \dots = p_n = 1$ (i.e., encounters are not ignored). For each activity type i , let the gain function $g_i(\tau_i)$ be parameterized by *commitment time* τ_i that represents the average time the individual is committed to participating in that type of activity. That is, if the commitment time for a type of activity is low, there is a low probability that the individual will complete the activity, and thus the average gain returned from this type of activity will also be low. Likewise, if the commitment time for the activity is high, it will likely be attended, and so the value returned will level off as it returns no value after it is attended. Hence, the analysis summarized in Fig. 3 also applies here. As ticket price

¹ Here, to be consistent with Arkes and Blumer (1985), we do not allow encounters to be ignored, and so initial costs are forced and the pure patch model predicts the optimal behavior. The combined prey-patch model better fits reality as ticket purchasing opportunities can be ignored.

increases for this set of individuals who are forced to buy the ticket, they become more strongly committed to attending the movie.²

This example matches classical economic intuition. As the price a consumer pays increases, the overall purchasing frequency decreases. The longer commitment to a higher-priced movie reflects reluctance to make more purchases. Over the long term, this behavior maximizes the net value the consumer holds collectively in goods and capital. Behaving in any other way would be reckless when viewed from a long-term opportunity-cost perspective. In particular, exploitation times are longer in order to accumulate more gain to justify returning to a search that will likely result in finding more “prey” that force high initial costs. Thus, the sunk-cost effect may seem nonsensical for any single purchase, but it is adaptive when considering an economic or ecological system view.

5.3 Escalation Behavior

The propensity to continue a costly task ad infinitum is an extreme form of the sunk-cost effect known as *escalation behavior* or *escalation error* (Kanodia et al. 1989; Staw 1981). In fact, the *Concorde fallacy* gets its name from the example given by Dawkins and Carlisle (1976) of continued government investment on a supersonic jet that has no promise of ever returning a profit. Here, we show how escalation behavior can be an adaptive rate-maximizing strategy under special circumstances. These examples presume that the optimal long-term rate of gain R^* is negative, and so they are a better fit for economic examples where prolonged deficits are allowed. However, as several studies show that foragers can dynamically adjust behavior to maximize the present estimate of long-term rate of gain (Giraldeau and Livoreil 1998; Giraldeau and Caraco 2000; Sih and Christensen 2001), these examples may also be consistent with short-term trends in behavior when foragers are temporarily subjected to high-cost environments. Additionally, this analysis may apply to similar mathematical models of value maximization when behaviors must choose from a set of costly options. For example, in spiders that choose to join existing webs or spin their own (Jakob 2004), small spiders face potentially dangerous competition with large spiders in new webs but also face high costs in building new webs. Also, studies of women in abusive relationships suggests their decisions to leave the relationship are affected by fear of violence toward pets that remain (Faver and Strand 2003). In both of these cases, the optimal behavior may be accompanied by costs that might seem prohibitive in isolation. As in Sect. 5.1, we use the simplified optimization objective given by Eq. 6.

In Fig. 4, a positive constant gain function is shown as the limit of asymptotic gain functions of increasing steepness. Because of the asymptotic upper limit, sharp increases in gain correspond to sharp decreases in marginal returns. So the optimal patch residence time decreases as steepness increases because marginal returns fall so quickly. In the limiting case, when the gain function is constant, the optimal behavior is to leave immediately because staying longer cannot result in any

² If the experimenters allowed for encounters to be ignored (i.e., if participants could choose to not purchase a ticket), movies with zero commitment times would also have zero ticket sales.

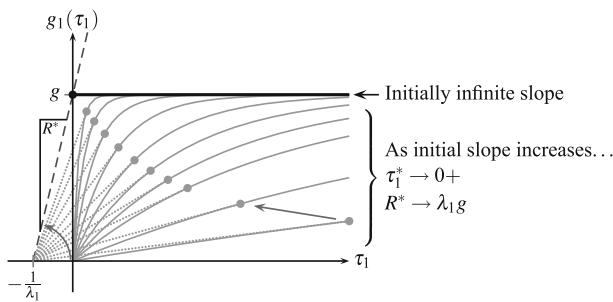


Fig. 4 Analysis of positive constant gain function. The gain function is the limit of smooth asymptotic gain functions of increasing steepness. The limiting optimal behavior is to leave the patch immediately. Staying longer returns no additional gain, and leaving immediately decreases the time between subsequent high-quality patches

additional gain. In particular, it is better to invest time searching for new patches that guarantee an entry reward g at regular rate λ_1 (i.e., long-term rate of gain $\lambda_1 g$) rather than staying in any single patch that provides no additional gain (i.e., long-term rate of gain $g/\infty \approx 0$).

The opposite effect is demonstrated in Fig. 5, which shows a negative gain function as the limit of asymptotic gain functions of decreasing steepness. As the steepness of each gain function is reduced, the optimal patch residence time increases without bound. So the optimal behavior for a constant gain function that is negative is to remain forever in each encountered patch. It is better to pay a single recognition cost once and remain in the patch forever rather than paying the recognition cost repeatedly at each new encounter.

As shown in the lower portion of Fig. 5, even when the gain function is initially decreasing and convex, the optimal behavior is not to immediately leave the patch unless the initial slope is sufficiently negative. It is better for the forager to pay additional costs within the patch rather than face the $-\lambda_1 g$ lifetime rate of gain

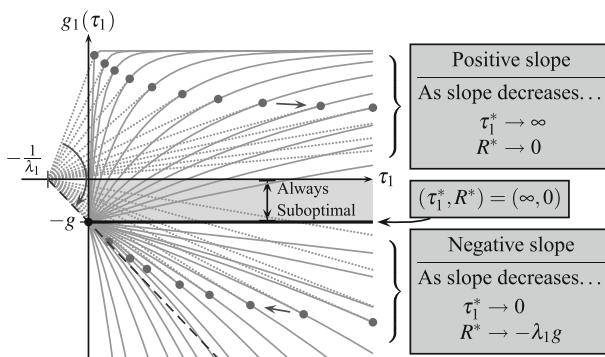


Fig. 5 Analysis of negative constant gain function. Escalation error can be viewed as the limit of increasingly shallow gain functions that are initially negative. Zero patch residence time is not restored until steepness is sufficiently negative. For this class of functions, optimal points will never be found in the shaded region between the horizontal axis and the constant gain function

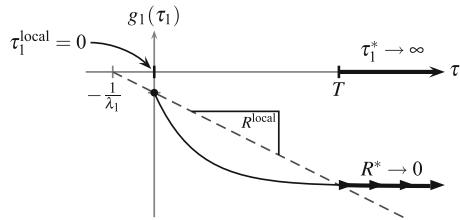


Fig. 6 Initially negative concave gain function has latent escalation behavior that dominates the early local maximum. Leaving the patch immediately returns a higher long-term gain than staying a short time longer, but it returns a lower gain than staying for any time greater than T . Hence, escalation behavior is optimal because it prevents paying any future entry costs

associated with immediately leaving the patch. However, escalation behavior is avoided in this case because the convex gain function eventually becomes too steep. A different effect occurs for initially negative and initially steep concave functions like the one shown in Fig. 6. Here, immediate withdrawal from the patch is a locally optimal behavior. That is, the objective function decreases as the exploitation time is increased slightly. However, as exploitation time increases, the objective function reaches its global minimum and then starts increasing. Moreover, there exists a time T such that exploiting these patches for any time $\tau_1 \geq T$ will be favorable to immediate withdrawal. Again, it is better to exploit patches with a shallow negative gain longer than to search for more patches that have a steep gain.

In the supersonic jet construction example (Dawkins and Carlisle 1976), the costly behavior may be adaptive because it prevents starting other costlier ventures. More generally, successful politicians in power may induce endless and massive spending in one area for no other reason than to avoid new spending in other areas. In a foraging context, a forager that must pay a large energetic cost to enter a patch (e.g., due to high predation risk, having to swim upstream, or high uncertainty about patch quality) may hesitate to ever leave that patch if the habitat contains few other types of patches. Similarly, social group members may resist the temptation to leave a low-quality group if joining any other group is accompanied by costly antagonistic violence (e.g., Faver and Strand 2003; Jakob 2004). In general, in a set of apparently bad choices, the choice that has the lowest eventual marginal losses will minimize the very long-term losses.

6 Conclusions

Optimal foraging theory introduces quantitative analysis into the study of behavior in a way that complements intuition. We have shown how explicitly modeling foraging costs can improve the accuracy of rate-maximizing OFT methods and answer the criticisms of Nonacs (2001) that rate maximization too often underestimates patch residence time. As verified by van Gils et al. (2003), the Bayesian forager modeled by Olsson and Holmgren (1998) accurately predicts observed patch residence times. This result is investigated by Olsson and Brown (2006) who derive a “foraging benefit of information” that is an additional reward

that Bayesian foragers receive for staying longer in patches to gather information. As these models investigate foraging behavior using numerical dynamic programming techniques, they lack the intuition of classical OFT. However, if the discovery penalty of the information gathering process is modeled as an initial in-patch cost, then relatively simple classical OFT methods also predict longer exploitation times.

We also show how explicitly including large initial in-patch costs in classical OFT methods explains why sunk-cost effects like those observed by Nolet et al. (2001) are rational rate-maximizing foraging behaviors. This result seems counter-intuitive, but it follows from using a farsighted decision-making model based on comparing present returns to the opportunity cost of not searching. When encounters with costly patches are likely, it is better to spend more time in each patch in order to reduce the accumulation rate of said costs. This interpretation may also explain human preference for high price (e.g., Arkes and Blumer 1985) and the escalation of costly behaviors ad infinitum (e.g., Dawkins and Carlisle 1976; Faver and Strand 2003). Sunk-cost behaviors are criticized because they are examples of making decisions based on past investments; however, they should be better understood as decisions that reduce the frequency of similar future investments.

Acknowledgments We thank Thomas A. Waite for his helpful insights and instruction and Ian M. Hamilton for his comments on this paper. We also thank two anonymous referees for their help in improving this paper. Additionally, the comments of three anonymous reviewers on a related submission have also been influential in the presentation of this work.

References

- Arkes HR, Ayton P (1999) The sunk cost and Concorde effects: are humans less rational than lower animals? *Psychol Bull* 125(5):591–600
- Arkes H, Blumer C (1985) The psychology of sunk cost. *Organ Behav Hum Decis* 35:124–140
- Charnov EL (1973) Optimal foraging: some theoretical explorations. PhD thesis, University of Washington
- Charnov EL (1976) Optimal foraging: the marginal value theorem. *Theor Popul Biol* 9(2):129–136
- Dawkins R, Carlisle TR (1976) Parental investment, mate desertion and a fallacy. *Nature* 262(5564): 131–133
- Faver CA, Strand EB (2003) To leave or to stay? *J Interpers Violence* 18(12):1367–1377. doi:[10.1177/0886260503258028](https://doi.org/10.1177/0886260503258028)
- Giraldeau LA, Caraco T (2000) Social foraging theory. Princeton University Press, Princeton
- Giraldeau LA, 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
- Houston AI, McNamara JM (1999) Models of adaptive behavior. Cambridge University Press, Cambridge
- Jakob EM (2004) Individual decisions and group dynamics: why pholcid spiders join and leave groups. *Anim Behav* 68(1):9–20 doi:[10.1016/j.anbehav.2003.06.026](https://doi.org/10.1016/j.anbehav.2003.06.026)
- Kanodia C, Bushman R, Dickhaut J (1989) Escalation errors and the sunk cost effect: an explanation based on reputation and information asymmetries. *J Account Res* 27(1):59–77
- Nolet BA, Langevoord O, Bevan RM, Engelaar KR, Klaassen M, Mulder RJW, Dijk SV (2001) Spatial variation in tuber depletion by swans explained by differences in net intake rates. *Ecology* 82(6):1655–1667 doi:[10.1890/0012-9658\(2001\)082\[1655:SVITDB\]2.0.CO;2](https://doi.org/10.1890/0012-9658(2001)082[1655:SVITDB]2.0.CO;2)
- Nonacs P (2001) State dependent behavior and the marginal value theorem. *Behav Ecol* 12(1):71–83
- Olsson O, Brown JS (2006) The foraging benefits of information and the penalty of ignorance. *Oikos* 112(2):260–273 doi:[10.1111/j.0030-1299.2006.13548.x](https://doi.org/10.1111/j.0030-1299.2006.13548.x)
- Olsson O, Holmgren NMA (1998) The survival-rate-maximizing policy for bayesian foragers: wait for good news. *Behav Ecol* 9(4):345–353

- Pavlic TP (2007) Optimal foraging theory revisited. Master's thesis, The Ohio State University, Columbus. http://www.ohiolink.edu/etd/view.cgi?acc_num=osu1181936683
- Pyke GH, Pulliam HR, Charnov EL (1977) Optimal foraging: a selective review of theory and tests. *Q Rev Biol* 52(2):137–154
- Schoener TW (1971) Theory of feeding strategies. *Annu Rev Ecol Syst* 2:369–404
- 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](https://doi.org/10.1006/anbe.2000.1592)
- Staw BM (1981) The escalation of commitment to a course of action. *Acad Manag Rev* 6(4):577–587
- Stephens DW, Krebs JR (1986) Foraging theory. Princeton University Press, Princeton
- van Gils JA, Schenk IW, Bos O, Piersma T (2003) Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches. *Am Nat* 161(5):777–793. doi:[10.1086/374205](https://doi.org/10.1086/374205)
- van Gils JA, de Rooij SR, van Belle J, van der Meer J, Dekking A, Piersma T, Drent R (2005) Digestive bottleneck affects foraging decisions in red knots shape *Calidris canutus*. I. prey choice. *J Anim Ecol* 74(1):105–119. doi:[10.1111/j.1365-2656.2004.00903.x](https://doi.org/10.1111/j.1365-2656.2004.00903.x)