Received: 14 June 2022 | Revised: 8 December 2022 | Accepted: 14 December 2022

DOI: 10.1111/ele.14160

LETTER

ECOLOGY LETTERS WILEY

Strategy maps: Generalised giving-up densities for optimal foraging

Emerson Arehart[1] | **Jody R. Reimer**[2] | **Frederick R. Adler**[2]

[1]Department of Biology, University of Pennsylvania, Philadelphia, Pennsylvania, USA

[2]Department of Mathematics and School of Biological Sciences, University of Utah, Salt Lake City, Utah, USA

Correspondence
Emerson Arehart, Department of Biology, 433 S University Ave, Philadelphia, PA 19104, USA.
Email: arehart@sas.upenn.edu

Funding information
The Modeling the Dynamics of Life Fund

Editor: Minus van Baalen

Abstract

Finding a common currency for benefits and hazards is a major challenge in optimal foraging theory, often requiring complex computational methods. We present a new analytic approach that builds on the Marginal Value Theorem and giving-up densities while incorporating the nonlinear effect of predation risk. We map the space of all possible environments into *strategy regions*, each corresponding to a discrete optimal strategy. This provides a generalised quantitative measure of the trade-off between foraging rewards and hazards. This extends a classic optimal diet choice rule-of-thumb to incorporate the hazard of waiting for better resources to appear. We compare the dynamics of optimal decision-making for three foraging life-history strategies: One in which fitness accrues instantly, and two with delays before fitness benefit is accrued. Foragers with delayed-benefit strategies are more sensitive to predation risk than resource quality, as they stand to lose more fitness from a predation event than instant-accrual foragers.

KEYWORDS

diet choice, giving up density, landscape of fear, Markov decision process, optimal foraging, stochastic dynamic programming

INTRODUCTION

Organisms must repeatedly make decisions that balance risk and reward. For example, foraging animals must choose whether to continue eating food that they have already found or move on in the hopes of finding something better. Optimal Foraging Theory is the study of optimising such a decision process. Due to its critical importance for understanding animal behaviour and potential applications in industrial, social and computational science, optimal foraging has been the subject of intense theoretical and empirical work for over 60 years (Holling, 1959; Mangel, 2015; Pyke, 1984; Stephens, 2008).

Classical Optimal Foraging Theory considers a forager in a region with heterogeneously distributed resources. The forager encounters a patch of resources, depletes resources in the patch at a diminishing rate, and must choose when to depart and seek out a new patch (MacArthur & Pianka, 1966). The Marginal Value Theorem provides a criterion for when a forager should abandon a patch and switch to another: exactly when the expected rate of return for remaining in a patch drops below the expected average rate across all patches (Charnov, 1976a). A related finding applies to Optimal Diet Choice, in which a forager must choose between different food types. By comparing the rates of return (the ratio of expected value to handling time) of different resource types in the environment, one can determine which types of food should be included in a forager's diet, and in what order of preference (Charnov, 1976b).

While elegant, these simple rate-maximisation foraging methods fail to capture the full complexity of real foraging behaviour (Hirvonen & Ranta, 1996; Mangel et al., 1988; Nonacs, 2001). One central issue is the tradeoff between gathering food and the risk of becoming food for something else. A variety of approaches have been developed to incorporate predation risk into optimal foraging (Houtman & Dill, 1998; Yearsley et al., 2001), but perhaps the most well-known is the concept of giving-up densities, or GUDs (Brown, 1988; Brown et al., 1999). Consider a forager depleting resources in a patch, such as a squirrel digging for seeds in an exposed area of dirt, where the risk of being attacked by a bird is high. The GUD is the density of seeds remaining – that is, rate of food return – at which the squirrel should abandon the patch because it estimates that continuing to forage in an exposed setting poses a hazard that is not worth a lower level of reward (Brown, 1988). The GUD approach has been applied to a range of experimental systems

(Bedoya-Perez et al., 2013), allowing comparison of risk (or perceived risk) in different environments.

Analysing GUDs poses several key challenges, including potentially nonlinear relationships between harvest rate, energetic gain and predation risk (Bedoya-Perez et al., 2013). In practice, GUD approaches have most commonly been employed qualitatively, for example comparing the "fear" foragers feel in open versus closed terrain (Brown et al., 1999). Quantitative analysis of fear behaviour requires finding a common currency for all factors affecting fitness (Krebs & Davies, 1978; Mangel, 2015; Mangel & Clark, 1986; McNamara & Houston, 1986). Previously, this required more complex mathematical tools such as Stochastic Dynamic Programming.

Stochastic Dynamic Programming (SDP) is a versatile tool for finding optimal strategies for complex, iterative decision processes (Bellman, 1957; Houston et al., 1993; Mangel, 2015). In SDP, decision processes, in which decisions must be made at a sequence of discrete times, are codified as recursive mathematical expressions. The value of any chain of decisions is computed by iterative calculation either forward or backward in time, depending on the structure of the problem. Foraging decision-making processes have been analysed with SDP (Brodin et al., 2017; Houston et al., 1988; Iwasa et al., 1984; Pirotta et al., 2019; Reimer et al., 2019b). SDP has also been fruitfully applied to finding optimal strategies in domains ranging from wildlife management/conservation reserve design to epidemiology and industrial optimisation (reviewed in Marescot et al. (2013)).

SDP methods often track the "state" of a forager, such as fat reserves, and factor the state into decision-making: For example, a starving bird may tolerate a high degree of predation risk in order to get a certain food reward, whereas a satiated one might only claim that food reward if predation risk is much lower. In such an analysis, the feedback between decisions and forager states must be tracked, greatly increasing the computational complexity of a problem (Mangel et al., 1988). Techniques exist for speeding up or approximating SDP for complex state-dependent SDP problems (Fackler, 2018; Nicol & Chadès, 2011; Reimer et al., 2019a).

Classical Optimal Foraging Theory produced many "rules of thumb" that might allow foragers to approximate optimal strategies (Houston et al., 1993; Pyke, 1984). Many such "rules" fell out of favour, as they failed to explain the full range of observed foraging behaviours. GUD methods are easy to interpret conceptually and yield useful insight into the fitness cost of predation but can be challenging to analyse quantitatively (Bedoya-Perez et al., 2013). SDP enables elaborate computations and can reveal optimal strategies for specific complex problems, but produces results for each specific case, rather than general rules.

Here, we present a mathematical framework that extends rate-based optimal foraging and optimal diet choice to include nonlinear trade-offs between food and predation risk. We use this framework to extend the rule of thumb from Gilliam and Fraser (1987) to incorporate the hazard from taking no resources, revealing additional complexity in optimal decision behaviour. Our method balances predation risk and expected reward, generalising GUDs (Brown, 1988; Brown et al., 1999) to multiple resource-type environments (as in Optimal Diet Choice theory) and without requiring predation-free refugia. This maps *all possible environments* to their corresponding optimal strategies, providing a baseline for understanding forager behaviour. Our framework provides a high-level view of how strategies change as a forager's environment changes, quantifying the value of information for the forager.

We demonstrate the versatility of our framework by comparing strategy dynamics for three biologically relevant fitness paradigms: EGGS, in which fitness is measured by total resources collected (analogous to the rate-maximising foragers from classic optimal foraging, but with predation risk included), AON (all or nothing), in which a forager collects resources on a foraging bout but gains no benefit unless it survives until a fixed terminal time, and PRE (periodic return), in which the forager engages in multiple, shorter trips, only receiving fitness benefits if it survives to the end of each trip. These fitness paradigms are special cases of the canonical fitness equations in Mangel (2015). The threshold for accepting a resource is much more sensitive to predation risk in the AON and PRE paradigms, due to hazard compounding over the course of a foraging bout. We close by discussing the implications of our approach for the value of information and learning.

METHODS

Consider an environment with multiple resource types and a forager encountering them in discrete time. The granularity of time steps can be chosen at whatever scale is appropriate for the problem. For foraging bees, a time step might be 1 s; for lions, a time step might be hours or days (Mangel et al., 1988). λ_i is the probability of encountering resource i on a given time step. The forager can choose to consume resource i if it appears, which occupies the animal for the handling time (h_i time steps), or it can skip the resource and advance one time step. This reflects the opportunity cost of taking resource i, which is resources that the organism would encounter if it was not busy handling resource i. If multiple resources appear at the same time, the forager can choose only one of the resources; if no resources appear, the forager moves forward one time step.

If resource i is consumed, the organism receives a nutritional reward with expected value v_i, and survives with probability $q_i^{h_i}$ (corresponding to a risk of predation or other hazard of $1 - q_i$ applied at each of

h_i time steps during the handling time). If the forager chooses to wait and collect no resources, it survives with probability q_0, again modelling predation risks as well as other generalised hazards, including the risk of starvation.

The forager determines whether to consume or skip resources based on a strategy Z which ranks resources: it will accept resources in the order dictated by Z, and reject any resources not included in Z. The foraging process continues until it is terminated either due to the death of the individual, or by reaching an assigned terminal time T. We assume that the forager can finish consuming any resource it encounters on or before time T, even if the processing time extends beyond T. See Appendix C for a graphical explanation of the foraging setup.

For the remainder of this paper, we consider environments in which there are only two resource types. Our results generalise to environments with more resource types, but are much easier to calculate and visualise with two. In a two-resource system, there are five possible strategies: accept nothing $\left(Z = \{\emptyset\}\right)$, only ever accept resource 1, ($Z = \{1\}$), only ever accept resource 2, ($Z = \{2\}$), accept resource 1 when available but take resource 2 if 1 is not available, ($Z = \{1, 2\}$), and the converse (accept resource 2 when available but take resource 1 if 2 is not available, $Z = \{2, 1\}$).

The strategy which maximises fitness depends on what we call the *fitness paradigm*, which determines how fitness benefits and costs affect a forager. Most classical optimal foraging theory is concerned with maximising foraging rates, but some organisms follow life-history strategies for which rate is not the correct metric. For example, if an individual worker bee must return to the hive in order to benefit from the resources it has collected, then the fitness cost of being killed by a predator depends on how many resources it has collected since leaving the hive. For a solitary parasitic wasp, which intersperses egg laying with foraging, the cost of predation is not dependent on the whole foraging bout, since it continuously converts food into offspring. These differences are captured by the total fitness function, $\Phi(Z)$, which summarises the expected overall fitness obtained by following a strategy Z. The goal of our method is to find the strategy Z that maximises $\Phi(Z)$, the total fitness that has accumulated by the end of a foraging bout.

For the first part of this paper, we will concern ourselves with the EGGS fitness paradigm, in which the organism instantaneously converts the reward v_i into "eggs," and thus into fitness, as long as it survives feeding on resource i. This is a special case of the "resource allocation" canonical equation of SDP, in which fitness is accumulated while the foraging process continues (Mangel, 2015). EGGS is closest to the classic rate-maximising approach to Optimal Foraging and Optimal Diet Choice, while still incorporating nonlinear hazards.

Mathematical framework

We use the probabilistic formulation of SDP to incorporate nonlinear effects of predation into an optimal foraging framework. Consistent with classical optimal foraging theory, however, we ignore state dependence. In our recursive equations, we maximise total foraging success while accounting for the differential risks of predation from each possible action (taking resource 1, taking resource 2 or waiting). We introduce F_t, the expected total reward accrued from time t up to and including the final time step T. The goal is to determine which Z maximises the total fitness $\Phi_{EGGS}(Z) = F_1$ in an environment as defined by λ_i, v_i, h_i and q_i.

At T, the forager encounters only resource 1 with probability $\lambda_1\left(1-\lambda_2\right)$, only resource 2 with probability $\lambda_2\left(1-\lambda_1\right)$, no resources with probability $\left(1-\lambda_1\right)(1-\lambda_2)$, and both resources simultaneously with probability $\lambda_1\lambda_2$. We calculate returns from a maximising strategy, starting with the expected total rewards accrued on the final time step, F_T:

$$F_T = \lambda_1\left(1-\lambda_2\right)v_1 q_1^{h_1} + \lambda_2\left(1-\lambda_1\right)v_2 q_2^{h_2} + \lambda_1\lambda_2 \max\left(v_1 q_1^{h_1}, v_2 q_2^{h_2}\right). \tag{1}$$

Any resource appearing on the final time step T should be accepted, since there are no future rewards the forager could forego by doing so. Going back one time step, the forager still encounters resources at the same rates, but the decision involves different considerations. The forager could forego a lesser reward at time $T-1$ in the hopes of getting something better at time T. Following this logic further backward in time yields a recursive relationship (Mangel et al., 1988):

$$\begin{aligned} F_t = \lambda_1\left(1-\lambda_2\right)\max\left(\left(v_1+F_{t+h_1}\right)q_1^{h_1}, F_{t+1}q_0\right) + \\ \lambda_2\left(1-\lambda_1\right)\max\left(\left(v_2+F_{t+h_2}\right)q_2^{h_2}, F_{t+1}q_0\right) + \\ \lambda_1\lambda_2\max\left(\left(v_1+F_{t+h_1}\right)q_1^{h_1}, \left(v_2+F_{t+h_2}\right)q_2^{h_2}, F_{t+1}q_0\right) + \\ \left(1-\lambda_1\right)\left(1-\lambda_2\right)F_{t+1}q_0, \end{aligned} \tag{2}$$

with $F_t = 0$ for $t > T$ (the forager may finish handling any resource it finds on or before time T, but it cannot accept any new resources after T). The four terms in Equation 2 correspond to the four possible situations the forager could face for making a decision: Only resource 1 available, only resource 2 available, both resources available and neither resource available, respectively. The expressions inside the max functions correspond to the reward from each option the forager can take. Recursively computing this relation back through F_1 determines a strategy for accepting or rejecting a resource at every time step, and a hierarchy for accepting resources if multiple resources appear at the same time. The strategy may change over time, especially towards the end of the foraging bout. However, to benefit from such a change in strategy, the forager would have to know exactly when the foraging bout will end, and the difference in total reward resulting from changing the strategy for only a few time steps is usually small compared to the overall expected fitness over many time steps. In general, we will instead seek the *stationary strategy*, which is the fixed decision rule that offers the best return applied across the foraging bout. The system converges to the stationary strategy for realistic foraging times (Mangel et al., 1988).

By finding the environmental parameter combinations where optimal strategies change, we can construct a global map of optimal strategies across parameter space. Determining the optimal strategy for a given environment is then as straightforward as locating the environmental parameters of interest on the strategy map.

Strategy maps

The optimal strategy Z for a specific environment (specified by λ_i, v_i, h_i and q_i) can be determined with SDP for all combinations of all possible discretised values of each parameter (Figure 1b). However, we would like to find an analytical way to map every possible combination of λ_i, v_i, h_i and q_i to one of these strategies. We will focus primarily on the relationship between v_i and the optimal strategy, but a similar approach can be applied to other environmental parameters.

Transitions between accepting and rejecting resource i occur where the expected value from a strategy including resource i equals the expected value from a strategy excluding it. This transition always occurs along a straight line in the $v_1 - v_2$ plane, as it depends on the ratio of v_1 to v_2 (Figure 1). The slopes of these lines are v_2^*, the transition value of v_2 when $v_1 = 1$, and $\frac{1}{v_1^*}$, where v_1^* is the transition value of v_1 when $v_2 = 1$. The transition in the order of preference between the two resources, $\tilde{v}$, occurs where the expected fitness value of accepting resource 1 exactly equals the expected fitness value of accepting resource 2. In Figure 1c, we plot the ratio of the fitness from the best strategy to the fitness from the second-best strategy across the same parameter ranges, showing the sensitivity of foraging success to choice of foraging strategy. If $\lambda_1 = \lambda_2 = 1$, meaning that both resources are encountered on every time step, v_1^* and v_2^* are undefined, because a forager should always take only the preferred resource (so the strategy map reduces to a single boundary between $Z = \{1\}$ and $Z = \{2\}$).

The following sections describe our methods for constructing and analysing strategy maps. We first derive exact expressions for the total fitness functions $\Phi(Z)$ for each possible strategy Z in the EGGS paradigm. By setting pairs of these expressions equal to each other and solving for v_i, we find v_1^*, v_2^* and $\tilde{v}$ analytically. We compare these expressions with “rules of thumb” from Optimal Foraging and Optimal Diet Choice literature, recovering and extending previous results on balancing resource quality with predation risk. We further illustrate the power and versatility of the strategy maps method by applying the same approach to compare EGGS to two other biologically-inspired fitness paradigms reflecting different life-history strategies.

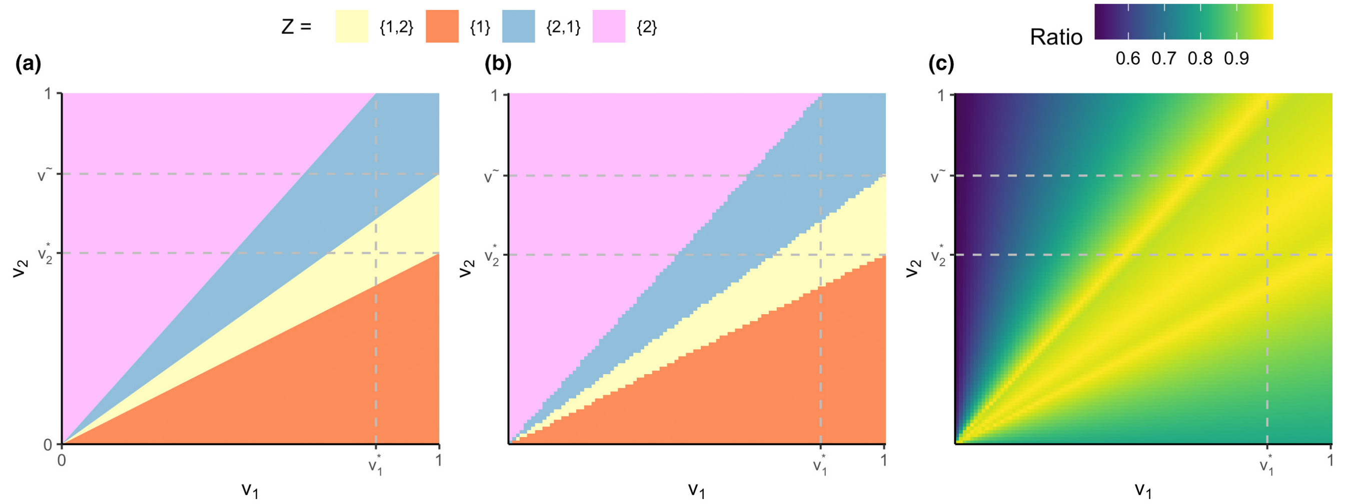

FIGURE 1 Strategy regions in the $v_1 - v_2$ plane, (a) as calculated by our analytic method and (b) as computed by SDP with a horizon of 5000 timesteps. Both methods yield the same mapping of optimal strategies, with the SDP method converging on the analytic results. In (b), each pixel represents one SDP calculation for a specific v_1, v_2 pair. (c) Shows the ratio of the second best strategy to the best strategy for the same parameter range as (a-b); the best two strategies become equal right at lines with slopes $\frac{1}{v_1^*}$, v_2^* and $\tilde{v}$. For all figures, $\lambda_1 = 0.3$, $\lambda_2 = 0.2$, $h_1 = 6$, $h_2 = 5$, $q_0 = 0.999$, $q_1 = .9981$ and $q_2 = .9983$.

Computing v_2^* and $\tilde{v}$

We find v_2^* by setting the total fitness functions $\Phi(\{1,2\})$ and $\Phi(\{1\})$ equal to each other and solving for v_2. We first derive expressions for the total fitness functions in terms of environmental parameters λ_i, v_i, h_i and q_i from the recursive expressions used to formulate such a problem for SDP.

Consider the case where the handling time is identical for both resources, $h = h_1 = h_2$. For strategy $Z = \{1,2\}$, the expected fitness reward per action (accepting a resource) is $f = \lambda_1 q_1^h v_1 + \lambda_2 (1-\lambda_1) q_2^h v_2$. At each time step, the probability of accepting a resource and surviving h time steps is $a = \lambda_1 q_1^h + \lambda_2 (1-\lambda_1) q_2^h$. The probability of not encountering any resource and surviving 1 time step is $b = (1-\lambda_1)(1-\lambda_2) q_0$. Iterating backward in time, the cumulative fitness at time t, F_t, is given by the recursion:

$$F_t = f + aF_{t+h} + bF_{t+1} \tag{3}$$

where $F_t = 0$ for $t > T$. Exploiting the recursive nature of (3), we can write this as a geometric series:

$$F_t = f \sum_{j=0}^{\operatorname{floor}\left(\frac{t}{h}\right)} \left(a^j \sum_{i=0}^{t-jh} \binom{i+j}{j} b^i \right) \tag{4}$$

where $\operatorname{floor}\left(\frac{t}{h}\right)$ is how many times h divides t, ignoring the remainder.

For fixed T, v_2^* exhibits transient behaviour towards the end of the foraging bout due to changes in opportunity cost as the forager runs out of time. Such changes in strategy only benefit a forager with fixed and exactly known T, which is biologically unrealistic. As long as $q_i \neq 1$, the expected length of the foraging bout is constrained by the forager's risk of death, so the expected fitness value will converge as $T - t$ becomes large (as the expected lifespan of the forager falls below T). We take advantage of this by taking the limit $\lim_{t\to\infty}$ of (4) and applying the identity for the sum of binomial coefficients:

$$\lim_{t\to\infty} F_t = \frac{f}{1-a-b} = \frac{\lambda_1 q_1^h v_1 + \lambda_2 (1-\lambda_1) q_2^h v_2}{1 - \lambda_1 q_1^h - \lambda_2 (1-\lambda_1) q_2^h - (1-\lambda_1)(1-\lambda_2) q_0}. \tag{5}$$

An alternate approach is necessary for environments with no hazard (i.e., if $q_i = 1$; see Appendix A).

For $h_1 \neq h_2$, the process is very similar. Setting a_i = probability of surviving and moving ahead h_i time steps,

$$F_t = f + a_1 F_{t+h_1} + a_2 F_{t+h_2} + bF_{t+1} \tag{6}$$

$$\begin{aligned} \lim_{t\to\infty} F_t &= \frac{f}{1-a_1-a_2-b} \\ &= \frac{\lambda_1 q_1^{h_1} v_1 + \lambda_2 (1-\lambda_1) q_2^{h_2} v_2}{1 - \lambda_1 q_1^{h_1} - \lambda_2 (1-\lambda_1) q_2^{h_2} - (1-\lambda_1)(1-\lambda_2) q_0}. \end{aligned} \tag{7}$$

The total fitness is exactly equal to this quantity: $\Phi(\{1,2\}) = \lim_{t\to\infty} F_t$. The same process yields $\Phi(\{1\})$, substituting the correct expressions for f, a, and b. Finally, we set the limiting expressions for $\Phi(\{1,2\})$ and $\Phi(\{1\})$ equal to each other and solve for v_2^*:

$$\Phi(\{1,2\}) = \Phi(\{1\}) \tag{8}$$

$$\begin{aligned} &\frac{\lambda_1 q_1^{h_1} + \lambda_2 (1-\lambda_1) q_2^{h_2} v_2^*}{1 - \lambda_1 q_1^{h_1} - \lambda_2 (1-\lambda_1) q_2^{h_2} - (1-\lambda_1)(1-\lambda_2) q_0} \\ &= \frac{\lambda_1 q_1^{h_1}}{1 - \lambda_1 q_1^{h_1} - (1-\lambda_1) q_0} \end{aligned} \tag{9}$$

$$v_2^* = \frac{q_1^{h_1} v_1}{q_2^{h_2}} \left(\frac{\lambda_1 \left(q_0 - q_2^{h_2}\right)}{1 - q_0 + \lambda_1 \left(q_0 - q_1^{h_1}\right)} \right). \tag{10}$$

Normalising $v_1 = 1$, this expression gives us the slope v_2^* as observed in Figure 1. Note that λ_2 disappears from this expression, consistent with previous findings that the availability of the lesser resource is irrelevant to the strategy for including it in a forager's diet (Charnov, 1976a; Holling, 1959). A similar process yields $\tilde{v}$:

$$\Phi(\{1,2\}) = \Phi(\{2,1\}) \tag{11}$$

$$\tilde{v} = \frac{q_1^{h_1} v_1}{q_2^{h_2}} \left(\frac{(1-\lambda_1)(1-\lambda_2)\left(q_0 - q_2^{h_2}\right) + q_2^{h_2} - 1}{(1-\lambda_1)(1-\lambda_2)\left(q_0 - q_1^{h_1}\right) + q_1^{h_1} - 1} \right). \tag{12}$$

These quantities define the transitions between different regions in our strategy map (Figure 1).

Foraging in patches: Finding GUTs (giving up times), GUDs (giving up densities) and perceived risk

The same functional form for total fitness can be used to calculate optimal patch departure time. Consider an environment with one type of patch, encountered at rate λ and with survival probability q applied for each time step in the patch. Instead of a fixed reward, the total return from foraging in a patch for h time steps is given by a function $v(h)$. As before, q_0 is the survival probability while not in a patch. We assume that $v(h)$ is a saturating function, reflecting the fact that the patch contains finite resources. Applying a method similar to (5) gives an expression for the long-term fitness for a strategy of foraging in each patch for h time steps:

$$\Phi = \frac{\lambda q^h}{1 - \lambda q^h - (1-\lambda) q_0} v(h). \tag{13}$$

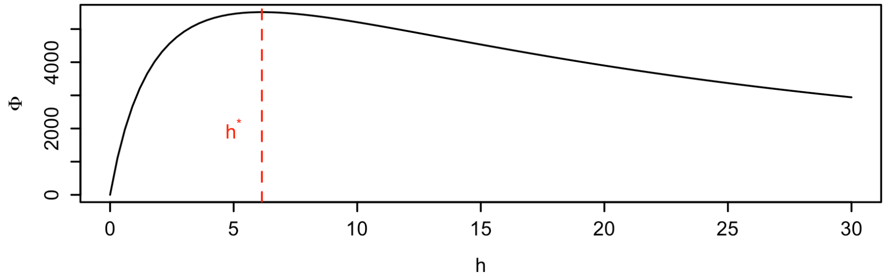

FIGURE 2 Total fitness in (13), Φ, plotted against patch residence time h. To obtain the optimal residence time h^*, solve numerically for the maximum of this curve, which occurs where $\frac{d\Phi}{dh} = 0$. For this figure, $v(h) = \frac{1-r^h}{1-r}$, $r = 0.9$, $\lambda = 0.8$, $q = 0.999$ and $q_0 = 0.9999$.

To find the optimal patch residence time h^*, often referred to as the GUT or giving up time (Charnov, 1976a), we solve for the value of h which maximises the expression in Equation (13). This expression is the convolution of the survival curve with the total fitness return curve. h^* occurs where the derivative of (13) with respect to h equals zero (Figure 2). The details depend on the functional form of $v(h)$, but in general this quantity can only be attained numerically (Stephens et al., 2008). This process is analogous to the procedure used in the Marginal Value Theorem, but includes nonlinear effects of hazard.

Experimental studies have used the GUD concept to compare the forager's perceived risk of predation in different environments (Bleicher, 2017; Brown, 1988; Brown et al., 1999). For given values of q_0, $v(h^*)$ and h^*, all of which can be acquired through GUD experiments, assume the forager is maximising fitness and set $\Phi = v(h^*)$ in Equation (13) to solve for q^*, the quantitative risk estimate that a forager is using to make decisions. This is in contrast to traditional GUD analysis, which primarily allow qualitative comparison between environments. By also incorporating the "background" survival rate q_0, we account for more realistic experimental/field conditions, whereas traditionally a refuge with $q_0 = 1$ (no hazard) is assumed.

To find the GUD, take the patch initial value and subtract the amount harvested over h^* time steps:

$$\mathrm{GUD} = v(0) - v(h^*). \tag{14}$$

For an environment with two patch types, applying this method to each type yields optimal patch residency times h_1^* and h_2^*. The Optimal Diet Choice strategy map approach developed in the previous section can then be applied to the resulting rewards $v_1\left(h_1^*\right)$, $v_2\left(h_2^*\right)$ and residency times h_1^*, h_2^*, determining thresholds for including or rejecting the patch types from the foraging strategy.

RESULTS

Analytic expressions for the boundaries between strategy regions allow us to explore how strategies change as a function of environmental parameters. In Figure 3, we plot v_2^* (the threshold value for including v_2 when $v_1 = 1$) against two such parameters, λ_1 and q_0, to illustrate the power of this method for exploring the structure of strategy space. Increasing λ_1, the availability of the preferred resource, means that the threshold for accepting a secondary resource becomes higher, as a forager's chance of missing preferred resources increases with λ_1. v_2^* increases with q_0, as the hazard of waiting for preferred resources decreases.

Revisiting and extending a rule of thumb for predation

A rule of thumb established by Gilliam and Fraser (1987) relates the income rate and the per-unit-time risk as a decision criterion for including a resource in the diet. This rule of thumb assumes no background hazard while waiting to encounter resources (in our notation, $q_0 = 1$). The Gilliam and Fraser rule of thumb is formulated as follows:

$$v_2^* = \frac{h_2\left(1-q_2\right)}{h_1\left(1-q_1\right)}. \tag{15}$$

In comparison, our expression for v_2^* in (10) yields the following when $q_0 = 1$:

$$v_2^* = \frac{q_1^{h_1}\left(1-q_2^{h_2}\right)}{q_2^{h_2}\left(1-q_1^{h_1}\right)} \tag{16}$$

Rewriting $q_1^{h_1} = \left(1-\varepsilon_1\right)^{h_1} \approx 1 - h_1\varepsilon_1$ and $q_2^{h_2} = \left(1-\varepsilon_2\right)^{h_2} \approx 1 - h_2\varepsilon_2$, where ε_i is the mortality risk from resource i, we find

$$v_2^* \approx \frac{\left(1-h_1\varepsilon_1\right)h_2\varepsilon_2}{\left(1-h_2\varepsilon_2\right)h_1\varepsilon_1} = \frac{\left(1-h_1\varepsilon_1\right)}{\left(1-h_2\varepsilon_2\right)} \frac{h_2\left(1-q_2\right)}{h_1\left(1-q_1\right)}. \tag{17}$$

The expressions in (15) and (17) are equivalent modulo the difference in how hazard applies to the forager. Equations (16) and (17) capture the correlational hazard structure of our model – once the forager accepts resource

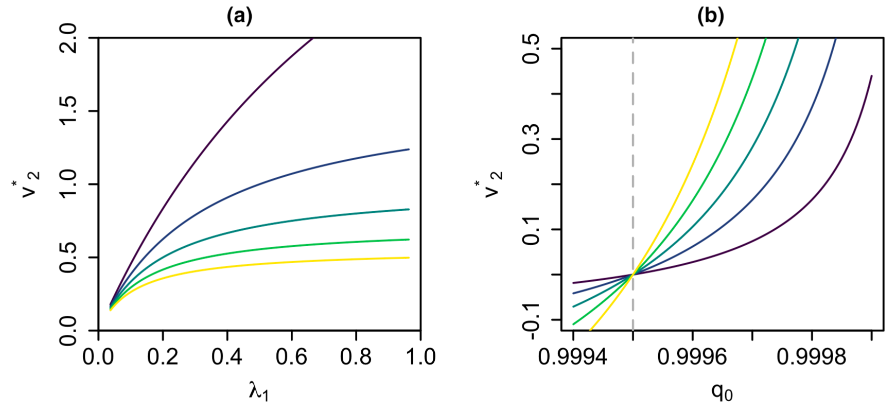

FIGURE 3 Sensitivity of v_2^*, the threshold value for including v_2 when $v_1 = 1$. (a) v_2^* is invariant to λ_2, but sensitive to λ_1. Darker colours (more blue) correspond to higher handling times h_2, scaling v_2^* to accommodate the additional foregone opportunity. $q_1 = q_2 = 0.9999$, and $q_0 = 0.99995$. h_2 ranges from 1 to 5. (b) v_2^* is plotted against survival while not foraging, q_0. The survival associated with each resource is fixed at $q_1 = 0.99999$ and $q_2 = 0.9995$. When $q_0 < q_2$ (vertical dashed grey line), v_2^* takes on negative values, as the forager should trade negative resource values for higher safety. Lighter colours associate with higher λ_1 values, again trading off the opportunity costs of missing a better resource while consuming a secondary resource. For (b), $h_1 = 6$ and $h_2 = 5$. λ_1 ranges from 0.1 to 0.5.

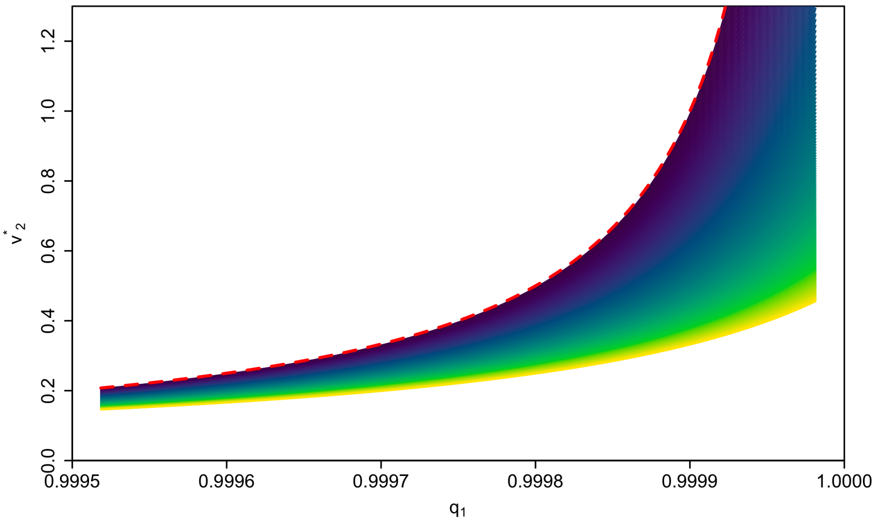

FIGURE 4 Gilliam and Fraser's rule of thumb (red line) compared to results from our method (Equation 10). When $q_0 = 1$, v_2^* is not sensitive to λ_1, and matches the rule of thumb (darker colours). However, as q_0 decreases, v_2^* increasingly diverges from the rule of thumb as λ_1 becomes larger (lighter colours). For this figure, $h_1 = 6$ and $h_2 = 5$.

i, it spends h_i time steps handling it and survives with probability $q_i^{h_i}$ – and balances the opportunity cost of missing one resource while accepting another (Figure 4). However, unlike the rule of thumb, our method can include the background hazard of waiting for better resources ($q_0 \neq 1$). Terms in (10) containing λ_1 cancel if $q_0 = 1$ (16), but with $q_0 \neq 1$, v_2^* becomes sensitive to λ_1. As λ_1 becomes sufficiently small, v_2^* decreases, because the tradeoff between the hazards of taking resources must balance the background hazard $1 - q_0$.

Comparing fitness paradigms

We next consider applications to finite foraging times and other fitness paradigms, exploring differences in the effect of predation risk as a function of the time an organism takes to redeem fitness benefit from resources collected. We compare three cases: EGGS, the instantaneous accrual paradigm from the above sections; the "all or nothing" (AON) paradigm, in which a forager must complete a foraging bout and safely return home before accruing any fitness benefit; and an intermediate version, in which the forager receives fitness benefits periodically throughout the bout, or PRE (periodic return). The AON paradigm only makes sense in finite time, since the probability of surviving an infinitely long foraging bout is zero for q_i less than 1. We develop approximations for EGGS, AON and PRE in finite time in order to compare them directly; for details on the derivation of these approximations, see Appendix B.

In the EGGS paradigm, a special case of the “resource allocation” canonical equation (Mangel, 2015), the expected total fitness for a strategy Z over T time steps, $\Phi_{EGGS}(Z,T)$, is the sum of possible outcomes in which the organism survives exactly t time steps, where the probability of survival until time t for strategy Z, denoted by $S_Z(t)$, is multiplied by α_Z, the expected reward per unit time conditional upon surviving, times t time steps:

$$\Phi_{EGGS}(Z,T) = \sum_{t=0}^{T} t\alpha_Z S_Z(t). \tag{18}$$

In the AON (all-or-nothing) paradigm, a special case of the “activity choice” canonical equation (Mangel, 2015), no fitness is gained if the forager dies before returning home, no matter how many resources it has collected. For example,

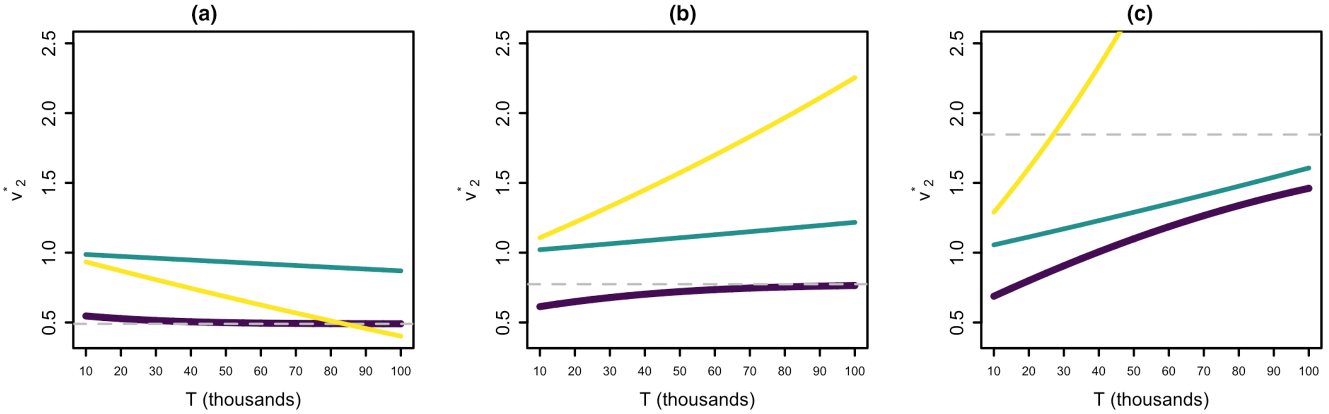

FIGURE 5 v_2^* against T, derived from the approximations in (32) and (34), for EGGS (purple), AON (yellow) and PRE (green). (a) $q_1 < q_2$, with $q_1 = 0.99992$ and $q_2 = 0.99995$, (b) $q_1 = q_2 = 0.99995$ and (c) $q_1 > q_2$, with $q_1 = 0.99998$ and $q_2 = 0.99995$, the values for EGGS converge to the exact value of v_2^* from (10) (dashed grey lines). For this figure, $\lambda_1 = 0.5$, $h_1 = 6$, $h_2 = 5$ and $q_0 = 0.99999$. For PRE, $\tau = \frac{T}{5}$.

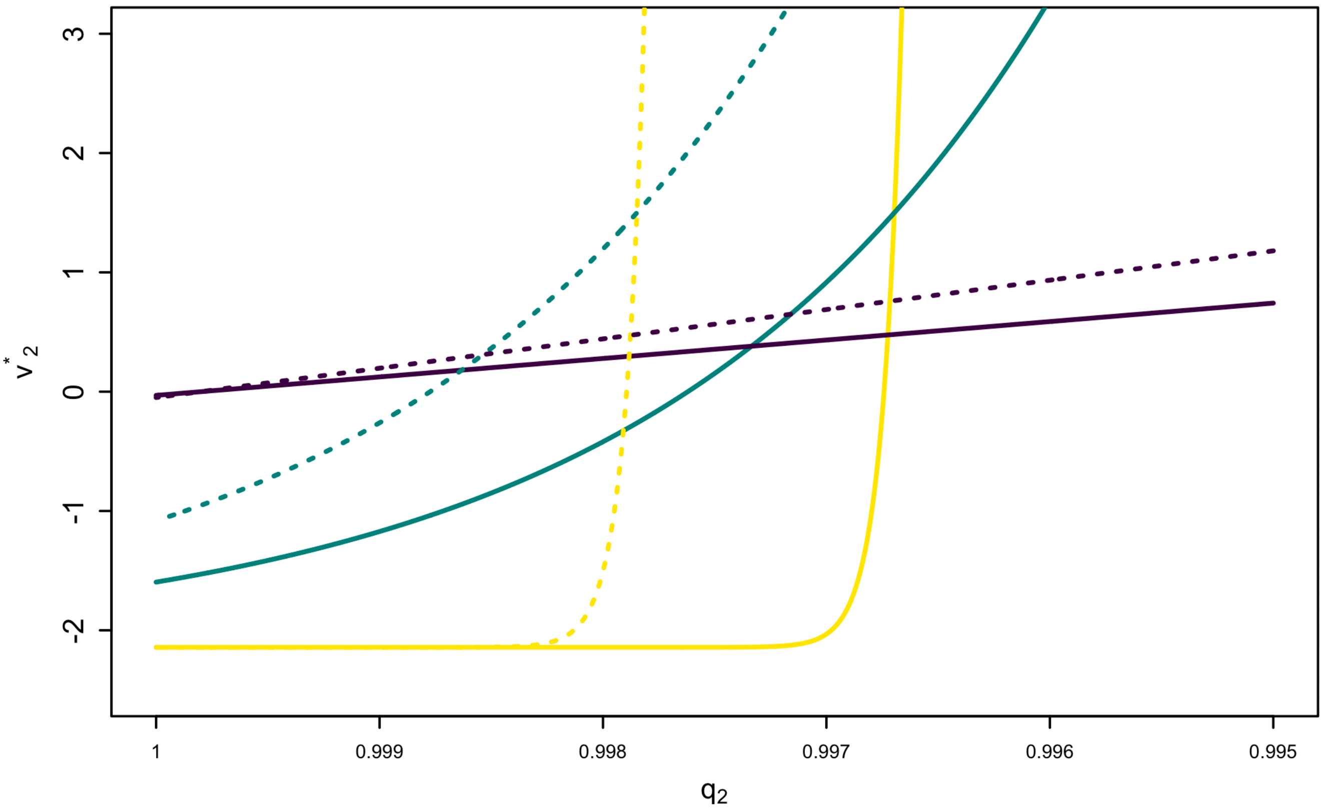

FIGURE 6 v_2^* plotted against survival q_2 for EGGS (purple), AON (yellow) and PRE (green). Solid lines correspond to $q_1 = 0.995$ and dashed lines to $q_1 = 0.997$. For EGGS, v_2^* increases linearly with increasing hazard. For AON and PRE, v_2^* remains low, sometimes even negative, when accepting resource 2 has significantly lower hazard, and then rapidly increases towards infinity. For this figure, $\lambda_1 = 0.3$, $h_1 = 6$, $h_2 = 5$, $q_0 = 0.999$ and $T = 10{,}000$. For PRE, $\tau = 500$.

an individual worker bee collecting resources for the hive will only contribute fitness to the colony if it returns safely from a foraging bout. If it survives until time T, the forager will stop seeking resources and safely return home (we assume no additional travel time for simplicity). Expected total fitness is T multiplied α_Z, the expected reward per unit time under strategy Z conditional upon surviving, multiplied by the probability of surviving until T, $S_Z(T)$:

$$\Phi_{AON}(Z, T) = T\alpha_Z S_Z(T). \tag{19}$$

The PRE paradigm is another special case of the "activity choice" canonical equation of SDP (Mangel, 2015). For example, think of a parent foraging in bouts in order to provision offspring. The PRE paradigm is mathematically almost identical to AON, except that the time horizon for each trip is a fraction τ of the total time T. By optimising the return of one such bout, we optimise the total fitness, so $\Phi_{PRE}(Z, T) = \Phi_{AON}(Z, \tau)$.

We can now directly compare the structure of strategy space for different fitness paradigms, and hence for different life-history strategies. v_2^* is plotted against total foraging time T (Figure 5). For EGGS, v_2^* converges towards the limit we found analytically in (10). In the AON paradigm, v_2^* diverges as T becomes large: any change in survivorship will have an increasing impact over an increasing amount of time, so the tradeoff necessary to balance it grows quickly with T. However, the direction and magnitude of that change depends on the relationship between survival for the two resources. If $q_2 < q_1$, then v_2^* increases rapidly. If $q_2 > q_1$, v_2^* increases much more slowly, and may actually decrease over time when $q_2 > > q_1$, as the forager may be willing to reduce total reward accrued in exchange for avoiding dangerous but valuable resources. PRE follows the same trends as AON, but at a slower rate.

The relationship between v_2^*, and survival for each resource q_i is linear for the EGGS paradigm, but nonlinear in the AON and PRE paradigms (Figure 6). In AON and PRE, v_2^* maintains a low value when q_2 is close to 1, before rapidly climbing as q_2 decreases. This transition occurs where q_2 starts to dominate the dynamics of the system. This effect becomes more dramatic and switch-like as T becomes large. This suggests that foragers in these paradigms will be extremely sensitive to change in hazard within specific parameter ranges but will show little variability otherwise. The negative values of v_2^* in Figure 6 reflect that the forager should include resource 2 in its diet even at a fitness cost due to the lower predation risk associated with it.

DISCUSSION

We have developed new methods for analysing optimal foraging behaviour, integrating nonlinear hazards (which previously required computational techniques such as stochastic dynamic programming) into an analytically tractable framework. We employ our methods to unify and extend classical results from rate-maximising optimal foraging theory, optimal diet choice theory, SDP and GUDs (Brown, 1988). By partitioning environmental parameter space into regions associated with specific optimal strategies, we provide a robust framework for analysing risk and reward across environments. This enables us to extend the rule of thumb from (Gilliam & Fraser, 1987) to account for hazards associated with not consuming any resource. We demonstrate the power of this approach by applying it to different fitness paradigms (methods for accumulating fitness which corresponds to different life-history strategies), illuminating structural differences in how optimal strategies depend on how fitness accrues.

Our framework does not directly incorporate the notion of "state" common to SDP, such as tracking an internal level of satiation. In SDP, including state is computationally intensive, especially if several state variables are considered simultaneously, and often requires discretisation of continuous state variable values (Mangel et al., 1988). Recent innovations reduce the computational burden (Fackler, 2018; Nicol & Chadès, 2011) or increase the analytical tractability (Reimer et al., 2019a) of state-dependent SDP models. We circumvent this challenge by optimising foraging success independent of state, while still incorporating handling times, stochastic resource encounters and the nonlinear effects of hazard (predation risk).

However, strategy maps can be incorporated into other modelling methods which track state, as a method for determining how a forager should optimally behave. For example, a learning forager may attempt to behave optimally based on an internal information state (estimates of environmental parameters). The strategy map equations return the strategy the forager should follow based on its current estimates. Accomplishing this task using SDP would involve computing the optimal strategy recursively from the current time step to the final time step, at every time step. The strategy map can also measure how much the forager's estimates need to be updated before it switches to the optimal strategy (Arehart and Adler, In prep). Accomplishing this with SDP would require approximation, and involve repeated computations across many time steps.

Empirical work on animal foraging has produced mixed results (Pierce & Ollason, 1987; Pyke, 1984; Rechten et al., 1983; Stephens, 2008). One area which has proven challenging to integrate into theory is the role of learning in optimal foraging. Foragers never possess complete information about their environments and obtaining reliable information, especially in complex and changing environments, can be costly. A forager may arrive at an environment with little or no prior knowledge. Depending on the complexity and variability of the environment, it may take a lot of sampling for a forager to develop realistic estimates for the quality or availability of resources. Learning organisms must partition effort between increasing the accuracy of estimates and maximising benefits given those estimates. This is the explore-exploit tradeoff (Eliassen et al., 2007; Sutton & Barto, 2018), in which foragers trade some reward value for better information. Our approach places quantitative bounds on the value of information: Since strategies are discrete, forager's estimates of environmental parameters need only fall in the correct strategy region for the correct optimal strategy to be determined.

Our framework could be helpful for designing foraging experiments. Using estimates of quantities such as the quality of resources and handling times or patch departure times, experimentalists could identify regions of parameter space where strategy is expected to change. Quantitative measurements of GUD can be used to explicitly calculate a forager's perceived risk by inputting h^*, the experimentally derived GUD, into Equation 13 and solving for transitions between "strategy regions." Strategy transition thresholds can even be mapped explicitly onto a landscape, similar to constructing a Landscape of Fear (Laundré et al., 2010). This could open up new avenues for analysing the spatial distribution and movement ecology of species across landscapes.

As researchers have increasingly turned their attention towards mechanisms for making decisions (Budaev et al., 2019; Fawcett et al., 2013; Pierce & Ollason, 1987), a need has arisen for a bridge between calculating optimal strategies and methods by which animals may actually make decisions, especially with imperfect information. The strategy maps method offers a dynamic decision rule which may have some functional analogue in how foragers actually make decisions. These concepts could be tested experimentally by presenting a forager with a series of diet choice decisions or through patch depletion. By changing environmental variables over time and tracking the changes in a forager's strategy, it may be possible to uncover the dynamics of a forager's decision threshold. Finally, many foragers make decisions and exploit resources collectively (Falcón-Cortés et al., 2019). Strategy maps could be combined with information updating rules as well as resource dynamics (such as depletion of resources or increased handling times resulting from forager exploitation) to model the collective behaviour of foragers.

AUTHOR CONTRIBUTIONS

EA and FA conceptualised the research, EA, JR and FA conducted the research, and EA, JR and FA wrote the manuscript.

FUNDING INFORMATION

The Modeling the Dynamics of Life Fund

PEER REVIEW

The peer review history for this article is available at https://publons.com/publon/10.1111/ele.14160.

DATA AVAILABILITY STATEMENT

This paper does not use or create any empirical data.

ORCID

Emerson Arehart https://orcid.org/0000-0002-0862-4219
Jody R. Reimer https://orcid.org/0000-0001-7742-2728
Frederick R. Adler https://orcid.org/0000-0002-9022-3157

REFERENCES

Bedoya-Perez, M.A., Carthey, A.J., Mella, V.S., McArthur, C. & Banks, P.B. (2013) A practical guide to avoid giving up on giving-up densities. *Behavioral Ecology and Sociobiology*, 67(10), 1541–1553.

Bellman, R. (1957) A markovian decision process. *Journal of Mathematics and Mechanics*, 6, 679–684.

Bleicher, S.S. (2017) The landscape of fear conceptual framework: definition and review of current applications and misuses. *PeerJ*, 5, e3772.

Brodin, A., Nilsson, J.-Å. & Nord, A. (2017) Adaptive temperature regulation in the little bird in winter: predictions from a stochastic dynamic programming model. *Oecologia*, 185(1), 43–54.

Brown, J.S. (1988) Patch use as an indicator of habitat preference, predation risk, and competition. *Behavioral Ecology and Sociobiology*, 22(1), 37–47.

Brown, J.S., Laundré, J.W. & Gurung, M. (1999) The ecology of fear: optimal foraging, game theory, and trophic interactions. *Journal of Mammalogy*, 80(2), 385–399.

Budaev, S., Jørgensen, C., Mangel, M., Eliassen, S. & Giske, J. (2019) Decision-making from the animal perspective: bridging ecology and subjective cognition. *Frontiers in Ecology and Evolution*, 7, 164.

Caswell, H. (1980) On the equivalence of maximizing reproductive value and maximizing fitness. *Ecology*, 61(1), 19–24.

Charnov, E. (1976a) Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2), 129–136.

Charnov, E.L. (1976b) Optimal foraging: attack strategy of a mantid. *The American Naturalist*, 110(971), 141–151.

Eliassen, S., Jørgensen, C., Mangel, M. & Giske, J. (2007) Exploration or exploitation: life expectancy changes the value of learning in foraging strategies. *Oikos*, 116(3), 513–523.

Fackler, P.L. (2018) Stochastic dynamic programming without transition matrices. Technical Report, JSTOR.

Falcón-Cortés, A., Boyer, D. & Ramos-Fernández, G. (2019) Collective learning from individual experiences and information transfer during group foraging. *Journal of the Royal Society Interface*, 16(151), 20180803.

Fawcett, T.W., Hamblin, S. & Giraldeau, L.-A. (2013) Exposing the behavioral gambit: the evolution of learning and decision rules. *Behavioral Ecology*, 24(1), 2–11.

Gilliam, J.F. & Fraser, D.F. (1987) Habitat selection under predation hazard: test of a model with foraging minnows. *Ecology*, 68(6), 1856–1862.

Hirvonen, H. & Ranta, E. (1996) Within-bout dynamics of diet choice. *Behavioral Ecology*, 7(4), 494–500.

Holling, C.S. (1959) Some characteristics of simple types of predation and parasitism. *The Canadian Entomologist*, 91(7), 385–398.

Houston, A., Clark, C., McNamara, J. & Mangel, M. (1988) Dynamic models in behavioural and evolutionary ecology. *Nature*, 332(6159), 29–34.

Houston, A.I., McNamara, J.M. & Hutchinson, J.M. (1993) General results concerning the trade-off between gaining energy and avoiding predation. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 341(1298), 375–397.

Houtman, R. & Dill, L.M. (1998) The influence of predation risk on diet selectivity: a theoretical analysis. *Evolutionary Ecology*, 12(3), 251–262.

Iwasa, Y., Suzuki, Y. & Matsuda, H. (1984) Theory of oviposition strategy of parasitoids. i. Effect of mortality and limited egg number. *Theoretical Population Biology*, 26(2), 205–227.

Krebs, J. & Davies, N. (1978) *Behavioural ecology: an evolutionary approach oxford*. UK: Blackwell Scientific Publications.

Laundré, J.W., Hernández, L. & Ripple, W.J. (2010) The landscape of fear: ecological implications of being afraid. *The Open Ecology Journal*, 3(1), 1–7.

Lucas, J.R. (1983) The role of foraging time constraints and variable prey encounter in optimal diet choice. *The American Naturalist*, 122(2), 191–209.

MacArthur, R.H. & Pianka, E.R. (1966) On optimal use of a patchy environment. *The American Naturalist*, 100(916), 603–609.

Mangel, M. (2015) Stochastic dynamic programming illuminates the link between environment, physiology, and evolution. *Bulletin of Mathematical Biology*, 77(5), 857–877.

Mangel, M. & Clark, C.W. (1986) Towards a unifield foraging theory. *Ecology*, 67(5), 1127–1138.

Mangel, M. and Clark, C.W. (1988). *Dynamic modeling in behavioral ecology*, volume 63. Princeton: Princeton University Press.

Marescot, L., Chapron, G., Chades, I., Fackler, P.L., Duchamp, C., Marboutin, E. et al. (2013) Complex decisions made simple: a primer on stochastic dynamic programming. *Methods in Ecology and Evolution*, 4(9), 872–884.

McNamara, J.M. & Houston, A.I. (1986) The common currency for behavioral decisions. *The American Naturalist*, 127(3), 358–378.

Nicol, S. & Chadès, I. (2011) Beyond stochastic dynamic programming: a heuristic sampling method for optimizing conservation decisions in very large state spaces. *Methods in Ecology and Evolution*, 2(2), 221–228.

Nonacs, P. (2001) State dependent behavior and the marginal value theorem. *Behavioral Ecology*, 12(1), 71–83.

Pierce, G.J. & Ollason, J. (1987) Eight reasons why optimal foraging theory is a complete waste of time. *Oikos*, 49, 111–118.

Pirotta, E., Mangel, M., Costa, D.P., Goldbogen, J., Harwood, J., Hin, V. et al. (2019) Anthropogenic disturbance in a changing environment: modelling lifetime reproductive success to predict the consequences of multiple stressors on a migratory population. *Oikos*, 128(9), 1340–1357.

Pyke, G.H. (1984) Optimal foraging theory: a critical review. *Annual Review of Ecology and Systematics*, 15(1), 523–575.

Rechten, C., Avery, M. & Stevens, A. (1983) Optimal prey selection: why do great tits show partial preferences? *Animal Behaviour*, 31(2), 576–584.

Reimer, J.R., Mangel, M., Derocher, A.E. & Lewis, M.A. (2019a) Matrix methods for stochastic dynamic programming in ecology and evolutionary biology. *Methods in Ecology and Evolution*, 10(11), 1952–1961.

Reimer, J.R., Mangel, M., Derocher, A.E. & Lewis, M.A. (2019b) Modeling optimal responses and fitness consequences in a changing arctic. *Global Change Biology*, 25(10), 3450–3461.

Stephens, D.W. (2008) Decision ecology: foraging and the ecology of animal decision making. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4), 475–484.

Stephens, D.W., Brown, J.S. & Ydenberg, R.C. (2008) *Foraging: behavior and ecology*. Chicago: University of Chicago Press.

Sutton, R.S. & Barto, A.G. (2018) *Reinforcement learning: an introduction*. Cambridge: MIT press.

Yearsley, J., Tolkamp, B.J. & Illius, A.W. (2001) Theoretical developments in the study and prediction of food intake. *Proceedings of the Nutrition Society*, 60(1), 145–156.

How to cite this article: Arehart, E., Reimer, J.R. & Adler, F.R. (2023) Strategy maps: Generalised giving-up densities for optimal foraging. *Ecology Letters*, 00, 1–13. Available from: https://doi.org/10.1111/ele.14160

APPENDIX A

A.1 | FINDING v_2^* WITH NO HAZARD

In cases where $q_i = 1$ for all i, F_t does not converge and the method used to derive Equation (10) breaks down. To find v_2^* in cases with no hazard, we return to the original SDP formulation with backward iteration, and identify v_2^* as the transition value at which the max condition in the second term of (1) turns into an equality:

$$v_2^* + F_{t+h_2} = F_{t+1}, \tag{20}$$

where F_t equals rewards accrued from time t through the final time step T, rather than from the beginning. We assume $v_1 = 1$ without loss of generality, normalising the value of v_2. We use the approximation $\alpha t = F_t$, where α is a constant for the per-unit-time return of a given strategy. Then $F_{t+h_i} = \alpha\left(t + h_i\right)$, and (2) becomes:

$$\begin{aligned}\alpha t = \lambda_1\left(1-\lambda_2\right)\max\Big(\left(1+\alpha\left(t+h_1\right)\right)q_1^{h_1}, \alpha(t+1)q_0\Big)+\\ \lambda_2\left(1-\lambda_1\right)\max\Big(\left(v_2^*+\alpha\left(t+h_2\right)\right)q_2^{h_2}, \alpha(t+1)q_0\Big)+\\ \lambda_1\lambda_2\max\Big(\left(v_1+\alpha\left(t+h_1\right)\right)q_1^{h_1}, \left(v_2^*+\alpha\left(t+h_2\right)\right)q_2^{h_2}, \alpha(t+1)q_0\Big)+\\ \left(1-\lambda_1\right)\left(1-\lambda_2\right)\alpha(t+1)q_0.\end{aligned} \tag{21}$$

Solving the system of Equations (20) and (21) gives

$$v_2^* = \frac{\lambda_1\left(h_2 - 1\right)}{1 + \lambda_1\left(h_1 - 1\right)}. \tag{22}$$

This is exactly the result obtained by discretising Holling's Disc Equation, which relates the role of handling time (h_i) and the encounter rate of the preferred resource (λ_1) to an optimal foraging strategy (Charnov, 1976b; Holling, 1959). The expression for v_2^* does not include λ_2, which also agrees with previous findings that the criterion for including a resource depends only upon the resources better than it, rather than upon the availability of the resource itself (Lucas, 1983; MacArthur & Pianka, 1966; Nonacs, 2001).

APPENDIX B

B.1 | COMPUTING v_2^* WITH CONVERGENT APPROXIMATION

It is possible to solve for Φ_{EGGS} and Φ_{AON} exactly for finite T (i.e., not taking the limit as $t \to -\infty$), using variations of equation (4). However, we use an approximation which is accurate for realistic choices of parameter values, and allows for analytical sensitivity analysis. To compare the differences between fitness paradigms in finite time, we derive time-dependent expressions to approximate the total fitness functions. Our approximations make the following two simplifying assumptions:

1. The forager may initiate the foraging bout in any action state with equal probability. States include encountering a resource (or no resources) or being at any of the h_i time steps during the handling time for resource i.
2. Whether or not the forager survives a given time step is determined by an independent random variable based on the average survival for a strategy. This makes handling time geometrically distributed, an approximation commonly employed in optimal foraging models.

Incorporating these assumptions eliminates the artefacts mentioned above that occur when foraging initiates or approaches termination, but still captures the overall behaviour of the system. We define p_i = probability of being in state i, where state i = consuming resource i, with $i = 0$ corresponding to not consuming a resource. To compute p_i, we first construct A_Z, the Markov transition matrix between states 0, 1 and 2, conditional on being alive ($q_i = 1$ for all i), corresponding to strategy Z. For example, $A_{\{1,2\}}$ is

$$A_{\{1,2\}} = \begin{pmatrix} (1-\lambda_1)(1-\lambda_2) & \frac{1}{h_1}(1-\lambda_1)(1-\lambda_2) & \frac{1}{h_2}(1-\lambda_1)(1-\lambda_2) \\ \lambda_1 & \frac{1}{h_1}(h_1+\lambda_1-1) & \frac{\lambda_1}{h_2} \\ \lambda_2(1-\lambda_1) & \frac{1}{h_1}\lambda_2(1-\lambda_1) & \frac{1}{h_2}(h_2+\lambda_2(1-\lambda_1)-1) \end{pmatrix} \tag{23}$$

where each column sums to 1. The stationary distribution of the system for strategy $\{1,2\}$, given by the right eigenvector of (23) corresponding to the eigenvalue of 1, is

$$p_0 = \frac{(1-\lambda_1)(1-\lambda_2)}{h_1\lambda_1 + h_2\lambda_2(1-\lambda_1) + (1-\lambda_1)(1-\lambda_2)} \tag{24}$$

$$p_1 = \frac{h_1\lambda_1}{h_1\lambda_1 + h_2\lambda_2(1-\lambda_1) + (1-\lambda_1)(1-\lambda_2)} \tag{25}$$

$$p_2 = \frac{h_2\lambda_2(1-\lambda_1)}{h_1\lambda_1 + h_2\lambda_2(1-\lambda_1) + (1-\lambda_1)(1-\lambda_2)}. \tag{26}$$

With no hazard ($q_i = 1$), the p_i sum to 1. For $q_i \neq 1$, the system becomes an absorbing Markov process, so finding the steady state is not informative (the steady state is death). We perturb A to find the effect of q_i on the system, following Caswell (1980). Starting with the standard eigenvalue relation $Ap = eA$, in which p is the right eigenvector of A and e is the corresponding dominant eigenvalue (we use nonstandard notation here because λ and v are already in use), we formally differentiate and left multiply by $w^\top$, where $w^\top$ is the transpose of $w = <1,1,1>$:

$$w^\top(\Delta A)p + w^\top A(\Delta p) = ew^\top(\Delta p) + w^\top p(\Delta e). \tag{27}$$

Solving for the first-order sensitivity, we obtain:

$$\Delta e = \frac{w^\top(\Delta A)}{w^\top p} = \frac{-\left(\Delta_0(1-\lambda_1)(1-\lambda_2)\right) + \Delta_1 h_1\lambda_1 + \Delta_2 h_2\lambda_2(1-\lambda_1)}{h_1\lambda_1 + h_2\lambda_2(1-\lambda_1) + (1-\lambda_1)(1-\lambda_2)}, \tag{28}$$

where $\Delta_i = 1 - q_i$. The similarity between this expression and the right eigenvalues (24–26) suggests our formulation for P_Z, which approximates the probability of survival per unit time:

$$P_Z = \sum_{i \in Z} p_i q_i, \tag{29}$$

the probability of being in each state included in Z multiplied by the survival per time step for that state. Because the p_i depend on strategy, so does survival. p_i and P_Z for other strategies are constructed in the same manner. This is a first-order approximation, and does not capture the correlational structure of the actual foraging process, in which taking a resource i requires a full h_i time steps. However, it captures the overall survival well when the h_i are small relative to the timescale of interest, and for q_i sufficiently close to 1 for Δe to remain small. To approximate the fitness gained over time, we compute α_Z, the expected per-time step payoff of strategy Z conditional on being alive:

$$\alpha_Z = \sum_{i \in Z} \frac{p_i v_i}{h_i}. \tag{30}$$

This is the probability of being in state i, p_i, multiplied by state i's payoff per time step, $\frac{v_i}{h_i}$. Substituting into the total fitness function from (18),

$$\Phi_{EGGS}(Z,T) = \sum_{t=0}^{T} t\alpha_Z (P_Z)^t = t\alpha \frac{1-P_Z^{T+1}}{1-P_Z} \tag{31}$$

as the partial sum of the geometric series. We set $\Phi_{EGGS}(\{1,2\},T) = \Phi_{EGGS}(\{1\},T)$ and solve for v_2 (which appears in α_Z) to obtain:

$$v^*_{2,EGGS} = \frac{h_2}{h_1 p_{2,\{1,2\}}} \left(p_{1,\{1\}} \frac{\left(1-P_{\{1\}}^{T+1}\right)\left(1-P_{\{1,2\}}\right)}{\left(1-P_{\{1,2\}}^{T+1}\right)\left(1-P_{\{1\}}\right)} - p_{1,\{1,2\}} \right), \tag{32}$$

where the subscript in $p_{i,Z}$ is p_i computed for strategy Z. For the AON paradigm, we substitute the values for P_Z and α_Z into (19), yielding

$$\Phi_{AON}(Z,T) = T\alpha_Z P_Z^T. \tag{33}$$

Solving for v_2^*:

$$v^*_{2,AON} = \frac{h_2}{h_1 p_{2,\{1,2\}}}\left(\left(\frac{P_{\{1\}}}{P_{\{1,2\}}}\right)^T - p_{1,\{1,2\}}\right). \tag{34}$$

The PRE paradigm is identical to AON, but applied to the foraging bout length τ instead of the total foraging time. To find $\tilde{v}$, set the expressions for strategies $\{1,2\}$ and $\{2,1\}$ equal, and solve for $\frac{v_2}{v_1}$.

APPENDIX C

C.1 | MODEL SETUP

Examples of the foraging process. Foragers encounter multiple, single or no resources in discrete time. If a forager accepts resource i, it receives reward v_i and advances h_i time steps, surviving with probability $q_i^{h_i}$. On the top row, a forager follows the strategy $\{2,1\}$, accepting yellow when available but taking blue if it appears with no yellow. The forager accepts blue (resource type 1) on the first time step, proceeding h_1 time steps while handling the resource, then encounters a yellow (resource type 2), proceeding h_2 time steps to find another yellow, and so on. On the bottom row, a forager follows strategy $\{2\}$, only taking yellow; the forager skips the first blue flower, then accepts a yellow, proceeding h_2 timesteps, where it encounters *no* resources, so must wait until another yellow appears. This forager is unlucky and does not survive past this resource.

C.2 | THE FORAGING ALGORITHM

1. At time t, encounter each resource i with probability λ_i.
2. If resources are encountered, either:
   a. Choose one resource to consume according to strategy Z. Survive with probability $q_i^{h_i}$ and advance h_i time steps. Total resources collected increase by v_i
   b. Forego all resources and survive with probability q_0 and advance 1 time step.
3. If no resources are encountered, survive with probability q_0 and advance 1 time step.
4. Repeat until $t = T$ or the forager dies.