Navigating the Landscape of Fear

Marissa Gee[1], Nicolas Gonzalez-Granda[2], Sunay Joshi[3], Nagaprasad Rudrapatna[4], Anne Somalwar[5], Stephen P. Ellner[6], and Alexander Vladimirsky[7]

[1]Department of Mathematics and Statistics, Kenyon College, Gambier, OH, gee1@kenyon.edu
[2]Division of Biostatistics and Health Data Science, School of Public Health, University of Minnesota, Minneapolis, MN, gonz0796@umn.edu
[3]Graduate Group in Applied Mathematics and Computational Science, University of Pennsylvania, Philadelphia, PA, sunayj@sas.upenn.edu
[4]Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX, nagaprasad.rudrapatna@alumni.duke.edu
[5]Graduate Group in Applied Mathematics and Computational Science, University of Pennsylvania, Philadelphia, PA, somalwar@sas.upenn.edu
[6]Department of Ecology and Evolutionary Biology, Cornell University, Ithaca, NY, spe2@cornell.edu
[7]Department of Mathematics, Cornell University, Ithaca, NY, vladimirsky@cornell.edu

August 15, 2024

Abstract

Animals searching for food must navigate through complex landscapes with varying terrain, food availability, predator activity, and shelter. One method for analyzing how foraging animals balance these competing interests is Optimal Foraging Theory, which posits that foragers make decisions that maximize some expected reward or utility. However, models have generally focused on simplified settings, especially regarding animal movement decisions, either ignoring spatial variability entirely or limiting the forager's options to choosing among a few discrete habitat patches (or patch types) that are modeled as spatially uniform. In this paper, we present a model of optimal foraging in a continuous landscape for an animal that is subject to predation. Foraging animals thus must choose not only where to gather food, but also how to (safely) travel across the landscape. Furthermore, we explicitly model stochastic predator interactions, allowing us to predict optimal foraging trajectories conditional on the presence or absence of an immediate threat from a predator. We illustrate our model with numerical examples, one a hypothetical landscape with two regions of high food abundance and two areas of high predation risk on the direct route between the feeding areas and the forager's overnight refuge, the other inspired by empirical data on foraging Samango monkeys (*Cercopithecus albogularis schwarzi*). We find that

the shape of a forager's utility function affects not only where it chooses to feed, but also the paths it takes to and from the optimal feeding ground. Thus, examining predicted optimal trajectories can provide additional information about what quantity, if any, animals optimize while foraging. We also develop and demonstrate a preliminary model for an animal that depletes the food supply in its local environment.

1 Introduction

Foraging animals make decisions every day that impact their chances of survival and reproductive success. Where, when, and for how long should they gather food? To what extent should they engage in anti-predator behaviors such as vigilance or seeking refuge if a predator is detected? These decisions depend on balancing the animal's short-term and long-term goals, and on the information available about the food abundance and predator density at different locations. Optimal Foraging Theory (OFT), [Stephens and Krebs, 1987] seeks to understand these decisions by viewing animals as rational decision makers, and interpreting observed choices of diet and foraging locations as the best among the available options for maximizing the organism's fitness (transmission of genes to subsequent generations). Understanding the basis for foraging decisions can shed light on how animals evaluate and trade off between conflicting goals when foraging, inform the design of experiments to test inferences about the basis for foraging decisions, and provide predictions of how behavior may change as a result of a changing environment [Reimer et al., 2019].

Early work in OFT assumed that a forager seeks to maximize its net rate of energy gain [Stephens and Krebs, 1987]. These models examined questions such as how a forager should decide which food items to eat (or attempt to eat) if they are encountered, and how it should exploit different food patches in a patchy environment (see [MacArthur and Pianka, 1966], [Emlen, 1966], [Charnov, 1976]). While early rate-maximization models focused entirely on energy gain, models of risk-sensitive foraging (e.g., [Caraco, 1980, Stephens, 1981, Houston and McNamara, 1982, Lima and Dill, 1990, Werner and Gilliam, 1984, Gilliam and Fraser, 1987] were able to incorporate the mortality risk associated with foraging decisions. However, rate-maximization theories cannot reflect how an animal's priorities might change with its *state* (e.g., its physical location, or how close

it is to starvation).

To address these limitations, [Mangel and Clark, 1986, McNamara and Houston, 1986] proposed models that explicitly track the forager's state, and include randomness in both food acquisition and possible death through predation. These models assume that a forager seeks to maximize the expected reward accrued during a modeled time period, where the reward encodes reproductive success during the time period, plus expected future success based on the forager's state at the final time. State-dependent optimal decisions are found numerically through stochastic dynamic programming (SDP), versions of which were later used to analyze many theoretical questions about animal behavior [Newman, 1991], [Houston et al., 1993]. The SDP framework has seen numerous empirical applications including birds deciding whether to flock or find a mate [Szekely et al., 1991], tuna foraging near ocean fronts [Kirby et al., 2000], and polar bears preparing for hibernation and reproduction [Reimer et al., 2019]. However, in almost all cases the forager's current energy reserve level is the only state variable. The choice of the next food item or feeding area does not take into account the animal's current location, or what path it will take to its chosen target.

In this paper, we take a first step in generalizing the state-based dynamic programming approach of classical optimal foraging theory in order to characterize optimal paths for an animal foraging in a spatially continuous habitat. That is, we try to characterize how a forager should navigate through a *landscape of fear*. The term "landscape of fear" has been used in multiple ways, but it generally refers to the spatial distribution of predation risk, real or perceived, throughout a habitat (e.g., [Brown, 1988, Brown et al., 1999, Pinti et al., 2022, Laundré et al., 2001, Gaynor et al., 2019]). Studies have examined how the introduction of predators affects the spatial distribution of prey animals [Laundré et al., 2001], [Raynor et al., 2021], spatio-temporal distributions of prey [Kohl et al., 2018], and the relative impact of predation and food availability on utilization of the environment [Coleman and Hill, 2014]. Optimal foraging theory is potentially well-posed to contribute to the analysis of these data, but must first overcome the challenges associated with finding optimal decisions about movement in a continuous habitat where food availability and predation risk are both spatially varying.

Some previous optimal foraging models sought to incorporate spatial dynamics, though with signif-

icant simplifications. [Rands, 2017] considers a forager collecting food and avoiding a predator in a discretized 1-dimensional habitat. Interactions between the forager and a predator are modeled explicitly, and the forager must decide not only when to feed and when to seek refuge, but also where to position itself when it is not optimal to do either of those things. [Fiksen et al., 1995] proposed a spatially explicit optimal foraging model of capelin in the Barents sea. The authors estimate food availability, temperature, and mortality due to predation for a complex environment to predict how the distribution of capelin will evolve in time. However, they focus on a timescale of months to years, and thus do not attempt to recover detailed capelin trajectories nor capture any short-term changes in their behavior induced by predator encounters. In contrast, we seek to solve a spatially explicit optimization problem in a continuous habitat with enough temporal resolution to characterize the specific paths taken by a forager and abrupt changes to those paths when they are spotted by a predator.

In all SDP-style ecological models, one of the vexing questions is the appropriate optimization criterion. One critique of early optimal foraging models was that the maximization of energy rate gain is not always a reasonable proxy for reproductive success [Stephens, 1981]. Previous work has shown that the shape of the utility/reward function can impact foraging behavior [Newman, 1991]. Here, we do not specify a single reward function but instead contrast foraging behaviors for several different functional forms of the reward.

In the following section, we propose a model of a forager's movement and energy dynamics based on differential equations, and a stochastic model of the impact of predators. We assume that the forager seeks to maximize an expected terminal utility and present three possible functional forms that encode different risk preferences. Treating this as an optimal control problem, we derive a set of Hamilton-Jacobi-Bellman partial differential equations that we solve numerically to find the forager's optimal policy. We use our model to analyze two examples: one mimicking the risk/reward trade-offs in an environment with two main feeding grounds and another inspired by an empirical study of foraging Samango monkeys in the Soutpansberg Mountains in South Africa [Coleman and Hill, 2014]. Finally, we introduce and discuss an extension of our model addressing the impact of the forager depleting food as it exploits its environment, and include an example illustrating this model in the

Parameter	Description
$\mu_s(\mathbf{x}, e, t)$	Rate at which the forager is spotted by predators
$\mu_k(\mathbf{x}, e, t)$	Rate at which a chasing predator kills the forager
$\mu_g(\mathbf{x}, e, t)$	Rate at which a predator gives up on a chase (equivalently, the forager escapes)
$K_m(\mathbf{x}, e, t)$	Metabolic cost
$f_m(\mathbf{x}, e, t)$	Speed of travel
$\mathbf{a}_m(\mathbf{x}, e, t)$	Direction of motion as a feedback control
$\psi(\mathbf{x}, t)$	Underlying food density in the environment
$F(\mathbf{x}, t)$	Food availability (rate of energy accumulation from foraging)
$U(e)$	Utility representing future reproductive success
Ω	Domain
E	Forager's maximum energy reserve
T	Planning horizon

Table 1: Model parameters. We formulate a continuous model of position and energy and assume the forager selects its direction of motion in order to maximize its expected terminal utility (see Figure 1b). Food availability is computed based on a weighted average of the food density near the forager's current location, and can be thought of as the rate of energy intake. The subscript m represents the forager's possible operating modes (Figure 1a) and can take the values 1 or 2. We allow the metabolic costs and speed of travel to be mode-dependent; the animal's full state (including mode) also influences its selection of the optimal direction of motion. As a result, the evolution of the forager's energy and position is strongly impacted by the random mode switches.

Supplementary Materials, Section C. While the model we present is not sophisticated enough to provide quantitative predictions about foraging animals, it can provide qualitative insight into what behavior would arise from optimizing different utility functions. This in turn can inform the design of experiments or observational studies that seek to infer what (if any) utility function a forager is maximizing.

2 Model

Forager Dynamics

We consider a foraging animal subject to predation and moving through a continuous, two-dimensional domain Ω. We assume that the forager plans up until a finite, positive time T, which we call the planning horizon. Reasonable planning horizons will depend on the problem being studied, but we will generally choose T to be on the order of one day or one night. One reason for this is that we explicitly model spatial interactions between predators and prey through chases, with each chase lasting on the order of minutes. Thus, timescales on the order of seasons, while theoretically possible,

will not generally be computationally tractable for our model. We track the forager's location in Ω and its current energy level as functions of time, denoted $\mathbf{y}(t)$ and $\xi(t)$ respectively.

As the forager moves through the environment, there is a chance that it is spotted by a predator. We model this as a random event occurring at a rate μ_s, i.e., the probability of the forager being spotted in a small time interval $[t, t+\tau]$ is approximately $\mu_s \tau$. If a forager is spotted, we assume it is not killed instantly; instead, a chase ensues in which the forager may be killed or may escape. Both of these alternatives also occur at random times and are governed by rates μ_k and μ_g respectively. We encode the current status of the forager relative to predators by assigning it one of two "operating modes:" foraging (mode 1) and escaping (mode 2). The current mode can thus be thought of as the state of a continuous-time Markov chain, illustrated in Figure 1a (where mode 3 is a dummy mode that represents the case that the forager has died). Further, the forager's complete state at time t can be viewed as $(\mathbf{y}(t), \xi(t), \eta(t))$ where $\eta(t)$ is a random variable that takes values in $\{1, 2\}$ and encodes the current operating mode. More generally, there is no reason why the rates must be constant, so μ_s, μ_g, and μ_k can be functions of the forager's full state and time, though we only consider location-dependence for our examples.

Since the operating mode impacts how the forager interacts with its environment, we must equip each mode with its own location and energy dynamics. While the transitions between these modes are stochastic, we assume that the dynamics within each mode are deterministic, meaning that this model falls within the framework of piecewise-deterministic Markov processes (PDMPs) [Davis, 1984]. We use a simplified model of motion wherein the forager's speed is determined by its position, and it can choose any direction to move in instantaneously.[1] This yields the following governing ODE for the animal's position:

$$\frac{d\mathbf{y}}{dt} = f_{\eta(t)}\left(\mathbf{y}(t), \xi(t), t\right)\mathbf{a}, \qquad t \in [0, T] \tag{1}$$

$$\mathbf{y}(0) = \mathbf{x}, \qquad \mathbf{x} \in \Omega \tag{2}$$

[1] These are standard *isotropic* dynamics and can represent the dependence of speed on vegetation density or other location-based terrain characteristics. Much more general anisotropic dynamics (e.g., to model the dependence of speed on the terrain slope in the chosen direction) can be handled similarly, yielding more general Hamilton-Jacobi-Bellman PDEs. While the general computational approach remains the same, it would require more data to calibrate the model and we choose this simpler setting to highlight the main ideas.

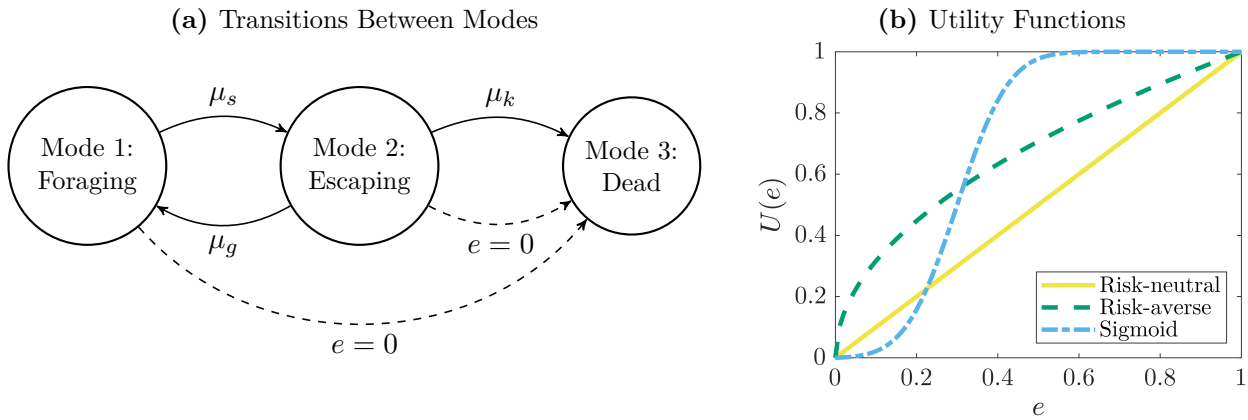

Figure 1: A forager maximizes its *expected* terminal utility due to random interactions with predators. (a) The forager's current mode can be foraging (mode 1, actively collecting food), escaping (mode 2, moving in the immediate presence of a predator), or dead (mode 3). Each mode specifies its own dynamics for energy and position. Transitions between modes are stochastic and governed by known rates: μ_s, the rate at which the forager is spotted by a predator; μ_g, the rate at which a pursuing predator gives up on a chase; and μ_k, the rate at which a pursuing predator successfully kills the forager. In general these rates may depend on time, the forager's current energy level, and its current location, though our examples in the Results section consider location-dependence only. The forager deterministically transitions to mode 3 if it ever starves (reaches $e = 0$). **(b)** Possible functional forms for a utility function representing future fitness. Although utility is only a function of the forager's final energy, the trajectory-dependence of the final state means that this function affects the optimal behavior at all times. Additionally, each utility function can be thought of as encoding specific risk preferences, discussed in detail in Section Utility Maximization.

where $f_m(\mathbf{x}, e, t)$ is a mode-dependent scalar speed function and $\mathbf{a}$ is a unit vector describing the chosen direction of motion. When we view this as an optimization problem, $\mathbf{a}$ will be the feedback control, i.e., the animal will select a travel direction as a function of its current state (position, energy, and mode) and time.

The forager continuously expends energy at a mode-dependent rate $K_m(\mathbf{x}, e, t)$ and gains energy by collecting food from the environment at some rate $F(\mathbf{x}, t)$. We assume that there is an underlying density of food $\psi(\mathbf{x}, t)$ in calories per unit area, and compute $F(\mathbf{x}, t)$ based on a weighted average of the available food in a small circle around the forager's current location (see the Supplementary Materials, Section Calculating Food Availability for details). Since we only allow the prey animal to

collect food while actively foraging in mode 1, the mode-dependent energy dynamics are

$$\frac{d\xi}{dt} = \begin{cases} \mathbb{1}_{[\eta(t)=1]} F(\mathbf{y}(t), t) - K_{\eta(t)}\left(\mathbf{y}(t), \xi(t), t\right), & \text{if } \xi(t) \in (0, E), \\ \min\left\{\mathbb{1}_{[\eta(t)=1]} F(\mathbf{y}(t), t) - K_{\eta(t)}(\mathbf{y}(t), \xi(t), t),\ 0\right\}, & \text{if } \xi(t) = E, \\ 0, & \text{if } \xi(t) = 0, \end{cases} \qquad t \in [0, T], \tag{3}$$

$$\xi(0) = e, \qquad e \in [0, E]. \tag{4}$$

Here $\mathbb{1}_{[\eta(t)=1]}$ represents the characteristic function that is 1 if the forager is in mode 1 and 0 otherwise. The second case ensures that the forager never exceeds its maximum attainable energy level E, and the third case ensures that once the energy reaches 0, the forager dies of starvation and can no longer gain energy.

Utility Maximization

We assume that the forager seeks to maximize a terminal reward, in our case the *expected* utility of its final energy level. The expectation accounts for the fact that interactions with the predator, and thus the forager's final state, are stochastic. The most common interpretation of the terminal reward is as a measure of or proxy for future reproductive success (e.g., *fitness*) as a function of an animal's state [McNamara and Houston, 1986]. While not a requirement, many fitness functions depend on the animal's energy level. Some functional forms that have been used in models of optimal foraging are step-functions [Rands, 2017], linear functions [Houston et al., 1993], [Fiksen et al., 1995], various concave functional forms [Houston et al., 1993], [Sirot, 2019], sigmoid curves [Houston et al., 1993], [Caraco and Chasin, 1984], and occasionally convex functions [Kirby et al., 2000]. Additionally, [Brown and Kotler, 2004] propose various measures of fitness based on a combination of survival probability and energy gained from foraging.

We do not prescribe a single utility function, but instead consider three possible functional forms for the utility, and study their impact on optimal foraging behavior.[2] We generally assume that

[2] [Newman, 1991] also examines the impact of the shape of a utility function on foraging behavior, though for a discrete patch-based model with depletable resources.

an animal earns utility 0 if it is eaten. Throughout the paper, we compare three commonly used functional forms for utility shown in Figure 1b and given by

$$U_{\text{linear}}(e) = \sigma \frac{e}{E} \tag{5}$$

$$U_{\text{concave}}(e) = \sigma \sqrt{\frac{e}{E}} \tag{6}$$

$$U_{\text{sigmoid}}(e) = \frac{\text{erf}\left(\frac{e-e_t}{\sqrt{2}\sigma}\right) - \text{erf}\left(\frac{-e_t}{\sqrt{2}\sigma}\right)}{\text{erf}\left(\frac{E-e_t}{\sqrt{2}\sigma}\right) - \text{erf}\left(\frac{-e_t}{\sqrt{2}\sigma}\right)} \tag{7}$$

where e_t is an energy threshold and σ is a shape parameter. The utility functions we consider can further be thought of as encoding risk preferences, where a risk-neutral planner has a linear utility (i.e., has no preference between two reward distributions having the same mean reward) and a risk-averse planner has concave utility (i.e., prefers distributions with lower variance, all else being equal). The sigmoid utility represents a planner that is risk-seeking at low expected utility and risk-averse at high expected utility. For steep curves (low values of σ) this approximates a step-function, the utility function associated with maximizing the probability of ending above a specific energy level.

We use a dynamic programming approach to compute the policy that maximizes the expected utility. To do this, we assign each mode a value function $u_m(\mathbf{x}, e, t)$ that represents the optimal expected reward-to-go for a forager in state $(\mathbf{x}, e)$ and mode m at time t. Formally, we define this as

$$u_m(\mathbf{x}, e, t) = \sup_{\mathbf{a}(\cdot)\in\mathcal{A}} \mathbb{E}(U(e(T))). \tag{8}$$

It can be shown that these value functions must solve a coupled system of Hamilton-Jacobi-Bellman PDEs (see the Supplementary Materials, Section Deriving the Governing HJB Equations for a detailed formal derivation). Each term in each PDE describes a part of how the optimal expected utility changes with respect to time and the forager's state. In both cases, all changes must sum to

zero.

$$\underbrace{\frac{\partial u_1}{\partial t}}_{\substack{\text{Change in}\\ \text{time}}} + \underbrace{\left([F-K_1]_*\right)\frac{\partial u_1}{\partial e}}_{\substack{\text{Change due to}\\ \text{gained/lost energy}}} + \underbrace{f_1|\boldsymbol{\nabla}_{\mathbf{x}} u_1|}_{\substack{\text{Change due to moving}\\ \text{in the optimal direction}}} + \underbrace{\mu_s(u_2-u_1)}_{\text{Change if spotted}} = 0 \tag{9}$$

$$\underbrace{\frac{\partial u_2}{\partial t}}_{\substack{\text{Change in}\\ \text{time}}} - \underbrace{K_2\frac{\partial u_2}{\partial e}}_{\substack{\text{Change due to}\\ \text{lost energy}}} + \underbrace{f_2|\boldsymbol{\nabla}_{\mathbf{x}} u_2|}_{\substack{\text{Change due to moving}\\ \text{in the optimal direction}}} + \underbrace{\mu_g(u_1-u_2)}_{\substack{\text{Change if}\\ \text{escaped}}} - \underbrace{\mu_k u_2}_{\substack{\text{Change if}\\ \text{killed}}} = 0. \tag{10}$$

We use $[F-K_1]_*$ to denote the quantity that is $F-K_1$ when the forager is not at its maximum energy, and $[F-K_1]_- = \min(F-K_1, 0)$ when it is. We impose the boundary and terminal conditions

$$u_m(\mathbf{x}, 0, t) = 0, \tag{11}$$

$$u_m(\mathbf{x}, e, T) = U(e), \tag{12}$$

which ensure that the reward is 0 if the animal starves and $U(e)$ if it survives the entire planning horizon. We solve this system backwards in time using an explicit finite differences discretization (outlined in detail in the Supplementary Materials, Section Numerical Methods).

3 Results

Balancing Risk and Reward

For our first example, we consider a forager that faces a trade-off between risk and reward. The environment is assumed to contain two food-rich regions (similar in structure to food patches in a discrete model), with the upper patch containing more food than the lower (Figure 2a). We further assume that the forager begins and ends its foraging bout within a refuge (outlined in white), which represents a burrow, den, or other area of increased cover. When in the refuge, the forager cannot be spotted or killed by a predator ($\mu_s(\mathbf{x}, e, t) = \mu_k(\mathbf{x}, e, t) = 0$) and has an increased chance of

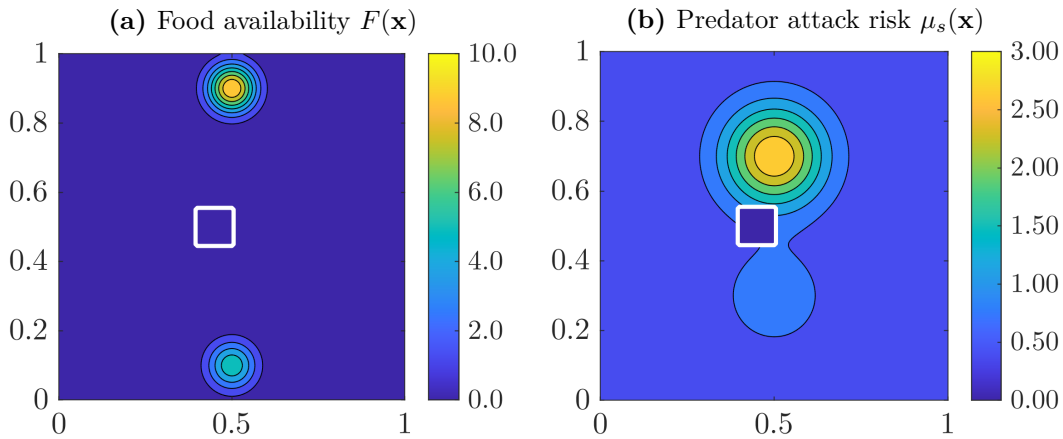

Figure 2: An inhomogeneous continuous environment similar in spirit to a discrete two-patch model. Food is concentrated in two small regions of the domain, where the upper is more lucrative (see the food availability $F(\mathbf{x})$ in subfigure **(a)**) but also more risky (see the predator attack risk $\mu_s(\mathbf{x})$ in subfigure **(b)**). In both panels, the white square represents the refuge, where the foraging bout must begin and end and where the forager can hide from predators without any risk but cannot collect food. Exact expressions for each quantity are provided in the Supplementary Materials E. The areas of highest risk lie along the most direct routes from the refuge to each feeding ground. This impacts the paths selected by a rational forager even when no predators are currently present. Both the choice of the feeding ground and the extent of a risk-avoidance detour that the forager is willing to take depend non-trivially on the initial energy, the planning horizon, the level of "risk premium" ρ (differentiating the upper and lower halves of the domain and defined formally in the text), and the forager's utility function. These dependencies are further illustrated in Figures 3 and 4.

escaping (μ_g is increased by a factor of 10), but also cannot collect any food ($F(\mathbf{x}, t) = 0$). Outside of the refuge, we choose the spotting rate $\mu_s(\mathbf{x}, e, t)$ so that the direct path to each food patch passes through an area of elevated spotting risk (Figure 2b). We will generally assume that this risk is higher along the path to the better food patch, with the animal aware of this *risk premium* ρ when it chooses which patch to visit.[3] We use a constant kill rate $\mu_k(\mathbf{x}, e, t) = 4.0$, give-up rate $\mu_g(\mathbf{x}, e, t) = 1.0$, speed $f_m(\mathbf{x}, e, t) = 2.0$, and metabolic cost $K_m(\mathbf{x}, e, t) = 3.0$. We allow the animal to reach a maximum energy level of $E = 10$ and assume the foraging bout lasts until time $T = 1$. Combined with our choice of K_m, this implies that any animal that begins at an energy level $e_0 > 3$ will not starve, even if it does not collect any food.

We focus on three questions: first, if the animal chooses to forage instead of staying in the safety of its refuge, which patch should it head to? Second, what is the optimal path to and from the

[3]More formally, we represent each region of elevated risk using a Gaussian function, and define the "risk premium" ρ to be the difference in their peak values. See Supplementary Materials E for detailed expressions.

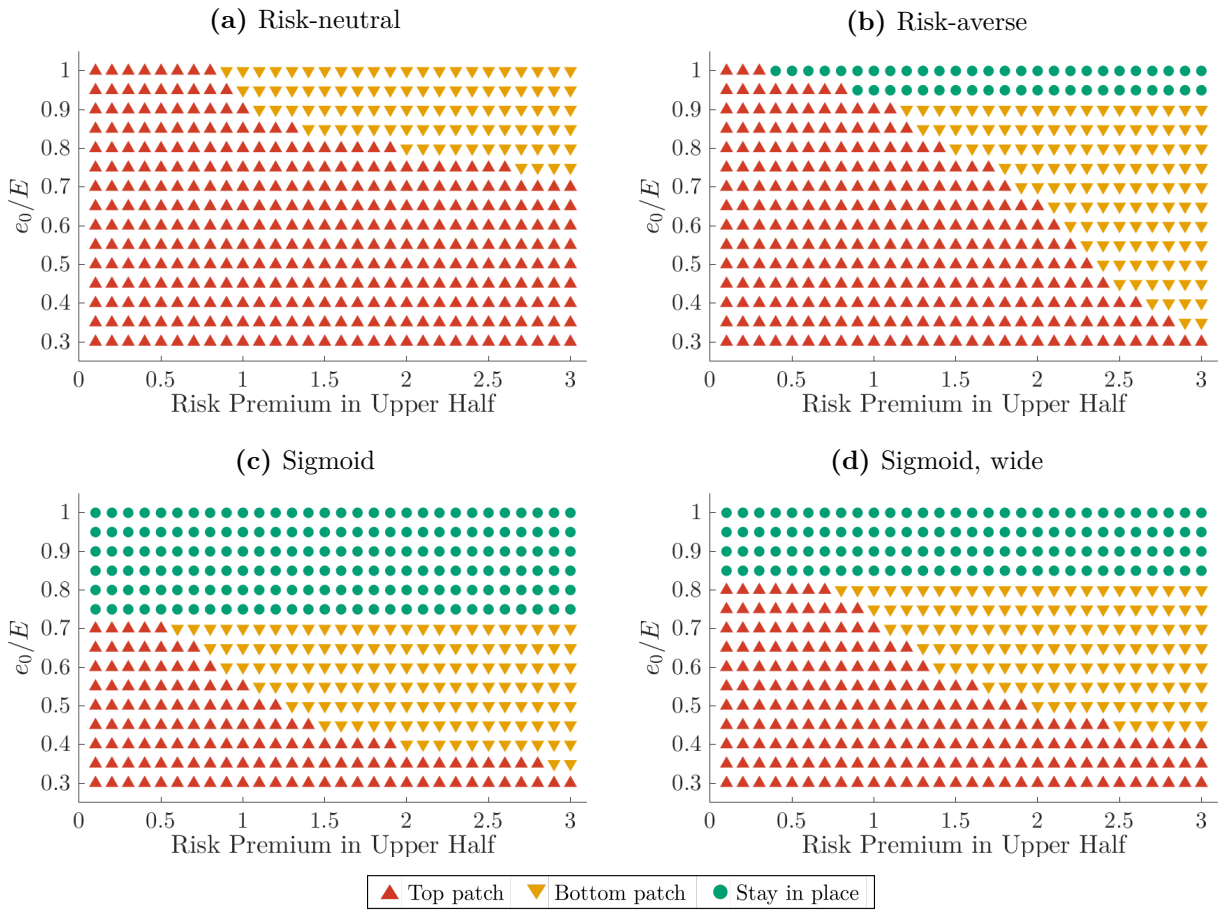

Figure 3: Preferred feeding ground for variable initial energy, risk premium, and utility function. Patch choice represents which food rich region the forager initially plans to visit (this plan may change due to the appearance of a predator). Across all utility functions, the forager is more likely to choose the riskier option when its initial energy is low (it is close to starvation), though this is mitigated by the level of the risk premium. There are significant regions of parameter space for which more than one utility function leads to the same patch choice, e.g., at low values of the initial energy and risk premium or for very high values of the risk premium. In these cases, it is natural to ask whether other aspects of the forager's behavior differ across utility functions as this has implications for experiments that seek to infer risk preference from behavior. Figure 4 examines the forager's planned *trajectories* for a range of initial energy levels.

chosen patch? Third, how does the optimal behavior depend on the shape of the forager's utility function? While the first of these questions could potentially be answered using a "traditional" (discrete) foraging model, the other two questions leverage the continuity of our model to extract additional information about the optimal behavior.

Figure 3 summarizes how the preferred foraging patch changes as we vary initial energy level and risk premium for four possible utility functions—risk-neutral, risk-averse, and two versions of a sigmoid

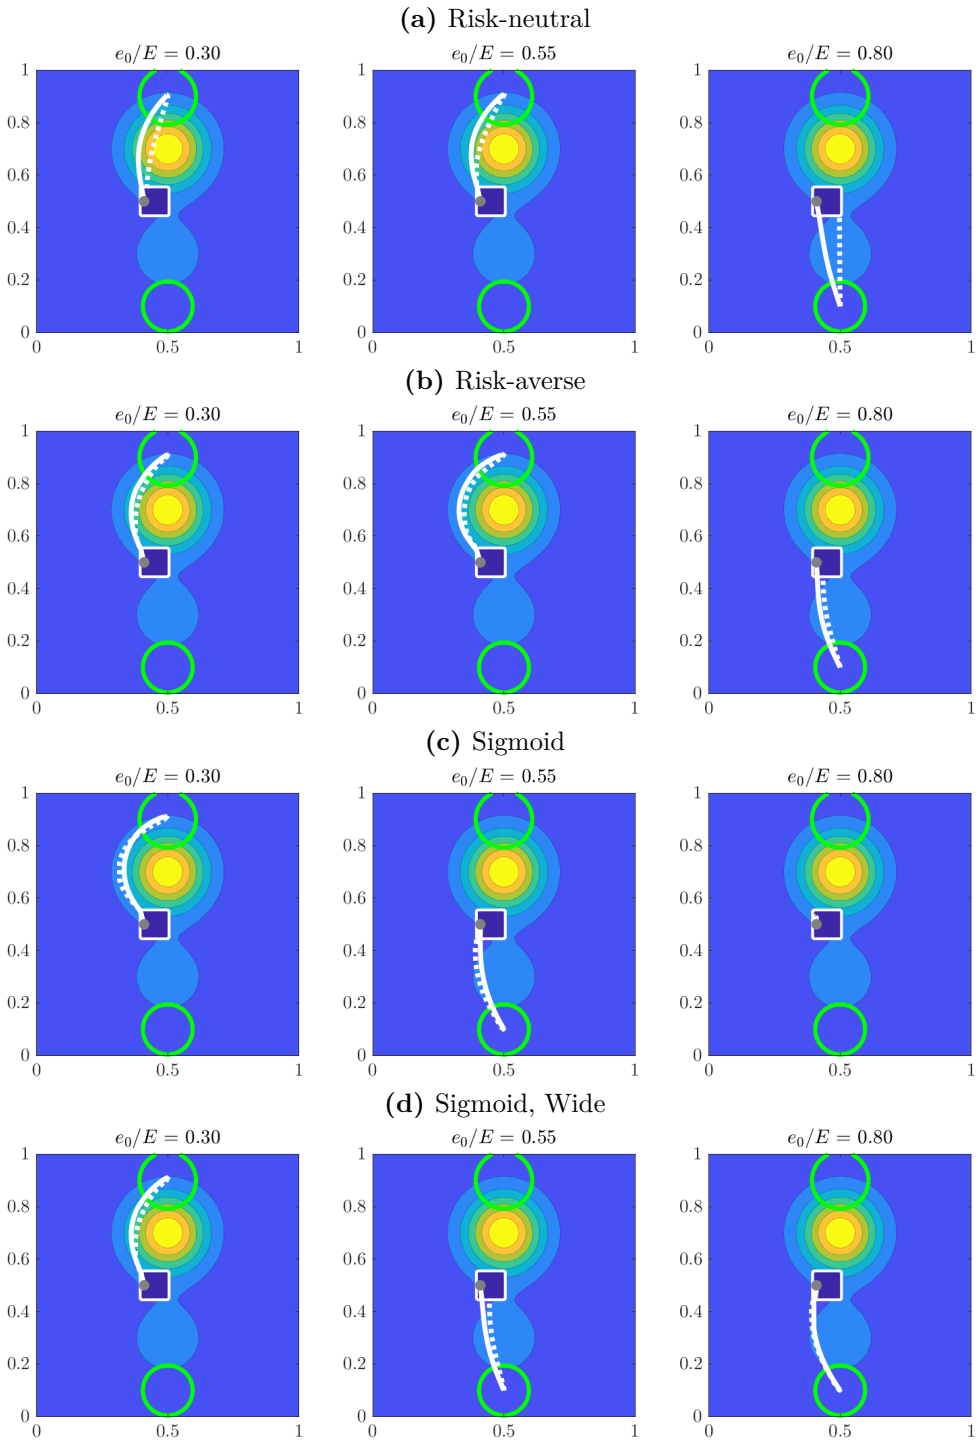

Figure 4: Preferred paths for selected initial energy levels and a fixed risk premium ($\rho = 2$). Paths to feeding grounds shown in solid white; return paths shown in dashed white. All subfigures present optimal trajectories chosen with possible predator encounters in mind—these are the paths followed if predators are feared but not really encountered along the way. As the initial energy increases, the forager chooses lower-risk options—longer detours, safer destinations, or even staying within the refuge. This matches the trend seen in Figure 3. Comparing different utility functions, foragers who choose the same feeding ground will sometimes take different paths to and from it, e.g., column 2 of **(a)** and **(b)** or column 1 of **(c)** and **(d)**). In those cases, risk preferences are only discernible if we consider factors other than patch choice.

utility.[4] For both sigmoid utility functions we set the threshold energy to $e_t = 0.3E$, which means the forager can easily surpass the inflection point of the curve for most initial energy levels. We vary the steepness of the sigmoid curve, with the default curve having a shape parameter of $\sigma = 0.1E$ and the wide curve having $\sigma = 0.2E$. The risk premium is varied by scaling the risk level in the upper half of the domain, while keeping the risk level in the bottom half the same across all simulations.

In general, the risk-neutral planner is the least conservative in its choice of food patch, and the narrow sigmoid planner is the most conservative. The wide sigmoid planner makes qualitatively similar patch choices to the narrow sigmoid planner, but it has a higher baseline risk tolerance. Indeed, when the risk premium is high and the initial energy is low, the wide sigmoid planner is less conservative than the risk-averse planner. This is because the wider sigmoid curve more closely approximates a linear, or risk-neutral, planner for energy levels near the inflection point of the curve. However, because the curve is not truly linear, the wide sigmoid planner is still more conservative in its patch choice than the truly risk-neutral planner.

All planners except the risk-neutral one sometimes make the choice to simply stay in place rather than collecting food; in those cases, the guaranteed survival outweighs any possible reward associated with venturing out. Since we do not vary the risk level in the bottom half of the domain, any animal that prefers staying at home to visiting the lower patch will do so for all levels of the risk-premium. However, it is still possible that the top (more risky) patch will be preferred over both of those for some values of ρ, as can be seen in Figure 3b for the risk-averse planner. These preferences could have consequences for how often foragers with differing risk preferences leave safety to collect food. Animals that are more risk-averse may wait until they are at lower energy reserves to begin foraging relative to their less risk-averse counterparts.

While the shape of the utility function U clearly impacts the optimal patch choice, there are significant regions of parameter space where the same patch is optimal for several different utility functions. This is especially true at low levels of the risk premium and at the lowest levels of initial energy. In experiments that test risk-preferences for foraging animals, it is thus important that the risk

[4] Here, the preferred patch is the patch that the animal initially plans to visit before any possible predator encounters have occurred. If the forager is spotted by a predator en route, it may return to the refuge before reaching its preferred patch and may visit a different patch on a subsequent excursion.

be substantial enough to motivate different behaviors. It also suggests that examining a range of starting energy states is important for discerning the animal's risk-preferences (as captured by the utility function).

However, the patch choice is not the full story of forager's optimal behavior, as two animals exploiting the same feeding ground could follow substantially different routes to and from it. A longer detour bypassing high predator density areas can often decrease the overall risk, but will also decrease the time available for feeding. In balancing the time spent on feeding versus traveling, the forager must take into account the food patch quality, predator density variations, its level of hunger, and the utility function. Figure 4 shows how the shape of optimal trajectories changes with utility function and initial energy level e_0 for a fixed value of the risk premium. In general, all foragers take safer (and longer) trajectories as e_0 increases. For the same initial energy level and patch choice, the risk-averse planner takes longer detours than the risk-neutral planner on both legs of the trip. The risk-neutral planner takes a longer detour on the way to the feeding ground than on the way back. The same phenomenon is less pronounced for the risk-averse planner and is energy-level dependent for sigmoidal planners. Thus, the observed routes to and from feeding locations provide additional valuable information for selecting a suitable utility function in models of foragers' behavior.

Foraging Samango Monkeys

We now consider an example based loosely on data from foraging Samango monkeys studied in [Coleman and Hill, 2014] and [Coleman et al., 2021]. Samango monkeys are small, arboreal primates that primarily forage for fruit in trees, travel in groups, and face predation from large birds of prey, snakes, and jaguars [Coleman et al., 2021]. One method that Samango monkeys use to avoid predators is to have a lookout, who will sound an alarm call if a predator is spotted, allowing the group to escape [Coleman and Hill, 2014]. While we do not model intra- or intergroup dynamics, we do assume that monkeys are aware of a predator when it appears and try to escape. In the study range, the primary predator encountered by the monkeys were eagles [Coleman and Hill, 2014], so we focus on just this predator-prey system. We track the position of the foraging monkeys in 2-dimensional space (though in reality they can also move up and down in the tree canopy).

Parameter	Symbol	Range	Units	Justification/Source
Metabolic cost	K_m	14	$\frac{\text{kcal}}{\text{h}}$	[Pontzer et al., 2014]
Speed	f_m	0.7	$\frac{\text{km}}{\text{h}}$	[Coleman et al., 2021]
Food density	ψ	$2.00 \times 10^6 - 2.20 \times 10^8$	$\frac{\text{kcal}}{\text{km}^2}$	[Coleman et al., 2021]
Food radius	r	0.05	km	Ad hoc
Harvest rate	α	1×10^{-6}	$\frac{1}{\text{h}}$	Ad hoc
Maximum energy	E	1400	kcal	[Coleman et al., 2021], [Zihlman and Bolter, 2015]
Planning horizon	T	6	h	[Coleman et al., 2021]
Spotting rate	μ_s	0.01 - 0.12	$\frac{1}{\text{h}}$	[Coleman and Hill, 2014]
Kill rate	μ_k	5.0	$\frac{1}{\text{h}}$	[Touchton et al., 2002]
Give up rate	μ_g	1.0 - 11	$\frac{1}{\text{h}}$	[Cordeiro, 2003]

Table 2: Parameter values and sources used for the stylized samango monkey example. Ranges are provided for parameters that vary across the environment, otherwise the parameter is assumed to be constant. The planning horizon of 6 hours roughly matches the amount of time samango monkeys spend traveling and foraging in one day [Coleman et al., 2021]. Food density, food radius, and harvest rate are used to calculate food availability (see Supplementary Materials Section A for details). When information directly pertaining to samango monkeys was not available, we used sources that study similar primates. While we aim to build a plausible model of the environment, we do not account for factors such as intergroup competition, the location of water, or other daily activities outside of foraging and traveling. We use the model to make qualitative predictions about the impact of the environment and utility function on foraging behavior (see Figure 6) rather than quantitative predictions about the trajectory of a specific animal.

We report our estimated model parameters in Table 2. All parameters are assumed constant except for food density ψ,[5] spotting rate μ_s, and give-up rate μ_g. The relative distributions of ψ and μ_s were estimated from color data from images in [Coleman and Hill, 2014] and then linearly rescaled using the parameters in Table 2. We set μ_g to be complementary to μ_s, i.e., μ_g is a linear rescaling of $1 - \mu_s / \max(\mu_s)$. Our versions of the resulting distributions are shown in Figure 5. The foraging area reported in [Coleman and Hill, 2014] is 0.547 km^2 but for simplicity we assume that the map has a width of 1 km and a height of 0.7 km, which results in a comparably sized domain once obstacles are accounted for. When simulating forager trajectories, starting locations were chosen from areas identified as being close to sleeping sites in [Coleman and Hill, 2014], and the planning horizon was on the scale of a single day. To reflect the fact that most monkeys are likely far from starvation, we use an initial energy of $e = 0.5E = 700$ kcal which means that the monkeys cannot die of starvation

[5] Recall that food density is used to compute food availability $F(\mathbf{x})$. See Supplementary Materials Section A for the details of this calculation.

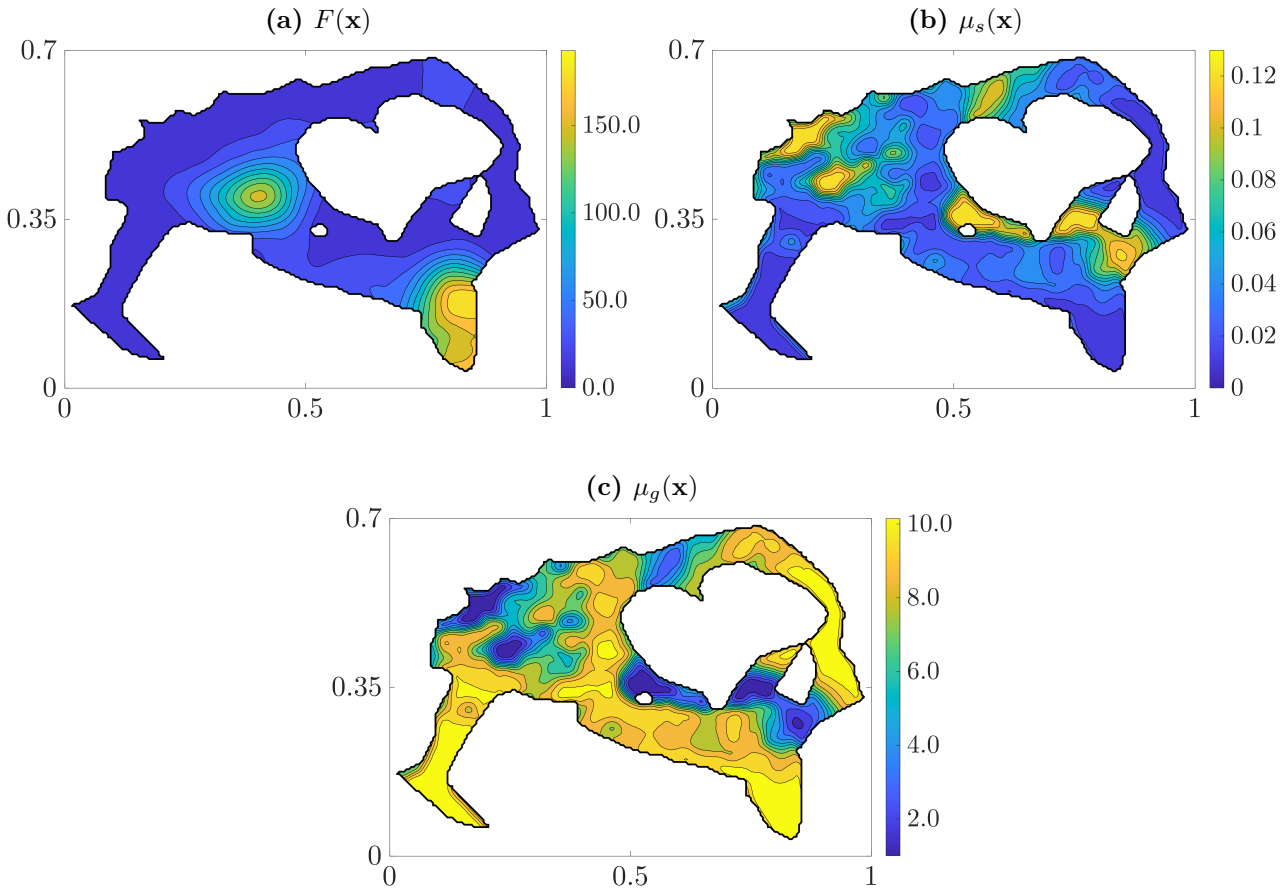

Figure 5: Stylized environment based on data for samango monkeys. Samango monkeys are small arboreal primates that mainly forage by collecting fruit in trees and are subject to predation from snakes, jaguars, and eagles [Coleman et al., 2021]. In [Coleman and Hill, 2014], the authors present detailed data about the monkeys' environment in the form of heat maps. To create our stylized environment, we estimate the relative distribution of parameters using color data from the images and linearly rescale it using the parameters shown in Table 2. **(a)** Food availability is calculated based on food density (see Supplementary Materials Section A for details). Distribution of food density is based on Figure 1b in [Coleman and Hill, 2014]. **(b)** Distribution of predators is based on Figure 1g in [Coleman and Hill, 2014], which measures alarm calls due to eagles. We apply a smoothing filter due to small high-frequency variations in color data. **(c)** We choose to make the give-up rate complementary to μ_s (i.e., μ_g is a linear rescaling of $1-\mu_s/\max(\mu_s)$) because we assume that eagles are reluctant to pursue monkeys far from areas they normally hunt.

within the time period we consider.

Figure 6 shows predicted optimal foraging trajectories for 1000 randomly chosen starting locations and the resulting habitat utilization. For all utility functions, optimal foraging paths avoid areas with high rates of predation, which is in line with the main finding of [Coleman and Hill, 2014] that predation risk was a consistent driver of habitat utilization. Additionally, our model predicts the

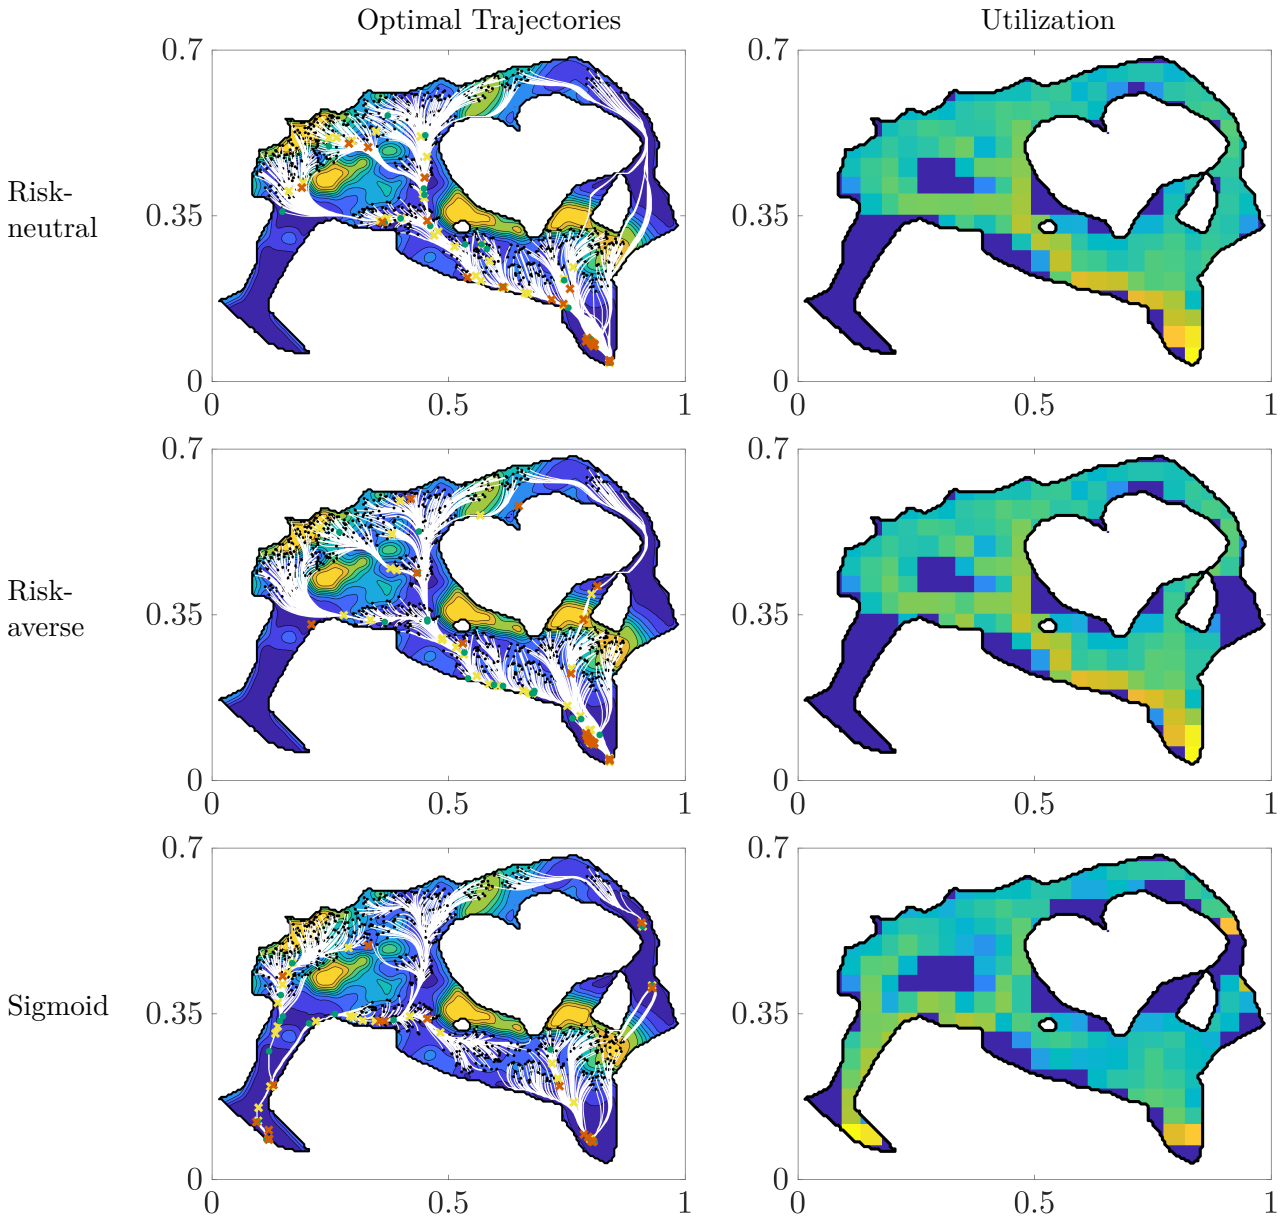

Figure 6: Monkeys' optimal foraging trajectories and habitat utilization for three utility functions. Initial positions (black dots) are randomly chosen from areas identified as being close to sleeping sites in [Coleman and Hill, 2014]. Initial energy is set to $e = 0.5E = 700$, meaning that the monkeys cannot starve over the time horizon we consider. Left column: optimal trajectories and interactions with predators, where yellow "x"s are predator spotting events, green dots are successful evasions, and red "x"s are deaths. Trajectories congregate in "channels" as the monkeys head to their chosen destination. Risk-neutral and risk-averse planners generally end up at an area of high food availability, while sigmoid planners prioritize reaching nearby feeding grounds with low spotting rate. Right column: utilization is computed as the log of total time spent in a region divided by the area of that region. A notable similarity between our simulated utilization and the utilization estimated in [Coleman and Hill, 2014] is the dedicated avoidance areas with high risk of being spotted by an eagle. Our model is not sophisticated enough to reliably predict overall utilization but it provides qualitative insight into areas where we could expect large differences in monkey behavior based on the foraging objective.

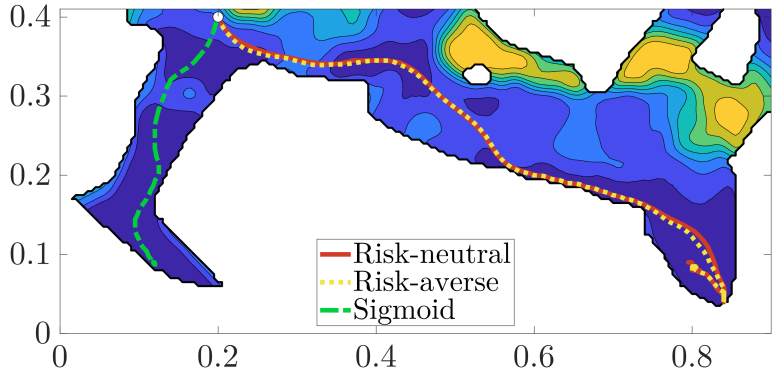

Figure 7: Sample "no predators encountered" trajectories under different utility functions. A monkey moves from the white dot to a chosen destination, where it spends the remaining time until the end of the planning horizon. Even though the monkey starts and remains in mode 1 (no predators), the possibility of encountering eagles still influences the chosen trajectory. The risk-neutral and risk-averse planners choose a feeding ground destination and follow very similar paths to it. (However, see Figure 8 for how differently they would handle an encounter with predators along the way.) Since the monkey begins above the energy threshold of the sigmoid curve and cannot fall below it in the time period we consider, the sigmoid-planner does not prioritize food as much and heads for the safest nearby location instead.

formation of "channels" within which the optimal trajectories congregate. These channels are very similar for the risk-neutral and risk-averse utility functions, with most trajectories converging on a region in the lower-right corner of the domain with high food availability. In contrast, under a sigmoidal utility function, optimal trajectories favor areas of low predation risk. Here, we focus on the regime in which the monkeys' energy reserves are safely above the threshold of starvation (and thus the inflection point of the utility function), thus they prioritize safety over taking risks for the sake of greater energy gain.

Because we explicitly model interactions with predators, we can also examine predicted optimal behavior with and without predator encounters under each utility function. Figure 7 shows optimal trajectories from a single starting location for all utility types and no predator interactions, which largely reflect the trends shown in Figure 6. Next, we consider predator interactions in two cases: the monkey is immediately spotted by an eagle and the monkey is spotted by an eagle after 0.5 hours. Even though the sigmoidal utility function leads to very different behavior in the absence of a predator, all planners take the same escape route if they are spotted immediately (Figures 8a, 8c, and 8e.) Conversely, though the risk-neutral and risk-averse planners have nearly identical trajectories in the absence of a predator, they take different escape routes when spotted at 0.5 hours (Figures

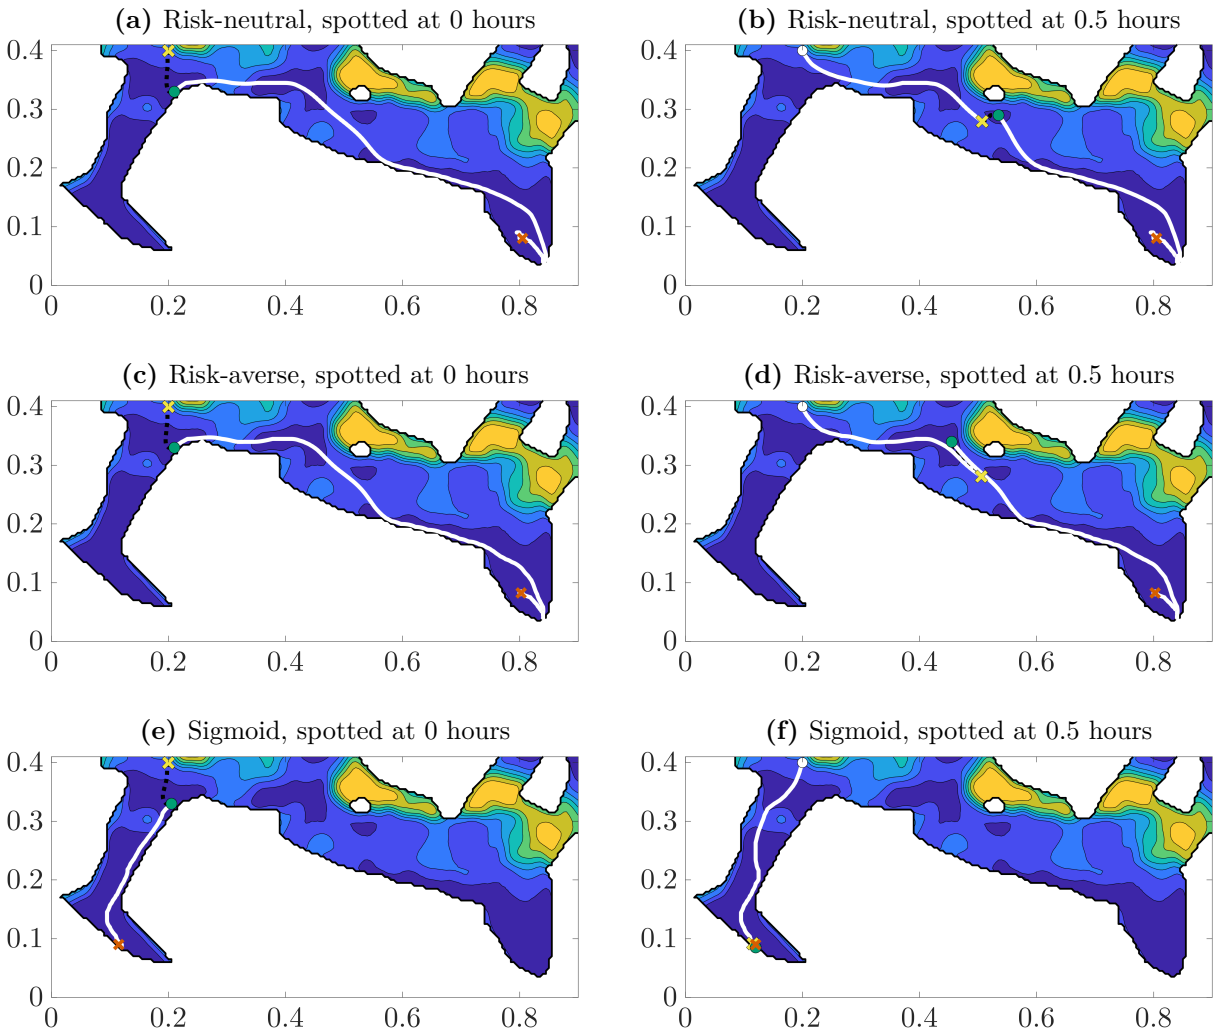

Figure 8: Sample trajectories with predator encounters across utility functions and spotting times. A monkey begins in the same starting location as Figure 7 but now it is spotted by an eagle either immediately (left column) or after 0.5 hours (right column). Yellow "x"s represent predator spottings and green dots represent successful escapes. In this scenario a monkey that is spotted immediately will follow the same escape trajectory (black dotted line), regardless of its utility function. Conversely, the risk-neutral and risk-averse planners (which follow near identical paths in the absence of a predator) take significantly different escape routes if spotted at 0.5 hours (subfigures **(b)** and **(d)**). The risk-neutral planner takes a slight detour to a safer location but does not greatly deviate from its planned path, while the risk-averse planner significantly backtracks along its previous route in order to escape. Here, we observe that monkeys that behave similarly in the absence of a predator may not react to a predator appearance the same way and vice versa.

8b and 8d). Thus, the arrival of a predator can cause distinct planners to converge upon a highly favorable escape route, but can also cause otherwise very similar planners to take drastically different escape measures.

4 Modeling Food Depletion

So far, we have assumed that foragers do not deplete food as they exploit their environments, as in the classical optimal diet model. This is a reasonable assumption when resources are abundant or regenerate quickly. However, some foragers meaningfully deplete resources as they exploit them, and must account for this in their foraging decisions. Failure to account for this in a model can lead to unrealistic predicted behaviors. A forager that experiences food depletion but does not plan for it will generally achieve a lower energy level than it initially predicted, and may adjust its path in order to compensate.[6] In principle, food depletion can be addressed by including the remaining food at each location as additional state variables in the dynamic programming model. Unfortunately, this approach quickly becomes computationally infeasible due to the "curse of dimensionality"—the computing time required to calculate the value function grows exponentially with the dimension of the state space. At the same time, it is disputed whether it is accurate to model foraging animals as making highly accurate long-term predictions about their environment [Railsback, 2022], so it is not obvious that we should model animals as exactly accounting for food depletion during an entire foraging bout.

Motivated by these considerations, we propose a "suboptimal" foraging model in which a forager makes decisions based on imperfect, short-term predictions for the future state of its environment. Here we give an intuitive description of the model and discuss its benefits and limitations, hoping that it serves as a starting point for future work (mathematical details are available in the Supplementary Materials, Section C). [Railsback, 2022] outlines a "suboptimal foraging" approach for individual based models (IBMs) where foragers evaluate the quality of a patch under the assumption that its state remains constant over a long period of time. We instead assume that the animal makes myopic estimates and decisions over a *short* time period (stage) relative to the duration of a foraging bout. Specifically, at the start of each "short stage" of length S, a forager makes plans to maximize its expected utility S time units later, on the assumption that food availability will not change from its current state. It then behaves accordingly for S time units, collecting and depleting food as it moves.

[6] We present an example of this behavior in the Supplementary Materials, Figure 10. When a forager collects less food than it expected, it remains at the feeding site longer, takes riskier paths back to the refuge, and ends up with a lower expected utility than a forager that does not deplete food.

At the end of each short stage, the forager becomes aware of the updated map of food availability, and makes new plans for the next short stage. This continues until the end of the time horizon T. This "multistage model" converts the intractable problem of optimally accounting for food depletion into a sequence of simpler, solvable problems for each stage.

To illustrate how this model can lead to sensible decisions, we consider a forager with two possible feeding areas (top and bottom), identical in food abundance and distance to a refuge, but with the bottom patch being safer (i.e., the forager is less likely to be killed there if it is spotted by a predator, see the Supplementary Materials, Figure 11 for details). A forager blind to food depletion would only exploit the safer area, but a depletion-aware forager in our multistage model recognizes that it is often better to exploit both, though favoring the safer area.

Figure 9 shows three realized trajectories under the multistage model for a risk-neutral forager. If not spotted by a predator (top row), the forager spends the first three stages exploiting the safer food patch. At the beginning of the fourth stage, the patch has been depleted to the extent that it is optimal to travel to the unexploited top patch, though taking a detour to remain closer to the refuge. If the forager is instead spotted while traveling between patches (middle row), it escapes towards the refuge to avoid the predator. When the forager is later spotted near the end of stage four, it misjudges the cumulative danger since it does not plan beyond the end of that stage and the remaining time is insufficient to reach the burrow anyway. As a result, it temporarily lingers in that patch (despite the predator's presence) and only returns to the refuge once the next stage begins. This behavior is an artefact of the modeling process and not biologically motivated. The bottom row of Figure 9 highlights a possible impact of an early predator spotting: when spotted at the end of stage two, the forager escapes to the refuge in stage three, and then takes advantage of its central position to travel to the unexploited upper patch. As a result, it arrives in the risky patch earlier, and spends more time there than in the safe patch over the total time horizon. Similar figures for concave (risk-averse) and sigmoid utility functions are available in the Supplementary Materials, Figures 12 and 13.

A key strength of the multistage model is that it incorporates the forager's feeding history in a computationally feasible manner. Indeed, the cost of solving the dynamic programming problem is

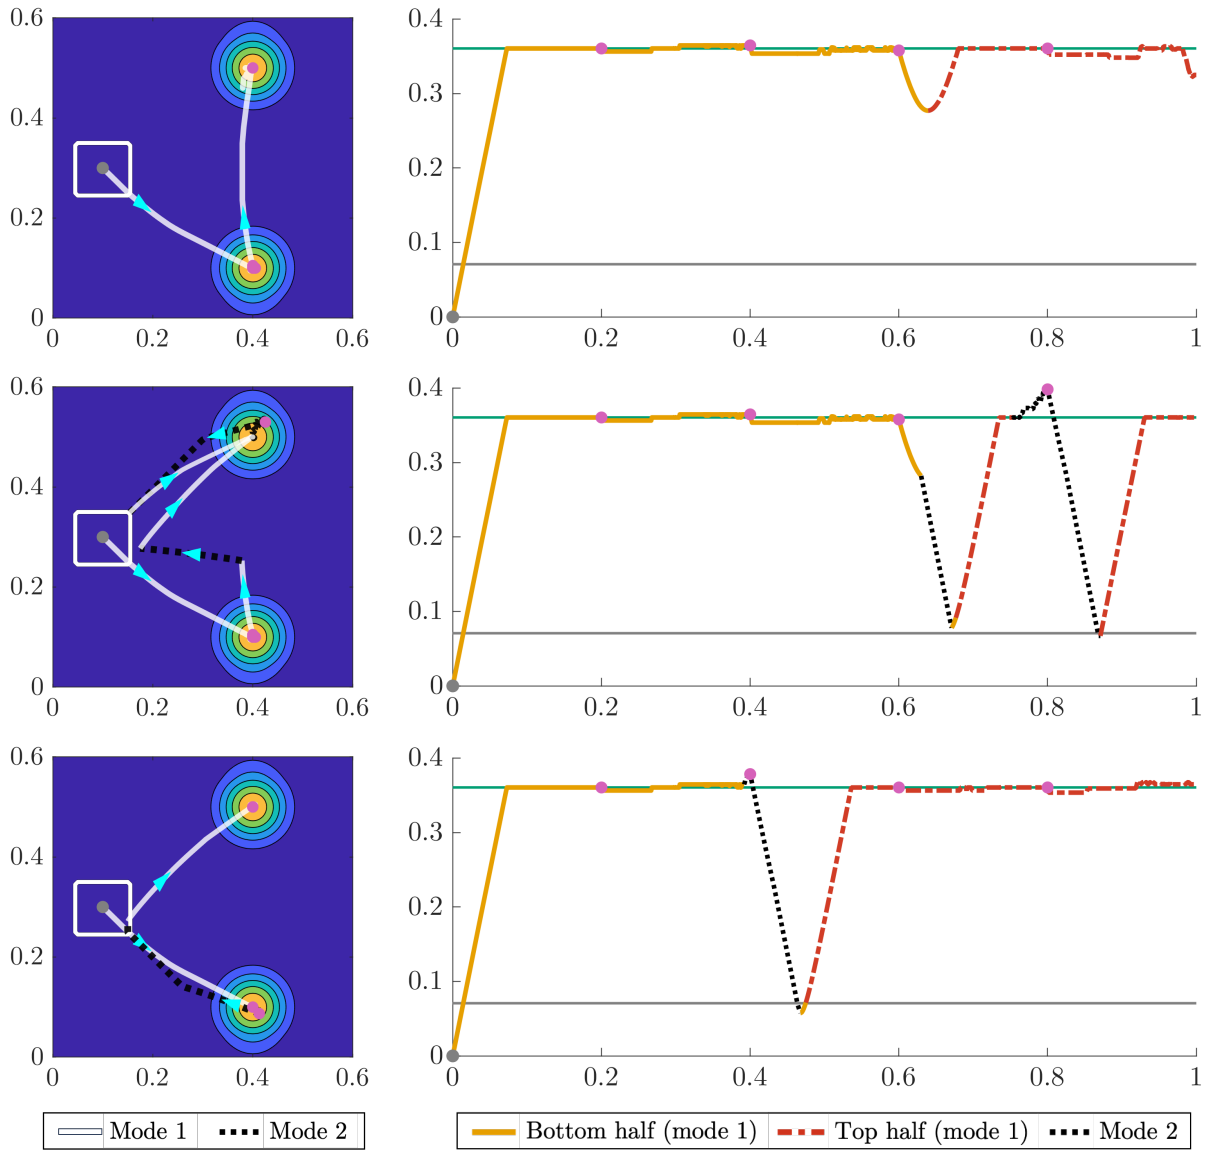

Figure 9: A myopic forager that periodically becomes aware of food depletion exploits both food patches (Risk-neutral). Three realized optimal trajectories according to the "multi-stage model" with five stages using a risk-neutral (linear) utility function. The forager begins in the center of the refuge (grey dot) and the end of each stage is marked with a magenta dot. The first column shows optimal trajectories in the forager's 2D environment, with blue arrows indicating the direction of travel. In the second column, we plot the forager's distance from the starting location over time for the same realized trajectories. The green line represents the distance to the center of each food patch and the grey line is the distance to the *corners* of the refuge. Across all realizations, the forager preferentially exploits the safer food patch, but eventually switches to the top patch as it becomes aware of the impact of food depletion. If not spotted by a predator, the forager travels directly from the first patch to the second (top row). However, if spotted by a predator it may return to the refuge (middle and bottom rows), and in some cases switch to the top patch earlier due to such a detour (bottom row). "Wobbling" in the trajectories is caused by small variations in food availability due to depletion.

the same as for the original model. However, because these solutions incorporate information specific to a single trajectory, they cannot be used to efficiently simulate trajectories for many starting locations or for many different random interactions with predators. Furthermore, the multistage model introduces additional modeling choices: the terminal condition and length for each stage. The forager's expected utility at the end of the stage is natural terminal condition, as we have already interpreted it as a measure of future reproductive success. However, it does not incorporate any information about the forager's future prospects in the short term, and may lead to unrealistic behavior, e.g., a forager that stays in place and starves because it cannot reach a food source within a single stage (even if it could reach such food over multiple stages of travel).

Careful choice of the stage length can alleviate this issue, but comes with its own challenges. Even when each stage is long enough for the forager to reach food, travel will take up proportionally more time for short stages, exaggerating the opportunity cost associated with traveling. Similarly, the forager will underestimate the risk of being killed by a predator for shorter stages or near the end of a stage, because it is less likely to be spotted and killed before the stage ends (e.g., stage four in the middle row of Figure 9). However, overly long stages inherit the drawback of the original single-stage model, that a forager might remain at a depleted food source when it is better to move on. Thus, stages must be long enough for the forager to reasonably estimate trade-offs between travel, foraging, and predation, but short enough that it receives timely information about depleting food. Despite these limitations, we believe that this approach is an important step in developing approximate models that can better represent situations where foragers impact their environment.

5 Discussion

Fear of predation has consequences at all scales in ecology. Exposing snowshoe hares to simulated predators during gestation reduced female survival by 30% and survival to weaning of offspring by more than 85% [MacLeod et al., 2018]. In general, the traits and actions evoked by predation risk have on average larger demographic impacts on prey populations than the actual deaths due to predation [Preisser et al., 2005]. Fear-induced responses by many individuals in a population can collectively affect community and large-scale ecosystem properties [Pinti et al., 2022]. For example,

simulated presence of large predators reduced raccoon foraging in an intertidal zone community, in part by causing raccoons to spend more time foraging elsewhere. Reduced raccoon foraging resulted in 59% to 97% increases in the abundance of their main prey species; moreover, it also affected the survival and abundance of those raccoon prey species' own prey and competitors [Suraci et al., 2016].

Mobile animals often try to reduce the predation risk (real or perceived) through their choices of feeding areas and movement paths while foraging (e.g., [Gilliam and Fraser, 1987, Lima and Dill, 1990, Brown and Kotler, 2004, Laundré et al., 2001, Brown et al., 1999, Coleman and Hill, 2014]). Optimal foraging models that incorporate predation risk have generally assumed discrete foraging "patches" varying in food availability and predation risk. The time resolution of such models is typically coarse and the physical proximity of patches is often ignored. An animal is essentially "teleporting" to and from its feeding grounds. while the predation risks en route to and at the chosen patch are lumped together. In contrast, the models introduced here treat foraging as a two-dimensional navigation problem in a complex, continuous landscape with food and risk varying continuously, and the forager's behavior explicitly changes whenever a predator is encountered. Using stochastic dynamic programming, we find the optimal behavior for all state configurations, thereby predicting where an animal should choose to feed, what paths it should take to and from those locations, and how it should move when trying to escape a predator.

Understanding better how foragers decide where to go and what to eat is not just scientifically interesting – it is essential for predicting how these populations may be impacted by habitat modifications and changes in the abundance or types of prey or predator species. While we can often directly observe *what* a forager chooses to do in different circumstances in the real world, its decision algorithm has to be inferred indirectly from the observed behaviors. Our theoretical and empirically motivated case studies show that a forager's chosen movement paths provide valuable information about its motivations, above and beyond what we could learn from just observing where it chooses to feed. In particular, these paths could be used to evaluate how the forager values different possible energy levels at the conclusion of a foraging bout (Figure 1b). For any assumed utility function, the dynamic programming solution makes it easy to numerically construct a "phase diagram" of how foraging behaviors change as a function of the forager's initial energy state and initial locations. Knowing

which conditions lead to qualitatively different behaviors for different utility functions can identify which field experiments will be most informative about the unobserved decision-making process.

In situations where one focal forager has a substantial effect on food availability in the environment, exact dynamic programming is computationally infeasible due to the "curse of dimensionality". To address this case, we propose a "suboptimal foraging model" [Railsback, 2022], in which the forager plans over a short time horizon but periodically updates its plans based on how it has impacted the environment. While this requires a careful selection of the frequency of such updates, we believe this is a useful step beyond modeling a single forager in an unchanging resource landscape.

One limitation of our model is the assumption of deterministic energy and location dynamics (Equations (1) through (4)). More generally, food uptake could be modeled as a series of random encounters with food items at a rate determined by the local food density, and death from starvation could be similarly treated as random (with an increasing probability as the energy level approaches zero). Randomness in movement paths could be modeled by a stochastic differential equation with additive Brownian motion reflecting an "exploration" aspect of foraging in addition to the "exploitation" already included in our model. Including these additional stochastic features would change the form of the governing Hamilton-Jacobi-Bellman PDEs, but but the general numerical approach for solving it would be quite similar [Fleming and Soner, 2006].

Our current approach also does not allow us to consider effects of interactions between animals [Agrawal et al., 2007], with or without food depletion. One possible method within a continuous framework would be to model large numbers of interacting individuals using mean-field games (MFGs) [Lasry and Lions, 2007]. Multi-population MFGs can be also used to model complex cooperative interactions within cohesive groups and the competitive inter-group interactions of many animals exploiting depletable food resources [Coleman, 2013, Coleman and Hill, 2014]. Furthermore, using approximate dynamic programming methods designed to cope with high-dimensional state spaces, rather than seeking exact solutions as we have done here, may expand the range of models that are computationally feasible [Powell, 2007].

A broader limitation of our modeling framework is the assumption that animals are perfect rational

optimizers. In reality, there are reasons to believe that foragers would not have evolved perfectly optimal strategies since those rules would not necessarily provide large selective advantages over simpler decision rules [Janetos and Cole, 1981]. In fact, there is some evidence that simple, "good enough" rules produce behavior that is more closely aligned with reality [Crowley et al., 1990]. Thus, the trajectories presented here should not be thought of as quantitative predictions of animal behavior, but rather as qualitative descriptions of what a perfect optimizer would do. These are useful for two reasons: providing a bound on how well an animal could do in practice and providing qualitative insights for the behavior of "somewhat rational" and "somewhat informed" optimizers. One extension that could more accurately characterize trade offs between maximizing utility and relying on strategies that are simple and broadly applicable is multi-objective optimization, where a forager seeks to optimize two or more objectives at once. Efficient methods also exist for recovering such "Pareto optimal" in the framework of dynamic programming [Mitchell and Sastry, 2003], [Kumar and Vladimirsky, 2010], [Désilles and Zidani, 2019].

Overall, we hope that this paper will motivate further work on modeling how mobile organisms choose to move in a landscape full of opportunities and threats. As animal tracking and biologging capabilities continue to improve, [Tomkiewicz et al., 2010, Kays et al., 2015], there are increasing opportunities to track animals over long periods of time and characterize qualitative relationships between animal state and behavior. These observations can in turn be compared to the optimal decision rules under different assumptions to learn more about how foragers trade off risks and opportunities. The consistently mediocre performance of species distribution models that seek to predict organism presence/absence or abundance using climate and habitat covariates (e.g., [Lee-Yaw et al., 2022, Liu et al., 2020]) suggests that mobile animals are not just "going where the climate suits my clothes" [Whitter, 1924]. Predicting how distribution and abundance will change in response to human impacts (either detrimental or intended to be helpful) will require a deeper understanding of how animals decide what to do where and when.

Acknowledgments

This research was partially funded by the National Science Foundation (awards DMS-1645643 and DMS-2111522) as well as the Air Force Office of Scientific Research (award FA9550-22-1-0528).

Statement of Authorship

All authors contributed to the formulation of the mathematical model. MG, SJ, NR, and AS developed the solution code and visualization. NGG also contributed to the visualization. MG, AS, and NGG contributed to data collection and processing. MG, SJ, AS, AV, and SE developed and refined the numerical experiments. MG, NGG, SJ, and AS were responsible for early drafting of the paper and results. MG, AV, and SE revised and developed the paper, and all authors contributed to final editing and approval.

Data and Code Availability

All data and code used to generate the figures in this manuscript are available online at `https://github.com/eikonal-equation/Navigating_LoF/tree/main`.

References

[Agrawal et al., 2007] Agrawal, A. A., Ackerly, D. D., Adler, F., Arnold, A. E., Cáceres, C., Doak, D. F., Post, E., Hudson, P. J., Maron, J., Mooney, K. A., Power, M., Schemske, D., Stachowicz, J., Strauss, S., Turner, M. G., and Werner, E. (2007). Filling key gaps in population and community ecology. *Frontiers in Ecology and the Environment*, 5(3):145–152.

[Brown, 1988] Brown, J. S. (1988). Patch use as an indicator of habitat preference, predation risk, and competition. *Behavioral ecology and sociobiology*, 22:37–47.

[Brown and Kotler, 2004] Brown, J. S. and Kotler, B. P. (2004). Hazardous duty pay and the foraging cost of predation. *Ecology Letters*, 7(10):999–1014.

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

[Caraco, 1980] Caraco, T. (1980). On foraging time allocation in a stochastic environment. *Ecology*, 61(1):119–128.

[Caraco and Chasin, 1984] Caraco, T. and Chasin, M. (1984). Foraging preferences: Response to reward skew. *Animal Behaviour*, 32(1):76–85.

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

[Coleman, 2013] Coleman, B. (2013). *Spatial and temporal determinants of samango monkey (Cercopithecus mitis erythrarchus) resource acquisition and predation avoidance behaviour*. PhD thesis, Durham University.

[Coleman and Hill, 2014] Coleman, B. T. and Hill, R. A. (2014). Living in a landscape of fear: the impact of predation, resource availability and habitat structure on primate range use. *Animal Behaviour*, 88:165–173.

[Coleman et al., 2021] Coleman, B. T., Setchell, J. M., and Hill, R. A. (2021). Seasonal variation in the behavioural ecology of samango monkeys (Cercopithecus albogularis schwarzi) in a southern latitude montane environment. *Primates*, 62(6):1005–1018.

[Cordeiro, 2003] Cordeiro, N. J. (2003). Two unsuccessful attacks by crowned eagles (Stephanoaetus coronatus) on white-throated monkeys (Cercopithecus mitis). *African Journal of Ecology*, 41(2):190–191.

[Crowley et al., 1990] Crowley, P., DeVries, D., and Sih, A. (1990). Inadvertent errors and error-constrained optimization: fallible foraging by bluegill sunfish. *Behavioral Ecology and Sociobiology*, 27:135–144.

[Davis, 1984] Davis, M. H. A. (1984). Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society. Series B (Methodological)*, 46(3):353–388.

[Désilles and Zidani, 2019] Désilles, A. and Zidani, H. (2019). Pareto front characterization for multiobjective optimal control problems using Hamilton–Jacobi approach. *SIAM Journal on Control and Optimization*, 57(6):3884–3910.

[Emlen, 1966] Emlen, J. M. (1966). The role of time and energy in food preference. *The American Naturalist*, 100(916):611–617.

[Fiksen et al., 1995] Fiksen, Ø., Giske, J., and Slagstad, D. (1995). A spatially explicit fitness-based model of capelin migrations the Barents Sea. *Fisheries Oceanography*, 4(3):193–208.

[Fleming and Soner, 2006] Fleming, W. H. and Soner, H. M. (2006). *Controlled Markov Processes and Viscosity Solutions*. Springer-Verlag, 2nd edition.

[Gaynor et al., 2019] Gaynor, K. M., Brown, J. S., Middleton, A. D., Power, M. E., and Brashares, J. S. (2019). Landscapes of fear: Spatial patterns of risk perception and response. *Trends in Ecology & Evolution*, 34(4):355–368.

[Gilliam and Fraser, 1987] Gilliam, J. F. and Fraser, D. F. (1987). Habitat selection under predation hazard: Test of a model with foraging minnows. *Ecology*, 68:856–1862.

[Houston and McNamara, 1982] Houston, A. and McNamara, J. (1982). A sequential approach to risk-taking. *Animal Behaviour*, 30(4):1260–1261.

[Houston et al., 1993] Houston, A. I., McNamara, J. M., and 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.

[Janetos and Cole, 1981] Janetos, A. C. and Cole, B. J. (1981). Imperfectly optimal animals. *Behavioral Ecology and Sociobiology*, 9:203–209.

[Kays et al., 2015] Kays, R., Crofoot, M. C., Jetz, W., and Wikelski, M. (2015). Terrestrial animal tracking as an eye on life and planet. *Science*, 348(6240):aaa2478.

[Kirby et al., 2000] Kirby, D. S., Fiksen, Ø., and Hart, P. J. B. (2000). A dynamic optimisation model for the behaviour of tunas at ocean fronts. *Fisheries Oceanography*, 9(4):328–342.

[Kohl et al., 2018] Kohl, M. T., Stahler, D. R., Metz, M. C., Forester, J. D., Kauffman, M. J., Varley, N., White, P. J., Smith, D. W., and MacNulty, D. R. (2018). Diel predator activity drives a dynamic landscape of fear. *Ecological Monographs*, 88(4):638–652.

[Kumar and Vladimirsky, 2010] Kumar, A. and Vladimirsky, A. (2010). An efficient method for multiobjective optimal control and optimal control subject to integral constraints. *Journal of Computational Mathematics*, 28(4):517–551.

[Lasry and Lions, 2007] Lasry, J.-M. and Lions, P.-L. (2007). Mean field games. *Jpn. J. Math*, 2:229–260.

[Laundré et al., 2001] Laundré, J. W., Hernández, L., and Altendorf, K. B. (2001). Wolves, elk, and bison: reestablishing the "landscape of fear" in Yellowstone National Park, U.S.A. *Canadian Journal of Zoology*, 79(8):1401–1409.

[Lee-Yaw et al., 2022] Lee-Yaw, J. A., McCune, J. A., Pironon, S., and Sheth, S. N. (2022). Species distribution models rarely predict the biology of real populations. *Ecography*, 2022(6):e05877.

[Lima and Dill, 1990] Lima, S. L. and Dill, L. M. (1990). Behavioral decisions made under the risk of predation: a review and prospectus. *Canadian Journal of Zoology*, 68(4):619–640.

[Liu et al., 2020] Liu, C., Wolter, C., Xian, W., and Jeschke, J. M. (2020). Species distribution models have limited spatial transferability for invasive species. *Ecology Letters*, 23(11):1682–1692.

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

[MacLeod et al., 2018] MacLeod, K. J., Krebs, C. J., Boonstra, R., and Sheriff, M. J. (2018). Fear and lethality in snowshoe hares: the deadly effects of non-consumptive predation risk. *Oikos*, 127:375–380.

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

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

[Mitchell and Sastry, 2003] Mitchell, I. M. and Sastry, S. (2003). Continuous path planning with multiple constraints. In *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, volume 5, pages 5502–5507. IEEE.

[Newman, 1991] Newman, J. A. (1991). Patch use under predation hazard: Foraging behavior in a simple stochastic environment. *Oikos*, 61(1):29–44.

[Pinti et al., 2022] Pinti, J., Visser, A., Serra-Pompei, C., Andersen, K., Ohman, M., and Kiørboe, T. (2022). Fear and loathing in the pelagic: How the seascape of fear impacts the biological carbon pump. *Limnology and Oceanography*, 67(6):1238 – 1256.

[Pontzer et al., 2014] Pontzer, H., Raichlen, D. A., Gordon, A. D., Schroepfer-Walker, K. K., Hare, B., O'Neill, M. C., Muldoon, K. M., Dunsworth, H. M., Wood, B. M., Isler, K., et al. (2014). Primate energy expenditure and life history. *Proceedings of the National Academy of Sciences*, 111(4):1433–1437.

[Powell, 2007] Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons.

[Preisser et al., 2005] Preisser, E. L., Bolnick, D. I., and Benard, M. F. (2005). Scared to death? the effects of intimidation and consumption in predator-prey interactions. *Ecology*, 86(2):501–509.

[Railsback, 2022] Railsback, S. F. (2022). Suboptimal foraging theory: How inaccurate predictions and approximations can make better models of adaptive behavior. *Ecology*, 103(8):e3721.

[Rands, 2017] Rands, S. A. (2017). Leaving safety to visit a feeding site: is it optimal to hesitate while exposed? *Royal Society Open Science*, 4(1):160910.

[Raynor et al., 2021] Raynor, J. L., Grainger, C. A., and Parker, D. P. (2021). Wolves make roadways safer, generating large economic returns to predator conservation. *Proceedings of the National Academy of Sciences*, 118(22):e2023251118.

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

[Sirot, 2019] Sirot, E. (2019). Adjustments in compound defensive strategies in response to variation in predation risk. *Animal Behaviour*, 147:53–60.

[Stephens, 1981] Stephens, D. (1981). The logic of risk-sensitive foraging preferences. *Animal Behaviour*, 29(2):628–629.

[Stephens and Krebs, 1987] Stephens, D. W. and Krebs, J. R. (1987). *Foraging Theory*. Princeton University Press.

[Suraci et al., 2016] Suraci, J. P., Clinchy, M., Dill, L. M., Roberts, D., and Zanette, L. Y. (2016). Fear of large carnivores causes a trophic cascade. *Nature Communications*, 7:10698.

[Szekely et al., 1991] Szekely, T., Sozou, P. D., and Houston, A. I. (1991). Flocking behaviour of passerines: a dynamic model for the non-reproductive season. *Behavioral Ecology and Sociobiology*, 28:203–213.

[Tomkiewicz et al., 2010] Tomkiewicz, S. M., Fuller, M. R., Kie, J. G., and Bates, K. K. (2010). Global positioning system and associated technologies in animal behaviour and ecological research. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1550):2163–2176.

[Touchton et al., 2002] Touchton, J. M., Hsu, Y.-C., and Palleroni, A. (2002). Foraging ecology of reintroduced captive-bred subadult harpy eagles (Harpia harpyja) on Barro Colorado Island, Panama. *Neotropical Ornithology*, 13(4):365–379.

[Werner and Gilliam, 1984] Werner, E. E. and Gilliam, J. F. (1984). The ontogenetic niche and species interactions in size-structured populations. *Annual Review of Ecology and Systematics*, 15:393–425.

[Whitter, 1924] Whitter, H. (1924). *Lonesome Road Blues (Okeh 20015-A)*. General Phonograph Corporation.

[Zihlman and Bolter, 2015] Zihlman, A. L. and Bolter, D. R. (2015). Body composition in Pan paniscus compared with Homo sapiens has implications for changes during human evolution. *Proceedings of the National Academy of Sciences*, 112(24):7466–7471.

A Calculating Food Availability

Here we describe the details of how we compute food availability $F(\mathbf{x}, t)$, which can be thought of as the rate of energy accumulation throughout the environment. Within our domain Ω, there may be terrain in which the animal cannot collect food, such as obstacles or refuge areas representing burrows or dens. The region in which the forager can collect food is called "forageable" and we denote it as Ω_F. We view total food collected per unit time at some location $\mathbf{x}$ as a weighted average of the harvestable food throughout $D_r(\mathbf{x})$, a circle of radius r centered at $\mathbf{x}$. Thus, when $D_r(\mathbf{x}) \subset \Omega_F$ we write this as

$$F(\mathbf{x}, t) = \int_{D_r(\mathbf{x})} \alpha B(\|\mathbf{x} - \mathbf{y}\|)\psi(\mathbf{y}, t)\, d\mathbf{y} \tag{13}$$

where α is a harvest rate, representing how efficiently the animal can access food resources, and $B : \mathbb{R} \to \mathbb{R}_{+,0}$ is a kernel function representing how the animal distributes its effort, with $\int_{D_r(\mathbf{x})} B(\|\mathbf{x} - \mathbf{y}\|)\, d\mathbf{y} = 1$. More generally, if $D_r(\mathbf{x}) \not\subset \Omega_F$, the animal spreads the same effort over a smaller area, $D_r(\mathbf{x}) \cap \Omega_F$. To account for this, we introduce a normalizing factor

$$Z(\mathbf{x}) = \int_{D_r(\mathbf{x}) \cap \Omega_F} B(\|\mathbf{x} - \mathbf{y}\|) d\mathbf{y}, \tag{14}$$

and derive a new expression for the rate of food accumulation:

$$F(\mathbf{x}, t) = \begin{cases} \int_{D_r(\mathbf{x}) \cap \Omega_F} \alpha \frac{B(\|\mathbf{x}-\mathbf{y}\|)}{Z(\mathbf{x}))} \psi(\mathbf{y})\, d\mathbf{y}, & \text{if } Z(\mathbf{x}) \neq 0, \\ 0, & \text{otherwise,} \end{cases} \tag{15}$$

where we note that the only case that leads to $Z(\mathbf{x}) = 0$ is the case where there is no forageable terrain near the forager, and thus it cannot collect any food. In the paper we focus on the case that B is a constant function but that does not have to be the case. Another reasonable choice for B could be a truncated Gaussian or normalized bump function, where the effort is then concentrated more around the forager's average location $\mathbf{y}(t)$.

B Deriving the Governing HJB Equations

We take a dynamic programming approach to determine the optimal policy for the foraging animal in each mode. For notational convenience, we will use $\mathbf{w} \in \mathbf{W} = \Omega \times [0, E] \times [0, T]$ to denote a generic state, and $\omega(t) = (\mathbf{y}(t), \xi(t), t)$ to denote the current state of the animal with combined dynamics $\frac{d\omega}{dt} = \left(\frac{d\mathbf{y}}{dt}, \frac{d\xi}{dt}, 1\right) = \mathbf{g}_{\eta(t)}(\omega)$. We now introduce the feedback control $\mathbf{a}(\mathbf{w}) = (\mathbf{a}_1(\mathbf{w}), \mathbf{a}_2(\mathbf{w})) : \mathbf{W} \mapsto \mathbb{S}^1 \times \mathbb{S}^1$, which prescribes the direction of motion in each mode as a function of the current state (for an animal to "use" policy $\mathbf{a}(\cdot)$ it must use $\mathbf{a}_1(\cdot)$ when it is in mode 1 and $\mathbf{a}_2(\cdot)$ when it is in mode 2.)

Let $J_m(\mathbf{w}, \mathbf{a}(\cdot))$ be the expected remaining reward for an animal in mode m, at state $\mathbf{w}$, and using policy $\mathbf{a}(\cdot)$. We can define J_m by explicitly computing the expected reward over the distribution of mode switches. Let $T_s > t$ be a random variable denoting the first time the animal is spotted by a predator after time some time t. Mode 1 ends at: T_s, if the animal is spotted for the end of the time horizon ($T_s < T$), any time at while $\xi(t) = 0$ (representing starvation), or at T if neither of those events occur first. We thus have the following expression for J_1:

$$J_1(\mathbf{w}, \mathbf{a}(\cdot)) = \mathbb{E}_{T_s}\left[\mathbb{1}_{[T_s \leq T]} J_2\left(\omega(T_s), \mathbf{a}(\cdot)\right) + \mathbb{1}_{[T_s > T]} J_1\left(\omega(T), \mathbf{a}(\cdot)\right)\right] \tag{16}$$

along with the boundary conditions $J_1(\mathbf{x}, 0, t, \mathbf{a}(\cdot)) = 0$ to represent death by starvation.

Similarly, define T_g and T_k to be the next times the animal escapes or is killed, respectively. Mode 2 ends if the predator gives up on the chase ($T_g \leq \min\{T_k, T\}$), the forager is killed ($T_k \leq \min\{T_g, T\}$), $\xi(t) = 0$ at any time (starvation), or the forager survives until T in mode 2 ($\min\{T_g, T_k\} > T$ and $\xi(t) > 0$). Thus, we have the following expression for J_2:

$$J_2(\mathbf{w}, \mathbf{a}(\cdot)) = \mathbb{E}_{(T_g, T_k)}\left[\mathbb{1}_{\left[T_g \leq \min\{T, T_k\}\right]} J_1\left(\omega(T_g), \mathbf{a}(\cdot)\right) + \mathbb{1}_{\left[\min\{T_g, T_k\} > T\right]} J_2\left(\omega(T), \mathbf{a}(\cdot)\right)\right] \tag{17}$$

with a similar starvation condition $J_2(\mathbf{x}, 0, t, \mathbf{a}(\cdot)) = 0$. We don't specify a reward for mode 3 because when the animal is dead the remaining reward is identically equal to 0. That is also the reason why equations (16) and (17) do not contain terms explicitly associated with switching to mode 3: any such terms are equal to 0. We will sometimes impose the additional constraint that the animal must

be located in the refuge at time T, i.e., $\mathbf{y}(T) \in \Omega_R$ where Ω_R is the set of points contained in the refuge. In those cases we use the terminal condition is $J_m(\mathbf{x}, e, T, \mathbf{a}(\cdot)) = U(e) \cdot \mathbb{1}_{[\mathbf{x} \in \Omega_H]}$ in both modes. For any areas through we the animal cannot travel, such as obstacles or points outside the domain, we impose the boundary condition $J_m(\mathbf{x}, e, t, \mathbf{a}(\cdot)) = 0$. Note that while we do not impose distinct terminal conditions for the different possible modes here, the modeling framework we propose would support that in general.

Let $u_1(\mathbf{w})$ and $u_2(\mathbf{w})$ be value functions for each mode, which represent the optimal expected reward-to-go for an animal at state $\mathbf{w}$ in mode 1 or 2 respectively. The value functions can be obtained by maximizing J_1 or J_2 over all admissible policies:

$$u_1(\mathbf{w}) = \sup_{\mathbf{a}(\cdot) \in \mathcal{A}} J_1(\mathbf{w}, \mathbf{a}(\cdot)) \tag{18}$$

$$u_2(\mathbf{w}) = \sup_{\mathbf{a}(\cdot) \in \mathcal{A}} J_2(\mathbf{w}, \mathbf{a}(\cdot)). \tag{19}$$

where $\mathcal{A}$ is the set of measurable functions from W to $\mathbb{S}^1 \times \mathbb{S}^1$. We now use Bellman's Optimality Principle to derive a system of governing PDEs for u_1 and u_2. Consider a small time interval τ, and assume that no mode switch occurs in $[t, t+\tau)$. The value functions must then satisfy the following optimality conditions:

$$u_1(\mathbf{w}) = \sup_{\mathbf{a}(\cdot) \in \mathcal{A}} \left\{ e^{-\mu_s(\mathbf{x})\tau} u_1\left(\mathbf{w}(t+\tau)\right) + \left(1 - e^{-\mu_s(\mathbf{x})\tau}\right) u_2\left(\mathbf{w}(t+\tau)\right) \right\} + o(\tau), \tag{20}$$

$$u_2(\mathbf{w}) = \sup_{\mathbf{a}(\cdot) \in \mathcal{A}} \left\{ e^{-\mu_k(\mathbf{x})\tau} \left(1 - e^{-\mu_g(\mathbf{x})\tau}\right) u_1\left(\mathbf{w}(t+\tau)\right) + e^{-\mu_k(\mathbf{x})\tau} e^{-\mu_g(\mathbf{x})\tau} u_2\left(\mathbf{w}(t+\tau)\right) \right\} + o(\tau) \tag{21}$$

where for a random event governed by rate μ we have approximated the probability that the event occurs over τ units as $1 - e^{-\mu\tau}$. We next perform a Taylor expansion. For the value function in mode

$\hat{m}$ and for an animal that is in mode m at time t we have

$$\begin{aligned} u_{\hat{m}}\left(\mathbf{w}(t+\tau)\right) &= u_{\hat{m}}\left(\mathbf{w} + \mathbf{g}_m(\mathbf{w})\tau + o(\tau)\right) \\ &= u_{\hat{m}}(\mathbf{w}) + \tau\Bigg(\boldsymbol{\nabla}_{\mathbf{x}} u_{\hat{m}}(\mathbf{w}) \cdot \mathbf{a}_m(\mathbf{w}) f_m(\mathbf{x}) \\ &+ \frac{\partial u_{\hat{m}}}{\partial \xi}\left(\left[\mathbb{1}_{[m=1]} F(\mathbf{x}) - K_m(\mathbf{x})\right]_*\right) + \frac{\partial u_{\hat{m}}}{\partial t}(\mathbf{w})\Bigg) + o(\tau). \end{aligned} \tag{22}$$

Substituting (22) into equations (20) and (21) and suppressing arguments, we find that the HJB PDEs for each mode are:

$$\mu_s(u_2 - u_1) + \frac{\partial u_1}{\partial t} + \left[F(\mathbf{x}) - K_m(\mathbf{x})\right]_* \frac{\partial u_1}{\partial e} + \max_{\mathbf{a}\in\mathbb{S}^1}\left\{\boldsymbol{\nabla}_{\mathbf{x}} u_1 \cdot \mathbf{a} f_1\right\} = 0 \tag{23}$$

$$\mu_g(u_1 - u_2) - \mu_k u_2 + \frac{\partial u_2}{\partial t} - K_2 \frac{\partial u_2}{\partial e} + \max_{\mathbf{a}\in\mathbb{S}^1}\left\{\boldsymbol{\nabla}_{\mathbf{x}} u_2 \cdot \mathbf{a} f_2\right\} = 0. \tag{24}$$

Using the fact that $\mathbf{a}_m(\mathbf{w}) = \boldsymbol{\nabla} u_m(\mathbf{w}) / \|\boldsymbol{\nabla} u_m(\mathbf{w})\|$ satisfies the maximization, we can rewrite this as

$$\mu_s(u_2 - u_1) + \frac{\partial u_1}{\partial t} + \left[F(\mathbf{x}) - K_m(\mathbf{x})\right]_* \frac{\partial u_1}{\partial e} + f_1 |\boldsymbol{\nabla}_{\mathbf{x}} u_1| = 0 \tag{25}$$

$$\mu_g(u_1 - u_2) - \mu_k u_2 + \frac{\partial u_2}{\partial t} - K_2 \frac{\partial u_2}{\partial e} + f_2 |\boldsymbol{\nabla}_{\mathbf{x}} u_2| = 0. \tag{26}$$

Finally, the boundary and terminal conditions for the value functions come directly from those for J_m:

$$u_m(\mathbf{x}, 0, t) = 0 \tag{27}$$

$$u_m(\mathbf{x}, e, T) = U(e). \tag{28}$$

C Details of Multistage Model

Here we present the mathematical details of the multistage foraging model proposed in the main text. For simplicity of presentation, we present this model for a food density function $\psi(\mathbf{x})$ that does not have any inherent time dependence, but the model supports time-dependent food densities as well. To formalize this approach, we first divide the planning horizon into N "stages" of fixed length

$S = T/(N-1)$ denoted by $[T_0, T_1]$, $[T_1, T_2]$, ..., $[T_{N-1}, T_N]$, where $T_0 = 0$ and $T_N = T$. For stage $n \in \{0, \ldots, N-1\}$, mode $m \in \{0, 1\}$, and $s \in [0, S]$ we define the value function $u_m^n(\mathbf{x}, e, s)$ which represents the optimal reward-to-go until the end of stage n from the current state and mode. To compute these value functions we alternate between solving a version of equations (25) and (26) over the interval $s \in [0, S]$ and simulating a realization of the forager's optimal trajectory over a single stage. To initialize, we solve for u_m^0 using the initial food density $\psi(\mathbf{x})$ to compute $F^0(\mathbf{x})$, the food availability during the first stage. Then, we trace a realization of the trajectory $\mathbf{y}(t)$ for $t \in [T_0, T_1]$ using the optimal feedback policy prescribed by u_m^0.

An animal truly collecting food according to our model should deplete the available food according to

$$\frac{\partial}{\partial t}\Psi(\mathbf{x}, t) = \begin{cases} -\beta \left(\mathbb{1}_{[m(t)=1]} \alpha \frac{B\left(\|\mathbf{x}-\mathbf{y}(t)\|\right)}{Z(\mathbf{y}(t))} \right) \Psi(\mathbf{x}, t), & \text{if } Z(\mathbf{y}(t)) \neq 0, \\ 0, & \text{otherwise,} \end{cases} \tag{29}$$

$$\Psi(\mathbf{x}, 0) = \psi(\mathbf{x}), \tag{30}$$

where $\Psi(\mathbf{x}, t)$ is the now time-dependent food density and

$$\beta = \begin{cases} \mathbb{1}_{[\xi(t)=E]} \min\left\{1, \frac{K_1(\mathbf{y}(t))}{F(\mathbf{y}(t))}\right\} & \text{if } \xi(t) = E, \\ 1 & \text{otherwise} \end{cases} \tag{31}$$

captures the fact that the animal can only eat a limited amount when it is at maximum energy. For a known trajectory $\mathbf{y}(t)$, solving for the resulting $\Psi(\mathbf{x}, t)$ is simple solving a linear ODE. In general, once we have solved for $u_m^0, \ldots, u_m^{n-1}$ and $\Psi(\mathbf{x}, t)$ for $t \in [T_0, T_n]$ we can compute u_m^n in the same manner by setting $\psi^n(\mathbf{x}) = \Psi(\mathbf{x}, T_n)$, using $\psi^n(\mathbf{x})$ to compute $F^n(\mathbf{x})$, and solving equations (25) and (26). This assumes that we have chosen a terminal reward for each stage, which we discuss in the main text.

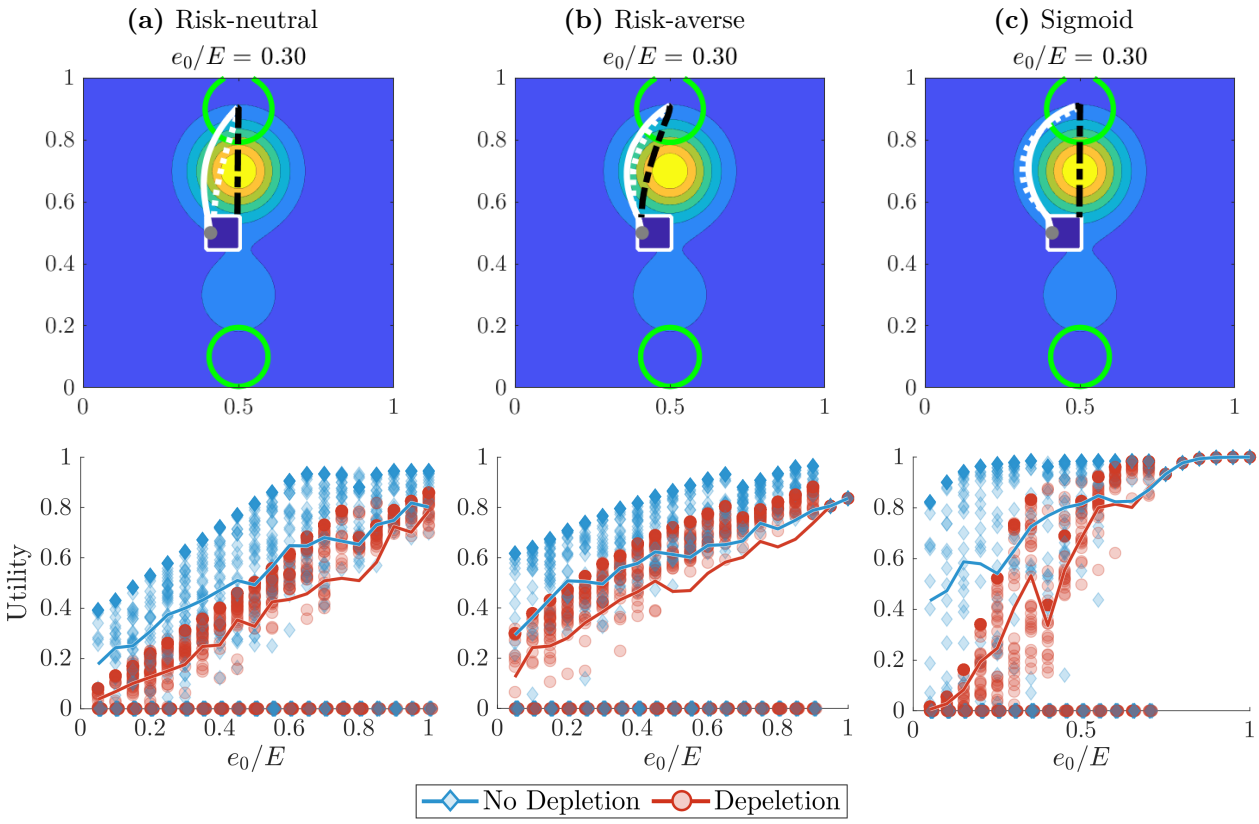

Figure 10: The impact of food depletion on a rational forager that *does not plan for it.* When the forager depletes food, it receives less energy than it planned for and spends longer at the feeding ground to compensate. It subsequently takes shorter, riskier return trips to the refuge (top row, black dot-dash line). Outbound trips (solid white line) are identical because no depletion has occurred yet. The bottom row shows realized terminal utilities with (orange circle) and without (blue diamond) food depletion, which vary for fixed e_0 due to random interactions with predators. A forager that dies via starvation or being killed by a predator receives a utility of 0 (points at $U = 0$ have been horizontally perturbed to indicate frequency). As expected, utility is on average lower when the forager depletes food (orange and blue solid lines). Overall, food depletion may significantly impact the forager's optimal trajectory and final utility in a way that is nonuniform across utility functions and initial energy levels. However, exactly accounting for the impact of food depletion on an optimal forager is computationally infeasible and may not be ecologically sound. We instead propose a model in which the forager makes inaccurate short-term predictions about its environment, but periodically updates them to account for the food it has depleted.

C.1 Example with Food Depletion

We illustrate our multistage approach on a simple environment with two food-rich areas and a safe refuge (Figure 11a, again similar in spirit to a discrete two-patch model). While the food availability and distance to the refuge are identical in both "patches," we think of the bottom patch as having "cover" that reduces the likelihood of the forager being killed there (Figure 11). We assume that

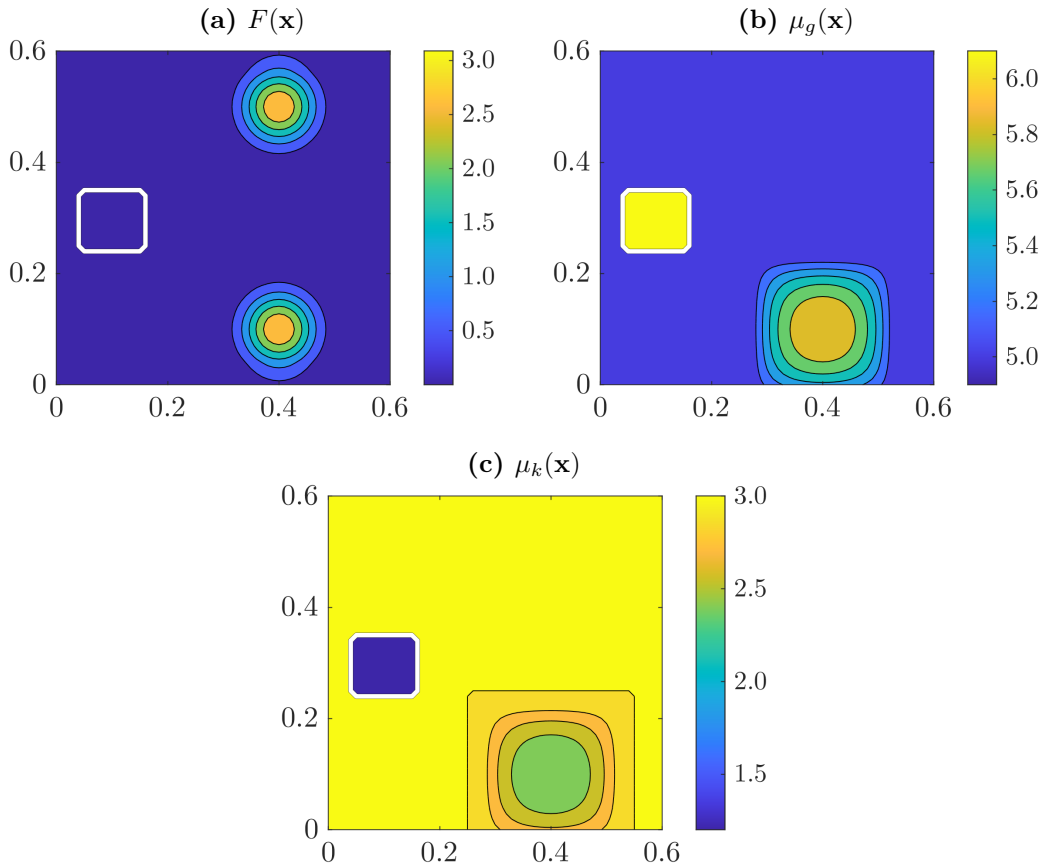

Figure 11: A "two-patch" environment where one patch is strictly better than the other. Two regions of equal food availability are equidistant from the refuge (subfigure **(a)**). The lower patch is assumed to have cover, which increases the chance of escaping (subfigure **(b)**) and decreases the chance of being killed (subfigure **(c)**) when there is a predator. A rational forager that starts in the refuge and does not deplete food will only exploit the lower patch, but our "multistage model" correctly predicts that it can be optimal to exploit both patches when depletion occurs (see Figures 9 - 13). Exact expressions for each quantity are provided in Section E.

the animal travels with constant speed $f_m(\mathbf{x}) = 5$, expends energy at a constant rate $K_m(\mathbf{x}) = 0.8$, and faces constant spotting rate $\mu_s(\mathbf{x}) = 3$ outside of the refuge. Inside of the refuge, the forager cannot be spotted or killed ($\mu_s(\mathbf{x}) = \mu_k(\mathbf{x}) = 0$) and the predator is more likely to give up on the chase ($\mu_g(\mathbf{x})$ increased by a factor of 10). Since the two patches are otherwise identical, a rational forager without knowledge of food depletion would only exploit the bottom patch (assuming it starts within the refuge) and would thus do poorly if significant food depletion actually occurs. However, as explained in the main text, a depletion-aware forager may sometimes exploit both regions, though still favoring the lower.

In addition to the risk-neutral forager discussed in the main text, Figures 12 and 13 present similar

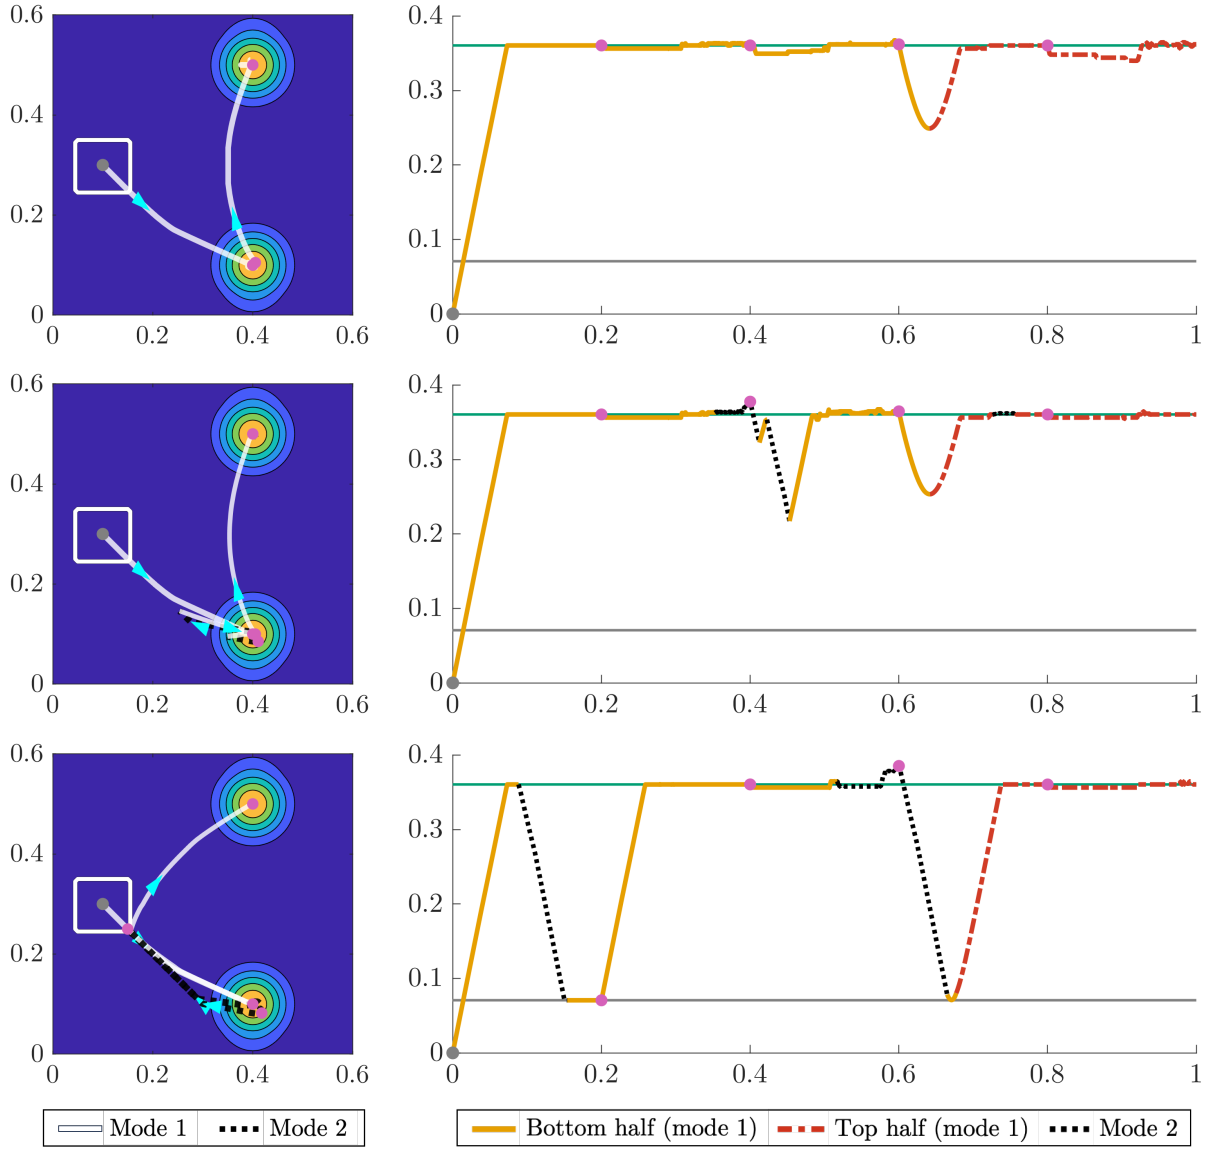

Figure 12: A myopic forager that periodically becomes aware of food depletion exploits both food patches (Risk-averse). Formatted as in Figure 9. In the absence of predators, the behavior of the risk-averse planner is similar to that of the risk-neutral planner, though it takes a longer detour to remain close to the refuge on the way to the second food patch (top row). If the forager is spotted by a predator but successfully escapes before reaching the refuge, it may return to its original foraging patch (middle row). In the third row, the forager escapes to the refuge twice, the first chase is early enough that it is optimal to return to the lower patch (bottom row). After the second chase, the forager switches to the upper patch due to the effects of food depletion.

stories but for risk-averse and sigmoid planners. In the absence of predators, the behavior of the risk-averse planner is similar to that of the risk-neutral planner, though it takes a longer detour between the patches to remain close to the refuge (Figure 12, top row). When not spotted by a predator, the sigmoid planner only exploits the safer patch before returning to the refuge to wait out the end

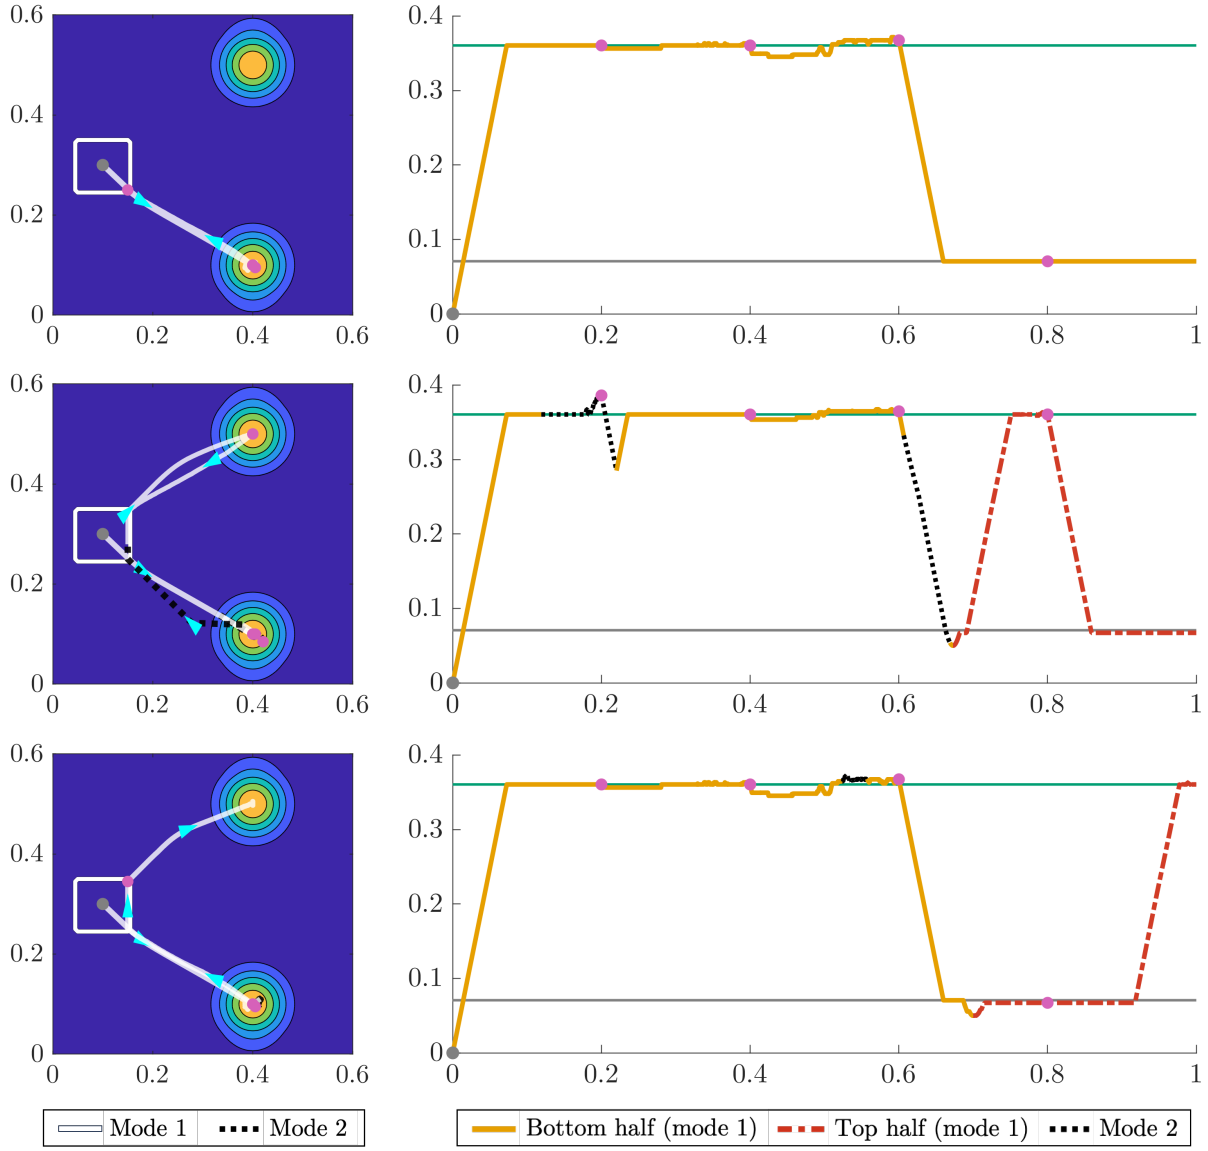

Figure 13: A myopic forager that periodically becomes aware of food depletion exploits both food patches (Sigmoid). Formatted as in Figure 9. The sigmoid planner is the most cautious of the foragers we model. In the absence of predators, it only exploits the safer patch (top row) before returning to the refuge at the beginning of the fourth stage, a clear distinction from the other two planners. However, the arrival of a predator can cause the forager to return to the refuge earlier than planned, and from there is may be optimal to switch to the upper patch (middle row). In the bottom row, a forager that has only a brief encounter with a predator later chooses to travel to the upper patch to make up for lost foraging time.

of the planning horizon. However, it if travels to the refuge to escape a predator or loses foraging time due to a chase it may exploit the upper patch as well (Figure 13, middle and bottom rows). In this setting, the multistage model can capture more nuanced behavior than a model without food depletion, desptie the limitations associated with stage length and terminal condition.

D Numerical Methods

D.1 HJB PDE Discretization

We solve equations (25) and (26) backwards in time using an explicit finite-difference scheme. We denote the grid spacing using $\Delta x = \Delta y = h$, Δe, and Δt. Further, we denote a point in our discretized (x, y, e, t) space as (x_i, y_j, e_k, t_l) where $x_i = ih$, $y_j = jh$, $e_k = k\Delta e$, and $t_l = l\Delta t$. While we focus on the case that $\Delta x = \Delta y$ for simplicity of presentation, the scheme can be generalized to the case that $\Delta x \neq \Delta y$ in a straightforward manner.

Let U_1 and U_2 to represent the solutions to the discretized system of equations, such that

$$U_1^{i,j,k,l} \approx u_1(x_i, y_j, e_k, t_l) \qquad \text{and} \qquad U_2^{i,j,k,l} \approx u_2(x_i, y_j, e_k, t_l). \tag{32}$$

We can then similarly write $\mu_s^{i,j} = \mu_s(x_i, y_j)$, $\mu_g^{i,j} = \mu_g(x_i, y_j)$, $\mu_k^{i,j} = \mu_k(x_i, y_j)$, $F^{i,j} = F(x_i, y_j)$, $f_m^{i,j} = f_m(x_i, y_j)$, and $K_m^{i,j} = K_m(x_i, y_j)$.

In addition to solving backwards in time, we use an *upwind* discretization in space and energy. First, we can define the following one-sided difference operators:

$$\mathcal{D}_x^{\pm} U_m^{i,j,k,l} = \frac{U_m^{i\pm1,j,k,l} - U_m^{i,j,k,l}}{\pm\Delta x} \tag{33}$$

with respect to x, where operators for y, e, and t can be defined in the same way. We can use these one-sided difference operators to construct upwind difference operators for the x, y, and e coordinates. Starting with x and y, the upwind operator is defined as

$$\mathcal{D}_x U_m^{i,j,k,l} = \max\left\{\mathcal{D}_x^{+} U_m^{i,j,k,l}, -\mathcal{D}_x^{-} U_m^{i,j,k,l}, 0\right\} \quad \text{and} \quad \mathcal{D}_y U_m^{i,j,k,l} = \max\left\{\mathcal{D}_y^{+} U_m^{i,j,k,l}, -\mathcal{D}_y^{-} U_m^{i,j,k,l}, 0\right\}. \tag{34}$$

For the energy dimension, the upwind operator can be defined a priori as long as the mode is known. If $m = 1$ the animal collects and expends energy, so the upwind direction depends on the sign of

$F^{i,j} - K_1^{i,j}$. Since both functions are known, we can define the following upwind operator:

$$\mathcal{D}_e U_1^{i,j,k,l} = \begin{cases} \mathcal{D}_e^+ U_1^{i,j,k,l} & F^{i,j} - K_1^{i,j} < 0 \\ \mathcal{D}_e^- U_1^{i,j,k,l} & else \end{cases}. \tag{35}$$

If $m = 2$ then the animal cannot collect food and the energy level must be nonincreasing so we simply have

$$\mathcal{D}_e U_2^{i,j,k,l} = \mathcal{D}_e^+ U_2^{i,j,k,l} \tag{36}$$

as our upwind approximation.

Putting all of this together, we can write the fully discretized equations as

$$\begin{aligned} 0 &= \mu_s^{i,j}\left(U_2^{i,j,k,l} - U_1^{i,j,k,l}\right) + \mathcal{D}_t^- U_1^{i,j,k,l} + \left(F^{i,j} - K_1^{i,j}\right)\mathcal{D}_e U_1^{i,j,k,l} \\ &\quad + f_1^{i,j}\sqrt{\left(\mathcal{D}_x U_1^{i,j,k,l}\right)^2 + \left(\mathcal{D}_y U_1^{i,j,k,l}\right)^2} \end{aligned} \tag{37}$$

$$\begin{aligned} 0 &= \mu_g^{i,j}\left(U_1^{i,j,k,l} - U_2^{i,j,k,l}\right) - \mu_k^{i,j} U_2^{i,j,k,l} + \mathcal{D}_t^- U_2^{i,j,k,l} - K_2^{i,j}\mathcal{D}_e U_2^{i,j,k,l} \\ &\quad + f_2^{i,j}\sqrt{\left(\mathcal{D}_x U_2^{i,j,k,l}\right)^2 + \left(\mathcal{D}_y U_2^{i,j,k,l}\right)^2}. \end{aligned} \tag{38}$$

We use $\mathcal{D}_t^- U_m^{i,j,k,l}$ as the derivative approximation in the time dimension in order to obtain an explicit update for $U_m^{i,j,k,l-1}$ in terms of quantities in the lth slice of the value function. Rearranging equations (37) and (38) we get the following update rule:

$$\begin{aligned} U_1^{i,j,k,l-1} &= U_1^{i,j,k,l} + \Delta t\Bigg[\mu_s^{i,j}\left(U_2^{i,j,k,l} - U_1^{i,j,k,l}\right) \\ &\quad + \left(F^{i,j} - K_1^{i,j}\right)\mathcal{D}_e U_1^{i,j,k,l} + f_1^{i,j}\sqrt{\left(\mathcal{D}_x U_1^{i,j,k,l}\right)^2 + \left(\mathcal{D}_y U_1^{i,j,k,l}\right)^2}\Bigg] \end{aligned} \tag{39}$$

$$\begin{aligned} U_2^{i,j,k,l-1} &= U_2^{i,j,k,l} + \Delta t\Bigg[\mu_g^{i,j}\left(U_1^{i,j,k,l} - U_2^{i,j,k,l}\right) - \mu_k^{i,j} U_2^{i,j,k,l} \\ &\quad - K_2^{i,j}\mathcal{D}_e U_2^{i,j,k,l} + f_2^{i,j}\sqrt{\left(\mathcal{D}_x U_2^{i,j,k,l}\right)^2 + \left(\mathcal{D}_y U_2^{i,j,k,l}\right)^2}\Bigg]. \end{aligned} \tag{40}$$

D.2 CFL Conditions

Since we are using an explicit discretization in time, the stability of the numerical scheme is governed by a Courant-Friedrichs-Lewy (CFL) condition. This condition is based on the dynamics of the problem, and since each mode has its own dynamics, each will also have its own CFL condition. For simplicity, we use the same time step for both modes in our numerical scheme, so our choice of Δt is based on the more restrictive of the two conditions. We now derive the CFL conditions for modes 1 and 2.

Mode 1: The analytic domain of dependence for a gridpoint (x_i, y_j, e_k, t_l) at the slice $t + \Delta t$ is the circle in (x, y, e)-space centered at $(0, 0, \tilde{e})$ with radius R, where $\tilde{e} = |F - K_1|\Delta t$ and $R = f_1^{i,j}\Delta t$. For the CFL condition to be satisfied, this circle must lie within the simplex generated by the 6 grid-points surrounding (x, y, e). In the first octant of (x, y, e)-space, the simplex is given by the boundary $\frac{x}{h} + \frac{y}{h} + \frac{z}{\Delta e} = 1$, where $h = \Delta x = \Delta y$. We now compute the maximum radius of a circle centered at $(0, 0, \tilde{e})$ parallel to the xy-plane lying within the simplex. Note that the intersection of the plane $e = \tilde{e}$ with the simplex in the first octant is given by $(\frac{h(\Delta e - \tilde{e})}{\Delta e}, \frac{h(\Delta e - \tilde{e})}{\Delta e}, \tilde{e})$. Thus the maximum radius is given by $R_{max} = \frac{h(\Delta e - \tilde{e})}{\sqrt{2}\Delta e}$. Rearranging $R \leq R_{max}$, we find

$$\left(f_1 + \frac{h|F - K_1|}{\sqrt{2}\Delta e}\right)\Delta t \leq \frac{h}{\sqrt{2}} \tag{41}$$

Mode 2:

This is the same as mode 1, except the (absolute value of) rate of change of energy is $|-K_2| = K_2$, and we use f_2 instead of f_1 for the speed in physical space. Thus the CFL condition becomes

$$\left(f_2 + \frac{hK_2}{\sqrt{2}\Delta e}\right)\Delta t \leq \frac{h}{\sqrt{2}} \tag{42}$$

Combining mode 1 + mode 2, we find the combined CFL condition

$$\max\left\{f_1 + \frac{h|F - K_1|}{\sqrt{2}\Delta e}, f_2 + \frac{hK_2}{\sqrt{2}\Delta e}\right\} \cdot \Delta t \leq \frac{h}{\sqrt{2}} \tag{43}$$

D.3 Path Tracing and Food Depletion

To trace the optimal trajectories, we simulate the position and energy dynamics using a forward Euler scheme with time step τ (which may differ from Δt use to solve the PDEs). We use a grid search over directions of motion to determine the optimal direction, θ. Slightly abusing notation, if (x_l, y_l, e_l, t_l) is the numeric position of the forager after l steps of size τ, the update rule for position is thus:

$$x_{l+1} = x_l + \cos(\theta_l) f_{m_l}(x_l, y_l)\tau \tag{44}$$

$$y_{l+1} = y_l + \sin(\theta_l) f_{m_l}(x_l, y_l)\tau \tag{45}$$

where

$$\theta_l \approx \underset{\theta \in [0,2\pi)}{\arg\max} \left\{ u_{m_l}(x_{l+1}, y_{l+1}, e_{l+1}, t_{l+1}) \right\} \tag{46}$$

For energy, without food depletiong, the update rule is given by

$$e_{l+1} = \max\left\{0, \min\left\{E, e_l + (F(x_l, y_l) - K_1(x_l, y_l))\tau\right\}\right\} \tag{47}$$

in mode 1 and

$$e_{l+1} = \max\left\{0, e_l - K_2(x_l, y_l)\tau\right\}. \tag{48}$$

in mode 2. We linearly interpolate all functions when evaluating them between gridpoints, except for $F(\mathbf{x})$, which we compute via numerical quadrature for each location.

When considering food depletion, the update for the energy is more complicated. To compute e_{l+1} in mode 1, we deplete the food in the environment according the animal's position (x_l, y_l). To deplete food at a given gridpoint (x, y), we consider two cases. First assume that the animal cannot reach E from its current energy level. Because of the CFL condition governing Δt this is equivalent to the condition $E - e_l > \Delta e$. In this case, we approximate the solution to equation (29) on the small

interval $[t_l, t_{l+1}]$ using a forward Euler step

$$\Psi_n(\mathbf{x}, t_{l+1}) = \Psi_n(\mathbf{x}, t_l)\left(1 - \alpha\frac{B(\|(x,y)-(x_l,y_l)\|)}{Z((x_l,y_l))}\tau\right) \tag{49}$$

which serves as our update equation for $\Psi_m(\mathbf{x}, t_{j+1})$. If $E - e_l \leq \Delta e$, it is possible that the animal will reach E in the next time step, and we instead use the update

$$\Psi_n(\mathbf{x}, t_{l+1}) = \Psi_n(\mathbf{x}, t_l)\left(1 - \left(\min\left\{1, \frac{E-(e_l - K_1(x_l,y_l)\tau)}{F(x_l,y_l)\tau}\right\}\right)\left(\alpha\frac{B(\|(x,y)-(x_l,y_l)\|)}{Z((x_l,y_l))}\right)\tau\right) \tag{50}$$

We make use of the fact that the food depleted across the entire environment is equivalent to the food accumulated by the animal, and set

$$e_{l+1} = e_l - K_1(x_l,y_l)\tau + \sum_{(x_i,y_j)\in P_l}\left((\Psi_n(x_i,y_j,t_l) - \Psi_n(x_i,y_j,t_{l+1})\right) \tag{51}$$

where P_l is the set of grid points for which $\Psi_n(x_i, y_j, t_l) \neq \Psi_n(x_l, y_l, t_{l+1})$. In mode 2 the animal does not eat so there is no food depletion and we can use the update from equation (48).

Since we use a forward Euler step to simulate the ODE for Ψ_n this imposes an additional CFL to ensure stability. We must choose τ such that

$$\alpha\frac{B(\|(x,y)-(x_l,y_l)\|)}{Z((x_l,y_l))}\tau < 1 \tag{52}$$

In practice, we assume that B is continuous and nonincreasing on $[0, r]$ and τ takes the form $\tau = \frac{\Delta t}{\gamma}$. We can then substitute the stricter condition

$$\alpha\frac{B(r)}{B(0)h^2}\frac{\Delta t}{\gamma} < 1. \tag{53}$$

and choose γ such that

$$\gamma > \alpha\frac{B(r)\Delta t}{B(0)h^2}. \tag{54}$$

E Formulas for Numerical Experiments

E.1 Balancing Risk and Reward Environment

In this example, we compute the food availability based on food density

$$\psi(\mathbf{x}) = 10000 \exp\left(-\left(2\sigma^2\right)^{-1}\left((x-0.5)^2 + (y-0.9)^2\right)\right) + 6000 \exp\left(-\left(2\sigma^2\right)^{-1}\left((x-0.5)^2 + (y-0.1)^2\right)\right)$$

with $\sigma = 0.05$. We assume that the animal has food radius $r = 0.01$, harvest rate $\alpha = 0.001$, and that the food kernel $B(d)$ is the characteristic function of $D_r(0)$. The animal's initial position is $(x_0, y_0) = (0.41, 0.5)$. For a fixed risk premium ρ, the spotting rate is computed as

$$\mu_s(\mathbf{x}) = 0.5 + (0.5+\rho) \exp\left(-\left(2\sigma^2\right)^{-1}\left((x-0.5)^2 + (y-0.7)^2\right)\right) + 0.5 \exp\left(-\left(2\sigma^2\right)^{-1}\left((x-0.5)^2 + (y-0.3)^2\right)\right)$$

for $\sigma = 0.1$.

E.2 Multistage Model Environment

We compute the food availability based on food density

$$\psi(\mathbf{x}) = 500 \sum_{j=1}^{2} \exp\left(-\left(2\sigma^2\right)^{-1}\left((x-\hat{x}_j)^2 + (y-\hat{y}_j)^2\right)\right)$$

for $(\hat{x}_j, \hat{y}_j) \in \{(0.4, 0.1), (0.4, 0.5))\}$ and $\sigma = 0.035$. We further assume that the animal has food radius $r = 0.05$, harvest rate $\alpha = 0.01$, and that the food kernel $B(d)$ is the characteristic function of $D_r(0)$. For all experiments, we set the initial energy to be $e_0 = 1.0$ and the initial position to $(x_0, y_0) = (0.1, 0.3)$. The give up rate and kill rate are computed as follows:

$$\mu_g(\mathbf{x}) = \begin{cases} 5 + \exp\left(2 - \left(\frac{1}{1-((x-\hat{x}_1)/0.15)^2} + \frac{1}{1-((y-\hat{y}_1)/0.15)^2}\right)\right), & (x,y) \in R_1, \\ 5, & else, \end{cases}$$

and

$$\mu_k(\mathbf{x}) = \begin{cases} 3 - 0.6 \exp\left(2 - \left(\frac{1}{1-((x-\hat{x}_1)/0.15)^2} + \frac{1}{1-((y-\hat{y}_1)/0.15)^2}\right)\right), & (x,y) \in R_1, \\ 3, & else, \end{cases}$$

where R_1 is the square of width 0.15 centered at $(0.4, 0.1)$.