INTERFACE

royalsocietypublishing.org/journal/rsif

Research

Cite this article: Kilpatrick ZP, Davidson JD, El Hady A. 2021 Uncertainty drives deviations in normative foraging decision strategies. *J. R. Soc. Interface* **18**: 20210337. https://doi.org/10.1098/rsif.2021.0337

Received: 24 April 2021
Accepted: 22 June 2021

Subject Category:
Life Sciences–Mathematics interface

Subject Areas:
biomathematics

Keywords:
foraging, normative theory, evidence accumulation, decision-making

Authors for correspondence:
Zachary P. Kilpatrick
e-mail: zpkilpat@colorado.edu
Jacob D. Davidson
e-mail: jdavidson@ab.mpg.de
Ahmed El Hady
e-mail: ahady@princeton.edu

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

Uncertainty drives deviations in normative foraging decision strategies

Zachary P. Kilpatrick[1,2], Jacob D. Davidson[3,4,5] and Ahmed El Hady[6]

[1]Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, USA
[2]Department of Physiology and Biophysics, University of Colorado School of Medicine, Aurora, CO, USA
[3]Department of Collective Behaviour, Max Planck Institute of Animal Behavior, 78464 Konstanz, Germany
[4]Department of Biology, and [5]Centre for the Advanced Study of Collective Behaviour, University of Konstanz, 78464 Konstanz, Germany
[6]Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08540, USA

ZPK, 0000-0002-2835-9416; JDD, 0000-0001-9206-085X; AEH, 0000-0002-0045-8017

Nearly all animals forage to acquire energy for survival through efficient search and resource harvesting. Patch exploitation is a canonical foraging behaviour, but there is a need for more tractable and understandable mathematical models describing how foragers deal with uncertainty. To provide such a treatment, we develop a normative theory of patch foraging decisions, proposing mechanisms by which foraging behaviours emerge in the face of uncertainty. Our model foragers statistically and sequentially infer patch resource yields using Bayesian updating based on their resource encounter history. A decision to leave a patch is triggered when the certainty of the patch type or the estimated yield of the patch falls below a threshold. The time scale over which uncertainty in resource availability persists strongly impacts behavioural variables like patch residence times and decision rules determining patch departures. When patch depletion is slow, as in habitat selection, departures are characterized by a reduction of uncertainty, suggesting that the forager resides in a low-yielding patch. Uncertainty leads patch-exploiting foragers to overharvest (underharvest) patches with initially low (high) resource yields in comparison with predictions of the marginal value theorem. These results extend optimal foraging theory and motivate a variety of behavioural experiments investigating patch foraging behaviour.

1. Introduction

Foraging is performed by many different species [1–5] and engages cognitive computations such as learning of resource distributions across spatio-temporal scales, route planning and decision-making [6]. Comparing species, one can ask how these integrated processes have been shaped by natural selection to optimize returns in the face of environmental and physiological constraints [6,7]. Foraging thus provides the opportunity to study and quantify how both evolution and neural circuitry shape a natural behaviour [8–11].

In natural landscapes, foraging involves a decision hierarchy that unfolds across multiple length and time scales, which consider both *where to forage* as well as *how long to exploit a certain resource* [12,13]. On long time scales, animals accumulate evidence to choose which of a collection of large areas they will dwell in and forage, during which their activity does not appreciably change the resource landscape. Following previous work, we refer to this as 'habitat choice' [12]. Within habitats, animals exploit resources on shorter time scales, while their activity depletes resources in visits to localized regions. We refer to this behaviour as 'patch exploitation' or 'patch leaving'. Questions of where to forage and how long to exploit local patches of resource constitute a multi-level framework for examining behaviour across spatial scales; still larger scales consider the home range of an individual, as well as the species range [13,14]. For both habitat choice and patch leaving, the local regions of

© 2021 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License http://creativecommons.org/licenses/by/4.0/, which permits unrestricted use, provided the original author and source are credited.

THE ROYAL SOCIETY PUBLISHING

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

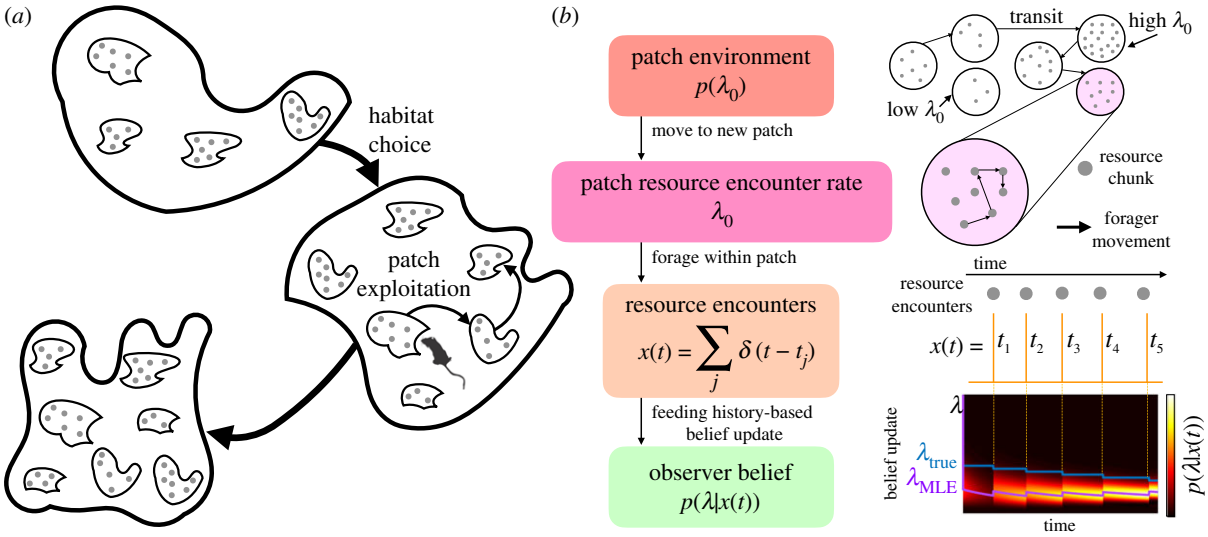

Figure 1. Patch-departure tasks and model. (*a*) Task environments: on long time scales, an animal decides between habitats whose resource yields change slowly; on shorter time scales, the animal exploits patches whose resources are depleted more quickly. In our model analysis, we assume that the forager is solving one of these problems at a time, but not both simultaneously. (*b*) Ideal observer foraging model: the initial yield of the patch is drawn from the distribution $p(\lambda_0)$, generating random resource encounter times $t_{1:K}$, and updating the belief of the current resource yield rate $\lambda(t)$ for the patch. We illustrate the movement of the forager and subsequent time series of resource encounters $x(t)$, resulting in a refinement of the posterior $p(\lambda|x(t))$. The maximum-likelihood estimate λ_{MLE} approaches the true λ_{true} yield rate over time.

an environment can be conceptualized as a 'patch'; the key difference is in whether or not the forager's activity impacts the resource availability in the landscape. As a decision problem, both are sequential choice processes, so the forager does not make a choice between discrete alternatives that are presented simultaneously, but rather only receives evidence from the current patch and must decide whether to stay or go [15]. Thus, habitat choice and patch leaving are related but differ in the length and time scales involved (figure 1*a*). However, most theoretical work has considered these two problems separately.

Regarding habitat choice, the optimal behaviour is clear: the forager should locate and spend as much time as possible in the habitat patch that maximizes fitness outcomes. Observational studies of habitat choice often consider the combined effects of multiple factors to ask how well they predict observed habitat use (e.g. [16]). Theoretical and experimental studies have investigated multiple factors, such as density-dependent effects predicted by the ideal free distribution when there are multiple foragers on a landscape [17] and how perceptual constraints may lead to deviations from optimal choices [18]. Our focus here is specifically on understanding an accumulation process by which individuals use information to reach decisions on habitat choice, and deriving this result in a mathematically tractable model that can more precisely parameterize strategy and performance in a wide range of environments.

Patch leaving considers shorter time scales, where the forager substantially depletes the patch during its visit and must subsequently decide when to leave the current patch in search of another. A basic result in behavioural ecology, the marginal value theorem (MVT), states that an animal can optimize the resource intake rate by leaving its current patch when the estimated within-patch resource yield rate falls below the global average resource intake rate of the environment [19]. While this theory has been validated in multiple behavioural studies [20–29], it does not explicitly indicate how beliefs about environmental features (e.g. resource distribution) are accumulated over time or how decision rules might be applied to these beliefs in the form of concrete dynamical equations. To explain how foraging decisions are shaped by the presence and reduction of uncertainty based on resource encounters, it is useful to have normative Bayesian models of patch leaving that ask how animals use limited information to make foraging decisions. However, those that do tend to consider a narrow range of environmental conditions [30,31–46]. General models that have been proposed are challenging to analyse mathematically [47], making it difficult to reveal how environmental parameters shape an optimal forager's strategy and yield in many different environments.

The aim of our study is thus to develop a Bayesian framework of foraging behaviour, treating decisions as a statistical inference problem and connecting normative theory of foraging decisions with mechanistic evidence accumulation models [48]. Normative strategies provide a best-case scenario for a particular objective, which can then be used as a touchstone for comparison with heuristics and actual animal behaviours. Animals may actually use evidence accumulation mechanisms that approximate Bayes' optimal strategies, such as energy measurement via the passage of food through the gut [49], so we consider our model to be an abstraction that such physical strategies can be measured against. We first define a general mathematical framework to model patch-foraging decisions that applies at different time scales with regards to search and depletion of the resource. These represent the different ecological decision cases described above: 'habitat choice' and 'patch leaving'. Because these cases are only separated in the time scale of resource depletion, we treat both with the framework of patch foraging as an evidence accumulation process whereby a threshold on available evidence triggers a decision.

Using several mathematically tractable cases in which probabilistic updating based on the receipt of resources within a patch can be modelled by stochastic differential equations (SDEs), we determine patch-leaving statistics via solutions to first-passage time problems. We thus obtain analytical expressions for optimal decision thresholds that

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

connect to observable quantities of interest, including patch residence time, travel time, resource consumption over time and patch yield rate over time. For both habitat choice and patch leaving, we show that, in uncertain or resource-poor environments, uncertainty causes even an ideal Bayesian observer to tend to stay too long in low-yielding patches (overharvesting) and not long enough in high-yielding patches (underharvesting). Past studies have suggested that such deviations can also arise from patch-discrimination limits [18] or behavioural state dependence [50]. Though our work is not the first to demonstrate that over- and underharvesting can arise from uncertainty in statistical inference models [51,52], we are able to show how these trends vary across a wide range of environments owing to our model's mathematical tractability. By establishing a general Bayesian framework for patch foraging at multiple scales, our study provides a platform to study behavioural and neural mechanisms of naturalistic decision-making akin to how trained decision-making behaviour is studied within systems neuroscience [8,53].

2. Sequential sampling model framework

The patch-foraging model framework, which describes both habitat choice and patch leaving, considers an animal searching its environment, which contains distributed resource patches (figure 1*a*). When the animal enters a patch, it consumes resources within the patch. In the case of habitat choice, we assume that resources are depleted slowly enough that the depletion is negligible, whereas for the patch leaving problem resources are depleted. We represent the decision process to leave a patch via a sequential sampling model for an ideal observer's posterior of its current patch's yield rate, $\lambda(t)$. This assumes that an animal learns over time the yield of the patch it is currently in and to decide if and when it should leave and search for another patch.

The initial yield rate λ_0^k determines the rate at which the animal initially encounters a resource in the patch and is drawn from the distribution $p(\lambda_0)$. We assume that the forager knows and initializes its belief with the prior $p(\lambda_0)$ when arriving in a patch (figure 1*b*). This simplifying assumption allows us to obtain tractable solutions. Considering randomly timed resource encounters within a patch, we use a Poisson rate $\lambda(t) = \lambda_0 - \rho K(t)$ generating exponentially distributed waiting times between encounters (as in random search [54]) that decreases with $K(t)$, the number of resource encounters so far, where ρ is the impact of each resource encounter on the underlying yield rate of the patch. Resource encounter history can be described by the summed sequence of encounters, each at time t_j: $x(t) = K'(t) = \sum_{j=1}^{K(t)} \delta(t - t_j)$. An ideal forager performs a Bayesian update of its belief about the current patch yield rate λ,

$$p(\lambda|x(t)) = p(x(t)|\lambda)\frac{p_0(\lambda + K(t)\rho)}{p(x(t))} \propto \frac{(\lambda/\rho + K)!}{(\lambda/\rho)!} e^{-\lambda t} p_0(\lambda + K\rho), \quad \lambda \geq 0. \tag{2.1}$$

In general, resource encounters both: (i) give evidence of higher yield rates λ, since encounters are more probable in high-yielding patches; and (ii) deplete the patch, decrementing the yield rate λ by ρ (figure 1*b*).

Varying ρ changes the rate of patch depletion relative to the time scale of the foraging process. Small relative values of ρ/λ_0 represent a large resource patch that the forager depletes very slowly. The limiting case $\rho/\lambda_0 \to 0$ represents the habitat choice problem. Alternatively, when this ratio is intermediate up to unity ($\rho/\lambda_0 \in [10^{-2}, 1]$), the forager considerably depletes the patch with each encounter. We refer to this as the patch leaving problem, and show that, in such cases, uncertainty in the patch yield can play a major role in shaping the departure strategy. We first consider the habitat choice problem in §3; following this, we consider the patch leaving problem in §4.

3. 'Habitat choice': minimizing time to find high-resource habitats

Habitat choice refers to patch use at scales where the forager's activity does not significantly affect the resource distribution or, in other words, that resource depletion occurs very slowly relative to the time needed for the search process. We represent this with the mathematically tractable yet representative limit of zero patch depletion. In this case, the optimal behaviour is to quickly locate a patch with the highest yield and remain there. Although in real environments habitats eventually deplete and the forager would leave, our theoretical treatment of a 'remain in the high-yielding patch' strategy can simply translate to a 'stay a long time in the initially high-yielding patch' strategy, with results applying similarly to both because of the separation of time scales: for habitat choice, the time needed to search and decide on a high-yield patch to remain in (which we denote as T_{arrive}) is much less than the time that would be needed to actually deplete the patch.

Upon entering a patch, the forager must use its experience of resource encounters to decide whether to stay in the patch or leave for another. We first consider a simplified *binary* environment where there are only two patch types—high-yield versus low-yield—and that the forager knows these possible patch types and their return rates. Here, the optimal behaviour is to infer whether or not it is currently in a high-yield patch, and, if so, to stay, but otherwise to leave. Uncertainty and stochasticity of resource encounters means that the forager will visit some low-yielding patches until it learns the yield rate and departs, and may also visit and depart from high-yielding patches if the type is incorrectly inferred. We then consider more general cases, and show that the general trends and optimal strategies from the simpler binary case still apply; this includes environments with multiple patch types and environments with continuous distributions of patch types where the forager has a threshold for accepting a patch as sufficiently dense with resources. With this approach, we can explicitly derive statistics associated with patch departures and examine how the efficient identification of high-quality habitats depends on environmental parameters like patch discriminability (e.g. λ_H/λ_L) and high-yield patch prevalence (p_H).

3.1. Two patch types

An environment with two possible patch types—high yielding and low yielding—is a mathematically tractable case that gives insight into optimal decision strategies and their resulting behavioural observables. Here, the probability distribution of patch types (likelihood that the next-visited patch is of a certain type) is $p_0(\lambda) = p_H\delta(\lambda - \lambda_H) + p_L\delta(\lambda - \lambda_L)$: H denotes the higher yielding patch and L denotes the lower yielding patch. As stated, we assume that the forager knows the values λ_H and

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

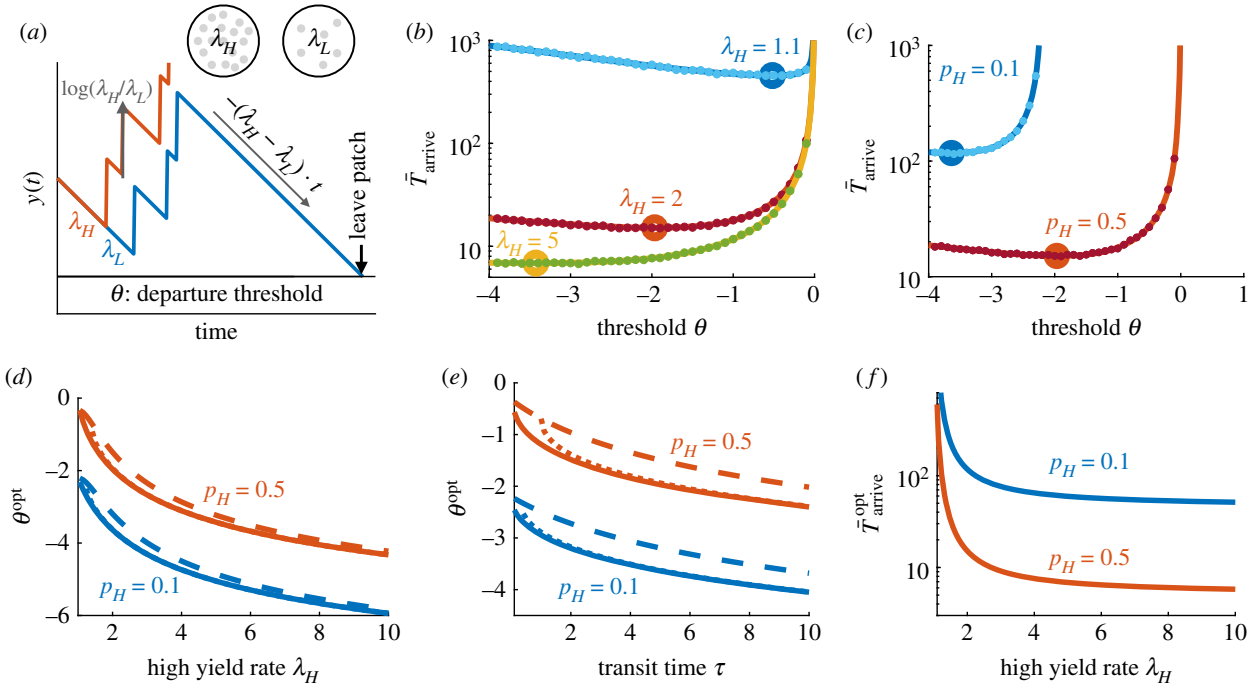

Figure 2. Statistics of habitat identification in environments with two patch types. (*a*) Habitat type belief $y(t) = \log(p(\lambda_H \mid x(t))/p(\lambda_L \mid x(t)))$ (as the animal decides between a habitat with a high, λ_H, and a low, λ_L, resource yield) increases with resource encounters and decreases between resource encounters until $y(t) = \theta$ and the observer departs the patch. (*b*,*c*) Mean time to arrive and remain in a high-yield habitat varies non-monotonically with departure threshold θ and decreases as the patch discriminability λ_H/λ_L and high-yield fraction p_H increase. Solid lines are equation (3.2). Dots are averages from 10^4 Monte Carlo simulations, each with independently drawn food encounter sequences. $\lambda_H = 2$ and $p_H = 0.5$ are fixed unless indicated. $\tau = 5$ and $\lambda_L = 1$ are fixed. (*d*,*e*) Departure threshold θ^{opt} minimizing the time to arrive in the high-yield patch decreases with λ_H/λ_L and τ. Solid lines are numerically obtained minima of equation (3.2), dotted lines are equation (3.3) and dashed lines are equation (3.4). In (*d*), dotted lines appear overlaid on solid lines because of the close fit. (*f*) The minimal mean time $\bar{T}^{\text{opt}}_{\text{arrive}}$ to arrive in a high-yield patch decreases with λ_H/λ_L and p_H.

λ_L and uses these as prior information to infer the type of the current patch. Using the limit of slow depletion ($\rho \to 0$) to represent habitat choice, the animal determines which patch type it is currently in using the log-likelihood ratio (LLR) $y(t) \equiv \log(p(\lambda_H \mid x(t))/p(\lambda_L \mid x(t)))$. With this, their belief update can be written as an SDE:

$$\frac{\mathrm{d}y}{\mathrm{d}t} = \log\frac{\lambda_H}{\lambda_L}\sum_{j=1}^{\infty}\delta(t-t_j) - (\lambda_H - \lambda_L), \tag{3.1}$$

with initial condition set by the prior $y(0) = \log((p_H)/(1-p_H))$. Resource encounters provide evidence for the high-yielding patch (first term) while elapsed time between resource encounters builds up evidence for the low-yielding patch (second term). Equation (3.1) has a simple form similar to classic evidence accumulation models of decision-making psychophysics [55,56], recently extended to foraging decisions [48].

The long-term resource intake rate is maximized if the forager finds and remains in a high-yielding patch (figure 2*a*). If the forager remains in a high-yielding patch, then the energy intake rate will reach λ_H in the limit of long time. Before locating and deciding to remain in a high-yielding patch, the forager may also visit low-yielding patches, leaving when its belief crosses the threshold (figure 2*a*), and may also visit and depart from high-yielding patches, if they are mistaken for low-yielding ones. The departure threshold sets the certainty that the forager obtains before leaving: a low threshold means high certainty of the patch type before leaving: while a high threshold will result in more departures. Too low a threshold can lead to too much time spent in low-yielding patches while gathering more evidence, while too high a threshold can lead to (incorrectly) departing from high-yielding patches before gathering enough evidence to distinguish their type. By setting the optimal threshold that balances uncertainty to minimize the time to arrive and remain in a high-yielding patch, we can ask how the environmental characteristics of relative patch yield, relative patch density and travel time influence behaviour and expected return of resources.

The forager's strategy is determined by the threshold θ on its belief (LLR), given by equation (3.1), shaping the time to find and remain in a high-yield patch $\bar{T}_{\text{arrive}}(\theta)$. This quantity can be computed from the patch-departure statistics of the forager, using first-passage time methods, given the prior $y(0) = \log(p_H/(1-p_H))$ [57,58]; τ is the mean travel time between patches, which we assume is known or determined from experience. Using this, the time to arrive and remain in a high-yield patch is

$$\bar{T}_{\text{arrive}}(\theta) = (1-p_H)\left[\frac{1+\mathrm{e}^{\theta}}{p_H-(1-p_H)\mathrm{e}^{\theta}}\right] \times \left[\frac{\log(p_H/(1-p_H)) - \theta}{(\lambda_H-\lambda_L) - \lambda_L\log(\lambda_H/\lambda_L)} + \tau\right]. \tag{3.2}$$

The minimum of this, corresponding to θ^{opt} (figure 2*b*,*c*), can be determined numerically (solid lines in figure 2*d*,*e*). An explicit approximation of θ^{opt} is obtained by differentiating equation (3.2), dropping higher order terms and solving for

$$\theta^{\text{opt}} \approx W_{-1}\left[-(1-p_H)\left(\frac{\lambda_H}{\lambda_L}\right)^{\lambda_L\tau}\mathrm{e}^{-(\lambda_H-\lambda_L)\tau}\mathrm{e}^{2p_H-1}\right] + A\tau + \log\frac{p_H}{1-p_H} + 1 - 2p_H, \tag{3.3}$$

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

where $A = (\lambda_H - \lambda_L) - \lambda_L \log \frac{\lambda_H}{\lambda_L}$ is defined for ease of notation ($A = 0$ for $\lambda_H = \lambda_L$, and $A > 0$ for $\lambda_H > \lambda_L$) and $W_{-1}(z)$ is the (-1)th branch of the Lambert W function (inverse of $z = We^W$). This approximation matches well with the numerically obtained minima of equation (3.2) (dotted lines in figure 2*d*,*e*), and can be further simplified using the approximation $W_{-1}(z) \approx \log(-z) - \log(-\log(-z))$, yielding

$$\theta^{\text{opt}} \approx \log(1 - p_H) - \log[A\tau - \log(1 - p_H) - 2p_H + 1], \quad (3.4)$$

indicating the scalings of the optimal threshold in limits of environmental parameters (dashed lines in figure 2*d*,*e*). Details on the derivation of the optimal threshold can be found in [58].

How should an animal best adapt its habitat search strategy to the statistics of the environment? When high-yield patches are rare (low p_H), travel times are large (high τ) or patches are easily discriminable (high λ_H relative to λ_L), the forager should gain higher certainty by deliberating longer before departing a patch; indeed, from equation (3.4) and figure 2*d*,*e* we see that the optimal threshold decreases with ρ_H, τ and λ_H. Increasing discriminability (λ_H/λ_L) or the high-yield patch fraction p_H decreases the minimal mean time $T^{\text{opt}}_{\text{arrive}}$ needed to arrive and remain in a high-yield patch (figure 2*f*), since this makes finding a high-yielding patch easier for the animal.

Furthermore, the outcome $\bar{T}_{\text{arrive}}(\theta)$ is most sensitive to the strategy (choice of threshold θ) in environments with low discriminability and a small fraction of desirable patches (figure 2*b*,*c*). In experiments, the value of T_{arrive} is an observable that can be used to infer the effective value of θ that an animal is using. The parameter sensitivity suggests that an animal's patch-selection strategy—i.e. the value of θ it is using—could be more precisely inferred when high-yield patches are rare or more difficult to identify. Note that the patch leaving rule of thresholding one's LLR is mathematically equivalent to thresholding the mean estimated resource yield rate since $\bar{\lambda} = (\lambda_H + e^{-y}\lambda_L)/(1 + e^{-y})$, analogous to previous patch-departure rules developed [31,41]. Next, we generalize this approach to environments with more than two patch types, so decisions use multiple LLRs, such that optimal decisions do not simply map to thresholding the estimated yield rate.

3.2. Multiple patch types

Animals may have to select from any number of patch types in an environment, which begs the question as to how decision and search strategies should extend to more general environments. With multiple patch types, decisions made by computing only two LLRs is sufficient to obtain near optimal performance in terms of minimizing the time to find and remain in a high-return patch. This result thus complements and extends previous work that has considered optimal strategies for two patch types [33,34] or more general models that are intractable to a mathematical study of how behaviour and yield vary with environmental and strategy parameters [47,59].

To model multiple patch types, consider environments with N patch types having resource yield rates $\lambda_1 > \lambda_2 > \cdots > \lambda_N \geq 0$ with patch fractions $p_1, p_2, \ldots, p_N$. Defining LLRs $y_j = \log(p(\lambda_1|x(t))/p(\lambda_{j+1}|x(t)))$ for $j = 1, \ldots, N-1$, yields the $N-1$-dimensional system fully describing an ideal observer's belief about the current patch type

$$\dot{y_j} = \log\frac{\lambda_1}{\lambda_{j+1}} \sum_{j=1}^{\infty} \delta(t - t_j) - (\lambda_1 - \lambda_{j+1}), \quad (3.5)$$

where $y_j(0) = \log(p_1/p_{j+1})$, and any likelihood can be recovered as $p(\lambda_{j+1}|x(t)) = e^{-y_j}/(1 + \sum_{k=1}^{N-1} e^{-y_k})$, $j = 0, 1, \ldots, N-1$, where $y_0 = 0$ for $j = 0$.

As in the binary case, the optimal strategy is to find and remain in the highest yielding patch (λ_1). We again represent patch leaving decisions by thresholding the probability of being in the high-yielding patch, such that when $p(\lambda_1|x(t)) = \phi \in (0, p_1)$ the forager exits the patch. We approximate this thresholding process by requiring $y_j \geq \theta$ (for $j = 1, 2, \ldots, N$) to remain in the patch, so the forager departs given sufficient evidence that it is not in the highest yielding patch (see figure 3*a* for three patch types).

We also consider how effective reduced strategies are, for which the observer only tracks the first L LLRs $y_1, y_2, \ldots, y_L$ and compares these with the threshold θ to decide when to leave the patch. Thus, we compute the mean time to arrive and remain in the high-yield patch, which depends on the escape probability $\pi_1(\theta)$ from the high-yielding patch and the mean time to visit each patch $\bar{T}_j(\theta, L)$ when escaping,

$$\bar{T}_{\text{arrive}}(\theta) = \frac{\pi_1(\theta, L)}{1 - \pi_1(\theta, L)}(\bar{T}_1(\theta, L) + \tau) + \frac{1 - p_1}{p_1}\frac{\sum_{j=2}^{N} p_j(\bar{T}_j(\theta, L) + \tau)}{1 - \pi_1(\theta)}, \quad (3.6)$$

where the patch-departure strategy depends on the number L of LLRs thresholded and the threshold θ used.

The mean high-yield patch arrival time $\bar{T}_{\text{arrive}}$ depends strongly on the high-yield patch resource yield rate λ_1, which decreases considerably as the patch becomes more discriminable (three patches: figure 3*b*; five patches: figure 3*d*). On the other hand, $\bar{T}_{\text{arrive}}$ depends weakly on the worst patch's yield rate λ_3 (figure 4*a*), so uncertainty among the less valuable patches has little effect on behaviour. In a related way, $\bar{T}_{\text{arrive}}$ is much more strongly affected by changes in the fraction of the high-yielding patch (p_1: figure 3*c*) than by changes in the balance of the mid- (λ_2) and low-yielding (λ_3) patches (figure 4*b*).

Again, the optimal threshold decreases when patches are more discriminable: as λ_1 increases the forager should gain a higher certainty before leaving (figure 3*b*). The average high-yield patch arrival time $\bar{T}_{\text{arrive}}$ depends weakly on the threshold near the optimum, but the optimum threshold shows a non-monotonic dependence on p_1. The lower optimum threshold for both low- and high-yield p_1 values represents that, in these cases, it is optimal to be more certain of the future harvest rate before leaving: for low p_1 this occurs because high-yield patches are rare (and thus there is a higher premium on distinguishing the high-yielding patch when actually in one), and for high p_1 this occurs because they are plentiful (one is more likely to land in a high-yielding patch, so one can afford to require more certainty to depart). Such an increased premium placed on information gathering for the reduction of uncertainty about the future yield of a patch in sparse environments has been identified in previous analyses of foraging models governed by statistical decision theory [59,60]. Between these cases, although the optimal threshold is slightly higher, the dependence is weak. Additionally, this demonstrates that, if the forager did not know p_1 (we assumed that this is known and is used to formulate the leaving decisions), the best strategy would be to err on the side of choosing a low threshold, because the sensitivity of $\bar{T}_{\text{arrive}}$ to threshold is relatively weak for choices too low but can be higher for choices too high (figure 3*c*).

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

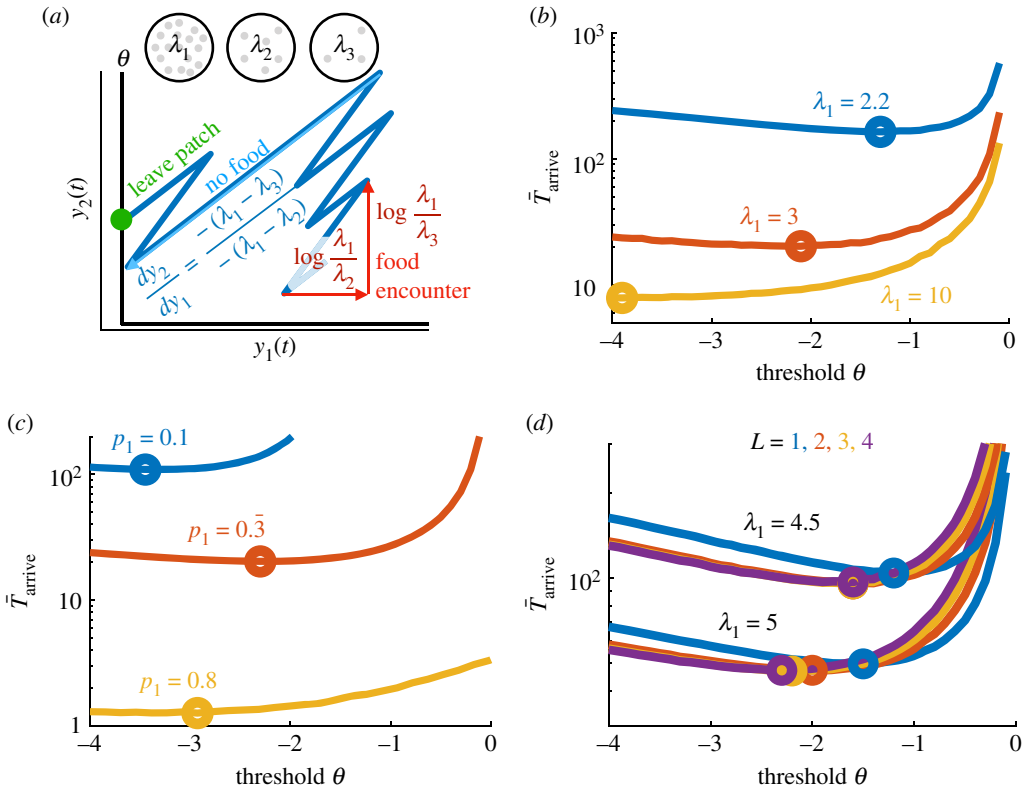

Figure 3. Optimal departure strategies for habitat choice in environments with multiple patch types. (*a*) Beliefs about three possible habitat types $y_1(t) = \log\left(p(\lambda_1|x(t))/p(\lambda_2|x(t))\right)$ and $y_2(t) = \log\left(p(\lambda_1|x(t))/p(\lambda_3|x(t))\right)$ increase with resource encounters and decrease between resource encounters until either reaches the departure threshold θ. (*b*) The mean time $\bar{T}_{\text{arrive}}$ to arrive and remain in the highest yielding patch λ_1 (equation (3.6)) (with $N = L = 3$) decreases with patch discriminability λ_1, as does the optimal departure threshold θ^{opt} (circles). Three patch types with $\lambda_2 = 2$, $\lambda_3 = 1$, $\tau = 5$, and $p_1 = p_2 = p_3 = 1/3$. (*c*) $\bar{T}_{\text{arrive}}$ decreases with the prevalence of the best patch p_1 (while $p_2 = p_3 = (1 - p_1)/2$), but θ^{opt} varies non-monotonically. (*d*) The mean time $\bar{T}_{\text{arrive}}$ to arrive and remain in the highest yielding patch λ_1 (see equation (3.6)) depends on how many LLRs (L) the forager uses to make a decision. Although the optimal time $\bar{T}^{\text{opt}}_{\text{arrive}}$ (curve minima – circles) decreases with L, the dependence is weak; $L = 2$ yields nearly identical mean optimal arrival times as $L = 4$. The optimal threshold which leads to $\bar{T}^{\text{opt}}_{\text{arrive}}$ decreases with L. Other parameters are $p_j = 1/N$, $\lambda_j = 6 - j$, $j = 1, 2, 3, 4$, $N = 5$. In (*c*) and (*d*), 10^6 Monte Carlo simulations are used to compute the curves $\bar{T}_{\text{arrive}}$.

In an environment with five patches, performance depends weakly on how many LLRs are used to make patch-leaving decisions for $L > 2$. It is sufficient to simply track the LLRs between the first three patches, which correspond to using $L = 2$ (figure 3*d*). This is because the forager only needs to know whether it is in one of the best patches or not, since the goal is to eventually settle in one such patch as a habitat. This demonstrates again that the key features of uncertainty that matter to the optimal forager are the discriminability and prevalence of the best and second best patch type.

3.3. Continuum limit: many patch types

Building on the N-patch case, we now consider a scenario where there is a continuous distribution of patch qualities ($N \to \infty$), so the resource yield rate for each patch λ is drawn from a continuous distribution $p_0(\lambda)$, which serves as a prior for the posterior $p(\lambda|x(t))$ with each patch visit (figure 5*a*). For any continuous probability distribution function $p_0(\lambda)$, the maximum λ will never be sampled, so arriving and remaining in the 'maximum' yielding patch is not possible. We therefore assume that the forager seeks patches with yield rates λ_θ or above, but deems lower yield rates to be insufficient. With this formulation, the forager updates an LLR based on a belief of whether the current patch is greater or less than λ_θ. Because this divides the continuous distribution into two categories, the mathematical treatment is similar to the binary case, but with added uncertainty because the patches in each category do not have the same return.

To model this, given a reference yield rate λ_θ, we represent decisions in an environment with a continuous distribution of patch qualities by tracking $P(\lambda > \lambda_\theta|x(t)) = \int_{\lambda_\theta}^{\infty} p(\lambda|x(t))\,\mathrm{d}\lambda$. For the case of an exponential prior $p_0(\lambda) = \alpha e^{-\alpha\lambda}$, given $K(t)$ resource encounters, we define $\rho(t) = \log\left(P(\lambda > \lambda_\theta|x(t))/P(\lambda < \lambda_\theta|x(t))\right)$ and state that the forager departs the patch when $\rho(t) \leq \hat{\theta}$ or when $P(\lambda > \lambda_\theta|x(t)) \leq \theta := 1/(1 + e^{-\hat{\theta}})$. Note that, to allow evidence accumulation, we require that $\theta < \alpha \int_{\lambda_\theta}^{\infty} e^{-\alpha\lambda}\,\mathrm{d}\lambda = e^{-\alpha\lambda_\theta} \equiv \phi$, which represents the fraction of patches where $\lambda \geq \lambda_\theta$.

Computing the probability of escaping a high-yield patch, $\pi_H(\theta; \lambda_\theta) = \int_{\lambda_\theta}^{\infty} p_0(\lambda)\pi(\theta; \lambda)\,\mathrm{d}\lambda$, and the mean time per visit to high- and low-yield patch types, $\bar{T}_H(\theta; \lambda_\theta) = \int_{\lambda_\theta}^{\infty} p_0(\lambda)\bar{T}(\theta; \lambda)\,\mathrm{d}\lambda$ and $\bar{T}_L(\theta; \lambda_\theta) = \int_0^{\lambda_\theta} p_0(\lambda)\bar{T}(\theta; \lambda)d\lambda$, when departing (using Monte Carlo sampling), we then use equation (3.2) to compute the time to arrive in a high-yielding patch (figure 5*b*). Placing a higher threshold λ_θ on the quality of an acceptably high-yielding patch increases the time to arrive in an acceptably high-yield patch. Moreover, the optimal threshold θ decreases, as more time must be spent in patches to discriminate a high-yielding patch, which become rarer as λ_θ increases. Increasing λ_θ corresponds to making sufficiently high-yield patches more discriminable and more rare.

With this formulation, the mathematical treatment in the case of a continuous distribution of patches is then the same as the binary case, and we can map corresponding

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

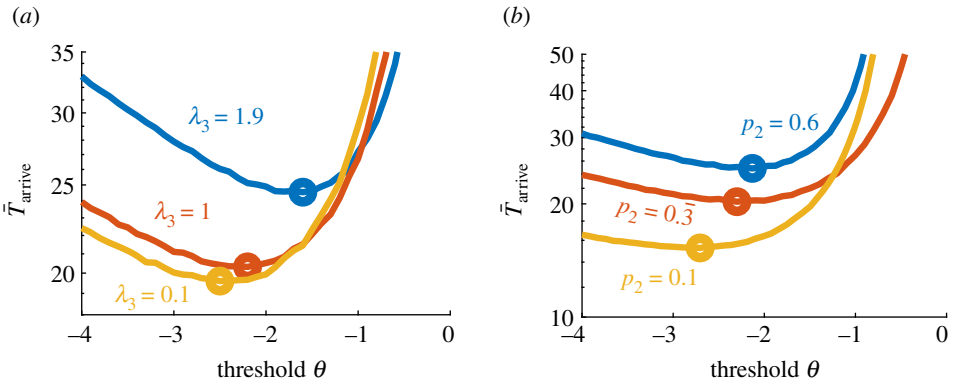

Figure 4. Insensitivity of patch-finding performance to worst patch statistics. (*a*) $\bar{T}_{\text{arrive}}$ increases with λ_3, as does θ^{opt}, since the worst patches become less easy to distinguish from the best ($\lambda_1 = 3$). Here $\lambda_2 = 2$, $\tau = 5$, $p_1 = p_2 = p_3 = 1/3$. (*b*) $\bar{T}^{\text{opt}}_{\text{arrive}}$ increases with p_2 (while $p_1 = 0.333$ and $p_3 = 1 - p_1 - p_2$), as does θ^{opt}. All curves for $\bar{T}_{\text{arrive}}$ computed from 10^6 Monte Carlo simulations.

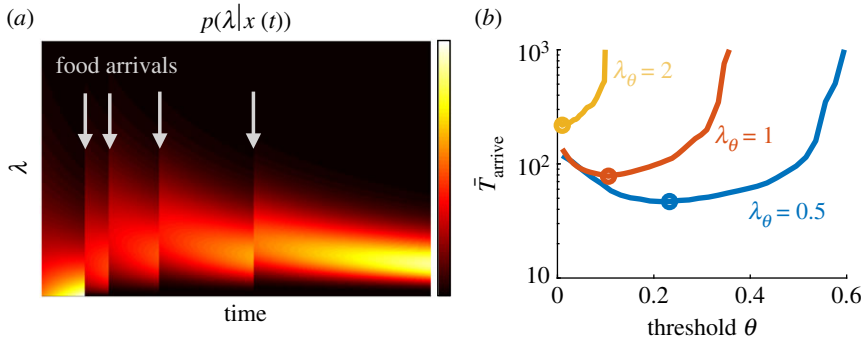

Figure 5. Habitat selection given a continuum of patch types. (*a*) The posterior $p(\lambda|x(t))$ of the estimated yield rate is shifted up by each resource encounter and decreases between resource encounters. (*b*) There is an optimal θ that minimizes the time to arrive and remain in a high-yielding patch ($\lambda > \lambda_\theta$), and the optimal time $\bar{T}^{\text{opt}}_{\text{arrive}}$ (circles) increases as the acceptable yield rate is increased. $\alpha = 1$ is used here. 10^6 Monte Carlo simulations for each λ_θ and θ.

results: setting a higher λ_θ is equivalent to decreasing p_H and concurrently increasing λ_H. Although we considered an exponential distribution for $p_0(\lambda)$, we note that if this distribution changes, then this will affect the relationship between λ_θ and the equivalent mapping onto the binary case (in terms of p_H and λ_H). Another possibility would be to further 'bin' the continuous distribution to correspond to three effective types, instead of two, as we did using a single threshold. The continuous case then would be treated analogously to a three-patch-type environment, and could be represented with two LLRs. However, further binning may not be necessary to achieve near-optimal decisions. Overall, this demonstrates that effective strategies for foraging environments with a continuum of patch types could be generated using particle filters that compute likelihoods over a finite set of patch types [61].

3.4. Summary of results: habitat choice problem

In general, we see that strategies that divide the environment into two patch types work well in efficiently finding the best or near-best patches, even in the presence of many patch types. The optimal time to arrive and remain in the highest yielding patch decreases as the high-yield patch discriminability increases and as high-yield patches become more common. Considering more than two patch types, the associated foraging strategies are most strongly coupled to environmental parameters of the highest and second highest yielding patch types. It is not necessary to compute LLRs associated with all possible types in order to efficiently find a high-yield patch—even considering only a single LLR gives reasonable results, and the average time to arrive in a high-yield patch is not strongly affected when the number of LLRs continues to increase beyond two. This suggests that animals select habitats by estimating a possible range of high-quality patches and then making patch-departure decisions based on whether patches meet those criteria or not.

4. Patch leaving: depletion- versus uncertainty-driven decisions

When the scale of a patch is smaller, the forager will significantly deplete the patch's resources during its visit. The decision is then not of which patch to remain in, but rather of when to leave the current patch in search of another. We therefore refer to this as patch exploitation (figure 1*a*; [12,62]). While the nature of the resource differs for different animals (and typical patch residence times can accordingly vary from seconds up to hours—for some examples, see [30,36,45,63,64]). These cases all have in common that each resource patch is small enough that availability within the patch is affected by the consumption of the forager.

The MVT sets the optimal time to leave a patch in order to maximize resource consumption over time: when the current patch yield rate equals the overall average yield rate for the environment [19]. However, the MVT is simply an optimal rule, and does not specify the mechanistic process of how an animal uses its experience to reach a decision to leave a

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

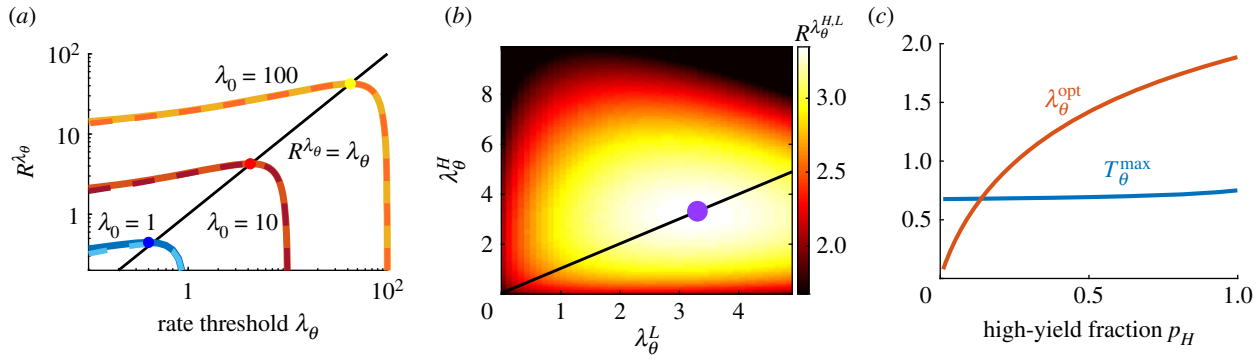

Figure 6. Departure strategies for the patch exploitation problem. (*a*) Rate of resource consumption as a function of strategy R^{λ_θ}, where the forager departs once the inferred yield rate falls to or below λ_θ. Solid lines are exact solution of equation (4.2) and dashed lines are the large m_0 approximation. Note optimal $\lambda_\theta \approx R^{\lambda_\theta}$. (*b*) Rate of resource consumption $R^{\lambda_\theta^{H,L}}$ in binary environments in which the observer knows the patch type (λ_H or λ_L) upon arrival. Optimal strategy (purple dot) takes $\lambda_\theta^H = \lambda_\theta^L \approx R^{\lambda_\theta^{H,L}} \approx 3.3$. $p_H = 0.5$, $m_0^H = 100$, $m_0^L = 50$ and $\rho = 0.1$. (*c*) Optimal waiting time $T_\theta^{\max}$ and departure threshold $\lambda_\theta^{\max}$ in a binary environment where the low-yield patch has zero chunks of resource and the high-yield patch initially has $m_0^H = 50$. $\rho = 0.1$. Transit time $\tau = 5$ throughout.

patch. Previous work has demonstrated that rewards in discrete chunks—instead of as a continuous rate—can affect the process an animal uses in decision-making [21,24,48]. From a Bayesian perspective, decisions should use available information about the resource distribution in the environment. If resource availability within a patch is discrete or uncertain, even an ideal observer may not be able to accurately infer the actual rate of return, and thus would not be able to implement the leaving rule prescribed by the MVT. Experiments show that, while the general trends predicted by the MVT hold in many cases, animals often deviate from an MVT-predicted strategy [50]. Moreover, in cases where patches contain very few items (e.g. 0 or 1 resource chunks), reward is not described by a rate function, and the MVT leaving rule does not apply.

Here we consider an animal that encounters resources in discrete chunks and infers the state of the environment and subsequently acts. This allows us to ask when the MVT rule is actually optimal versus when it does not apply, when deviations from the MVT occur owing to uncertainty and how a forager can incorporate prior knowledge about the resource distribution in the environment to reach a patch leaving decision. We first treat the simple case of homogeneous patch types to establish the basic theoretical approach. Then, we consider an environment with two patch types to show how the inference procedure affects decisions in different environmental configurations, which we refer to as the 'depletion-dominated' versus 'uncertainty-dominated' regimes.

4.1. Homogeneous environments

To show how discreteness of resources affects decisions [21,24], we first consider the simple case of a homogeneous environment with a single patch type. An ideal forager with prior knowledge of the initial yield rate λ_0 can track time and resource encounters to determine the current yield rate $\lambda(t)$, and then depart the patch when the inferred value of $\lambda(t)$ falls below some threshold λ_θ. Prior knowledge of the initial patch yield can be used in order to infer

$$\lambda(t) = \lambda_0 - K(t)\rho, \tag{4.1}$$

which represents the true underlying value of $\lambda(t)$. Using this in a patch leaving decision strategy is equivalent to departing after a fixed number of resource encounters [48].

Using this inference strategy, we calculate the long-term resource intake rate by assuming that λ_θ is an integer multiple of ρ. With this, the number of chunks consumed before departure is $m_\theta \equiv K(T(\lambda_\theta)) := (\lambda_0 - \lambda_\theta)/\rho$. Linearity of expectations allows us to compute the mean departure time as the sum of mean exponential waiting times between resource encounters $T_{\lambda_\theta} = [H_{m_0} - H_{m_0 - m_\theta}]/\rho$, where H_n is the nth harmonic number. Thus, we can approximate the long-term resource intake rate given λ_θ as

$$R^{\lambda_\theta} \approx \frac{m_0\rho - \lambda_\theta}{\log(\rho m_0) - \log\lambda_\theta + \rho\tau}, \tag{4.2}$$

which is valid for $m_0 \gg 1$. There is an interior optimum m_θ that maximizes the long-term resource consumption rate, which we can estimate by computing the approximate critical point equation of equation (4.2)

$$\frac{m_0\rho - \lambda_\theta}{\lambda_\theta} = \log\frac{\rho m_0}{\lambda_\theta} + \rho\tau. \tag{4.3}$$

For large m_0 (many chunks per patch), $\lambda_\theta = R^{\lambda_\theta}$, i.e. the leaving threshold is equal to the overall average rate of return in the environment; this is the optimum prescribed in the MVT [19]. Using the exact formula in equation (4.2), we can numerically determine the optimal threshold for small m_0 (a few chunks per patch). This shows that, when there are only a few chunks per patch, the true optimal threshold is close, but not exactly equal, to R^{λ_θ} (figure 6*a*). Note that the current estimate of the encounter rate is monotonically related to the future expected harvest rate, so a patch-departure rule based on either can easily be mapped to the other.

4.2. Binary environments

In binary environments, the forager estimates the underlying yield rate of the current patch, $\hat{\lambda}(t)$, and departs when this falls below a threshold λ_θ. Constant threshold strategies are our focus owing to their relative simplicity, but we note that alternatively dynamic programming could be used to determine optima of a more general class of departure strategies [65]. We assume that the forager uses knowledge that two different patch types exist to estimate the yield rate of the current patch; this involves using prior information to discriminate the patch type (high or low), combined with resource encounters which decrement the estimated

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

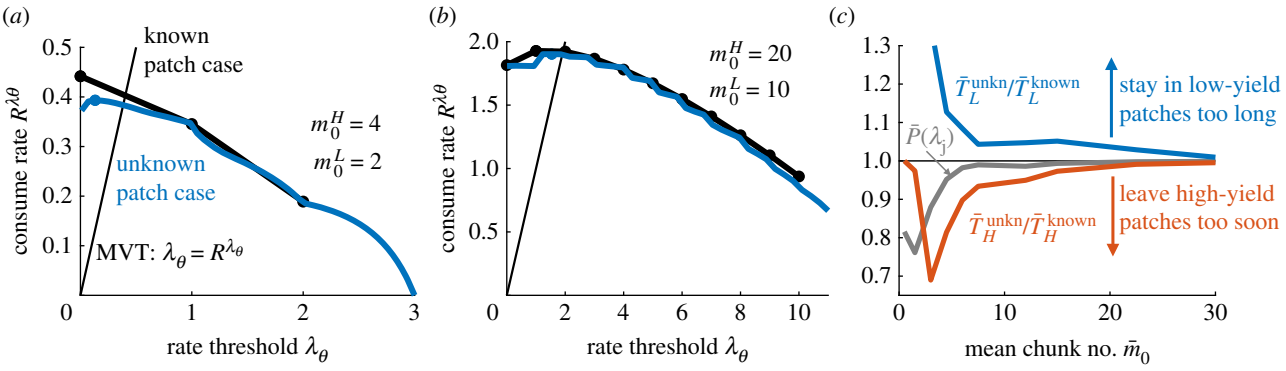

Figure 7. Depleting environments with 'known' versus 'unknown' patch resource values. In the 'known' case, the forager knows the patch type and initial resource yield rate upon arrival, whereas in the 'unknown' case the forager must infer the path type. Patch-leaving strategies in the unknown case approach those of the known case in the high-chunk count limit. (*a*) Total resource intake rate R^{λ_θ} as a function of the estimated yield rate λ_θ at which the observer departs a patch. For low initial resource levels m_0^H and m_0^L, R^{λ_θ} near the optimum (dots) deviates between the case of initially known patch type (black) versus initially unknown (blue). Both thresholds are less than the MVT prediction owing to the discreteness of resource encounters. (*b*) Resource intake rates of the known/unknown cases converge at higher initial resource levels. (*c*) Deviations in optimal resource intake rate between known and unknown cases are due to less certainty $\bar{P}(\lambda_j)$ (grey line) about the patch yield rate upon departure, leading the observer to stay in low-yield patches too long (blue line) and leave high-yield patches too soon (red line). Other parameters are $\rho = 1$, $\tau = 5$ and $p_H = 0.5$. Mean rates and departure times in the unknown patch case were computed from 10^6 Monte Carlo simulations.

yield rate. The belief can be determined according to a non-autonomous SDE for an LLR,

$$\frac{\mathrm{d}y}{\mathrm{d}t} = \sum_{j=1}^{K_{\max}} \log \frac{\lambda_H - (j-1)\rho}{\lambda_L - (j-1)\rho} \cdot \delta(t - t_j) - (\lambda_H - \lambda_L), \tag{4.4}$$

where $y(0) = \log(p_H/(1-p_H))$ as in the case of habitat choice.

We focus on two different scenarios, which we refer to as: (i) depletion-dominated regime, where the initial yield rate of the patch is known, and therefore leaving decisions are based solely on depletion, and (ii) uncertainty-dominated regime, where the type of patch is not known upon entry, and optimal leaving decisions must consider uncertainty in the estimate of the current yield rate of the patch. For the uncertainty-dominated regime, we consider first the cases where low-return patches contain zero resources, and then generalize to different amounts of resources per patch type.

Depletion-dominated regime. To represent what we term the depletion-dominated regime, we assume that the forager arrives in a patch and immediately knows the patch type λ_j ($j \in \{H, L\}$) in which it resides (e.g. owing to information provided by conspecifics or visual cues). In this case, the forager can make an accurate estimate of the true underlying yield rate of the patch, and therefore there is no uncertainty. Thus, leaving decisions are driven by depletion of the patch, as determined by when the estimated yield rate falls below some level λ_θ^j.

Following our calculations from the homogeneous case, the long-term resource intake rate depends on the initial resource chunk count in patches of type j, m_0^j, and the departure thresholds, so in the large m_0^j limit

$$R^{\lambda_\theta^{H,L}} \approx \frac{p_H(m_0^H\rho - \lambda_\theta^H) + p_L(m_0^L\rho - \lambda_\theta^L)}{p_H\tilde{T}_H(\lambda_\theta^H) + p_L\tilde{T}_L(\lambda_\theta^L) + \rho\tau}, \tag{4.5}$$

where $\tilde{T}_j(\lambda_\theta^j) = \log(\rho m_0^j) - \log\lambda_\theta^j$ and the critical point equations for each partial derivative $\partial_{\lambda_\theta^j} R^{\lambda_\theta^{H,L}} = 0$ imply

$$p_H(m_0^H - \lambda_\theta^H) + p_L(m_0^L\rho - \lambda_\theta^L) = \lambda_\theta^j\left[p_H(\log(\rho m_0^H) - \log\lambda_\theta^H) + p_L(\log(\rho m_0^L) - \log\lambda_\theta^L) + \rho\tau\right]. \tag{4.6}$$

This can be rewritten as $R^{\lambda_\theta^{H,L}} = \lambda_\theta^j$ ($j = H, L$). The aforementioned equation shows that, like the homogeneous case, an optimal strategy when there are many chunks per patch is to depart as the inferred yield rate equals the mean rate of resource encounters for the environment; additionally, the optimal threshold only depends on the average yield rate for the environment, and not the individual patch types (MVT; figure 6*b*). When there are a few chunks per patch, the optimal threshold may slightly differ from this value (see results for the homogeneous case in figure 6*a*). The depletion-dominated regime is similar to a homogeneous environment: since the forager knows the initial yield rate of the patch it is currently in, it can accurately infer the true underlying yield rate, and depart based on depletion of the patch. As in the homogeneous case, the optimal decision strategy can be formulated equivalently as either leaving when the estimated rate of return falls below a threshold or as counting—leaving after consuming a certain amount of resources.

Uncertainty-dominated regime—empty low-yield patch. In the 'uncertainty-dominated' regime, the forager does not know the initial yield rate of the patch upon entry. However, we assume that it has prior knowledge of the types of patches in the environment, i.e. that it knows the values of λ_H and λ_L. We first consider the tractable scenario where the low-yielding patch is empty ($\lambda_L = 0$). Such situations occur if certain regions of the environment appear to have food (e.g. fruiting vegetation) but on closer inspection turn out to be empty (e.g. already foraged or rotten). The optimal strategy is for the observer to first wait a finite time T_θ to depart if no resources are encountered, but if resources are encountered before $t = T_\theta$ to consume those resources until the inferred yield rate drops to $\lambda_\theta = (m_0 - m_\theta)\rho$. Early/late decisions are thus driven by uncertainty/depletion. Assuming $m_0 \gg 1$, we can continuously approximate the long-term resource intake rate and find that it is maximized using a waiting time T_θ that is insensitive to p_H. However, the threshold λ_θ depends on both p_H and travel time, because these parameters affect the overall average rate of resources available in the environment (figure 6*c*). The optimal threshold $\lambda_\theta^{\text{opt}}$ is not specified by the MVT, since uncertainty drives the forager to spend non-zero time in empty patches, adding extraneous time to the foraging process.

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

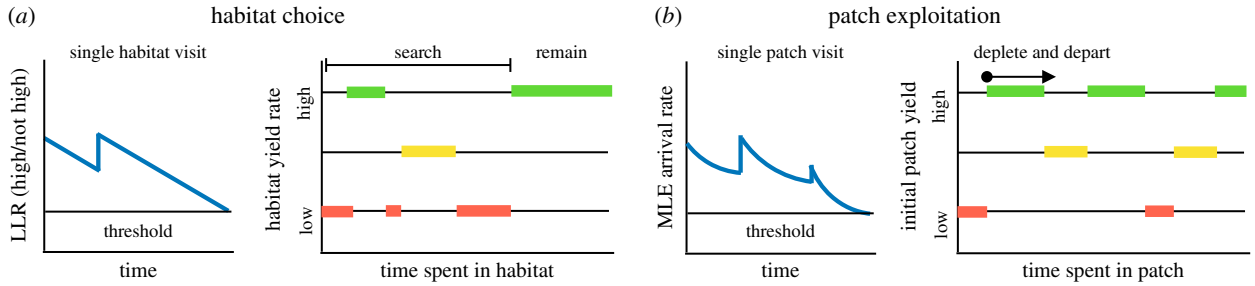

Figure 8. Summarized taxonomy of foraging strategies. See table 1 for details. In different environments with three patch types (low: red, medium: yellow, high: green yielding), the different time series of decision variables (for a single patch decision) and patch visit time intervals. (*a*) In habitat choice, an animal must determine whether its current patch is of the highest yielding type, departing if the probability that it is not in the highest reaches some threshold, and undergoing a sequence of patch visits until finding and remaining in a high-yielding patch. (*b*) An ideal forager performing patch exploitation infers the yield rate of its current patch and departs when the resource yield rate reaches a threshold, continuing patch visits indefinitely in large environments.

Many resource chunks. Next, we generalize to examine binary environments in which $m_0^H > m_0^L$ are arbitrary integers. In this case, the belief $y(t) = \log\,(P(\lambda_H - K(t)\rho|x(t))/P(\lambda_L - K(t)\rho|x(t)))$ evolves according to equation (4.4). The forager estimates the current yield rate of the patch from this belief,

$$\tilde{\lambda}(t) = \frac{\lambda_H + \mathrm{e}^{-y}\lambda_L}{1 + \mathrm{e}^{-y}} - \rho K(t), \tag{4.7}$$

and an optimal strategy is to depart when $\tilde{\lambda}(t) \leq \lambda_\theta$. The threshold λ_θ should be tuned to $\lambda_\theta^{\mathrm{opt}}$ so the long-term resource intake rate

$$R^{\lambda_\theta} = \frac{p_H \bar{m}_H + p_L \bar{m}_L}{p_H \bar{T}_H + p_L \bar{T}_L + \tau}$$

is maximized. We can compute departure times $\bar{T}_H$ and $\bar{T}_L$ numerically via Monte Carlo sampling. For an environment where the overall availability of resources is low and there are few resource chunks per patch, the optimal strategy when the patch type is known is to fully deplete each patch before leaving—this is represented by a threshold of $\lambda_\theta = 0$ for the inferred return rate. However, this is only optimal when the patch type is known; in the case of an unknown patch type, the forager has uncertainty in whether or not there are remaining resources in the patch, and this causes the optimum threshold to be non-zero. In both cases, the discreteness of resource encounters causes the optimal threshold to be lower than predicted by the MVT, although within this range of threshold values the average resource intake actually received is similar (figure 7*a*). In the case of high resource availability and many chunks per patch, the optimal thresholds are similar whether or not the forager knows the patch type upon arrival (figure 7*b*; compare black/blue curves) and coincide with the optimal threshold predicted by the MVT.

Comparing cases, we see that foragers in sparser environments (lower average initial resource amount $\bar{m}_0 = (\bar{m}_0^H + \bar{m}_0^L)/2$) stay in low-yield patches too long and leave high-yield patches too soon in comparison with observers that immediately know their patch type, owing to their uncertainty about their current patch before and at the time of departure (figure 7*c*). Uncertainty thus drives animals to underexploit (overexploit) high (low)-yielding patches when high- and low-yield patches are different enough that it is optimal to spend significantly more time in high- versus low-yield patches, but similar enough as to not be immediately distinguishable. Optimal leaving decisions in the uncertainty-dominated regime must use a rate-estimation process, because of the associated uncertainty in the true yield rate of the current patch.

5. Discussion

Patch foraging is a rich and flexible behaviour where an animal enters a patch of resources, harvests them and then leaves to search for another patch. An animal's behaviour can be quantified by its patch residence time distribution, travel time distribution, the amount of resources consumed and the movement pattern between patches. In this work, we used principles of probabilistic inference to establish a normative theory of patch-leaving decisions. With this general framework, we showed how foraging at different temporal and spatial scales is connected by a similar decision problem: 'habitat choice' refers to larger scales when foragers do not significantly deplete a resource, and 'patch exploitation' refers to smaller scales when the forager's activity depletes the patch. For habitat choice the optimal behaviour is to quickly locate and remain in a high-yielding habitat, while for patch exploitation it is optimal to use prior information along with reward encounters to estimate the current underlying yield rate to determine when to leave the patch (figure 8 and table 1).

In ecological contexts, these activities are part of a behavioural hierarchy, where an animal must decide *where to forage* and *how long to exploit a certain resource*.

In the case of habitat choice, the forager should use its experience of reward encounters to determine whether to stay or leave; in our model an optimal forager departs a habitat when its LLR for the probability of high-yield versus other habitats falls below a threshold. Optimal decisions are based on inference of habitat quality, with uncertainty being the driving factor in habitat departure times; while this is related to resource intake rates, it is not the same, because of how prior information can be used in patch inference. We showed that, with multiple different patch types, it is not necessary to track LLRs for all patch types—behaviour is most strongly affected by inference related to the highest and second-highest yield habitat types. The optimal time to arrive and remain in a high-yielding habitat is lower when patches are more discriminable, or when high-yield patches are more common. While T_{arrive} is an experimentally observable quantity, an animal's internal decision threshold is not; our model connects these quantities, and thus can be used to infer the decision rules an animal is using (figure 2). For example, a similar approach has been very informative to

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

Table 1. Detailed taxonomy of departure decision strategies. Departure strategies and observable trends depend on the environment and task: habitat selection or patch exploitation (see also figure 8). Columns describe the important aspects of the optimal decision strategy for each case, along with key model results.

environment	decision strategy and dependencies	equations	figures
habitat selection			
objective: minimize time to find highest yielding habitat			
known: resource yield rates of each habitat type			
N-habitat types	— depart habitat when likelihood of being in highest yielding habitat falls below a threshold — optimal strategy and arrival time depend on fraction and discriminability of high-yield habitats	$N = 2$: equation (3.1); $N \geq 3$: equation (3.5)	$N = 2$: figure 2; $N \geq 3$: figure 3
continuum of habitat types	— categorize habitats as high or low-yielding and depart habitat if likelihood of a high-yield falls below a threshold — time to identify high-yielding habitat is non-monotonic in departure threshold, and much longer when high-yield patches are rare		figure 5
patch exploitation			
objective: maximize mean resource intake rate R over a long time (several patches)			
known: initial yield rates of each patch type			
1-patch type	— depart when yield rate $\lambda(t)$ falls to a threshold value λ_θ — matches MVT except when there are very few chunks per patch, in which case the forager should empty the patch	equation (4.3)	figure 6*a*
2-patch types: patch type known on arrival	— depart when resource yield rate $\lambda_j(t)$ reaches a threshold — represents 'depletion-dominated' regime; recovers MVT	equations (4.5) and (4.6)	figure 6*b*
2-patch types: empty low-yield patch	— wait a time T_θ, then depart patch if no resources found; if resources are encountered by $t < T_\theta$, use threshold on inferred yield rate to make leaving decision (similar to single-patch-type case) — 'uncertainty-dominated' regime deviates from MVT — optimal wait time and departure threshold λ_θ increase with prevalence of high-yielding patch		figure 6*c*
2-patch types: both high- and low-yield patches have resources	— decision via threshold on current estimated yield rate $\tilde{\lambda}(t)$; choose optimal threshold λ_θ that maximizes long-term resource intake rate — optimal return differs from known case given few resources per patch, converges to known patch case when resource density is high — forager stays in low-yield patches too long, leaves high-yield patches too soon when there are few resources per patch (uncertainty-dominated regime)	equation (4.7)	figure 7

infer the parameters underlying two-choice decision tasks [66]. We showed that it is optimal to have a lower threshold—and thus gain a higher certainty before leaving—when travel times are large, high-yield patches are rare or high-yield patches are easier to discriminate. Analogous results have been found in observations of animals seeking habitats, where dispersal costs affect whether or not animals search for a new habitat [67]. These results give quantitative predictions that can be used to interpret experiments, for example to examine whether the animal is minimizing the time to reach the highest yield habitat in a heterogeneous environment. Moreover, we showed how behaviour in the general case where many habitat types exist can be understood by mapping results onto the tractable case of only two different patch types (figures 3 and 5). By varying systematically the percentage of high-yielding habitats and the discriminability (ratio of high- to low-yield rate), the model predicts how this affects the minimal time to arrive at the highest yield habitat, and connects this to a process that could implement such computations.

For patch exploitation, when the forager depletes patches in its habitat, in most cases the long-term intake rate is maximized by departing a patch when the in-patch estimated resource yield

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

rate matches the average return rate of the environment (i.e. by implementing the MVT rule). However, this does not apply when resources within a patch are limited so there is more uncertainty about the yield of the current patch upon departure (figure 7). Often in Nature, the environment is volatile and animals make foraging decisions while uncertain about local resource availability [68]. Our model predicts that, if there is high uncertainty about the patch type, this causes even an ideal Bayesian forager to stay longer in low-yield patches and shorter in high-yield patches than predicted by the MVT.

Our theoretical treatment of patch-leaving decisions builds on previous Bayesian models of foraging [31–39,41]. Our approach goes further than previous work by providing a step-by-step derivation of the normative strategies associated with a continuum of different environmental conditions, systematically identifying the dependence of observable behaviours (e.g. patch-departure times) on environmental parameters. During habitat choice, the minimal arrival time to the high-yield habitat scales with the probability of high-yield patches in the environment. On the other hand, we have shown that, in the case of depleting patches, the amount of time a forager overstays or understays in a patch scales with the density of resources in the patch. We are also able to infer the optimal threshold an animal should use. While this cannot be measured directly in experiments, our observations do reveal environmental parameter regimes under which performance (e.g. foraging yield) is sensitive to changes in strategy. This not only informs the design of behavioural foraging experiments, so as to determine task parameters that best reveal an animal's strategy, but optimal yields can also be compared with those obtained by animals in the wild to see how finely tuned their foraging strategies are.

Analysis of experiments shows that animals forage in ways that suggest they use Bayesian reasoning [40–46], using prior knowledge of their environment to modulate foraging behaviour [64,69,70]. For example, bumblebees [30] and Inca doves [71] adjust their foraging strategies in response to the predictability of the environment, as a Bayesian forager would, but this is not a universal trend [72]. Patch-leaving decisions may deviate from Bayes' optimality as animals become risk-averse in variable environments [73]. Other Bayesian models have considered patch-foraging decisions, even in rich multiple patch environments [34,47], but this work does not necessarily systematically vary environmental and strategy parameters to explore how the sensitivity of performance changes. Our work is also sufficiently mathematically tractable to suggest a mechanistic implementation that the forager can use to implement optimal decision rules. Moreover, our theoretical approach applies not only to patch exploitation, but also to habitat choice—where the MVT does not apply—and thus enables connections across these multiple scales of behaviour [13].

Although we used a constant threshold value based on either the belief or estimated yield rate, other work has examined cases where optimal decision strategies involve time-dependent decision thresholds [65,74–76]. Typically, these results arise in the context of multi-trial experiments in which the quality of evidence on each trial varies stochastically and is initially unknown. In the habitat-choice context, the quality of evidence is fixed across habitat visits, fitting the assumptions of classic, constant threshold optimal policies. We would therefore expect an analysis allowing for a dynamic threshold to yield the same results as we obtained here. On the other hand, when the animal performed patch exploitation in uncertain binary environments, we projected a higher dimensional description of the patch value to a single scalar estimate of the patch yield rate. In this case, a constant-threshold implementation may not be purely optimal. Leveraging methods from dynamic programming commonly used to set optimal decision policies [75] would be a fruitful next step in ensuring the optimality of our patch-leaving decision strategies. A common theme in previous Bayesian models that use dynamic programming [41,47] and our approach is that the forager should use the expected future return, not the current return, to make departure decisions. However, note that current and future return should be tightly monotonically related, especially in the limit of short times into the future. An advantage to the constant-threshold treatment is that it enables simple explicit quantitative relations that can be used to interpret experimental data (figure 2). For example, experiments evaluating how animals value the quality of evidence could help us delineate whether the animal is using a constant threshold or not.

Effective search is integral to survival in Nature [77], and search behaviour can give information on individual decision strategies. One can consider search within and between patches. Although we assumed random timings of resource encounters, an extension of our model could take into account different spatial arrangements. Within patches, an animal may perform random or systematic searches. For example, a recent study found that rats can solve the stochastic travelling salesman problem using a nearest neighbour algorithm [78]. A similar approach could ask how an animal's search and navigation pattern interacts with different patch-leaving decisions to create an effective foraging strategy that also considers memory of specific patch locations. Indeed, the explicit consideration of spatial movement may be necessary to understand foraging decisions. Previous work found that, when rats must physically move to perform foraging, the observed behaviour differed from tasks that 'simulate' foraging by presenting sequential choices or that consider a visual search [79]. It is an open question as to how the animal integrates aspects of spatial movement with economic valuations of future reward.

Our model assumes that animals know the initial yield rate of each type of patch in the environment. In a real-life context or in an experimental set-up, the animal would learn the environmental parameters, which we could model by considering another level in the inference hierarchy whereby the patch-type quality and fraction are learned along with the transit time distribution. Although we considered a single forager acting alone, another important extension will be to consider interactions between animals, either through predator–prey interactions which affect foraging decisions [80], social foraging of groups [52,81] or even competitive foraging [17,18,82]. In our model, the forager only receives direct (non-social) information about resource availability; in collective foraging, an individual receives both social and non-social information [83,84]. This can significantly affect foraging decisions, for example in the case where an individual must balance resource-seeking with group cohesion. Building on our modelling approach, foragers could share social information either by cooperating in the inference of the patch quality or by signalling to each other when to depart a patch as a threshold is reached.

To conclude, our model establishes a formal framework for the quantitative analysis of a natural behaviour—patch foraging (involving both habitat choice and patch exploitation)—that can be studied with the same formal rigor as

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

many trained behavioural tasks. Such validated *behavioural algorithms* are crucial for the systematic design of future experiments and interpretation of data on animal behaviour [85]. By comparing with theoretical optimal strategies, experiments and data can be used to understand the decision strategies an animal is employing and relate these to recorded animal movement and neural data. Future work will build on this model framework to generate testable hypotheses on the role of social interactions and the neural mechanistic underpinnings of foraging behaviour.

Data accessibility. No data were collected for or used in this paper.

Authors' contributions. Z.P.K., J.D.D. and A.E.H. conceived the project; Z.P.K. performed the model simulations and analysis and drew the figures; J.D.D. and A.E.H. performed a literature review of past foraging models and experiments; all authors wrote the manuscript.

Competing interests. We declare we have no competing interests.

Funding. Z.P.K. was supported by a CRCNS grant no. (R01MH115557) and NSF (DMS-1853630). A.E.H. was supported by NIH grant no. 1R21MH121889-01. J.D.D. was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2117—422037984 and the Heidelberg Academy of Science.

Acknowledgements. We thank Jan Drugowitsch for helpful discussions regarding the optimality of decision threshold policies across patch visits.

References

1. Calhoun AJ, Chalasani SH, Sharpee TO. 2014 Maximally informative foraging by Caenorhabditis elegans. *eLife* **3**, e04220. (doi:10.7554/eLife.04220)
2. Landayan D, Feldman DS, Wolf FW. 2018 Satiation state-dependent dopaminergic control of foraging in *Drosophila*. *Sci. Rep.* **8**, 1–9. (doi:10.1038/s41598-018-24217-1)
3. Marsh B, Schuck-Paim C, Kacelnik A. 2004 Energetic state during learning affects foraging choices in starlings. *Behav. Ecol.* **15**, 396–399. (doi:10.1093/beheco/arh034)
4. Carter EC, Redish AD. 2016 Rats value time differently on equivalent foraging and delay-discounting tasks. *J. Exp. Psychol.: General* **145**, 1093. (doi:10.1037/xge0000196)
5. Rose LM. 1994 Sex differences in diet and foraging behavior in white-faced capuchins (Cebus capucinus). *Int. J. Primatol.* **15**, 95–114. (doi:10.1007/BF02735236)
6. Hills TT. 2006 Animal foraging and the evolution of goal-directed cognition. *Cogn. Sci.* **30**, 3–41. (doi:10.1207/s15516709cog0000_50)
7. Raine NE, Ings TC, Dornhaus A, Saleh N, Chittka L. 2006 Adaptation, genetic drift, pleiotropy, and history in the evolution of bee foraging behavior. *Adv. Study Behav.* **36**, 305–354. (doi:10.1016/s0065-3454(06)36007-x)
8. Mobbs D, Trimmer PC, Blumstein DT, Dayan P. 2018 Foraging for foundations in decision neuroscience: insights from ethology. *Nat. Rev. Neurosci.* **19**, 419. (doi:10.1038/s41583-018-0010-7)
9. Hayden BY, Walton ME. 2014 Neuroscience of foraging. *Front. Neurosci.* **8**, 81. (doi:10.3389/fnins.2014.00081)
10. Hall-McMaster S, Luyckx F. 2019 Revisiting foraging approaches in neuroscience. *Cogn. Affect. Behav. Neurosci.* **19**, 225–230. (doi:10.3758/s13415-018-00682-z)
11. Calhoun AJ, Hayden BY. 2015 The foraging brain. *Curr. Opin. Behav. Sci.* **5**, 24–31. (doi:10.1016/j.cobeha.2015.07.003)
12. Stephens DW. 2008 Decision ecology: foraging and the ecology of animal decision making. *Cogn. Affect. Behav. Neurosci.* **8**, 475–484. (doi:10.3758/CABN.8.4.475)
13. McGarigal K, Wan HY, Zeller KA, Timm BC, Cushman SA. 2016 Multi-scale habitat selection modeling: a review and outlook. *Landscape Ecol.* **31**, 1161–1175. (doi:10.1007/s10980-016-0374-x)
14. Meyer CB, Thuiller W. 2006 Accuracy of resource selection functions across spatial scales. *Divers. Distrib.* **12**, 288–297. (doi:10.1111/j.1366-9516.2006.00241.x)
15. Kacelnik A, Vasconcelos M, Monteiro T, Aw J. 2011 Darwin's 'tug-of-war' vs. starlings' 'horse-racing': how adaptations for sequential encounters drive simultaneous choice. *Behav. Ecol. Sociobiol.* **65**, 547–558. (doi:10.1007/s00265-010-1101-2)
16. Boyce MS, Vernier PR, Nielsen SE, Schmiegelow FK. 2002 Evaluating resource selection functions. *Ecol. Modell.* **157**, 281–300. (doi:10.1016/S0304-3800(02)00200-4)
17. Kennedy M, Gray RD. 1993 Can ecological theory predict the distribution of foraging animals? A critical analysis of experiments on the ideal free distribution. *Oikos* **68**, 158–166. (doi:10.2307/3545322)
18. Abrahams MV. 1986 Patch choice under perceptual constraints: a cause for departures from an ideal free distribution. *Behav. Ecol. Sociobiol.* **19**, 409–415. (doi:10.1007/BF00300543)
19. Charnov EL. 1976 Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.* **9**, 129–136. (doi:10.1016/0040-5809(76)90040-X)
20. Waage JK. 1979 Foraging for patchily-distributed hosts by the parasitoid, *Nemeritis canescens*. *J. Anim. Ecol.* **48**, 353–371. (doi:10.2307/4166)
21. McNamara J. 1982 Optimal patch use in a stochastic environment. *Theor. Popul. Biol.* **21**, 269–288. (doi:10.1016/0040-5809(82)90018-1)
22. Mellgren RL. 1982 Foraging in a simulated natural environment: there's a rat loose in the lab. *J. Exp. Anal. Behav.* **38**, 93–100. (doi:10.1901/jeab.1982.38-93)
23. Cuthill I, Kacelnik A. 1990 Central place foraging: a reappraisal of the 'loading effect'. *Anim. Behav.* **40**, 1087–1101. (doi:10.1016/S0003-3472(05)80175-7)
24. Rita H, Ranta E. 1998 Stochastic patch exploitation model. *Proc. R. Soc. Lond. B* **265**, 309–315. (doi:10.1098/rspb.1998.0297)
25. Driessen G, Bernstein C. 1999 Patch departure mechanisms and optimal host exploitation in an insect parasitoid. *J. Anim. Ecol.* **68**, 445–459. (doi:10.1046/j.1365-2656.1999.00296.x)
26. Ward JF, Austin RM, Macdonald DW. 2000 A simulation model of foraging behaviour and the effect of predation risk. *J. Anim. Ecol.* **69**, 16–30. (doi:10.1046/j.1365-2656.2000.00371.x)
27. Eliassen S, Jørgensen C, Mangel M, Giske J. 2009 Quantifying the adaptive value of learning in foraging behavior. *Am. Nat.* **174**, 478–489. (doi:10.1086/605370)
28. Taneyhill DE. 2010 Patch departure behavior of bumble bees: rules and mechanisms. *Psyche* **2010**, 872736. (doi:10.1155/2010/872736)
29. Zhang F, Hui C. 2014 Recent experience-driven behaviour optimizes foraging. *Anim. Behav.* **88**, 13–19. (doi:10.1016/j.anbehav.2013.11.002)
30. Biernaskie JM, Walker SC, Gegear RJ. 2009 Bumblebees learn to forage like Bayesians. *Am. Nat.* **174**, 413–423. (doi:10.1086/603629)
31. McNamara JM, Green RF, Olsson O. 2006 Bayes' theorem and its applications in animal behaviour. *Oikos* **112**, 243–251. (doi:10.1111/j.0030-1299.2006.14228.x)
32. Olsson O, Holmgren NMA. 1998 The survival-rate-maximizing policy for Bayesian foragers: wait for good news. *Behav. Ecol.* **9**, 345–353. (doi:10.1093/beheco/9.4.345)
33. McNamara J, Houston A. 1980 The application of statistical decision theory to animal behaviour. *J. Theor. Biol.* **85**, 673–690. (doi:10.1016/0022-5193(80)90265-9)
34. Olsson O. 2006 Bayesian foraging with only two patch types. *Oikos* **112**, 285–297. (doi:10.1111/j.0030-1299.2006.13549.x)
35. Pierre J-S, Green RF. 2008 A Bayesian approach to optimal foraging in parasitoids. *Behavioral ecology of insect parasitoids: from theoretical approaches to field applications*, pp. 357–383. Oxford, UK: Blackwell.
36. Olsson O, Holmgren NMA. 1999 Gaining ecological information about Bayesian foragers through their behaviour. I. Models with predictions. *Oikos* **81**, 251–263. (doi:10.2307/3546740)
37. Van Gils JA. 2010 State-dependent Bayesian foraging on spatially autocorrelated food

Downloaded from https://royalsocietypublishing.org/ on 16 July 2021

distributions. *Oikos* **119**, 237–244. (doi:10.1111/j.1600-0706.2009.17497.x)

38. Rodriguez-Girones MA, Vasquez RA. 1997 Density-dependent patch exploitation and acquisition of environmental information. *Theor. Popul. Biol.* **52**, 32–42. (doi:10.1006/tpbi.1997.1317)
39. Ellison AM. 2004 Bayesian inference in ecology. *Ecol. Lett.* **7**, 509–520. (doi:10.1111/j.1461-0248.2004.00603.x)
40. Alonso JC, Alonso JA, Bautista LM, Muñoz-Pulido R. 1995 Patch use in cranes: a field test of optimal foraging predictions. *Anim. Behav.* **49**, 1367–1379. (doi:10.1006/anbe.1995.0167)
41. Green RF. 1980 Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging. *Theor. Popul. Biol.* **18**, 244–256. (doi:10.1016/0040-5809(80)90051-9)
42. Olsson O, Wiktander U, Holmgren NM, Nilsson SG. 1999 Gaining ecological information about Bayesian foragers through their behaviour. II. A field test with woodpeckers. *Oikos* **87**, 264–276. (doi:10.2307/3546741)
43. Valone TJ. 2006 Are animals capable of Bayesian updating? An empirical review. *Oikos* **112**, 252–259. (doi:10.1111/j.0030-1299.2006.13465.x)
44. Killeen PR, Palombo G-M, Gottlob LR, Beam J. 1996 Bayesian analysis of foraging by pigeons (Columba livia). *J. Exp. Psychol.: Animal Behav. Process.* **22**, 480. (doi:10.1037/0097-7403.22.4.480)
45. Nonacs P, Soriano JL. 1998 Patch sampling behaviour and future foraging expectations in Argentine ants, Linepithema humile. *Anim. Behav.* **55**, 519–527. (doi:10.1006/anbe.1997.0615)
46. Valone TJ, Brown JS. 1989 Measuring patch assessment abilities of desert granivores. *Ecology* **70**, 1800–1810. (doi:10.2307/1938113)
47. Green RF. 2006 A simpler, more general method of finding the optimal foraging strategy for Bayesian birds. *Oikos* **112**, 274–284. (doi:10.1111/j.0030-1299.2006.13462.x)
48. Davidson JD, El Hady A. 2019 Foraging as an evidence accumulation process. *PLoS Comput. Biol.* **15**, e1007060. (doi:10.1371/journal.pcbi.1007060)
49. Higginson AD, Fawcett TW, Houston AI, McNamara JM. 2018 Trust your gut: using physiological states as a source of information is almost as effective as optimal Bayesian learning. *Proc. R. Soc. B* **285**, 20172411. (doi:10.1098/rspb.2017.2411)
50. Nonacs P. 2001 State dependent behavior and the marginal value theorem. *Behav. Ecol.* **12**, 71–83. (doi:10.1093/oxfordjournals.beheco.a000381)
51. Oaten A. 1977 Optimal foraging in patches: a case for stochasticity. *Theor. Popul. Biol.* **12**, 263–285. (doi:10.1016/0040-5809(77)90046-6)
52. Clark CW, Mangel M. 1984 Foraging and flocking strategies: information in an uncertain environment. *Am. Nat.* **123**, 626–641. (doi:10.1086/284228)
53. Gabay AS, Apps MAJ. 2020 Foraging optimally in social neuroscience: computations and methodological considerations. *Soc. Cogn. Affect. Neurosci.* **30**, nsaa037. (doi:10.1093/scan/nsaa037)
54. Iwasa Y, Higashi M, Yamamura N. 1981 Prey distribution as a factor determining the choice of optimal foraging strategy. *Am. Nat.* **117**, 710–723. (doi:10.1086/283754)
55. Gold JI, Shadlen MN. 2007 The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574. (doi:10.1146/annurev.neuro.29.051605.113038)
56. Brunton BW, Botvinick MM, Brody CD. 2013 Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**, 95–98. (doi:10.1126/science.1233912)
57. Gardiner CW. 2004 *Handbook of stochastic methods*, vol. 3. New York, NY: Springer.
58. Kilpatrick ZP, Davidson JD, Hady AE. 2020 Normative theory of patch foraging decisions. (http://arxiv.org/abs/2004.10671).
59. Olsson O. 2006 The foraging benefits of information and the penalty of ignorance. *Oikos* **112**, 260–273. (doi:10.1111/j.0030-1299.2006.13548.x)
60. Dall SR, Giraldeau L-A, Olsson O, McNamara JM, Stephens DW. 2005 Information and its use by animals in evolutionary ecology. *Trends Ecol. Evol.* **20**, 187–193. (doi:10.1016/j.tree.2005.01.010)
61. Glaze CM, Filipowicz AL, Kable JW, Balasubramanian V, Gold JI. 2018 A bias–variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nat. Hum. Behav.* **2**, 213–224. (doi:10.1038/s41562-018-0297-4)
62. Stephens DW, Brown JS, Ydenberg RC. 2008 *Foraging: behavior and ecology*. Chicago, IL: University of Chicago Press.
63. Cuthill IC, Kacelnik A, Krebs JR, Haccou P, Iwasa Y. 1990 Starlings exploiting patches: the effect of recent experience on foraging decisions. *Anim. Behav.* **40**, 625–640. (doi:10.1016/S0003-3472(05)80692-X)
64. Van Gils JA, Schenk IW, Bos O, Piersma T. 2003 Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches. *Am. Nat.* **161**, 777–793. (doi:10.1086/374205)
65. Drugowitsch J, Moreno-Bote R, Pouget A. 2014 Optimal decision-making with time-varying evidence reliability. In *Advances in neural information processing systems* (eds Z Ghahramani, M Welling, C Cortes, N Lawrence, KQ Weinberger), pp. 748–756. Red Hook, NY: Curran Associates, Inc.
66. Ratcliff R, McKoon G. 2007 The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922. (doi:10.1162/neco.2008.12-06-420)
67. Morris DW. 2003 Toward an ecological synthesis: a case for habitat selection. *Oecologia* **136**, 1–13. (doi:10.1007/s00442-003-1241-4)
68. Nishimura K. 1992 Foraging in an uncertain environment: patch exploitation. *J. Theor. Biol.* **156**, 91–111. (doi:10.1016/S0022-5193(05)80658-7)
69. Klaassen RH, Nolet BA, van Gils JA, Bauer S. 2006 Optimal movement between patches under incomplete information about the spatial distribution of food items. *Theor. Popul. Biol.* **70**, 452–463. (doi:10.1016/j.tpb.2006.04.002)
70. Hughes JJ, Ward D, Perrin MR. 1995 Effects of substrate on foraging decisions by a Namib Desert gerbil. *J. Mammal.* **76**, 638–645. (doi:10.2307/1382372)
71. Valone TJ. 1991 Bayesian and prescient assessment: foraging with pre-harvest information. *Anim. Behav.* **41**, 569–577. (doi:10.1016/S0003-3472(05)80897-8)
72. Marschall EA, Chesson PL, Stein RA. 1989 Foraging in a patchy environment: prey-encounter rate and residence time distributions. *Anim. Behav.* **37**, 444–454. (doi:10.1016/0003-3472(89)90091-2)
73. Kacelnik A, Bateson M. 1996 Risky theories—the effects of variance on foraging decisions. *Am. Zool.* **36**, 402–434. (doi:10.1093/icb/36.4.402)
74. Cisek P, Puskas GA, El-Murr S. 2009 Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* **29**, 11 560–11 571. (doi:10.1523/JNEUROSCI.1844-09.2009)
75. Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A. 2012 The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* **32**, 3612–3628. (doi:10.1523/JNEUROSCI.4010-11.2012)
76. Malhotra G, Leslie DS, Ludwig CJ, Bogacz R. 2018 Time-varying decision boundaries: insights from optimality analysis. *Psychon. Bull. Rev.* **25**, 971–996. (doi:10.3758/s13423-017-1340-6)
77. Nathan R. 2008 An emerging movement ecology paradigm. *Proc. Natl Acad. Sci. USA* **105**, 19 050–19 051. (doi:10.1073/pnas.0808918105)
78. Jackson BJ, Fatima GL, Oh S, Gire DH. 2019 Many paths to the same goal: metaheuristic operation of brains during natural behavior. *bioRxiv* 697607. (doi:10.1101/697607)
79. Wikenheiser AM, Stephens DW, Redish AD. 2013 Subjective costs drive overly patient foraging strategies in rats on an intertemporal foraging task. *Proc. Natl Acad. Sci. USA* **110**, 8308–8313. (doi:10.1073/pnas.1220738110)
80. Greene CH. 1986 Patterns of prey selection: implications of predator foraging tactics. *Am. Nat.* **128**, 824–839. (doi:10.1086/284608)
81. Giraldeau L-A, Caraco T. 2018 *Social foraging theory*, vol. 73. Princeton, NJ: Princeton University Press.
82. Harper DGC. 1982 Competitive foraging in mallards: 'ideal free' ducks. *Anim. Behav.* **30**, 575–584. (doi:10.1016/S0003-3472(82)80071-7)
83. Templeton JJ, Giraldeau L-A. 1996 Vicarious sampling: the use of personal and public information by starlings foraging in a simple patchy environment. *Behav. Ecol. Sociobiol.* **38**, 105–114. (doi:10.1007/s002650050223)
84. Pérez-Escudero A, Polavieja GGD. 2011 Collective animal behavior from Bayesian estimation and probability matching. *PLoS Comput. Biol.* **7**, e1002282. (doi:10.1371/journal.pcbi.1002282)
85. Hein AM, Altshuler DL, Cade DE, Liao JC, Martin BT, Taylor GK. 2020 An algorithmic approach to natural behavior. *Curr. Biol.* **30**, R663–R675. (doi:10.1016/j.cub.2020.04.018)