

Research



Cite this article: Kikuchi DW. 2024

Who innovates? Abundance of novel and familiar food changes which animals are most persistent. *Proc. R. Soc. B* **291**: 20231936. <https://doi.org/10.1098/rspb.2023.1936>

Received: 29 August 2023

Accepted: 11 December 2023

Subject Category:

Behaviour

Subject Areas:

behaviour, cognition, theoretical biology

Keywords:

innovation, novelty, foraging, neophobia, diet, dynamic programming

Author for correspondence:

David W. Kikuchi

e-mail: david.kikuchi@oregonstate.edu

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.6992004>.

Who innovates? Abundance of novel and familiar food changes which animals are most persistent

David W. Kikuchi

Department of Integrative Biology, Oregon State University, Corvallis, OR, USA

DWK, 0000-0002-7379-2788

Novel behaviours are the raw material of cultural evolution, yet we do not have a clear picture of when they are likely to arise. I use a state-dependent model to examine how individual age and energy reserves interact with the abundance of known and novel prey to promote dietary innovation (incorporating a new food item into the diet). I measure innovativeness as persistence in attempting to capture novel prey. I find a trend towards greater persistence among younger individuals. Decreased abundance of known prey and increased abundance of novel prey also favour persistence. However, many exceptions to these trends occur. These exceptions are critical because they may explain inconsistencies among studies of animal innovation. Care must be taken in experiments to control for multiple factors relevant to an animal's energy budget and foraging opportunities. We may learn more about innovation in experimental contexts by (i) manipulating the abundance of novel and familiar food resources, (ii) directly measuring animal age and condition, and—where possible—(iii) fitting nonlinear models to innovative behaviour. Results indicate that selection for persistence may also favour neophilia.

1. Introduction

Ecological innovation happens when an animal learns to use a new resource or to exploit a current resource in a new way [1,2]. Innovations are hypothesized to impact diverse ecological and evolutionary processes [3–5]. Given these myriad consequences, we need to know when animals are most likely to innovate. However, there is no consensus on the factors that govern innovation. Possibilities include an animal's state, ecological circumstances and personality [6,7]. Evidence supports multiple hypotheses, although with a low degree of consistency [6,8]. Differences between methodologies may explain some variation, but given the sheer quantity of variation, it seems plausible that many hypotheses may be only conditionally true. For example, 'hungry individuals should innovate' [7] could be true when ecological circumstances are poor, whereas 'individuals with extra resources should innovate' [9] could be true when ecological circumstances are good. We need theory to reconcile the conflicting predictions of simple linear hypotheses across a broad range of individual states and ecological conditions.

In this study, I integrate individual state (age and energy reserves) with ecological variables (known resource availability and ecological opportunity to innovate) into a theoretical model that predicts innovativeness. I focus on innovations in foraging behaviour where an animal adds a novel item to its diet. Dietary innovations are widely observed (e.g. [10,11]) and subjected to experimental study (e.g. [12,13]). They are also directly connected to community ecology. Contemporary challenges such as biological invasion or reintroduction can often be recast as problems where innovation changes the outcome of predator–prey interactions [14–16]. Theory on dietary innovations will therefore be a useful tool for empiricists. Furthermore, the foundations of community ecology are based on consumer–resource dynamics [17]. Knowing when new

resources are likely to be added to a consumer's diet will help further weave innovation and social learning into ecological theory [18–20].

'Innovation' also encompasses novel behavioural patterns [1,11], sometimes called 'technical innovations' [11,21]. Technical innovations are thought to be more cognitively demanding than opportunistic dietary innovations, and they depend on behavioural repertoire diversity, motor development and physical strength [22]. Technical innovations will need separate theory that considers such traits and has a larger range of behavioural possibilities than simply accepting a novel prey item.

The present study is built on two branches of theory: models of risk-taking (reviewed in [23,24]) and the exploration–exploitation trade-off [25,26]. In formal models of risk-taking, a risk-prone animal prefers a high-variance option over a low-variance option when the two options have equal expected rewards [23,27]. The mean and variances are known to the animal. If an animal chooses the high variance option, it may go hungry, but it could alternatively experience a string of good luck during which it has (temporarily) higher gains. It is often assumed in risk-taking models that an animal has a level of energy reserves it must maintain to survive [27]. If the relationship between energy reserves and survival is nonlinear, the value of the two options will not be equal. An animal will be favoured to choose either the high- or low-risk option depending on, among other things, its energy reserves. Similarly, energy reserves may determine an animal's choice between a known food option or an unknown food option. However, an animal exploring unknown food also faces uncertainty (i.e. the parameters that describe an option, such as mean and variance, are not known), which is distinct from a known risk. To model how an animal would best manage uncertainty, I use theory developed to solve the exploration–exploitation trade-off.

The exploration–exploitation trade-off is a general problem of whether to use only known information to make decisions (being 'greedy' by maximizing current expected reward [28]), or invest in exploration that may carry future benefits because it improves the information upon which decisions are based. An instance of this problem is the 'bandit' model, where an actor chooses between different options with unknown rates of return, learning from each choice it makes [29]. Bandit models have been used to model processes such as learning the value of foraging patches [30,31], exploring potentially unpalatable prey [32] and neophobia [33]. A key feature of bandit models is that the actor's estimate of the average payout of each option changes as it samples, which makes it distinct from most models of risk-taking (but see [34]). Sherratt & Morand-Ferron [33] made a bandit model of neophobia that examines how an animal would optimally respond to novel prey. Their approach captured the exploration–exploitation trade-off, but did not include energy reserves as a state variable. By combining bandit models and models of risk-taking, we can ask how animals with different ages and energy reserves might choose between familiar and novel options.

I model a scenario where a predator can forage on a known, familiar type of prey, or explore an unknown, novel type of prey. I include information about the abundance of each prey in the predator's decision-making process. There are two specific objectives of this study. First, I quantify how innovative predators are at different ages and energy reserves. Second, I determine how innovativeness changes with the

abundance of familiar and novel prey. The specific aspect of 'innovativeness' that I focus on is persistence. I define persistence as the number of attempts on the novel prey without success, i.e. investment in innovation without reward. These unsuccessful attempts are part of the exploration process for the predator. Assuming all predators are equal in their capacity to capture the novel prey, the most persistent are the most likely to innovate (i.e. succeed) across multiple attempts. I make no assumptions about intrinsic innovativeness as an aspect of intelligence or creativity. In many systems, persistence increases the odds of an animal becoming an 'innovator'—i.e. successfully capturing novel prey [35]. In empirical studies of innovation in natural populations, dietary innovations are often operationalized as reports of species feeding on novel prey items (e.g. [10]), so this definition is compatible with existing literature. Persistence is also relatively easy to measure in an experimental setting.

2. Methods

I assume a predator makes a sequence of foraging decisions over a finite time period. Time is divided into discrete time steps, where time step is denoted by t , and bounded so that $0 < t \leq t_{\max}$. The predator has energy reserves x . It is unable to assimilate more energy if its reserves reach x_{\max} . It dies if its reserves fall below 0 (so it achieves 0 fitness). The predator loses m energy units each time step due to metabolism. I assume that predator fitness is determined by its reproductive output. Reproduction only occurs if it reaches t_{\max} with energy reserves $0 < x \leq x_{\max}$ (i.e. semelparity). Its fitness is a linear function of x at t_{\max} , so that if it survives to t_{\max} , it has fitness equal to x , otherwise it has fitness of 0.

I assume there are two types of prey that the predator can potentially forage on. In every time step, it can hunt a known, familiar prey (type 1), although its hunting success is not guaranteed. For example, a crow could seek rodents, but not reliably encounter rodents during a time of scarcity. However, sometimes, the predator randomly encounters an unfamiliar, novel prey (type 2) that it can choose to pursue, abandoning its quest for familiar prey during that time step. For example, investing time in trying to handle a novel, toxic cane toad will rob a crow of the opportunity to attack rodents. Thus, in time steps when it encounters a novel prey, the predator is briefly faced with a simultaneous choice (figure 1). By allowing simultaneous choice between the two prey, I can vary their abundances independently from each other.

I assume that if the predator does not encounter a novel prey in a particular time step (which happens with probability $1 - q_2$), it always hunts the familiar prey. When the predator hunts for the familiar prey, its probability of successfully locating it is q_1 . I assume q_1 is known to the predator. Furthermore, because the question of interest is how abundance affects innovation, I make the simplifying assumption that if the known prey is encountered, it is always killed. However, at each time step in which it encounters a novel prey (probability q_2), the predator must make a choice of whether to pursue the novel prey or the familiar prey. The predator cannot attack the novel prey unless it is randomly encountered. The probability of encountering novel prey q_2 is known to the predator. I assume that q_1 and q_2 are fixed. This assumption means that the decisions the predator makes between 0 and t_{\max} take place on a faster timescale than changes to prey populations.

In contrast to q_2 , the probability of successfully killing the novel prey s is uncertain and must be inferred. Building on a model by Sherratt [32], I assume that the predator uses Bayesian inference to estimate s . Bayesian inference uses prior beliefs and experiential information to estimate a posterior probability,

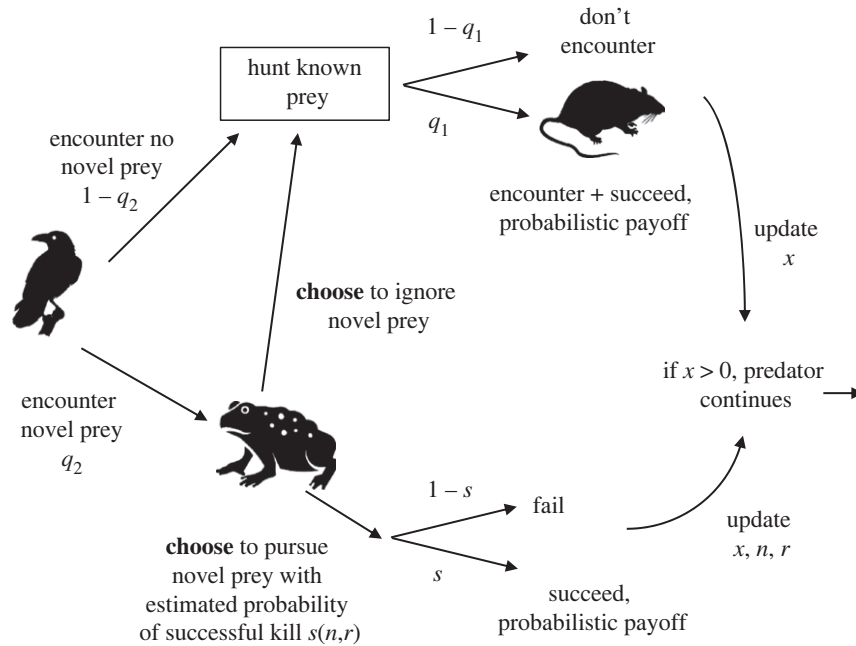


Figure 1. The sequence of events and decisions a predator (crow) faces in the model. First, it does or does not encounter a novel prey. Second, it chooses between hunting known (rodent) and novel (cane toad) prey, or defaults to hunting known prey. Third, it either succeeds or fails in capturing prey. Fourth, its energy reserves (and informational state, if relevant) are updated according to the outcome of events. See main text and table 1 for variable definitions.

which makes it ideal for adaptive models of learning [36,37]. I assume a Beta distribution to represent the predator's prior belief about s . The Beta distribution has parameters α_s and β_s and, initially, an expectation equal to $\alpha_s/(\alpha_s + \beta_s)$. After observing r successes in n trials, the Beta distribution can be updated, yielding a posterior distribution for the expectation of s : $\pi_s(n,r) = (\alpha_s + r)/(\alpha_s + \beta_s + n)$. Assuming priors of $\alpha_s = \beta_s = 1$, we can conveniently write the function for the predator's estimate of success with novel prey as $\pi_s(r,n) = (1 + r)/(2 + n)$. These prior values yield a naive expectation of $\pi_s = 0.5$, but beliefs about s are rapidly modified by experience (figure 2). This method of modelling learning about the probability of success in foraging has been discussed elsewhere [30–32,36,38,39]. For this reason, I do not explore it in detail, e.g. by adopting alternative priors.

Because the predator knows the probability of a successful hunt if it chooses the familiar prey (encounter = q_1 , kill probability = 1), but must learn about the probability of making a kill if it chooses the novel prey (encounter = q_2 , kill probability estimated as $\pi_s(n,r)$), the predator's choice constitutes a one-armed bandit problem [28,40,41]. Note that because the predator loses m energy per time step, if it fails to capture prey, it has a net fitness loss. Thus, novel prey that are difficult to capture could, for some applications, be viewed as having a defence that penalizes attacks.

When the predator encounters a novel prey, I assume the predator chooses the behaviour (attack known, attack novel) that maximizes its expected fitness, given its current age t , its current energy reserves x , and its current estimate of s expressed as $\pi_s(n,r)$. The optimal choice depends not only on the prey type that yields the highest expected reward at the current time step t , but also on the potential future benefits that the predator might realize between time $t + 1$ and t_{\max} from learning more about the novel prey. For example, the predator may expect that the novel prey is no better than the familiar prey, but if the novel prey turns out to be easier to capture than expected, the predator will profit from this information in the future. To find optimal actions in a series of sequential decisions based on state, where future states are determined by those decisions, dynamic programming is the standard method [27,36,42].

Stochastic dynamic programming uses a recursion equation to describe the current expected payoffs as functions of future states




Figure 2. Change in the predator's estimated probability of success on the novel prey $\pi_s(n,r)$ as a function of the number of prey it has sampled. In the solid line, all of its attempts are unsuccessful. In the dashed line, it experiences a stochastic sequence of successes and failures that approaches the true expected value of success (in this case, 0.5).

that may arise due to decisions made or random events that occur. I write the payoff to the predator at each timestep as $z(t,x,n,r)$, where z is the expected fitness given state variables t , x , n , r . At t_{\max} , $z = x$. The dynamic programming equation that solves for z can be used to determine the predator's optimal action in any state. It is as follows:

$$z(t,x,n,r) = (1 - q_2) \cdot z_1(t,x,n,r) + q_2 \cdot \max[z_1(t,x,n,r), z_2(t,x,n,r)]. \quad (2.1)$$

The first term describes the probability of not encountering the novel prey (in which case the predator pursues the known prey) multiplied by $z_1(t,x,n,r)$, the expected payoff for attacking the known prey. The second term describes the probability of encountering novel prey multiplied by the maximum of

the expected payoffs for attacking familiar prey or novel prey ($z_2(t, x, n, r)$).

To calculate $z_1(t, x, n, r)$, recall that the predator has chance q_1 of encountering (and killing) the known prey. If it does, it obtains a benefit with possible values b_{1a} , b_{1b} and b_{1c} , which have probabilities p_{1a} , p_{1b} and p_{1c} , respectively. I assume that all these quantities are known to the predator. I use this distribution of values to avoid artefacts that may arise when running a dynamic programming algorithm on a grid of discrete values [27]. Recall that whether or not the predator is successful, it pays metabolic cost m . The value $z_1(t, x, n, r)$ does not include changes to n or r because the predator gains no information by sampling the known prey.

The calculation of $z_2(t, x, n, r)$ is similar. The predator has the chance s of successfully capturing the prey. If it does, it obtains a benefit with possible values b_{2a} , b_{2b} and b_{2c} , which have probabilities p_{2a} , p_{2b} and p_{2c} , respectively. Regardless, however, it will improve its estimate of s , as reflected by updates to n , r if it is successful, or n alone if it is unsuccessful. Again, the predator always pays the metabolic cost m . Equations for $z_1(t, x, n, r)$ and $z_2(t, x, n, r)$ are

$$\begin{aligned} z_1(t, x, n, r) = & q_1 \cdot [p_{1a}z(t+1, x-m+b_{1a}, n, r) \\ & + p_{1b}z(t+1, x-m+b_{1b}, n, r) \\ & + p_{1c}z(t+1, x-m+b_{1c}, n, r)] \\ & + (1-q_1) \cdot z(t+1, x-m, n, r), \end{aligned} \quad (2.2)$$

$$\begin{aligned} z_2(t, x, n, r) = & \pi_s(n, r) \cdot [p_{2a}z(t+1, x-m+b_{2a}, n+1, r+1) \\ & + p_{2b}z(t+1, x-m+b_{2b}, n+1, r+1) \\ & + p_{2c}z(t+1, x-m+b_{2c}, n+1, r+1)] \\ & + (1-\pi_s(n, r)) \cdot z(t+1, x-m, n+1, r). \end{aligned} \quad (2.3)$$

I assume that if $z_1(t, x, n, r) = z_2(t, x, n, r)$ (within a floating point margin of error), then the predator chooses randomly. A description of all terms used in this study is given in table 1.

To solve the dynamic programming equation, one works backwards from t_{\max} , considering the predator's fitness at all possible combinations of the state variables $t = t_{\max}$, x , n and r . The resulting values are used to find the optimal actions at earlier time steps. This process is called backwards iteration [27,36]. To calculate quantities in addition to the optimal actions and their associated fitness values, one performs forward iterations [27,36], which are Monte Carlo simulations where one starts a predator at a particular combination of t , x , n and r , and then measures the quantities of interest as it encounters prey and makes decisions at each additional time step. The predator's behaviour during the forward iterations is governed by expected payoffs looked up from the array of backwards iterations.

To measure how persistent predators are in different states, in forward iterations, I set the actual probability of success to 0, and set q_2 to 1 so that novel prey are always encountered. For an individual in state t, x , the possible number of novel prey it can attack ranges from 0 to $t_{\max} - t$ (the number of remaining time steps). At each starting combination of t, x where an individual attacks its first novel prey, I find the total number of novel prey n it samples over the remaining time. This includes breaks where it samples known prey, which have stochastic payoffs. I average n across 30 runs. This gives the mean persistence of an individual starting in state t, x in the face of repeated failure when it attacks novel prey.

The central question of this paper is how state-dependent decisions interact with ecology. To answer this, I repeat the analysis of persistence over a grid of values for q_1 and q_2 from 0.1 to 0.9. I interpret the parameter q_1 as representing the abundance of the known prey (it could also represent the probability of both encountering and successfully consuming the known prey, or the probability that the known prey provides a nutritional reward). The essential point is that q_1 governs how much the predator can

Table 1. Parameters and variables used in this study.

term	definition
t	number of time steps over which the animal's behaviour is evaluated, e.g. its lifespan ($0 \leq t \leq t_{\max}$)
x	animal's energy reserves ($0 < x \leq x_{\max}$)
m	animal's metabolic rate
q_1	probability of encountering known prey
q_2	probability of encountering novel prey
s	true probability of successfully killing novel prey (must be inferred)
$\pi_s(n, r)$	posterior expectation of s
n	number of novel prey the predator has encountered and attempted to attack
r	number of attacks on novel prey that have resulted in a kill
α_s, β_s	priors for successfully killing novel prey
b_{1a}, b_{1b}, b_{1c}	possible payoffs for killing known prey
p_{1a}, p_{1b}, p_{1c}	probabilities of each payoff for killing known prey
b_{2a}, b_{2b}, b_{2c}	possible payoffs for killing novel prey
p_{2a}, p_{2b}, p_{2c}	probabilities of each payoff for killing novel prey
$z(t, x, n, r)$	expected payoff for a given set of state variables t, x, n, r
$z_1(t, x, n, r)$	expected payoff for attacking the known prey, given t, x, n, r
$z_2(t, x, n, r)$	expected payoff for attacking the novel prey, given t, x, n, r

rely on the known prey to increase its energy reserves. When q_1 is very low, the predator risks starving if it does not have any other options. With the novel prey, encounter rate is described by q_2 and hunting success is described by s , so q_2 is straightforwardly interpreted as abundance of the novel prey.

(a) Sensitivity to the fitness function

It is important to test sensitivity to the quantity that is maximized in a dynamic programming equation [27]. When fitness is proportional to energy reserves (semelparity), it can incentivize much risk-taking even late in life to achieve maximal reserves x at t_{\max} . Maximizing reserves at the terminal time is not realistic for animals that innovate on timescales shorter than their entire lives. Instead, it might be more useful to model a fixed period of time during which an animal must simply survive. For example, an animal might need to survive the winter by foraging primarily on familiar foods until new ones become available. Alternatively, it might need to exploit an abundant and known, yet transient resource like an irruption of butterflies, while exploring novel prey to switch to when the butterflies die [43]. Furthermore, a short, survival-maximizing timescale will be easier to work with in many experimental settings as long as it corresponds to the actual ecology of the study organism. Some early risk-taking may be favoured by survival maximization, but not so much as with semelparity.

To implement a step function for survival at t_{\max} , I assumed the terminal fitness function has the value 1 if $0 < x \leq x_{\max}$ and 0




Figure 3. Heatmaps of persistence behaviour as a function of time and energy reserves at four combinations of known (q_1) and novel (q_2) prey abundance. The colour key gives the maximum number of prey sampled (the lower bound is approx. 1 plus the number for the colour above).

otherwise. All other aspects of this model remained identical to the one above. This means that as above, between $t = 1$ and $t = t_{\max} - 1$, energy reserves were constrained between 0 and x_{\max} (because anytime x falls below 0, the predator dies and achieves 0 fitness).

Across simulations, I used the same parameter values for consistency, except those that were the targets of specific analyses: $t_{\max} = 30$, $x_{\max} = 20$, $m = 1$, $\alpha = \beta = 1$, $b_{1a} = b_{2a} = 1$, $b_{1b} = b_{2b} = 2$, $b_{1c} = b_{2c} = 3$, $p_{1a} = p_{2a} = 0.1$, $p_{1b} = p_{2b} = 0.8$, $p_{1c} = p_{2c} = 0.1$. The distribution of benefits was chosen to be approximately normal with a mean of 2. Note that there is no difference to the distribution of benefits provided by familiar and novel prey. It is only the predator's uncertainty about the novel prey that causes its behaviour towards the two to differ.

(b) Interpretation of the predictions made by a dynamic programming algorithm

The behaviours of an animal predicted by a dynamic programming algorithm may be dizzyingly complex. Each decision integrates a massive number of calculations. However, using dynamic programming to model behaviour does not require that an animal be able to make such calculations. Instead, the algorithm predicts behaviours that might result from many generations of selection in particular contexts—such as what to do when one encounters something new. The target of selection may in fact be a simple ‘rule of thumb’ that approximates more complex rules [44], or intricate systems of neural and hormonal feedback that can produce behaviours of arbitrary complexity in response to an animal's internal and external environment.

3. Results

(a) Persistence with novel prey: tendency to innovate

Forward iterations reveal complex relationships between persistence, age and energy reserves. Figure 3 shows four cases that capture all these relationships, while electronic supplementary material, figure S1, gives a more continuous picture. A result that emerges is that generally, young individuals persist more in sampling novel prey, despite repeated

failure (left sides of the panels in figure 3 are more shaded). Conversely, older individuals are usually less persistent because they have less time remaining to benefit from anything they learn about novel prey. All individuals with energy reserves close to 0 are also less persistent because they cannot afford extensive sampling without risking starvation (light or white regions on bottoms of figure 3 panels).

Persistence generally decreases with the abundance of known prey, and increases with the abundance of novel prey (in figure 3 and electronic supplementary material, figure S1). In some cases, there is a local or global maximum in persistence for younger individuals with low to moderate resource levels (e.g. figure 3, top left panel; electronic supplementary material, figure S1). This ‘children in the middle class’ tendency to persist stems from a lack of incentive to innovate among individuals with somewhat higher energy reserves, and a lack of necessary buffering against risk in individuals with lower energy reserves. For some combinations of known and novel prey abundance, there are multiple intervals of energy reserves x over which individuals of the same age t are either innovative or not. For example, in the top left panel of figure 3, individuals at $t = 1$ are innovative at moderate-to-low and high energy reserves, whereas this pattern occurs at older ages in the top right panel. Additionally, sometimes older individuals with moderate energy reserves persist more than young ones with the same energy reserves (e.g. the bottom of the shaded region in the top right panel of figure 3). This probably happens because young individuals are under selection to avoid the lower bound on energy reserves (starvation) and still have the chance for a run of good luck with known prey. Older predators choose the novel prey because it has higher potential rewards. That older individuals ever persist more than younger ones is somewhat counterintuitive, because older individuals have less time to benefit from any favourable discovery about the novel prey. Recall from equation (2.1) that whether an individual decides to explore novel prey or pursue known prey depends on the balance between expected payoff from the known prey and




Figure 4. The density of innovators recruited at each time and energy reserve combination when 1000 individuals start at each possible x value. Arrows show the starting values of x for individuals that became innovators at later time steps.

the potential value of the novel prey. The model shows that which of these quantities is greater can shift nonlinearly with age.

A region of t, x space where individuals are predicted to innovate does not necessarily mean that innovation will actually occur there. Each individual must begin life at $t=0$, and depending on the energetic reserves it starts life with, may or may not be likely to enter a region of t, x space that favours high persistence. To illustrate this point, I simulated individuals starting life at $t=0$ with initial energy reserves equally distributed along the interval $[1, x_{\max}]$. I measured behaviour over 1000 runs for each starting x . Figure 4 shows where most innovations occurred in t, x space, and corresponds to the ecological conditions in the top left panel of figure 3. Individuals that started in with energy reserves that favoured attacking novel prey quickly became innovators at early time steps (dark left part of shaded regions in figure 4). At later time steps, there were few individuals from those starting conditions left to become innovators. Instead, new innovators at later time steps mostly come from starting states that did not favour innovation (arrows in figure 4). Movement into regions that favour innovation gives those regions dark peripheries and pale interiors.

This analysis shows that the life history of a species (e.g. how it allocates resources to individual offspring) may influence the emergence of innovations in a population. A species that prioritizes provisioning of young at independence may have a different tendency to innovate than one where young achieve independence with minimal energy reserves. This will be important for some ecological applications, for example when innovations can be discovered within subpopulations of an expanding metapopulation.

(b) Changes to the fitness function: terminal survival, or fitness proportional to lifespan

In many cases, it seems probable that an animal will be adapted to explore and innovate on shorter timescales than its entire life. When this is true, the animal is assumed to maximize survival rather than terminal energy reserves. In examining this fitness function, I found that the most obvious difference was that once an animal accrued energy reserves $x > m(t_{\max} - t)$, it chose randomly (electronic supplementary material, figure S2). Achieving this threshold of energy reserves

can be seen as the survival-maximizer's 'goal'. This leads to greater persistence among individuals with high energy reserves, at least in animals that are too young to have reached the threshold. At high abundances of known prey, individuals either choose known prey, or choose randomly.

4. Discussion

The major contribution of this study is to make predictions about which individuals (in terms of age and condition) are most likely to innovate as a function of the probability of encountering known and novel prey. Individuals tend to be more persistent with novel prey when (i) they are younger, (ii) known prey are less abundant and (iii) novel prey are more abundant. However, there are many exceptions to these general trends. They may have important implications for how we interpret inconsistencies between empirical studies or design theoretical studies. First, the effect of the abundance of known prey tends to trump the abundance of novel prey (electronic supplementary material, figure S1), so new theory or experiments need to take this into account. Second, and critically for empirical work, there can be multiple states that favour innovation in a population, such as young individuals with high energy reserves and young individuals with low-to-moderate energy reserves (top left figure 3). These nonlinear wrinkles may, in some cases, explain meta-analyses that fail to find clear support for simple hypotheses such as 'younger, hungrier individuals innovate' or 'individuals in better condition innovate'. The present study reveals that these hypotheses, which may appear mutually exclusive, can be true in the same population at the same time. Experiments would have higher chances of finding a linear relationship between age and energy by choosing individuals with particular ages or conditions, and carefully manipulating the prey abundance that subjects experience. This could be done by creating conditions that mimic, for example, the bottom panels of figure 3. Such experiments would maximize our power to relate these variables to innovation. If a linear effect of age or energy reserves is found where they are predicted by the model, experiments could subsequently test levels of prey abundance that are predicted to break this pattern. By using both these approaches, we could confirm the roles of age, state, and prey abundance in innovation, and then begin to explore complex patterns generated by their interactions. Third, when an animal maximizes terminal survival, if it has sufficient energy reserves not to starve by the end of the time period, it may guess randomly (it could also do nothing, which I do not consider above). Results that appear noisy or inconclusive, or disagree with the results of different studies, may not rule out that a particular state variable influences innovation, but rather warrant further scrutiny of the fitness function.

In some systems, it may be challenging to manipulate how an animal perceives the ecological availability of prey, or to do so over the relevant values. This study suggests that it would nevertheless be worthwhile to test for nonlinear relationships between innovation and state. This is critical to understanding what factors are most important in promoting innovation across study systems. To illustrate why, consider a recent meta-analysis that found personality is an important determinant of innovation, where more exploratory or neophilic individuals innovated more [6]. It seems probable

that this is true in some species. However, linear relationships are much more often tested in experiments than nonlinear ones. If the relationship between personality and innovation is linear, but the relationship between age or condition and innovation is nonlinear, meta-analysis will reveal more significant effects of personality because only linear relationships are tested. Most researchers wisely do not test nonlinear hypotheses without an *a priori* reason to do so, but I argue that this study provides such a reason, at least for testing the effects of age or condition on innovation. Testing nonlinear hypotheses will require approaches that penalize the number of parameters in the best-fitting model [45,46]. It will also place a premium on detailed knowledge of the natural history and behaviour of the study species to ensure that the experiment meets the assumptions of the hypothesis being tested. When testing nonlinear relationships, the availability of known and novel food resources should be reported along with subjects' ages and conditions.

This model also makes predictions about how an individual's level of dietary innovation may change with prey diversity. Due to the rapid expansion of dimensionality when including many state variables in dynamic programming, explicitly modelling a predator's learning about multiple prey is challenging [36]. However, modelling multiple prey may not be necessary to answer the question of when an animal should stop adding new items to its diet. When multiple novel prey are available to explore, each successful innovation may result in diminishing returns on future innovations. In the present model, if a predator repeatedly samples a novel prey item until its expected success is fairly certain, that prey item could be added to the category of 'known' prey. This would have an effect much like increasing the value of q_1 . Increasing q_1 tends to decrease an animal's persistence with novel prey. We might therefore expect that innovation decreases with the number of innovations an animal has already acquired until it ceases altogether. This is predicted to be true irrespective of the number of additional, novel types of prey available as long as diversity is not inversely correlated with the abundance of each option. Individual specialization in falcons may arise through such a process. Several species of falcon have been recorded to each prey upon dozens of kinds of birds, but individual falcons are often highly selective (e.g. [47,48]). This may happen because experience makes falcons more successful at capturing a particular species of prey [48]. Once a falcon can meet its energy needs, it may have little incentive to incorporate additional bird species into its diet.

There are many other factors that have been hypothesized to influence innovation that the present model does not include. The model does not consider the social network an individual occupies. Some hypotheses for how state influences innovation operate through social interactions. For example, the hypothesis that energy reserves govern innovation has been interpreted to mean that more dominant individuals—presumably in better condition—will be more innovative [6]. Conversely, others have hypothesized that individuals that are more successful in scramble competition against conspecifics—presumably dominant—may have less motivation to innovate [7]. If social networks have no influence on foraging other than to enhance or limit energy reserves, then the present model may produce useful predictions (although it would be helpful to directly measure condition, rather than using network position as a proxy). However, in contexts where social

network position is correlated with other behaviours such as cultural conformity [49], the assumptions of this model do not hold. Additionally, some authors assume that the amount of effort that an individual invests in innovation trades off with investment in social learning (e.g. [18,50]). When this is true, it may also violate the assumptions of this model, as social learning could reduce the predicted level of investment in innovation. It could also lead to producer–scrounger dynamics when some individuals innovate and others specialize in copying [51]. Furthermore, future work that considers changes in prey abundance on the same time-scale as behavioural innovation—for example when a novel prey rapidly invades a new range or known prey become depleted—would be welcome.

This model is closely related to that of Sherratt & Morand-Ferron [33]. Sherratt & Morand-Ferron [33] used a dynamic approach to model an animal's neophobia or neophilia towards a novel object. Their state variables included an animal's age, number of encounters with the novel object, and experiences (good or bad) with the novel object. Their model used Bayesian inference to estimate the probability of encountering the novel object and its expected payoff. They found that when animals maximized terminal payoff, younger animals should always be more neophilic, in the sense that they should be more willing to interact with novel objects that carry a higher cost–benefit ratio. They also found that the age at which an animal should ignore novel objects increases with its lifespan and decreases with the cost–benefit ratio. Their result that younger individuals are more willing to take risks with a novel object is similar to what I find in most, but not all, circumstances (e.g. Figure 3; electronic supplementary material, figure S1). However, Sherratt & Morand-Ferron [33] found that when animals pay a survival cost to making mistakes and accrue fitness continuously (i.e. iteroparity), young and old individuals were more risk-prone than middle-aged ones.

I find different behaviour from Sherratt & Morand-Ferron [33] with respect to the cost–benefit ratio. When I decrease the mean benefits provided by known and novel prey to 1 (so that the cost–benefit ratio increases because animals must pay the metabolic cost m), animals *increase* in willingness to sample unfamiliar prey at nearly all ages (as measured by persistence; electronic supplementary material, figure S3). This contrasting result with Sherratt & Morand-Ferron [33] occurs because in the present model, individuals face starvation when potential rewards are low. It is preferable to gamble repeatedly on an unknown option than choose one that, on average, leads to starvation. Despite some differences between my results and those of Sherratt & Morand-Ferron [33], there appears to be strong similarity between neophobic behaviour and innovativeness, at least at the ultimate level of analysis used in these models. This is because the appropriate decision when encountering a novel object for the first time (neophobia/neophilia) or when repeatedly attempting to attack a novel prey (persistence) depends on current and expected future payoffs.

Models that predict complex patterns as a consequence of multiple interacting factors risk of fostering a sense of experimental nihilism. However, I hope this study instead helps us design more powerful experiments to test specific hypotheses by accounting for precisely those factors predicted to generate variability in innovativeness. Specifically, controlling for the abundance of known and novel prey could increase the

ability of experiments to detect linear relationships between age, energy reserves and innovation.

Ethics. This work did not require ethical approval from a human subject or animal welfare committee.

Data accessibility. Supplementary material is available online [52].

Declaration of AI use. I have not used AI-assisted technologies in creating this article.

Conflict of interest declaration. I declare I have no competing interests.

Author contributions. D.W.K.: conceptualization, formal analysis, investigation, methodology, writing—original draft.

Funding. This work was partially funded by the German Research Foundation (DFG): SFB TRR 212 (NC3)—project no. 316099922.

Acknowledgements. Thanks to Tom Sherratt and Bryan Lynn for reading earlier drafts of the manuscript and providing very helpful comments.

References

- Greenberg R. 2003 The role of neophobia and neophilia in the development of innovative behaviour of birds. In *Animal innovation* (eds SM Reader, KN Laland), pp. 175–196. Oxford, UK: Oxford University Press. (doi:10.1093/acprof:oso/9780198526223.003.0008)
- Reader SM, Laland KN. 2003 *Animal innovation*. Oxford, UK: Oxford University Press Oxford.
- Cohen TM, Kumar RS, Nair M, Hauber ME, Dor R. 2020 Innovation and decreased neophobia drive invasion success in a widespread avian invader. *Anim. Behav.* **163**, 61–72. (doi:10.1016/j.anbehav.2020.02.012)
- Laland KN, Toyokawa W, Oudman T. 2020 Animal learning as a source of developmental bias. *Evol. Dev.* **22**, 126–142. (doi:10.1111/ede.12311)
- Whitehead H, Laland KN, Rendell L, Thorogood R, Whiten A. 2019 The reach of gene–culture coevolution in animals. *Nat. Commun.* **10**, 2405. (doi:10.1038/s41467-019-10293-y)
- Amici F, Widdig A, Lehmann J, Majolo B. 2019 A meta-analysis of interindividual differences in innovation. *Anim. Behav.* **155**, 257–268. (doi:10.1016/j.anbehav.2019.07.008)
- Laland KN, Reader SM. 1999 Foraging innovation is inversely related to competitive ability in male but not in female guppies. *Behav. Ecol.* **10**, 270–274. (doi:10.1093/beheco/10.3.270)
- Perry S, Carter A, Smolla M, Akçay E, Nöbel S, Foster JG, Healy SD. 2021 Not by transmission alone: the role of invention in cultural evolution. *Phil. Trans. R. Soc. B* **376**, 20200049. (doi:10.1098/rstb.2020.0049)
- Kummer H, Goodall J. 1985 Conditions of innovative behaviour in primates. *Phil. Trans. R. Soc. Lond. B* **308**, 203–214. (doi:10.1098/rstb.1985.0020)
- Lefebvre L, Whittle P, Lascaris E, Finkelstein A. 1997 Feeding innovations and forebrain size in birds. *Anim. Behav.* **53**, 549–560. (doi:10.1006/anbe.1996.0330)
- Overington SE, Morand-Ferron J, Boogert NJ, Lefebvre L. 2009 Technical innovations drive the relationship between innovativeness and residual brain size in birds. *Anim. Behav.* **78**, 1001–1010. (doi:10.1016/j.anbehav.2009.06.033)
- Burkart JM, Strasser A, Foglia M. 2009 Trade-offs between social learning and individual innovativeness in common marmosets, *Callithrix jacchus*. *Anim. Behav.* **77**, 1291–1301. (doi:10.1016/j.anbehav.2009.02.006)
- Quinn JL, Cole EF, Reed TE, Morand-Ferron J. 2016 Environmental and genetic determinants of innovativeness in a natural population of birds. *Phil. Trans. R. Soc. B* **371**, 20150184. (doi:10.1098/rstb.2015.0184)
- Beckmann C, Shine R. 2011 Toad's tongue for breakfast: exploitation of a novel prey type, the invasive cane toad, by scavenging raptors in tropical Australia. *Biol. Invasions* **13**, 1447–1455. (doi:10.1007/s10530-010-9903-8)
- Cabrera-Guzmán E, Crossland MR, Pearson D, Webb JK, Shine R. 2015 Predation on invasive cane toads (*Rhinella marina*) by native Australian rodents. *J. Pest Sci.* **88**, 143–153. (doi:10.1007/s10340-014-0586-2)
- Parrott ML, Doody JS, McHenry C, Clulow S. 2019 Eat your heart out: choice and handling of novel toxic prey by predatory water rats. *Aust. Mammalogy* **42**, 235–239. (doi:10.1071/AM19016)
- McPeck MA. 2022 *Coexistence in ecology: a mechanistic perspective*. Princeton, NJ: Princeton University Press.
- Borofsky TM, Feldman MW. 2022 Success-biased social learning in a one-consumer, two-resource model. *Theor. Popul. Biol.* **146**, 29–35. (doi:10.1016/j.tpb.2022.05.004)
- Borofsky TM, Feldman MW. 2022 Static environments with limited resources select for multiple foraging strategies rather than conformity. *Ecol. Monogr.* **92**, e01483. (doi:10.1002/ecm.1483)
- Kikuchi DW, Simon MW. 2023 Social learning of innovations in dynamic predator-prey systems. *Am. Nat.* **201**, 895–907. (doi:10.1086/724491)
- Sol D, Sayol F, Ducatez S, Lefebvre L. 2016 The life-history basis of behavioural innovations. *Phil. Trans. R. Soc. B* **371**, 20150187. (doi:10.1098/rstb.2015.0187)
- Tebbhich S, Griffin AS, Peschl MF, Sterelny K. 2016 From mechanisms to function: an integrated framework of animal innovation. *Phil. Trans. R. Soc. B* **371**, 20150195. (doi:10.1098/rstb.2015.0195)
- 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)
- Kacelnik A, El Mouden C. 2013 Triumphs and trials of the risk paradigm. *Anim. Behav.* **86**, 1117–1129. (doi:10.1016/j.anbehav.2013.09.034)
- Cohen JD, McClure SM, Yu AJ. 2007 Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Phil. Trans. R. Soc. B* **362**, 933–942. (doi:10.1098/rstb.2007.2098)
- Mehlhorn K, Newell BR, Todd PM, Lee MD, Morgan K, Braithwaite VA, Hausmann D, Fiedler K, Gonzalez C. 2015 Unpacking the exploration–exploitation tradeoff: a synthesis of human and animal literatures. *Decision* **2**, 191–215. (doi:10.1037/dec0000033)
- Houston AI, McNamara JM. 1999 *Models of adaptive behaviour: an approach based on state*. Cambridge, UK: Cambridge University Press.
- Sherratt TN, O'Neill E. 2023 Signal detection models as contextual bandits. *R. Soc. Open Sci.* **10**, 230157. (doi:10.1098/rsos.230157)
- Gittins JC. 1979 Bandit processes and dynamic allocation indices. *J. R. Stat. Soc. B* **41**, 148–164.
- 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)
- Krebs JR, Kacelnik A, Taylor P. 1978 Test of optimal sampling by foraging great tits. *Nature* **275**, 27–31. (doi:10.1038/275027a0)
- Sherratt TN. 2011 The optimal sampling strategy for unfamiliar prey. *Evolution* **65**, 2014–2025. (doi:10.1111/j.1558-5646.2011.01274.x)
- Sherratt TN, Morand-Ferron J. 2018 The adaptive significance of age-dependent changes in the tendency of individuals to explore. *Anim. Behav.* **138**, 59–67. (doi:10.1016/j.anbehav.2018.01.025)
- McNamara JM. 1996 Risk-prone behaviour under rules which have evolved in a changing environment. *Am. Zool.* **36**, 484–495. (doi:10.1093/icb/36.4.484)
- Griffin AS, Guez D. 2014 Innovation and problem solving: a review of common mechanisms. *Behav. Processes* **109**, 121–134. (doi:10.1016/j.beproc.2014.08.027)
- Clark CW, Mangel M. 2000 *Dynamic state variable models in ecology: methods and applications*. Oxford, UK: Oxford University Press on Demand.
- Dall SRX, Giraldeau L, Olsson O, McNamara J, Stephens D. 2005 Information and its use by animals in evolutionary ecology. *Trends Ecol. Evol.* **20**, 187–193. (doi:10.1016/j.tree.2005.01.010)
- Kikuchi DW, Sherratt TN. 2015 Costs of learning and the evolution of mimetic signals. *Am. Nat.* **186**, 321–332. (doi:10.1086/682371)
- Aubier TG, Sherratt TN. 2020 State-dependent decision-making by predators and its consequences for mimicry. *Am. Nat.* **196**, E127–E144. (doi:10.1086/710568)
- Jones PW. 1978 On the two armed bandit with one probability known. *Metrika* **25**, 235–239. (doi:10.1007/BF02204367)

41. Woodroffe M. 1976 On the one arm bandit problem. *Sankhyā: Indian J. Stat. Series A* **38**, 79–91.
42. Bellman R. 1952 On the theory of dynamic programming. *Proc. Natl Acad. Sci. USA* **38**, 716–719. (doi:10.1073/pnas.38.8.716)
43. Langham GM. 2004 Specialized avian predators repeatedly attack novel color morphs of *Heliconius* butterflies. *Evolution* **58**, 2783–2787.
44. Davies NB, Krebs JR, West SA. 2012 *An introduction to behavioural ecology*. New York, NY: John Wiley & Sons.
45. Bolker BM. 2008 *Ecological models and data in R*. Princeton, NJ: Princeton University Press.
46. Burnham KP, Anderson DR. 2004 Multimodel inference: understanding AIC and BIC in model selection. *Sociol. Methods Res.* **33**, 261–304. (doi:10.1177/0049124104268644)
47. Warkentin IG, Sodhi NS, Espie RHM, Poole AF, Oliphant LW, James PC. 2020 Merlin (*Falco columbarius*), version 1.0. *Birds of the World*. (doi:10.2173/bow.merlin.01)
48. White CM, Clum NJ, Cade TJ, Hunt WG. 2020 Peregrine falcon (*Falco peregrinus*), version 1.0. *Birds of the World*. (doi:10.2173/bow.perfal.01)
49. Aplin LM, Farine DR, Mann RP, Sheldon BC. 2014 Individual-level personality influences social foraging and collective behaviour in wild birds. *Proc. R. Soc. B* **281**, 20141016. (doi:10.1098/rspb.2014.1016)
50. Wakano JY, Aoki K. 2007 Do social learning and conformist bias coevolve? Henrich and Boyd revisited. *Theor. Popul. Biol.* **72**, 504–512. (doi:10.1016/j.tpb.2007.04.003)
51. Lehmann L, Feldman MW. 2009 Coevolution of adaptive technology, maladaptive culture and population size in a producer–scrounger game. *Proc. R. Soc. B* **276**, 3853–3862. (doi:10.1098/rspb.2009.0724)
52. Kikuchi DW. 2024 Who innovates? Abundance of novel and familiar food changes which animals are most persistent. Figshare. (doi:10.6084/m9.figshare.c.6992004)