

## RESEARCH ARTICLE

## Evolution of innate behavioral strategies through competitive population dynamics

Tong Liang<sup>1,2</sup>, Braden A. W. Brinkman<sup>2\*</sup>

**1** Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York, United States of America, **2** Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, New York, United States of America

\* [braden.brinkman@stonybrook.edu](mailto:braden.brinkman@stonybrook.edu)**OPEN ACCESS**

**Citation:** Liang T, Brinkman BAW (2022) Evolution of innate behavioral strategies through competitive population dynamics. *PLoS Comput Biol* 18(3): e1009934. <https://doi.org/10.1371/journal.pcbi.1009934>

**Editor:** Iaroslav Ispolatov, Universidad de Santiago de Chile, CHILE

**Received:** July 20, 2021

**Accepted:** February 18, 2022

**Published:** March 14, 2022

**Copyright:** © 2022 Liang, Brinkman. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** Simulation code used in this study, processed data, and scripts for plotting the figures in this work are available at <https://github.com/bradenbrinkman/Evolution-dynamics>.

**Funding:** TL and BAWB thank Stony Brook University (<https://www.stonybrook.edu/>) for financial support. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** The authors have declared that no competing interests exist.

## Abstract

Many organism behaviors are innate or instinctual and have been “hard-coded” through evolution. Current approaches to understanding these behaviors model evolution as an optimization problem in which the traits of organisms are assumed to optimize an objective function representing evolutionary fitness. Here, we use a mechanistic birth-death dynamics approach to study the evolution of innate behavioral strategies in a simulated population of organisms. In particular, we performed agent-based stochastic simulations and mean-field analyses of organisms exploring random environments and competing with each other to find locations with plentiful resources. We find that when organism density is low, the mean-field model allows us to derive an effective objective function, predicting how the most competitive phenotypes depend on the exploration-exploitation trade-off between the scarcity of high-resource sites and the increase in birth rate those sites offer organisms. However, increasing organism density alters the most competitive behavioral strategies and precludes the derivation of a well-defined objective function. Moreover, there exists a range of densities for which the coexistence of many phenotypes persists for evolutionarily long times.

## Author summary

The innate, or instinctual, behavioral strategies that populations of organisms employ to navigate their environments and fend for survival are shaped over epochs of evolutionary selection, in contrast to individual behaviors that can change within an individual’s lifetime based on experience and sensory input. Understanding how the interplay between organism and their environment shapes which behavior strategies emerge as the most successful for a population’s survival is a major problem in mathematical biology. Often, evolution is modeled as an optimization process that selects for behaviors that optimize the “fitness” of organisms in their environment. However, the fundamental evolutionary events are stochastic birth and death events, and the most successful organisms that emerge under these dynamics are not always those predicted by fitness-based approaches. In this work, we use agent-based stochastic simulations and mean-field approximations of a mechanistic population dynamics model to investigate the evolution of a population’s innate foraging strategies. In particular, we investigate when an emergent fitness function

can be derived and how competition between individuals for resources alters the most successful behavioral strategies and precludes the derivation of a simple fitness function that predicts the most successful behavioral strategies.

## Introduction

Foraging for resources—shelter, food, etc.—is one of the fundamental problems an organism must solve in order to survive and reproduce, ensuring the continuity of its species [1]. Many of the strategies organisms employ to solve such problems appear to be instinctual or innate; for example, birds and fish know where to migrate even if they have never visited their migration grounds before [2], and many complex species such as fish, reptiles, birds, and insects have been found to employ search strategies similar to the much simpler *C. Elegans* [3]. The fact that many disparate species have converged upon similar strategies suggests there may be unifying principles that can be captured by simple quantitative models to help explain the emergence of these behavioral patterns.

Perhaps the most well-known candidate principle is “survival of the fittest:” species with adaptations that give them any edge for outcompeting other species in their environments are those that go on to proliferate. On its own this principle is somewhat tautologous, as the organisms deemed fittest for their environment are those observed to survive and proliferate [4]. In order to render this problem well-posed for quantitative modeling, a common approach is to define a “fitness function,” a function of a population’s phenotypic parameters whose output is a proxy for that phenotypic configuration to survive and reproduce in its environment. Often the fitness function is modeled as the fitness of a single representative agent, rather than an entire population. This idea has given rise to modeling evolution as an optimization process that maximizes an organism’s fitness [4], i.e., selects for the phenotype parameters that maximize the fitness function. Fitness functions can range from relatively straightforward and concrete quantities, such as an organism’s rate of producing offspring or the time it takes to find shelter or food [5–8], or they can be relatively abstract, such as the amount of information an organism’s sensory organs convey to its nervous system [9–12]—the idea being that the more efficiently an organism’s nervous system can process information it acquires about its environment, the more successful it will be at finding food, shelter, mates, etc.

Modeling evolution as an optimization process has had some notable successes in predicting features of biological organisms [12–17]; however, some caveats and cautions are in order. Firstly, one can always find an objective function and set of constraints for which an observed solution is optimal [18]. Secondly, there are examples in which evolution is demonstrably not optimal [19], exhibits significant deviations from optimality [16], or can even be repeatedly reversible [20, 21]. These and other issues with modeling evolution as optimization have been discussed extensively by Doebeli *et al.* [4]. Doebeli *et al.* emphasize that the fundamental events in evolution are stochastic birth-death processes, which can give rise to complex evolutionary dynamics that cannot always be captured by optimizing fitness. For instance, an approach that focuses on birth-death events of entire populations can yield the distribution of successful phenotypes, including phenotypes that would be considered “sub-optimal” by optimization approaches yet are able to coexist with phenotypes predicted to be optimal.

In this paper we take the mechanistic approach, simulating the dynamics of a population of organisms with different innate behavioral strategies for exploring their environment in search of locations with favorable conditions for reproduction. In particular, non-consumable resources are distributed randomly throughout the environment, and each organism’s

behavioral strategy consists of the set of rates at which the organism will leave locations of each resource level in search of other locations with potentially greater amounts of the resource. Each location in the environment can only support a single individual, resulting in competitive population dynamics. The most competitive phenotypes in our model will be the set of relocation rates that enable organisms to proliferate, allowing these phenotypes to persist in the population over long time scales. Importantly, the birth rates of organisms in our model depend only on an individual's location within the environment, not that individual's phenotype (i.e., locations with plentiful resources confer larger birth rates, regardless of the phenotype of the organism). Thus, the reproductive competitiveness of a phenotype is only determined by how that phenotype enables our model organisms to outcompete other phenotypes in discovering high-resource locations. There is no explicit fitness function that the dynamics are assumed or enforced to optimize, but we show that in the limit of low interaction density we can derive a growth rate that is a function of the relocation rate parameters and serves as an emergent fitness function. However, we also show that in general this function does not describe the most successful phenotypes when interactions between individuals are frequent, or when demographic fluctuations are significant.

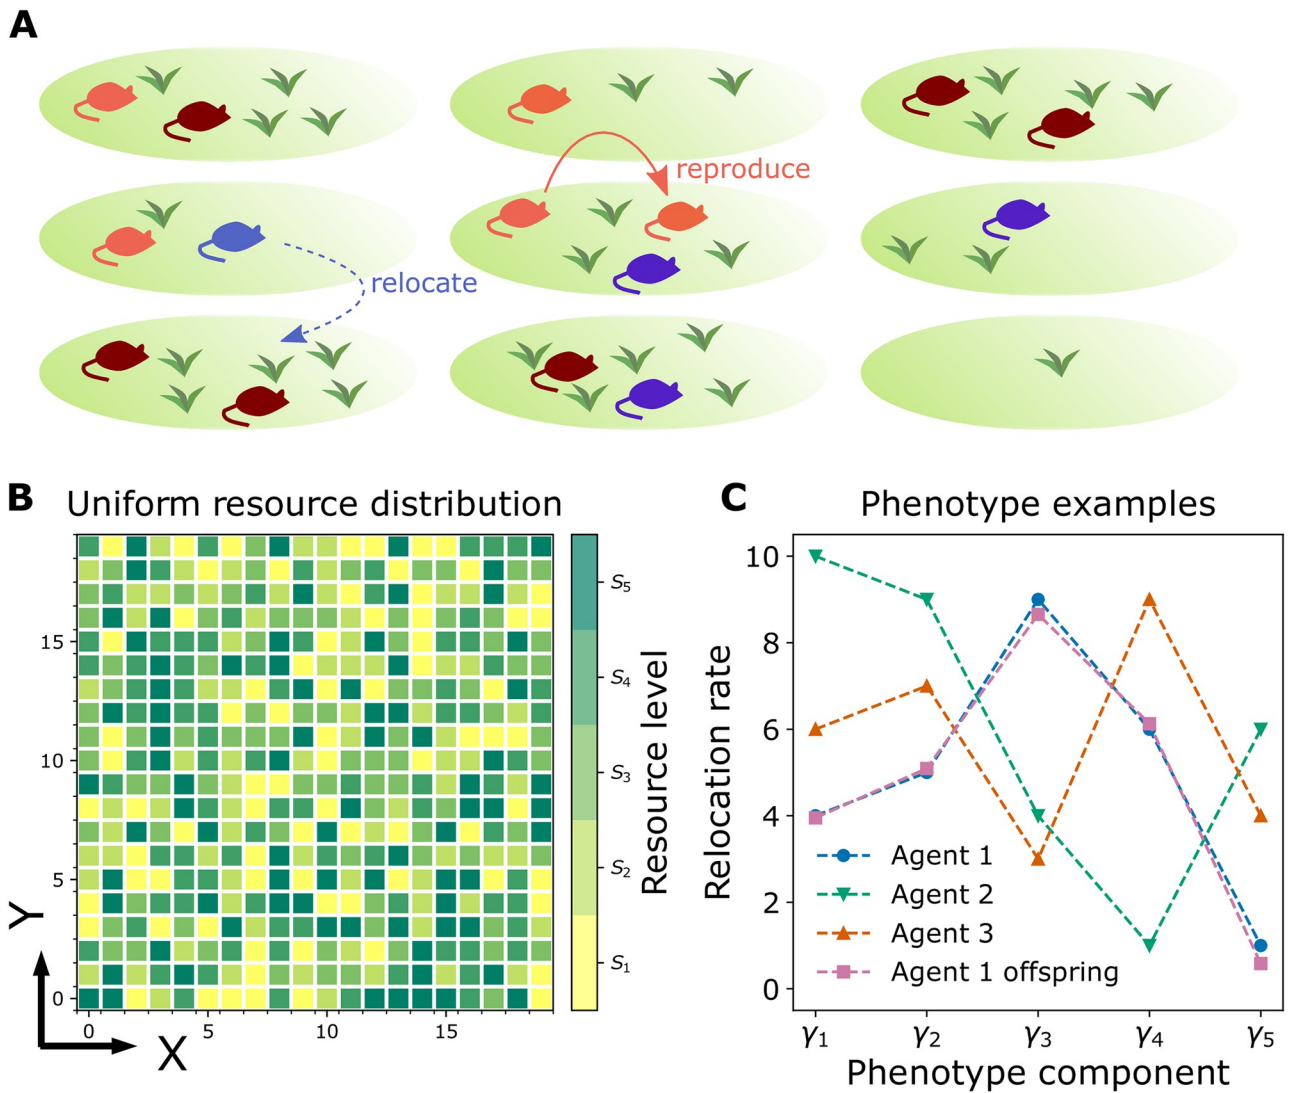
This paper is organized as follows: in [Results](#) we first define our stochastic model and its formulations as an agent-based simulation and a corresponding analytic description in terms of a stochastic master equation (Mechanistic model of evolutionary dynamics). We then show that the environmental structure and the population dynamics of the simulated model indeed select for particular phenotypes, even though birth rates depend only on spatial location and *not* the phenotype. Importantly, we find that while some components of the phenotype experience selection pressures strong enough to select for a very narrow range of values, others experience relative weak selection pressure, resulting in distributions of nearly equally competitive phenotypes (Distributions of competitive phenotypes are selected by birth-death dynamics). We show that we can understand the origins of weak selection pressures using a minimal mean-field model, in which we can derive the growth rate as a function of phenotype, which acts as an effective fitness function, when organism density is low (A minimal mean-field model explains variability in competitive phenotypes). In this limit we also study how the exploration-exploitation trade-off depends on the scarcity of high-resource sites and their corresponding birth rate (Distribution of resources shapes exploration-exploitation trade-offs). At higher density the interactions between individual conspecifics cannot be neglected and alters the dominant phenotypes: at very high densities there is a reversal in the most competitive phenotypes, and there is a range of intermediate densities at which many phenotypes coexist for evolutionarily long times (Ranges of population density allow for long-lasting phenotype coexistence). We conclude in the [Discussion](#) by addressing the interpretation of our results in the context of mechanistic and normative approaches to the evolution of biological structure (Summary and interpretation of results) and comparing our results to those of related studies (Comparisons to related work), as well as discussing the limitations of our approach (Limitations of the study) and future directions for extensions of this work (Future directions).

## Results

### Mechanistic model of evolutionary dynamics

To investigate the evolution of foraging behavioral strategies, we focus on modeling a population of organisms that must navigate their environment in search of locations favorable for reproduction. For simplicity, we will assume the favorability of each location is determined by the amount of a non-depleting “resource” at this location. Additionally, each location has a

fixed carrying capacity, allowing only a single individual to occupy the location at any given moment. Organisms are able to detect the resource level only at their current locations, and must decide whether to remain where they are or explore in search of regions with more resources (see Fig 1A); this is referred to as the “exploration-exploitation trade-off” [22]. In this model the organisms’ movements are not deterministic, but instead each individual will attempt to relocate stochastically at rates determined by the amount of resource at the individual’s current location; the value of these rates is determined by the individual’s phenotype. The



**Fig 1. Schematic diagram of the mechanistic evolutionary dynamics model.** **A.** The model consists of a population of agents foraging in an environment with a random distribution of resources, in which each agent can either asexually produce an offspring (as shown by the orange arrow) or relocate to one of the vertical or horizontal neighboring locations (as shown by the dashed blue arrow). Different colors here distinguish agents with different phenotypes. **B.** The environment is a two dimensional discrete lattice on which each location has a predetermined resource level  $S_k$ ,  $k \in \{1, 2, 3, 4, 5\}$ , chosen from a distribution (a discrete uniform distribution shown here). A  $20 \times 20$  subset of the environment is shown here. Each site on the lattice is colored according to its resource level  $S_k$ . Agents at locations with resource level  $S_k$  will reproduce at a birth rate  $b_k$ . **C.** Each agent has a phenotype vector with 5 components, corresponding to the relocation or hopping rates of the agent at 5 different resource levels. These components are initially chosen randomly for each agent and are shaped over the course of evolution. In the example shown here, agent 2 has a relocation rate  $\gamma_3 \approx 4$  when it is at a site with resource level  $S_3$ . In the agent-based stochastic model, an offspring inherits a noisy copy of its parent’s phenotype, as shown here for the phenotypes of agent 1 and its offspring.

<https://doi.org/10.1371/journal.pcbi.1009934.g001>

resource level at each location is predetermined and sampled from a discrete distribution; see Fig 1B, depicting a discrete uniform resource distribution. The set of relocation rates (used interchangeably with the term “hopping rates”) for each resource level comprises an organism’s phenotype in this model; see Fig 1C.

In addition to determining an organism’s relocation rate, the resource level at each site directly determines the organisms’ birth rates at that location. Importantly, the birth rates *do not* depend on the organisms’ phenotype—the only way the phenotype influences the organism’s reproductive success is by improving the organism’s odds of encountering high-resource sites and staying there long enough to produce sufficiently large numbers of offspring. That is, the selective advantage of a phenotype is due only to its dynamical success in the environment, not any intrinsic advantage. For simplicity, we model all reproduction as asexual. In the agent-based model to be defined shortly, the offsprings’ phenotypes are taken to be noisy copies of the parents’ phenotype (see “Agent 1” and “Agent 1 offspring” in Fig 1C). In the mean-field model discussed in the next section, the offsprings’ phenotypes will be direct copies of the parents’.

The final ingredient the model needs is a means of selection. Typically, one would allow for stochastic death of individuals, such that phenotypes that do not confer a sufficiently competitive advantage will eventually die out, leaving the more competitive phenotypes to dominate the population. However, a technical downside of stochastic death is that extinction becomes not only possible but inevitable, potentially complicating investigation of long-term phenotype distributions. To avoid the technical complications of stochastic death, we instead induce selection by subjecting the population to a periodically-occurring “catastrophe” that randomly culls individuals until the population is reduced to its initial size. In doing so, the population is never at risk for extinction, but selection still occurs because the phenotypes that are able to produce large numbers of offspring have better odds of surviving each catastrophe.

To formalize this schematic picture, we first build an agent-based stochastic model. In this model the population consists of individual organisms, or “agents,” each of which is assigned a location within a  $2d$  environment, modeled as a discrete lattice, and the values of its hopping rates  $\gamma$  at each resource level. In the simulations we investigate here, we assume the resource level of site at location  $\mathbf{x}$ , labeled  $s_{\mathbf{x}}$ , is drawn from 5 distinct resource levels,  $s_{\mathbf{x}} \in \{S_1, S_2, S_3, S_4, S_5\}$ . Thus each agent has 5 corresponding hopping rates  $\gamma = \{\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5\}$ , one for each of the five resource levels. The values of each of these hopping rates are initially chosen randomly from the range  $[\gamma_{\min}, \gamma_{\max}]$ , where  $\gamma_{\min}$  is the minimum rate at which an organism relocates and  $\gamma_{\max}$  is the fastest rate an organism can relocate—we interpret this upper bound on relocation rates as being set by metabolic constraints that we do not model or allow to change during the evolution of the population.

The state of the population of agents is updated in continuous time. At any given moment, an agent can either relocate to or asexually give birth to an offspring into one of the vertical or horizontal neighboring locations, as schematically shown in Fig 1A (because only one agent is permitted per site, an agent cannot give birth to a child on its current site). The target site is chosen randomly with a probability  $1/4$  in a  $2d$  environment. The per-capita birth rate of an agent,  $b(s_{\mathbf{x}})$ , is a function of the local resource level  $s_{\mathbf{x}}$  at location  $\mathbf{x}$ . Because the resources are not depleted in this model the actual numerical values of the  $s_{\mathbf{x}}$  are unimportant; the meaningful values are the birth rates  $b(s_{\mathbf{x}})$ . However, we assume that the values of the birth rates are ordered according to resource level, such that a site  $\mathbf{x}$  with resource level  $s_{\mathbf{x}} = S_1$  has the lowest birth rate and a site with resource level  $s_{\mathbf{x}} = S_5$  has the highest birth rate. Again, we stress that the birth rates are entirely independent of an agent’s phenotype  $\gamma$ .

Agents with very different phenotypes will navigate the environment in different ways. An agent with a small value of  $\gamma_1$ , for example, will tend to remain in sites of resource level  $S_1$  for

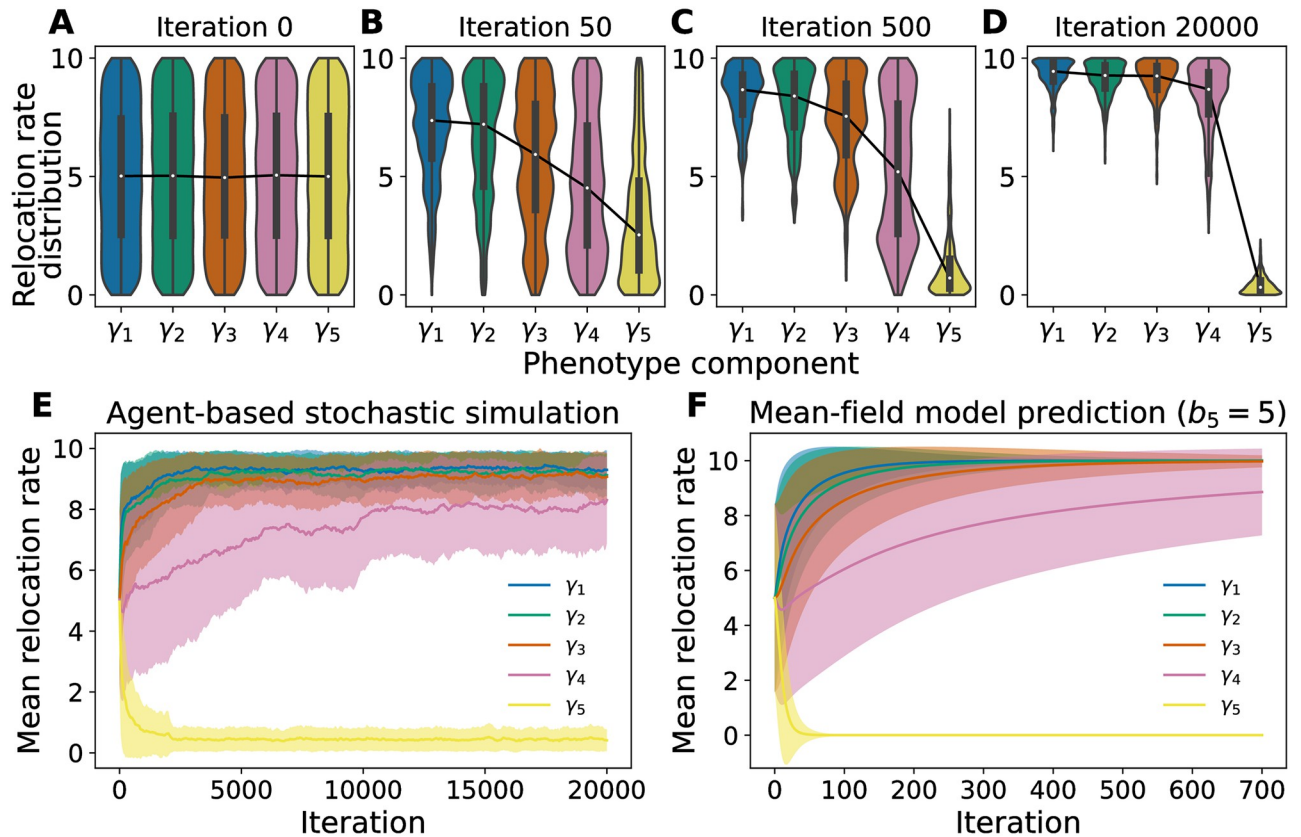
long periods of time, while agents with large values of  $\gamma_1$  will tend to leave sites of this resource level rather quickly. Because the birth rate in sites of resource level  $S_1$  is low, we might expect that agents with small  $\gamma_1$  may be less successful at proliferating than agents with larger  $\gamma_1$ , and will therefore be less likely to survive the periodic catastrophes that occur, creating a selection pressure toward larger values of  $\gamma_1$ . Similarly, agents with small (large) values of  $\gamma_5$  will tend to remain in these high-resource sites for long (short) periods of time. Agents with small  $\gamma_5$  may therefore produce more offspring because  $b(S_5)$  is the highest, increasing their odds of surviving the catastrophes. This line of thinking makes it seem intuitively “obvious” that the most competitive phenotype ought to have large hopping rates out of low resource sites and minimal hopping out of the highest resource sites. Moreover, we might expect we can identify an objective function that yields this phenotype as its optimum. In the next sections we present the results of our stochastic simulations as well as the numerical solutions of the mean-field model, showing that this intuition is largely correct when the population density is low—in which case we can indeed derive an effective objective function—and when the distribution of resource levels are drawn from a discrete uniform distribution, but this picture is not always correct when high resource sites are rare or population density is high.

### Distributions of competitive phenotypes are selected by birth-death dynamics

For the results presented in this section, we simulate our agent-based model using the Gillespie algorithm [23], starting with 500 agents on a  $128 \times 128$   $2d$  lattice. This corresponds to an initial population density of 3.1%. Each lattice site belongs to one of the 5 resource levels  $\{S_1, S_2, S_3, S_4, S_5\}$ , which is drawn from a discrete uniform distribution. Initially, each agent is assigned a spatial location  $\mathbf{x} = (x, y)$  and a phenotype vector  $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5)$  uniformly at random. Catastrophes occur after a simulation runtime  $T$ , which defines an iteration of the simulation. While the catastrophe randomly picks which agents to cull, it always reduces the total population back to 500 agents. In Distribution of resources shapes exploration-exploitation trade-offs we consider non-uniform distributions, and in Ranges of population density allow for long-lasting phenotype coexistence we investigate different lattice sizes as a means of changing the initial population density. The details of the simulation setups are presented in the [Methods](#) Section.

[Fig 2A–2D](#) displays the evolution of the distribution of phenotypes over several thousand catastrophes, with snapshots at (A) 0, (B) 50, (C) 500, and (D) 20000 iterations; these distributions represent averages over 20 trials of the dynamics starting from random initial conditions generated from the same distribution. The initially uniform distributions of each of the five hopping rates (phenotype components) are gradually reshaped by births and deaths. At long times, the distributions are approximately stable; we do not expect the distributions to converge to a single value for each  $\gamma_k$  due to the reproduction noise that accompanies each birth. See [S1A–S1C Fig](#) for additional examples for different realizations of the resource levels randomly assigned to each site. In [Fig 2E](#) we plot the evolution of the mean value of each hopping rate over the 20000 catastrophes, averaged over 20 trials. The shaded regions of the plot correspond to the width of the distributions (the standard deviation of each hopping rate component). These dynamics are qualitatively reproduced by a mean-field model, shown in [Fig 2F](#), and extensively discussed in [A minimal mean-field model explains variability in competitive phenotypes](#).

We see that the long-time distribution appears to favor high hopping rates in sites with resource levels  $S_1$  to  $S_3$  and low hopping rates for sites with the highest resource level  $S_5$ , reflecting our intuition that the most competitive phenotypes would be those that result in



**Fig 2. Emergent competitive phenotypes from evolutionary birth-death dynamics.** A-D. Violin plots of distributions of phenotype components in the agent-based stochastic simulations at iterations 0, 50, 500, and 20000, respectively. Each violin plot combines data from 20 independent trials. The connected white dots are the median values of each relocation rate component, and the black boxes confine the first and third quartiles of each phenotype component (i.e., box plot). The colored areas are the kernel density estimations of the underlying distributions. Panel A shows a uniform relocation rate distribution for each phenotype component as the initial condition at iteration 0. The phenotype distributions evolve as the birth-death dynamics play out in the agent-based model and approach a steady shape after around 20000 iterations as shown in panel D. E. The evolutionary trajectory of each phenotype component, averaged over 20 independent trials. The shaded area represents the standard deviations of phenotype components. The stochastic birth-death dynamics selects out competitive phenotypes whose first three phenotype components  $\gamma_{1,2,3}$  saturate at hopping rate of around 9, close to the maximum possible hopping rate of  $\gamma_{\max} = 10$ , and whose last phenotype component  $\gamma_5$  saturates at hopping rate around 0.5, close to the minimum possible hopping rate of  $\gamma_{\min} = 0$ . In contrast, the fourth component  $\gamma_4$  approaches a value around 8 with a relatively large variability throughout the simulation. F. Predictions of the mean relocation rate of each phenotype component and their standard deviations using the coupled 5-state mean-field model, discussed in A minimal mean-field model explains variability in competitive phenotypes. Note that the mean-field model uses a slightly different birth rate  $b_5 = 5$ , while the stochastic simulations use  $b_5 = 4$ , and the 5-state mean-field model converges at a much shorter timescale (700 iterations vs. 20000 iterations). These quantitative differences between the full agent-based model and the simplified 5-state model are discussed in detail in the Discussion, S5 and S6 Figs.

<https://doi.org/10.1371/journal.pcbi.1009934.g002>

agents leaving their location as fast as possible when at low resource sites and staying put as long as possible when at high resource sites. Notably, however, the distribution of the hopping rates  $\gamma_4$  at the second-highest resource level  $S_4$  is not narrow, but spans much of the entire possible range of hopping rate values. This indicates that the selection pressure on this phenotype component is relatively weak compared to the selection pressure on the other components. While the average value of  $\gamma_4$  is closer to  $\gamma_{\max}$  than  $\gamma_{\min}$ , the large variance demonstrates that the most competitive phenotype does not conform tightly to our intuitive predictions, and the situation is more complex than might be naively expected. To gain insight into our results and investigate how selection pressure depends on our environments' properties, namely the distribution of resources and the values of the birth rates tied to each resource level, we develop a deterministic minimal mean-field model that enables us to derive an effective objective

function in a low-density limit. This fitness function predicts the optimal phenotypes and explains the lack of selection pressure on  $\gamma_4$  observed in our stochastic simulations.

### A minimal mean-field model explains variability in competitive phenotypes

Agent-based models are often difficult to formulate directly as mathematical models, owing to the combination of having individual with private properties (here, their phenotype) and the fact that agents may be born or culled at any time. However, a formalism that can approximate the agent-based model outlined here is a master equation formalism that keeps track of the probability of configurations of the number of agents with each phenotype at each spatial location. From such a formalism one can then derive a system of differential equations for the expected number of agents with each phenotype and location. Such a model serves as the starting point for the reduced model we investigate in this section. We focus only on the reduced mean-field model here; see the [Methods](#) for details on the full master equation model and the approximate reduction to the minimal mean-field model described below.

When agent density is low, such as in the simulations of the previous section (initial density of 3.1%), we may assume that events in which an agent is blocked from relocating to another site by another occupying agent are rare and may be neglected. In this approximation the growth dynamics of the different phenotypes decouple, and we can describe the dynamics of the population as a set of linear differential equations describing the expected number of agents at each site. To simplify the mean-field model further, we use the fact that each site belongs to one of the five resource levels and only count the number of agents at each level, eliminating the spatial dependence of the environment. This reduces the model to a 5-state model whose solution is the time course of the expected number of agents of each phenotype at each resource level; see the [Methods](#) for a detailed derivation of this reduction. The resulting system of differential equations reads

$$\frac{dn_{k\gamma}}{dt} = w_k \sum_i b_i n_{i\gamma} + w_k \sum_i \gamma_i n_{i\gamma} - \gamma_k n_{k\gamma} \quad (1)$$

where  $n_{k\gamma}$  is the expected number of agents at a site with resource level  $k$  and phenotype  $\gamma$ . The parameter  $b_k$  is the per-capita birth rate of resource level  $k$ , a constant of the environment, and  $\gamma_k$  is the hopping rate out of state  $k$ , determined by the phenotype. Again, we note that in this approximation the dynamics of different phenotypes are decoupled, such that there is a system of differential equations of the form [Eq \(1\)](#) for each phenotype  $\gamma$ . For practical calculations we must discretize the phenotype space, giving a finite system of differential equations. Finally, the weights  $w_k$  correspond to the probability that a site in the agent-based model has resource level  $S_k$ . For resource levels drawn from a discrete uniform distribution, which we investigate first,  $w_k = 1/5$  for all  $k = 1$  to 5.

The first term,  $w_k \sum_i b_i n_{i\gamma}$ , in [Eq \(1\)](#) is the number of new agents being birthed into state  $k$  by agents in all other states. The second term  $w_k \sum_i \gamma_i n_{i\gamma}$  is the number of agents hopping into resource level  $k$  from all other resource levels, formally including resource level  $k$  itself to account for the possibility of an agent hopping from a site of resource level  $k$  to a neighboring site that also has resource level  $k$ . The last term  $-\gamma_k n_{k\gamma}$  is the number of agents hopping out of resource level  $k$ .

The system of differential [Eq \(1\)](#) can be written down in matrix form as  $\frac{d\mathbf{n}_\gamma(t)}{dt} = \mathbb{M}_\gamma \mathbf{n}_\gamma(t)$ , which has a solution in terms of the matrix exponential,  $\mathbf{n}_\gamma(t) = \exp(t\mathbb{M}_\gamma) \mathbf{n}_\gamma(0)$ . The between-catastrophe dynamics of each phenotype are therefore dominated by the eigenvectors and



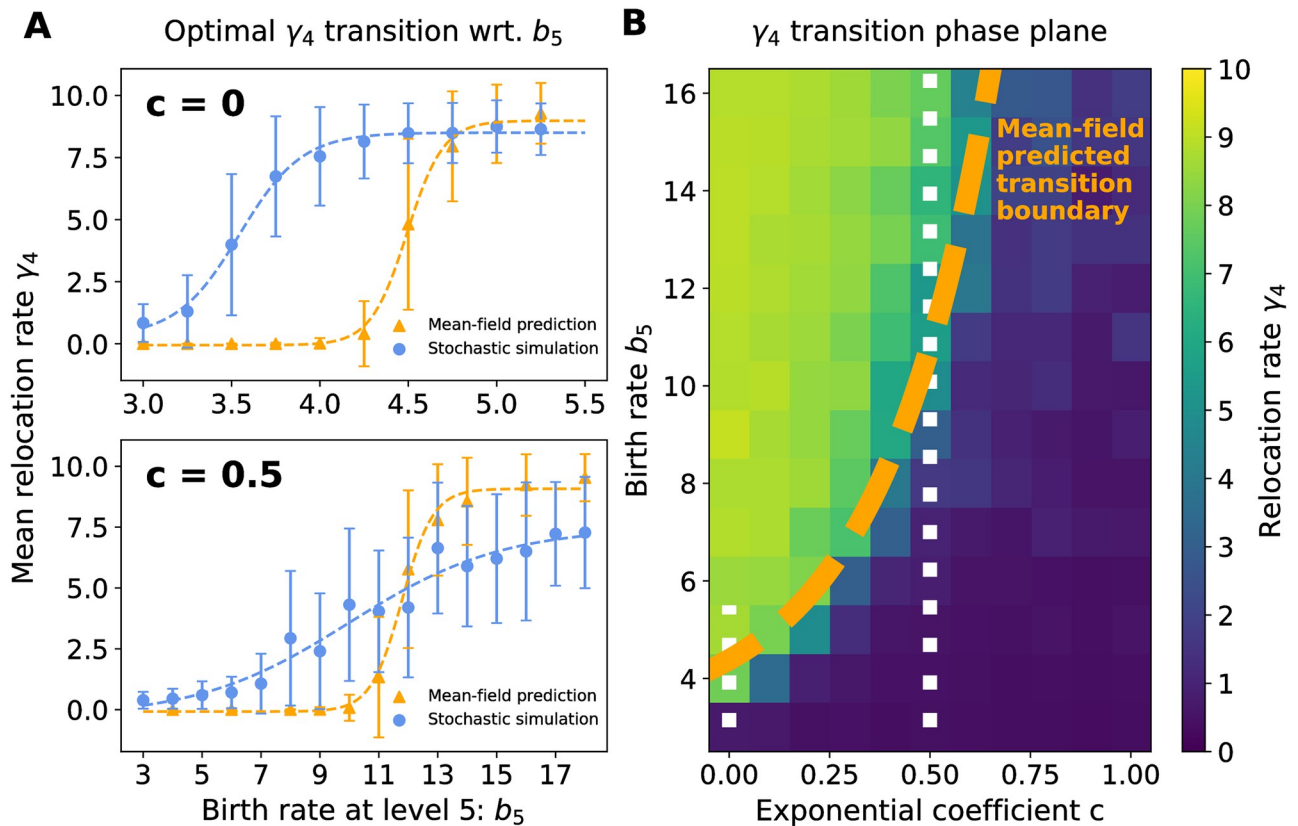
eigenvalues of  $\mathbb{M}_\gamma$ . Although the dynamics of the phenotypes are independent, they are coupled through the periodic catastrophes, which reduce the population down to its original size after a time  $T$ . Let  $\Lambda_\gamma$  be the maximum eigenvalue of  $\mathbb{M}_\gamma$ , which we interpret as a function of the phenotype parameters  $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5)$ . In the [Methods](#) we show that the matrix  $\mathbb{M}_\gamma$  has purely real eigenvalues and only the maximum eigenvalue  $\Lambda_\gamma$  is positive. Therefore, the phenotype  $\gamma$  that yields with the largest maximum eigenvalue  $\Lambda_\gamma$  overall will be the only phenotype to survive after a large number of catastrophes, with coexistence possible only if different phenotypes have the same value of  $\Lambda_\gamma$ . Because the hopping rates  $\gamma$  are just parameters in this formulation of the model, the eigenvalues of  $\mathbb{M}_\gamma$  will simply be functions of the phenotype parameters (as well as environmental parameters), and therefore the maximal eigenvalue serves as an *effective objective function* of the system in the low-density limit.

If the maximal eigenvalue of this minimal model acts like an effective objective function, then we would expect the lack of selection pressure on the hopping rate  $\gamma_4$  to be reflected in an insensitivity of this eigenvalue to varying  $\gamma_4$ , i.e., there would be a range of values of  $\gamma_4$  for which the maximum eigenvalue approximately takes on its largest value. This is indeed what we find *but only when the birth rate in the highest resource site  $b_5$  is close to a particular critical value  $B_{5c}$* . As shown in [S2 Fig](#), when  $b_5 > B_{5c}$  we find that the competitive population dynamics drive  $\gamma_4 \rightarrow \gamma_{\max}$ , and the eigenvalue  $\Lambda_\gamma$  is insensitive to variations in the value of  $\gamma_4$  around this point. When  $b_5 < B_{5c}$  and is close to  $b_4$ , the birth rate in sites of resource level  $S_4$ , we instead find that selection drives  $\gamma_4 \rightarrow \gamma_{\min}$ , again with little variation in the maximal eigenvalue. This makes sense: if the birth rates  $b_4$  and  $b_5$  are close then there is little advantage in leaving a site with resource level  $S_4$  to search for a site of level  $S_5$ . That is, the maximal eigenvalue predicts a transition from  $\gamma_4 \rightarrow \gamma_{\min}$  to  $\gamma_4 \rightarrow \gamma_{\max}$  as  $b_5$  increases from  $b_4$  and crosses  $B_{5c}$ . Near the transition at  $b_5 \approx B_{5c}$  the maximum of the eigenvalue is less sensitive to  $\gamma_4$ —equivalently, fluctuations in  $\gamma_4$  are large, as is often observed near phase transitions and bifurcations [24], such that selection is weak, resulting in coexistence of a range of  $\gamma_4$  values near this transition, at least in a finite population. See [S3A–S3C Fig](#) for the distributions of phenotypes around this transition point.

[Fig 3A](#), upper panel, demonstrates that the predicted transition occurs in both our minimal mean-field model and our full agent-based simulations. The critical value  $B_{5c}$  is different for the full stochastic simulation and the mean-field model, but this is to be expected due to the simplifications of the mean-field model as well as the lack of demographic noise. As seen in [Fig 3A](#), upper panel, the distributions of  $\gamma_4$  are indeed narrower when  $b_5$  is far from the critical value  $B_{5c}$ , reflecting that the selection pressure indeed increases when there is a clear advantage to exploring ( $b_5 > B_{5c}$ ) versus exploiting ( $b_5 < B_{5c}$ ). In principle, we might expect to see similar transitions in the other hopping rates  $\gamma_1$ ,  $\gamma_2$ , and  $\gamma_3$  if we adjust the value of  $b_4$ ,  $b_3$ , etc. Instead of varying these parameters to investigate whether such transitions exist, we shift focus to another ecologically important scenario in which these transitions will occur: the effect of the distribution of resources throughout the environment.

### Distribution of resources shapes exploration-exploitation trade-offs

The distribution of resources throughout an environment will have an impact on the most successful phenotypes. For instance, when high resource sites are scarce, it may be advantageous to settle for lower resource locations. To explore this effect, we vary the distribution of resource levels within the random environment, in particular by selecting distributions in which the high-resource sites are much rarer than lower resources sites. In the previous sections, the probability that any site has resource level  $S_k$ , denoted  $w_k$ , was chosen to be  $w_k = 1/5$ ; i.e., a discrete uniform distribution of the 5 resource levels. A simple choice of distribution in which



**Fig 3. The most competitive phenotypes are shaped by the rarity of the highest-resource sites.** The abundance of the five resource levels distributed with frequency  $w_k \propto e^{-ck}$ , where  $k \in \{1, 2, 3, 4, 5\}$  and the exponential coefficient  $c$  determines the rarity of levels. The most competitive phenotype changes as a function of birth rates and resource rarity: **A, top**. For resources drawn from discrete uniform distribution ( $c = 0$ ), the fourth component of the most competitive phenotype  $\gamma_4$  transitions from low to high as the birth rate  $b_5$  increases. All other phenotype components remain steady as  $\gamma_{1,2,3} \rightarrow \gamma_{\max}$  and  $\gamma_5 \rightarrow \gamma_{\max}$ . The agent-based stochastic simulation and the mean-field prediction transitions are qualitatively similar, though the transition occurs at different values of  $b_5$ . **A, bottom**. A similar transition in  $\gamma_4$  occurs for a resource distribution with exponential coefficient  $c = 0.5$ , which has rare high resource level sites. The transition persists but at a higher birth rate  $b_5$ . **B**. A heatmap plot of the fourth phenotype component  $\gamma_4$  as a function of the exponential coefficient  $c$  and the birth rate  $b_5$ . The orange dashed line is the predicted transition boundary from the mean-field model, where  $\gamma_4 = 5$  maximizes the effective fitness function derived in the low-density limit. The two white dotted lines show the range of  $b_5$  for  $c = 0$  and  $0.5$  in panel A.

<https://doi.org/10.1371/journal.pcbi.1009934.g003>

high resources sites become rarer than low resource sites is an exponential probability  $w_k = Ae^{-ck}$ , where  $A^{-1} = \sum_{k=1}^5 e^{-ck}$  is the normalization constant and  $c$  is a parameter that determines the steepness of the distribution. The factor  $w_k$  is the same factor that appears in Eq (1). Note that  $c = 0$  results in  $w_k = 1/5$ ; a positive  $c$  results in a distribution that resource level rarity increases with  $k$ , that is, the high resource levels become rarer than the low resource levels.

As shown in Fig 3A, lower panel, our agent-based simulations confirm that the transition from an average value of  $\langle \gamma_4 \rangle = \gamma_{\min} = 0$  to  $\gamma_{\max} = 10$  still exists when  $c$  is, e.g.,  $0.5$ , but occurs at a much larger value of  $b_5 \approx 11$ . Indeed, as seen in Fig 3B, there exists a transition line in the  $b_5$ - $c$  plane that separates  $\langle \gamma_4 \rangle \approx \gamma_{\min}$  from  $\gamma_{\max} = 10$ . The orange dashed curve is the boundary predicted by evaluating the maximal eigenvalue of the low-density model  $\Lambda_\gamma$  as  $c$  varies from 0 to 1 and solving for the value of  $b_5$  that yields  $\gamma_4 = (\gamma_{\max} + \gamma_{\min})/2 = 5$  as the most competitive phenotype. This curve represents the line along which the population achieves a balance of the exploration-exploitation costs. See S1D–S1F Fig for examples of the phenotype distributions in different realizations of environments with exponentially distributed resource levels and S3D–S3F Fig for the phenotype distributions around the transition point.

Intuitively, agents staying at low resource levels need sufficient incentives to leave their current locations and start to explore new sites, at the risk of missing out reproductive opportunities while exploring. These results demonstrate that rarity of high resource sites can drive selection of phenotypes that settle for sites with lower resource levels.

Importantly, here the rarity of sites is only due to the intrinsic environmental distribution of resources. However, high resource sites could also become *dynamically* rare and unavailable as they become occupied, preventing multiple individuals from exploiting the same sites. So far we have been working within a low-density regime, such that it is relatively rare for organisms' movements to be blocked by other individuals, and such that a sufficient fraction of high resource sites remain available for discovery and occupation. We now turn to investigating how the competitive phenotypes change with organism density.

### Ranges of population density allow for long-lasting phenotype coexistence

In order to investigate how population density impacts the most competitive phenotype changes, we return to our agent-based stochastic simulations with resource levels drawn from a discrete uniform distribution. The simulations presented previously had 500 agents on a  $128 \times 128$  2d lattice, resulting in an initial occupancy level of 3.1%. To study the effect of population density we keep the initial number of agents fixed and decrease the lattice size  $L \times L$  to  $L = 96, 80, 64, 48, 40, 35, 32, 28, 26, 25, 24,$  and  $23,$  giving initial occupancies of 5.4%, 7.8%, 12.2%, 21.7%, 31.2%, 40.8%, 48.8%, 63.8%, 74.0%, 80.0%, 86.8%, and 94.5%, respectively. One could of course also vary density by increasing the initial agent number and keeping the lattice size at  $128 \times 128,$  but increasing the number of agents dramatically increases the simulation load.

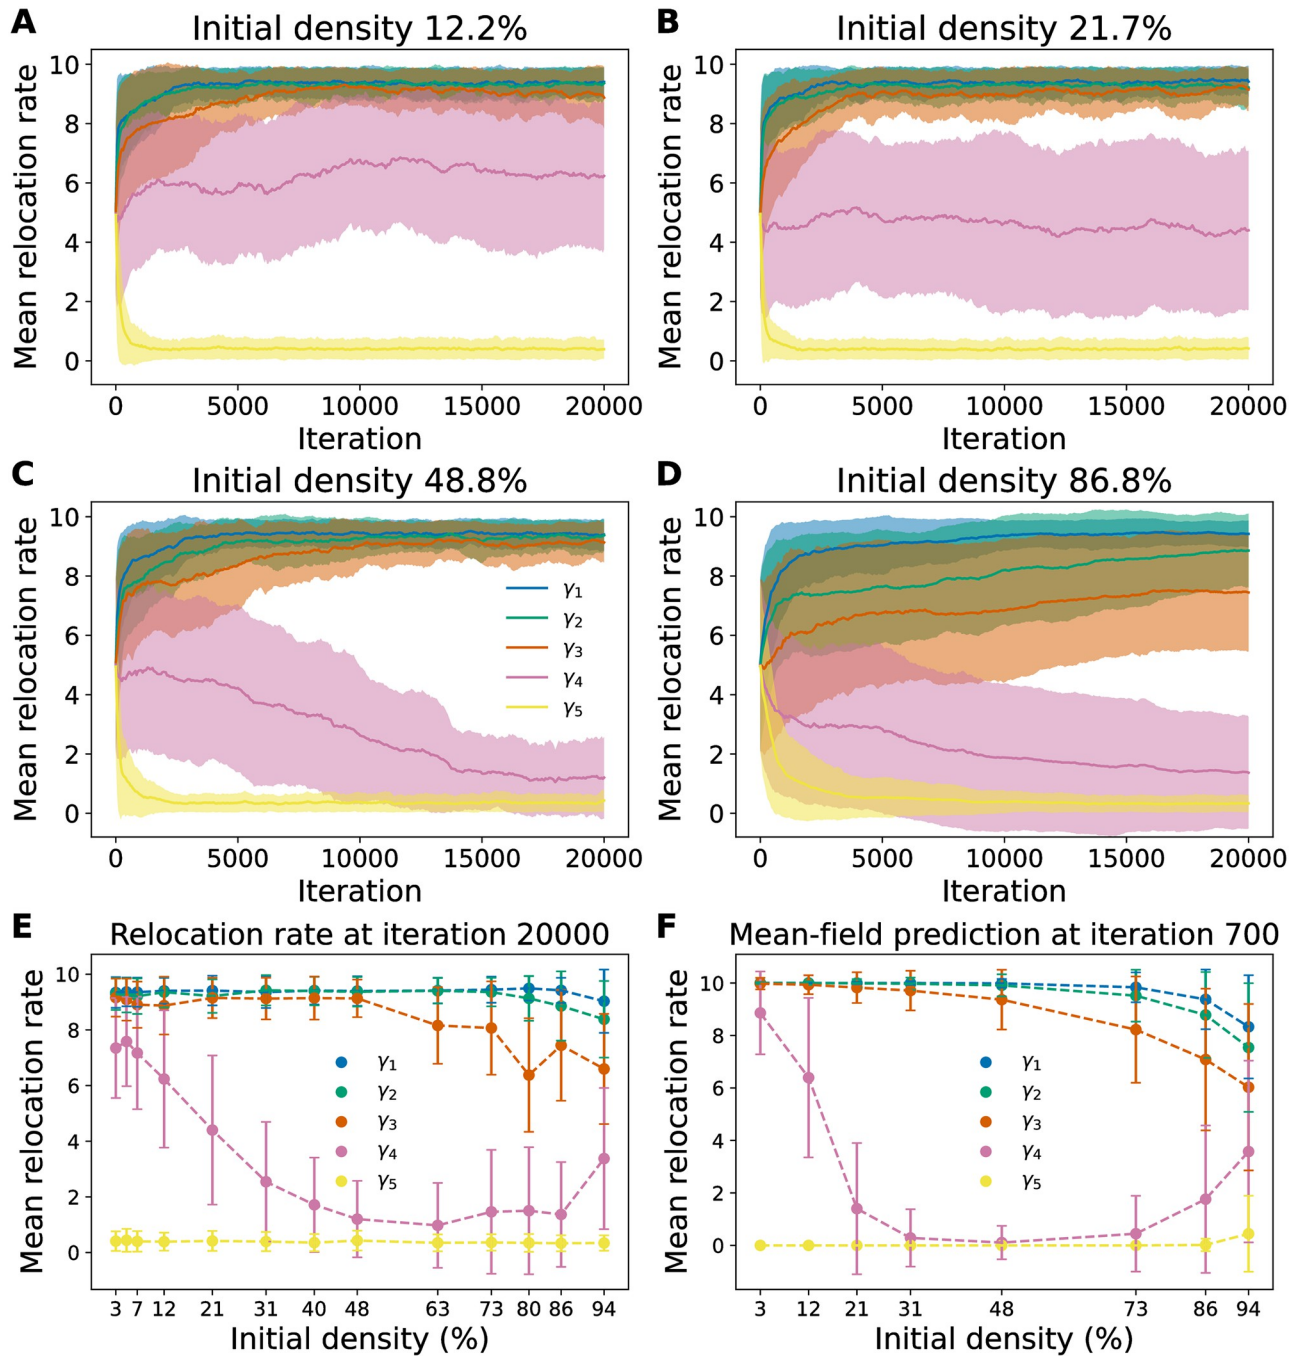
For each density we ran 20 independent trials of the agent-based stochastic simulation, the results of which are summarized in Fig 4. We found that as density—and hence interactions among agents—increases, the most competitive phenotypes also change. Notably, these changes qualitatively mimic the effects of increasing the rarity of high resource sites: the mean hopping rate  $\langle \gamma_4 \rangle$  saturates at a value  $\gamma_{\min} < \langle \gamma_4 \rangle \approx 5 < \gamma_{\max}$  with a large variance at 12.2% initial density, and drops to  $\langle \gamma_4 \rangle \approx \gamma_{\min}$  with small variance at higher densities. See S4 Fig for examples of the phenotype distributions at increased densities, which show this transition take place.

The results observed in our stochastic simulations are replicated by a 5-state mean-field model that incorporates the effects of site occupation:

$$\frac{dn_{k\gamma}}{dt} = w_k \left( 1 - \frac{\sum_{\gamma'} n_{k\gamma'}}{w_k M} \right) \sum_i (b_i + \gamma_i) n_{i\gamma} - \sum_i w_i \left( 1 - \frac{\sum_{\gamma'} n_{i\gamma'}}{w_i M} \right) \gamma_k n_{k\gamma}, \tag{2}$$

where  $n_{k\gamma}$  is again the number of agents at sites of resource level  $k$  and  $M = L^2$  is the total number of sites, with  $w_k M$  the expected number of sites of resource level  $k,$  which sets a maximum occupancy for each state. Note that when  $\frac{\sum_{\gamma'} n_{k\gamma'}}{w_k M} \ll 1$  this model reduces to Eq (1). Importantly, Eq (2) is nonlinear, and unlike the low-density approximation of Eq (2), all phenotypes are dynamically coupled. Not only does this preclude writing down an analytical solution, but the dynamical coupling between phenotypes also suggests there may be no effective objective function of a single arbitrary set of the phenotype parameters  $\gamma = (\gamma_1, \dots, \gamma_5)$  that predicts the most competitive phenotype simply by maximizing it with respect to these values.

As in the low-density model, we solve Eq (2) for a discretized set of phenotypes. Solving the high density model is comparatively computationally demanding due to the necessity of solving for all phenotypes simultaneously, but is still feasible for the set of discretized phenotypes

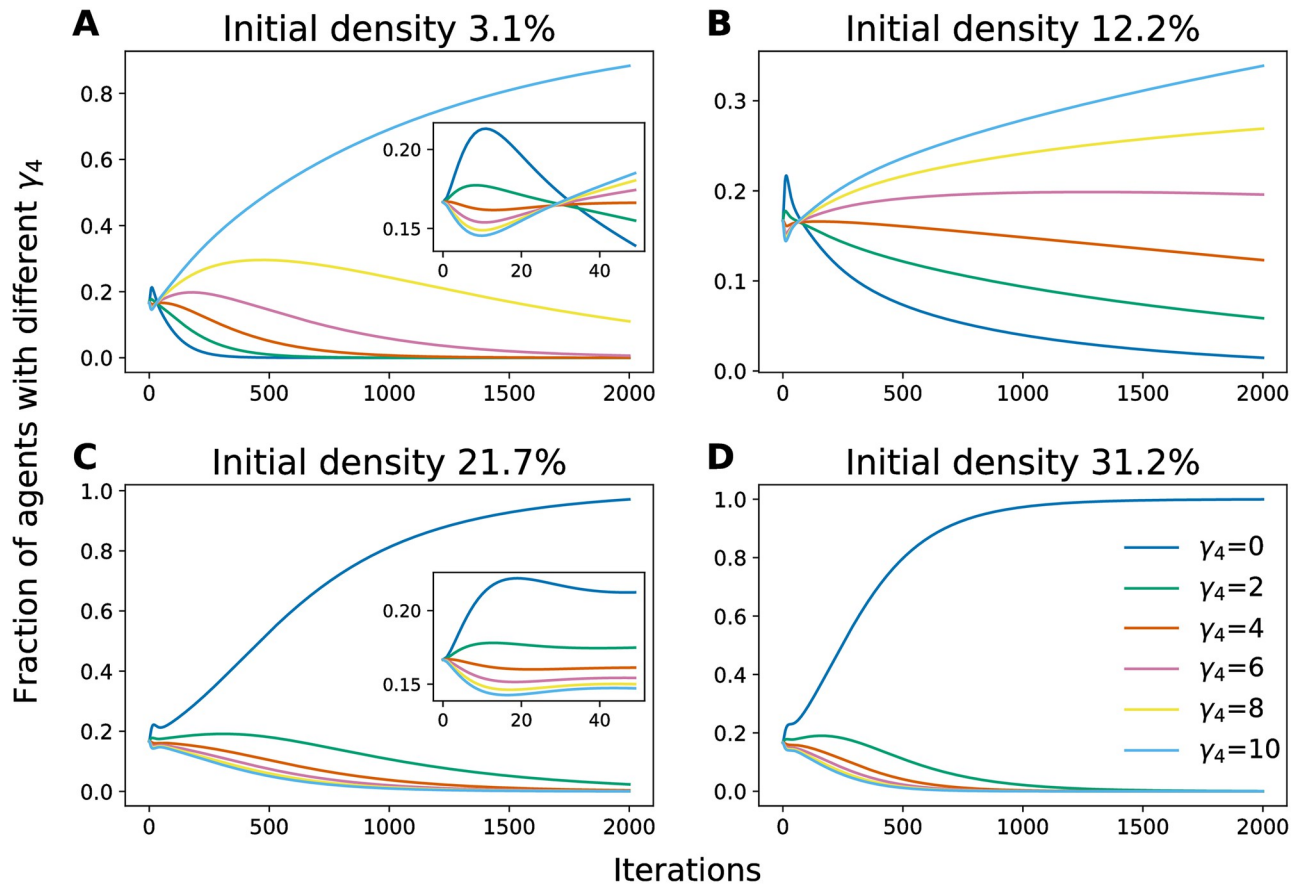


**Fig 4. Higher initial densities alter the most competitive phenotypes.** A–D. The evolution course for populations of agents starting with different initial densities: 12.2%, 21.7%, 48.8%, and 86.8%, respectively; c.f. Fig 2E at 3.1% density. The initial density is varied by fixing the initial population size at 500 and decreasing the environment size  $L \times L$  with  $L = 128, L = 64, 48, 32,$  and  $24$ . We observe a systematic shift of the most competitive phenotypes; in particular, the fourth phenotype component  $\gamma_4$  decreases as the initial density increases. E. Mean values of  $\gamma_4$  after 20000 iterations in the agent-based model as a function of the initial density, showing the change in the most competitive phenotype. At ultrahigh densities, e.g. 94%, there is no space for agents to move or give birth, resulting in little change in the initially uniform phenotype distribution. F. Mean-field model predictions of the most competitive phenotypes as a function of initial density after 700 iterations. A birth rate of  $b_5 = 5$  is used in the mean-field model so that  $\langle \gamma_4(t) \rangle$  converges to 10 at low densities.

<https://doi.org/10.1371/journal.pcbi.1009934.g004>

chosen in this study. As shown in Fig 4F, the mean-field predictions are consistent with our full stochastic simulations, and we indeed observe changes in the most competitive phenotype as a function of initial density  $N_0/M$ . At low densities the value of  $\langle \gamma_4 \rangle$  is close to  $\gamma_{\max}$  (after 700 catastrophes), and as  $N_0/M$  increases we see a coexistence region in which multiple phenotypes are present at an observation time  $T_{\text{obs}}$ , before  $\langle \gamma_4 \rangle$  drops to  $\gamma_{\min}$  at densities  $N_0/M \gtrsim 20\%$ . As  $N_0/M$  increases towards 1, we observe an increase in  $\langle \gamma_4 \rangle$ . This extreme high-density limit represents the scenario in which the population is so dense that there is little room for agents to give birth or move, leading to an extremely slow change in the composition of the phenotypes. In the extreme limit  $N_0/M = 1$  we have  $\langle \gamma_4(t) \rangle = (\gamma_{\max} + \gamma_{\min})/2$  for all times  $t$ , due to the fact that phenotypes are initially present in equal proportions and there is no room for any agent to move or give birth.

The mean-field model is useful for dissecting the dynamics of the phenotype changes. Initially, we hypothesized that the dominance of  $\gamma_4 = \gamma_{\min}$  phenotypes at high densities is due to the  $S_5$  resource sites filling up, leaving any  $\gamma_4 = \gamma_{\max}$  agents that did not find an  $S_5$  site to search fruitlessly for resource-level 5 sites instead of producing offspring, allowing the  $\gamma_4 = \gamma_{\min}$  phenotypes to gain the reproductive advantage. However, as shown in Fig 5, it turns out that early



**Fig 5. Numerical solutions of the coupled mean-field model at different initial densities.** A–D show the evolution of the fraction of the population with each value of  $\gamma_4$  in the mean-field model; insets at 3.1% and 21.7% initial densities show a zooming in of the initial transient dynamics from iteration 0 to 50. Initially in the evolution,  $\gamma_4 = \gamma_{\max} = 10$  agents show a reproductive disadvantage, intuitively because agents with larger  $\gamma_4$  tend to leave  $S_4$  resource levels (with a relatively high birth rate) more often and produce less offspring, while later in the evolutionary process this disadvantage is overcome by  $\gamma_4 = \gamma_{\max}$  agents in low densities as they eventually find enough available  $S_5$  sites to overcome their initial disadvantage. However, as density increases, other agents can quickly fill up  $S_5$  sites and leave little opportunities for  $\gamma_4 = \gamma_{\max}$  agents to overcome that initial disadvantage. Note that  $b_5 = 5$  is used in these numerical solutions so that  $\langle \gamma_4(t) \rangle$  converges to 10 at low densities.

<https://doi.org/10.1371/journal.pcbi.1009934.g005>

in the evolutionary dynamics the  $\gamma_4 = \gamma_{\max}$  agents lag behind the other phenotypes, as the  $\gamma_4 = \gamma_{\max}$  agents leave  $S_4$  sites more frequently and explore the neighboring locations while agents with lower  $\gamma_4$  gain an advantage by tending to stay at  $S_4$  longer and produce offspring at a relatively high birth rate  $b_4$ . At low densities the  $\gamma_4 = \gamma_{\max}$  agents eventually discover  $S_5$  sites and their reproduction rate is able to overcome their slow reproductive start. However, at higher densities the  $\gamma_4 = \gamma_{\max}$  agents are increasingly less able to overcome the other phenotypes' head start, and are eventually outcompeted completely by the  $\gamma_4 = \gamma_{\min}$  phenotypes at sufficiently high densities.

Unlike the low-density model, we cannot tractably derive an effective objective function similar to  $\Lambda_\gamma$  in the high-density regime—in particular, it is not clear whether there exists a function  $C(\gamma)$  of a single phenotype vector  $\gamma$  whose optima predict the most competitive phenotypes. We discuss this issue further in Evolution as an optimization or learning process. Even if no such cost function for the phenotypes exists, one might wonder whether the dynamics for the (expected) population sizes can be written as a dynamical system that follows the gradient of some cost function  $C(\mathbf{n})$ ,  $\frac{dn_{k,\gamma}}{dt} = -\frac{\partial}{\partial n_{k,\gamma}} C(\mathbf{n})$ . In evolutionary models, the negative of the cost function  $-C(\mathbf{n})$  is often interpreted as fitness. However, the vector field defined by the right hand side of Eq (2) is not conservative, and therefore cannot be written as the gradient of an objective function  $C(\mathbf{n})$  (see The mean-field dynamics are not gradient-based in Methods). The evolution of the population sizes therefore cannot be interpreted as optimization problem, at least using gradient dynamics.

## Discussion

### Summary and interpretation of results

In this work we have used a population dynamics approach to explore the evolution of the innate foraging behaviors of a population of organisms exploring an environment with randomly distributed non-consumable resources. These organisms have limited sensory capabilities and no memory, and must take actions to either continue exploiting the resources at their current location or leave in search of locations with more resources (i.e., more favorable to reproduction). The phenotype of the organisms comprises the rates at which the agents attempt to relocate. While an “intuitive” prediction—the agents should leave low-resource sites as quickly as possible and not move from the highest resources sites—is indeed correct when the population density is low and high resource sites are readily available or have high enough birth rates to compensate their scarcity, this intuitive prediction becomes incorrect when the difficulty of finding high resource sites increases due to intrinsic scarcity or occupation by other organisms in high-density populations. In particular, as the initial population density varies, we not only observe a change in the dominant phenotype at long times, but we also observe long-lived periods of phenotype coexistence.

We have used a combination of stochastic agent-based simulations and a simplified mean-field model to gain insight into the basic phenomenology of the evolutionary dynamics. In the low-density limit we derived an effective fitness function that predicts the dominant phenotype at long times, but could not identify such a function that predicts the dominant phenotypes in higher density environments. Similarly, the resulting dynamics of the average population sizes are not interpretable as gradient-based learning dynamics, as might be obtained from a normative approach in which the organism phenotypes are predicted by optimizing a hypothetical objective function. Thus, there is no guarantee that one can in general derive a function of the phenotype parameters that predicts the optimal phenotype—or distribution of phenotypes—without running the dynamics of the model.

We will discuss some of these points in detail in the following subsections, along with detailed remarks on Comparisons to related work and Limitations of the study, and we close with a brief discussion of Future directions.

**Comparison of the agent-based and mean-field models.** Overall, the outcomes of the full stochastic agent-based model and the minimal 5-state mean-field model we used to gain further insight were in qualitative agreement. This is often the case in statistical physics models; for example, mean-field approximations of Ising models of magnetism correctly predict the qualitative features of the phase diagram of the model, but overestimate the critical temperature at which a transition from zero to non-zero net magnetization of the material occurs [24], similar to the overestimate of the critical birth rate for the transition shown in Fig 3. This discrepancy may also be due to the well-mixed assumption made during the reduction of the full  $2d$  model to the 5-state model, discussed in the next paragraph. Despite some quantitative discrepancies in predicting the values of the critical birth rates  $B_{5,c}(c)$ , the low-density mean-field model accurately predicts the transition boundary from  $\langle \gamma_4 \rangle \approx 10$  to  $\langle \gamma_4 \rangle \approx 0$  as a function of the resource scarcity  $c$ , as shown in Fig 3.

The primary source of quantitative discrepancies between stochastic models and mean-field models is the neglect of fluctuations and the necessity of choosing a carrying capacity that takes continuous inputs (which gives rise to the factor  $1 - \sum_{\gamma} n_{k\gamma}/(w_k M)$  in Eq (2)). In many mean-field approximations of master equation models, upon which our 5-state model is based, the variance of these fluctuations eventually invalidate the mean-field approximation at long times or when the mean population sizes are small. As discussed in Limitations of the study, we avoid these causes of mean-field breakdown by using a periodic catastrophe to reduce the population size to a fixed level, rather than stochastic death. Additional discrepancies are caused by the reduction from the mean-field approximation for the full  $2d$  environment to the 5-state model (see Numerical solution and analysis of the 5-state mean-field model for details). Some of the discrepancy in Fig 3 may also reflect the effects of population density, particularly in our simulations with scarce resources, as the transition occurs at relatively large values of  $b_5$  (compared to the discrete uniform distribution of resource levels), which could lead to the population growing large enough before each catastrophe that interactions between individuals are not negligible, reducing the validity of the low-density model used to predict the transition boundary in Fig 3B.

The largest qualitative discrepancy between the stochastic simulations and the mean-field model is the timescale of the evolutionary dynamics. Saturation of the hopping rates in the stochastic model appears to take  $\sim 30$  times longer than in the mean-field model, despite the birth rates, hopping rates, and catastrophe period parameters being nearly the same across the models. We believe this discrepancy is due to the well-mixing assumption of the mean-field model, which posits that the spatial variability of agents throughout the environment is negligible, allowing our reduction to the 5-state model; we quantify the magnitude of this approximation in S5 Fig. This conjecture is supported by direct simulations of a 5-state agent-based model, which shows a convergence time similar to the mean-field model predictions (see S6 Fig). We therefore expect that in a mean-field model that takes into account spatial variation of agents and resources we would get convergence times more similar to the full  $2d$  agent-based simulations.

**Metastability of phenotype coexistence.** In both our stochastic simulations and mean-field model we observed long-lived coexistence of different phenotypes (Figs 4 and 5). In the low-density limit of the mean-field model we could show analytically that this coexistence is almost-surely transient: only the phenotypes with the largest eigenvalues  $\Lambda_{\gamma}$  will survive in the limit of infinitely many catastrophes,  $\ell T \rightarrow \infty$ . The exceptions in which coexistence is possible in this long time limit are fine-tuned conditions under which multiple phenotypes have the

exact same values of  $\Lambda_\gamma$ . Therefore, outside of this special fine-tuned condition, at low-density coexistence is *metastable* in the sense that phenotypes may coexist for long durations of time but the ultimate fixed point is survival of only a single phenotype (we remind the reader that extinction is not possible in our setup; in practice the fate of any model with a finite number of agents and stochastic death is extinction [25]). In the stochastic agent-based model the reproduction noise prevents any one phenotype from being the only survivor, but the distribution of phenotypes will be narrow around a dominant phenotype if it exists. In contrast, long-lived coexistence is represented by a wide distribution of phenotype that persists across a large number of catastrophes. In both the mean-field model and the agent-based simulations the time-scale of this phenotype coexistence depends strongly on the value of  $b_5$ : when  $b_5 \approx B_{5c}$  the distribution of  $\gamma_4$  remains wide for a long time, while the distribution quickly converges when  $b_5 \ll B_{5c}$  or  $b_5 \gg B_{5c}$ . This long-lived phenotype coexistence is again somewhat fine-tuned, requiring  $b_5$  to be close to  $B_{5c}$ .

In the high density regime (Eq (2)), on the other hand, we observed a range of densities for which multiple phenotypes coexist for long times and do not show clear signs of a trend towards a single dominant phenotype. This form of phenotype coexistence appears more robust—requiring less fine-tuning of the initial density—than the coexistence observed in the low-density regime when the birth rate  $b_5$  is close to the critical value  $B_{5c}$ . However, it is difficult, if not analytically intractable, to determine whether the phenotype distribution would ultimately converge to a single dominant phenotype if run for a much longer time, or if there is a broader range of parameter space for which coexistence occurs in the infinite time limit. Inspection of the average hopping rate dynamics  $\langle \gamma_k(t) \rangle$  in the mean-field model (Fig 4F) seems to suggest there will be fine-tuned values of the initial density  $N_0/M$  at which two or more phenotypes may coexist in the long-time limit, but deviations from these values result in a single surviving phenotype in the long time limit, rendering coexistence of phenotypes in the high-density regime similar to low-density coexistence near  $b_5 \approx B_{5c}$ . However, because the mean-field model does not include reproduction noise, it is possible that the reproduction noise could stabilize coexistence states, preventing any one phenotype from dominating at long enough times in the full stochastic model—the trends of  $\langle \gamma_4(t) \rangle$  in Fig 4A and 4B could be reflecting such a behavior.

**Evolution as an optimization or learning process.** In this work we investigated the evolution of innate foraging behaviors using evolutionary population dynamics [4], in contrast to the more common normative approaches in which evolution is assumed to act like an optimization process that extremizes some objective function subject to constraints [6, 26, 27]. One of our original motivations for doing so was to better understand conditions under which we might expect such normative approaches to be mimicked naturally by the population dynamics, either through the population dynamics being interpretable as gradient dynamics or by deriving an effective “fitness function” that predicts the dominant phenotype, or conditions under which typical normative approaches might need to consider “multi-agent” interactions to properly predict emergent phenotypes, a growing trend in deep reinforcement learning [28]. While normative methods have proven valuable for developing insight into biological structure in many areas of application [12, 17], conditions under which evolutionary dynamics can be expected to optimize abstract quantities such as “information” have not been extensively explored, and our goal was to gain insight into such limitations with a simple population dynamics model.

We found that at low initial population densities we could effectively capture the evolution of the agent phenotypes with a mean-field model that neglects interactions between different agents, and hence different phenotypes. Effectively, agents of different phenotypes only interact through the periodic catastrophe we use in our model to drive phenotype selection. In this



regime it turns out that the overall growth rates of each phenotype are independent, and the catastrophes ultimately select for the phenotype with the largest growth rate. Thus, in this regime the maximum eigenvalue of the linear dynamics serves as an effective fitness function for the population that predicts the dominant phenotype at long times. This result is non-trivial, because *a priori* it was not obvious that the catastrophes, which nonlinearly couple the phenotypes, would not change the predictions of the linear eigenvalues. We might therefore hypothesize that organisms that evolved under low-density conditions in which individuals were not often competing for resources or space may be more amenable to normative approaches that aim to predict the optimal phenotype through the use of objective functions of only a single representative agent. This said, in a more realistic stochastic model with asynchronous death rather than a synchronous selecting catastrophe, there is no mechanism (other than extinction or limited resources) for keeping the population size from eventually growing large enough to invalidate the low-density approximation. In an environment with a very large carrying capacity the low-density approximation may still be valid on the relevant evolutionary timescales, and the eigenvalues (growth rates) of such a mean-field model could still predict the phenotype that comprises the largest fraction of the population.

When the population density is sufficiently high that interactions between agents cannot be neglected, then interactions between individuals can alter which phenotypes dominate the population at long times, or lead to long-lived coexistence, and there does not appear to be an obvious single-agent objective function that can be derived to predict the phenotype distributions observed in our simulations or mean-field model. In light of this, one might wonder whether the dynamics of the expected numbers of agents of each phenotype might follow gradient dynamics, but this is not the case for the model studied here. This suggests it may not be possible to derive a function of a single representative phenotype parameter to exist in the interacting regime.

While the conditions under which evolution acts as an optimization process are not generally clear, and may not be valid in some circumstances [20, 29–31], evolution can perhaps still be interpreted as a kind of learning process. While the individual agents in our model do not possess any ability to learn, the population effectively “learns” the structure of the environment over the course of its evolution: the phenotype structure of the population after a long time depends on the structure of the environment and the dynamics of the agents. Even though there is not necessarily any well-defined fitness or objective function being optimized, the adaptation of the population to the environment displays some of the hallmarks of learning.

## Comparisons to related work

Several lines of work have used population dynamics approaches based on master equation models to study the evolution of species under various ecological situations [25, 32], including competition for resources [33–36], prey and predation [37–40], evolutionary arms races [41, 42], and more. Our model follows this methodology, and may be viewed as a generalization of logistic population growth to a spatial environment with multiple possible phenotypes competing for the same limited resources.

Because our model considers only organisms of the same trophic level (i.e., similar organisms compete for a common pool of resources, as opposed to, e.g., predator and prey dynamics), it is similar to the context of the neutral theory of biodiversity [35, 43–45]. The neutral theory of biodiversity assumes that the relative abundances of different species occupying the same environment is not necessarily due to relative evolutionary advantages or disadvantages, but potentially simply due to stochasticity of birth-death dynamics—or “luck.” This idea has been a subject of debate in ecology [35], but neutral models have been successful in explaining

some observed trends in experimental data, although a more complete picture of ecological dynamics is thought to require non-neutral effects [33, 43]. While our model population has no intrinsic differences in terms of their raw birth rates, and hence may superficially be considered neutral, different phenotypes do have emergent advantages over one another because some phenotypes are certainly more successful than others at finding high resource sites and producing more offspring, and is therefore effectively not neutral.

Beyond master-equation based models, there are many other studies that take normative approaches to understanding the behaviors and structures of organisms performing similar foraging tasks. Work in this vein may take several forms of abstraction, ranging from models in which the quantities to be optimized are directly relevant quantities—such as expected numbers of offspring [5] or the expected time it takes organisms to find a stockpile of food [6, 7]—to more abstract “rewards” such as the information an organism has extracted from its environment [11, 17, 27, 46, 47]. In some cases these normative models consider just a single representative agent and its associated parameters that determine its dynamics over the course of a specified task, and hence phenotype. More recent work has considered the dynamics of multiple agents performing tasks through the lens of evolutionary game theory [17, 48–51], sometimes taking the limit of a continuous population density, effectively reducing the individual agent-based dynamics to a coarse-grained statistical description [27]. Many similar approaches are being developed in machine learning applications [28]. We note that the dynamics in these normative approaches are typically a single instantiation of a task, such as a single catastrophe period in our model, not a full evolutionary dynamics, though some studies do use evolution-inspired “genetic algorithms” [52], in which the parameters of a population are remixed and inherited by the next generation after a certain amount of simulation time (i.e., there are typically no mechanistic birth-death events as in the model considered in this work). Relatedly, multiple independent lines of work have begun formulating Bayesian approaches to bridge mechanistic and normative approaches to understanding biological structure [6, 9, 10]. As investigation of biological evolution from this multitude of approaches continues to develop, it would be interesting for future studies to directly investigate the alignment of the outcomes of agent-based birth-death dynamics and normative approaches in such multi-agent settings, especially as increases in computational efficiency allow for simulation and investigation of increasingly complex agents.

### Limitations of the study

As with any model, there are many aspects of the phenomena we seek to study that cannot be fully captured by our approach or assumptions, and warrant discussion. We focus on discussing how some of the important model simplifications could potentially be relaxed, and consider possible implications for empirical data at the end of this section.

**Disparity between evolutionary timescales and individual agent lifetimes.** One of the most challenging aspects of modeling the evolutionary process is that the actions of individual organisms take place on time scales orders of magnitude shorter than the timescales over which significant phenotypic change occurs. This discrepancy is even more extreme for timescales over which accumulated mutations and reproduction result in radiations of new species. This separation of timescales presents an obstacle to modeling the evolution of complex behaviors and structures, motivating our choice of modeling agents with a very simple repertoire of available actions. Even relatively “simple” organisms like *C. Elegans* are capable of much more complicated behaviors than simply relocating based on the conditions at the organism’s current location [22]. While it is possible to consider models of more complex agents, for example with internal states that process environmental information, accounting for both the fast

internal timescales and long evolutionary timescales in a single model is a significant challenge. Other modeling strategies for bridging these timescales will be necessary for the foreseeable future, such as the timescale separation approximation of Mann [27].

**Selection by catastrophes versus stochastic death.** We have implemented selection in our model by an unorthodox choice (in population dynamics models) of implementing a periodically occurring catastrophe that reduces the population down to a fixed size, rather than the more common and realistic scenario of having stochastic death in which agents randomly perish according to an asynchronous death rate (or set of death rates). This said, this method of removing agents is similar to approaches used in evolutionary game theoretic models [17, 48–51]. The reasons for this modeling choice in our population dynamics approach are largely technical: in any stochastic model with a finite population size and stochastic death, the ultimate fate of the population is always extinction. While this may not be a problem in practice for simulations, as the extinction time can be quite large [25] and only a few simulations might terminate in extinction, it complicates analytic analyses. In particular, the mean-field approximation in a model with stochastic death will always predict that the extinction fixed point is unstable and the non-zero population size is stable, which is qualitatively invalidated at long times: the variance of the stochastic process grows with time, owing to the fact that eventually a large enough fluctuation will occur to drive the system to the extinction fixed point, from which it cannot escape. By only allowing agents to perish via a periodically occurring catastrophe that resets the population size to the initial number of agents we allow our model to remain close to the mean-field regime throughout the entirety of the simulation. One can imagine a hybrid scenario in which the catastrophes do not occur periodically, but with stochastic timing (e.g., with Poisson-distributed arrival times), which would retain the technical advantages of the period catastrophes while also re-introducing stochasticity to the selection that could be interesting to investigate in future iterations of this model.

**Sexual versus asexual reproduction.** We considered only asexual reproduction in our model for simplicity, in which the phenotypes inherited by offspring are noisy copies of the parent's phenotype (or direct copies in the mean-field model). One could introduce sexual reproduction into the model by allowing agents at adjacent sites to mate and produce an offspring onto one of the unoccupied sites adjacent to the parent that gives birth, with the child's phenotype determined by some combination of the parents' phenotypes.

The most significant impact that adding sexual reproduction would have on our results is the complication of the low-density approximation that we used to derive the effective cost function  $\Lambda_\gamma$ . Because sexual reproduction requires interactions between individuals, a mean-field model that includes sexual reproduction is necessarily nonlinear. As a consequence, deriving an effective cost function like  $\Lambda_\gamma$  in this case is not straightforward, similar to the high density regime studied in this work. This could suggest that single-agent normative models may be most applicable not simply to organisms that evolved in low density conditions, but also organisms that reproduce primarily through asexual reproduction—or perhaps organisms that reproduce via external fertilization, such as many aquatic animals [53–56]. However, a more thorough analysis of phenotype dynamics in populations with sexual selection is required to test these possibilities.

**Complexity of agents and environment in population dynamics models.** The organisms we have considered in this work are rather simple in terms of their ability to process sensory information and take actions: an organism can only detect the resource level at its current location and move to a neighboring location at a fixed attempt rate. The agents have no memory, nor can they change their behavioral strategies, capabilities that even “simple” organisms like *C. Elegans* have been found to display after short periods of local exploration [22].

Even within the restricted repertoire of sensory processing and behavioral actions, we consider only relatively low-dimensional phenotype spaces, which in turn imposes a restriction on the environmental structure: keeping the phenotype space constrained to 5 components requires we have only 5 resource levels. We could in principle have more resource levels while retaining a low-dimensional phenotype space by defining a function  $\gamma(s)$  parametrized by a small set of abstract parameters, but the choice of this functional form imposes additional structure, which could result in different competitive behavioral strategies for different choices.

Implementing even more complex agents will thus require even larger phenotype spaces; while this is feasible in single-agent normative approaches to evolutionary modeling, having a large number of dynamic complex organisms is extremely computationally expensive. We therefore opted for the more parsimonious choice of a simpler environmental structure with a 1-to-1 correspondence with the (directly interpretable) hopping rate parameters.

**Depleting versus non-depleting resources.** In this study we have considered only non-depleting resources (also called “non-destructive resources”), such that the resource levels  $S_k$  are not modified by the dynamics of the organisms in the environment. This restricts the interpretation of these resource levels to “conditions favorable for reproduction,” such as temperature or shelter, or consumables that are so plentiful that consumption by agents does not appreciably change the resource level of a site. Consumption of resources with replacement (e.g., continued growth of food sources or rainfall) would allow for a feedback loop between organisms and the environment, as the actions of the organisms change the environment, which in turn can alter which behavioral phenotypes are most successful at finding sites with plentiful resources. In such a scenario the dominant phenotype could be constantly changing, unless the statistics of resource consumption and agent dynamics achieve an equilibrium. Such a scenario could give rise to the evolutionary dynamics similar to those observed in populations of turtle ants, in which the head shapes of soldier ants have repeatedly reappeared through evolution [20], though the specific case of soldier ants likely involves additional environmental selective pressures that change with time, not just changing resource distributions. Modeling such systems from a normative perspective would require defining a fitness landscape that changes with time. Investigating these possibilities within frameworks like ours will be an interesting avenue for future work studying the evolution of organism behaviors and sensory capabilities.

**Metabolic costs of reproduction and relocation.** In our model we have not explicitly implemented any metabolic penalty for giving birth or hopping to a new location. The phenomenological effects of such metabolic costs are implicit in the resource level dependence of the birth rates and the imposition of maximum hopping rate  $\gamma_{\max}$ . A more complex agent-based model could give each agent its own metabolic budget. Agents would acquire energy from the environment—e.g., high resource sites would provide more energy per unit time than low resource sites, possibly through consumption of more plentiful resources—and then agents would spend that energy to give birth (incurring a large metabolic cost) or relocate, with the metabolic cost of moving depending on the speed (hopping rate) at which the agents attempt to move. In such a formulation the hopping rate would not need to have an intrinsic maximum per se, though there would be an effective maximum determined by the ability of the agents to acquire enough food to support the high metabolic costs of moving quickly. A metabolic budget could be incorporated into a master equation model or mean-field model by treating the energy as an additional variable on par with the agent position and phenotype, i.e., by counting the number of agents at a particular location with a particular phenotype configuration and specified amount of energy.

We expect that in an environment with non-consumable resources that provide agents with a fixed amount of energy per unit time, with higher rates conferred by higher resource levels, one could choose the energy rates and metabolic costs of reproduction and relocation in such a way that the results of our current study are qualitatively reproduced.

**Comparison to empirical data.** Owing to many of the simplifications of our model discussed above, it is difficult to directly compare our results to empirical data. Studies investigating the evolution of innate behaviors have focused on the interplay of the innate behavior evolution and learning within an individual's lifetime [57–59], a relationship that is difficult to study in this sort of population-dynamics model due to the disparate timescales of evolution and within-life learning.

The most promising means of comparing our model to empirical data would be to design and perform an experiment that closely mimics the set-up of the model. For instance, one could imagine an experiment with a population of organisms that a) are capable of self-propelled motion, b) have some means of detecting local nutrient level/temperature/etc., c) have tools for genetic modification to create mutants representing an initial distribution the phenotype behaviors, and d) have relatively short reproduction times. These organisms would be placed into a controlled environment in which nutrients or reproductively favorable conditions are randomly distributed; consumable resources would need to be replenished at rates comparable to the rate of uptake by the organisms. The organisms would be given time to explore their environment and discover locations of high nutrient content, after which they are culled back to the initial population size, mimicking the catastrophe in our model. By using tracking software to monitor the organisms' relocation rates out of each nutrient level, one could observe over a large number of cullings whether there is a change in the distribution of phenotyped behaviors over the course of the experiment.

This said, the implementation of such an experiment may not be pragmatic for the foreseeable future, and would likely be limited to single-celled organisms. A more promising avenue might be to compare the qualitative predictions of future extensions of this work modeling more complex organisms. For example, experimental studies have observed selection-induced genetic drift in a mixed population of wild-type mice and mice with a mutation that disrupts their circadian rhythm [60]. One could model such an experiment by incorporating some of the more complex features discussed above into a formalism like the one presented in this work, with additional modifications to introduce a day/night cycle that interacts with agents' internal circadian rhythms. It would be of great interest to model similar scenarios that investigate the relative reproductive success of organisms with different perturbations to their sensory modalities, to see if the outcomes can be assessed using various normative predictions commonly used in neuroscience.

## Future directions

In this work we have focused on agents with relatively simple sensory capabilities and behaviors in order to tractably compare agent-based simulations with mean-field approximations and calculate an effective fitness function in a tractable limit of the mean-field model. Two main avenues for future investigations of the models that we have not attempted to elucidate here include i) performing more advanced mathematical analyses of master equation models, for example by investigating the effect of stochastic fluctuations on the stability of phenotype coexistence in more detail, and ii) pushing the boundaries of agent and environment complexity, along with implementing other more realistic features of evolutionary dynamics.

Investigating the effects of stochasticity requires tools to go beyond the mean-field models introduced here. Several tools exist in the mathematical and statistical physics literature for this purpose, such as the Wentzel–Kramers–Brillouin (WKB) approximation [25, 32] or field-theoretic formulations of the master equation model [37–39, 61]. Because these cases explicitly deal with beyond-mean-field effects it would be more natural to use stochastic death in this setting, rather than the periodic catastrophe introduced in this work to control the efficacy of the mean-field approximation. For example, the WKB approximation could allow for a direct calculate of the large deviation functions determining the extinction transition. This large deviation function may serve as an effective cost function of the population, potentially allowing for the derivation of such a function, even in the high-density regime [61].

In terms of increasing organism complexity, a long-term goal of this line of work would be to endow agents with “brains”—small neural networks that take in some sensory information from the environment and then determine how the agent will attempt to move. In such a context evolution would act on the structure of the brain, similar to the case of Braitenberg vehicles [62, 63] (as opposed to determining the synaptic weights in neural network models such as in [64]). In practice, this will be a challenging direction due to the large phenotype space and the separation of timescales involved in sensory processing and evolution. A more tractable approach would be to endow the organisms with a simple “receptive field” that detects sensory information in its immediate environment, down to a modest modification of the current model that allows for agents to detect resource levels at neighboring sites in addition to its current location. Straightforward increases in environmental complexity include adding spatial correlations between resources levels at different locations, as well as introducing sensory signals distributed through the environment that are correlated with resource levels that agents cannot detect directly. Modeling the impact of consumable resources is also of ecological importance, and could be implemented in our models, e.g., by introducing a non-motile “species” at each site that is consumed by agents and replenished at a fixed rate. In this way one could also introduce more realistic metabolic limitations: actions such as giving birth or relocating could incur energy costs that must be replenished by consuming food in the environment. This is one way we could formally remove the artificial restriction of a maximum hopping rate  $\gamma_{\max}$ , allowing the organisms to move as fast as possible between sites if they can acquire the necessary energy to maintain such a strategy.

Finally, one of the goals of this work was to attempt to derive functions that could serve as effective cost functions in normative approaches to modeling the biological structures and behavioral strategies that emerge from the more fundamental evolutionary dynamics. It would be interesting to perform a more thorough comparison of the predictions of birth-death dynamics and normative approaches to identify regimes where the predictions diverge, especially in approaches in which the objective functions are abstract quantities such as information-theoretic quantities related to random variables in an organism’s environment [9–11]. Doing so will provide valuable insight into the relationships between fundamental evolutionary processes and the more parsimonious normative approaches useful for modeling complex organisms.

## Methods

### Details of the agent-based stochastic simulations

In the agent-based stochastic simulation, we use an  $L \times L$  lattice with periodic boundary conditions, with a standard value of  $L = 128$ . The resource level at each location is drawn from a discrete uniform distribution with 5 possible levels, labeled  $S_1, S_2, S_3, S_4,$  and  $S_5$ . Each agent has a phenotype vector  $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5)$  where the  $i^{\text{th}}$  component corresponds to

the hopping or relocation rate when the agent is at a site with resource level  $S_i$ . The birth rate of agents at spatial coordinate  $\mathbf{x}$  within the environment is a function of the resource level at that location,  $b(s_{\mathbf{x}}) = b(S_i) \equiv b_i$  if  $s_{\mathbf{x}} = S_i$ . Initially, 500 agents are introduced to the system, each of which has a random initial location and random initial phenotype vector with each component ranging from  $\gamma_{\min} = 0$  to  $\gamma_{\max} = 10$ . This bound is set considering that the relocation rate cannot go negative and there are metabolic costs that constrain agents from relocating too frequently. An offspring inherits its parent's phenotype vector with Gaussian noise added to each component. The noise term follows  $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbb{I}_5)$ , with  $\sigma = 0.05$  (arbitrary units) in all the stochastic simulations throughout this study. If the noise would push a relocation rate component outside of the range  $[\gamma_{\min}, \gamma_{\max}]$  then the component is set to the closest boundary value. Each location can hold at most 1 agent at any given time and thus the reproduction process has to leave the offspring at one of the vertical or horizontal neighboring locations in practice—if the agent attempts to give birth into an occupied site then this reproduction attempt is considered as unsuccessful. Similarly, hopping may only occur from an agent's position to one of the four adjacent sites, with the hop being unsuccessful if the target site is occupied.

Importantly, a system-wide periodic random culling is applied to the system after a duration of  $T = 0.3$  (arbitrary time units), which reduces the number of agents back to its initial state. This periodic catastrophe removes agents in a random manner without distinguishing its phenotype or location, functioning as the selection in the evolutionary dynamics. We expect that the most competitive phenotypes will be selected out after a series of growth and catastrophe events. A full iteration of the simulation consists of stochastic growth and relocation, as well as a random culling at the end. The next iteration starts off with the survived agents from the last iteration as the new initial condition.

As the carrying capacity of each location is set to 1, the  $128 \times 128$  2d lattice can hold at most 16384 agents. For a typical run, in 0.3 units of simulation time the number of agents grows from 500 to around 850 before the random culling brings the number of agents back to 500. The initial density of agents is defined as initial number of agents divided by total possible number of agents in the system, which is  $500/16384 \approx 3.1\%$  for this standard parameter set. When investigating the high density regime, we kept the number of agents to be 500 in all the initial conditions and reduced the lattice size from  $128 \times 128$  to smaller size to achieve higher initial densities.

To reduce the potential variability from stochastic simulations, we run 20 independent simulations with the same set of model parameters and resource background for all the agent-based stochastic simulations in this study. Then the averaged phenotype components cross all trials and their standard deviation were calculated as the solid lines and shaded areas in Figs 2E and 4A–4D.

The stochastic agent-based model was simulated in the Python programming language using a direct method Gillespie algorithm [23]. The Gillespie algorithm computes the total rates of the possible events that may take place in the system to determine when and what the next “action” is. Here “action” refers to either birth or relocation of an agent in the system. The total rate is the sum of the birth and relocation rates of all the agents in the system; the algorithm first determines the time of the next action and then probabilistically chooses one of the available actions based on the relative rate of each action compared to the total rate. That is, at each step, the time of the next birth or relocation event is drawn from an exponential distribution with mean  $1/\text{total rate}$ . The probability that the next event is agent  $A$  giving birth or relocating is  $A$ 's birth rate/total rate or  $A$ 's relocation rate/total rate. The target location of the birth or relocation event is then randomly chosen from one of the four possible neighboring locations of the agent's current position. If the chosen target location is already occupied, the

birth or relocation attempt is considered unsuccessful. Note that in this implementation of the algorithm, hopping or birth from site  $\mathbf{x}$  to site  $\mathbf{y}$  is a two-step process: the rates going into the total rate are just  $b(s_{\mathbf{x}})$  and  $\gamma_{\mathbf{x}}$ , and *then* if one of these events is chosen to occur the target site is drawn with probability  $1/2d = 1/4$  in our two-dimensional environment. However, an alternative way we could have implemented these dynamics is via a one-step process in which we include the birth and hopping rates from site  $\mathbf{x}$  to site  $\mathbf{y}$  for all possible target sites as separate rates into the total rate (all of which happen to have the same numerical values). In this second possible implementation the birth and hopping rates must be scaled down by the number of target sites ( $2d = 4$ ) in order to match the results of the two-step implementation. This technical detail is important when developing the mean-field model, as the one-step implementation corresponds more directly to the master equation formalism we introduce in Master equation model of the stochastic dynamics and derivation of the mean-field approximation, and scaling the birth and hopping rates is required to make the derived mean-field model match the agent-based simulations.

Multiple Python packages were used in the stochastic simulations, mean-field model analyses as well as data visualizations in this study, most notably, Numpy [65], Scipy [66], and Matplotlib [67]. In the stochastic simulations, each stochastic trajectory was run in series on a single CPU core. A computing server with two Intel Xeon Processor E5–2690 v4 CPU (14 cores for each CPU unit) was used to generate 20 stochastic trajectories for the same set of parameters in parallel with the aid of the Python Multiprocessing module. Typically, each run takes approximately 35 minutes for a system starting with 500 agents and  $b_5 = 4$  in Fig 2. Higher  $b_5$  led to more agents being produced, resulting in a higher simulation load, and thus simulations with higher  $b_5$  value take longer to run. For example, the simulations with  $b_5 = 16$  depicted in Fig 3 took over an hour to run.

## Master equation model of the stochastic dynamics and derivation of the mean-field approximation

**Full model in a  $d$ -dimensional environment with random resources.** Formalizing agent-based stochastic simulations as mathematical models that allow for tractable analytic analyses is often difficult [68]. Instead of keeping track of every agent individually, an equivalent description that lends itself to a standard class of mathematical models is one in which we only keep track of the number of agents with each phenotype at each spatial location in the environment. In order to write down such a model we need to make a few simplifications:

1. We discretize the phenotype space  $\gamma$  so that there are a finite number of discrete phenotypes and the number of agents with each phenotype can be enumerated. In principle, we could work with a continuous phenotype space and define a density of agents with each phenotype, but for practical numerical calculations using the mean-field model it is necessary to discretize the phenotypes.
2. We omit reproduction noise. In the mean-field model we will formally include all possible phenotypes in the initial conditions, so we do not need reproduction noise to introduce any phenotypes that were not initially present, as is the case in the agent-based simulations.

Now, let  $n_{\mathbf{x},\gamma}$  be the number of agents at location  $\mathbf{x}$  with a particular phenotype  $\gamma$  (i.e., a particular choice of values for each of the relocation rate components) and denote by  $\mathbf{n}$  the collection of all of these variables into a single vector. The probability that the state of the system has a particular configuration  $\mathbf{n}$  at time  $t$ ,  $p(t, \mathbf{n})$ , is the solution to the system of differential



equations

$$\frac{dp(t, \mathbf{n})}{dt} = \sum_{\mathbf{x}, \gamma} \sum_{\mathbf{y} \in V(\mathbf{x})} \left\{ f_{\mathbf{x}}(\mathbf{n} - \hat{\mathbf{e}}_{\mathbf{x}, \gamma}) b(s_{\gamma}) n_{\gamma} p(t, \mathbf{n} - \hat{\mathbf{e}}_{\mathbf{x}, \gamma}) - f_{\mathbf{y}}(\mathbf{n}) b(s_{\mathbf{x}}) n_{\mathbf{x}, \gamma} p(t, \mathbf{n}) \right. \\ \left. + f_{\mathbf{x}}(\mathbf{n} - \hat{\mathbf{e}}_{\mathbf{x}, \gamma}) \gamma_{\mathbf{y}} (n_{\gamma} + 1) p(t, \mathbf{n} + \hat{\mathbf{e}}_{\mathbf{y}, \gamma} - \hat{\mathbf{e}}_{\mathbf{x}, \gamma}) - f_{\mathbf{y}}(\mathbf{n}) \gamma_{\mathbf{x}} n_{\mathbf{x}, \gamma} p(t, \mathbf{n}) \right\}. \tag{3}$$

We have suppressed the conditional dependence of  $p(t, \mathbf{n})$  on the configuration of resource levels  $s_{\mathbf{x}}$ . The notation  $\hat{\mathbf{e}}_{\mathbf{x}, \gamma}$  represents a vector with a 1 at the element labeled by  $(\mathbf{x}, \gamma)$  and 0 elsewhere. The bare birth rate per capita at a site  $\mathbf{x}$  is  $b(s_{\mathbf{x}})$ , explicitly denoting that the birth rate depends only on the resource level of the current location. We assume that each site has a maximum occupancy of a single agent, so all births occur via an agent at some site  $\mathbf{y}$  birthing a new agent into an adjacent site  $\mathbf{x}$ . The bare hopping rate out of a site  $\mathbf{y}$  is denoted  $\gamma_{\mathbf{y}}$ ; the value of the hopping rate is equal to the hopping agent's value of the hopping rate at resource level  $s_{\mathbf{y}}$ . Both birth rates and hopping rates into a site  $\mathbf{x}$  are modulated by a factor  $f_{\mathbf{x}}(\mathbf{n})$ , which is 0 if site  $\mathbf{x}$  is occupied and 1 if unoccupied. The possible target sites for hopping or birth by an agent at site  $\mathbf{x}$  are identified by the notation  $\mathbf{y} \in V(\mathbf{x})$ , which means that site  $\mathbf{y}$  is in the ‘vicinity’ of  $\mathbf{x}$ . By assumption,  $\mathbf{x} \notin V(\mathbf{x})$ . For our  $2d$  simulations,  $V(\mathbf{x})$  consists of the nearest neighbor sites  $\mathbf{y} = \mathbf{x} \pm (1, 0)^T$  and  $\mathbf{y} = \mathbf{x} \pm (0, 1)^T$ .

The initial condition  $p(0, \mathbf{n})$  represents the initial configuration of the system. In many applications the initial condition is deterministic, but one can choose an initial distribution as well. In particular, we can ensure every phenotype is represented by choosing an initial distribution of  $N_0$  agents and assigning each agent's phenotype and location by drawing from the possible phenotypes and sites with a uniform probability. This eliminates the need for reproduction noise in this model.

The dynamics of Eq (3) run from  $t = 0$  up to  $t = T$ , at which point the catastrophe occurs.  $N_0$  agents are randomly selected out of the agents alive at time  $t = T$  and retain their spatial position and phenotype as we eliminate the remaining agents from the environment. The distribution of selected  $N_0$  agents forms the initial distribution for the next iteration. This procedure repeats indefinitely, but we expect the distributions of the total number of agents of each phenotype to converge after a large number of catastrophes, though it may not be the case that the full spatial distributions  $p(0, \mathbf{n})$  and  $p(T, \mathbf{n})$  converge.

By summing Eq (3) over all possible configurations  $\mathbf{n}$ , changing dummy indices, and using the fact that  $p(t, \mathbf{n}) = 0$  if any  $n_{\mathbf{x}, \gamma} < 0$ , one can show that the right hand side vanishes, as it must in order for probability to be conserved.

In practice, solving probability flow equations such as Eq (3) is intractable except for very simple cases. Moreover, implementing the catastrophe at this level of description is challenging. A more common approach that lends itself to analytical and numerical analyses is to instead derive a system of equations for the moments of the distribution and make a mean-field approximation to close the system of equations (as moments are normally coupled to higher order moments). For example, the ‘mean-field approximation’ corresponds to deriving a system of equations for the first moments of the distribution and assuming all higher order cumulants vanish, such that averages over functions of  $\mathbf{n}$  can be replaced by those functions evaluated at the average  $\langle \mathbf{n} \rangle$ . For Eq (3) the mean-field system of equations is

$$\frac{d\langle n_{\mathbf{x}, \gamma} \rangle}{dt} = \sum_{\mathbf{y} \in V(\mathbf{x})} \left[ f_{\mathbf{x}}(\langle \mathbf{n} \rangle) b(s_{\gamma}) \langle n_{\gamma} \rangle + f_{\mathbf{x}}(\langle \mathbf{n} \rangle) \gamma_{\mathbf{y}} \langle n_{\gamma} \rangle - f_{\mathbf{y}}(\langle \mathbf{n} \rangle) \gamma_{\mathbf{x}} \langle n_{\mathbf{x}, \gamma} \rangle \right], \tag{4}$$

where we employed the mean-field approximation to replace products such as  $\langle f_x(\mathbf{n})n_{x,\gamma} \rangle$  with the factored moments  $f_x(\langle \mathbf{n} \rangle)\langle n_{x,\gamma} \rangle$ . Note that in order for this approximation to work with our modulation function  $f_x(\mathbf{n})$  we must specify the values that this function  $f_x$  takes for continuous-valued inputs, as  $\langle \mathbf{n} \rangle$  will not be an integer. In this work we make the choice

$$f_x(\mathbf{n}) = 1 - \sum_{\gamma'} n_{x,\gamma'}, \tag{5}$$

which satisfies the desired boundary conditions of  $f_x = 1$  when the site is unoccupied and 0 when there is an agent at the site. This choice amounts to the assumption that  $\langle f_x(\mathbf{n})n_{x,\gamma} \rangle / \langle n_{x,\gamma} \rangle$  varies approximately linearly with  $\sum_{\gamma'} \langle n_{x,\gamma} \rangle$ . Although it appears that we should also impose a threshold to prevent this choice of  $f_x(\mathbf{n})$  from becoming negative, this is unnecessary in practice as long as the initial conditions respect the occupancy of the system, as once  $f_x$  reaches 0 it blocks any further increases in occupancy that would drive it to be negative. This is true in both the full master equation model and the mean-field model (Eq 4).

These mean-field equations are still challenging to analyze due to the spatial variability of the system. We simplify our analysis further by deriving an approximate system for the total number of agents at sites of each resource level, detailed in the next section.

**Reduction to a 5-state mean-field model.** Rather than keeping track of the expected number of agents of each phenotype at every location in the environment we will instead focus only on the number of agents of each phenotype at sites of each resource level  $S_k$ . The total expected number of agents at resource level  $k$  is

$$\langle n_{k,\gamma} \rangle = \sum_x \delta(s_x, S_k) \langle n_{x,\gamma} \rangle,$$

where  $\delta(s_x, S_k)$  is an indicator function that is 1 if  $s_x = S_k$  and 0 otherwise. We will denote spatial sites by bold  $\mathbf{x}$  and  $\mathbf{y}$  indices while states will be denoted by italic  $k$  and  $i$  indices.

We can derive the system of equations for  $\langle n_{k,\gamma} \rangle$  by multiplying Eq (4) by  $\delta(s_x, S_k)$  and summing over  $\mathbf{x}$ . The restricted sum over  $\mathbf{y} \in V(\mathbf{x})$  poses a complication for this approach, and will require an approximation to be discussed below. We detail the calculation for each term of Eq (4) in turn.

The left-hand side of Eq (4) is trivial, and simply gives  $\frac{d}{dt} \langle n_{k,\gamma} \rangle$ , so we consider the birth rate term  $\sum_x \delta(s_x, S_k) f_x(\langle \mathbf{n} \rangle) b(s_x) \langle n_{x,\gamma} \rangle$  first. Note that we may write  $b(s_y) = \sum_i \delta(s_y, S_i) b_i$ , which essentially just says that  $b(s_y) = b_i$  if  $s_y = S_i$ . Proceeding,

$$\begin{aligned} & \sum_x \sum_{\mathbf{y} \in V(\mathbf{x})} \delta(s_x, S_k) f_x(\langle \mathbf{n} \rangle) b(s_y) \langle n_{\mathbf{y},\gamma} \rangle \\ &= \sum_x \delta(s_x, S_k) f_x(\langle \mathbf{n} \rangle) \sum_{\mathbf{y} \in V(\mathbf{x})} \left\{ \sum_i \delta(s_y, S_i) b_i \right\} \langle n_{\mathbf{y},\gamma} \rangle \\ &= \sum_i b_i \sum_x \delta(s_x, S_k) f_x(\langle \mathbf{n} \rangle) \sum_{\mathbf{y} \in V(\mathbf{x})} \delta(s_y, S_i) \langle n_{\mathbf{x},\gamma} \rangle. \end{aligned}$$

The sum over  $\mathbf{y}$  depends on the current value of  $\mathbf{x}$ , and so cannot be easily evaluated. One way to circumvent this problem would be to make the extreme approximation of allowing the target sites to include not just nearest neighbors but *all* sites in the environment (a usage sometimes used synonymously with “mean-field approximation,” but this is due to the fact that the mean-field approximation becomes exact on all-to-all graphs in many models). A slightly less drastic approximation is to assume that the environment is “well-mixed,” such that the sum over a local set of sites should just be a scaled down version of the sum over all sites.

Specifically, we will assume

$$\sum_{y \in V(\mathbf{x})} A_y \approx \frac{2d}{M} \sum_y A_y, \tag{6}$$

where the sum on the right-hand side is over all sites  $y$  and  $A_y$  is a general observable. The ratio  $2d/M$  is the number of target sites (nearest-neighbors) in  $d$  dimensions to the total number of sites in the environment. For  $A_y = \delta(s_y, S_i) \langle n_{y,\gamma} \rangle$  we therefore obtain

$$\sum_{y \in V(\mathbf{x})} \delta(s_y, S_i) \langle n_{y,\gamma} \rangle \approx \frac{2d}{M} \sum_y \delta(s_y, S_i) \langle n_{y,\gamma} \rangle = \frac{2d}{M} \langle n_{i,\gamma} \rangle.$$

In S5 Fig we quantify the standard deviation of the error of this assumption.

This approximation of the sum no longer depends on  $\mathbf{x}$ , so we may now evaluate the sum over the modulation function  $f(\langle \mathbf{n} \rangle)$ . For our choice of the modulation function we have

$$\begin{aligned} \sum_{\mathbf{x}} \delta(s_{\mathbf{x}}, S_k) f_{\mathbf{x}}(\langle \mathbf{n} \rangle) &= \sum_{\mathbf{x}} \delta(s_{\mathbf{x}}, S_k) \left( 1 - \sum_{\gamma'} \langle n_{\mathbf{x},\gamma'} \rangle \right) \\ &= \left( w_k M - \sum_{\gamma'} \langle n_{k,\gamma'} \rangle \right) \\ &= w_k M \left( 1 - \frac{\sum_{\gamma'} \langle n_{k,\gamma'} \rangle}{w_k M} \right), \end{aligned}$$

where  $\sum_{\mathbf{x}} \delta(s_{\mathbf{x}}, S_k) = w_k M$  is the expected number of sites of resource level  $k$ , with  $M$  the total number of sites and  $w_k$  the frequency of resource level  $k$ . We assume the number of sites is large enough that the actual number of sites of each level is equal to its expected value. Thus, the birth rate term evaluates to

$$\sum_{\mathbf{x}} \sum_{y \in V(\mathbf{x})} \delta(s_{\mathbf{x}}, S_k) f_{\mathbf{x}}(\langle \mathbf{n} \rangle) b(s_y) \langle n_{y,\gamma} \rangle \approx 2d w_k \left( 1 - \frac{\sum_{\gamma'} \langle n_{k,\gamma'} \rangle}{w_k M} \right) \sum_i b_i \langle n_{i,\gamma} \rangle.$$

Next we consider the first hopping rate term. For this term it will be useful to use the identity  $\gamma_{\mathbf{x}} = \sum_i \delta(s_{\mathbf{x}}, S_i) \gamma_i$ , which like the similar identify for the birth rates just says that  $\gamma_{\mathbf{x}} = \gamma_i$  if  $s_{\mathbf{x}} = S_i$ . Then,

$$\begin{aligned} &\sum_{\mathbf{x}} \sum_{y \in V(\mathbf{x})} \delta(s_{\mathbf{x}}, S_k) f_{\mathbf{x}}(\langle \mathbf{n} \rangle) \gamma_y \langle n_{y,\gamma} \rangle \\ &= \sum_{\mathbf{x}} \delta(s_{\mathbf{x}}, S_k) f_{\mathbf{x}}(\langle \mathbf{n} \rangle) \sum_{y \in V(\mathbf{x})} \left\{ \sum_i \delta(s_y, S_i) \gamma_i \right\} \langle n_{y,\gamma} \rangle \\ &= \sum_{\mathbf{x}} \delta(s_{\mathbf{x}}, S_k) f_{\mathbf{x}}(\langle \mathbf{n} \rangle) \sum_i \gamma_i \sum_{y \in V(\mathbf{x})} \delta(s_y, S_i) \langle n_{y,\gamma} \rangle \\ &\approx 2d w_k \left( 1 - \frac{\sum_{\gamma'} \langle n_{k,\gamma'} \rangle}{w_k M} \right) \sum_i \gamma_i \langle n_{i,\gamma} \rangle, \end{aligned}$$

where we used the same approximations for the sums over  $y$  and  $\mathbf{x}$  as in the approximation of the birth rate term.

For the final term we have

$$\begin{aligned} \sum_x \sum_{y \in V(x)} \delta(s_x, S_k) f_y(\langle \mathbf{n} \rangle) \gamma_x \langle n_{xy} \rangle &= \sum_{y \in V(x)} f_y(\langle \mathbf{n} \rangle) \gamma_k \sum_x \delta(s_x, S_k) \langle n_{xy} \rangle \\ &\approx \frac{2d}{M} \sum_y f_y(\langle \mathbf{n} \rangle) \gamma_k \langle n_{k,y} \rangle, \end{aligned}$$

where the equality in the first line follows from using the fact that  $\delta(s_x, S_k) \gamma_x = \delta(s_x, S_k) \gamma_k$ , and we used our approximation for the sum over nearest-neighbor sites as a fraction of the sum over all sites in going to the second line. For our choice of  $f$  the sum over  $y$  is

$$\begin{aligned} \sum_y f_y(\langle \mathbf{n} \rangle) &= \sum_y \left( 1 - \sum_{y'} \langle n_{y,y'} \rangle \right) \\ &= \left( M - \sum_{y'} \sum_y \langle n_{y,y'} \rangle \right) \\ &= \left( \sum_i w_i \right) M - \sum_{y'} \sum_i \langle n_{i,y'} \rangle \\ &= \sum_i w_i M \left( 1 - \frac{\sum_{y'} \langle n_{i,y'} \rangle}{w_i M} \right), \end{aligned}$$

where we used the fact that  $\sum_x \langle n_{x,y} \rangle = \sum_i \langle n_{i,y} \rangle$  (the total number of agents obtained by summing over all sites must be equal to the total number of agents obtained by summing over all states) and the total frequency of each state must be  $\sum_i w_i = 1$ . The final term of the mean-field equation is thus

$$\sum_x \sum_{y \in V(x)} \delta(s_x, S_k) f_y(\langle \mathbf{n} \rangle) \gamma_x \langle n_{xy} \rangle \approx \sum_i 2dw_i \left( 1 - \frac{\sum_{y'} \langle n_{i,y'} \rangle}{w_i M} \right) \gamma_k \langle n_{k,y} \rangle.$$

Putting everything together, the system of differential equations for our 5-state model is

$$\begin{aligned} \frac{d\langle n_{k,y} \rangle}{dt} &\approx 2dw_k \left( 1 - \frac{\sum_{y'} \langle n_{k,y'} \rangle}{w_k M} \right) \sum_i (b_i + \gamma_i) \langle n_{i,y} \rangle \\ &\quad - \sum_i 2dw_i \left( 1 - \frac{\sum_{y'} \langle n_{i,y'} \rangle}{w_i M} \right) \gamma_k \langle n_{k,y} \rangle. \end{aligned} \tag{7}$$

This is nearly our final result for the mean-field model. One additional correction is necessary to align it with our agent-based simulations.

**Conditional-event correction and the low-density approximation.** The 5-state model Eq (7) describes the mean-field approximation to Eq (3), which would correspond to a direct Gillespie simulation of the master Eq (3) that treats all of the possible events at any moment on equal footing when calculating the total rate to determine the next event. For example, birth from site  $\mathbf{x}$  to  $\mathbf{x} \pm (0, 1)^T$  and birth from  $\mathbf{x} \pm (1, 0)^T$  are all treated as separate events, even though they have the same rate (when the target sites are all unoccupied). The same is true for the different hopping directions.

However, as detailed in Details of the agent-based stochastic simulations, our implementation of the agent dynamics uses the fact that the birth and hopping rates only depend on an agent’s current location, not its target location, to calculate the total Gillespie rate using only a single instance of the birth or hopping rates. Then, *conditional on a birth or hopping event*

being determined to occur, the target site is chosen randomly from one of the possible targets with equal probability.

These two implementations of the Gillespie dynamics are *not* equivalent. However, they are simply related. The conditional birth/hopping implementation essentially reduces the birth and hopping rates by a factor of the number of target sites. In a  $d$ -dimensional lattice this means we should rescale the rates  $b_i \rightarrow b_i/2d$  and  $\gamma_i \rightarrow \gamma_i/2d$  in Eq (7), yielding the mean-field dynamics that correspond to our implementation of the stochastic simulations,

$$\frac{d\langle n_{k,\gamma} \rangle}{dt} \approx w_k \left( 1 - \frac{\sum_{\gamma'} \langle n_{k,\gamma'} \rangle}{w_k M} \right) \sum_i (b_i + \gamma_i) \langle n_{i,\gamma} \rangle - \sum_i w_i \left( 1 - \frac{\sum_{\gamma'} \langle n_{i,\gamma'} \rangle}{w_i M} \right) \gamma_k \langle n_{k,\gamma} \rangle. \tag{8}$$

The mean-field equations given in the main text follow from this equation. In particular, the low-density approximation follows from this system of equations by assuming that  $\sum_{\gamma'} \langle n_{k,\gamma'} \rangle \ll w_k M$ , giving

$$\frac{d\langle n_{k,\gamma} \rangle}{dt} \approx w_k \sum_i (b_i + \gamma_i) \langle n_{i,\gamma} \rangle - \gamma_k \langle n_{k,\gamma} \rangle. \tag{9}$$

for each phenotype  $\gamma$ . As this is a linear equation, we may also write it in matrix form,

$$\frac{d}{dt} \langle \mathbf{n}_\gamma \rangle = \mathbb{M}_\gamma \langle \mathbf{n}_\gamma \rangle,$$

where  $\langle \mathbf{n}_\gamma \rangle$  is a  $5 \times 1$  vector of the expected number of agents of phenotype  $\gamma$  at each of the 5 resource levels and  $\mathbb{M}_\gamma$  is the transition matrix involving the birth and hopping rates. This equation can be formally solved in terms of the matrix exponential:  $\langle \mathbf{n}_\gamma(t) \rangle = \exp(t\mathbb{M}_\gamma) \langle \mathbf{n}_\gamma(0) \rangle$ , for each  $\gamma$ .

Note that although the phenotypes have become dynamically independent in this low-density approximation, the catastrophe still couples them together, as the fraction by which the population size is reduced depends on all of the phenotypes. We give our mathematical analysis of the effect of catastrophes in the next section.

**Catastrophes in the mean-field model.** We introduce catastrophes into the dynamics of Eq (8) by running the mean field dynamics up to time  $t = T$ , at which point the catastrophe occurs and we reset the population to its initial size  $N_0$  by uniformly reducing the size of each phenotype:

$$\mathbf{n}_\gamma^{(\ell+1)}(0) = \frac{N_0}{\sum_{i,\gamma'} n_{i,\gamma'}^{(\ell)}(T)} \mathbf{n}_\gamma^{(\ell)}(T), \tag{10}$$

where  $\ell$  indexes the number of catastrophes that have occurred and we have dropped the angled brackets for convenience. Though this equation holds for any density, in the low-density limit we can formally solve Eq (9) and use this iterative map to derive the effective objective function  $\Lambda_\gamma$ , the growth rate of the phenotypes, as we show below.

**At low-density the most competitive phenotype has the largest maximum eigenvalue.** We now prove that the most competitive phenotype in the reduced mean field model is that whose hopping rates maximize the largest eigenvalue of  $\mathbb{M}_\gamma$ , denoted  $\Lambda_\gamma$ .

Our proof begins by using the explicit solution of Eq (9) in terms of the matrix exponential, together with the reset condition Eq (10), to relate the final population sizes just before the

next catastrophe by a system of iterative maps:

$$\mathbf{n}_\gamma^{(\ell+1)}(T) = \frac{N_0}{\sum_{i,\gamma'} \mathbf{n}_{i\gamma'}^{(\ell)}(T)} \exp(T\mathbb{M}_\gamma) \mathbf{n}_\gamma^{(\ell)}(T), \tag{11}$$

where  $\mathbf{n}_\gamma^{(\ell)}(T)$  is the vector of population sizes at the end of the  $\ell^{\text{th}}$  iteration (i.e., just prior to the  $\ell^{\text{th}}$  catastrophe). As  $\ell \rightarrow \infty$  this iterative map may converge to a fixed point,  $\mathbf{n}_\gamma^*(T)$ , which is the stable distribution of the population in this dynamical system. The fixed point equation may be rewritten

$$\left( \frac{\sum_{i,\gamma'} \mathbf{n}_{i\gamma'}^*(T)}{N_0} \mathbb{I} - \exp(T\mathbb{M}_\gamma) \right) \mathbf{n}_\gamma^*(T) = \mathbf{0} \tag{12}$$

for each phenotype  $\gamma$ , and where  $\mathbb{I}$  is the identity matrix. We expect the trivial solution  $\mathbf{n}_\gamma^*(T) = \mathbf{0}$  to be obtained for phenotypes that are outcompeted by other phenotypes, such that these phenotypes comprise smaller and smaller fractions of the total population, eventually vanishing as  $\ell \rightarrow \infty$ .

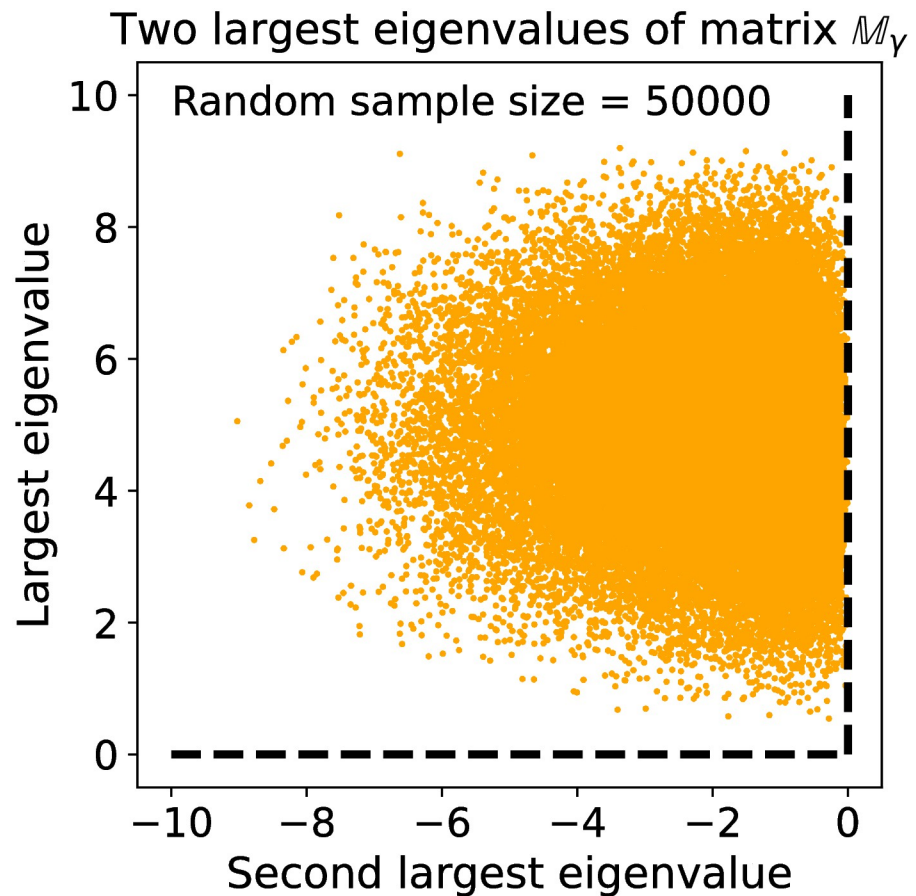
Because extinction is not possible in this model, at least one phenotype must achieve a non-zero population size at the fixed point. For a non-trivial fixed point to exist, the determinant of  $(\sum_{i,\gamma'} \mathbf{n}_{i\gamma'}^*(T))/N_0 \mathbb{I} - \exp(T\mathbb{M}_\gamma)$  must vanish for some phenotypes  $\gamma$ . This implies that  $(\sum_{i,\gamma'} \mathbf{n}_{i\gamma'}^*(T))/N_0$  is an eigenvalue of the matrices  $\exp(T\mathbb{M}_\gamma)$  and  $\mathbf{n}_\gamma^*(T)$  is the corresponding eigenvector. The eigenvalues and eigenvectors of  $\exp(T\mathbb{M}_\gamma)$  are related to the eigenvalues  $\lambda_\gamma$  and eigenvectors of the much simpler  $\mathbb{M}_\gamma$ , which can be efficiently computed ([69]). As empirically observed in Fig 6, only one of the eigenvalues of  $\mathbb{M}_\gamma$  is positive; the rest are negative. As the self-consistency condition requires that  $\sum_{k,\gamma} \mathbf{n}_{k\gamma}^*(T) = N_0 e^{\lambda_\gamma T}$ , the only possible choice of eigenvector is that which corresponds to the largest, positive eigenvalue, which we denote  $\Lambda_\gamma$ , as otherwise  $\sum_{k,\gamma} \mathbf{n}_{k\gamma}^*(T)$  would be less than  $N_0$ , which is not possible in this model. (Moreover, the eigenvectors with negative eigenvalues often have negative components, which does not make sense as a population count).

Because we assume that all phenotypes are present in equal proportions in the mean field model, it follows that the phenotype with the largest overall maximum eigenvalue  $\Lambda_\gamma$  will be the sole survivor in the  $\ell \rightarrow \infty$  limit. The fact that  $(\sum_{i,\gamma'} \mathbf{n}_{i\gamma'}^*(T))/N_0$ , a single scalar number, must be an eigenvalue of  $\exp(T\mathbb{M}_\gamma)$  for every  $\gamma$  implies that the only way multiple phenotypes may coexist in the low-density limit is if their largest eigenvalues  $\Lambda_\gamma$  have the same value. In such a case the relative proportions of those phenotypes will depend on the overlap of  $\mathbf{n}(0)$  with the eigenvector corresponding to  $\Lambda_\gamma$ , which could be different for different phenotypes even if their eigenvalues are the same, leading to different relative proportions in the long time limit.

### The mean-field dynamics are not gradient-based

Here, we briefly prove the claim that the mean-field model cannot be written as a gradient-based dynamics  $\frac{d\mathbf{n}}{dt} = -\frac{\partial}{\partial \mathbf{n}} C(\mathbf{n})$  for some effective fitness function  $-C(\mathbf{n})$ . The proof uses the fact that the right-hand side of Eq (2) is not conservative, and hence cannot be written as such a gradient. Let the right-hand-side of Eq (2) be

$$Q_{k\gamma} \equiv w_k \left( 1 - \frac{\sum_{\gamma'} \mathbf{n}_{k\gamma'}}{w_k M} \right) \sum_i (b_i + \gamma_i) n_{i\gamma} - \sum_i w_i \left( 1 - \frac{\sum_{\gamma'} \mathbf{n}_{i\gamma'}}{w_i M} \right) \gamma_k n_{k\gamma}.$$



**Fig 6. Empirical demonstration that the low-density coefficient matrix  $M_\gamma$  has only one positive eigenvalue.** We numerically calculate the eigenvalues of the matrix  $M_\gamma$  by randomly drawing all birth and relocation rates from a uniform distribution in  $[0, 10)$ . Across 50000 samples shown, we find that only the largest eigenvalue of this matrix is positive and all the other eigenvalues are always negative.

<https://doi.org/10.1371/journal.pcbi.1009934.g006>

If this vector field were conservative, then it must be the case that

$$\frac{\partial Q_{k\gamma}}{\partial n_{jz}} - \frac{\partial Q_{jz}}{\partial n_{k\gamma}} = 0,$$

as if  $Q_{k\gamma} = \partial C(\mathbf{n}) / \partial n_{k\gamma}$  then this quantity would vanish by equality of mixed partial derivatives. We will demonstrate this is not the case. The derivative is

$$\begin{aligned} \frac{\partial Q_{k\gamma}}{\partial n_{jz}} &= -\frac{\delta_{jk}}{M} \sum_i (b_i + \gamma_i) n_{i\gamma} + \frac{\gamma_k n_{k\gamma}}{M} \\ &\quad + w_k \left( 1 - \frac{\sum_{\gamma'} n_{k\gamma'}}{w_k M} \right) (b_j + \gamma_j) \delta_{\gamma z} - \sum_i w_i \left( 1 - \frac{\sum_{\gamma'} n_{i\gamma'}}{w_i M} \right) \gamma_k \delta_{jk} \delta_{\gamma z}. \end{aligned}$$

Then, the difference becomes

$$\begin{aligned} \frac{\partial Q_{k\gamma}}{\partial n_{jz}} - \frac{\partial Q_{jz}}{\partial n_{k\gamma}} &= \frac{\delta_{jk}}{M} \sum_i (b_i + \gamma_i)(n_{iz} - n_{i\gamma}) + \left( \frac{\gamma_k n_{k\gamma} - \gamma_j n_{jz}}{M} \right) \\ &+ w_k \left( 1 - \frac{\sum_{\gamma'} n_{k\gamma'}}{w_k M} \right) (b_j + \gamma_j) \delta_{jz} - w_j \left( 1 - \frac{\sum_{\gamma'} n_{j\gamma'}}{w_j M} \right) (b_k + \gamma_k) \delta_{jz}. \end{aligned}$$

This quantity is non-zero in general, and hence the vector field  $Q_{k\gamma}$  is not conservative and cannot be written as a gradient of a cost function that predicts the steady-state configuration of the population. This holds true even in the low-density limit ( $M \gg \sum_{\gamma'} n_{k\gamma'}$ ), for which the difference reduces to

$$\frac{\partial Q_{k\gamma}}{\partial n_{jz}} - \frac{\partial Q_{jz}}{\partial n_{k\gamma}} = \left[ w_k (b_j + \gamma_j) - w_j (b_k + \gamma_k) \right] \delta_{jz}.$$

### Numerical solution and analysis of the 5-state mean-field model

With the simplified 5-state mean-field model we can evaluate the evolution of agents by numerically solving the system of differential equations in Eq (8). To do so, we discretize the continuous phenotype space  $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5)$  with  $\gamma_i \in [0, 10]$  into a space where each component can take discrete values from the set  $\{0, 2, 4, 6, 8, 10\}$ . This discretized phenotype space contains  $6^5 = 7776$  possible phenotypes.

We numerically solve Eq (9) with the ODE solver from Scipy package [66], using explicit Runge-Kutta method of order 5(4). In practice, we solve the full mean-field model Eq (2) for the fully coupled phenotype system, even in the low density limit. The number of agents is reset following Eq (10) after a duration of simulation time  $T = 0.3$ , repeated for 2000 iterations. A typical run takes around 2 hours to finish. Note that in all the numerical evaluations of the mean-field model, the birth rates are taken to be  $[0, 1, 2, 3, 5]$  except for Fig 3 where the change of  $b_5$  is explicitly investigated. This slightly different choice of  $b_5 = 5$  in the mean-field model compared to the  $b_5 = 4$  in the stochastic model is to compensate the minor quantitative difference of the critical value of  $b_5$  at which the transition from  $\langle \gamma_4 \rangle = \gamma_{\max}$  to  $\gamma_{\min}$  occurs, as shown in Fig 3A, upper panel.

The mean phenotype and its standard deviation can be calculated from the fraction of agents with different phenotypes

$$f_{\gamma}(t) = \frac{N_{\gamma}(t)}{\sum_{\gamma'} N_{\gamma'}(t)},$$

where  $N_{\gamma}(t) = \sum_i n_{i\gamma}(t)$  is the total number of agents with phenotype  $\gamma$ . Because there are no deaths in our model the numerator will never go to zero, and  $f_{\gamma}$  is therefore bounded between 0 and 1, so we can treat it as a probability and calculate the mean phenotype by

$$\langle \gamma(t) \rangle = \sum_{\gamma} \gamma f_{\gamma}(t),$$

for each component of the vector  $\gamma$ . We similarly estimate the standard deviation by

$$\sigma_{\gamma}(t) = \sqrt{\sum_{\gamma} \gamma^2 f_{\gamma}(t) - \langle \gamma(t) \rangle^2}.$$



## Supporting information

**S1 Fig. Different realizations of the resource distributions give consistent phenotype distributions.** In evaluating the phenotype distributions shown in the main text, we average over 20 trials of the evolutionary dynamics, but we keep the random realization of the resource distribution the same across these 20 trials. To demonstrate that the specific realization of resources throughout the environment does not strongly shape the ultimate phenotype distribution, we run the agent-based stochastic simulations for additional unique random realizations of the resources throughout the environment. **A-C.** Phenotype distribution at 20000 iterations, for 3 different realizations of resource levels drawn from a discrete uniform distribution. **D-F.** Phenotype distributions at 20000 iterations, for 3 different realizations drawn from an exponential distribution  $w_k \propto e^{-ck}$  with  $c = 0.5$ . In both cases, different realizations of the resource give us consistent phenotype distributions.

(TIF)

**S2 Fig. Maximum eigenvalue  $\Lambda_\gamma$  as an effective objective function.** As discussed in the main text, the maximum eigenvalue behaves as an effective objective function and predicts the transition of  $\gamma_4$  when  $b_5$  is varied. It also explains the large variability of the  $\gamma_4$  component, as the maximum eigenvalue is insensitive to  $\gamma_4$  when  $b_5 \approx B_{5c}$ . **A.** The maximum eigenvalue  $\Lambda_\gamma$  of the low-density coefficient matrix  $\mathbb{M}_\gamma$  predicts the most competitive foraging strategies  $\gamma_{\text{opt}}$ . Typically, the optimal values are  $\gamma_{1,2,3} = 10$  and  $\gamma_5 = 0$ , but the value of  $\gamma_4$  depends on the birth rate at resource level  $S_5$ ,  $b_5$ . When  $b_5$  is above the critical value  $B_{5c}$ , the phenotype with  $\gamma_4 \rightarrow \gamma_{\text{max}}$  has the largest maximum eigenvalue and thus are predicted to be the optimal; however when  $b_5$  is below  $B_{5c}$ , the phenotype with  $\gamma_4 \rightarrow \gamma_{\text{min}}$  is predicted to dominate the population instead. The same transition in the full agent-based stochastic simulation is shown in Fig 3. **B.** The derivatives of the eigenvalue predict a range of  $\gamma_4$  values yield similar maximum eigenvalues when  $b_5$  is close to  $B_{5c}$ . The effective objective function is insensitive to changes of  $\gamma_4$ , resulting in a large variability of  $\gamma_4$  as seen in the simulation results shown in Fig 2E and 2F. When the derivative is positive, larger  $\gamma_4$  leads to larger  $\lambda_\gamma$  and thus agents are better off leaving their current location to explore (“go”). When the derivative is negative, smaller  $\gamma_4$  leads to larger  $\lambda_\gamma$  and thus agents are better off exploiting their current location (“stay”).

(TIF)

**S3 Fig. Phenotype distributions across the transition points.** Changing the birth rate at the highest resource level  $b_5$  leads to a transition in the most competitive phenotypes, primarily in the fourth component  $\gamma_4$ , as shown in Fig 3B. Here we show the full phenotype distributions for some representative cases around and at this transition, for both discrete uniform and exponentially distributed resource levels. **A-C.** Phenotype distribution at 20000 iterations for a discrete uniform resource distribution, for  $b_5$  of 4, 5, and 6, respectively. **D-F.** Phenotype distribution at 20000 iterations for an exponential resource distribution  $c = 0.5$ , for  $b_5$  of 4, 10, and 16, respectively. In all the cases,  $b_1 = 0$ ,  $b_2 = 1$ ,  $b_3 = 2$ ,  $b_4 = 3$ .

(TIF)

**S4 Fig. Phenotype distributions of agent-based stochastic simulations at different initial densities.** In the full agent-based stochastic simulation we observed that there exists a range of initial densities with which different phenotypes coexist for an evolutionarily long time, at least up to 20000 iterations of the full stochastic simulations shown in Fig 4. Here we present the violin plots of the relocation rate distributions of the full agent-based stochastic simulations at these different initial densities, showing the coexistence of phenotypes at a range of initial densities, indicated by a large width of the distribution. At lower densities, coexistence occurs when  $b_5 \approx B_{5c}$ , while some degree of coexistence is observed for several initial densities

without tuning  $b_5$ . Note that the large degree of coexistence at 94% density is mostly due to the environment being so dense that opportunities for relocation and birth are extremely rare. (TIF)

**S5 Fig. Spatial covariance and the well-mixed assumption error at different initial density levels.** In deriving the 5-state mean-field model from the  $2d$  master equation model, we made a “well-mixed approximation” (Eq (6)). The error introduced by this assumption (the difference of left and right hand sides of Eq (6)) has a zero mean when averaged across sites and here we show how the standard deviation of this error changes over different initial densities. To assess the spatial variation neglected by our 5-state approximation, we evaluated the spatial cross-covariance of the agent numbers across different sites and then averaged across trials. This cross-covariance is localized around the nearest neighbor sites at low-to-medium densities, but weakens at higher densities. Averaging the covariance of the 4 nearest neighbors exhibits qualitatively similar behavior as the standard deviation of the well-mixed assumption error as the initial density changes. **A.** The spatial covariance of the number of agents at location  $(x, y)$  with the number of agents at sites  $(x + i, y + j)$ , averaged across all sites with periodic boundary conditions and then averaged over 20 trials of independent simulations. The center of the covariance map (the variance) is suppressed for visibility. **B.** Green line and circle markers show the averaged spatial covariance of the 4 nearest neighbors in panel A (i.e., the mean of the covariances at  $(1, 0)$ ,  $(-1, 0)$ ,  $(0, 1)$ , and  $(0, -1)$ ) at different initial density levels. Blue line and triangle markers show the standard deviation of the well-mixed assumption error (the difference of left and right hand sides of Eq (6)) across 20 trials at different initial density levels. Both quantities are small for very low or high densities, with peaks in a density range around 31% to 48%. (TIF)

**S6 Fig. The well-mixed assumption introduces a mismatch in timescales.** The full agent-based model has a population of agents navigating on a  $2d$  environment. For analytical tractability, this full model is reduced to a 5-state model that eliminates spatial dependencies. However, it is observed in Fig 2E and 2F that the full model and the reduced model have a qualitative mismatch in the order of magnitude of the convergence speed, as well as a quantitative mismatch in the value of the transition point of  $b_5$ . Here we show the comparison of the two models at a shorter timescale (700 iterations) with different  $b_5$  values (4 and 5). The results suggest that the well-mixed assumption introduces the discrepancy in timescales, as simulations of a stochastic 5-state model agree with the mean-field 5-state model. **A.** The mean relocation rates of the full  $2d$  stochastic simulation with  $b_5 = 4$  have not converged by 700 catastrophes. **B.** Agent-based simulations of a 5-state model with  $b_5 = 4$  show a faster convergence of  $\gamma_5$  to 0 and  $\gamma_{1,2,3}$  to high values near  $\gamma_{\max} = 10$ , within the same 700 catastrophes shown in A (note the different scales of the horizontal axes in Fig 2E and 2F). **C.** The 5-state mean-field model with  $b_5 = 5$  (which predicts  $\gamma_4 \rightarrow 10$  at long times) converges within 700 catastrophes, though the fate of  $\gamma_4$  is different from the stochastic 5-state model with  $b_5 = 4$ . **D.** The stochastic agent-based 5-state model with  $b_5 = 5$  largely converges within 700 iterations and agrees slightly better with the mean-field model’s prediction for  $\gamma_4$ . (TIF)

## Acknowledgments

We thank Ken Dill, Piotr Sokół, Jiahao Hu, and Wenbin Lu for helpful discussions, and the Python community for building the tools used in this work.

## Author Contributions

**Conceptualization:** Tong Liang, Braden A. W. Brinkman.

**Data curation:** Tong Liang.

**Formal analysis:** Tong Liang, Braden A. W. Brinkman.

**Funding acquisition:** Braden A. W. Brinkman.

**Investigation:** Tong Liang, Braden A. W. Brinkman.

**Methodology:** Tong Liang, Braden A. W. Brinkman.

**Project administration:** Braden A. W. Brinkman.

**Resources:** Braden A. W. Brinkman.

**Software:** Tong Liang.

**Supervision:** Braden A. W. Brinkman.

**Validation:** Tong Liang, Braden A. W. Brinkman.

**Visualization:** Tong Liang, Braden A. W. Brinkman.

**Writing – original draft:** Tong Liang, Braden A. W. Brinkman.

**Writing – review & editing:** Tong Liang, Braden A. W. Brinkman.

## References

1. Stephens DW, Brown JS, Ydenberg RC. Foraging: behavior and ecology. University of Chicago Press; 2008.
2. Marra RGPP. Birds of two worlds: the ecology and evolution of migration. JHU Press; 2005.
3. López-Cruz A, Sordillo A, Pokala N, Liu Q, McGrath PT, Bargmann CI. Parallel Multimodal Circuits Control an Innate Foraging Behavior. *Neuron*. 2019; 102(2):407–419.e8. <https://doi.org/10.1016/j.neuron.2019.01.053> PMID: 30824353
4. Doebeli M, Ispolatov Y, Simon B. Point of view: Towards a mechanistic foundation of evolutionary theory. *Elife*. 2017; 6:e23804. <https://doi.org/10.7554/eLife.23804> PMID: 28198700
5. Zylberberg J, DeWeese MR. How should prey animals respond to uncertain threats? *Frontiers in computational neuroscience*. 2011; 5:20. <https://doi.org/10.3389/fncom.2011.00020> PMID: 21559347
6. Kilpatrick ZP, Davidson JD, Hady AE. Normative theory of patch foraging decisions. *bioRxiv*. 2020;
7. Kilpatrick ZP, Davidson JD, Hady AE. Uncertainty drives deviations in normative foraging decision strategies. *bioRxiv*. 2021;
8. Bressloff PC. First-passage processes and the target-based accumulation of resources. *Phys Rev E*. 2021; 103:012101. <https://doi.org/10.1103/PhysRevE.103.012101> PMID: 33601498
9. Park IM, Pillow JW. Bayesian efficient coding. *BioRxiv*. 2017; p. 178418.
10. Młynarski W, Hledík M, Sokolowski TR, Tkačik G. Statistical analysis and optimality of neural systems. *Neuron*. 2021; 109(7):1227–1241. <https://doi.org/10.1016/j.neuron.2021.01.020> PMID: 33592180
11. Karpas ED, Shklarsh A, Schneidman E. Information socialtaxis and efficient collective behavior emerging in groups of information-seeking agents. *Proceedings of the National Academy of Sciences*. 2017; 114(22):5589–5594. <https://doi.org/10.1073/pnas.1618055114> PMID: 28507154
12. Laughlin SB. A Simple Coding Procedure Enhances a Neuron's Information Capacity. *Z Naturforsch*. 1981; 36c:910–912. <https://doi.org/10.1515/znc-1981-9-1040> PMID: 7303823
13. Bialek W, Owen WG. Temporal filtering in retinal bipolar cells. Elements of an optimal computation? *Biophysical journal*. 1990; 58(5):1227–1233. [https://doi.org/10.1016/S0006-3495\(90\)82463-2](https://doi.org/10.1016/S0006-3495(90)82463-2) PMID: 2291942
14. Van Steveninck RdR, Bialek W, Potters M, Carlson R. Statistical adaptation and optimal estimation in movement computation by the blowfly visual system. In: *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*. vol. 1. IEEE; 1994. p. 302–307.

15. Bautista LM, Tinbergen J, Wiersma P, Kacelnik A. Optimal foraging and beyond: how starlings cope with changes in food availability. *The American Naturalist*. 1998; 152(4):543–561. <https://doi.org/10.1086/286189> PMID: 18811363
16. Pérez-Escudero A, Rivera-Alba M, de Polavieja GG. Structure of deviations from optimality in biological systems. *Proceedings of the National Academy of Sciences*. 2009; 106(48):20544–20549. <https://doi.org/10.1073/pnas.0905336106> PMID: 19918070
17. Celani A, Vergassola M. Bacterial strategies for chemotaxis response. *Proceedings of the National Academy of Sciences*. 2010; 107(4):1391–1396. <https://doi.org/10.1073/pnas.0909673107> PMID: 20080704
18. Diedrichsen J, Shadmehr R, Ivry RB. The coordination of movement: optimal feedback control and beyond. *Trends in Cognitive Sciences*. 2010; 14(1):31–39. <https://doi.org/10.1016/j.tics.2009.11.004> PMID: 20005767
19. Rodrigues JV, Shakhnovich EI. Adaptation to mutational inactivation of an essential gene converges to an accessible suboptimal fitness peak. *eLife*. 2019; 8:e50509. <https://doi.org/10.7554/eLife.50509> PMID: 31573512
20. Powell S, Price SL, Kronauer DJC. Trait evolution is reversible, repeatable, and decoupled in the soldier caste of turtle ants. *Proceedings of the National Academy of Sciences*. 2020; 117(12):6608–6615. <https://doi.org/10.1073/pnas.1913750117> PMID: 32152103
21. Buskirk SW, Rokes AB, Lang GI. Adaptive evolution of nontransitive fitness in yeast. *eLife*. 2020; 9:e62238. <https://doi.org/10.7554/eLife.62238> PMID: 33372653
22. Calhoun AJ, Chalasani SH, Sharpee TO. Maximally informative foraging by *Caenorhabditis elegans*. *eLife*. 2014; 3:e04220. <https://doi.org/10.7554/eLife.04220> PMID: 25490069
23. Gillespie DT. Stochastic simulation of chemical kinetics. *Annu Rev Phys Chem*. 2007; 58:35–55. <https://doi.org/10.1146/annurev.physchem.58.032806.104637> PMID: 17037977
24. Goldenfeld N. *Lectures on Phase Transitions and the Renormalization Group*. Westview Press; 1992.
25. Assaf M, Meerson B. WKB theory of large deviations in stochastic populations. *Journal of Physics A: Mathematical and Theoretical*. 2017; 50(26):263001. <https://doi.org/10.1088/1751-8121/aa669a>
26. Lehman J, Clune J, Misevic D, Adami C, Altenberg L, Beaulieu J, et al. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial life*. 2020; 26(2):274–306. [https://doi.org/10.1162/artl\\_a\\_00319](https://doi.org/10.1162/artl_a_00319) PMID: 32271631
27. Mann RP. Evolution of heterogeneous perceptual limits and indifference in competitive foraging. *PLOS Computational Biology*. 2021; 17(2):1–20. <https://doi.org/10.1371/journal.pcbi.1008734> PMID: 33621223
28. Sutton RS, Barto AG. *Reinforcement learning: An introduction*. MIT press; 2018.
29. Dobbs IM, Molho I. Evolution and sub-optimal behaviour. *Journal of Evolutionary Economics*. 1999; 9(2):187–209. <https://doi.org/10.1007/s001910050080>
30. Martin TL, Huey RB. Why “Suboptimal” Is Optimal: Jensen’s Inequality and Ectotherm Thermal Preferences. *The American Naturalist*. 2008; 171(3):E102–E118. <https://doi.org/10.1086/527502> PMID: 18271721
31. Vercken E, Wellenreuther M, Svensson EI, Mauroy B. Don’t Fall Off the Adaptation Cliff: When Asymmetrical Fitness Selects for Suboptimal Traits. *PLOS ONE*. 2012; 7(4):1–9. <https://doi.org/10.1371/journal.pone.0034889> PMID: 22509364
32. Bressloff PC. Stochastic switching in biology: from genotype to phenotype. *Journal of Physics A: Mathematical and Theoretical*. 2017; 50(13):133001. <https://doi.org/10.1088/1751-8121/aa5db4>
33. Jabot F, Chave J. Analyzing tropical forest tree species abundance distributions using a nonneutral model and through approximate Bayesian inference. *The American Naturalist*. 2011; 178(2):E37–E47. <https://doi.org/10.1086/660829> PMID: 21750378
34. Haegeman B, Loreau M. A mathematical synthesis of niche and neutral theories in community ecology. *Journal of theoretical biology*. 2011; 269(1):150–165. <https://doi.org/10.1016/j.jtbi.2010.10.006> PMID: 20946903
35. Azae S, Suweis S, Grilli J, Volkov I, Banavar JR, Maritan A. Statistical mechanics of ecological systems: Neutral theory and beyond. *Rev Mod Phys*. 2016; 88:035003. <https://doi.org/10.1103/RevModPhys.88.035003>
36. Bazzani A, Castellani G, Giampieri E, Sala C. Master equation and relative species abundance distribution for Lotka-Volterra models of interacting ecological communities. *Master Equation and Relative Species Abundance Distribution for Lotka-Volterra Models of Interacting Ecological Communities*. 2016; p. 37–47. PMID: 29513351

37. Butler T, Reynolds D. Predator-prey quasicycles from a path-integral formalism. *Phys Rev E*. 2009; 79:032901. <https://doi.org/10.1103/PhysRevE.79.032901> PMID: 19392001
38. Butler T, Goldenfeld N. Robust ecological pattern formation induced by demographic noise. *Phys Rev E*. 2009; 80:030902. <https://doi.org/10.1103/PhysRevE.80.030902> PMID: 19905053
39. Shih HY, Goldenfeld N. Path-integral calculation for the emergence of rapid evolution from demographic stochasticity. *Phys Rev E*. 2014; 90:050702. <https://doi.org/10.1103/PhysRevE.90.050702> PMID: 25493725
40. Xue C, Goldenfeld N. Coevolution Maintains Diversity in the Stochastic “Kill the Winner” Model. *Phys Rev Lett*. 2017; 119:268101. <https://doi.org/10.1103/PhysRevLett.119.268101> PMID: 29328693
41. Dieckmann U, Law R. The dynamical theory of coevolution: a derivation from stochastic ecological processes. *Journal of mathematical biology*. 1996; 34(5):579–612. <https://doi.org/10.1007/BF02409751> PMID: 8691086
42. Dobramysl U, Mobilia M, Pleimling M, Täuber UC. Stochastic population dynamics in spatially extended predator–prey systems. *Journal of Physics A: Mathematical and Theoretical*. 2018; 51(6):063001. <https://doi.org/10.1088/1751-8121/aa95c7>
43. Wootton JT. Field parameterization and experimental test of the neutral theory of biodiversity. *Nature*. 2005; 433(7023):309–312. <https://doi.org/10.1038/nature03211> PMID: 15662423
44. Alonso D, Etienne RS, McKane AJ. The merits of neutral theory. *Trends in ecology & evolution*. 2006; 21(8):451–457. <https://doi.org/10.1016/j.tree.2006.03.019> PMID: 16766082
45. Allouche O, Kadmon R. Demographic analysis of Hubbell’s neutral theory of biodiversity. *Journal of Theoretical Biology*. 2009; 258(2):274–280. <https://doi.org/10.1016/j.jtbi.2008.12.024> PMID: 19168081
46. Agarwala EK, Chiel HJ, Thomas PJ. Pursuit of food versus pursuit of information in a Markovian perception–action loop model of foraging. *Journal of theoretical biology*. 2012; 304:235–272. <https://doi.org/10.1016/j.jtbi.2012.02.016> PMID: 22381540
47. Harpaz R, Schneidman E. Social interactions drive efficient foraging and income equality in groups of fish. *Elife*. 2020; 9:e56196. <https://doi.org/10.7554/eLife.56196> PMID: 32838839
48. Chow SS, Wilke CO, Ofria C, Lenski RE, Adami C. Adaptive radiation from resource competition in digital organisms. *Science*. 2004; 305(5680):84–86. <https://doi.org/10.1126/science.1096307> PMID: 15232105
49. McGill BJ, Brown JS. Evolutionary game theory and adaptive dynamics of continuous traits. *Annu Rev Ecol Evol Syst*. 2007; 38:403–435. <https://doi.org/10.1146/annurev.ecolsys.36.091704.175517>
50. Olson RS, Hintze A, Dyer FC, Knoester DB, Adami C. Predator confusion is sufficient to evolve swarming behaviour. *Journal of The Royal Society Interface*. 2013; 10(85):20130305. <https://doi.org/10.1098/rsif.2013.0305> PMID: 23740485
51. Adami C, Schossau J, Hintze A. Evolutionary game theory using agent-based methods. *Physics of life reviews*. 2016; 19:1–26. <https://doi.org/10.1016/j.plrev.2016.11.005> PMID: 27617905
52. Goldberg D. Genetic algorithms in search, optimization, and machine learning. Boston, MA: Addison-Wesley; 1989.
53. Thomas FI, Kregting LT, Badgley BD, Donahue MJ, Yund PO. Fertilization in a sea urchin is not only a water column process: effects of water flow on fertilization near a spawning female. *Marine Ecology Progress Series*. 2013; 494:231–240. <https://doi.org/10.3354/meps10601>
54. Murua H. Fish reproduction assortment: a wonderful diversity. *Environmental biology of fishes*. 2014; 97(3):329–333. <https://doi.org/10.1007/s10641-013-0154-2>
55. Alonzo SH, Stiver KA, Marsh-Rollo SE. Ovarian fluid allows directional cryptic female choice despite external fertilization. *Nature Communications*. 2016; 7(1):1–8. <https://doi.org/10.1038/ncomms12452> PMID: 27529581
56. Hussain YH, Guasto JS, Zimmer RK, Stocker R, Riffell JA. Sperm chemotaxis promotes individual fertilization success in sea urchins. *Journal of Experimental Biology*. 2016; 219(10):1458–1466. PMID: 26994183
57. Mery F, Burns JG. Behavioural plasticity: an interaction between evolution and experience. *Evolutionary Ecology*. 2010; 24(3):571–583. <https://doi.org/10.1007/s10682-009-9336-y>
58. Kadow ICG. State-dependent plasticity of innate behavior in fruit flies. *Current Opinion in Neurobiology*. 2019; 54:60–65. <https://doi.org/10.1016/j.conb.2018.08.014>
59. Robinson GE, Barron AB. Epigenetics and the evolution of instincts. *Science*. 2017; 356(6333):26–27. <https://doi.org/10.1126/science.aam6142> PMID: 28385970
60. Spoelstra K, Wikelski M, Daan S, Loudon ASI, Hau M. Natural selection against a circadian clock gene mutation in mice. *Proceedings of the National Academy of Sciences*. 2016; 113(3):686–691. <https://doi.org/10.1073/pnas.1516442113> PMID: 26715747

61. Deroulers C, Monasson R. Field-theoretic approach to metastability in the contact process. *Physical Review E*. 2004; 69(1):016126. <https://doi.org/10.1103/PhysRevE.69.016126> PMID: 14995686
62. Shaikh D, Rañó I. Braitenberg vehicles as computational tools for research in neuroscience. *Frontiers in Bioengineering and Biotechnology*. 2020; 8. <https://doi.org/10.3389/fbioe.2020.565963> PMID: 33042967
63. Dvoretzkii S, Gong Z, Gupta A, Parent J, Alicea B. Braitenberg Vehicles as Developmental Neurosimulation. arXiv preprint arXiv:200307689. 2020;.
64. Cueva CJ, Wei XX. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. In: *International Conference on Learning Representations*; 2018. Available from: <https://openreview.net/forum?id=B17JT0e0->.
65. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, et al. Array programming with NumPy. *Nature*. 2020; 585:357–362. <https://doi.org/10.1038/s41586-020-2649-2> PMID: 32939066
66. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*. 2020; 17:261–272. <https://doi.org/10.1038/s41592-019-0686-2> PMID: 32015543
67. Hunter JD. Matplotlib: A 2D graphics environment. *IEEE Annals of the History of Computing*. 2007; 9(03):90–95.
68. Bonabeau E. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the national academy of sciences*. 2002; 99(suppl 3):7280–7287. <https://doi.org/10.1073/pnas.082080899> PMID: 12011407
69. Horn RA, Johnson CR. *Matrix Analysis*. 2nd ed. USA: Cambridge University Press; 2012.