

Exploration-exploitation tradeoffs dictate the optimal distributions of phenotypes for populations subject to fitness fluctuations

Andrea De Martino,^{1,2} Thomas Gueudré,³ and Mattia Miotto⁴

¹*Soft and Living Matter Laboratory, CNR-NANOTEC, 00185 Rome, Italy*

²*Italian Institute for Genomic Medicine, 10126 Turin, Italy*

³*DISAT, Politecnico di Torino, 10129 Turin, Italy*

⁴*Department of Physics, Sapienza University, 00185 Rome, Italy*



(Received 1 October 2018; published 17 January 2019)

We study a minimal model for the growth of a phenotypically heterogeneous population of cells subject to a fluctuating environment in which they can replicate (by exploiting available resources) and modify their phenotype within a given landscape (thereby exploring novel configurations). The model displays an exploration-exploitation trade-off whose specifics depend on the statistics of the environment. Most notably, the phenotypic distribution corresponding to maximum population fitness (i.e., growth rate) requires a nonzero exploration rate when the magnitude of environmental fluctuations changes randomly over time, while a purely exploitative strategy turns out to be optimal in two-state environments, independently of the statistics of switching times. We obtain analytical insight into the limiting cases of very fast and very slow exploration rates by directly linking population growth to the features of the environment.

DOI: [10.1103/PhysRevE.99.012417](https://doi.org/10.1103/PhysRevE.99.012417)

I. INTRODUCTION

The exploration-exploitation trade-off scenario constitutes a paradigm for the optimal balance between the risky search for new resources and the safe exploitation of available ones that occurs in a variety of systems [1]. As a generic example, one may consider a population occupying a patch of space in a land in which the availability of an essential resource fluctuates in time and across patches. By remaining on a certain patch for a sufficiently long time the population will be able to exploit the resource available in that patch to the fullest. That benefit, however, has to be weighed against the cost of the opportunities that are missed by not searching for a better patch. The central question concerns which balance of exploitation (stay) and exploration (go) will provide the population with the highest fitness (e.g., the fastest growth rate) in the long run. The optimal strategy is obviously interlocked with details such as the statistics of resources and can be challenging to analyze at a quantitative level [2–4]. Still, fitness maximization is very often found to require a nonzero exploration rate.

An especially significant effort to understand this trade-off is ongoing for biological systems, as seen, e.g., in the recent interest about the “ecology of cancer growth” [5,6] (the strikingly diverse distributions of cell strains observed throughout different types of cancers) and its relationship to the timing of drug administration [6]. Microbial systems have also been a natural testing ground for the exploration-exploitation scenario for many years. It is empirically known that, in fluctuating environments, microbes tend to display a high degree of phenotypic heterogeneity driven by stochasticity in the regulation of gene expression and metabolism [7–12]. The ability to explore the space of allowed phenotypes ultimately provides an effective route to hedge against

environmental noise [13,14], favoring, e.g., the persistence of a subpopulation of resistant but slow-growing bacteria within a population subject to high doses of antibiotics [15,16]. Starting with Ref. [17], several mathematical models have shown that switching between different phenotypes at the individual cell level can be advantageous in rapidly changing conditions, depending essentially on (1) the statistics of environmental fluctuations and (2) the specific coupling between the environment and the allowed phenotypes [18–28]. Such models capture the physical and mathematical complexity of these systems starting from minimal assumptions about the environment and/or the space of feasible phenotypes. In more structured cases, the spectrum of viable behaviors appears to be even richer [29].

Here, inspired by recent work on single-cell physiology [30] and by the growth-entropy balance that appears to underlie part of the empirical observations [31], we characterize the exploration-exploitation trade-off in a model for the growth of a phenotypically heterogeneous population in a fluctuating environment. In short, we assume that each phenotype is represented by an intrinsic or constitutive growth rate and that the landscape of phenotypes accessible to cells is described by a given probability distribution. Over time, cells modify their phenotype due, e.g., to stochastic fluctuations in intracellular composition or regulatory processes that effectively cause cells to perform random walks in the phenotypic landscape (the exploration part). In turn, the cellular replication rate is determined by the coupling to an externally varying environment. While fast phenotypes are in principle favored (the exploitation part), the environment is subject to fluctuations that can punish them (as, e.g., in Ref. [15]). In such conditions, the balance between exploration of the phenotypic space and exploitation of fast phenotypes ultimately controls both the

overall fitness of the population and its structure (i.e., how individuals distribute over accessible phenotypes).

We show that the optimal evolutionary strategy (yielding maximum fitness for the population) can indeed require a nonzero exploration rate as suggested by the general explore-exploit paradigm. The gain due to exploration is particularly marked in the most unpredictable environments. On the other hand, in presence of more regular scenarios (e.g., periodic changes), an optimal population will adopt simpler strategies, such as maintaining two phenotypically distinct populations.

Our analysis will focus on universal observables, relying on numerics for the general case. The limiting cases of very fast and very slow search rates will instead be characterized by approximate analytical arguments.

II. MODEL DEFINITIONS

We consider a population of cells evolving in time. The phenotype of each cell is assumed to be fully characterized by a single variable λ , which we call the “constitutive replication rate” (CRR), taking on values in $[0, \lambda_{\max}]$. For sakes of simplicity, different values of λ will effectively correspond to different cellular phenotypes. To account for the fact that some phenotypes might be easier to attain than others, the space of allowed phenotypes is assumed to be described by a probability density $q(\lambda)$ such that $q(\lambda)d\lambda$ represents the fraction of phenotypes with CRR between λ and $\lambda + d\lambda$. The density of cells having CRR in $[\lambda, \lambda + d\lambda]$ at time t is instead denoted by $n(\lambda, t)$. In turn, $N(t) = \int n(\lambda, t) d\lambda$ represents the total number of cells in the population at time t . Following, e.g., Ref. [31], we assume that n changes due to (a) replication events and (b) diffusion in the phenotypic space, whereby cells change their CRR from λ to λ' . If the rate of the latter process is given by $W(\lambda \rightarrow \lambda')$, $n(\lambda, t)$ evolves according to

$$\frac{dn(\lambda, t)}{dt} = f(\lambda, t)n(\lambda, t) + \int [W(\lambda' \rightarrow \lambda)n(\lambda', t) - W(\lambda \rightarrow \lambda')n(\lambda, t)] d\lambda', \quad (1)$$

where f denotes the instantaneous replication rate (IRR) of cells with CRR λ . To couple the system to an external environment, we assume that the IRR depends both on the CRR λ and on the state of an exogenously varying medium, which, for sake of simplicity, will be described by the single time-dependent variable x . To focus on a relevant case, we consider a fluctuating environment in which x describes, in rough terms, the threshold fitness for replication under randomly occurring shocks. This corresponds to the choice

$$f(\lambda, t) = \begin{cases} \lambda & \text{if } \lambda \leq x(t) \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

according to which cells with CRR smaller than $x(t)$ can replicate at time t , while replication is inhibited for the others.

To study the impact of randomness in the environment, we look at various scenarios, ranging from the most predictable (switching periodically between two fixed states) to the most random (switching after a random time and to a random value). More specifically, the threshold x will fluctuate in time by switching between the value $x = \lambda_{\max}$, in which case all cells in the population can replicate, and a value $x = \lambda^* <$

λ_{\max} , in which case replication can take place only for cells with $\lambda \leq \lambda^*$. We consider two choices for λ^* . In the first case, λ^* is a constant kept fixed throughout the dynamics, so that x takes the values λ^* and λ_{\max} alternately, leading to a two-state environment (“const- x ” case). In the second case, λ^* is sampled independently at every switch from a uniform distribution on the interval $[x_{\min}, \lambda_{\max}]$, leading to an environment with a continuum of states (“rand- x ” case). For simplicity, we set $\lambda^* = x_{\min}$ in the const- x environment. Switches from the nonselective environment where all cells replicate to the selective one where only some do ($ns \rightarrow s$) and vice versa ($s \rightarrow ns$) are assumed to occur either periodically, i.e., after fixed times ω_{ns} and ω_s respectively (“const- t ” case) or at exponentially distributed random times with means equal to ω_{ns} or ω_s respectively (“rand- t ” case). (We, however, expect all our results to be qualitatively robust to changes in the distributions from which times and thresholds are drawn.)

Ultimately, for the process $x(t)$ we shall consider all possible mixtures of the above recipes for the threshold x and the switching times (i.e., const- t , const- x ; rand- t , const- x , etc.). In what follows, we begin by analyzing the simpler case of symmetric environment with $\omega_{ns} = \omega_s = \omega$, representative examples of which are sketched in Fig. 1. The asymmetric case with $\omega_{ns} \neq \omega_s$ will be dealt with in Sec. III E.

Introducing the population density

$$p(\lambda, t) = \frac{n(\lambda, t)}{N(t)}, \quad (3)$$

we recast Eq. (1) as

$$\begin{aligned} \frac{dp(\lambda, t)}{dt} &= [f(\lambda, t) - \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t)]p(\lambda, t) \\ &+ \int [W(\lambda' \rightarrow \lambda)p(\lambda', t) - W(\lambda \rightarrow \lambda')p(\lambda, t)] d\lambda' \end{aligned} \quad (4)$$

with

$$\mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) = \int_0^{\lambda_{\max}} f(\lambda, t)p(\lambda, t) d\lambda. \quad (5)$$

Furthermore, we assume that transition rates satisfy a detailed-balance condition of the form

$$W(\lambda \rightarrow \lambda')q(\lambda) = W(\lambda' \rightarrow \lambda)q(\lambda'), \quad (6)$$

with $q(\lambda)$ the density of phenotypes, and introduce the mean waiting time τ characterizing transitions via

$$\int W(\lambda \rightarrow \lambda') d\lambda' = \frac{1}{\tau}. \quad (7)$$

We further assume that only transitions from phenotype λ to phenotypes $\lambda \pm \delta\lambda$ are allowed, with equal probability and small $\delta\lambda$ (“diffusive transition kernel”). This choice provides the most natural route to model the effects induced at phenotypic level by small random fluctuations in intracellular composition, as they are unlikely to cause major gains or losses in terms of CRR. One easily shows (see Appendix and Ref. [31]) that Eq. (4) in this case can be approximated with

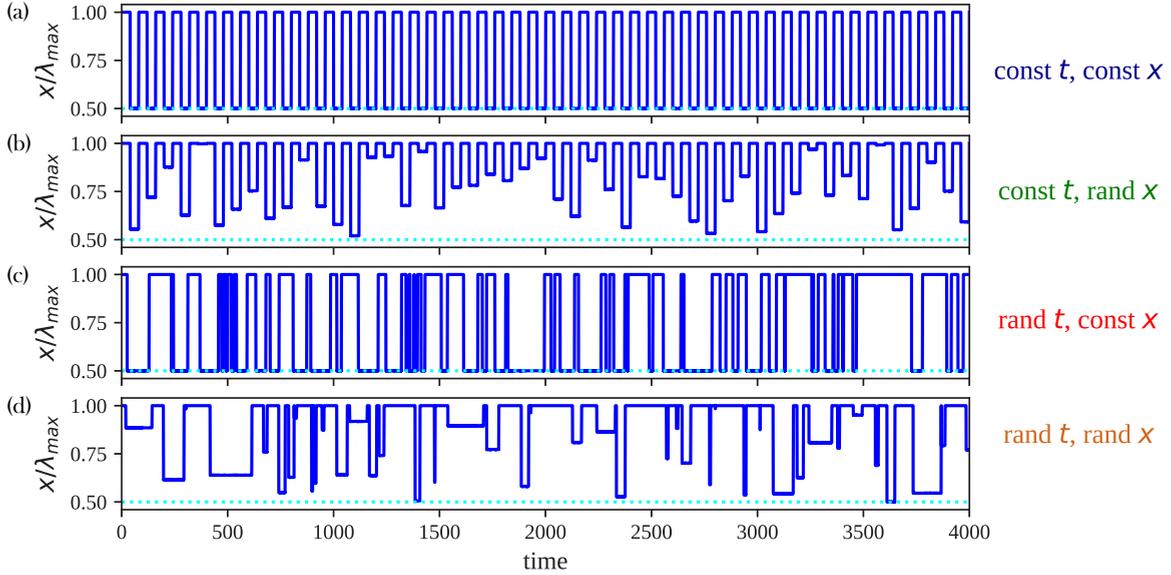


FIG. 1. Representative behavior of the threshold x as a function of time [see Eq. (2)] in the four environments we consider: (a) a periodic two-state environment where x switches (in this case) between the values λ_{\max} and $x_{\min} = \lambda_{\max}/2$; (b) a periodically switching environment where x takes on random values drawn uniformly from $[x_{\min}, \lambda_{\max}]$; (c) a two-state environment where switches occur at exponentially distributed random times; and (d) an environment where x behaves as in panel (b) but in which switches occur at exponentially distributed random times. In this example, the characteristic switching times ω_{ns} and ω_s are taken to be equal and fixed to 40 (a.u.).

the nonlinear Fokker-Planck equation

$$\frac{dp(\lambda, t)}{dt} = [f(\lambda, t) - \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t)]p(\lambda, t) + D \left(\frac{\partial^2 p(\lambda, t)}{\partial \lambda^2} - \frac{\partial}{\partial \lambda} \left\{ p(\lambda, t) \frac{\partial}{\partial \lambda} [\ln q(\lambda)] \right\} \right), \quad (8)$$

where $D = \frac{(\delta\lambda)^2}{2\tau}$ is the diffusion coefficient in the phenotypic space.

We finally have to specify a form for the phenotypic landscape $q(\lambda)$. To focus on a realistic case, we set

$$q(\lambda) = \frac{a+1}{\lambda_{\max}} \left(1 - \frac{\lambda}{\lambda_{\max}} \right)^a, \quad (9)$$

where the exponent $a \geq 0$ modulates the steepness of $q(\lambda)$. In short, the larger a , the more heterogeneous the landscape, with slow phenotypes being increasingly more frequent than fast ones as a increases. The above choice is based on recent studies showing that functions like (9) describe the CRR landscape underlying genome-scale models of bacterial metabolic networks, with values of a extracted from genome-scale models of *E. coli* lying between 200 and 300 depending on the specifics of the environment [31,32]. To focus on tractable extremes, we shall consider explicitly the cases $a=0$ [uniform $q(\lambda)$] and $a=20$ [strongly heterogeneous $q(\lambda)$].

The setup just described generalizes that considered in Refs. [31,33] to the case in which the instantaneous replication rate f depends on the coupling of cells to a fluctuating environment. The structure of a population governed by (8) emerges from the balance between the term that rewards fast-growing states (which are, however, sensitive to environmental shocks) and the diffusion term favoring states with larger entropy in the phenotypic space (but slower replication rates).

In the following, we characterize the above setting from the viewpoints of

(1) How the interplay between replication and diffusion (i.e., the trade-off between exploration and exploitation) affects the growth rate of the population as a whole and

(2) The emergent asymptotic structure of the population, i.e., how cells distribute over the one-dimensional phenotypic space $[0, \lambda_{\max}]$ at long times.

It is important to note that, in symmetric environments, two different timescales rule the time evolution of $p(\lambda, t)$: the mean switching time between different environments (ω) and the mean time to transition between different phenotypes (τ). The latter is inversely proportional to the diffusion constant D . The system's behavior is ultimately modulated by the ratio ω/τ . To explore the full range of this ratio, it is convenient to fix one timescale, e.g., ω , and use the other (i.e., D) as a control parameter. The limiting cases $\omega \ll \tau$ (in which exploration occurs on much longer timescales than exploitation) and $\omega \gg \tau$ (in which exploration occurs on much shorter timescales than exploitation) correspond to $D \rightarrow 0$ and $D \gg 1$, and we shall refer to these as the ‘‘exploitation’’ and ‘‘exploration’’ limits, respectively.

III. RESULTS

A. Dynamical patterns of population structure under symmetric switching

The nonlinear Fokker-Planck equation (8) can be solved numerically for any choice of the environment, of the diffusion coefficient and of the prior phenotypic density $q(\lambda)$. After a short transient, $p(\lambda, t)$ appears to settle in qualitative robust, environment-dependent patterns, a sample of which is shown in Fig. 2. Different types of distributions emerge across the various environments, including bimodal distributions in

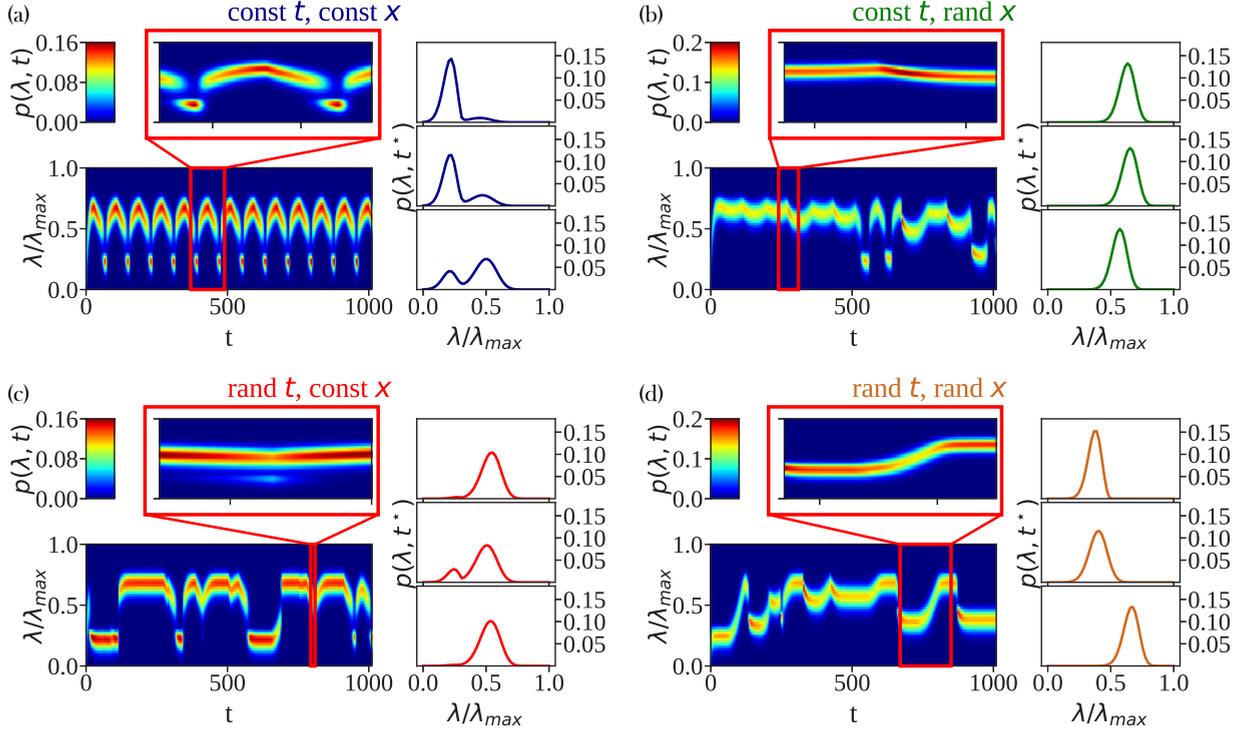


FIG. 2. Color maps showing representative probability densities $p(\lambda, t)$ obtained by solving Eq. (8) numerically in the different kinds of environment described in Fig. 1. Panels to the right of each map depict the density profile at different time points within the zoomed-in region, at time increasing from top to bottom. Results are shown for (a) const- t and const- x environment, (b) const- t and rand- x environment, (c) rand- t and const- x environment, and (d) rand- t and rand- x environment. Parameter values: $a = 20$, $x = 0.3\lambda_{\max}$, $D = 10^{-3}$.

which most of the population occupies the two peaks alternately [Fig. 2(a)] or in which one peak always dominates over the other [Fig. 2(c)], unimodal distributions with fluctuating positions [Fig. 2(b)], and unimodal distributions in which peaks drift in a specific direction [Fig. 2(d)]. While all of these can occur in every type of environment, both their frequency of occurrence and the relative intensities of the peaks appear to be strongly environment-dependent.

Such patterns provide hints about the way in which the population copes with environmental fluctuations. An important feature observed from data is that, independently of whether switches occur periodically or randomly, adaptation to two-state environments (const- x) is achieved more efficiently by structuring the population in a bimodal form, while complex environments (rand- x) favor unimodal distributions. We shall see in the following that such a scenario is indeed correct even asymptotically, although it can be modulated by the strength of diffusion.

B. Population growth rate and statistics at long times under symmetric switching

As we are mostly interested in understanding how the system behaves in the long-time limit, we focus on the long-term population structure as well as on the growth rate

$$\Lambda \equiv \lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{N(t)}{N(0)} \quad (10)$$

$$= \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t [\mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t')] dt'. \quad (11)$$

[The second equality follows directly from Eq. (1) and from the fact that $N(t) = \int n(\lambda, t) d\lambda$.] Λ will be used as a proxy for the long-term evolutionary success of the population. Figure 3 shows, for all environments, the stationary probability distributions $p(\lambda)$ obtained by averaging over time after Λ has reached its stationary value, for representative values of the parameters (in particular for $x_{\min} = 0.3\lambda_{\max}$, describing a strong negative perturbation which can be evaded only by cells whose CRR is at most 30% of the maximum), different values of D , and for $a = 0$ (corresponding to a uniform phenotypic landscape, top panels) and $a = 20$ (a strongly heterogeneous landscape with a predominance of slow growing states, bottom panels).

Generically, at sufficiently small values of D , phenotypes tend to concentrate close to λ_{\max} [see Figs. 3(a) and 3(b) and Figs. 3(d) and 3(e)]. This situation reproduces the “exploitation” limit $D \rightarrow 0$, where (8) reduces to the replicator dynamics

$$\frac{dp(\lambda, t)}{dt} = [f(\lambda, t) - \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t)]p(\lambda, t). \quad (12)$$

A population whose phenotypic diffusion occurs on exceedingly long timescales (compared to those characterizing environmental fluctuations) can only grow exploiting resources available from the environment and is therefore maximally sensitive to environment-derived shocks. In such a case, the population growth rate is significantly smaller than λ_{\max} [see Figs. 3(c) and 3(f), due to the growth-curbing effect of environmental fluctuations. (We shall analyze this limit at quantitative level in the following.)

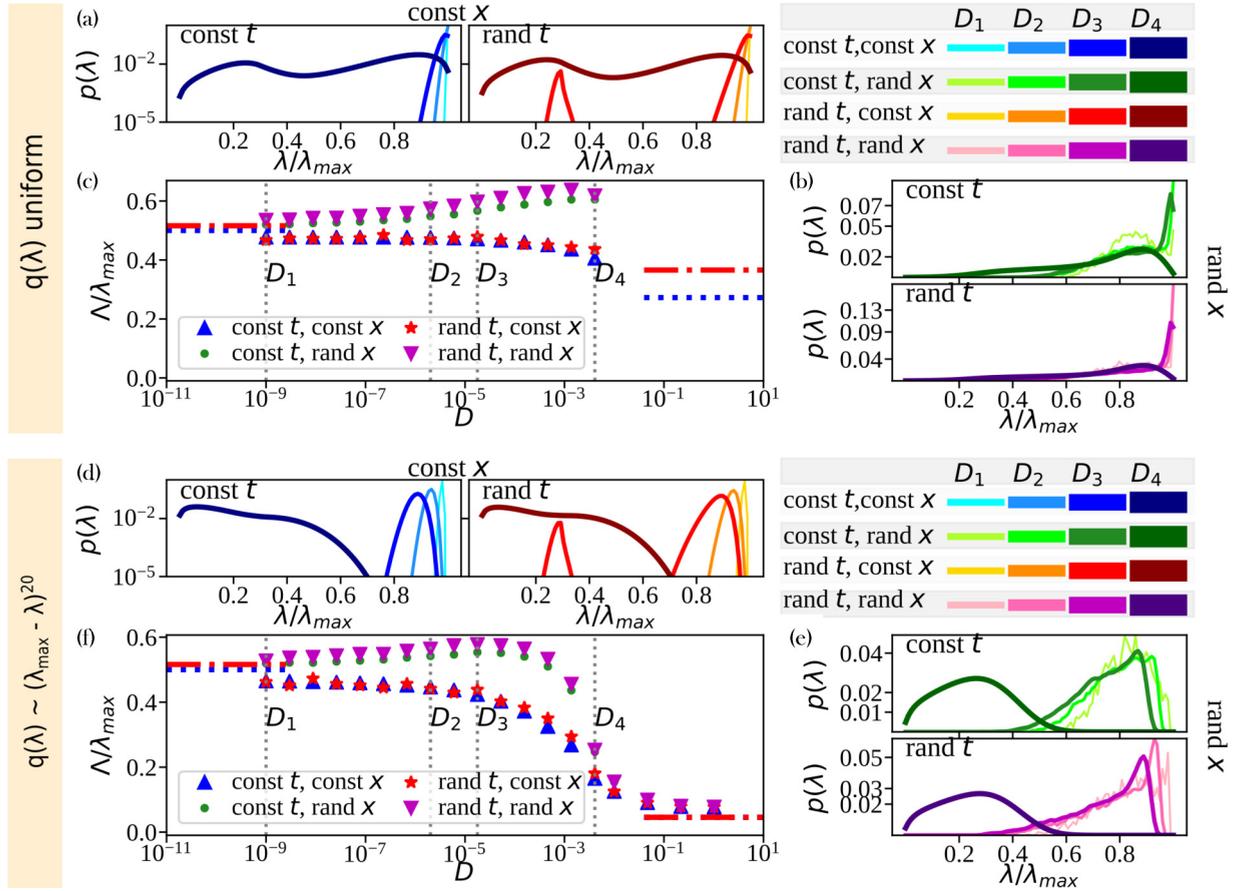


FIG. 3. (a), (b) Asymptotic, time-averaged phenotypic distributions obtained for a population evolving according to Eq. (8) with a diffusive kernel in a uniform background phenotypic landscape $q(\lambda)$ in const- x (a) and rand- x (b) environments for different time scenarios and values of D , marked by increasing color shades and line widths. (c) Asymptotic population growth rate Λ (in units of λ_{max}) as a function of D for the four types of environment. Vertical dotted lines mark the values of D studied in panels (a) and (b). Horizontal lines at small and large D stand for the analytical estimates for Λ obtained in the const- x regime [dotted blue line, Eqs. (24) for small D and (33) for large D] and the rand- x regime [dot-dashed red line, Eqs. (29) for small D and (33) for large D], respectively. (d)–(f) Same as panels (a)–(c) but with $q(\lambda)$ as in Eq. (9) (with $a = 20$) rather than uniform. Displayed curves are averaged over 100 independent realizations of the dynamics performed with $x_{min} = 0.3\lambda_{max}$.

Upon increasing D (and therefore the relevance of diffusion in the phenotypic space), distributions start to acquire nontrivial traits, including bimodality [see Figs. 3(a) and 3(d)] and extended tails [see Figs. 3(b) and 3(e)]. The population growth rate Λ then increases with D with respect to the small-diffusion limit in complex (rand- x) environments, where the population structure develops tails. In such cases, Λ has a well-defined maximum at a specific value of D (which depends, as in Ref. [2], on the characteristic time of environmental switches), marking the existence of an optimal trade-off between diffusion (exploration) and growth (exploitation) in the given environment. On the other hand, the population growth rate decreases continuously with D , albeit slowly, in the simpler two-state (const- x) environments, implying that any amount of exploration is detrimental to fitness in such contexts.

When diffusion dominates the dynamics (larger values of D), Λ appears to drop rapidly in all environments. In such a case, which is close to the purely “exploration” limit $D \rightarrow \infty$ that is analyzed in detail below, cells explore the phenotypic

space very efficiently, continuously redistributing their CRR among allowed states. The asymptotic behavior is hence dominated by the background provided by $q(\lambda)$. Indeed, the phenotypic distribution evolves towards its stationary limit $q(\lambda)$ due to the detailed balance constraint (6).

These results suggest that phenotypic diffusion can indeed be tuned to cope optimally with environmental fluctuations so as to ensure a significant gain in terms of fitness, provided the selective threshold of the environment changes randomly over time. In such a case, the fitness advantage appears to be slightly more marked when a is smaller. Still, the qualitative scenario just described is robust to changes in a . Correspondingly, the population structures into an extended unimodal distribution of phenotypes. On the other hand, in an environment fluctuating between two well-defined states, bimodal phenotypic distributions occur but exploration does not appear to provide a significant fitness advantage.

Note that a similar qualitative scenario for Λ is obtained for weaker environmental perturbations (i.e., larger x_{min}), the main effect induced by increasing x_{min} being (expectedly) that

of reducing the gap in Λ as a function of D between const- x and rand- x environments without modifying the overall behavior of individual cases.

C. Exploitation limit (case of symmetric switching)

To characterize our model in greater detail, it is convenient to focus on its limiting behaviors starting from the case $D \rightarrow 0$ (i.e., exploitation much faster than exploration), in which (8) reduces to (12). Here the population is dominated by the cells carrying the largest CRR. Intuitively, though, the coupling to the environment limits the reproductive efficiency of fast-growing phenotypes and ultimately introduces cutoffs to the CRR that are represented in the population. The parameter controlling this effect is x_{\min} . Numerical results indeed show (see Fig. 4) that, while the statistics of switching times does not appear to qualitatively influence the long-time limit, const- x and rand- x regimes produce qualitatively different asymptotics for the phenotypic distribution depending on whether $x_{\min} < \lambda_{\max}/2$ [Fig. 4(a)] or $x_{\min} > \lambda_{\max}/2$ [Fig. 4(b)]. Specifically, for $x_{\min} > \lambda_{\max}/2$ (weaker perturbation) the population concentrates around $\lambda = x_{\min}$ in every environment with an overall fitness $\Lambda > \lambda_{\max}/2$. For $x_{\min} < \lambda_{\max}/2$, instead, $p(\lambda)$ displays a peak at $\lambda = \lambda_{\max}$ in const- x environments, while an extended set of phenotypes is represented in the population when x is random and uniform. In both cases, the population growth rate Λ settles close to $\lambda_{\max}/2$. In other words, more efficient phenotypes appear to dominate the population when the perturbation is stronger and the overall growth rate is slower, while a weaker perturbation leading to a larger population growth rate seems to select for less efficient phenotypes.

A key observation to understand these results is that, independently of whether environmental switches occur at fixed times or at random times, in a symmetric environment with $\omega_{\text{ns}} = \omega_{\text{s}}$ cells will spend on average half the time in the “favorable,” not selective environment with $x = \lambda_{\max}$ and the other half in the selective environment with $x < \lambda_{\max}$. The statistics of switching times should therefore not be expected to influence outcomes at least as long as averages are concerned. On the other hand, because x is a random variable, the IRR f [Eq. (2)] will also be randomly fluctuating in time, with a mean value given by

$$\mathbb{E}_t f(\lambda, t) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(\lambda, t') dt'. \quad (13)$$

It is now convenient to discuss the const- x and rand- x cases separately.

1. Const- x (two-state) environments

For the “const- x ” case (two-state environment with x oscillating between λ_{\max} and a constant value x_{\min}), f will equal λ at all times if $\lambda < x_{\min}$, while for $\lambda > x_{\min}$ it will be equal to λ for approximately half the time and to zero for the other half. This implies that

$$\mathbb{E}_t f(\lambda, t) \simeq \begin{cases} \lambda & \text{if } \lambda < x_{\min} \\ \lambda/2 & \text{if } \lambda \geq x_{\min} \end{cases}. \quad (14)$$

The mean IRR therefore displays a discontinuity at the threshold x_{\min} , and the value of λ for which it attains a maximum

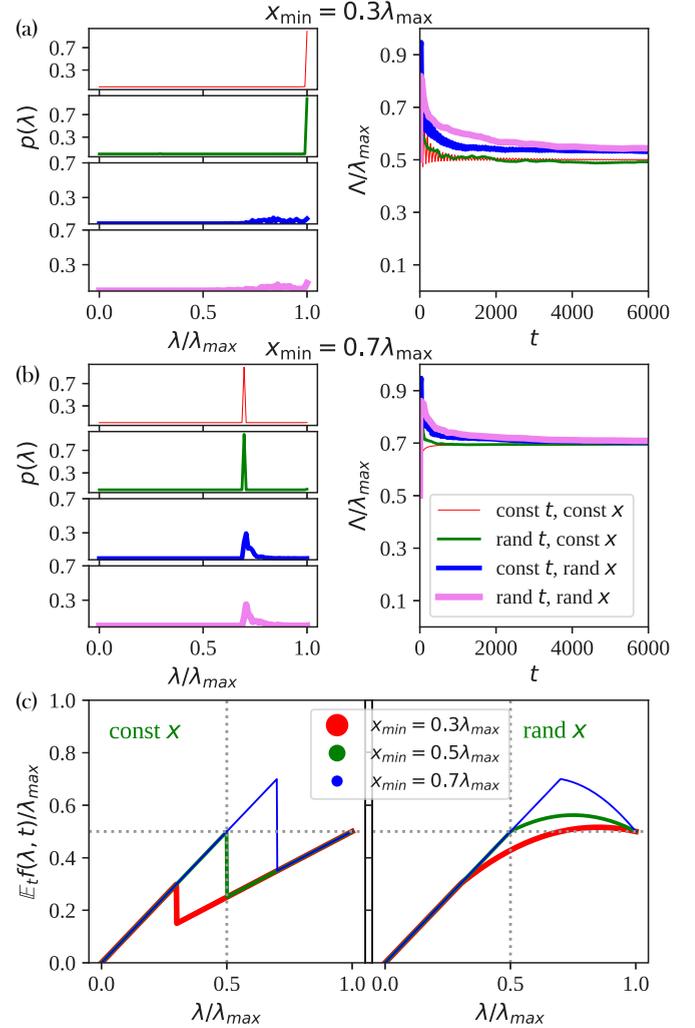


FIG. 4. (a), (b) Long-time phenotypic distributions (left) and time evolution of the population growth rate Λ (right) in the absence of diffusion in the different environments (represented by different colors and line widths) and for $x = 0.3\lambda_{\max}$ (a) and $x = 0.7\lambda_{\max}$ (b). In the former case ($x < \lambda_{\max}/2$), the distribution can achieve the highest possible CRR. In turn, the long-term fitness Λ sets around $\lambda_{\max}/2$. In the latter case ($x > \lambda_{\max}/2$), the distribution peaks around the threshold CRR, while the population achieves a growth rate Λ larger than $\lambda_{\max}/2$. Curves are averaged over 100 realizations of the dynamics. (c) Time-averaged mean $\mathbb{E}_t f(\lambda, t)$ of f as a function of the CRR for const- x (left) and rand- x (right) environments and for three different values of x_{\min} . One sees that the position of the maximum depends both on the chosen threshold and on the specific environment.

depends on the value of x_{\min} [see Fig. 4(c), left panel]. In specific, for $x_{\min} > \lambda_{\max}/2$ (resp. $x_{\min} < \lambda_{\max}/2$), the mean IRR has a maximum for $\lambda = x_{\min}$ (resp. $\lambda = \lambda_{\max}$), where $\mathbb{E}_t f(\lambda, t) = x_{\min}$ [resp. $\mathbb{E}_t f(\lambda, t) = \lambda_{\max}/2$]. Hence, at long times, we expect the population to grow at the fastest IRR achievable, with a phenotypic distribution $p(\lambda)$ peaked at $\lambda = x_{\min}$ (resp. $\lambda = \lambda_{\max}$) for $x_{\min} > \lambda_{\max}/2$ (resp. $x_{\min} < \lambda_{\max}/2$). This is in agreement with the numerical evidence shown in Figs. 4(a) and 4(b) (as well as in Fig. 3) for const- x environments.

Based on the above reasoning we can approximate $p(\lambda, t)$ with the bimodal function

$$p(\lambda, t) \simeq \alpha(t)\delta(\lambda - \lambda^*) + (1 - \alpha(t))\delta(\lambda - \lambda_{\max}), \quad (15)$$

with $0 \leq \alpha(t) \leq 1$ a time-dependent coefficient quantifying the fraction of cells with CRR equal to $\lambda^* < \lambda_{\max}$. (For sakes of simplicity, we shall henceforth omit to indicate explicitly the dependence of α on time.) We can then use (12), which in discrete time takes the form

$$p(\lambda, t + \delta t) \simeq \{1 + [f(\lambda, t) - \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t)]\delta t\} p(\lambda, t) \quad (16)$$

to evolve the above ansatz for small time intervals δt during which the environment does not change.

In nonselective conditions ($x = \lambda_{\max}$), one can use the fact that

$$\mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) = \alpha\lambda^* + (1 - \alpha)\lambda_{\max} \quad (17)$$

to arrive at

$$p(\lambda, t + \delta t) \simeq (\alpha - \delta\alpha_{\text{ns}})\delta(\lambda - \lambda^*) + (1 - \alpha + \delta\alpha_{\text{ns}})\delta(\lambda - \lambda_{\max}), \quad (18)$$

where $\delta\alpha_{\text{ns}} = (\lambda_{\max} - \lambda^*)(1 - \alpha)\delta t$.

In a selective environment ($x = \lambda^*$), instead,

$$\mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) = \alpha\lambda^*, \quad (19)$$

and one finds

$$p(\lambda, t + \delta t) \simeq (\alpha + \delta\alpha_s)\delta(\lambda - \lambda^*) + (1 - \alpha - \delta\alpha_s)\delta(\lambda - \lambda_{\max}), \quad (20)$$

with $\delta\alpha_s = \lambda^*(1 - \alpha)\delta t$.

This shows that, at every switch, the population distribution will tend to shift from one threshold to the other, but the speed with which the two peaks grow or shrink is different. In particular, one has

$$\frac{\delta\alpha_{\text{ns}}}{\delta\alpha_s} = \frac{(\lambda_{\max} - \lambda^*)}{\lambda^*}. \quad (21)$$

This implies that $\delta\alpha_{\text{ns}} < \delta\alpha_s$ for $\lambda^* > \lambda_{\max}/2$. Hence the peak growing at speed $\delta\alpha_s$ is favored, and the probability density will peak around λ^* in the long run. On the other hand, $\delta\alpha_{\text{ns}} > \delta\alpha_s$ when $\lambda^* < \lambda_{\max}/2$, causing the population to concentrate around λ_{\max} . In other terms,

$$p(\lambda) \simeq \begin{cases} \delta(\lambda - \lambda^*) & \text{if } \lambda^* > \lambda_{\max}/2 \\ \delta(\lambda - \lambda_{\max}) & \text{if } \lambda^* < \lambda_{\max}/2 \end{cases} \quad (22)$$

in agreement with the numerical picture for the two-state (const- x) environment shown in Fig. 4.

This result can be used to obtain an analytical approximation for Λ . In fact, considering that the system spends roughly half the time in the nonselective environment ($x = \lambda_{\max}$) and the other half in the selective one ($x = x_{\min}$), we have [see (10)]

$$\Lambda \simeq \frac{1}{2} \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) + \frac{1}{2} \mathbb{E}_{\lambda \leq x_{\min}} f(\lambda, t) \quad (23)$$

$$\simeq \begin{cases} \lambda^* & \text{if } \lambda^* > \frac{\lambda_{\max}}{2} \\ \frac{\lambda_{\max}}{2} & \text{if } \lambda^* \leq \frac{\lambda_{\max}}{2} \end{cases}, \quad (24)$$

where

$$\mathbb{E}_{\lambda \leq z} f(\lambda, t) = \int_0^z f(\lambda, t) p(\lambda, t) d\lambda, \quad (25)$$

and we used the fact that $\mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) = \lambda^*$ [resp. $\mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) = \lambda_{\max}$] for $\lambda^* > \lambda_{\max}/2$ (resp. $\lambda^* < \lambda_{\max}/2$), while $\mathbb{E}_{\lambda \leq x_{\min}} f(\lambda, t) = \lambda^*$ [resp. $\mathbb{E}_{\lambda \leq x_{\min}} f(\lambda, t) = 0$] for $\lambda^* > \lambda_{\max}/2$ (resp. $\lambda^* < \lambda_{\max}/2$).

In Fig. 3(c) we show that the value of Λ estimated numerically agrees with the one just derived in the limit $D \rightarrow 0$ (horizontal blue line) for $\lambda^* = x_{\min}$. Note that Λ , Eq. (24), corresponds to the maximum of the time-averaged IRR $\mathbb{E}_t f(\lambda, t)$ [see Fig. 4(c)], confirming how, for small D (when diffusion is much slower than environmental fluctuations), fitness is ultimately limited by the environment alone.

2. Rand- x environments

In the ‘‘rand- x ’’ case (x oscillating between λ_{\max} and a random value λ^* uniformly chosen from $[x_{\min}, \lambda_{\max}]$), f will again equal λ roughly half the time, while for the other half it will be randomly zero or λ depending on x_{\min} . In particular, $\text{Prob}\{f = \lambda\} \equiv \text{Prob}\{x > \lambda\} = 1 - \phi$, with

$$\phi = \frac{\lambda - x_{\min}}{\lambda_{\max} - x_{\min}}. \quad (26)$$

One therefore finds

$$\mathbb{E}_t f(\lambda, t) \simeq \begin{cases} \lambda & \text{if } \lambda < x_{\min} \\ \lambda(1 - \phi) & \text{if } \lambda \geq x_{\min} \end{cases}, \quad (27)$$

from which one sees that $\mathbb{E}_t f(\lambda, t)$ attains a maximum value \bar{f}_{\max} given by

$$\bar{f}_{\max} = \frac{1}{2} \frac{(\lambda_{\max} - \frac{1}{2} x_{\min})^2}{\lambda_{\max} - x_{\min}}, \quad (28)$$

at $\lambda = \lambda_{\max} - \frac{1}{2} x_{\min}$ if $x_{\min} < \frac{2}{3} \lambda_{\max}$, while $\bar{f}_{\max} = x_{\min}$ at $\lambda = x_{\min}$ if $x_{\min} > \frac{2}{3} \lambda_{\max}$. In complete analogy with the previous case, the population concentrates around phenotypes λ for which $\mathbb{E}_t f(\lambda, t)$ is maximum, while for the asymptotic growth rate of the population Λ one finds

$$\Lambda \simeq \bar{f}_{\max} \quad (29)$$

(see Fig. 4). The results displayed in Fig. 3(c) (red horizontal line for $D \rightarrow 0$) indeed support this conclusion.

D. Exploration limit (case of symmetric switching)

In the limit $D \rightarrow \infty$ (and more generally whenever diffusion occurs on timescales much faster than those of environmental fluctuations), the growth term in Eq. (4) is negligible with respect to the diffusion one and population is rapidly redistributed according to the underlying phenotypic landscape described by $q(\lambda)$. As a consequence $p(\lambda) \rightarrow q(\lambda)$ asymptotically. It is again possible to derive an approximate expression for Λ from Eq. (10) following the lines traced in the previous section. One finds, in analogy with (23),

$$\Lambda \simeq \frac{1}{2} \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) + \frac{1}{2} \langle\langle f \rangle\rangle, \quad (30)$$

where

$$\langle\langle f \rangle\rangle = \int_{x_{\min}}^{\lambda_{\max}} [\mathbb{E}_{\lambda \leq x} f(\lambda, t)] \pi(x) dx \quad (31)$$

and $\pi(x)$ stands for the probability distribution of the threshold x . Specifically, $\pi(x) = \delta(x - x_{\min})$ in the const- x case and $\pi(x) = (\lambda_{\max} - x_{\min})^{-1}$ for $x \in [x_{\min}, \lambda_{\max}]$ in the rand- x case. Note that, because $p(\lambda) \simeq q(\lambda)$ and $f = \lambda$ (resp. $f = 0$) for $\lambda < x$ (resp. $\lambda > x$), we have

$$\mathbb{E}_{\lambda \leq x} f(\lambda, t) \simeq \int_0^x \lambda q(\lambda) d\lambda = \frac{\lambda_{\max}}{(a+2)} \left\{ 1 - \left[1 - (a+1) \frac{x}{\lambda_{\max}} \right] \left(1 - \frac{x}{\lambda_{\max}} \right)^{a+1} \right\}. \quad (32)$$

Substituting this into (31) and then in Eq. (30) one obtains

$$\Lambda \simeq \begin{cases} \frac{\lambda_{\max}}{(a+2)} \left[1 - \frac{(a+2)}{2} \frac{x_{\min}}{\lambda_{\max}} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+1} - \frac{1}{2} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+2} \right] & \text{(const-}x \text{ environment)} \\ \frac{\lambda_{\max}}{(a+2)} \left[1 - \frac{1}{2} \frac{x_{\min}}{\lambda_{\max}} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+1} - \frac{1}{(a+3)} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+2} \right] & \text{(rand-}x \text{ environment)} \end{cases}. \quad (33)$$

These formulas confirm the intuitive picture according to which the more the underlying distribution of phenotypes q concentrates on small values of CRR (i.e., the larger the value of a), the slower the population grows at fast phenotypic diffusion. Figures 3(c) and 3(f) (horizontal lines at $D \gg 1$) show that the agreement between the long-term population growth rate computed numerically and the theoretical estimate given above is excellent in both const- x and rand- x environments.

E. Case of asymmetric switching times

We have so far assumed that the characteristic times for switching between selective and nonselective environments are identical. This leaves a single environmental timescale in the problem and simplifies the analysis because of the fact that the population spends on average half the time in the selective regime and the other half in the nonselective one. We now want to address the extension of our results to asymmetric switching times.

Numerical results (see Fig. 5) reproduce the qualitative picture derived in the symmetric case, with some (noteworthy) modifications. In the first place, when the mean time spent in the nonselective environment is larger, the advantage provided by diffusion in complex environments is diminished while the exploitation limit yields higher fitness with respect to the symmetric case. Vice versa, exploration can be tuned to obtain a higher fitness for the population when the mean time spent in the selective environment is larger. The fitness achieved in the exploration limit is, however, smaller than the symmetric case. Perhaps most interestingly, in two-state environments with random switching times (rand- t , const- x) the population can still structure in a bimodal fashion, but the weight of the slower part of the distribution (smaller λ) reflects the (mean) time spent in the selective state (i.e., it increases with ω_s). This behavior fully corresponds to the classical ‘‘bet-hedging’’ scenario described, e.g., in Ref. [14]. In other types of environments, though, other population structures are favored.

The key that allows us to easily generalize the fast and slow diffusion limits lies in the observation that, instead of spending on average half the time in each environmental state (selective or nonselective), the population now spends a fraction $p_{\text{ns}} = \frac{\omega_{\text{ns}}}{\omega_{\text{ns}} + \omega_s}$ of time in the nonselective state and a fraction

$p_s = 1 - p_{\text{ns}}$ of time in the selective one. Therefore, the time average of f in the const- x environment (14) now reads

$$\mathbb{E}_t f(\lambda, t) \simeq \begin{cases} \lambda & \text{if } \lambda < x_{\min} \\ p_{\text{ns}} \lambda & \text{if } \lambda \geq x_{\min} \end{cases}. \quad (34)$$

The mean IRR displays again a discontinuity at the threshold x_{\min} , but now the value of λ for which it attains a maximum depends on both x_{\min} and p_{ns} . In specifics, for $x_{\min} > p_{\text{ns}} \lambda_{\max}$ (resp. $x_{\min} < p_{\text{ns}} \lambda_{\max}$), the mean IRR has a maximum for $\lambda = x_{\min}$ (resp. $\lambda = \lambda_{\max}$), where $\mathbb{E}_t f(\lambda, t) = x_{\min}$ (resp. $\mathbb{E}_t f(\lambda, t) = p_{\text{ns}} \lambda_{\max}$).

Analytical approximation for the population fitness Λ that account for asymmetry in the environment can be easily obtained along the lines of Secs. III C and III D. In particular, in the exploitation limit and with a two-state (const- x) environment, expressions (23) and (24) generalize to

$$\Lambda \simeq p_{\text{ns}} \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) + p_s \mathbb{E}_{\lambda \leq x_{\min}} f(\lambda, t) \quad (35)$$

$$\simeq \begin{cases} x_{\min} & \text{if } x_{\min} > p_{\text{ns}} \lambda_{\max} \\ p_{\text{ns}} \lambda_{\max} & \text{if } x_{\min} \leq p_{\text{ns}} \lambda_{\max} \end{cases}. \quad (36)$$

Likewise, in the exploitation limit for the rand- x case one finds that Eq. (27) takes the form

$$\mathbb{E}_t f(\lambda, t) \simeq \begin{cases} \lambda & \text{if } \lambda < x_{\min} \\ \lambda(1 - p_s \phi) & \text{if } \lambda \geq x_{\min} \end{cases}. \quad (37)$$

One sees that $\mathbb{E}_t f(\lambda, t)$ now attains a maximum value \bar{f}_{\max} given by

$$\bar{f}_{\max} = \frac{1}{4p_s} \frac{(\lambda_{\max} - p_{\text{ns}} x_{\min})^2}{\lambda_{\max} - x_{\min}}, \quad (38)$$

at $\lambda = \frac{\lambda_{\max} - p_{\text{ns}} x_{\min}}{2p_s}$ if $x_{\min} < \frac{\lambda_{\max}}{1+p_s}$, while $\bar{f}_{\max} = x_{\min}$ at $\lambda = x_{\min}$ if $x_{\min} > \frac{\lambda_{\max}}{1+p_s}$. As before, $\Lambda \simeq \bar{f}_{\max}$.

Finally, in the exploration limit asymmetric environments turn Eq. (30) into

$$\Lambda \simeq p_{\text{ns}} \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t) + p_s \langle\langle f \rangle\rangle, \quad (39)$$

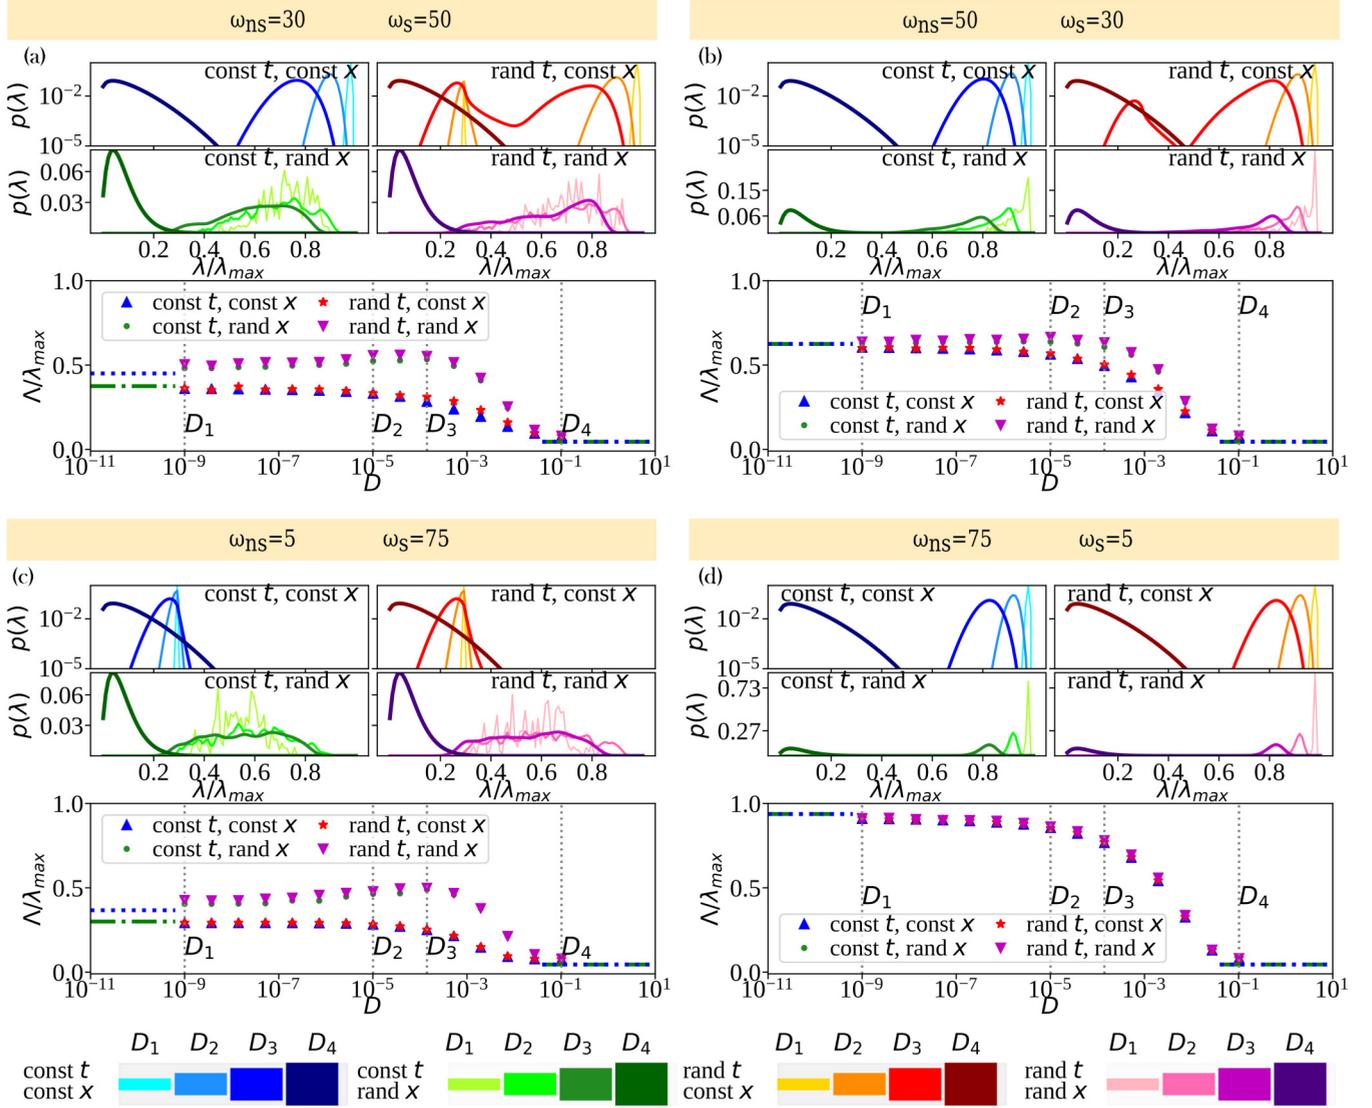


FIG. 5. (a) Asymptotic, time-averaged phenotypic distributions (top panels) and asymptotic population growth rate Λ (in units of λ_{\max}) as a function of D (bottom panel) obtained as in Figs. 3(d)–3(f) but in the presence of an asymmetric environment with characteristic switching times $\omega_{\text{s}} = 30$ and $\omega_{\text{ns}} = 50$ time units. As in Fig. 3, $x_{\min} = 0.3\lambda_{\max}$ and results obtained for different values of D are marked by increasing color shades and line widths. Dotted blue and dot-dashed green horizontal lines at small and large D show the analytical estimates of Λ in the exploration and exploitation limits, obtained in the const- x and rand- x regimes, respectively. (b)–(d) Same as panel (a) but with different choices of ω_{ns} and ω_{s} . Displayed curves are averaged over 100 independent realizations of the dynamics.

which allows us to generalize Eq. (33) as

$$\Lambda \simeq \begin{cases} \frac{\lambda_{\max}}{(a+2)} \left[1 - p_{\text{s}}(a+2) \frac{x_{\min}}{\lambda_{\max}} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+1} - p_{\text{s}} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+2} \right] & (\text{const-}x \text{ environment}) \\ \frac{\lambda_{\max}}{(a+2)} \left[1 - p_{\text{s}} \frac{x_{\min}}{\lambda_{\max}} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+1} - \frac{2p_{\text{s}}}{(a+3)} \left(1 - \frac{x_{\min}}{\lambda_{\max}} \right)^{a+2} \right] & (\text{rand-}x \text{ environment}) \end{cases}. \quad (40)$$

Figure 5 (see green and blue horizontal lines) shows that the above expressions for Λ provide an excellent agreement with numerical results in both the exploration and exploitation limits.

IV. DISCUSSION

Empirical data on phenotypic distributions, quantified, e.g., from protein expression data, display a rich spectrum of behaviors ranging from unimodal to bimodal depending on the applied stress, organism, etc. (see, e.g., Ref. [34] for evidence regarding *E. coli*). The question of when one type of distribution is favored therefore appears to be subtle and

possibly requires a case-by-case answer. Our results are in line with previous work in suggesting that the population structure is tightly linked to the specific features of the environment. In particular, when the strength of the coupling between the environment and phenotypes takes on two distinct levels (e.g., high or low, corresponding to the const- x case), bimodal distributions arise but exploration does not yield a fitness

advantage to the population. On the other hand, under the more complex scenario in which the coupling strength varies randomly (rand- x case), the exploration-exploitation trade-off leads to a nonzero optimal search rate, and unimodal phenotypic distributions are generically preferred. This picture is in complete agreement with the results obtained in Ref. [29], where the theoretical benefit of a bimodal distribution of stress response proteins was found to be highest in two-state environments, while more variable and structured environments allow for the selection of unimodal distributions. In addition, we have found that adding a small amount of diffusion to a purely exploitative strategy always leads to an increase of fitness in rand- x environments, while it is always detrimental in const- x environments. (More generally, diffusion appears to be broadly beneficial in rand- x environments.) Therefore, both the way a population is distributed across its phenotypic space and its fitness directly reflect its history in coping with the random environment.

At the quantitative level, the fitness gain given by exploration also appears to be linked to the structure of the underlying phenotypic landscape $q(\lambda)$. In particular, in the more realistic case in which $q(\lambda)$ is strongly heterogeneous, with rare fast phenotypes among a multitude of slow ones [31], a diffusive search dynamics can provide a significant fitness advantage. More generally, it appears to be possible to set the exploration rate within an optimal range for any environment when (1) losses caused by fast diffusion (high D) are avoided, while (2) losses that are to be faced by exploring the phenotypic landscape in two-state (const- x) environments are not too large with respect to the $D \rightarrow 0$ limit. A rather broad range of values of D fit this criterion, suggesting that, while possibly helpful in certain conditions, a tight regulation of the phenotypic exploration rate may be unnecessary as long as the key assumptions made here hold.

From a physical viewpoint, our model ultimately relies on Markovianity and detailed balance. These ingredients provide in our view the most elementary way to encode for the effects of fully unbiased random changes in cellular physiology at the level of a complex macroscopic parameter such as the growth rate. However, they are likely to fail in many biologically realistic contexts, and moving beyond them would be important. Another limiting modeling choice we made concerns the assumption that faster-growing cells susceptible to environmental shocks do not replicate, as we are implicitly postulating that they survive the shock. While this may be unrealistic in some situations, we note that the introduction of an explicit cellular death rate would effectively rescale the “replicator” term in Eq. (4). The qualitative scenario we describe should therefore persist. Finally, we focused on a diffusive transition kernel in which only small changes in CRR are allowed. While this is a reasonable choice in biological contexts, when significant phenotypic re-arrangements can occur the emergent scenario may be different. For instance, this is likely to be the case when transition rates follow a Gibbs kernel, in which the $W(\lambda \rightarrow \lambda')$ affecting (1) is proportional to the density of states with CRR λ' :

$$W(\lambda \rightarrow \lambda') = \frac{q(\lambda')}{\tau}. \quad (41)$$

In particular, in this situation diffusion may turn out to be more efficient in improving population fitness than under a

diffusive kernel, most notably so in homogeneous landscapes. On the other hand, justifying a kernel like (41) for biological modeling would necessarily require assumptions more extreme, and possibly less realistic, than those made here.

At a more speculative level, this work could shed some light on the origin of phenotypically heterogeneous cell populations such as tumors and may point to educated strategies to control their diversity. For instance, more heterogeneous populations are more likely to evolve in complex environments, suggesting, e.g., that higher intratumoral heterogeneity may be the result of highly variable microenvironments. On the other hand, if the “shocks” are taken to be caused by a therapeutic protocol, our study suggests that subjecting the population to a single repeated dose is effective in quenching its fitness irrespective of the timing of administration.

ACKNOWLEDGMENTS

Work was supported by the European Union’s Horizon 2020 Research and Innovation Staff Exchange program MSCA-RISE-2016 under Grant Agreement No. 734439 INFERNET.

APPENDIX: DERIVATION OF EQ. (8)

We start by noting that

$$\dot{n}(\lambda, t) = \dot{p}(\lambda, t)N(t) + p(\lambda, t)\dot{N}(t), \quad (A1)$$

where, from Eq. (1),

$$\frac{\dot{N}(t)}{N(t)} = \int_0^{\lambda_{\max}} f(\lambda, t)p(\lambda, t)d\lambda = \mathbb{E}_{\lambda \leq \lambda_{\max}} f(\lambda, t). \quad (A2)$$

A comparison between Eqs (1) and (A1) immediately yields the first term on the r.h.s of Eq. (8). To get the second term, we assume a diffusive transition kernel. Making use of the detailed balance condition (6) one finds

$$\begin{aligned} & \int [W(\lambda' \rightarrow \lambda)p(\lambda', t) - W(\lambda \rightarrow \lambda')p(\lambda, t)]d\lambda' \\ &= W(\lambda \rightarrow \lambda + \delta\lambda)q(\lambda) \left[\frac{p(\lambda + \delta\lambda, t)}{q(\lambda + \delta\lambda)} - \frac{p(\lambda, t)}{q(\lambda)} \right] \\ & \quad + W(\lambda - \delta\lambda \rightarrow \lambda)q(\lambda - \delta\lambda) \left[\frac{p(\lambda - \delta\lambda, t)}{q(\lambda - \delta\lambda)} - \frac{p(\lambda, t)}{q(\lambda)} \right] \\ & \simeq \frac{(\delta\lambda)^2}{2\tau} \frac{\partial}{\partial \lambda} \left[q(\lambda) \frac{\partial}{\partial \lambda} \frac{p(\lambda, t)}{q(\lambda)} \right], \end{aligned} \quad (A3)$$

where the last step follows after a second-order expansion in $\delta\lambda$ and we imposed that transitions from λ to $\lambda \pm \delta\lambda$ happen with the same probability [implying that $W(\lambda \rightarrow \lambda \pm \delta\lambda) = (2\tau)^{-1}$; see (7)]. Defining $D = \frac{(\delta\lambda)^2}{2\tau}$, the second term in Eq. (8) is immediately recovered.

- [1] J. D. Cohen, S. M. McClure, and J. Y. Angela, Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration, *Philos. Trans. R. Soc. B* **362**, 933 (2007).
- [2] T. Gueudré, A. Dobrinevski, and J. P. Bouchaud, Explore or Exploit? A Generic Model and an Exactly Solvable Case, *Phys. Rev. Lett.* **112**, 050602 (2014).
- [3] T. Gueudré and D. Martin, The balance of growth and risk in population dynamics, *Europhys. Lett.* **121**, 68005 (2018).
- [4] T. Dannemann, D. Boyer, and O. Miramontes, Lévy flight movements prevent extinctions and maximize population abundances in fragile Lotka-Volterra systems, *Proc. Natl. Acad. Sci. USA* **115**, 3794 (2018).
- [5] K. S. Korolev, J. B. Xavier, and J. Gore, Turning ecology and evolution against cancer, *Nat. Rev. Cancer* **14**, 371 (2014).
- [6] N. McGranahan and C. Swanton, Clonal heterogeneity and tumor evolution: Past, present, and the future, *Cell* **168**, 613 (2017).
- [7] I. R. Henderson, P. Owen, and J. P. Nataro, Molecular switches—The ON and OFF of bacterial phase variation, *Mol. Microbiol.* **33**, 919 (1999).
- [8] B. Hallet, Playing Dr Jekyll and Mr Hyde: Combined mechanisms of phase variation in bacteria, *Curr. Opin. Microbiol.* **4**, 570 (2001).
- [9] M. W. Van Der Woude and A. J. Bäumlner, Phase and antigenic variation in bacteria, *Clin. Microbiol. Rev.* **17**, 581 (2004).
- [10] M. Thattai and A. van Oudenaarden, Stochastic gene expression in fluctuating environments, *Genetics* **167**, 523 (2004).
- [11] J. L. Radzikowski, S. Vedelaar, D. Siegel, A. D. Ortega, A. Schmidt, and M. Heinemann, Bacterial persistence is an active σ^S stress response to metabolic flux limitation, *Mol. Syst. Biol.* **12**, 882 (2016).
- [12] J. H. van Heerden, M. T. Wortel, F. J. Bruggeman, J. J. Heijnen, Y. J. Bollen, R. Planqué, J. Hulshof, T. G. O'Toole, S. A. Wahl, and B. Teusink, Lost in transition: Start-up of glycolysis yields subpopulations of nongrowing cells, *Science* **343**, 1245114 (2014).
- [13] M. Acar, J. T. Mettetal, and A. Van Oudenaarden, Stochastic switching as a survival strategy in fluctuating environments, *Nat. Gen.* **40**, 471 (2008).
- [14] A. J. Grimbergen, J. Siebring, A. Solopova, and O. P. Kuipers, Microbial bet-hedging: The power of being different, *Curr. Opin. Microbiol.* **25**, 67 (2015).
- [15] N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, and S. Leibler, Bacterial persistence as a phenotypic switch, *Science* **305**, 1622 (2004).
- [16] I. Keren, N. Kaldalu, A. Spoering, Y. Wang, and K. Lewis, Persist cells and tolerance to antimicrobials, *FEMS Microbiol. Lett.* **230**, 13 (2004).
- [17] R. Levins, *Evolution in Changing Environments: Some Theoretical Explorations* (Princeton University Press, Princeton, 1968).
- [18] E. Kussell and S. Leibler, Phenotypic diversity, population growth, and information in fluctuating environments, *Science* **309**, 2075 (2005).
- [19] M. J. Gander, C. Mazza, and H. Rummeler, Stochastic gene expression in switching environments, *J. Math. Biol.* **55**, 249 (2007).
- [20] P. Visco, R. J. Allen, S. N. Majumdar, and M. R. Evans, Switching and growth for microbial populations in catastrophic responsive environments, *Biophys. J.* **98**, 1099 (2010).
- [21] J. Müller, B. A. Hense, T. M. Fuchs, M. Utz, and C. Pötzsche, Bet-hedging in stochastically switching environments, *J. Theor. Biol.* **336**, 144 (2013).
- [22] I. Cvijović, B. H. Good, E. R. Jerison, and M. M. Desai, Fate of a mutation in a fluctuating environment, *Proc. Natl. Acad. Sci. USA* **112**, E5021 (2015).
- [23] A. Melbinger and M. Vergassola, The impact of environmental fluctuations on evolutionary fitness functions, *Sci. Rep.* **5**, 15211 (2015).
- [24] P. Patra and S. Klumpp, Emergence of phenotype switching through continuous and discontinuous evolutionary transitions, *Phys. Biol.* **12**, 046004 (2015).
- [25] A. Skanata and E. Kussell, Evolutionary Phase Transitions in Random Environments, *Phys. Rev. Lett.* **117**, 038104 (2016).
- [26] A. Mayer, T. Mora, O. Rivoire, and A. M. Walczak, Transitions in optimal adaptive strategies for populations in fluctuating environments, *Phys. Rev. E* **96**, 032412 (2017).
- [27] P. G. Hufton, Y. T. Lin, and T. Galla, Phenotypic switching of populations of cells in a stochastic environment, *J. Stat. Mech.* (2018) 023501.
- [28] K. Wienand, E. Frey, and M. Mobilia, Evolution of a Fluctuating Population in a Randomly Switching Environment, *Phys. Rev. Lett.* **119**, 158301 (2017).
- [29] J. Garcia-Bernardo and M. J. Dunlop, Phenotypic diversity using bimodal and unimodal expression of stress response proteins, *Biophys. J.* **110**, 2278 (2016).
- [30] S. Taheri-Araghi, S. D. Brown, J. T. Sauls, D. B. McIntosh, and S. Jun, Single-cell physiology, *Ann. Rev. Biophys.* **44**, 123 (2015).
- [31] D. De Martino, F. Capuani, and A. De Martino, Growth against entropy in bacterial metabolism: The phenotypic trade-off behind empirical growth rate distributions in *E. coli*, *Phys. Biol.* **13**, 036005 (2016).
- [32] D. De Martino, F. Capuani, and A. De Martino, Quantifying the entropic cost of cellular growth control, *Phys. Rev. E* **96**, 010401 (2017).
- [33] D. De Martino and D. Masoero, Asymptotic analysis of noisy fitness maximization, applied to metabolism & growth, *J. Stat. Mech.* (2016) 123502.
- [34] Y. Taniguchi, P. J. Choi, G. W. Li, H. Chen, M. Babu, J. Hearn, A. Emili, and X. S. Xie, Quantifying *E. coli* proteome and transcriptome with single-molecule sensitivity in single cells, *Science* **329**, 533 (2010).