Phenotypic Heterogeneity in Lag Reflects an Evolutionarily Stable Bet-Hedging Strategy

Abir George [1,2] Ned S. Wingreen [1,3,*] and Gautam Reddy [2,†]
[1]*Lewis-Sigler Institute for Integrative Genomics, Princeton University, Princeton, New Jersey 08544, USA*
[2]*Joseph Henry Laboratories of Physics, Princeton University, Princeton, New Jersey 08544, USA*
[3]*Department of Molecular Biology, Princeton University, Princeton, New Jersey 08544, USA*

(Received 11 November 2025; accepted 12 June 2026; published 9 July 2026)

Single-cell experiments in yeast reveal two distinct heritable phenotypes—"arresters" and "recoverers"—when a clonal population experiences a negative shift in its growth environment. Recoverers exhibit a variable yet finite lag before resuming growth in the new environment, whereas arresters remain in a nongrowing, arrested state until more favorable conditions return. Although the diversification of individual cells into arresters and recoverers is a robust phenomenon, it remains unclear whether this coexistence constitutes an evolutionarily stable strategy. Here, we demonstrate that a heterogeneous strategy composed of both arrester and recoverer phenotypes maximizes long-term population fitness across a broad spectrum of growth-lag trade-offs. Our analysis employs a dynamic programming framework to identify the fitness-maximizing distribution of phenotypes for populations that stochastically switch between preferred and nonpreferred environments. We propose a minimal model incorporating metabolism, growth, and enzyme allocation to explain the physiological origin of a power-law growth-lag trade-off that favors phenotypic heterogeneity. The theory predicts a nontrivial relationship between the fraction of recoverers and their lag time, which aligns with existing data from wild yeast strains, evolved isolates, and variations in preshift growth conditions. This relationship suggests an evolutionary "rheostat"-like mechanism that enables populations to rapidly adapt to changing environmental conditions.

DOI: 10.1103/rt5t-qdr8

I. INTRODUCTION

Microbes typically live in environments characterized by rapid and unpredictable changes in nutrient availability, temperature, pH, and other critical factors. To persist and thrive under these uncertain conditions, microbes have evolved diverse adaptive strategies. These include mechanisms for physiological adaptation and gene regulation, as well as phenotypic heterogeneity—whereby multiple distinct phenotypes coexist within genetically identical populations. Notably, phenotypic heterogeneity can act as a bet-hedging mechanism [1], allowing microbial populations to hedge against environmental uncertainty, thus enhancing their long-term success. A well-known example of bet-hedging is the formation of antibiotic-tolerant cells in *Escherichia coli* [2–4]. Within a clonal population, a small fraction of cells grows very slowly, or not at all, which makes them more likely to survive antibiotics and other stressors. These "persister" cells serve as a reservoir for repopulation once favorable conditions are restored. Bet-hedging through phenotypic diversification has been extensively explored in ecological and evolutionary theory, with numerous theoretical studies analyzing its implications [5–13].

Recent experiments in yeast have revealed a distinctive form of phenotypic heterogeneity in response to nutrient shifts. Upon a shift between environments with different nutrients, microbes undergo a period of slow or halted growth, termed the lag phase, during which cells express the machinery that enables growth in the new environment [14]. Yeast (including *Saccharomyces cerevisiae* and its relatives) displays a lag period when transferred from a preferred carbon source like glucose to alternative nutrients such as maltose, galactose, or ethanol. Quantitative characterization of this lag phase has revealed considerable variability in the duration and dynamics of lag across natural yeast strains [15–23]: while some strains rapidly resume growth upon a shift, others display many hours of lag before gene expression can reorganize to resume growth.

Beyond this population-level variability observed across strains, single-cell experiments have also found significant heterogeneity in lag. Individual cells from a clonal yeast population display marked variation in lag duration, ranging from immediate resumption of growth to prolonged growth arrest [15,18]. A standard experimental protocol is to transfer isogenic yeast cells grown, for example, in glucose to maltose and track single-cell growth using time-lapse microscopy (Fig. 1). Cells show two distinct modes of phenotypic response: while some cells (termed "recoverers") resume growth after a few hours, a fraction of cells (termed "arresters") remain in a nongrowing, arrested state for the entire duration (≈24 hours) of the experiment [Figs. 1(b) and

*Contact author: wingreen@princeton.edu
†Contact author: greddy@princeton.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

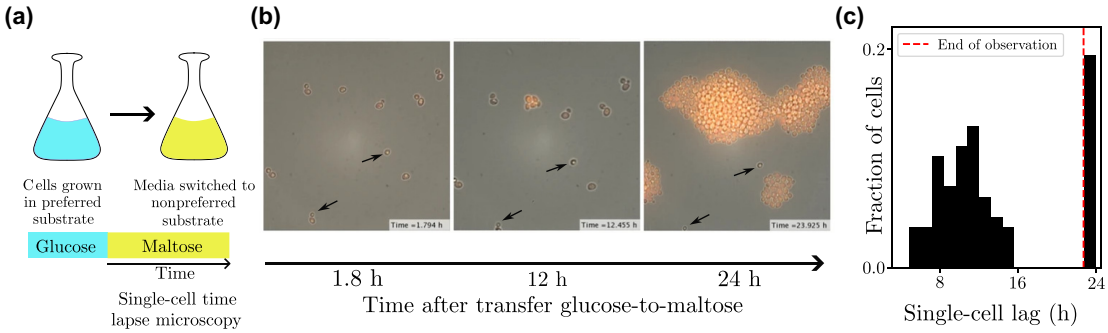


FIG. 1. Heterogeneity in single-cell lag durations following a glucose-to-maltose shift. (a) Schematic of experimental protocol: yeast cells are pregrown in glucose, then the media is switched to maltose, where cells are tracked using time-lapse microscopy. (b) Time-lapse images of a yeast population following a sudden switch from glucose to maltose. MAL12-yECitrine fluorescence indicates growth resumption. Images taken at 1.8, 12, and 24 h after the switch. Cells exhibit heterogeneous behavior in lag phase recovery, including a subpopulation that remains arrested (indicated by black arrows). Images reproduced from Ref. [18]. (c) Distribution of single-cell lag durations across the population. The red dashed line marks the end of the observation period (23.5 h), beyond which lag durations were not measured.

1(c)]. The arrested state is characterized by deformations of subcellular structures [15] and is reversible; that is, arresters resume growth when reintroduced into a glucose-rich environment. While this bimodal response is consistently observed across wild strains and evolved isolates as well as growth media and the time spent in the preshift environment, the ratios of recoverers to arresters and the distributions of lag durations within the recoverer subpopulation both vary substantially across these scenarios [18,21]. A natural question is whether this coexistence of arrester and recoverer states reflects a stable evolutionary strategy.

If rapid recovery from lag entailed a negligible physiological cost (for example, in preshift growth rate), then we should expect cells to prefer a homogeneous strategy that recovers as quickly as possible. However, the observation of an arrested state suggests that cells face intrinsic limitations when adapting to maltose or other less-preferred carbon sources. Extensive experimental work has characterized trade-offs between growth and adaptability in microbes [15,17,19,24–29]. When such trade-offs are present, the coexistence of distinct phenotypes can emerge as an evolutionarily stable strategy [15,26–28,30,31]. Thus, one possible explanation for the observed coexistence is that individual yeast cells are constrained by some nontrivial physiological trade-off that prevents them from simultaneously achieving rapid preshift growth and rapid postshift recovery. This hypothesis is supported by measurements that show that arresters grow faster than recoverers in the preshift environment [15].

To determine the evolutionarily stable strategy in the presence of such a growth-lag trade-off, we leverage a novel optimization framework based on dynamic programming. The framework leads to a general decision-making principle that can be used to find fitness-maximizing bet-hedging strategies for a population fluctuating between a preferred and a nonpreferred environment. We show that, depending on environmental statistics and the physiological trade-off, the population either adopts a bet-hedging strategy composed of two "specialist" phenotypes (arresters and recoverers) that respectively are well-suited for growth and rapid recovery, or picks a "generalist" phenotype that balances both traits. We propose a minimal metabolic model that explains the physiological origin of a power-law-like growth-lag trade-off. Furthermore, the model predicts that populations experiencing different environmental statistics follow a characteristic relationship between the ratio of recoverers and arresters and the lag duration of recoverers. We test this prediction using existing experimental data, which hints at an underlying evolutionary mechanism that enables rapid adaptation to changes in environmental statistics.

II. THEORETICAL FRAMEWORK

To quantify the long-term fitness benefits of phenotypic heterogeneity, we developed a cellular decision-making framework based on dynamic programming [32]. We first describe the general framework before specializing to the growth-lag scenario. Another application of this framework to a growth-death scenario is presented in Appendix B 4.

Consider a microbial population where cells adopt one of K phenotypic states, labeled by k, each of which is characterized by its growth rate $\gamma(s,k)$ in state s. The state s can represent, for example, the time since a nutrient switch, the concentration of a nutrient, or any history-dependent quantity tracked by the cell. The instantaneous growth rate of a population of N cells at time t is

$$\frac{1}{N}\frac{dN}{dt} = \sum_{k=1}^{K} \phi_k(t)\gamma(s(t),k) \equiv \overline{\gamma}(s(t),\boldsymbol{\phi}(t)), \tag{1}$$

where ϕ_k is the fraction of cells with phenotype k and $\boldsymbol{\phi} = (\phi_1, \phi_2, \ldots, \phi_K)$ is the phenotype distribution. The long-term growth rate Γ of the population is then

$$\Gamma = \lim_{T\to\infty} \frac{1}{T}\int_0^T \langle\overline{\gamma}(s(t),\boldsymbol{\phi}(t))\rangle dt, \tag{2}$$

where the angled brackets denote an expectation over changes in state.

We focus on populations that transition repeatedly between preferred (P) and nonpreferred (NP) environments, such as yeast shifting between glucose and maltose. Importantly, we assume that in P the population can rapidly reorganize gene expression and thereby adjust its phenotype distribution,

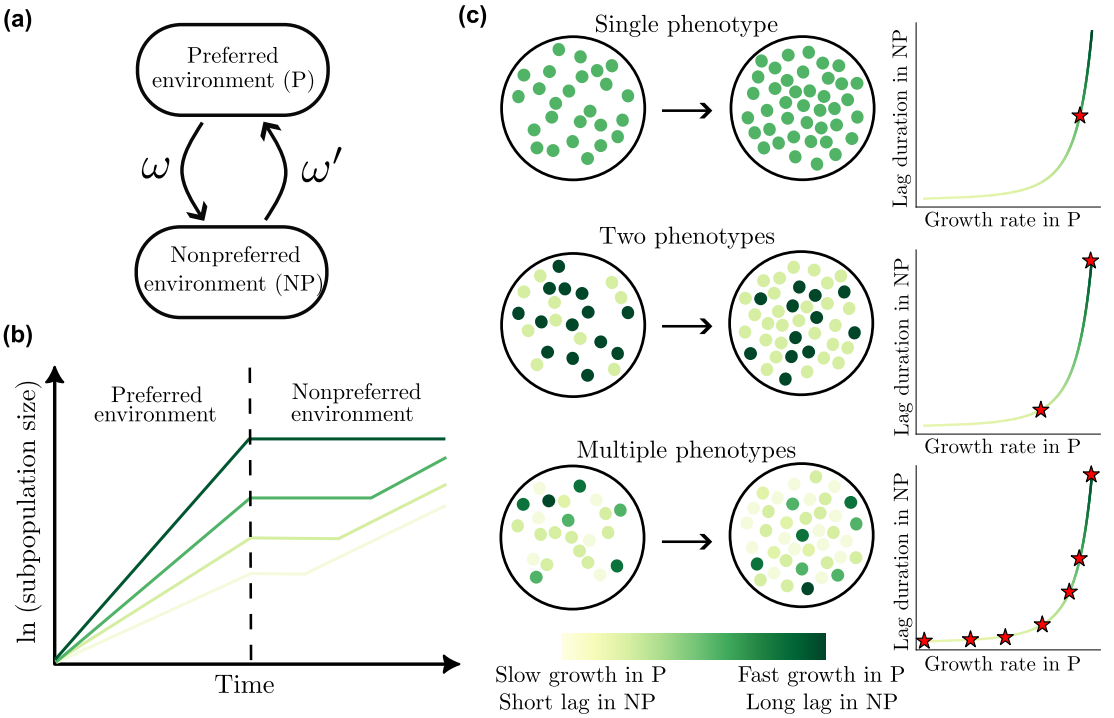


FIG. 2. Phenotypic strategies in fluctuating environments. (a) Model of environmental switching with transitions between a preferred (P) and a nonpreferred (NP) environment occurring at rates ω and ω', respectively. (b) Growth dynamics of pure phenotype subpopulations in a fluctuating environment. Distinct phenotypes grow at different rates in the preferred environment (P) and experience varying lag durations upon transition to the nonpreferred environment (NP), with a trade-off between growth rate in P and lag duration in NP. Cells that grow slowly in P recover more quickly and resume growth earlier in NP. As a result, these phenotypes increase in relative abundance during NP. (c) Schematic showing possible strategies for a population of cells: a single phenotype (top), two phenotypes (middle), or multiple phenotypes (bottom). Disks represent individual cells colored by their phenotype, corresponding to a position on the growth-lag trade-off curve (right). Red stars indicate the phenotypes that comprise the strategy.

whereas cells remain phenotypically locked upon the switch to NP. This environmental switching model is schematically illustrated in Figs. 2(a) and 2(b). Fluctuating environments will select for phenotype distributions that best balance overall population growth in P and NP. A key conceptual question is when and whether this best population-level strategy consists of one, two, or multiple phenotypes [Fig. 2(c)].

Our objective is thus to identify the optimal phenotype distribution $\boldsymbol{\phi}^*$ in P that maximizes the long-term growth rate Γ in the face of environmental fluctuations. Given state s, and our assumption of rapid phenotype adjustment in P, we show that

$$\boldsymbol{\phi}^*(s) = \arg\max_{\boldsymbol{\phi}} \left\{ \overline{\gamma}_{\mathrm{P}}(s,\boldsymbol{\phi}) + \omega(s) \left\langle \int_0^\infty d\tau\, \Pi(\tau) \left(\int_0^\tau dt\, \bar{\gamma}_{\mathrm{NP}}(s(t),\boldsymbol{\phi}(t)) \right) \right\rangle \right\}, \tag{3}$$

where $\Pi(\tau)$ is the probability density of the duration τ spent in NP before switching back to P, $\omega(s)$ is the instantaneous rate of switching from P to NP and we have included the subscripts P and NP to $\bar{\gamma}$ for clarity. Equation (3) has an intuitive interpretation: the optimal strategy balances the instantaneous growth rate in P (first term) and the expected long-term growth rate if the environment were to transition to NP, scaled by the probability that the transition occurs (second term).

III. RESULTS

Optimal strategies given a growth-lag trade-off. In our growth-lag scenario, cells with phenotype k have a growth rate γ_k in P and a lag duration ℓ_k in NP with no growth during lag. After recovering from lag, cells grow in NP at a common rate γ'. Phenotypes experience a growth-lag trade-off: phenotypes that grow faster (slower) in P exhibit longer (shorter) lag durations in NP [Fig. 2(b)]. To incorporate the effect of lag, the state $s(t)$ in NP tracks the time since the most recent transition from P to NP. For simplicity, we drop any dependence of ϕ on s in P, noting, however, that the results also apply to scenarios where ϕ^* depends on the nutrient in P or historical information that could be used to anticipate a switch. The optimization objective then becomes

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}} \left\{ \overline{\gamma}_{\mathrm{P}}(\boldsymbol{\phi}) + \omega \int_0^\infty d\tau\, \Pi(\tau) \int_0^\tau dt\, \overline{\gamma}_{\mathrm{NP}}(t,\boldsymbol{\phi}(t)) \right\}, \tag{4}$$

from which we derive the specific optimization objective

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}} \left\{ \sum_{k=1}^K \phi_k \gamma_k + \omega\omega' \int_0^\infty dt e^{-\omega' t} \times \frac{\sum_{k=1}^K \phi_k \gamma' \Theta(t-\ell_k) e^{\gamma'[t-\ell_k]_+}}{\sum_{k=1}^K \phi_k e^{\gamma'[t-\ell_k]_+}} \right\}, \tag{5}$$

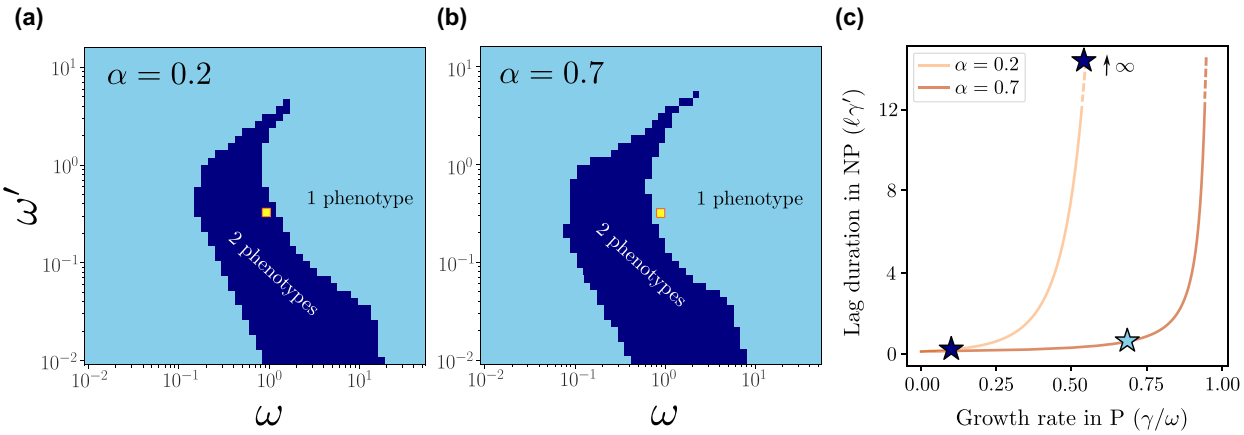


FIG. 3. Optimal phenotypic strategies consist of at most two distinct phenotypes. (a), (b) Phase diagrams of the optimal number of phenotypes over the space of environmental switching rates, ω and ω', for fixed trade-off curves parametrized by (a) $\alpha = 0.2$ and (b) $\alpha = 0.7$. Light-blue regions correspond to a single optimal generalist phenotype, while dark blue regions correspond to two coexisting specialists, an arrester and a recoverer. (c) Different trade-off curves between growth rate in P and lag duration in NP from A and B, with optimal phenotypes for the (ω, ω') values indicated by yellow squares. For increasing α, the optimal strategy transitions from a two-phenotype (specialist) solution (dark-blue stars) to a single-phenotype (generalist) solution (light-blue stars).

where ω' is an assumed constant rate of switching from NP back to P so that $\Pi(\tau) = \omega' e^{-\omega'\tau}$, $\Theta(t - \ell_i)$ is the Heaviside function, and $[t - \ell_i]_+ = \max(0, t - \ell_i)$. Integrating Eq. (5) by parts, the objective function can be simplified to

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}} \left\{ \sum_{k=1}^{K} \phi_k \tilde{\gamma}_k + \frac{1}{\eta} \int_0^\infty dt' e^{-\frac{t'}{\eta}} \ln \left[\sum_{k=1}^{K} \phi_k e^{[t'-\ell'_k]_+} \right] \right\}, \tag{6}$$

where $\tilde{\gamma}_k = \gamma_k/\omega$ and $\ell'_k = \gamma' \ell_k$ represent normalized growth rates and lag durations, respectively, and $\eta = \gamma'/\omega'$ is the average number of generations of a no-lag subpopulation in NP.

A key factor determining a population's optimal strategy in a fluctuating environment is the form of the growth-lag trade-off. To capture the range of possible behaviors, we consider a family of growth-lag trade-offs:

$$\gamma_P(\ell) = \gamma_{\max} - (\gamma_{\max} - \gamma_{\min}) \left(\frac{\ell_{\min}}{\ell} \right)^{\alpha}, \tag{7}$$

where lag $\ell \geqslant \ell_{\min}$, and the growth rate in P lies in the range $\gamma_{\min} \leqslant \gamma_P \leqslant \gamma_{\max}$. The parameter α defines a family of trade-off curves, with larger α corresponding to larger γ_P for the same lag ℓ. This parametric form always includes phenotypes that have maximum growth in P at the cost of infinite lag ($\ell = \infty$) in NP.

Numerical result: at most two phenotypes, with one arrester. We numerically computed the optimal phenotype distribution over a continuous set of phenotypes constrained by the above growth-lag trade-off curve [Eq. (7)]. Although one might have expected the optimal strategy to exploit a broad range of phenotypes, we found that the optimal distribution contains at most two distinct phenotypes. This restriction to no more than two phenotypes applies across the whole parameter space defined by the environmental transition rates (ω, ω'), as illustrated in Figs. 3(a) and 3(b).

Within the two-phenotype regime, the optimal strategy consistently comprises one phenotype with finite lag—a recoverer—and another phenotype with infinite lag ($\ell = \infty$)—an arrester. The recoverer is adapted to resume growth in the nonpreferred environment (NP), while the arrester is specialized for maximal growth in the preferred environment (P). The optimal strategy assigns a nonzero probability to the arrester phenotype, even though it never grows in NP. This is perhaps a surprising result: it implies that under certain environmental conditions, it is beneficial for a population to maintain a subpopulation that entirely forfeits the ability to grow in one environment in order to maximize performance in the other, even though both environments can support growth. Similarly, the recoverer phenotype sacrifices growth in P in order to perform better in NP. The two phenotypes in this bet-hedging strategy can thus be viewed as specialists, each appropriately suited for a distinct environmental state.

It is clear that the phase diagram is re-entrant; that is, for increasing ω/ω', the optimal strategy goes from generalist to specialist, and back to generalist. For small ω/ω', the generalist strategy corresponds to having only the arrester in the population since the environment is mostly P. By contrast, for $\omega/\omega' \approx 1$, the population should optimally adopt a specialist strategy (with deviations with respect to this ratio for $\omega, \omega' \ll 1$ and $\omega, \omega' \gg 1$). When ω/ω' increases beyond this, the optimal strategy reverts to a generalist strategy with a recoverer phenotype. This makes sense as the population spends more time in NP, and the recoverer phenotype is able to maximally contribute to the population growth.

As shown in Fig. 3(c), varying the shape of the growth-lag trade-off via the power-law exponent α can qualitatively change the optimal strategy for fixed (ω, ω'). For example, for the (ω, ω') values in the yellow squares in Figs. 3(a) and 3(b), at low α the optimal strategy is two specialists (one being an arrester), while at higher α, the optimal strategy is a single generalist phenotype that balances growth in P and lag in NP.

Phenotypic distributions with a few discrete phenotypes are optimal. The observation that the optimal phenotype

distribution contains a few (here, at most two) discrete phenotypes can be explained in the special case when the time spent in the nonpreferred environment (NP) is fixed at a constant value t. In this case, the optimization objective in Eq. (4) reduces to (with $t' = \gamma' t$)

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}} \left\{ \sum_{k=1}^{K} \phi_k \tilde{\gamma}_k + \ln\left(\sum_{k=1}^{K} \phi_k e^{[t'-\ell'_k]_+} \right) \right\}, \tag{8}$$

with the constraints $\sum_{k=1}^{K} \phi_k = 1$ and $\phi_k \geqslant 0$ for all k. The Karush-Kuhn-Tucker (KKT) theorem provides a set of conditions satisfied by $\boldsymbol{\phi}^*$ subject to these constraints. The KKT conditions imply that the growth rates and lag durations of any pair of phenotypes (indexed by m and n) with nonzero probability are related according to

$$\frac{e^{[t'-\ell'_m]_+} - e^{[t'-\ell'_n]_+}}{\tilde{\gamma}_m - \tilde{\gamma}_n} = C, \tag{9}$$

for some constant C. This restrictive condition may be true for a continuous set of phenotypes in certain special scenarios, for example, when $\gamma_P(\ell) \propto C_1 - C_2\, e^{-\gamma' \ell}$, but that will not be true in most scenarios, including for the family of growth-lag trade-offs in Eq. (7). Some geometric intuition can be obtained by noting that solutions $\boldsymbol{\phi}^*$ for an objective of the form Eq. (8) will lie on the vertices of a region formed by the intersection of a hyperplane with normal vector $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \ldots, \gamma_K)$ and the probability simplex. These vertices are points with nonzero probability for at most two phenotypes. For complex dwell time distributions in NP, it is possible that the optimal phenotypic distribution has more than two phenotypes. However, we consistently find at most two phenotypes for the exponentially distributed dwell times considered in our numerical simulations, presumably because this distribution has a single typical timescale.

The geometry of the growth-lag trade-off determines whether a generalist or specialist strategy is preferred. When is a single generalist phenotype favored over a mixture of two specialist phenotypes? Optimizing Eq. (4) for the single best phenotype solution, the lag ℓ_g of the generalist phenotype (provided it is finite) satisfies $d\gamma_P/d\ell|_{\ell_g} = \omega\gamma' F(\ell_g)$, where $F(\ell) = \int_\ell^\infty d\tau\, \Pi(\tau)$ for general growth-lag trade-offs γ_P and dwell-time distributions $\Pi(\tau)$ (Appendix D). Considering phenotypes in the neighborhood of ℓ_g, we find that a mixed strategy is preferred if (Appendix D)

$$-\left.\frac{d^2\gamma_P}{d\ell^2}\right|_{\ell_g} < \omega\gamma'[\Pi(\ell_g) + \gamma' F(\ell_g)]. \tag{10}$$

This condition implies that the curvature of the growth-lag trade-off around the generalist solution determines whether a pure or a mixed strategy is favorable. Moreover, the solutions on the two sides of the transition point [i.e., when equality holds in Eq. (10)] are discontinuous: extreme specialists are preferred whenever a mixed strategy is optimal (Appendix D).

The optimality of the arrester phenotype. A consistent observation in our numerical simulations is the selection of a perfect arrester phenotype (with infinite lag) whenever a mixed strategy is preferred. The selection of this arrester phenotype can be explained by a heuristic argument based on how the growth-lag trade-off scales at long lags and the tail probabilities of dwell times in NP (see Appendix E 1 for a longer discussion). Since the distribution of dwell times in NP decays exponentially, the marginal cost of increasing the lag of the arrester phenotype also falls off at the same exponential rate. On the other hand, the marginal benefit in preshift growth rate declines more slowly according to a power-law given by Eq. (7). Thus, for long lags, any increase in lag of the arrester delivers a net positive marginal benefit, implying that it is always beneficial to increase the lag indefinitely.

A minimal metabolic model reproduces a power-law-like trade-off between preshift growth rate and postshift lag duration. How might a power-law-like growth-lag trade-off arise? In a model where different nutrients are metabolized by independent pathways, it is not clear why a strong growth-lag trade-off should pertain to a nutrient switch. The allocation of catalytic enzymes for the two pathways could be independent, allowing growth and lag to be decoupled and separately optimized. However, studies in bacteria [24,25] reveal that cells primed to grow quickly on a glycolytic nutrient experience a lag upon switching to a gluconeogenic nutrient. In particular, Basan *et al.* [24] report a reciprocal relationship between the steady-state growth rate in different glycolytic substrates and the corresponding adaptation times after a shift to acetate. Schink *et al.* [25] provide a mechanistic interpretation associated with flux reversal in the glycolysis and gluconeogenesis pathways. Guided by these observations, we developed a minimal metabolic model (Fig. 5) that results in a power-law-like trade-off between preshift growth rates and postshift lag durations.

In our model, two nutrients are processed along the same pathway, but the metabolic flux is controlled by catalytic enzymes that favor opposing directions [Fig. 4(a)]. A key feature of this model is that after a nutrient switch, the imbalance in enzyme abundances that enabled flux in one direction must be reversed to allow flux in the other. This reversal in turn requires biomass production, which is slow until flux can be reversed. The coupling between flux reversal and biomass production consequently leads to a long lag. Thus, this minimal model is sufficient to obtain a growth-lag trade-off.

We quantify growth in the preferred environment (P) by the steady-state growth rate μ_P and define lag, ℓ, in the nonpreferred environment (NP) as the time for resumed growth to exceed half the NP steady-state growth rate μ_{NP} (i.e., ℓ is the first-passage time for post-shift growth rate $\geqslant \mu_{NP}/2$). The model tracks two central metabolites, g and n, interconverted by two opposing enzyme activities: φ catalyzes $g \to n$ and ζ catalyzes $n \to g$ as shown in Fig. 4(a). Metabolic fluxes are Michaelis–Menten, growth rate via biomass production μ is a cooperative function of g and n, and enzyme levels relax toward environment-specific set points via

$$\frac{d\varphi}{dt} = \mu\left(\varphi_{ss}^{(\text{nutrient})} - \varphi\right), \quad \frac{d\zeta}{dt} = \mu\left(\zeta_{ss}^{(\text{nutrient})} - \zeta\right),$$

where the instantaneous growth rate is $\mu = (\frac{n}{n+K_{BM}})^\nu (\frac{g}{g+K_{BM}})^\nu$, and $\varphi_{ss}^{(\text{nutrient})}$ and $\zeta_{ss}^{(\text{nutrient})}$ are the steady-state concentrations of enzymes φ and ζ for a specified nutrient, respectively. For simplicity, the system is taken to be symmetric with respect to the two metabolites, except for the externally supplied nutrients. The rate constant r governs the influx into metabolites g and n, contributing rG and rN

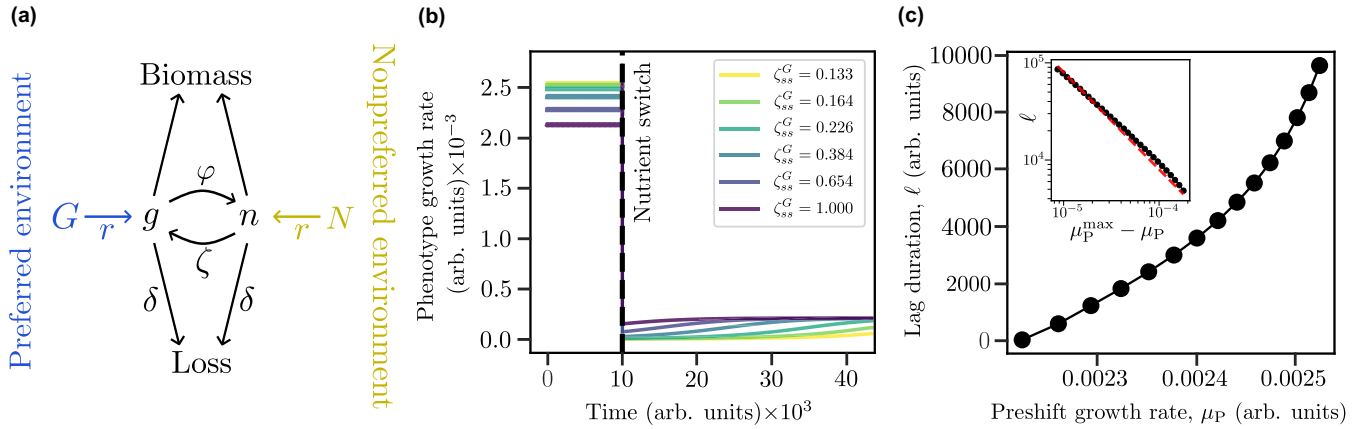


FIG. 4. A minimal metabolic model with a growth-lag trade-off. (a) Schematic showing two central metabolites g and n where enzymes φ catalyze $g \to n$, and enzymes ζ catalyze $n \to g$. In the preferred environment, nutrient G is present and N is absent. Allocating enzymes to high φ achieves a high growth rate in this environment. However, after a switch to the nonpreferred environment, pre-existing high φ prolongs the reaction $g \to n$, which delays net flux reversal, producing a lag in growth until φ falls and ζ rises. (b) Time courses around the nutrient switch (vertical dashed line) showing the growth rates of different enzyme-allocation phenotypes. We define "lag" to end when the postshift growth rate exceeds half the steady-state growth rate μ_{NP} in the nonpreferred environment. (c) Growth-lag trade-off curve generated by varying preshift enzyme allocation, keeping $\varphi_{\mathrm{ss}}^{G} = 1$ fixed while decreasing ζ_{ss}^{G} from 1 to $10^{-1.8}$. This monotonically increases preshift growth rate μ_{P} but prolongs postshift lag ℓ. The inset shows the predicted power-law scaling with exponent -1 (red dashed line) for $\nu = 2$. Details of the model are given in Appendix G, and the parameter values are listed in Table I.

from G and N, respectively. We initialize the system at steady state in P with G at a fixed nonzero concentration G_{P} and $N = 0$, then impose an instantaneous environmental switch to NP by changing nutrient availability so that $G = 0$ and N has nonzero concentration $N_{\mathrm{NP}} < G_{\mathrm{P}}$.

After the nutrient switch, high pre-existing φ suppresses the now-desirable net flux from $n \to g$ and leads to excess accumulation of n. This depletes cellular resources and delays the reversal of net flux. Slow biomass production gradually rebalances the enzymes, decreasing φ and increasing ζ, allowing the growth rate μ to rise to its steady-state value in NP. This rise time is an effective postshift lag.

For each value of $\varphi_{\mathrm{ss}}^{G}$, we vary ζ_{ss}^{G} to obtain a growth-lag trade-off. The values of $\varphi_{\mathrm{ss}}^{N}$ and ζ_{ss}^{N} are kept fixed. As described below, we find that the metabolic model operates in at least three qualitatively distinct regimes for different values of $\varphi_{\mathrm{ss}}^{G}$, depending on how the system produces new enzymes ζ after the switch from P to NP. Importantly, we obtain a power-law growth-lag trade-off for intermediate values of $\varphi_{\mathrm{ss}}^{G}$. In Fig. 4(b), we show specific time courses of the growth rate for different enzyme allocations (recoverer-like versus arrester-like) in this intermediate regime of $\varphi_{\mathrm{ss}}^{G}$.

First, when $\varphi_{\mathrm{ss}}^{G}$ is sufficiently small (and preshift growth is minimal), the depletion of g is gradual such that there is some growth (and thus production of ζ) for a short period after the switch from P to NP. The brief production of ζ is sufficient to avoid a prolonged lag phase as any nonzero ζ helps reverse the flux from $n \to g$, promoting growth and subsequent production of ζ. That is, the system avoids a lag phase at the expense of minimal preshift growth.

Second, when $\varphi_{\mathrm{ss}}^{G}$ is sufficiently large, the value of ζ_{ss}^{G} that maximizes growth in P is nonzero (Fig. 10). This is because a large $\varphi_{\mathrm{ss}}^{G}$ rapidly depletes g even with influx from G in P. It is beneficial to oppose some of the flux from $g \to n$ via ζ to maintain a pool of g and n and promote growth. The preshift growth rate μ_{P} decreases quadratically with small deviations of $\Delta\zeta_{\mathrm{ss}}^{G}$ from the growth-optimal allocation; that is, $\Delta\mu_{\mathrm{P}} \sim -(\Delta\zeta_{\mathrm{ss}}^{G})^2$. Since ζ_{ss}^{G} is nonzero, the postshift lag ℓ is finite even when preshift growth μ_{P} is maximal and ℓ decreases linearly with increasing $\Delta\zeta_{\mathrm{ss}}^{G}$: $\Delta\ell \sim -\Delta\zeta_{\mathrm{ss}}^{G}$. These two relations then lead to a quadratic relationship $\Delta\mu_{\mathrm{P}} \sim -\Delta\ell^2$ near the growth-optimal ζ_{ss}^{G} for a given $\varphi_{\mathrm{ss}}^{G}$. While the lag at maximal growth is finite, maintaining a large pool of enzymes $\varphi_{\mathrm{ss}}^{G}$ and ζ_{ss}^{G} would incur a significant energetic cost. That is, the system could optimize preshift growth and still have finite lag after the switch, but only by paying the energetic and resource cost of simultaneous expression of φ and ζ, which would lead to futile catalytic cycles [25].

Third, for intermediate values of $\varphi_{\mathrm{ss}}^{G}$, the growth-optimal allocation is $\zeta_{\mathrm{ss}}^{G} = 0$ (Fig. 10). The rapid depletion of g after the switch and the lack of any conversion from $n \to g$ leads to infinite lag when $\zeta_{\mathrm{ss}}^{G} = 0$. Any small, nonzero ζ_{ss}^{G} decreases growth rate linearly, $\Delta\mu_{\mathrm{P}} \sim -\zeta_{\mathrm{ss}}^{G}$. Importantly, we show analytically that $\ell \sim (\zeta_{\mathrm{ss}}^{G})^{1-\nu}$ to first order in ζ_{ss}^{G} (for $\nu > 1$). This relation arises from the slow production of ζ after the switch. The two relationships, $\Delta\mu_{\mathrm{P}} \sim -\zeta_{\mathrm{ss}}^{G}$ and $\ell \sim (\zeta_{\mathrm{ss}}^{G})^{1-\nu}$, together imply a power-law-like growth-lag trade-off $\Delta\mu_{\mathrm{P}} \sim -\ell^{-\alpha}$ with exponent $\alpha = (\nu - 1)^{-1}$ (see Appendix G

TABLE I. Values for parameters used in the metabolic model.

Parameters	Values
G_{P}, N_{P}	2.0, 0.0 arb. units
G_{NP}, N_{NP}	0.0, 1.0 arb. units
r, δ	1.0, 3.0 arb. units of 1/time
K_g, K_n, K_{BM}	1.25, 2.0, 0.75 arb. units
ν	2

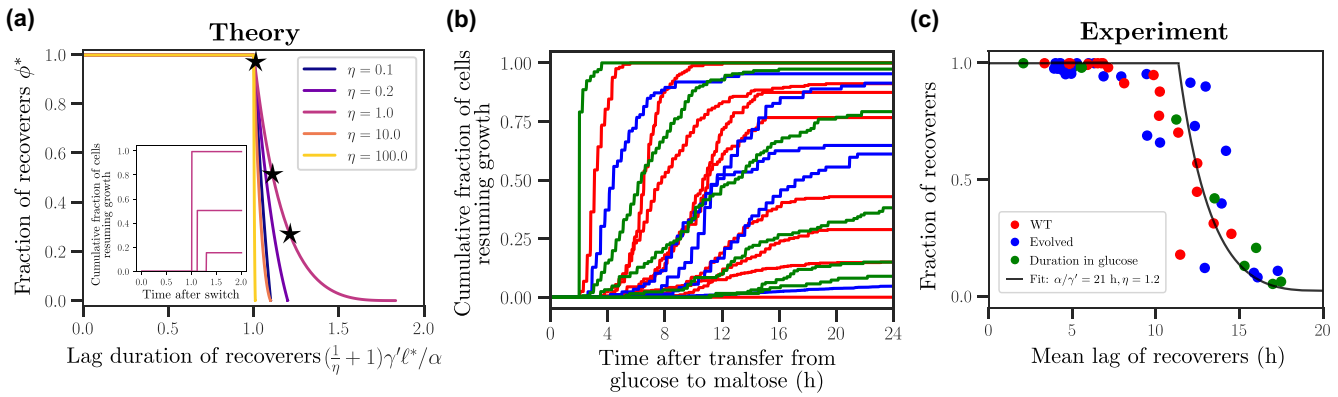


FIG. 5. Comparison of theoretical predictions and experimental data for fraction of recoverers versus lag duration. (a) Theoretical relationship between the optimal fraction of recoverers and their lag duration for different values of $\eta = \gamma'/\omega'$. The lag duration is rescaled using parameters η, γ', and α. The inset shows Theoretical curves showing the single-cell lag duration after a switch from the preferred environment to the nonpreferred environment. Different curves correspond to different points (black stars) on the curve for $\eta = 1$. Units for the x axis match those of the main figure. (b) Experimental measurements of yeast single-cell lag duration after a switch from glucose to maltose, recorded via time-lapse microscopy. Each curve corresponds to the cumulative fraction of cells resuming growth after the switch. Red curves represent measurements for different wild strains, blue curves represent measurements for evolved isolates that experienced repeated cycling of glucose-maltose media, and green curves correspond to populations that experienced different preshift durations in glucose. Only a subset of these traces is shown here due to limited space. See Fig. 7 for more. Data from Refs. [16,17]. (c) Experimental plot of the fraction of recoverers against the mean lag duration of the recoverers. Red, blue, and green dots represent, respectively, wild strains, evolved isolates, and variations based on preshift durations in glucose. The solid black curve is a fit to our theoretical prediction [Eq. (11)] relating the recoverer fraction to the lag duration, yielding $\alpha/\gamma' = 21$ h and $\eta = 1.2$.

for more details). This predicted power-law relationship was verified in simulations [Fig. 4(c)].

Theory predicts a relationship between the fraction of recoverers and their lag. Our theory implies that for every choice of environmental statistics (specified by ω and ω' in our model), there is a unique phenotype distribution that maximizes long-term growth. In other words, optimality requires that different features of the phenotype distribution (such as the fraction of recoverers and their lag) are not independent but rather co-vary with variations in environmental statistics. To derive these constraints, we now consider a simplified model where the phenotype distribution is restricted to at most two phenotypes. In this model, the phenotype distribution is parametrized by two quantities: the lag of the recoverer (denoted ℓ) and the fraction of recoverers (denoted ϕ). The remaining fraction $1-\phi$ of arresters has infinite lag. For such a two-phenotype system, cells can select ϕ in P by stochastically switching between recoverer and arrester phenotypes. The difference in preshift growth rates between the arrester and recoverer phenotypes and the switching rates fix ϕ at steady-state (Appendix F).

Optimality of the phenotype distribution predicts that ℓ and ϕ co-vary according to a nontrivial relationship across different environmental statistics. For exponential dwell times in NP and the growth-lag trade-off in Eq. (7), we derived an analytical relationship between the optimal fraction ϕ^* of recoverers and their optimal lag ℓ^* (for details, see Appendix E 2). Whenever $0 < \phi^* < 1$, the relationship is given by

$$\frac{\gamma'\ell^*}{\alpha} = \frac{\eta\, {}_2F_1\left(2, \frac{1}{\eta}+1; \frac{1}{\eta}+2; -\frac{1-\phi^*}{\phi^*}\right)}{\phi^*(1+\eta)\, {}_2F_1\left(1, \frac{1}{\eta}; \frac{1}{\eta}+1; -\frac{1-\phi^*}{\phi^*}\right)}, \qquad (11)$$

where ${}_2F_1$ is a hypergeometric function, and recall that $\eta = \gamma'/\omega'$. As $\phi^* \to 1$, we have $\gamma'\ell^*/\alpha = \eta/(1+\eta)$ and as $\phi^* \to 0$, $\gamma'\ell^*/\alpha = \min\{1, \eta\}$, which delineate the two boundaries of the coexistence region.

Figure 5(a) illustrates the predicted relationship between the optimal fraction of recoverers (ϕ^*) and their optimal lag (ℓ^*) for different values of η. We obtain a family of curves that show a consistent inverse relationship: ℓ^* initially increases while maintaining a single recoverer population ($\phi^* = 1$) with decreasing switching rate ω from P to NP. For a sufficiently small value of ω, the population splits into two phenotypes ($0 < \phi^* < 1$) while the lag of the recoverer continues to grow until the population again consists of a single phenotype ($\phi^* = 0$). The ratio between the values of ℓ^* at the two boundaries of the coexistence region, $\phi^* \to 0$ and $\phi^* \to 1$, is at most two (achieved when $\eta = 1$), imposing a constraint on the possible shapes in this family of curves.

An important point is that the relationship between the rescaled lag (in units of α/γ') and fraction of recoverers is only mediated by the parameter η (the expected number of generations in NP of a no-lag phenotype). It is agnostic to factors such as genetic variability among different strains and dwell time distributions in the preferred environment, making it broadly applicable across diverse experimental conditions and yeast strains. We can thus quantitatively test this prediction for various environmental conditions and genetically distinct yeast populations using existing datasets.

Experimental evidence for a relationship between fraction of recoverers and their lag. We analyzed experimental data from single-cell measurements of yeast populations undergoing a glucose-to-maltose transition [Fig. 5(b)]. These data report the cumulative fraction of cells resuming growth after a switch from media containing glucose to media

containing maltose for up to 24 hours after the switch [17]. In Fig. 5(b), each individual red curve represents lag measurements for a different wild strain of yeast, blue curves represent measurements taken for evolved isolates of a long-lag strain in parallel cultures after experiencing 6 to 8 cycles of alternating glucose and maltose-containing media, and green curves represent measurements for populations that experienced different glucose pregrowth durations. From the data, we obtained distributions of lag durations for each tracked population.

A consistent feature is the bimodality of these distributions, delineating two distinct subpopulations of recoverers, which resumed growth before the experiment ended, and arresters, which remain arrested. We calculated the average lag duration of the recoverers and their corresponding fraction ϕ within each population. We observed a robust inverse relationship where populations with longer lags had smaller fractions of recoverers [Fig. 5(c)]. This empirical trend matches well with predictions from our two-phenotype model. Treating η and α/γ' as fitting parameters, we obtained best-fit values of $\eta = 1.2$ and $\alpha/\gamma' = 21$ h. Using the experimental estimate $\gamma' = 0.1\ \mathrm{h}^{-1}$ from Ref. [17], these values correspond to $\alpha = 2.1$ and $\omega' = 0.08\ \mathrm{h}^{-1}$. See Fig. 9 for the phase diagram for the optimal number of phenotypes with these physiological and environmental parameters.

IV. DISCUSSION

Microbial populations repeatedly exposed to environmental shifts must reconcile fast growth in the current niche with the capacity to resume growth after a switch. In this study, we investigated the role of phenotypic heterogeneity in microbial populations, focusing on the adaptation of yeast to changing carbon sources. The central experimental observation motivating this work is that following a switch from a glucose-rich environment to a maltose-rich environment (or other environments), a substantial fraction of genetically identical yeast cells remain arrested in lag phase for the entire duration of the experiment [15–18]. This persistent growth arrest suggests that cells in this subpopulation occupy a distinct physiological state, raising the question of whether this heterogeneity might reflect an evolved strategy to optimize population-level fitness in fluctuating environments. To address this, we developed a general framework based on dynamic programming to determine the optimal bet-hedging strategies for microbial populations transitioning between environments where cells can vs cannot readily reorganize gene expression. We found that the optimal phenotype distribution consists of at most two discrete phenotypes, even when a continuous range of phenotypes is allowed. Under certain conditions, the population adopts a bet-hedging strategy composed of two distinct specialists: one phenotype minimizes lag to recover quickly in the nonpreferred environment (NP), while the other maximizes growth in the preferred environment (P) but never recovers in NP. In other regimes, a single optimal generalist phenotype emerges that balances growth and lag.

A key prediction from our model is a nontrivial relationship between the fraction of recoverers and their lag. This relationship is a consequence of imposing optimality of the phenotype distribution for a given physiological trade-off across different environmental transition rates. This prediction matches single-cell datasets spanning wild strains, evolved lines, and varying pregrowth nutrient durations, suggesting an underlying constraint in the genetic circuit that mediates phenotypic diversification.

Two assumptions underpin our results that populations benefit from having phenotypic heterogeneity with distinct discrete phenotypes: First, the population can retune its phenotype distribution rapidly in the preferred environment (P), but phenotypes are locked during lag in the nonpreferred environment (NP). Second, a monotonic trade-off links higher preshift growth to longer postshift lag. Other ingredients of the model, such as exponentially distributed dwell times, the exact form of the trade-off, and the perfect arrester limit ($\ell \to \infty$), allow analytic tractability but are not required for heterogeneity to be favored. More complex dwell-time distributions could create regimes where more than two distinct specialists make up the optimal strategy, each tuned to a different characteristic timescale of environmental shifts. We note that single-cell lag variability is not always strongly bimodal. In some bacterial contexts, lag times can span wide distributions, sometimes with long tails. For example, *Escherichia coli* can exhibit broad lag-time distributions following nutrient transitions or starvation, with growth resumption dominated by short-lag subpopulations, while long-lag cells can contribute to stress tolerance [33,34]. Our prediction of at most two discrete phenotypes should be interpreted as a statement about the optimal strategy within our framework's assumptions, namely, that there are two types of environments (P and NP) and that the distribution of times spent in the NP environment has a single timescale. Relaxing either assumption could lead to a broader optimal phenotype distribution.

Mechanistically, a minimal metabolic model motivates why a trade-off may exist: allocating enzymes to drive flux in the preshift direction (fast growth) delays net flux reversal and consequently delays resumed growth after the switch. Systematic variation of preshift enzyme allocation reproduces a power-law growth-lag trade-off curve consistent with the phenomenology encapsulated by Eq. (7). Prior work has also shown that growth-lag trade-offs can arise from constraints on flux reversal and proteome allocation in central metabolism [24,25,35,36], although the variation in preshift growth rates in these cases is due to different preshift glycolytic substrates rather than heterogeneous growth rates on a single substrate.

Our framework extends classical bet-hedging theory [4,7,8,13] in a key way: We couple preshift growth and postshift lag through an explicit metabolic constraint that encodes how proteome allocation before the switch creates a bottleneck in the speed of adaptation after the switch. This makes the trade-off mechanistic rather than just phenomenological. Our work is thus novel in that we link cellular physiology to population-level growth outcomes in fluctuating conditions.

Lag heterogeneity following carbon shifts in *S. cerevisiae* shares a key similarity with persister formation in *E. coli*, competence in *Bacillus subtilis*, and sporulation in filamentous fungi: all involve a slow-growing subpopulation that sacrifices immediate proliferation for future advantage. In these systems, phenotypic diversity provides a benefit to the population by being prepared for unfavorable conditions. The "recoverer" state in our model plays this protective role,

maintaining the population (i.e., against competitors) until favorable conditions return.

Our analysis suggests that the recoverer lag-fraction relationship can function as an evolutionary "rheostat," providing populations with a mechanism to tune their balance between fast growth and lag. In our metabolic model, such a rheostat could be wired through regulatory control of proteome allocation. Changes in the expression of enzymes that bias flux toward preshift growth would continuously shift both the growth rate in the preferred environment and the lag after a nutrient switch. This would enable populations to smoothly adjust their phenotypic distribution in response to the statistics of environmental change, without requiring rewiring of their metabolic circuitry. More broadly, any modulation (genetic or otherwise) of upstream regulatory pathways that control the phenotype distribution will then satisfy an optimality constraint that can be tuned to the rate of environmental change, offering a convenient evolutionary rheostat for rapidly adapting to new environments. In short, with a single control knob, closely related strains constrained by the same trade-off could flexibly adapt to different environmental fluctuation statistics.

Promising directions for future experimental work include (i) varying ω and ω' (i.e., in a microfluidic setup) to test the predicted re-entrant generalist-specialist-generalist behavior, (ii) adjusting preshift enzyme allocation (genetically or via growth history) to test the coupling between growth and lag, and (iii) measuring the survival of long-lag arresters to test for a "perfect arrester" state. Together, these experiments would show whether the recoverer fraction-lag relation functions as an evolutionary rheostat across strains.

ACKNOWLEDGMENTS

We thank Piyush Nanda and Andrew Murray for many fruitful discussions. This work was supported in part by Princeton University through the Center for the Physics of Biological Function, and in part by the National Institutes of Health under Award No. R01GM082938 (N.S.W.). G.R. was partially supported by a joint research agreement between NTT Research Inc. and Princeton University.

The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

DATA AVAILABILITY

The data and code used to generate all figures are available online [37].

APPENDIX A: EXPERIMENTAL DATA

Data for single-cell lag phase measurements were taken from New *et al.* [17] and Cerulus *et al.* [16]. In short, cells were pregrown in glucose, harvested, and transferred to media containing maltose. The lag duration of individual cells was measured using automated time-lapse microscopy, defined as the time from transfer until the first observable sign of division (bud emergence or resumption of bud growth). For experimental details, refer to New *et al.* [17]. To generate the distribution of single-cell lag durations [Fig. 1(b)], data were specifically taken from Dataset S1 of New *et al.* [17]. It contains single-cell lag phase measurements for 18 strains of *S. cerevisiae*, and we picked the distribution for strain "DBVPG1373" for Fig. 1(b) in the main text. Figure 6 contains the lag distributions of other strains from their dataset, which highlights a bimodality in lag durations. For Fig. 5(b) in the main text, we took single-cell lag data for wild yeast strains and evolved strains from Dataset S1 and Dataset S2 of New *et al.* [17] [Figs. 7(a) and 7(b)], respectively, and for varying preshift durations in glucose from Cerulus *et al.* [16] [Fig. 7(c)], to estimate the lag and fraction of recoverers in a population. Raw single-cell data from Ref. [16] presented in Fig. 7(c) were shared through personal communication by Cerulus and Verstrepen.

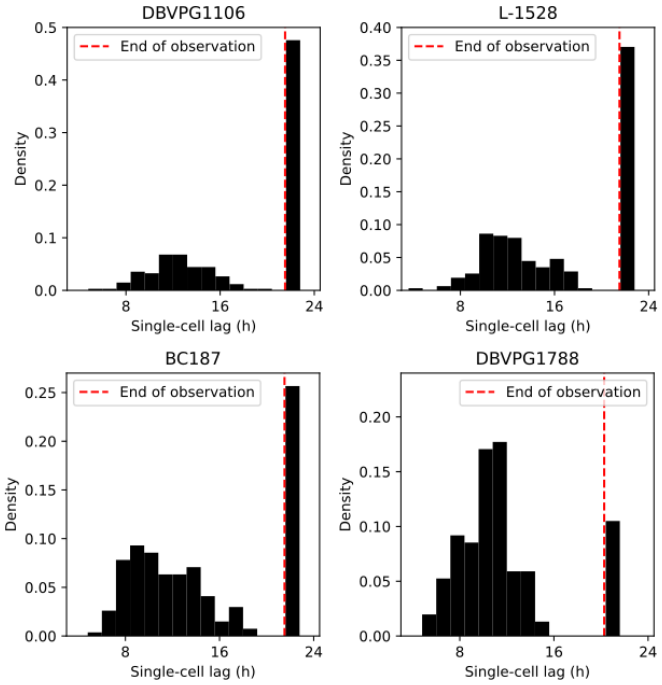


FIG. 6. Single-cell lag distributions of different yeast strains.

Most single-cell lag distributions are bimodal, i.e., composed of a group of recoverer cells that recover from lag and arrester cells, which do not resume growth until the end of the experiment. We compute the mean lag duration of the recoverer population for each strain and plot it against its fraction in the population.

APPENDIX B: CELLULAR DECISION-MAKING FRAMEWORK

1. Phenotype heterogeneity in fluctuating environments

We consider a scenario where a cellular population experiences transitions between a preferred environment (P) and a nonpreferred environment (NP). Cells can take any one of K possible phenotypes. For example, a cell with a phenotype that allows for rapid growth in P may experience a longer lag phase or greater death rate when the environment switches to NP. In P, cells grow and rapidly switch their phenotypes so as to maintain a phenotype distribution, $\boldsymbol{\phi} = (\phi_1, \phi_2, \ldots, \phi_K)$ (with $\sum_k \phi_k = 1$), within the population. The selection of a

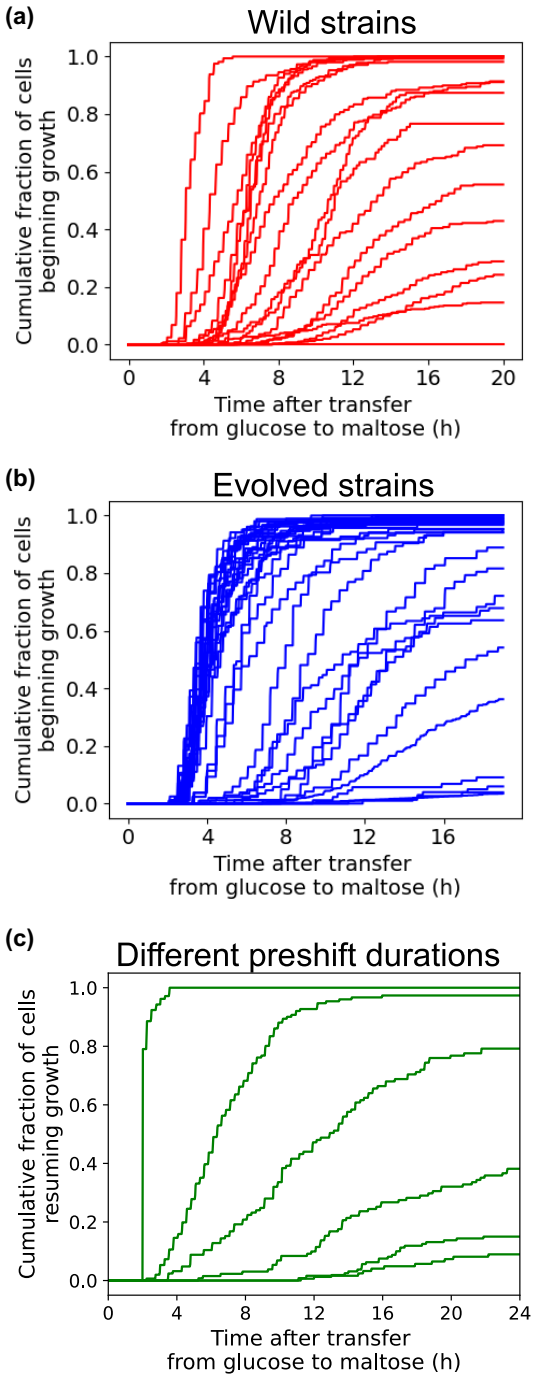


FIG. 7. (a) Cumulative fraction of cells resuming growth after the nutrient switch for wild strains of yeast. (b) Cumulative fraction of cells resuming growth after the nutrient switch for evolved populations of yeast. (c) Cumulative fraction of cells resuming growth after the nutrient switch for populations experiencing varying preshift durations in glucose.

phenotype distribution could occur through stochastic switching between phenotype states [38]; we discuss this further in Appendix F. The distinction between P and NP is that cells cannot choose their phenotype distribution in NP. This corresponds to the case where cells experience a nongrowing phase (lag, starvation, or death) in NP. In this nongrowing phase, the population cannot readily express genes and synthesize new proteins, and thus cannot arbitrarily choose its phenotype distribution.

The state of the environment is a stochastic process. The instantaneous rate at which the population switches from P to NP (denoted ω) or from NP to P (denoted ω') could have an arbitrary dependence on the history of environmental states (for example, the time since the last switch from NP to P). Cells in P can anticipate the transition from P to NP and choose their phenotype distribution $\boldsymbol{\phi}$ so as to maximize the long-term growth rate of the population. In the main text, we present a simplified version where cells make decisions that are independent of history. In this simplified version, the environment switches from P to NP and NP to P at constant rates ω and ω', respectively. Here, we present a general version of the framework where cellular decisions are allowed to be history-dependent.

2. Model setup

We developed a general optimization framework based on dynamic programming to investigate how phenotype heterogeneity influences population fitness under the scenario described in the previous section. This framework allows us to model cellular decision-making using methods similar to those used in reinforcement learning and control theory.

Let $s(t)$ represent the state at time t and $\gamma(s(t), k)$ denote the instantaneous growth rate of a cell with phenotype k in state $s(t)$. $s(t)$ includes information about whether the current environment is P or NP, and possibly additional information about the growth medium in P (which will affect instantaneous growth rate), the time since the most recent shift from P to NP (which is important to incorporate lag), or historical information the cell uses to better anticipate future environmental transitions. The probability of transitioning from state s to s' in interval dt is given by the transition probability $P(s(t+dt) = s'|s(t) = s)$. This transition probability is time-translation invariant, i.e., there is no dependence on absolute time except through s. Note that the framework applies for a general state space $s(t)$, and it can thus capture potentially complex historical dependencies.

Given $s(t)$ and the phenotype distribution $\boldsymbol{\phi}(t)$, the instantaneous growth rate of a population of N cells is

$$\frac{1}{N}\frac{dN}{dt} = \sum_{k=1}^{K} \phi_k(t)\gamma(s(t), k) \equiv \bar{\gamma}(s(t), \boldsymbol{\phi}(t)). \tag{B1}$$

If the initial number of cells is N_0 and the number of cells after time T is N_T, the long-term growth rate is

$$\Gamma_T \equiv \frac{1}{T}\ln\frac{N_T}{N_0} = \frac{1}{T}\int_0^T \langle\bar{\gamma}(s(t), \boldsymbol{\phi}(t))\rangle dt, \tag{B2}$$

where the angled brackets denote an expectation over changes in environmental state. Our goal is to find the optimal strategy to select $\boldsymbol{\phi}$ in P such that Γ_T is maximized as $T \to \infty$.

Instead of directly maximizing $T^{-1}\int_0^T \langle\bar{\gamma}(s(t), \boldsymbol{\phi}(t))\rangle dt$ as $T \to \infty$, we find the strategy that maximizes $\int_0^\infty e^{-\mu t}\langle\bar{\gamma}(s(t), \boldsymbol{\phi}(t))\rangle dt$ as $\mu \to 0$. This alternative objective is commonly used to remove time as a state variable when computing the optimal strategy. No singularities arise in the

limit $\mu \to 0$ as the decision to pick a particular $\boldsymbol{\phi}$ matters only until the next time the population encounters P.

A population in P selects its phenotype distribution $\boldsymbol{\phi}$ at every interval dt (where $1/dt$ is much larger than the typical growth rate of the cells in P). We now derive our central optimization objective. We write down a (Bellman) dynamic programming equation for the long-term growth rate given that the population is in state s in environment P, assuming the population makes the best possible choice for the phenotype distribution $\boldsymbol{\phi}$. Let us call this optimal long-term growth rate $\Gamma_{\mathrm{P}}(s)$. We have

$$\Gamma_{\mathrm{P}}(s) = \max_{\boldsymbol{\phi}} \left\{ \int_0^\infty dt e^{-\mu t} \langle \overline{\gamma}(s(t), \boldsymbol{\phi}(t)) \rangle_{s(0)=s, \boldsymbol{\phi}(0)=\boldsymbol{\phi}} \right\}, \tag{B3}$$

where the expectation (represented by the angled brackets) is over possible future states given the current state $s(0) = s$ and the phenotype distribution $\boldsymbol{\phi}$. Splitting the integral,

$$\Gamma_{\mathrm{P}}(s) = \max_{\boldsymbol{\phi}} \left\{ \overline{\gamma}_{\mathrm{P}}(s, \boldsymbol{\phi}) dt + \int_{dt}^\infty dt e^{-\mu t} \langle \overline{\gamma}(s(t), \boldsymbol{\phi}(t)) \rangle_{s(dt), \boldsymbol{\phi}(dt)|s(0)=s, \boldsymbol{\phi}(0)=\boldsymbol{\phi}} \right\}, \tag{B4}$$

where the expectation is now over the state of the population at dt given that at $t = 0$ it is at $s(0) = s$, $\boldsymbol{\phi}(0) = \boldsymbol{\phi}$. Setting $t = t' + dt$, we have

$$\Gamma_{\mathrm{P}}(s) = \max_{\boldsymbol{\phi}} \left\{ \overline{\gamma}_{\mathrm{P}}(s, \boldsymbol{\phi}) dt + e^{-\mu dt} \int_0^\infty dt' e^{-\mu t'} \langle \overline{\gamma}(s(t'), \boldsymbol{\phi}(t')) \rangle_{s(0), \boldsymbol{\phi}(0)|s(-dt)=s, \boldsymbol{\phi}(-dt)=\boldsymbol{\phi}} \right\}, \tag{B5}$$

where we used the fact that the transition matrix $P(s(t + dt), \boldsymbol{\phi}(t + dt)|s(t), \boldsymbol{\phi}(t))$ is time-translation invariant.

Given that the population is in P at state s, there are two possibilities: after dt, either the environment stays in P with probability $1 - \omega(s)dt$, or switches to NP with probability $\omega(s)dt$. If the environment stays in P and the state transitions to s', then the expected long-term growth rate after dt is $\Gamma_{\mathrm{P}}(s')$. If the environment transitions to NP and the state transitions to s', we denote the expected long-term growth rate by $\Gamma_{\mathrm{NP}}(s', \boldsymbol{\phi})$ (which we specify further below). We split the integral in the above equation based on the probability of these two possibilities,

$$\Gamma_{\mathrm{P}}(s) = \max_{\boldsymbol{\phi}} \{ \overline{\gamma}_{\mathrm{P}}(s, \boldsymbol{\phi}) dt + \{1 - [\omega(s) + \mu]dt\} \langle \Gamma_{\mathrm{P}}(s') \rangle + \omega(s) dt \langle \Gamma_{\mathrm{NP}}(s', \boldsymbol{\phi}) \rangle \}, \tag{B6}$$

where we have dropped the subscripts specifying that the expectations are over the state and phenotype distribution at the subsequent time step, i.e., s', $\boldsymbol{\phi}'$, given s at the current time step. Importantly, the choice of $\boldsymbol{\phi}$ does not matter if the environment stays in P. Moreover, the term involving $\langle \Gamma_{\mathrm{P}}(s') \rangle$ can be ignored when computing the optimal strategy. To obtain the optimal strategy $\boldsymbol{\phi}^*$, we simply have

$$\boldsymbol{\phi}^*(s) = \arg\max_{\boldsymbol{\phi}} \{ \overline{\gamma}(s, \boldsymbol{\phi}) + \omega(s) \langle \Gamma_{\mathrm{NP}}(s', \boldsymbol{\phi}) \rangle \}. \tag{B7}$$

Now, we unpack $\Gamma_{\mathrm{NP}}(s', \boldsymbol{\phi})$. Let $\Pi(\tau)$ be the probability density that the environment stays in NP for duration τ and let s'' be the state that the environment ends up in after the switch back to P after interval τ. We have

$$\Gamma_{\mathrm{NP}}(s', \boldsymbol{\phi}) = \int_0^\infty d\tau \, \Pi(\tau) \left(\int_0^\tau dt \, e^{-\mu t} \langle \overline{\gamma}_{\mathrm{NP}}(s(t), \boldsymbol{\phi}(t)) \rangle_{s(t)} + e^{-\mu \tau} \langle \Gamma_{\mathrm{P}}(s'') \rangle_{s''} \right), \tag{B8}$$

where $s(t)$ and $\boldsymbol{\phi}(t)$ are now the state and distribution of phenotypes at time t after the switch to NP given that the state and distribution were s' and $\boldsymbol{\phi}$ when the switch occurred. The first term on the right-hand side does depend on μ, but it does not contribute to the calculation of $\boldsymbol{\phi}^*$ if $\Pi(\tau)$ decays to zero on timescales much smaller than $1/\mu$. The second term on the right-hand side contains $\langle \Gamma_{\mathrm{P}}(s'') \rangle$ which is the expected long-term growth rate once the environment switches back to P. This term does not depend on $\boldsymbol{\phi}$ and can therefore be ignored when finding the optimal $\boldsymbol{\phi}$. We arrive at

$$\boldsymbol{\phi}^*(s) = \arg\max_{\boldsymbol{\phi}} \left\{ \overline{\gamma}_{\mathrm{P}}(s, \boldsymbol{\phi}) + \omega(s) \left\langle \int_0^\infty d\tau \, \Pi(\tau) \left(\int_0^\tau dt \, \bar{\gamma}_{\mathrm{NP}}(s(t), \boldsymbol{\phi}(t)) \right) \right\rangle_{s(t)} \right\}. \tag{B9}$$

Interchanging integrals and defining $F(t) = \int_t^\infty d\tau \Pi(\tau)$, we rewrite the above as

$$\boldsymbol{\phi}^*(s) = \arg\max_{\boldsymbol{\phi}} \left\{ \overline{\gamma}_{\mathrm{P}}(s, \boldsymbol{\phi}) + \omega(s) \left\langle \int_0^\infty dt F(t) \overline{\gamma}_{\mathrm{NP}}(s(t), \boldsymbol{\phi}(t)) \right\rangle_{s(t)} \right\}. \tag{B10}$$

The average growth rate after time t in NP is

$$\overline{\gamma}_{\mathrm{NP}}(s(t),\boldsymbol{\phi}(t)) = \frac{\sum_{k=1}^{K}\phi_k\gamma_{\mathrm{NP}}(s(t'),k)e^{\int_0^t dt'\gamma_{\mathrm{NP}}(s(t'),k)}}{\sum_{k'=1}^{K}\phi_{k'}e^{\int_0^t dt''\gamma_{\mathrm{NP}}(s(t'),k')}} \tag{B11}$$

$$= \frac{d}{dt}\ln\left(\sum_{k=1}^{K}\phi_k e^{\int_0^t dt'\gamma_{\mathrm{NP}}(s(t'),k)}\right). \tag{B12}$$

Plugging this expression into Eq. (B10) and integrating by parts, we finally obtain the objective:

$$\boldsymbol{\phi}^*(s) = \arg\max_{\boldsymbol{\phi}}\left\{\overline{\gamma}_{\mathrm{P}}(s,\boldsymbol{\phi}) + \omega(s)\left\langle\int_0^\infty d\tau\,\Pi(\tau)\ln\left(\sum_{k=1}^{K}\phi_k e^{\int_0^\tau dt\gamma_{\mathrm{NP}}(s(t),k)}\right)\right\rangle_{s(t)}\right\} \quad \text{subject to } \sum_k\phi_k = 1 \text{ and } \phi_k \geqslant 0 \text{ for all } k. \tag{B13}$$

This expression has an intuitive interpretation: the optimal strategy balances the instantaneous growth rate in P and the expected long-term growth rate if the environment were to transition to NP. We now show how this framework can be used to find the optimal decision-making strategy in two scenarios involving a growth-lag trade-off and a growth-death trade-off.

3. Balancing growth and lag

Equation (B13) is applicable generally when cells transition between preferred and nonpreferred environments. In this section, we discuss the situation where a population of cells experiences rapid growth in the preferred environment (P) and undergoes lag phase before resuming growth in the nonpreferred environment (NP). In P, cells grow such that $\gamma(\mathrm{P},k) = \gamma_k$. In NP, cells experience lag before they can resume growth: $\gamma(\tau_{\mathrm{NP}},k) = \gamma'\Theta(\tau_{\mathrm{NP}} - \ell_k)$, where τ_{NP} is the time spent in NP since the last switch from P and ℓ_k is the lag of phenotype k. We consider $\Pi(\tau) = \omega' e^{-\omega' t}$ and also ignore any other dependence on the state s. Interchanging the integrals over τ and t in Eq. (B13), we get

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}}\left\{\sum_{i=1}^{K}\phi_i\gamma_i + \omega\,\omega'\int_0^\infty dt e^{-\omega' t}\ln\left[\sum_{i=1}^{K}\phi_i e^{\gamma'[t-\ell_i]_+}\right]\right\}. \tag{B14}$$

Making substitutions: $t' = \gamma' t$, $\ell_i' = \ell_i\gamma'$, $\tilde{\gamma}_i = \gamma_i/\omega$, and $\eta = \gamma'/\omega'$, we now have

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}}\left\{\sum_{i=1}^{K}\phi_i\tilde{\gamma}_i + \frac{1}{\eta}\int_0^\infty dt' e^{-\frac{t'}{\eta}}\ln\left[\sum_{i=1}^{K}\phi_i e^{[t'-\ell_i']_+}\right]\right\}. \tag{B15}$$

4. Balancing growth and death

In an alternative paradigm, rather than experiencing a lag phase in the nonpreferred environment (NP), cells are subjected to lethal conditions and die at phenotype-specific rates. In this framework, phenotype variation is characterized not by differences in lag durations, but by distinct death rates ξ_k associated with each phenotype in NP. In P, cells grow such that $\gamma(\mathrm{P},k) = \gamma_k$ and in NP, cells die such that $\gamma(\mathrm{NP},k) = \xi_k$, where $\xi_k(\leqslant 0)$ is the death rate for phenotype k. Similar to the previous section, we use Eq. (B13) and modify the second term concerned with long-term growth in NP,

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}}\left\{\sum_{i=1}^{K}\phi_i\gamma_i + \omega\,\omega'\int_0^\infty dt e^{-\omega' t}\ln\left[\sum_{i=1}^{K}\phi_i e^{\xi_i t}\right]\right\}. \tag{B16}$$

After making substitutions: $t' = \omega' t$, $\overline{\xi_i} = \xi_i/\omega'$, and $\tilde{\gamma}_i = \gamma_i/\omega$, we have

$$\boldsymbol{\phi}^* = \arg\max_{\boldsymbol{\phi}}\left\{\sum_{i=1}^{K}\phi_i\tilde{\gamma}_i + \int_0^\infty dt' e^{-t'}\ln\left[\sum_{i=1}^{K}\phi_i e^{\overline{\xi_i}t'}\right]\right\}. \tag{B17}$$

We determine the optimal phenotype distribution $\boldsymbol{\phi}^*$ numerically.

5. Numerical optimization

We assume the population contains a distribution over K phenotypes, where phenotype k (where k takes on values from 1 to K) has a growth rate γ_k ($\gamma_{\min} \leqslant \gamma_k \leqslant \gamma_{\max}$) in P. In the growth-lag scenario, phenotypes are characterized with lags ℓ_k in NP and all phenotypes have the same growth rate γ' in NP when cells recover from the lag phase. In the growth-death scenario, phenotypes are characterized by death rates ξ_k in NP. The phenotype distribution is given by $\boldsymbol{\phi} = (\phi_1,\ldots,\phi_k,\ldots,\phi_K)$.

Step-by-step procedure to run the optimization scheme:

(1) First, we pick the parametric form for the trade-off between growth and lag (or death). This gives us growth rates $\{\gamma_i\}$ in P and the corresponding lag durations $\{\ell_i\}$ (or death rates $\{\xi_i\}$) in NP.

(2) Choose transition rates ω and ω'.

(3) Initialize $\boldsymbol{\phi} = \{\phi_i\}$ such that $\phi_i = 1/K$. To prevent having to enforce an extra constraint of $\sum_{i=1}^{K} \phi_i = 1$ during optimization, we parametrize the probabilities as $\phi_i = e^{\beta_i}/(1 + \sum_{k=2}^{K} e^{\beta_k})$ for $2 \leqslant i \leqslant K$ and $\phi_1 = 1/(1 + \sum_{k=2}^{K} e^{\beta_k})$. Here $\{\beta_i\}$ are unconstrained parameters that we can directly optimize over.

(4) Using the Adam optimizer in the Optax library applied to a JAX implementation in Python, we numerically find $\{\beta_i\}$ that maximizes $\Gamma_{\mathrm{P}}(\boldsymbol{\phi})$ from Eq. (B15). We then obtain the optimal phenotype distribution $\boldsymbol{\phi}^*$.

APPENDIX C: OPTIMAL PHENOTYPE DISTRIBUTION CONSISTS OF A FEW DISCRETE PHENOTYPES

We consider our optimization problem where the population selects a phenotype distribution $\boldsymbol{\phi}$ to maximize a long-term growth objective of the general form

$$\Gamma(\boldsymbol{\phi}) = \max_{\boldsymbol{\phi}} \left\{ \sum_{i=1}^{K} \phi_i \tilde{\gamma}_i + \frac{1}{\eta} \int_0^\infty dt' e^{-\frac{t'}{\eta}} \times \ln \left[\sum_{i=1}^{K} \phi_i e^{[(t'-\ell_i')]_+} \right] \right\}, \tag{C1}$$

with the constraints:

(1) $\phi_i \geqslant 0 \; \forall \; i$,

(2) $\sum_{i=1}^{K} \phi_i = 1$.

For this constrained nonlinear optimization problem, we use the Karush-Kuhn-Tucker (KKT) conditions to make predictions about the structure of the optimal solution. We construct the corresponding Lagrangian function,

$$\mathcal{L}(\boldsymbol{\phi}, \lambda, \mu) = -\Gamma(\boldsymbol{\phi}) + \sum_{i=1}^{K} \mu_i \phi_i + \lambda \left(\sum_{i=1}^{K} \phi_i - 1 \right), \tag{C2}$$

for some Lagrange multiplier λ and KKT multipliers $\mu_i \geqslant 0$. To maximize $\Gamma(\boldsymbol{\phi})$ given the constraints, the stationarity condition requires

$$-\frac{\partial \mathcal{L}}{\partial \phi_k} = \frac{\partial \Gamma}{\partial \phi_k} - \mu_k - \lambda = 0. \tag{C3}$$

Let $g_k(\boldsymbol{\phi}) = \partial\Gamma(\boldsymbol{\phi})/\partial\phi_k$ denote the marginal value of phenotype k. Then the support of the optimal solution (i.e., the set of phenotypes with $\phi_k > 0$) is determined by the set of indices where $g_k = \lambda$. We get this from the complementary slackness condition which requires that $\mu_i \phi_i = 0$ for $i = 1, 2, \ldots, K$. Geometrically, this corresponds to the intersection points of a horizontal line at height λ with the graph of g_k as a function of phenotype lag ℓ_k'. Let us examine the explicit form of g_k:

$$g_k(\boldsymbol{\phi}) = \frac{\partial \Gamma(\boldsymbol{\phi})}{\partial \phi_k} = \tilde{\gamma}_k + \int_0^\infty dt' \Pi(t') \frac{e^{[t'-\ell_k']_+}}{\sum_{i=1}^{K} \phi_i e^{[t'-\ell_i']_+}}. \tag{C4}$$

In our framework, the rescaled lag durations ℓ' and preferred-environment growth rates $\tilde{\gamma}$ are strictly ordered such that: $\ell_1' < \ell_2' < \cdots < \ell_K'$ and $\tilde{\gamma}_1 < \tilde{\gamma}_2 < \cdots < \tilde{\gamma}_K$. This implies that the first term in $g_k(\boldsymbol{\phi})$ monotonically increases with k. In the second term, the dwell time distribution $\Pi(t')$ acts as a discount factor such that for larger t', $\Pi(t')$ is smaller since the distribution monotonically decreases with t'. This implies that the contribution to the integral for larger ℓ_k' is reduced which in turn means that the second term in $g_k(\boldsymbol{\phi})$ monotonically decreases with k. $g_k(\boldsymbol{\phi})$ is composed of the sum of two terms with opposing monotonicities. If g_k is strictly monotonic in ℓ_k', then the horizontal line intersects the curve at most once, implying that the optimal strategy is supported on a single phenotype. The number of times g_k intersects with the horizontal line determines the number of phenotypes with nonzero probability. For well-behaved $\Pi(t)$, we expect at most a few discrete phenotypes, although the precise number depends on $\Pi(t)$ and the growth-lag trade-off.

Single timescale in NP

In this section, we present a geometric argument for why we expect at most two phenotypes when the distribution of dwell times in NP is unimodal. This geometric argument complements the formal argument in the main text using the KKT theorem. Suppose that the time spent in NP is a constant t instead of being exponentially distributed. We have for K phenotypes (with arbitrary growth rates)

$$\Gamma(\boldsymbol{\phi}) = \sum_k \gamma_k^{\mathrm{P}} \phi_k + \ln \left(\sum_k \phi_k e^{\gamma_k^{\mathrm{NP}}(t)} \right), \tag{C5}$$

where γ_k^{P} is the growth rate in P and $\gamma_k^{\mathrm{NP}}(t)$ is the effective growth rate in NP for phenotype k. In the growth-lag scenario $\gamma_k^{\mathrm{NP}}(t) = \gamma'[t - \ell_k]_+$, while in the growth-death scenario $\gamma_k^{\mathrm{NP}}(t) = \xi_k t$.

Consider $\boldsymbol{\phi}$ such that $\sum_k \phi_k \gamma_k^{\mathrm{P}} = y$ and $\sum_k \phi_k e^{\gamma_k^{\mathrm{NP}}(t)} = x$. The two constraints $\phi_k \geqslant 0$ and $\sum_k \phi_k = 1$ define a closed feasible region in $\boldsymbol{\phi}$. This closed feasible region in $\boldsymbol{\phi}$ translates to a closed feasible region D in the two-dimensional (x, y) plane. Consider the family of curves $y = \Gamma - \ln x$ parametrized by Γ. As Γ is decreased from infinity, the optimal feasible solution is due to the curve that just touches D. Now we show that the boundary points of D correspond to ϕ such that at most two of the phenotypes have nonzero probability mass.

First note that the intersection of the hyperplane $\sum_k \phi_k e^{\gamma_k^{\mathrm{NP}} t} = x$ with the probability simplex $\phi_k \geqslant 0$, $\sum_k \phi_k = 1$ is a convex hull formed by points on line segments joining two vertices of the probability simplex. These are points that have nonzero probability for at most two of the phenotypes. Next, for a given value of x, the top and bottom boundaries of D correspond to the maximum and minimum value of y taken over this convex hull. However, the maximum and minimum of y is always achieved at the vertices of the convex hull. The optimal solution therefore is such that at most two of the phenotypes have nonzero mass.

While this argument does not generalize to arbitrary distributions over t, it helps rationalize why at most two phenotypes are largely sufficient when the distribution over t has one typical timescale (such as for an exponential distribution).

APPENDIX D: UNDER WHAT SCENARIOS IS A BIMODAL SOLUTION FAVORED?

1. Growth-lag paradigm

We aim to obtain a condition on the curvature of the growth-lag trade-off relation $\gamma(\ell)$ for when the one-state solution splits into a two-state solution. We first obtain the best one-state solution and then ask if there is a better two-state solution such that the two states have lags that are in an ε-neighborhood of the best single-state solution.

Consider the general two-state objective function

$$\begin{aligned}\Gamma(\phi,\ell_1,\ell_2)&=[\phi\gamma_{\rm P}(\ell_1)+(1-\phi)\gamma_{\rm P}(\ell_2)]\\&\quad+\omega\int_0^\infty dt\,\Pi(t)\ln(\phi e^{[\gamma'(t-\ell_1)]_+}+(1-\phi)e^{[\gamma'(t-\ell_2)]_+}).\end{aligned}\tag{D1}$$

We set $\omega=\gamma'=1$ noting that growth rates in P are measured in units of ω and time spent in NP and lag durations are measured in units of $1/\gamma'$. The optimal single-state solution (with $\phi=1$ in Eq. (D1)) denoted by Γ_1 is obtained by maximizing

$$\Gamma_1=\max_\ell\left\{\gamma_{\rm P}(\ell)+\int_0^\infty dt\,\Pi(t)(t-\ell)\right\},\tag{D2}$$

which leads to the optimality condition and single-state growth rate

$$\left.\frac{d\gamma_{\rm P}}{d\ell}\right|_{\ell_0}=F(\ell_0),\quad \Gamma_1=\gamma_{\rm P}(\ell_0)+I(\ell_0),\tag{D3}$$

where ℓ_0 is the lag that maximizes the right-hand side (r.h.s.) of Eq. (D2), and we denote $F(\ell)=\int_\ell^\infty dt\,\Pi(t)$ and $I(\ell)=\int_0^\infty dt\,\Pi(t)(t-\ell)=\int_\ell^\infty dt F(t)$ (the latter is obtained via integration by parts). We only consider cases where the best single-state lag ℓ_0 is finite; that is, $\Gamma_1>\gamma_{\rm P}(\infty)$.

We now calculate the best two-state solution such that $\ell_0-\varepsilon\leqslant\ell_1,\ell_2\leqslant\ell_0+\varepsilon$ where ε is a small positive constant. To do this, we expand each term in Eq. (D1) to second order. Define $\varepsilon_1=\ell_1-\ell_0$, $\varepsilon_2=\ell_2-\ell_0$. All γ terms and its derivatives are evaluated at ℓ_0. The first term on the r.h.s. is

$$\begin{aligned}\phi\gamma_{\rm P}(\ell_1)+(1-\phi)\gamma_{\rm P}(\ell_2)&=\gamma+\frac{d\gamma_{\rm P}}{d\ell}[\varepsilon_1\phi+\varepsilon_2(1-\phi)]\\&\quad+\frac{1}{2}\frac{d^2\gamma_{\rm P}}{d\ell^2}\left[\varepsilon_1^2\phi+\varepsilon_2^2(1-\phi)\right].\end{aligned}\tag{D4}$$

We split the second term into three parts:

$$\begin{aligned}&\int_0^\infty dt\,\Pi(t)\ln[\phi e^{t-\ell_1}+(1-\phi)e^{t-\ell_2}]\\&=\int_{\ell_1}^{\ell_2}dt\,\Pi(t)\ln[\phi e^{t-\ell_1}+(1-\phi)]\\&\quad+\int_{\ell_2}^\infty dt\,\Pi(t)(t-\ell_2)\\&\quad+\int_{\ell_2}^\infty dt\,\Pi(t)\ln[\phi e^{\ell_2-\ell_1}+(1-\phi)].\end{aligned}\tag{D5}$$

The first term in Eq. (D5) to second order in ε is

$$\int_{\ell_1}^{\ell_2}dt\,\Pi(t)\ln[\phi e^{t-\ell_1}+(1-\phi)]\approx\frac{1}{2}\phi\Pi(\ell_0)(\varepsilon_2-\varepsilon_1)^2.\tag{D6}$$

The second term in Eq. (D5) to second order in ε is

$$\begin{aligned}\int_{\ell_2}^\infty dt\,\Pi(t)(t-\ell_2)&=\int_{\ell_2}^\infty dt F(t)\approx I(\ell_0)\\&\quad-\varepsilon_2F(\ell_0)+\frac{1}{2}\varepsilon_2^2\Pi(\ell_0).\end{aligned}\tag{D7}$$

The third term in Eq. (D5) to second order in ε is

$$\begin{aligned}&\int_{\ell_2}^\infty dt\,\Pi(t)\ln[\phi e^{\ell_2-\ell_1}+(1-\phi)]\\&\quad\approx(\varepsilon_2-\varepsilon_1)\phi F(\ell_0)-\varepsilon_2(\varepsilon_2-\varepsilon_1)\phi\Pi(\ell_0)\\&\quad\quad+\frac{1}{2}(\varepsilon_2-\varepsilon_1)^2\phi(1-\phi)F(\ell_0).\end{aligned}\tag{D8}$$

Collecting the zeroth-order terms, we get $\gamma_{\rm P}(\ell_0)+I(\ell_0)$, which is precisely Γ_1. The first-order terms cancel out exactly once we use the stationarity condition $d\gamma_{\rm P}/d\ell=F(\ell_0)$ at ℓ_0. After adding up the contributions from the second-order terms, the two-state objective Γ_2 to second order in ε is

$$\begin{aligned}\Gamma_2=\Gamma_1+\max_{\phi,\varepsilon_1,\varepsilon_2}\left\{\frac{\Gamma_1''}{2}\left[\varepsilon_1^2\phi+\varepsilon_2^2(1-\phi)\right]\right.\\\left.+\frac{\phi(1-\phi)F}{2}(\varepsilon_2-\varepsilon_1)^2\right\},\end{aligned}\tag{D9}$$

where $\Gamma_1''=\gamma''+\Pi(\ell_0)$ is the second derivative of the single-state objective at $\ell=\ell_0$, $F=F(\ell_0)$ and the optimization has constraints $0<\phi<1$ and $|\varepsilon_1|,|\varepsilon_2|\leqslant\varepsilon$. Moreover, $\varepsilon_1\neq\varepsilon_2$ (since otherwise that would be a single-state solution).

Taking the maximum over ϕ in Eq. (D9), we get for the optimum ϕ^*,

$$\phi^*=\frac{1}{2}\left(1-\frac{\Gamma_1''}{F}\frac{\varepsilon_2+\varepsilon_1}{\varepsilon_2-\varepsilon_1}\right).\tag{D10}$$

Note that $\Gamma_1''<0$ since ℓ_0 is an interior point that attains the maximum of the single-state objective. Denote $A=-\Gamma_1''/F>0$ (since $F>0$). Plugging ϕ^* into Eq. (D9) and simplifying, we get

$$\begin{aligned}\Gamma_2=\Gamma_1+\frac{F}{8}\max_{\varepsilon_1,\varepsilon_2}\left\{(\varepsilon_2-\varepsilon_1)^2-2A\left(\varepsilon_2^2+\varepsilon_1^2\right)\right.\\\left.+A^2(\varepsilon_2+\varepsilon_1)^2\right\},\end{aligned}\tag{D11}$$

with the constraints $\varepsilon_1\neq\varepsilon_2$ and $|\varepsilon_1|,|\varepsilon_2|\leqslant\varepsilon$. There are three possibilities for the optimal $\varepsilon_1,\varepsilon_2$: (1) they are both interior points, i.e., they belong to the open interval $(-\varepsilon,\varepsilon)$, (2) one of them takes a boundary value of $\pm\varepsilon$ and the other is an interior point, or (3) both of them are at the boundary, in which case, without loss of generality, $\varepsilon_2=-\varepsilon_1=\varepsilon$ (since they cannot be equal). We now consider each case separately.

If they are both interior points, we can take the derivatives of Eq. (D11) with respect to ε_1 and ε_2 and equate them to zero. We get

$$2(\varepsilon_2^*-\varepsilon_1^*)+2A^2(\varepsilon_1^*+\varepsilon_2^*)-4A\varepsilon_2^*=0,\tag{D12}$$

$$-2(\varepsilon_2^*-\varepsilon_1^*)+2A^2(\varepsilon_1^*+\varepsilon_2^*)-4A\varepsilon_1^*=0.\tag{D13}$$

Solving for the two, we get $\varepsilon_1^* = \varepsilon_2^*$, which violates the $\varepsilon_1 \neq \varepsilon_2$ constraint.

The second case is that one of them is on the boundary, say ε_2^*. Then, taking the derivative of Eq. (D11) with respect to ε_1 and setting it to zero, we get

$$\varepsilon_1^* = \frac{1-A}{1+A}\varepsilon_2^*. \tag{D14}$$

The term being maximized in (34) is then

$$-\frac{4A(1-A)^2}{(1+A)^2}(\varepsilon_2^*)^2 < 0, \tag{D15}$$

since $A > 0$. The single-state solution outcompetes the best two-state solution in this scenario.

This leaves the third case that $\varepsilon_2 = -\varepsilon_1 = \varepsilon$, which gives for the term being maximized in Eq. (D11) as $4(1-A)\varepsilon^2$. Note that $\phi^* = 1/2$ in this case, which satisfies the constraint $0 \leqslant \phi \leqslant 1$. When $A < 1$, the two-state solution outcompetes the single-state solution. Thus,

$$\Gamma_2 = \Gamma_1 + \frac{F}{2}\max\{0, (1-A)\varepsilon^2\}. \tag{D16}$$

The transition from a single-state to two-state solution occurs at $A = 1$, or in terms of the trade-off and switching statistics,

$$\left.\frac{d^2\gamma_P}{d\ell^2}\right|_{\ell_0} + \Pi(\ell_0) + F(\ell_0) = 0. \tag{D17}$$

Note that this is a sufficient condition, but not a necessary one.

a. Test for critical curvature. For a chosen growth-lag trade-off with form $\gamma_P(\ell) = \gamma_{\max} - (\gamma_{\max} - \gamma_{\min})(\frac{\ell_{\min}}{\ell})^\alpha$, we identify the single-phenotype optimum (ℓ_0, γ_0). At this point, we construct small parabolas with varying curvatures to test when the solution transitions from a single-state solution to a two-state solution. After performing a rotation, the parabola with vertex (ℓ_0, γ_0) is given by

$$\begin{aligned}&(-\sin\theta)(\ell-\ell_0) + \cos\theta(\gamma-\gamma_0)\\ &\quad = \chi[\cos\theta(\ell-\ell_0) + \sin\theta(\gamma-\gamma_0)]^2.\end{aligned} \tag{D18}$$

Taylor expanding γ around γ_0 and substituting in the equation above, we get

$$\begin{aligned}&(-\sin\theta)(\ell-\ell_0) + \cos\theta\left(\gamma'(\ell-\ell_0) + \frac{\gamma''}{2}(\ell-\ell_0)^2\right)\\ &\quad = \chi(\ell-\ell_0)^2[\cos\theta + \gamma'\sin\theta]^2.\end{aligned} \tag{D19}$$

Since $\gamma' = \sin\theta/\cos\theta$, this equation reduces to

$$\gamma'' = \frac{2\chi}{\cos^3\theta}. \tag{D20}$$

In the previous section, we derived the condition for the transition from a single-state solution to a two-state solution, $-[\gamma'' + \Pi(\ell_0)]/F = 1$. Assuming $\Pi(t) = e^{-t}$, the condition becomes

$$\gamma'' = 2e^{-\ell_0}. \tag{D21}$$

Finally, by equating Eqs. (D20) and (D21) and noting that $e^{-\ell_0} = \tan\theta$, we get a relation for the critical curvature parametrized by χ_{crit} for the transition,

$$\chi_{\text{crit}} = \sin\theta\cos^2\theta. \tag{D22}$$

We confirmed this result numerically, which is described in Figs. 8(a) and 8(b).

b. Power-law growth-lag trade-off and exponential dwell times in NP. When $\gamma(\ell) = -\ell^{-\alpha}$ and $\Pi(\ell) = \frac{e^{-\ell/\eta}}{\eta}$ (recall $\eta = \gamma'/\omega'$), ℓ_0 is obtained by solving

$$\frac{\alpha}{\ell_0^{1+\alpha}} = e^{-\ell_0/\eta}. \tag{D23}$$

Using Eq. (D17), we have

$$-\frac{\alpha(\alpha+1)}{\ell_0^{\alpha+2}} + \left(1+\frac{1}{\eta}\right)e^{-\ell_0/\eta} > 0, \tag{D24}$$

which gives

$$\alpha < \left(1+\frac{1}{\eta}\right)\ell_0 - 1. \tag{D25}$$

Plugging back $\eta = \gamma'/\omega'$ and reintroducing units, $\ell_0 \to \gamma'\ell_0$, we have

$$\alpha < (\gamma' + \omega')\ell_0 - 1. \tag{D26}$$

2. Growth-death paradigm

We now consider a similar analysis for the growth-death paradigm. We analyze two cases that are distinguished by the sign of the curvature of the growth-death trade-off. In the first case, we show that the optimal phenotype distribution always consists of at most two specialist phenotypes, regardless of the distribution of dwell times in NP. In the second case, analogous to our analysis in the growth-lag scenario, we obtain a sufficient condition on the curvature of the growth-death trade-off that favors a heterogeneous, specialist strategy compared with a generalist strategy.

Let R denote the set of physiologically plausible γ, ξ values (recall that the $\gamma \geqslant 0, \xi \leqslant 0$). We define the boundary ∂R of this region as those phenotypes for which there is no other phenotype in R where both γ and ξ are larger.

Case 1. We first consider R such that there are no physiologically plausible phenotypes on the line joining any two phenotypes on ∂R. In this case, we can show that the optimal solution is at most bimodal, regardless of the distribution over intervals in the NP environment.

Consider K phenotypes on ∂R such that $\gamma_1 > \gamma_2 > \cdots > \gamma_K$ and $\xi_1 < \xi_2 < \cdots < \xi_K$. We have

$$\Gamma(\boldsymbol{\phi}) = \sum_k \gamma_k\phi_k + \int_0^\infty dt\ \Pi(t)\ln\left(\sum_k \phi_k e^{\xi_k t}\right), \tag{D27}$$

where $\Pi(t)$ is the distribution over intervals in NP. Say $\sum_k \gamma_k\phi_k = x$. We show that for each x that is feasible, the optimal solution is a mixture of the first and Kth phenotypes.

The feasible points given $\sum_k \gamma_k\phi_k = x$ is the convex hull formed by points on the edges of the probability simplex. Let the vertices of the convex hull be $\boldsymbol{q}_1, \boldsymbol{q}_2, \ldots, \boldsymbol{q}_L$. Any feasible point ϕ (given x) can be written as $\sum_\mu u_\mu \boldsymbol{q}_\mu$ where $u_\mu \geqslant 0$, $\sum_\mu u_\mu = 1$. Consider the vector $\boldsymbol{w}(t)$ whose kth component is $w_k(t) = e^{\xi_k t}$. The last term in Eq. (D27) can then be written as $\int_0^\infty dt\ \Pi(t)\ln[\sum_\mu u_\mu \boldsymbol{q}_\mu \cdot \boldsymbol{w}(t)]$.

Whenever the hyperplane $\sum_k \gamma_k\phi_k = x$ intersects the probability simplex ($\gamma_1 \geqslant x \geqslant \gamma_K$), it always intersects the edge

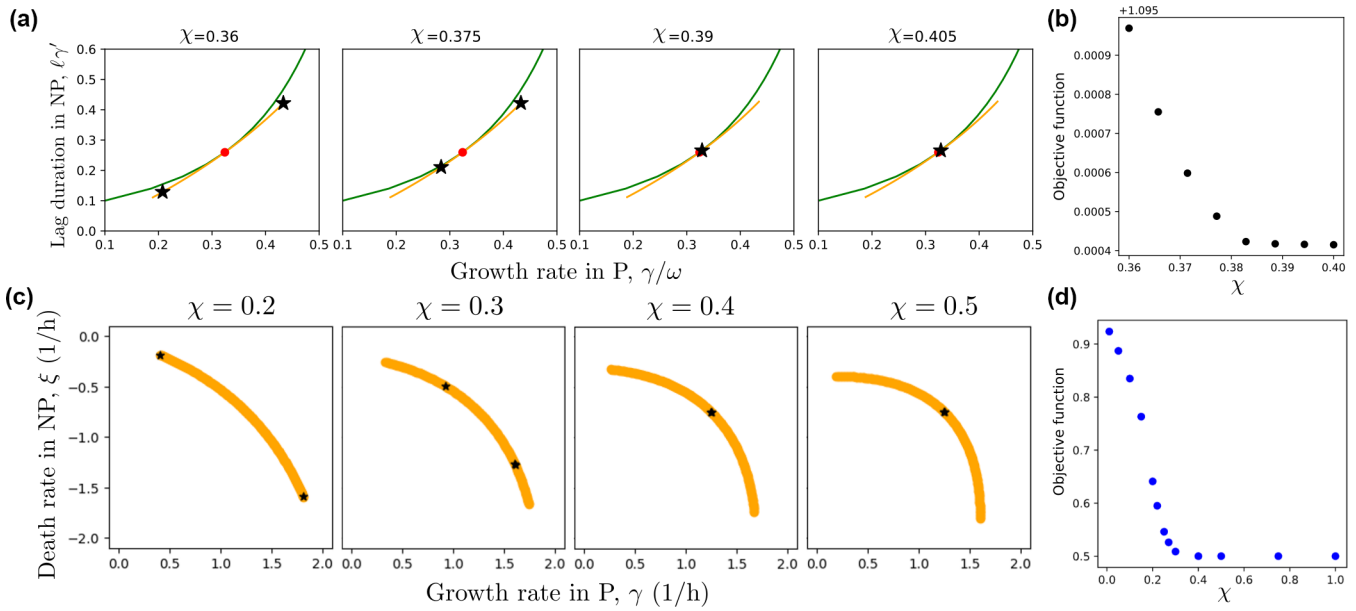


FIG. 8. Transition from a two-state solution to a single-state solution in the (a), (b) growth-lag system and (c), (d) growth-death system. (a) The green curve represents the original power-law growth-lag trade-off for which the single-phenotype optimum (γ_0, ℓ_0) is represented by the red circle. Taking this point as the vertex, we construct small parabolas with varying curvatures parametrized by χ and numerically determine the optimal solution. For a specific value of $\chi = \chi_{\text{crit}}$, there is a transition from a two-state solution to a single-state solution. For the example shown above, this transition occurs close to the analytic prediction of $\chi_{\text{crit}} = 0.3822$. (b) Maximum long-term population growth rate estimate for parabolic trade-offs with different curvatures, χ in the growth-lag system. (c) Analogous to panel (a): Optimal solutions for different small parabola growth-death trade-offs with varying χ are tested. The analytic prediction for χ_{crit} is 0.3535 for this example. (d) Maximum long-term population growth rate estimate for parabolic trade-offs with different curvatures, χ in the growth-death system.

joining the vertex $\phi_1 = 1$ and the vertex $\phi_k = 1$. Say this point of intersection is $\boldsymbol{q}_1$. We have

$$\boldsymbol{q}_1 \cdot \boldsymbol{w}(t) = \frac{x - \gamma_K}{\gamma_1 - \gamma_K} e^{\xi_1 t} + \frac{\gamma_1 - x}{\gamma_1 - \gamma_K} e^{\xi_K t}. \tag{D28}$$

We show that $\boldsymbol{q}_1 \cdot \boldsymbol{w}(t) \geqslant \boldsymbol{q}_\mu \cdot \boldsymbol{w}(t)$ for all μ and t. This implies that, for every feasible x,

$$\Gamma(\phi) \leqslant x + \int_0^\infty dt\, \Pi(t) \ln[\boldsymbol{q}_1 \cdot \boldsymbol{w}(t)], \tag{D29}$$

with equality when $u_1 = 1$, i.e., when all probability mass is on the first and Kth phenotypes.

To prove $\boldsymbol{q}_1 \cdot \boldsymbol{w}(t) \geqslant \boldsymbol{q}_\mu \cdot \boldsymbol{w}(t)$, first consider an intermediate phenotype i with growth rates γ_i, ξ_i. Let $\boldsymbol{q}_\mu$ be the intersection of the hyperplane with the line joining $\phi_1 = 1$ and $\phi_i = 1$ $(\gamma_1 \geqslant x \geqslant \gamma_i)$:

$$\boldsymbol{q}_\mu \cdot \boldsymbol{w}(t) = \frac{x - \gamma_i}{\gamma_1 - \gamma_i} e^{\xi_1 t} + \frac{\gamma_1 - x}{\gamma_1 - \gamma_i} e^{\xi_i t}. \tag{D30}$$

A short calculation gives

$$\boldsymbol{q}_1 \cdot \boldsymbol{w}(t) - \boldsymbol{q}_\mu \cdot \boldsymbol{w}(t) = \frac{\gamma_1 - x}{\gamma_1 - \gamma_i} \left[\left(\frac{\gamma_i - \gamma_K}{\gamma_1 - \gamma_K} \right) e^{\xi_1 t} + \left(\frac{\gamma_1 - \gamma_i}{\gamma_1 - \gamma_K} \right) e^{\xi_K t} - e^{\xi_i t} \right]. \tag{D31}$$

Since e^t is convex, one obtains

$$\boldsymbol{q}_1 \cdot \boldsymbol{w}(t) - \boldsymbol{q}_\mu \cdot \boldsymbol{w}(t) \geqslant \frac{\gamma_1 - x}{\gamma_1 - \gamma_i} (e^{\beta t} - e^{\xi_i t}), \tag{D32}$$

where

$$\beta = \frac{\gamma_i - \gamma_K}{\gamma_1 - \gamma_K} \xi_1 + \frac{\gamma_1 - \gamma_i}{\gamma_1 - \gamma_K} \xi_K. \tag{D33}$$

Thus $\boldsymbol{q}_1 \cdot \boldsymbol{w}(t) \geqslant \boldsymbol{q}_\mu \cdot \boldsymbol{w}(t)$ if

$$\frac{\xi_1 - \xi_K}{\gamma_1 - \gamma_K} \geqslant \frac{\xi_i - \xi_K}{\gamma_i - \gamma_K}. \tag{D34}$$

By the assumed shape of ∂R, this holds for all i. A similar argument applies when the hyperplane intersects the edge joining $\phi_k = 1$ and $\phi_i = 1$ $(\gamma_K \leqslant x \leqslant \gamma_i)$. Hence $\boldsymbol{q}_1 \cdot \boldsymbol{w}(t) \geqslant \boldsymbol{q}_\mu \cdot \boldsymbol{w}(t)$ for every μ.

Case 2. Now suppose ∂R does not satisfy the conditions of Case 1, so solutions with more than two modes become possible. Nonetheless, we can state when more than one phenotype is favored. The best single-phenotype strategy corresponds to the point where a line of slope -1 is tangent to ∂R in the (γ, ξ) plane. If the curvature of ∂R near this point is small enough, a heterogeneous strategy outcompetes the single-phenotype strategy.

To see this, consider an infinitesimal parabola $w = \chi z^2$ around the best phenotype (γ, ξ). The z axis is along the tangent; the w axis points away and downward. We look at solutions on this parabola with $|z| \leqslant \varepsilon \ll 1$ for various χ.

As in Case 1,

$$\Gamma(\boldsymbol{\phi}) = \sum_k \gamma_k \phi_k + \int_0^\infty dt\, \Pi(t) \ln \left(\sum_k \phi_k e^{\xi_k t} \right). \tag{D35}$$

Rescale t and ξ_k so that $\langle t\rangle = 1$. Since solutions are at most bimodal, let one phenotype have probability ϕ and lie at z_1, the other at z_2. Writing $\gamma_i = \gamma + \Delta\gamma_i$ and $\xi_i = \xi + \Delta\xi_i$ for $i = 1, 2$, a second-order Taylor expansion gives

$$\Gamma(\phi) - \gamma - \xi \approx \Gamma_\varepsilon(\phi) = \phi\mu_1 + (1-\phi)\mu_2 - \frac{\tau_2^2}{2}[\phi\nu_1 + (1-\phi)\nu_2]^2, \tag{D36}$$

with $\tau_2^2 = \langle t^2\rangle$, $\mu_i = \Delta\gamma_i + \Delta\xi_i + \tau_2^2(\Delta\xi_i)^2/2$, and $\nu_i = \Delta\xi_i$. Optimizing over ϕ yields

$$\phi^* = \frac{\mu_1 - \mu_2}{\tau_2^2(\nu_1 - \nu_2)^2} - \frac{\nu_2}{\nu_1 - \nu_2}, \tag{D37}$$

for $0 \leqslant \phi^* \leqslant 1$. Substituting back,

$$\Gamma_\varepsilon(\phi^*) = \frac{(\mu_1 - \mu_2)^2}{2\tau_2^2(\nu_1 - \nu_2)^2} + \frac{\mu_2\nu_1 - \mu_1\nu_2}{\nu_1 - \nu_2}. \tag{D38}$$

With $\Delta\gamma_i = (z_i - \chi z_i^2)/\sqrt{2}$ and $\Delta\xi_i = -(z_i + \chi z_i^2)/\sqrt{2}$ and retaining terms up to second order,

$$\Gamma_\varepsilon(\phi^*) \approx \frac{\kappa^2(z_1 + z_2)^2}{\tau_2^2} - \kappa z_1 z_2, \tag{D39}$$

where $\kappa = \tau_2^2/4 - \sqrt{2}\,\chi$.

For $\kappa < 0$, $a > 1$ and the optimum reduces to the single best phenotype; in this case,

$$\Gamma_\varepsilon(\phi^*) \approx \frac{\kappa^2 z_1^2}{\tau_2^2}\left(\frac{1 - a^2}{a^2}\right) < 0, \tag{D40}$$

which is suboptimal.

If $\kappa > 0$, a bimodal solution with $z_1 z_2 < 0$ beats the single-phenotype solution (provided $0 \leqslant \phi^* \leqslant 1$). One shows the optimum occurs at $z_1 = -z_2 = \varepsilon$:

$$\Gamma_\varepsilon(\phi^*) \approx \frac{\kappa^2 z_1^2}{\tau_2^2}[(a + x)^2 + 1 - a^2], \tag{D41}$$

where $x = z_2/z_1$ and $a = 1 - \tau_2^2/2\kappa$. Since $a < -1$ when $\kappa > 0$, $x = -1$ maximizes Eq. (D41). Then $\phi^* = 1/2$ and $\Gamma_\varepsilon(p^*) \approx \kappa z_1^2 > 0$, maximized at $z_1 = \varepsilon$.

Hence, a heterogeneous solution is favored if $\kappa > 0$, i.e.,

$$\chi < \frac{\langle t^2\rangle}{4\sqrt{2}}, \tag{D42}$$

where $\langle t^2\rangle$ is the second moment of the NP dwell-time distribution. If the curvature χ of ∂R is below this threshold, a mixed two-phenotype strategy (specialists) outperforms the single-phenotype strategy (generalist). This is confirmed numerically, where we plot the maximum long-term growth rate for different χ values as shown in Fig. 8(d). For χ above the threshold, the optimal growth rate stays constant as the single phenotype solution does not change with χ. For χ below the threshold, we see that higher values for the long-term growth are achievable with two-phenotype strategies.

We chose parabolic curves centered at a specific point with different curvatures χ for the growth-death trade-off. In Fig. 8(c), as χ increases, the optimal solution transitions from a bimodal phenotype distribution to a single-phenotype distribution.

APPENDIX E: TWO-PHENOTYPE SYSTEM IN THE GROWTH-LAG PARADIGM

1. Why do arresters have infinite lag?

Assuming that the population can have two phenotypes with lags ℓ_1 and ℓ_2 such that $\ell_{min} \leqslant \ell_1 < \ell_2$ and the fraction of the population with lag ℓ_1 in P is ϕ. We consider growth-lag trade-offs such that the growth rate in P of a cell with lag ℓ in NP is

$$\gamma_{\rm P}(\ell) = \gamma_{\max} - (\gamma_{\max} - \gamma_{\min})\left(\frac{\ell_{\min}}{\ell}\right)^\alpha. \tag{E1}$$

Consider the term to be maximized over ϕ in Eq. (B15), which we denote as $\Gamma(\phi, \ell_1, \ell_2)$. We consider $\Pi(t) = \omega' e^{-\omega' t}$. Interchanging the integrals over τ and t, we get

$$\Gamma(\phi, \ell_1, \ell_2) = \phi\gamma_{\rm P}(\ell_1) + (1-\phi)\gamma_{\rm P}(\ell_2) + \omega\int_0^\infty e^{-\omega' t}\overline{\gamma}_{\rm NP}(t)dt, \tag{E2}$$

where $\overline{\gamma}_{\rm NP}(t)$ is the average growth rate at time t after the environment is switched from P to NP given that the population began with fraction ϕ cells in phenotype 1. Specifically,

$$\bar{\gamma}_{\rm NP}(t) = \gamma'\begin{cases} 0 & t < \ell_1 \\ \frac{\phi e^{\gamma'(t-\ell_1)}}{\phi e^{\gamma'(t-\ell_1)} + (1-\phi)} & \ell_1 < t \leqslant \ell_2 \\ 1 & t \geqslant \ell_2. \end{cases} \tag{E3}$$

Rescaling time: $t \to \gamma' t$, and defining $\eta = \gamma'/\omega'$, we have

$$\Gamma(\phi, \ell_1, \ell_2) = \max_\phi \left\{\phi\gamma_{\rm P}(\ell_1) + (1-\phi)\gamma_{\rm P}(\ell_2) + \omega\left(\int_{\ell_1}^{\ell_2} dt e^{-\frac{t}{\eta}}\frac{\phi e^{t-\ell_1}}{\phi e^{t-\ell_1} + (1-\phi)} + e^{-\frac{\ell_2}{\eta}}\right)\right\}. \tag{E4}$$

Shifting time as $t \to t - \ell_1$ and after a few simplifications, Γ reduces to the objective function

$$f(\phi, \ell_1, \ell_2) = -\frac{\phi}{\ell_1^\alpha} - \frac{1-\phi}{\ell_2^\alpha} + ke^{-\frac{\ell_1}{\eta}}\left[\int_0^x dt e^{-\frac{t}{\eta}}\frac{\phi e^t}{\phi e^t + (1-\phi)} + \int_{\ell_2-\ell_1}^\infty dt e^{-\frac{t}{\eta}}\frac{1-\phi}{\phi e^t + (1-\phi)}\right], \tag{E5}$$

where k is a positive constant obtained by absorbing ω with $(\gamma_{\max} - \gamma_{\min})\ell_{\min}^\alpha$, and $\ell_2 = \ell_1 + x$. Substituting $\phi = 1/(1 + e^{-\beta})$, we derive an equivalent form for the objective function,

$$f(\beta, \ell_1, x) = -\frac{1}{(1 + e^{-\beta})\ell_1^\alpha} - \frac{e^{-\beta}}{(1 + e^{-\beta})(\ell_1 + x)^\alpha} + ke^{-\frac{\ell_1}{\eta}}\left[\int_0^x dt e^{-\frac{t}{\eta}}\frac{e^t}{e^t + e^{-\beta}} + \eta e^{-\frac{x}{\eta}}\right]. \tag{E6}$$

Taking the derivative of f with respect to x yields

$$\frac{\partial f}{\partial x} = \frac{\alpha e^{-\beta}}{(1 + e^{-\beta})(\ell_1 + x)^{\alpha+1}} - ke^{-\frac{\ell_1 + x}{\eta}}\left(\frac{e^{-\beta}}{e^x + e^{-\beta}}\right). \tag{E7}$$

The first term on the right-hand side is always positive and decays with x as a power-law, whereas the second term is always negative and decays exponentially. For sufficiently large x, the first term dominates and the gradient is positive. That is, it is preferable to indefinitely increase ℓ_2. While this argument does not prove that the arrester solution is the preferred phenotype under all scenarios, it helps rationalize why an arrester phenotype with an infinite lag is selected.

2. Relating the lag duration of recoverers and the fraction of recoverers

We now consider a situation where the population can choose the fraction of recoverers ϕ and their lag duration $\ell \geqslant \ell_{\min}$ [and consequently their growth rate $\gamma(\ell)$]. This scenario corresponds to $\ell_1 = \ell$ and $\ell_2 = \infty$ from Appendix E 1. The average growth in NP $\overline{\gamma}_{\rm NP}(t)$ satisfies $\overline{\gamma}_{\rm NP}(t) = 0$ for $t \leqslant \ell$. For $t > \ell$, we have

$$\overline{\gamma}_{\rm NP}(t) = \gamma' \frac{\phi e^{\gamma'(t-\ell)}}{1-\phi+\phi e^{\gamma'(t-\ell)}}. \tag{E8}$$

Substituting $t \to t - \ell$ and simplifying, we obtain

$$V(\phi, \ell) = \gamma' e^{-\omega'\ell} \int_0^\infty dt\, \frac{(1-\phi)e^{-\omega' t}}{(1-\phi)+\phi e^{-\gamma' t}}. \tag{E9}$$

The integral on the right-hand side can be expressed in terms of the hypergeometric function ${}_2F_1(a, b; c; x)$. Denoting $\eta \equiv \gamma'/\omega'$, we have

$$V(\phi, \ell) = \eta\, e^{-\omega'\ell}\, {}_2F_1\left(1, \frac{1}{\eta}; \frac{1}{\eta}+1; -\frac{1-\phi}{\phi}\right). \tag{E10}$$

When $0 < \phi < 1$ and $\ell > \ell_{\min}$, we can optimize for ϕ and ℓ by taking the derivative of the growth function and setting it to zero. For a population in state s in P, we obtain the optimality relation

$$\frac{\partial V}{\partial \phi} \Big/ \frac{\partial \bar{\gamma}_{\rm P}}{\partial \phi} = \frac{\partial V}{\partial \ell} \Big/ \frac{\partial \bar{\gamma}_{\rm P}}{\partial \ell} = -\frac{1}{\omega(s)}. \tag{E11}$$

The first equation above expresses a relationship between the optimal ℓ^* and ϕ^* independent of $\omega(s)$. We then obtain

$$\ell^* = \frac{\alpha\, \eta\, {}_2F_1\left(2, \frac{1}{\eta}+1; \frac{1}{\eta}+2; -\frac{1-\phi^*}{\phi^*}\right)}{\gamma' \phi^* (\eta+1)\, {}_2F_1\left(1, \frac{1}{\eta}; \frac{1}{\eta}+1; -\frac{1-\phi^*}{\phi^*}\right)}. \tag{E12}$$

3. Inferring the environmental parameters from experimental data

Treating η and α/γ' as fitting parameters, we obtained best-fit values of $\eta = 1.2$ and $\alpha/\gamma' = 21$ h [see Fig. 5(c)]. Using the experimental estimate $\gamma' = 0.1\ \text{h}^{-1}$ from Ref. [17], these values correspond to $\alpha = 2.1$ and $\omega' = 0.08\ \text{h}^{-1}$. To proceed with constructing the phase portrait for the optimal solution, we use the fact that the maximum growth rate $\gamma_{\rm P}$ on glucose is $0.5\ \text{h}^{-1}$. The phase portrait can be seen in Fig. 9. The red region indicates where the yeast system exhibits a bimodal strategy of specialists. This corresponds to $\omega \in [0.07, 0.59]\ \text{h}^{-1}$.

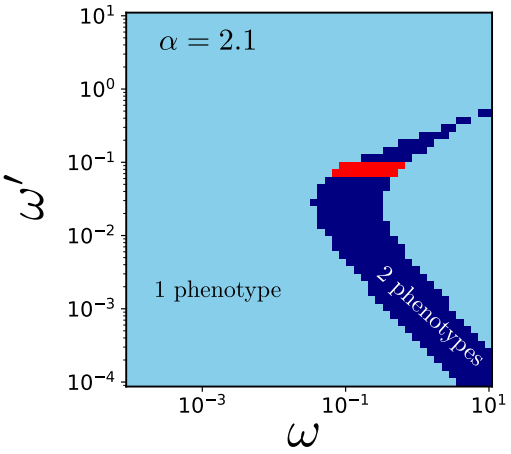


FIG. 9. Phase diagram of the optimal number of phenotypes over the space of environmental switching rates, ω and ω', for a trade-off with $\alpha = 2.1$, which is the value inferred from experimental data. The red region corresponds to where the yeast system exhibits a bimodal specialist strategy.

APPENDIX F: SWITCHING BETWEEN TWO PHENOTYPES IN THE PREFERRED ENVIRONMENT

While setting up the model, we made an assumption that cells rapidly switch between different phenotypes to maintain a constant phenotype distribution in the preferred environment (P). This switching occurs on timescales shorter than any doubling time and typical time spent in P, but slower than the time it takes to switch between P and NP. This ensures that cells cannot adapt while the environment is switching. In this section, we illustrate the balance between growth and switching that ensures a constant phenotype distribution in P.

Suppose the population is composed of two phenotypes with growth rates γ_1 and γ_2 ($\gamma_2 > \gamma_1$). $N_{\rm tot}$ is the total number of initial cells, of which N_1 cells adopt phenotype 1 and $N_2 (= N_{\rm tot} - N_1)$ cells adopt phenotype 2. The time evolution of $N_{\rm tot}$ and N_1 are governed by

$$\frac{dN_{\rm tot}}{dt} = \gamma_1 N_1 + \gamma_2 (N_{\rm tot} - N_1), \tag{F1}$$

$$\frac{dN_1}{dt} = \gamma_1 N_1 + \psi_{21}(N_{\rm tot} - N_1) - \psi_{12} N_1, \tag{F2}$$

where ψ_{21} and ψ_{12} are the switching rates from phenotypes 2 to 1 and 1 to 2, respectively. The quantity that remains constant in the preferred environment is $\phi = N_1/N_{\rm tot}$ such that $d\phi/dt = 0$. This gives us

$$\frac{d\phi}{dt} = \frac{d(N_1/N_{\rm tot})}{dt} = \frac{N_{\rm tot}(dN_1/dt) - N_1(dN_{\rm tot}/dt)}{N_{\rm tot}^2} = 0. \tag{F3}$$

Substituting Eqs. (F1) and (F2) and simplifying, we get

$$\phi^2 - (1 + \hat{\psi}_{21} + \hat{\psi}_{12})\phi + \hat{\psi}_{21} = 0, \tag{F4}$$

where we have defined $\hat{\psi}_{21} = \psi_{21}/(\gamma_2 - \gamma_1)$, $\hat{\psi}_{12} = \psi_{12}/(\gamma_2 - \gamma_1)$. The solution for ϕ is given by

$$\phi = \frac{(1 + \hat{\psi}_{21} + \hat{\psi}_{12}) \pm \sqrt{(1 + \hat{\psi}_{21} + \hat{\psi}_{12})^2 - 4\hat{\psi}_{21}}}{2}. \tag{F5}$$

Tuning $\hat{\psi}_{21}$ and $\hat{\psi}_{12}$ allows for tuning ϕ to any value between 0 and 1. For example, when the rate of switching is much larger than the difference in growth rates, i.e., $\hat{\psi}_{21}, \hat{\psi}_{12} \gg 1$, we can approximate the steady-state ϕ using Eq. (F5) as

$$\phi \approx \frac{\hat{\psi}_{21}}{\hat{\psi}_{21} + \hat{\psi}_{12}}. \tag{F6}$$

The sum $\hat{\psi}_{21} + \hat{\psi}_{12}$ sets the rate at which the population reaches steady-state, whereas the ratio $\hat{\psi}_{21}/\hat{\psi}_{12}$ sets the value of ϕ at this steady state.

APPENDIX G: A MINIMAL METABOLIC MODEL ILLUSTRATING CONSTRAINTS ON SPEED OF ADAPTATION

Microorganisms encounter environments in which the available carbon source fluctuates between preferred and nonpreferred nutrients. Certain nutrients (i.e., glucose) support rapid growth by activating the corresponding metabolic pathways (i.e., glycolysis), whereas other nutrients (i.e., glycerol, acetate) require induction of alternative pathways (i.e., gluconeogenesis). This results in a characteristic lag phase prior to growth resumption. Fast growth on the preferred nutrient and rapid adaptation (short lag) on the nonpreferred nutrient are inherently antagonistic, reflecting a growth-lag trade-off. We present a minimal model, inspired by central metabolism, which captures a growth-lag trade-off that could arise due to antagonism in core pathways.

1. Model structure and assumptions

The schematic of the model is shown in Fig. 4(a), where

(1) G and N denote the concentrations of the metabolite pools in the preferred environment (P) and nonpreferred environment (NP), respectively;

(2) g and n denote intracellular concentrations of the central metabolites;

(3) r is the maximal forward flux rate of the irreversible reactions in both P and NP;

(4) δ is the loss rate constant for g and n;

(5) φ and ζ are the abundances of the irreversible enzymes that catalyze the forward ($g \to n$) and reverse ($n \to g$) reactions, respectively.

We assume Michaelis–Menten kinetics for each enzymatic step. Growth rate linked to biomass production, μ, is assumed to depend on the levels of the metabolites g and n:

$$\mu = \left(\frac{n}{n + K_{\mathrm{BM}}}\right)^{\nu} \left(\frac{g}{g + K_{\mathrm{BM}}}\right)^{\nu}, \tag{G1}$$

where K_{BM} is the half-saturation constant, and the exponent $\nu > 0$ captures effective cooperative mechanisms in metabolism.

The time evolution of each metabolite pool is given by mass-action and Michaelis–Menten-type terms. In the preferred environment, when G is present and N is absent, the equations are

$$\frac{dg}{dt} = rG \; - \; \varphi \, \frac{g}{K_g + g} \; + \; \zeta \, \frac{n}{K_n + n} - \nu \, \mu - \delta \, g,$$

$$\frac{dn}{dt} = \varphi \, \frac{g}{K_g + g} \; - \; \zeta \, \frac{n}{K_n + n} - \nu \, \mu - \delta \, n.$$

In the nonpreferred environment, when G is absent and N is present, the equations are

$$\frac{dg}{dt} = - \, \varphi \, \frac{g}{K_g + g} \; + \; \zeta \, \frac{n}{K_n + n} - \nu \, \mu \; - \; \delta \, g,$$

$$\frac{dn}{dt} = rN \; + \; \varphi \, \frac{g}{K_g + g} \; - \; \zeta \, \frac{n}{K_n + n} - \nu \, \mu - \delta \, n.$$

We fix the concentrations of G and N to be 2 arb. units and 1 arb. unit in the preferred and nonpreferred environments, respectively.

2. Enzyme dynamics and regulation

We model the synthesis and dilution of the irreversible enzymes, φ and ζ, by simple first-order kinetics toward nutrient-specific steady-state levels:

$$\frac{d\varphi}{dt} = \mu\left(\varphi_{\mathrm{ss}}^{(\mathrm{nutrient})} - \varphi\right), \tag{G2}$$

$$\frac{d\zeta}{dt} = \mu\left(\zeta_{\mathrm{ss}}^{(\mathrm{nutrient})} - \zeta\right), \tag{G3}$$

where

(1) μ is the instantaneous growth rate;

(2) $\varphi_{\mathrm{ss}}^{(\mathrm{nutrient})}$ and $\zeta_{\mathrm{ss}}^{(\mathrm{nutrient})}$ are the steady-state enzyme levels where the "nutrient" may be G or N;

(3) $\varphi_{\mathrm{ss}}^{G} > \zeta_{\mathrm{ss}}^{G}$ biases flux in the direction $g \to n$ when G is the nutrient, and vice versa after a nutrient shift to N being the nutrient.

3. Origin of the power-law growth-lag trade-off

Upon a sudden shift from P to NP, the pre-existing high level of φ suppresses flux from $n \to g$ and dissipates excess accumulation of n. This depletes cellular resources and delays the reversal of net flux. Because φ decays only on the timescale $\approx 1/\mu$, flux reversal stalls until φ falls and ζ rises toward ζ_{ss}^{N} so that the growth rate μ can rise. This leads to a lag before growth can resume.

The metabolic model operates in at least three qualitatively distinct regimes, delineated by $\varphi_{\mathrm{ss}}^{G}$, depending on how the system produces new enzymes ζ after the switch from P to NP. Here, we expand on the third scenario discussed in the main text, where $\varphi_{\mathrm{ss}}^{G}$ has intermediate values such that the value of ζ_{ss}^{G} that maximizes preshift growth rate is zero (Fig. 10). We show below that the growth-lag trade-off in this regime is a power-law with exponent $1/(\nu - 1)$.

Suppose $\varphi = \varphi_{\mathrm{ss}}^{G}$ and $\zeta = \zeta_{\mathrm{ss}}^{G}$ are the steady-state values of the enzymes in P. At steady state, g and n are found by solving

$$G - \varphi v_1(g) + \zeta v_1(n) - \delta g - \nu \mu(g, n) = 0,$$

$$\varphi v_1(g) - \zeta v_1(n) - \delta n - \nu \mu(g, n) = 0,$$

where $r = 1$, $v_1(x) = x/(x + K)$ (assume $K_g = K_n = K$), $\mu(g, n) = [v_2(g) v_2(n)]^{\nu}$ and $v_2(x) = x/(x + K_{\mathrm{BM}})$.

We consider a parameter regime where, given $\varphi_{\mathrm{ss}}^{G}$, μ_{P} is maximized when $\zeta_{\mathrm{ss}}^{G} = 0$ and consequently ℓ is infinite. We calculate how μ_{P} and ℓ decrease with ζ_{ss}^{G} when $0 < \zeta_{\mathrm{ss}}^{G} \ll 1$.

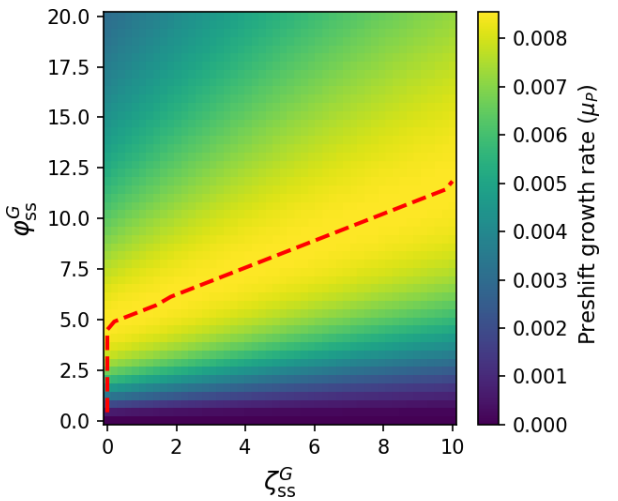


FIG. 10. Heatmap of preshift growth rate μ_P for different φ_{ss}^G and ζ_{ss}^G values. Other parameters in the model are fixed. Above a certain value of φ_{ss}^G, the optimal preshift growth rate is achieved at a nonzero ζ_{ss}^G (dashed red curve).

The maximal growth rate when $\zeta_{ss}^G = 0$ is found by solving the steady-state equations

$$0 = G - \varphi_{ss}^G v_1(g^*) - \delta g^* - \nu\mu(g^*, n^*),$$
$$0 = \varphi_{ss}^G v_1(g^*) - \delta n^* - \nu\mu(g^*, n^*),$$

where g^*, n^* are the values at steady state in this scenario. Let $\Delta g = g - g^*$ and $\Delta n = n - n^*$ be the deviations from these values when $0 < \zeta_{ss}^G \ll 1$. Δg and Δn are of order ζ_{ss}^G. Expanding the steady-state equations to first order when $0 < \zeta_{ss}^G \ll 1$, Δg and Δn satisfy

$$\zeta_{ss}^G v_1(n^*) = \left(\delta + \nu\frac{\partial\mu_P}{\partial g} + \varphi_{ss}^G v_1'(g^*)\right)\Delta g + \left(\nu\frac{\partial\mu_P}{\partial n}\right)\Delta n,$$
$$\zeta_{ss}^G v_1(n^*) = \left(\varphi_{ss}^G v_1'(g^*) - \nu\frac{\partial\mu_P}{\partial g}\right)\Delta g + \left(-\delta - \nu\frac{\partial\mu_P}{\partial n}\right)\Delta n,$$

where the derivatives are evaluated at g^*, n^*. Since $\Delta\mu_P = \frac{\partial\mu_P}{\partial g}\Delta g + \frac{\partial\mu_P}{\partial n}\Delta n$ (to first order in ζ_{ss}^G), we solve Δg and Δn by solving the above pair of equations. This gives an expression for $\Delta\mu_P$:

$$\Delta\mu_P = \zeta_{ss}^G \begin{pmatrix} \frac{\partial\mu_P}{\partial g} & \frac{\partial\mu_P}{\partial n} \end{pmatrix} A^{-1} \begin{pmatrix} v_1(n^*) \\ v_1(n^*) \end{pmatrix}, \text{ where}$$
$$A = \begin{pmatrix} \delta + \nu\frac{\partial\mu_P}{\partial g} + \varphi_{ss}^G v_1'(g^*) & \nu\frac{\partial\mu_P}{\partial n} \\ \varphi_{ss}^G v_1'(g^*) - \nu\frac{\partial\mu_P}{\partial g} & -\delta - \nu\frac{\partial\mu_P}{\partial n} \end{pmatrix}, \quad \text{(G4)}$$

which we can write as $\Delta\mu_P = \zeta_{ss}^G b$, where b is a scalar obtained from the vector-matrix-vector product on the right-hand side of the above equation.

After the switch to NP (starting from the steady state in P with $0 < \zeta_{ss}^G \ll 1$), we have

$$\frac{dg}{dt} = -\varphi v_1(g) + \zeta v_1(n) - \delta g - \nu\mu, \quad \text{(G5)}$$
$$\frac{dn}{dt} = N + \varphi v_1(g) - \zeta v_1(n) - \delta n - \nu\mu, \quad \text{(G6)}$$

where

$$\frac{d\varphi}{dt} = \mu\left(\varphi_{ss}^N - \varphi\right), \quad \text{(G7)}$$
$$\frac{d\zeta}{dt} = \mu\left(\zeta_{ss}^N - \zeta\right), \quad \text{(G8)}$$

with initial conditions $\varphi = \varphi_{ss}^G$ and $\zeta = \zeta_{ss}^G$. The approximations are valid soon after the switch. Dividing these two equations, we get $d\varphi/d\zeta = (\varphi_{ss}^N - \varphi)/(\zeta_{ss}^N - \zeta)$. Integrating, we get

$$\frac{\varphi_{ss}^N - \varphi}{\varphi_{ss}^N - \varphi_{ss}^G} = \frac{\zeta_{ss}^N - \zeta}{\zeta_{ss}^N - \zeta_{ss}^G}. \quad \text{(G9)}$$

As discussed earlier, we work in the limit where after the switch, g depletes rapidly due to both the transformation from g to n catalyzed by φ, and due to loss at rate δ. In quantitative terms, the first and third terms on the r.h.s. of (G5) contribute to the outflux of g.

During the lag phase, both φ and ζ do not change significantly from φ_{ss}^G and ζ_{ss}^G as biomass production (i.e., μ) is minimal in this phase. When φ_{ss}^G is large, g drops rapidly to near zero until it reaches a steady state. Since $v_1(g) \approx g$ when $g \ll K$ and $v_1(g) \approx 1$ when $g \gg K$, either g initially decreases exponentially or linearly with rate φ_{ss}^G depending on the value of g^*. Moreover, when $g \ll 1$, $\mu \ll v_1(g)$ as $\mu \sim g^\nu$ and $v_1(g) \sim g$ for small g and $\nu > 1$. That is, $dg/dt \approx -(\varphi + \delta)g + \zeta v_1(n)$. At steady state, we have $(\varphi + \delta)g \approx \zeta v_1(n)$.

Moreover, in steady state, from (G6) we see that all terms except N and δn are negligible so that $n \approx N/\delta$. Thus,

$$g \approx \left(\frac{\zeta}{\delta + \varphi_{ss}^G}\right) v_1(N/\delta), \quad \text{(G10)}$$

where we used (G9) to eliminate φ and ignored all terms of second order in ζ. Plugging this expression for g into μ in (G8), we have

$$\frac{d\zeta}{dt} \approx \zeta^\nu \left(\frac{v_1(N/\delta)v_2(N/\delta)}{\delta + \varphi_{ss}^G}\right)^\nu. \quad \text{(G11)}$$

Integrating this from $\zeta = \zeta_{ss}^G$ to ∞ (or some large value) for $\nu > 1$, we get an estimate of lag ℓ as the time it takes for ζ to diverge:

$$\ell \approx \frac{\left(\zeta_{ss}^G\right)^{1-\nu}}{\nu - 1}\left(\frac{\delta + \varphi_{ss}^G}{v_1(N/\delta)v_2(N/\delta)}\right)^\nu. \quad \text{(G12)}$$

Plugging in $\zeta_{ss}^G = \Delta\mu_P/b$, we have

$$\Delta\mu_P = A\ell^{-\frac{1}{\nu-1}}, \quad \text{(G13)}$$

for some constant A.

[1] M. Ackermann, A functional perspective on phenotypic heterogeneity in microorganisms, Nat. Rev. Microbiol. **13**, 497 (2015).

[2] E. Rotem, A. Loinger, I. Ronin, I. Levin-Reisman, C. Gabay, N. Shoresh, O. Biham, and N. Q. Balaban, Regulation of phenotypic variability by a threshold-based mechanism underlies bacterial persistence, Proc. Natl. Acad. Sci. USA **107**, 12541 (2010).

[3] B. Van den Bergh, J. E. Michiels, T. Wenseleers, E. M. Windels, P. V. Boer, D. Kestemont, L. De Meester, K. J. Verstrepen, N. Verstraeten, M. Fauvart, and J. Michiels, Frequency of antibiotic application drives rapid evolutionary adaptation of *Escherichia coli* persistence, Nat. Microbiol. **1**, 16020 (2016).

[4] E. Kussell, R. Kishony, N. Q. Balaban, and S. Leibler, Bacterial persistence: A model of survival in changing environments, Genetics **169**, 1807 (2005).

[5] D. M. Wolf, V. V. Vazirani, and A. P. Arkin, Diversity in times of adversity: Probabilistic strategies in microbial survival games, J. Theor. Biol. **234**, 227 (2005).

[6] M. Thattai and A. van Oudenaarden, Stochastic gene expression in fluctuating environments, Genetics **167**, 523 (2004).

[7] R. Levins, Theory of fitness in a heterogeneous environment. I. The fitness set and adaptive function, Am. Nat. **96**, 361 (1962).

[8] R. Levins, Theory of fitness in a heterogeneous environment. II. Developmental flexibility and niche selection, Am. Nat. **97**, 75 (1963).

[9] 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).

[10] O. Rivoire and S. Leibler, The value of information for populations in varying environments, J. Stat. Phys. **142**, 1124 (2011).

[11] B. Xue and S. Leibler, Benefits of phenotypic plasticity for population growth in varying environments, Proc. Natl. Acad. Sci. USA **115**, 12745 (2018).

[12] B. Xue, P. Sartori, and S. Leibler, Environment-to-phenotype mapping and adaptation strategies in varying environments, Proc. Natl. Acad. Sci. USA **116**, 13847 (2019).

[13] E. Kussell and S. Leibler, Phenotypic diversity, population growth, and information in fluctuating environments, Science **309**, 2075 (2005).

[14] M. D. Rolfe, C. J. Rice, S. Lucchini, C. Pin, A. Thompson, A. D. Cameron, M. Alston, M. F. Stringer, R. P. Betts, J. Baranyi, *et al.*, Lag phase is a distinct growth phase that prepares bacteria for exponential growth and involves transient metal accumulation, J. Bacteriol. **194**, 686 (2012).

[15] L. E. Bagamery, Q. A. Justman, E. C. Garner, and A. W. Murray, A putative bet-hedging strategy buffers budding yeast against environmental instability, Curr. Biol. **30**, 4563 (2020).

[16] B. Cerulus, A. Jariani, G. Perez-Samper, L. Vermeersch, J. M. Pietsch, M. M. Crane, A. M. New, B. Gallone, M. Roncoroni, M. C. Dzialo, S. K. Govers, J. O. Hendrickx, E. Galle, M. Coomans, P. Berden, S. Verbandt, P. S. Swain, and K. J. Verstrepen, Transition between fermentation and respiration determines history-dependent behavior in fluctuating carbon sources, eLife **7**, e39234 (2018).

[17] A. M. New, B. Cerulus, S. K. Govers, G. Perez-Samper, B. Zhu, S. Boogmans, J. B. Xavier, and K. J. Verstrepen, Different levels of catabolite repression optimize growth in stable and variable environments, PLoS Biol. **12**, e1001764 (2014).

[18] L. Vermeersch, G. Perez-Samper, B. Cerulus, A. Jariani, B. Gallone, K. Voordeckers, J. Steensels, and K. J. Verstrepen, On the duration of the microbial lag phase, Curr. Genet. **65**, 721 (2019).

[19] G. Perez-Samper, B. Cerulus, A. Jariani, L. Vermeersch, N. Barrajón Simancas, M. M. M. Bisschops, J. van den Brink, D. Solis-Escalante, B. Gallone, D. De Maeyer, E. van Bael, T. Wenseleers, J. Michiels, K. Marchal, P. Daran-Lapujade, and K. J. Verstrepen, The Crabtree effect shapes the Saccharomyces cerevisiae lag phase during the switch between different carbon sources, mBio **9**, e01331-18 (2018).

[20] Y. Li, S. Venkataram, A. Agarwala, B. Dunn, D. A. Petrov, G. Sherlock, and D. S. Fisher, Hidden complexity of yeast adaptation under simple evolutionary conditions, Curr. Biol. **28**, 515 (2018).

[21] L. Vermeersch, L. Cool, A. Gorkovskiy, K. Voordeckers, T. Wenseleers, and K. J. Verstrepen, Do microbes have a memory? History-dependent behavior in the adaptation to variable environments, Front. Microbiol. **13**, 1004488 (2022).

[22] A. Murugan, K. Husain, M. J. Rust, C. Hepler, J. Bass, J. M. Pietsch, P. S. Swain, S. G. Jena, J. E. Toettcher, A. K. Chakraborty, *et al.*, Roadmap on biology in time varying environments, Phys. Biol. **18**, 041502 (2021).

[23] G. Lambert and E. Kussell, Memory and fitness optimization of bacteria under fluctuating environments, PLoS Genet. **10**, e1004556 (2014).

[24] M. Basan, T. Honda, D. Christodoulou, M. Hörl, Y.-F. Chang, E. Leoncini, A. Mukherjee, H. Okano, B. R. Taylor, J. M. Silverman, C. Sanchez, J. R. Williamson, J. Paulsson, T. Hwa, and U. Sauer, A universal trade-off between growth and lag in fluctuating environments, Nature (London) **584**, 470 (2020).

[25] S. J. Schink, D. Christodoulou, A. Mukherjee, E. Athaide, V. Brunner, T. Fuhrer, G. A. Bradshaw, U. Sauer, and M. Basan, Glycolysis/gluconeogenesis specialization in microbes is driven by biochemical constraints of flux sensing, Mol. Syst. Biol. **18**, MSB202110704 (2022).

[26] A. Mukherjee, J. Ealy, Y. Huang, N. C. Benites, M. Polk, and M. Basan, Coexisting ecotypes in long-term evolution emerged from interacting trade-offs, Nat. Commun. **14**, 3805 (2023).

[27] D. A. F. Fuentes, P. Manfredi, U. Jenal, and M. Zampieri, Pareto optimality between growth-rate and lag-time couples metabolic noise to phenotypic heterogeneity in *Escherichia coli*, Nat. Commun. **12**, 3204 (2021).

[28] F. Schreiber, S. Littmann, G. Lavik, S. Escrig, A. Meibom, M. M. M. Kuypers, and M. Ackermann, Phenotypic heterogeneity driven by nutrient limitation promotes growth in fluctuating environments, Nat. Microbiol. **1**, 16055 (2016).

[29] G. Chure and J. Cremer, An optimal regulation of fluxes dictates microbial growth in and out of steady state, eLife **12**, e84878 (2023).

[30] M. Zhu and X. Dai, Shaping of microbial phenotypes by trade-offs, Nat. Commun. **15**, 4238 (2024).

[31] O. Shoval, H. Sheftel, G. Shinar, Y. Hart, O. Ramote, A. Mayo, E. Dekel, K. Kavanagh, and U. Alon, Evolutionary trade-offs, Pareto optimality, and the geometry of phenotype space, Science **336**, 1157 (2012).

[32] D. Bertsekas, *Dynamic Programming and Optimal Control*, Vol. I, 4th ed. (Athena Scientific, Belmont, MA, 2017).

[33] S. Moreno-Gámez, D. J. Kiviet, C. Vulin, S. Schlegel, K. Schlegel, G. S. van Doorn, and M. Ackermann, Wide lag

time distributions break a trade-off between reproduction and survival in bacteria, Proc. Natl. Acad. Sci. USA **117**, 18729 (2020).

[34] S. Boulineau, F. Tostevin, D. J. Kiviet, P. R. t. Wolde, P. Nghe, and S. J. Tans, Single-cell dynamics reveals sustained growth during diauxic shifts, PLoS One **8**, e61686 (2013).

[35] M. Mori, S. Schink, D. W. Erickson, U. Gerland, and T. Hwa, Quantifying the benefit of a proteome reserve in fluctuating environments, Nat. Commun. **8**, 1225 (2017).

[36] O. Kotte, B. Volkmer, J. L. Radzikowski, and M. Heinemann, Phenotypic bistability in *Escherichia coli*'s central carbon metabolism, Mol. Syst. Biol. **10**, MSB135022 (2014).

[37] A. George, Phenotypic heterogeneity in lag reflects an evolutionarily stable bet-hedging strategy, Zenodo, 2025, https://doi.org/10.5281/zenodo.17582890.

[38] T. M. Norman, N. D. Lord, J. Paulsson, and R. Losick, Stochastic switching of cell fate in microbes, Annu. Rev. Microbiol. **69**, 381 (2015).