🔓 OPEN ACCESS

Effective bet-hedging through growth rate dependent stability

Daan H. de Groot[a,b,1,2] (iD), Age J. Tjalma[b,c,1] (iD), Frank J. Bruggeman[b] (iD), and Erik van Nimwegen[a,2] (iD)

Microbes in the wild face highly variable and unpredictable environments and are naturally selected for their average growth rate across environments. Apart from using sensory regulatory systems to adapt in a targeted manner to changing environments, microbes employ bet-hedging strategies where cells in an isogenic population switch stochastically between alternative phenotypes. Yet, bet-hedging suffers from a fundamental trade-off: Increasing the phenotype-switching rate increases the rate at which maladapted cells explore alternative phenotypes but also increases the rate at which cells switch out of a well-adapted state. Consequently, it is currently believed that bet-hedging strategies are effective only when the number of possible phenotypes is limited and when environments last for sufficiently many generations. However, recent experimental results show that gene expression noise generally decreases with growth rate, suggesting that phenotype-switching rates may systematically decrease with growth rate. Such growth rate dependent stability (GRDS) causes cells to be more explorative when maladapted and more phenotypically stable when well-adapted, and we show that GRDS can almost completely overcome the trade-off that limits bet-hedging, allowing for effective adaptation even when environments are diverse and change rapidly. We further show that even a small decrease in switching rates of faster-growing phenotypes can substantially increase long-term fitness of bet-hedging strategies. Together, our results suggest that stochastic strategies may play an even bigger role for microbial adaptation than hitherto appreciated.

microbial adaptation | phenotypic heterogeneity | bet-hedging | growth rate dependent stability | microbial population dynamics

Many microbial organisms exhibit a remarkable ability to adapt their internal state to environments that are highly variable and can change in unpredictable ways. For example, not only will there be different types of carbon sources, nitrogen sources, and amino acids available in different environments but also the concentrations of all these nutrients may vary over orders of magnitude. In addition, general variables such as temperature, pH, osmotic pressure, and oxygen availability will vary, and cells may have to withstand a wide array of specific stresses such as antibiotics or reactive oxygen species. It is remarkable that microbes appear to be able to adapt to this enormous number of possible combinations of environmental variables, not only because it requires coordinating the expression of many genes but also because it seems unlikely that the microbes can have been specifically selected for adapting to all these environments. For example, *Escherichia coli* is able to adapt its gene expression in order to allow it to grow in fully deuterated water, a highly unnatural condition in which the rates of most reactions involving water molecules are significantly altered, even though it almost certainly never encountered such a condition in the wild (1).

It is well known that microbes have evolved sensory regulatory machinery that can sense a large variety of internal and environmental variables and adapt gene expression patterns in response. In principle, the more information an organism gathers about its environment, the better it can adapt to it (2, 3). However, sensory regulatory strategies face a number of limitations. First, sensing is limited to those environmental variables for which the organism has evolved sensors, which is likely only a subset of the many environmental variables that affect optimal gene expression states. Second, given the small number of molecules involved, there are fundamental thermodynamic limits on the accuracy with which cells can gather information about their environment (4, 5). Third, even if cells are able to gather accurate information about the state of environmental variables, the processing and integrating of this information so as to optimally set gene expression levels is a nontrivial regulatory problem. Moreover, the regulation of gene expression is itself also significantly affected by thermodynamic noise. Finally, there may be intrinsic costs associated with sensory regulatory machinery, be it through the cost of expressing proteins that do not directly contribute to growth (6, 7) or due to energetic costs (8).

Significance

How is it possible that microbial populations can robustly adapt to an immense variety of environments when their regulatory circuitry is extremely noisy and of limited accuracy? We propose that the remarkable adaptability of microbial populations can be explained by the combined effects of two experimentally observed phenomena: 1) that gene regulation is inherently noisy, causing individual cells to randomly switch between different phenotypic states, and 2) that the slower cells grow, the noisier they become. As a consequence, whenever cells encounter a new environment, they stochastically explore new phenotypes, only to stabilize after reaching a fast-growing phenotype. We propose that this growth rate dependent stability may play an important role in adaption of all microbes.

Apart from adapting gene expression in a regulated manner, it is known that due to the inherent noise in gene expression processes, isogenic cells also stochastically switch between different phenotypes. In this way, cells explore alternative phenotypes, and since subpopulations that happen to venture onto fast-growing phenotypes will automatically expand because of their higher growth rate, the population will appear to effectively "adapt" its phenotype to the environment. This general mechanism, which in principle allows microbes to adapt to a wide variety of unexpected environmental changes, including those that they are unable to sense, is typically referred to as a bet-hedging strategy (9–16). Although the term "bet-hedging strategy" is of course anthropomorphic, we will adopt it here to refer to situations where stochastic phenotype switching of individual cells is adaptive at the population level.

However, the long-term fitness that can be attained with such bet-hedging strategies, i.e., the long-term average population growth rate (17, 18), is limited by an intrinsic trade-off: Increasing the stochastic phenotype-switching rate speeds up adaptation to new environments, but it also decreases the long-term growth rate in each environment since it increases the rate at which cells switch out of well-adapted phenotypes (19). Due to this inherent trade-off, as demonstrated both by theoretical modeling (9, 19, 20) and by experimental approaches (21, 22), bet-hedging strategies are effective only when durations of environments are large relative to the doubling times of the cells, and when the number of possible environments that populations need to anticipate is limited.

Several recent studies have observed that gene expression noise generally increases at low growth rates (23–25). Since many phenotype switches may be driven by gene expression fluctuations (26–30), these observations suggest that there may be intrinsic coupling between the growth rate of cells and their phenotypic stability. That is, stochastic phenotype-switching rates may naturally be higher for slow-growing cells than for fast-growing cells. Intuitively, it seems that such growth rate dependent stability (GRDS) could benefit bet-hedging strategies since it would reduce the rate at which well-adapted cells switch to maladapted phenotypes, while at the same time increasing the rate at which maladapted cells explore alternatives. Indeed, as has been shown by Kaneko et al. in a nonevolutionary setting (31, 32), when phenotypic stability increases with growth rate, the distribution of phenotypes in the population is shifted toward faster-growing phenotypes. In conceptually related work, Schreier et al. (33) proposed a general abstract model of phenotypic adaptation in which cells randomly diffuse through phenotype space until they reach a desired phenotype, although no concrete mechanisms were proposed for how cells would sense that they had reached a desired phenotype.

Here, we systematically investigate the effect of GRDS on the performance of bet-hedging. In particular, we extend the basic model of a population evolving in a changing environment introduced in ref. 19, and using a combination of analytical solutions and numerical simulations, we determine how GRDS affects the long-term average growth rates that stochastically switching populations can achieve. We show that even a small growth rate dependence of the phenotype-switching rates immediately increases the average growth rate of bet-hedging populations and that GRDS can completely resolve the inherent trade-off of traditional bet-hedging strategies when the ratio of switching rates of slow- and fast-growing cells is sufficiently high. We also find that GRDS can improve average population growth rates through two qualitatively distinct mechanisms depending

on whether environment durations are short or long relative to the doubling time of adapted phenotypes. Taken together, our results show that GRDS dramatically expands the range of scenarios for which stochastic bet-hedging strategies can attain a high long-term average growth rate.

Results

Model Setup. To investigate the effects of growth rate dependent stability (GRDS) on bet-hedging strategies for microbial populations growing in changing environments, we extend the general model introduced by Kussell and Leibler (19). We consider a population of cells that switch stochastically between $n + 1$ discrete phenotypes and grow in an environment that switches stochastically between m discrete environment types. For a given realization of the stochastic environment switches, the sequence of environments is described by the function $\mathcal{E}(t)$, denoting which environment type is present at time t. The average time that an environment of type J remains before it switches to another environment type is τ_J. The order in which environment types occur is determined by a Markov chain: upon a switch, the probability that environment J is followed by environment I is denoted by b_{IJ}, where $\sum_I b_{IJ} = 1$ and $b_{II} = 0$. In environment J, cells with phenotype i grow at a rate $\mu_i^{(J)}$ and switch stochastically from phenotype j to i at rate $\phi_{ij}^{(J)}$.

The population is assumed to be sufficiently large such that the population dynamics can be modeled by a set of deterministic differential equations. The state of the population is described by an $(n + 1)$-dimensional vector $\boldsymbol{s}$, containing the number of cells of each phenotype, whose dynamics are described by

$$\frac{d}{dt}\boldsymbol{s} = A^{\mathcal{E}(t)}\boldsymbol{s}, \qquad [1]$$

where $A^{\mathcal{E}(t)}$ is the time evolution matrix of environment $\mathcal{E}(t)$. The components $A_{ij}^{(K)}$ of the time evolution matrix of environment K are given by

$$A_{ij}^{(K)} = \left[\mu_i^{(K)} - \sum_{k \neq i} \phi_{ki}^{(K)}\right]\delta_{ij} + \phi_{ij}^{(K)}, \qquad [2]$$

where δ_{ij} is the Kronecker delta function that is one when $i = j$ and zero otherwise.

Up to this point, this model is identical to the general model used in (19) for modeling a bet-hedging population in fluctuating environments. In the classical bet-hedging scenario of ref. 19, the switching rates $\phi_{ij}^{(K)}$ are assumed independent of the environment, i.e., $\phi_{ij}^{(K)} = \phi_{ij}$ for all K. To investigate the effects of GRDS, we extend this classical model by allowing the switching rates $\phi_{ij}^{(K)}$ to be functions of the current growth rate of the cell, i.e., $\phi_{ij}^{(K)} = \phi_{ij}f_j\left(\mu_j^{(K)}\right)$, and generally assume that $f_j(\mu)$ is a decreasing function of the growth rate μ. That is, given the same phenotype j, more slowly growing cells are more "restless" than faster-growing cells and switch their phenotype more often. However, the relative rates of switching to different phenotypes i are always the same.

We will quantify the strength of GRDS by the parameter r, which is the ratio of switching rates between the fastest- and

slowest-growing phenotypes, i.e., when phenotype j achieves its maximal growth rate in environment J and its minimal growth rate in environment K, we get $\phi_{ij}^{(K)} = r\phi_{ij}^{(J)}$. Note that the case where $r = 1$ thus corresponds to classical bet-hedging without GRDS.

The total number of cells at time t is given by $N(t) = \sum_i s_i(t)$, and we are generally interested in the average growth rate G of the population over a long sequence of environments, i.e.,

$$ G = \lim_{t \to \infty} \frac{1}{t} \log \left(\frac{N(t)}{N(0)} \right), \qquad [3] $$

which quantifies the "fitness" of a given strategy (3, 17, 18). In particular, we will compare the average growth rates G that are obtained with classical bet-hedging with those obtained with GRDS.

A Toy Example That Qualitatively Illustrates the Benefits of GRDS. We use a toy example of this general model to illustrate the differences in behavior between classical bet-hedging and bet-hedging with GRDS (Fig. 1). In this example, there are only three environments and phenotypes (shown as purple, red, and green); in each environment, one phenotype is optimal and leads to a growth rate of $\mu_1 = 1.0$, while the other phenotypes have a growth rate of $\mu_0 = 0$. For the population dynamics shown here, all cells start out in the green phenotype and encounter the purple environment, followed by the red environment; both environments have a duration of $T = 10$. We assume that cells in the optimal phenotype switch with a rate ϕ, while the other cells switch at a rate $r\phi$, where $r = 10$ for bet-hedging with GRDS (and $r = 1$ for classical bet-hedging). The switching rate ϕ was optimized to maximize G for both strategies separately.

The inherent trade-off of classical bet-hedging is illustrated in Fig. 1A. Because the switching rate is independent of the environment, a cell in a given phenotype is equally likely to switch independent of whether this phenotype is well or badly adapted to the current environment. Thus, increasing the switching rate will not only increase the rate at which cells in maladapted phenotypes explore alternative phenotypes but will also increase the rate at which cells with optimal phenotype switch to worse phenotypes. The trade-off thus arises because a high long-term growth rate requires a relatively low switching rate, while fast adaptation requires a high switching rate. Indeed, at the switching rate ϕ that optimizes this trade-off, the adaptation takes about half the environment duration, but speeding this up would require a higher switching rate which would decrease the long-term population growth rate (Fig. 1B, blue curve).

GRDS can largely resolve this trade-off by decoupling the rate of exploration by nonadapted cells, i.e., $r\phi$, from the rate of switching of well-adapted cells, i.e., ϕ. This makes it possible to speed up the adaptation to a new environment, while the same long-term population growth rate is reached (Fig. 1B, red curve). Note that the population in each environment stabilizes with a similar fraction of cells in suboptimal phenotypes as with the classical bet-hedging strategy (Fig. 1 C, *Bottom*). This shows that, at least in this setting, GRDS mainly works because it allows for a "panic mode": Immediately after an environment change, the growth rate of the majority of cells drops so that their phenotype-switching rates increase, quickly generating a heterogeneous population of cells that explore different phenotypes. Moreover, these exploring cells stabilize only once they find the phenotype that supports fast growth. Since these adapted cells are both more stable and grow faster, the population quickly becomes dominated by optimized cells again (Fig. 1 C, *Bottom*).
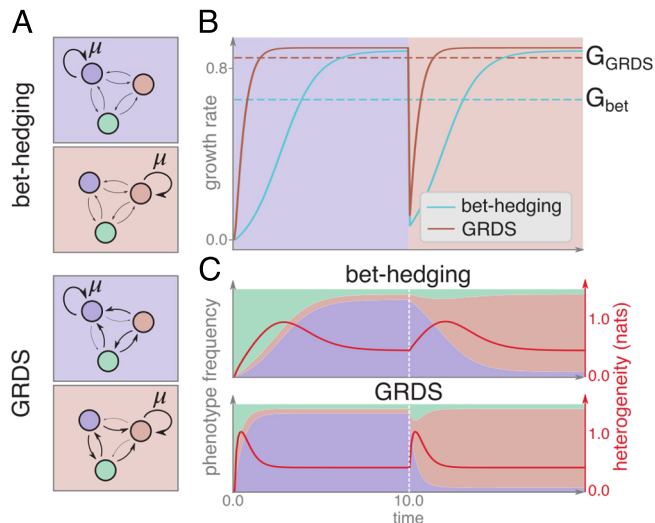

Fig. 1. A toy example illustrates how GRDS increases the effectiveness of bet-hedging. (A) Growth was simulated in a sequence of two environments (purple and red background) for cells with three different phenotypes: optimal for the purple, red, or green environment. The growth of a phenotype is indicated by the arrow marked by μ, while the other arrows indicate switching between phenotypes. Bolder lines indicate higher rates. Random phenotype-switching rates are either constant (bet-hedging) or growth rate dependent (GRDS). The optimal tuning of the switching rates ϕ to maximize the average growth rates G resulted in $\phi = 0.13$ for bet-hedging without GRDS, and $\phi = 0.26$ for bet-hedging with GRDS. (B) Population growth rates as a function of time for bet-hedging without GRDS (blue curve) and with GRDS (red curve), starting from an initial condition with all cells in the green phenotype. The average growth rates for both strategies are indicated by dashed lines. (C) Time courses of the fractions of the population in each of the phenotypes (colors) for the bet-hedging (*Top*) and GRDS (*Bottom*) strategies. The red curves show time courses of the population heterogeneity, defined as the entropy of the distribution of phenotypes in the population. The parameters used for these simulations, as described in the Model Setup section, were $n + 1 = 3$, $T = 10$, $\mu = 1.0$, and $r = 10$.

A Minimal Model of Bet-Hedging with GRDS. To quantify the effect of GRDS on bet-hedging strategies, we first analyze a model in which the number of parameters is reduced to a minimum and which can be solved analytically. In this minimal model, we set the number of environments and phenotypes both to $n + 1$. We assume that each environment lasts for a fixed time T and then switches to another environment with uniform probability, i.e., $b_{IJ} = 1/n$ for all $I \neq J$. In addition, we assume that there are only two possible growth rates, a "fast" growth rate μ_1 and a "slow" growth rate μ_0, i.e., $\mu_1 > \mu_0$, and that in each environment I, only cells in a "good" phenotype $i = I$ grow at the fast rate μ_1 and cells in the n ("bad") other phenotypes grow at the slow rate μ_0. Regarding phenotype switching, we assume that whenever a cell switches its phenotype, it is equally likely to switch to any of the n other phenotypes. Finally, to tune the overall switching rates and amount of GRDS, we assume that cells in the fast-growing phenotype $i = I$ switch out of their phenotype at a total rate ϕ, whereas cells in any of the n slow-growing phenotypes switch out of their phenotype at a total rate $r\phi$ with $r \geq 1$. Thus, formally, the switching rates are given by $\phi_{ij}^{(J)} = \phi/n$ for all $i \neq j$ when $j = J$ and $\phi_{ij}^{(J)} = r\phi/n$ for all $i \neq j$ and $j \neq J$.

Because switching rates to and from all bad phenotypes are equal in this model, we can describe the state of the population by the number of cells $s_g(t)$ and $s_b(t)$ in the good and bad phenotypes, respectively (more details are provided in *SI Appendix*, section 1.B.1). Because the growth rates μ_1, μ_0 and environment duration T are the same in each environment, we

can restrict ourselves to solving models where the bad phenotype does not grow at all, i.e., $\mu_0 = 0$, and measure time in units such that $\mu_1 = 1$, without loss of generality (*SI Appendix*, section 1.B). Solutions for any other setting of the growth rates μ_1 and μ_0 can then be obtained by rescaling and shifting the resulting time dynamics. This leaves us with the following differential equations for the number of cells with good and bad phenotypes

$$\frac{ds_g(t)}{dt} = (1 - \phi)s_g(t) + \frac{\phi r}{n}s_b(t),$$

$$\frac{ds_b(t)}{dt} = -\frac{\phi r}{n}s_b(t) + \phi s_g(t),$$

[4]

where ϕ is the switching rate of adapted ("good") phenotype cells and $r\phi$ the switching rate of cells in a bad phenotype. Note that because a cell in a bad phenotype can switch to n other phenotypes and only one of these is the good phenotype, the effective switching rate of cells from a bad to the good phenotype is $\phi r/n$. Eq. 4 would not change if instead of one good phenotype and n bad ones, we assumed that there were K good phenotypes and nK bad ones. This suggests interpreting the probability $1/n$ more generally as the probability that, under a stochastic phenotype switch, a cell in a bad phenotype will switch to a good phenotype. That is, the parameter n quantifies how rare fast-growing phenotypes are and thereby quantifies the complexity of the fluctuating environment.

When the environment changes, a new phenotype will become optimal so that all cells that were in the good phenotype now find themselves in a bad phenotype, whereas some cells that were in a bad phenotype may coincidentally find themselves in the new good phenotype. In our analysis of the minimal model, we will make the approximation that all cells that were in a bad phenotype in the previous environment have a probability $1/n$ to find themselves in the good phenotype of the new environment. Thus, if we denote the fraction of adapted cells at the end of an environment as $p(T)$, the initial fraction of adapted cells in the next environment is $(1 - p(T))/n$. In *SI Appendix*, section 1.C.2 we explain why, under mild conditions on the model parameters, this approximation is justified.

The dynamics of this system can be solved analytically (*SI Appendix*, sections 1.C and 1.D) yielding an expression for the long-term average growth rate G that is fully determined by the four parameters T, n, r, and ϕ. Of these parameters, n and T parametrize the regulatory problem that the microbial population faces: The complexity of the environment is set by n, and T sets the number of generations between environmental changes which determines the relative importance of fast adaptation versus a high stationary growth rate. In turn, the strength of GRDS r and the switching rate ϕ set the behavior of the cellular population.

To further analyze this minimal model, we assume that natural selection has optimized the switching rate ϕ to maximize the average growth rate G. In Fig. 2, we systematically investigate how the resulting average growth rate G varies with n, T, and the strength of GRDS r. We vary n and T on the vertical axes of Fig. 2 *A* and *B*, while the strength of GRDS increases along the horizontal axes. The growth rates for classical bet-hedging correspond to $r = 1$ and are thus shown along the vertical axes in these plots. We see that, as derived previously (19), the fitness of a bet-hedging population decreases with the complexity of the environment n and increases with the environment duration T for classical bet-hedging. Increasing GRDS by moving away from the vertical axes at $r = 1$ in Fig. 2 *A* and *B*, we see that GRDS can dramatically increase the average growth rate and that increasing r is always beneficial. That is, at least in this minimal model,
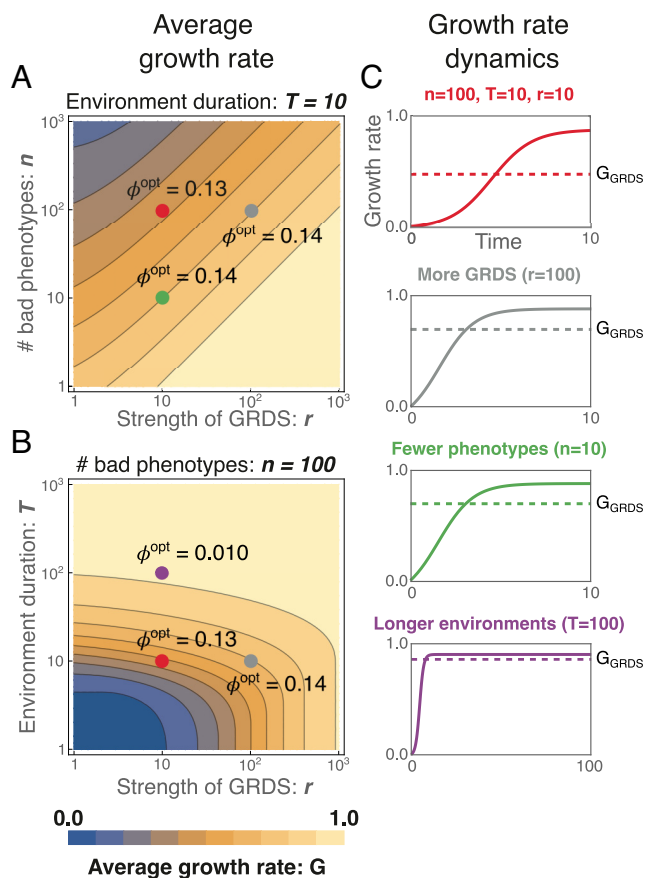

Fig. 2. Effect of GRDS on the average growth rate of bet-hedging strategies as a function of environment complexity and duration. The contour plots in (*A*) and (*B*) show the average population growth rate (G) at optimal switching rate ϕ as a function of the strength of GRDS (horizontal axes, both panels) and either the number of environments (*A*) or the environment duration (*B*). Since $\mu_1 = 1$ and $\mu_0 = 0$, $0 \leq G \leq 1$ and the contours occur at integer multiples of 0.1. The optimal switching rates ϕ at four example parameter settings (colored dots) are indicated in the contour plots. All axes are on logarithmic scales. Contour plots for additional parameter settings are shown in *SI Appendix*, Figs. S3 and S4. (*C*) Population growth rate versus time over the course of one environment for the parameter sets indicated by colored dots in *A* and *B*. The dashed lines indicate the average growth rate G. Note that the gray, green, and purple parameter settings each differ from the red parameter setting by a change in one parameter.

evolution would favor making the ratio r between the switching rates at slow and fast growth as large as possible. Moreover, provided that the strength of GRDS r is made sufficiently large, the average growth rate G can approach the theoretical optimum $G = 1$ arbitrarily closely.

To illustrate the dependence of the average growth rate G on the parameters, we start from a reference parameter set $n = 100$, $T = 10$, and $r = 10$ (red dots in Fig. 2 *A* and *B*) and increase r by a factor of ten (gray dots Fig. 2 *A* and *B*), decrease n by a factor of ten (green dot in Fig. 2*A*), or increase T by a factor of ten (purple dot in Fig. 2*B*). When the strength of GRDS r is increased by a factor of ten, this has a similar effect on growth rate G as decreasing the number of bad phenotypes n by ten, and more generally, the approximately straight diagonal contours in Fig. 2*A* show that G effectively depends only on the ratio r/n. This can be understood by noting that the population dynamics of Eq. 4 depends on r and n only through the ratio r/n. Although the initial fraction of good phenotype cells after an environment switch does depend directly on n and not r, we find that this is negligible when T is not small.

Long environment durations. Of the three parameter changes in Fig. 2, the optimal switching rate ϕ^{opt} changes significantly only under a change in environment duration T and is largely insensitive to changes in the other parameters. Indeed, as shown in *SI Appendix*, Fig. S7, we find that when T is sufficiently large (i.e., around $T = 10$ generations or larger), and r is not too large (i.e., $r < nT$, as discussed below), the optimal switching rate is just the inverse of the environment duration:

$$\phi^{\mathrm{opt}} \approx \frac{1}{T} \qquad \text{(large T, } r < nT\text{),} \qquad \text{[5]}$$

and in *SI Appendix*, section 2.B.1, we mathematically prove that this relationship holds exactly in the limit of long environment durations. Thus, the optimal phenotype-switching rate for cells in the fast-growth phenotype exactly equals the rate at which the environment changes. This extends a classical result for conventional bet-hedging which states that the probability of using a strategy should equal the probability that the strategy will become useful in the future (13, 19, 20, 34). Note that in this large T parameter regime, the optimal switching rate is independent of r and thus the same for classical bet-hedging and bet-hedging with GRDS. This shows that the benefits of GRDS derive not from greater stability of cells in the optimal phenotype, so that GRDS does not diminish the subpopulations that are preadapted for possible future environment changes. Rather, the benefits of GRDS are due to the increased switching rate of the slow-growing cells. That is, GRDS allows slow-growing cells to "panic" and rapidly explore different phenotypes until a fast-growing phenotype is found. Such an adaptive transition of the population from stable to explorative and back is impossible without GRDS.

We also derived an analytical approximation for G in the parameter regime where environment durations are sufficiently long for the population to reach its steady-state distribution of phenotypes (*SI Appendix*, section 2.B.3):

$$G \approx 1 - \underbrace{\frac{1}{T}}_{\text{diversity cost}} - \underbrace{\frac{1}{T}\log\left[\frac{Tn}{2}\right]}_{\text{delay cost}} + \underbrace{\frac{1}{T}\log\left[\frac{1+r}{2}\right]}_{\text{GRDS effect}}. \qquad \text{[6]}$$

The terms in this equation have intuitive interpretations: First, at the optimal switching rate $\phi^{\mathrm{opt}} = 1/T$, the long-term growth rate is $1 - 1/T$ because a fraction $1/T$ of the population will be in nongrowing phenotypes. In ref. 19, this is referred to as the "diversity cost" of bet-hedging. Second, at the end of a given environment, the nongrowing cells will be equally distributed over the n nongrowing phenotypes, so that the fraction of cells in the "good" phenotype just after an environment switch will be $1/(nT)$. The "delay cost" $\log[nT/2]/T$ corresponds to the reduction of the average growth rate from having to expand this small subpopulation. The final term $\log[(1+r)/2]/T$ is unique to bet-hedging with GRDS and quantifies the extent to which GRDS can compensate the intrinsic delay and diversity costs of bet-hedging. We have validated numerically that (6) provides an excellent approximation to G as long as environment durations are not short, i.e., $T \geq 10$ generations (*SI Appendix*, Fig. S8). *SI Appendix*, Fig. S8 also shows that once r becomes so large that the delay costs are fully compensated, i.e., when $r \geq Tn$, Eq. 6 starts to overestimate G, and the true growth rate saturates toward $G = 1$ when r increases further. As discussed in the next section, in this regime, the optimal switching rate no longer equals $1/T$ and the approximation breaks down.

We derived an analogous expression to Eq. 6 for the more general model, including differing average environment durations τ_J and transition probabilities b_{IJ}, using the same assumptions as were used in ref. 19 for the case without GRDS (*SI Appendix*, section 2.B.3). In this more general setting, GRDS still increases the average growth rate with $\log[(1+r)/2]/\tau$, giving

$$G \approx \mu_{\mathrm{m}} - \underbrace{\frac{1}{\tau}}_{\text{diversity cost}} - \underbrace{\sum_I \frac{p_I}{\tau}\log\left[\frac{\tau_I}{2}\right] - \frac{S_{\mathrm{env}}}{\tau}}_{\text{delay cost}}$$
$$+ \underbrace{\frac{1}{\tau}\log\left[\frac{1+r}{2}\right]}_{\text{GRDS effect}}, \qquad \text{[7]}$$

where μ_{m} is the growth rate of the adapted phenotype in each environment, $\tau = \sum_I p_I \tau_I$ is the average environment duration, p_I is the probability of environment I occurring, and $S_{\mathrm{env}} = -\sum_{I,J} b_{IJ} p_J \log(b_{IJ})$ is the environment entropy, which quantifies how hard it is on average to predict the next environment given the current environment. In Eq. 6, this entropy term was $\log(n)$ because we assumed that after each environment, all other environments were equally likely to occur next. In general, the change of environments can be much more predictable, for example, when environment I is always followed by environment J (i.e. $b_{IJ} = 1$), which decreases the delay cost for bet-hedging since switching rates can be adapted to this predictability.

Eq. 7 shows that GRDS can compensate the intrinsic costs of bet-hedging, including the uncertainty about the environment captured by S_{env}. The logarithmic dependence also shows that even a slight growth rate dependence of the phenotype-switching rates can already provide a substantial fitness advantage. Finally, Eq. 7 implies that a population employing GRDS outgrows a population employing classical bet-hedging by a factor $(1+r)/2$ over the course of each environment, independent of the number of environment types, the transition rates b_{IJ} between them, or their durations τ_I.

Short environment durations. For classical bet-hedging, it is known that, when environment durations are short relative to the doubling time of the fastest-growing cells, bet-hedging is ineffective because natural selection has no time to expand the subpopulation of fast-growing cells. However, Fig. 2B shows that, even when $T = 1$, bet-hedging with GRDS can reach close to maximal fitness provided that r is made sufficiently large. Interestingly, in this small T regime, the optimal strategy is to switch as fast as possible, i.e., $\phi^{\mathrm{opt}} \to \infty$. In this limit of very fast switching, the steady-state fraction of cells in the fast-growth phenotype is given by $p_{\phi\to\infty} = r/(n + r)$. Due to the fast switching, this steady state is reached very quickly and, consequently, the average population growth rate becomes

$$G = \frac{r}{n + r}. \qquad \text{[8]}$$

Although an infinite switching rate is not realistic, this strategy of fast switching is already effective when ϕ is large compared to the rate at which environments switch ($1/T$) and compared to the growth rate of adapted cells ($\mu = 1$).

GRDS can thus aid adaptation through two qualitatively different strategies. When environment durations T are large, fast-growing cells come to dominate the population by outgrowing the slow-growing cells, and the optimal strategy is for fast-growing cells to switch relatively infrequently, i.e., at

the same rate as the environment switches. In contrast, when environment durations T are short, fast-growing cells can still come to dominate the population when r is sufficiently large simply because the phenotypes of faster-growing cells are more stable than those of more slowly growing cells, i.e., as previously identified in refs. 31 and 32. In this regime, fitness is optimized by maximizing switching rates.

As shown in *SI Appendix*, Figs. S5 and S6, the boundary between the long-T regime where $\phi = 1/T$ is optimal and the short-T regime where $\phi \to \infty$ is optimal approximately corresponds to the line $T = r/n$ when $T \geq 10$ generations. When T is small, high average growth rates can be achieved only using the short-T strategy and require $r \gg n$. Of course, there are also parameter regimes where no high average growth rate G can be achieved, even with GRDS. In particular, when both T is small and GRDS is small relative to the number of environments (i.e., $r < n$), high average growth rates cannot be achieved.

General Model: GRDS Almost Always Increases Population Fitness. Our calculations so far have quantified the benefits of GRDS in the simplified case with only two growth rates: a fast growth rate for adapted phenotypes and a slow growth rate for nonadapted phenotypes. Next, we investigated the more general settings including different environment durations τ_I, arbitrary environment switching rates b_{IJ}, and arbitrary growth rates μ_j^I of each phenotype j in each environment I. Following analogous assumptions previously made in ref. 19, we study the general model in the regime where environment durations are relatively large and, to simplify notation, assume that there are equally many environments as phenotypes, with phenotype i being optimal in environment i. In *SI Appendix*, section 2.B.1, we derive that with GRDS, the optimal switching rate from phenotype j to k in environment i is given by

$$\phi_{kj}^i = \frac{b_{kj}}{\tau_j} e^{\delta(\mu_j^i - \mu_j^j)}, \quad [9]$$

where $\delta \geq 0$ controls the strength of GRDS. This result generalizes the result of ref. 19 that each phenotype-switching rate ϕ_{kj} should match the corresponding environment switching rate b_{kj}/τ_j. With GRDS, i.e., when $\delta > 0$, we find that switching out of phenotype j is increased whenever the current growth rate μ_j^i is less than the growth rate μ_j^j in the environment j where phenotype j has the optimal growth rate. This also implies that, in each environment j, the optimal switching rates ϕ_{kj}^j out of the optimal phenotype j are the same with and without GRDS. Consequently, when switching rates are optimized, the fractions of preadapted individuals upon an environment switch are the same with and without GRDS. This further confirms our finding with the minimal model that, when environment durations are relatively large, the benefit that GRDS provides comes from increasing the switching rates of suboptimal phenotypes.

Next, we asked under what conditions on the parameters of the general model, including cases where the switching rates have not been optimized, a small amount of GRDS is beneficial. Assuming again that environment durations are relatively large, we show in *SI Appendix*, section 2.B.2 that whenever switching rates are not optimized, GRDS can generally bring switching rates closer to their optima. We further show that, when switching rates are optimized, GRDS is beneficial whenever, averaged over all environment switches $j \to i$, we have $\langle \mu_j^j \rangle > \langle \mu_j^i \rangle$. That is, GRDS is beneficial when, on average, the growth rate of the

current optimal phenotype j decreases when the environment switches from j to a new environment i. Although it is possible to design parameter settings that do not obey this constraint, this requires a careful tuning of parameters that is very unlikely to occur by chance (*SI Appendix*, section 2.B.2). Indeed, when we randomly sample growth rates μ_j^i from different distributions, GRDS is always beneficial (*SI Appendix*, Fig. S9). Notably, these derivations apply to the general model and are thus valid for an arbitrary number of phenotypes and environments for arbitrary growth rates. This includes cases where the growth rate of the optimal phenotype in one environment is still lower than the growth rate of the worst phenotype in another environment. It also holds regardless of whether all environments have the same duration or whether the phenotype-switching rates are optimized. This thus strongly suggests that GRDS is generically beneficial when environment durations are not short.

Finally, to quantify the extent of the growth rate improvement when switching rates are not optimized and to explore whether the benefits of GRDS also extend to the regime of short environment durations, we numerically computed the effect of GRDS on the average growth rate for many different parameter sets (Fig. 3). The parameter sets were picked as follows: We systematically varied the number of environments between 5 and 20, chose the environment switching probabilities b_{IJ} uniformly at random, and varied the average environment duration from $T = 1$ to $T = 40$. The number of phenotypes m was chosen to be equal, twice, or half the number of environments. For each environment, the growth rate of the fastest-growing phenotype was drawn randomly over a range, such that one unit of time on average corresponds to one doubling of the best phenotype. All other phenotypes were assigned growth rates at random, chosen uniformly over a range below the maximal growth rate in that environment. The ranges from which the
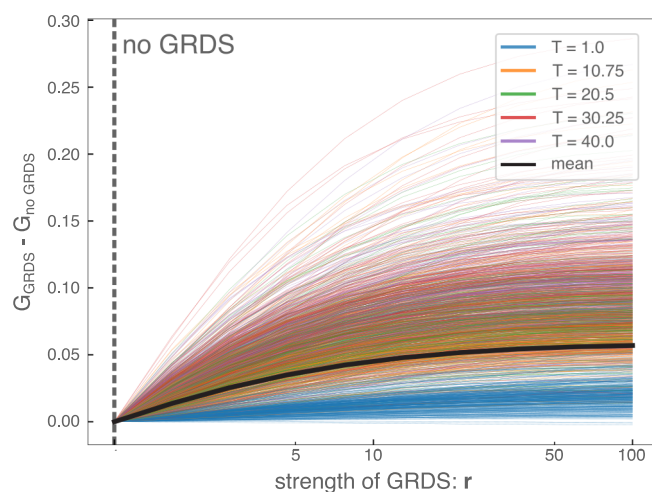

Fig. 3. Population growth rates rise with increasing strength of GRDS r, even when switching rates are not optimized. Each colored line corresponds to a different randomly picked parameter set and shows how the average population growth rate G increases when the strength of GRDS r is increased. Different colors correspond to different average environment durations T, where $T = 1$ corresponds roughly to a single doubling of the fastest-growing phenotype in each environment (*SI Appendix*, section 3). The bolder black line indicates the average. The vertical axis shows the difference between the average population growth rate with a certain strength of GRDS and the average growth rate without GRDS. The rate of switching from phenotype j to i decreases with the growth rate of phenotype j, and the ratio of the highest and lowest switching rate from j to i in different environments is r (*Methods* for details). The random sampling of the other parameters is described in *SI Appendix*, section 3.

growth rates were sampled were chosen such as to ensure the occurrence of environments in which the fastest growth rate was still slower than the slowest growth rate in other environments. The switching rate ϕ_{ij} for each pair of phenotypes was chosen randomly over a thousand-fold range. To model the effect of GRDS, we then add growth rate dependence to these randomly chosen switching rates, such that we get rates ϕ_{ij}^K that are higher in environments where phenotype j grows slow and lower where phenotype j grows fast. In particular, we let the logarithm $\log(\phi_{ij}^K)$ decrease linearly with growth rate μ_j^K such that ϕ_{ij}^K varies by a factor r across environments and equals ϕ_{ij} on average (*Methods* and *SI Appendix*, section 3 for details). That is, as before, r quantifies the strength of GRDS (plotted on the horizontal axis in Fig. 3), with $r = 1$ corresponding to traditional bet-hedging.

For almost all parameter sets (1,872 out of 1,875), a small growth rate dependence immediately increases the long-term growth rate that is achieved by a randomly switching population (Fig 3). The three exceptions all correspond to cases where environment durations are very short. Although the size of the fitness benefit of GRDS as a function of r varies across the randomly chosen parameter settings and is consistently smaller when the average environment duration is short, essentially all curves show an initial steep increase followed by a slower but nonsaturating increase with r, similar to the dependence observed for the minimal model with optimal switching rates, Eqs. **6** and **8**. These simulation results confirm that, even without optimized switching rates, GRDS is generically beneficial across a very wide range of parameter settings, including relatively short environment durations.

Discussion

Although it has long been observed that microbial populations of isogenic cells can exhibit significant phenotypic variability, over the last two decades, it has become increasingly appreciated that such phenotypic variability is pervasive and involves both continuous fluctuations in gene expression and stochastic switching between discrete phenotypic states (9, 10, 12, 14, 35–39). A large body of theoretical work has established that such phenotypic heterogeneity can be beneficial in fluctuating environments, which led to the suggestion that microbial populations may be employing "bet-hedging" strategies (35, 40, 41). Indeed, it has been shown that a bet-hedging strategy can be evolved in laboratory evolution experiments (14, 39). Moreover, it was shown that a bet-hedging strategy can be fixed in a population even if this population is of finite size (13) and that it leads to higher average growth rates especially if the population is colonizing new, unknown environments (42).

However, previous theoretical work has also suggested that bet-hedging strategies can be effective only if environments are not too diverse and environment durations are relatively long (9, 19, 20). These fairly restrictive bounds on the benefits of bet-hedging raise the question of to what extent the pervasive phenotypic heterogeneity that is observed in microbial populations can be attributed to bet-hedging. Here, we have shown that these bounds on the effectiveness of bet-hedging strategies disappear when we account for one additional ingredient: growth rate dependent stability (GRDS). With GRDS, phenotype-switching rates decrease with growth rate, and we have shown that, as the ratio between switching rates of slow- and fast-growing cells increases, the intrinsic costs of bet-hedging can

be compensated, and average population growth rates can approximate the theoretical maximum.

There are clear conceptual similarities between GRDS and stress-induced mutagenesis (SIM), which proposes that genotypic adaptation is aided when the mutation rate increases in cells that experience stress (43–45). Indeed, it may seem that much of the mathematical framework that we developed above could be applied to a model of growth rate dependent genotype "switching" as well. However, there are important differences between GRDS and SIM that make extending our analytical framework to SIM difficult. First, genotype changes through mutation of course occur on a much slower time scale than phenotype changes. Second, while growth rate directly reflects fitness, in SIM the mutation rate is proposed to be coupled to "stress," and it is unclear to what extent the stresses that increase the mutation rate systematically lower growth rates. In fact, as far as we can tell, it is currently not clear to what extent stress would increase mutation rate, and a recent in-depth study of mutation in single cells did not uncover large variation in mutation rates across single cells (46). In addition, an important assumption in our work is that when phenotype-switching rates are increased through GRDS, cells will switch their phenotype more often. This is hard to imagine in the context of SIM since a higher mutation rate will not only lead to more genotype switches but also to more mutations per genotype switch, which will drastically increase the probability of deleterious mutations.

Bet-hedging strategies are typically contrasted to sensing strategies, where cells detect features of their environment and change their phenotype in particular directions in response. Since GRDS requires that cells must be able to "sense" their own growth rate in some way, one may wonder to what extent GRDS corresponds to a partial sensing strategy. However, since the same growth rate μ can occur in different environments depending on the cell's phenotype, GRDS does not allow cells to directly sense their environment. With GRDS, the switching rates $\phi_{ij}^k = \phi_{ij}f_j(\mu_j^k)$ decrease with growth rate μ according to some decreasing function $f_j(\mu)$. Thus, for each phenotype j, the function $f_j(\mu)$ allows cells to decide whether their current growth rate μ is low or high relative to growth rates that can be achieved in other environments with the same phenotype j, but it cannot inform cells as to whether an other phenotype i may achieve a higher growth rate in the current environment. And crucially, which phenotypes cells will switch to is just as random with GRDS as it is in traditional bet-hedging since the relative switching rates to different phenotypes stay the same. Therefore, GRDS could be summarized by saying that it allows cells to sense something about themselves, i.e., whether they are doing "well" or "poorly" given their current phenotype, but it does not allow cells to sense anything about their current environment.

There is in fact significant evidence supporting that phenotype-switching rates tend to decrease with growth rate. In a number of studies, it has been observed that gene expression noise levels decrease with growth rate (23–25, 31), and metabolic heterogeneity has also been observed to increase with nutrient limitation (47–52). Since phenotype switches are often ultimately driven by fluctuations in gene expression or metabolic state (11, 26–30, 53), phenotype-switching rates will generally increase with gene expression noise levels. Notably, if we assume that a particular phenotype switch occurs under a particular rare fluctuation in gene expression, then even a small change in the noise level can have a large effect on the phenotype-switching rate. For example, when noise levels differ two-fold between fast- and slow-growing cells, and the gene expression fluctuation that

is required for a phenotype switch corresponds to 4 standard deviations in fast-growing cells, then the same fluctuation would correspond to 2 standard deviations in slow-growing cells, leading to a $e^{4^2/2}/e^{2^2/2} \approx 400$ fold increase in switching rate in slow-growing cells.

It is currently not clear what mechanisms underlie the decrease of gene expression noise with growth rate. Analysis of genome-wide noise in *E. coli* across different growth conditions has shown that while relative noise levels of different genes are highly condition-dependent and are driven by noise propagation through the regulatory network, absolute noise levels decrease systematically with growth rate in a way that appears to affect all genes (25). This suggests that the overall decrease in gene expression noise with growth rate results from mechanisms that affect all genes. However, this still leaves many possible mechanisms, including fluctuations in chromosome copy numbers across the cell cycle, fluctuations in transcription initiation rates due to variations in the RNA polymerase concentration, fluctuations in transcription elongation rates due to variation in nucleotide concentrations, fluctuations in translation initiation rates due to variation in the ribosome concentration, fluctuations in translation elongation rates due to variation in charged tRNA concentrations, fluctuations in the dilution rate due to variation in growth rate, intrinsic Poissonian fluctuations in all steps of the gene expression process, unequal division of proteins at cell division, and so on. Although a number of models have been proposed that show how some of these sources of noise may explain a decrease in noise with growth rate, e.g., refs. 23, 32, 54 and 55, these models make many simplifying assumptions and consider only some of the mechanisms listed above. As it is currently unknown which mechanisms are most important for determining expression noise levels in realistic settings, it is thus not yet clear which mechanisms drive the observed decrease in noise with growth rate.

We hypothesize that one important contributor to the decrease of noise with growth rate is that the growth rate sets the dilution rate of most intracellular molecules and thus also the rate at which intracellular fluctuations are diluted. Although both the frequency and amplitude of some intracellular noise may naturally increase with growth rate, thereby compensating for the increased dilution rate, this may not apply to all sources of noise, such as fluctuations in extracellular levels of metabolites and stressors. Indeed, in a study of regulatory circuits with positive feedback in *E. coli*, we have recently shown that because signaling molecules are diluted more quickly at higher growth rates, the sensitivity of these regulatory circuits to external signals generally decreases with growth rate (56), supporting that faster dilution may dampen fluctuations in the internal states of cells. Of course, since GRDS generally increases long-term average growth rates, the decrease of fluctuations with growth rate may even be an adaptation that has evolved through natural selection.

Another interesting question is whether GRDS also applies to conditions where instead of different phenotypes growing at different rates, none of the cells are growing, but cells with different phenotypes die at different rates. For example, whether GRDS would predict that cells are more likely to switch to persister states under treatment with antibiotics. In order for GRDS to function in such settings, one would need that fast-dying cells are more likely to change their phenotype than cells in a more slowly dying phenotype. Although this could be the case, we are unaware of any evidence suggesting that slowly dying cells are more phenotypically stable than cells that die more rapidly.

However, since GRDS predicts that cells are more likely to switch their phenotype at a lower growth rate, including to persister phenotypes, GRDS can explain experimental observations that persister fractions increase when cells are stressed before antibiotic treatment (57).

Finally, we have so far implicitly assumed that sensing/regulation and bet-hedging are mutually exclusive strategies, but these strategies can of course act in parallel and may in fact be deeply entangled. By comparing native and synthetic promoters, we have previously shown that natural selection has acted to increase the noise levels of native *E. coli* promoters (58). Moreover, expression noise in *E. coli* results to a substantial extent from noise propagating through the regulatory network so that noise levels are highly condition dependent, with noise in more-regulated promoters being both higher on average and more variable (25). These observations indicate that gene regulation and expression noise are intimately coupled and that the fluctuations in gene expression that we call "noise" at least to some extent result from fluctuations in environmental conditions that propagate through the gene regulatory network. This suggests a strategy in which sensing, regulation, and bet-hedging are all acting in concert, with sensing and regulation being used to constrain the subspace of phenotypes that is explored by stochastic phenotype switching.

Materials and Methods

The general setup of our model has been introduced in the "Model Setup" part of the Results section (*SI Appendix*, section 2.A.1). We have studied this model through several methods: By simplifying it to a toy example and a minimal model, by mathematically studying a limit where environment changes are relatively infrequent and switching rates are low, and by simulating the general model for many random samples of the model parameters. We will here briefly describe each of these methods.

Toy Example. For Fig. 1, we simulated the general model using a Python-script with only $n + 1 = 3$ phenotypes in a sequence of two environments with a duration of $T = 10$. The growth rate of the optimal phenotype was set to $\mu_1 = 1$, while the other phenotypes did not grow: $\mu_0 = 0$; the strength of GRDS was $r = 10$. For bet-hedging cells, we allow for only one global switching rate: $\phi_{ij}^{(K)} = \phi/2$. With GRDS, this switching rate becomes $r\phi/2$ from nongrowing phenotypes. In both cases, ϕ was numerically optimized to maximize the average growth rate.

Minimal Model. All results related to the minimal model were obtained with Mathematica, and an analytical expression was obtained for the average growth rate G as a function of the parameters (n, T, r, ϕ). Optimization of the switching rate ϕ to maximize the average growth rate G was done numerically.

Analytical Derivations. In *SI Appendix*, section 2.B.1, we analytically derive the optimal switching rates for a general model with GRDS; in *SI Appendix*, section 2.B.2, we show that GRDS almost always increases the average population growth rate G, and in *SI Appendix*, section 2.B.3, we approximate the fitness benefit as a function of the strength of the growth rate dependence r. These proofs are possible only if we apply the same approximations as proposed in ref. 19, which entail:

1. The duration of environments is long enough compared to division times and switching times, such that the phenotype distribution in the population has relaxed to a stationary distribution before the next environment switch.
2. The switching rates are small compared to the differences between the fastest growth rate and other growth rates in an environment, such that the stationary phenotype distribution can be well approximated by determining

the dominant eigenvector of the time evolution matrix with perturbation theory.

Numerical Simulations. The numerical calculations for the general model were done in Python. As detailed in *SI Appendix*, section 3, we randomly pick a number of environments m, a number of phenotypes n, an average environment duration τ, a growth rate for each phenotype in each environment $\mu_i^{(K)}$, and parameters b_{IJ} that determine the random sequence of environment types. Then, we randomly choose a switching rate ϕ_{ij} for each pair of phenotypes i, j and implement GRDS of strength r as follows. We start by taking the range $R_\phi(r) = [\log(\phi_{ij}) - 0.5\log(r), \log(\phi_{ij}) + 0.5\log(r)]$. Then, we determine the range of growth rates that cells of phenotype j can achieve in the different environments: $R_\mu = [\min_K\{\mu_j^{(K)}\}, \max_K\{\mu_j^{(K)}\}]$, and we let t^r be the linear map from R_μ to $R_\phi(r)$ that maps the minimal growth rate to the maximal switching rate and vice versa. Now, the switching rate from phenotype j to i in environment I with GRDS of strength r is determined by $\log(\phi_{ij}^{(K)}) = t^r(\mu_j^{(K)})$. Based on the growth rate in the current environment, the switching rates between two phenotypes are thus linearly interpolated in log-scale between an upper bound and a lower bound, where the factor difference between the upper and lower bound is r. In addition, to allow an unbiased comparison between different strengths of GRDS, we rescale these switching rates such that the average switching rate between two phenotypes ($\frac{1}{m}\sum_{K=1}^{m}\phi_{ij}^{(K)}$) is equal to the initially drawn switching rate ϕ_{ij} for all strengths of GRDS. For the results presented in Fig. 3, we did these simulations for 10 different values of r, where $r = 1$ corresponds to the case with no GRDS.

Given this complete set of parameters, we compute the average population growth rate by simulating a sequence of environments with a total duration that exceeded $n^2\tau$ to ensure that we sufficiently sampled the possible switches between different environments.

Data, Materials, and Software Availability. All data and figures presented in this work can be reproduced with the Python and Mathematica scripts deposited in github at https://github.com/dhdegroot/GRDS-code-repository.

Author affiliations: [a]Biozentrum and Swiss Institute of Bioinformatics, University of Basel, Basel 4056, Switzerland; [b]Systems Biology Lab, Amsterdam Institute of Molecular and Life Sciences, Vrije Universiteit, Amsterdam 1081HZ, The Netherlands; and [c]Institute for Atomic and Molecular Physics, Amsterdam 1098XG, The Netherlands

1. C. Opitz, E. Ahrné, K. N. Goldie, A. Schmidt, S. Grzesiek, Deuterium induces a distinctive *Escherichia coli* proteome that correlates with the reduction in growth rate. *J. Biol. Chem.* **294**, 2279–2292 (2019).
2. M. C. Donaldson-Matasci, C. T. Bergstrom, M. Lachmann, The fitness value of information. *Oikos* **119**, 219–230 (2010).
3. O. Rivoire, S. Leibler, The value of information for populations in varying environments. *J. Stat. Phys.* **142**, 1124–1166 (2011).
4. H. C. Berg, E. M. Purcell, Physics of chemoreception. *Biophys. J.* **20**, 193–219 (1977).
5. W. Bialek, S. Setayeshgar, Physical limits to biochemical signaling. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 10040–10045 (2005).
6. H. Dong, L. Nilsson, C. G. Kurland, Gratuitous overexpression of genes in *Escherichia coli* leads to growth inhibition and ribosome destruction. *J. Bacteriol.* **177**, 1497–1504 (1995).
7. M. Scott, C. W. Gunderson, E. M. Mateescu, Z. Zhang, T. Hwa, Interdependence of cell growth and gene expression: Origins and consequences. *Science* **330**, 1099–1102 (2010).
8. C. C. Govern, P. R. ten Wolde, Optimal resource allocation in cellular sensing systems. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 17486 LP–17491 (2014).
9. M. Thattai, A. Van Oudenaarden, Stochastic gene expression in fluctuating environments. *Genetics* **167**, 523–530 (2004).
10. N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, S. Leibler, Bacterial persistence as a phenotypic switch. *Science* **305**, 1622 LP–1625 (2004).
11. M. Kærn, T. C. Elston, W. J. Blake, J. J. Collins, Stochasticity in gene expression: From theories to phenotypes. *Nat. Rev. Genet.* **6**, 451–464 (2005).
12. W. K. Smits, O. P. Kuipers, J. W. Veening, Phenotypic variation in bacteria: The role of feedback regulation. *Nat. Rev. Microbiol.* **4**, 259–271 (2006).
13. O. D. King, J. Masel, The evolution of bet-hedging adaptations to rare scenarios. *Theor. Popul. Biol.* **72**, 560–575 (2007).
14. H. J. Beaumont, J. Gallie, C. Kost, G. C. Ferguson, P. B. Rainey, Experimental evolution of bet hedging. *Nature* **462**, 90–93 (2009).
15. A. M. Simons, Modes of response to environmental change and the elusive empirical evidence for bet hedging. *Proc. R. Soc. B: Biol. Sci.* **278**, 1601–1609 (2011).
16. S. F. Levy, N. Ziv, M. L. Siegal, Bet hedging in yeast by heterogeneous, age-correlated expression of a stress protectant. *PLoS Biol.* **10**, 1001325 (2012).
17. C. R. Lewontin, D. Cohen, On population growth rate in a randomly varying environment. *Proc. Natl. Acad. Sci. U.S.A.* **62**, 1056–1060 (1969).
18. S. Karlin, U. Lieberman, Random temporal variation in selection intensities: Case of large population size. *Theor. Popul. Biol.* **6**, 355–382 (1974).
19. E. Kussell, S. Leibler, Phenotypic diversity, population growth, and information in fluctuating environments. *Science* **309**, 2075–2078 (2005).
20. B. Gaál, J. W. Pitchford, A. J. Wood, Exact results for the evolution of stochastic switching in variable asymmetric environments. *Genetics* **184**, 1113–1119 (2010).
21. M. Acar, J. T. Mettetal, A. Van Oudenaarden, Stochastic switching as a survival strategy in fluctuating environments. *Nat. Genet.* **40**, 471–475 (2008).
22. A. M. New *et al.*, Different levels of catabolite repression optimize growth in stable and variable environments. *PLoS Biol.* **12**, 17–20 (2014).
23. L. Keren *et al.*, Noise in gene expression is coupled to growth rate. *Genome Res.* **25**, 1893–1902 (2015).
24. D. J. Kiviet *et al.*, Stochasticity of metabolism and growth at the single-cell level. *Nature* **514**, 376–379 (2014).
25. A. Urchueguía *et al.*, Genome-wide gene expression noise in *Escherichia coli* is condition-dependent and determined by propagation of noise through the regulatory network. *PLoS Biol.* **19**, e3001491 (2021).
26. S. R. Biggar, G. R. Crabtree, Cell signaling can direct either binary or graded transcriptional responses. *EMBO J.* **20**, 3167–3176 (2001).
27. M. Thattai, B. I. Shraiman, Metabolic switching in the sugar phosphotransferase system of *Escherichia coli*. *Biophys. J.* **85**, 744–754 (2003).
28. T. M. Norman, N. D. Lord, J. Paulsson, R. Losick, Stochastic switching of cell fate in microbes. *Ann. Rev. Microbiol.* **69**, 381–403 (2015).
29. A. Raj, A. van Oudenaarden, Nature, nurture, or chance: Stochastic gene expression and its consequences. *Cell* **135**, 216–226 (2008).
30. A. Eldar, M. B. Elowitz, Functional roles for noise in genetic circuits. *Nature* **467**, 167–173 (2010).
31. A. Kashiwagi, I. Urabe, K. Kaneko, T. Yomo, Adaptive response of a gene network to environmental changes by fitness-induced attractor selection. *PLoS One* **1**, 1–10 (2006).
32. C. Furusawa, K. Kaneko, A generic mechanism for adaptive growth rate regulation. *PLoS Comput. Biol.* **4**, 0035–0042 (2008).
33. H. I. Schreier, Y. Soen, N. Brenner, Exploratory adaptation in large random networks. *Nat. Commun.* **8**, 1–9 (2017).
34. J. L. Kelly, A new interpretation of information rate. *Bell Sys. Tech. J.* **35**, 917–926 (1956).
35. M. Slatkin, Hedging one's evolutionary bets. *Nature* **250**, 704–705 (1974).
36. E. M. Ozbudak, M. Thattai, H. H. Lim, B. I. Shraiman, A. Van Oudenaarden, Multistability in the lactose utilization network of *Escherichia coli*. *Nature* **427**, 737–740 (2004).
37. H. Maamar, D. Dubnau, Bistability in the *Bacillus subtilis* K-state (competence) system requires a positive feedback loop. *Mol. Microbiol.* **56**, 615–624 (2005).
38. D. B. Kearns, R. Losick, Cell population heterogeneity during growth of *Bacillus subtilis*. *Genes Dev.* **19**, 3083–3094 (2005).
39. P. B. Rainey *et al.*, The evolutionary emergence of stochastic phenotype switching in bacteria. *Microb. Cell Fact.* **10**, 1–7 (2011).
40. J. J. Bull, Evolution of phenotypic variance. *Evolution* **41**, 303–315 (1987).
41. P. Haccou, Y. Iwasa, Optimal mixed strategies in stochastic environments. *Theor. Popul. Biol.* **47**, 212–243 (1995).
42. P. V. Martín, M. A. Muñoz, S. Pigolotti, Bet-hedging strategies in expanding populations. *PLoS Comput. Biol.* **15**, e1006529 (2019).
43. L. Hadany, T. Beker, Fitness-associated recombination on rugged adaptive landscapes. *J. Evol. Biol.* **16**, 862–870 (2003).
44. Y. Ram, L. Hadany, The evolution of stress-induced hypermutation in asexual populations. *Evolution* **66**, 2315–2328 (2012).
45. M. Lukačišinová, S. Novak, T. Paixão, Stress-induced mutagenesis: Stress diversity facilitates the persistence of mutator genes. *PLoS Comput. Biol.* **13**, e1005609 (2017).
46. L. Robert *et al.*, Mutation dynamics and fitness effects followed in single cells. *Science* **359**, 1283–1286 (2018).
47. S. H. Kopf *et al.*, Heavy water and 15N labelling with NanoSIMS analysis reveals growth rate-dependent metabolic heterogeneity in chemostats. *Environ. Microbiol.* **17**, 2542–2556 (2015).
48. F. Schreiber *et al.*, Phenotypic heterogeneity driven by nutrient limitation promotes growth in fluctuating environments. *Nat. Microbiol.* **1**, 1–7 (2016).
49. N. Nikolic *et al.*, Cell-to-cell variation and specialization in sugar metabolism in clonal bacterial populations. *PLoS Genet.* **13**, 1–24 (2017).
50. M. Zimmermann *et al.*, Substrate and electron donor limitation induce phenotypic heterogeneity in different metabolic activities in a green sulphur bacterium. *Environ. Microbiol. Rep.* **10**, 179–183 (2018).
51. F. Schreiber, M. Ackermann, Environmental drivers of metabolic heterogeneity in clonal microbial populations. *Curr. Opin. Biotechnol.* **62**, 202–211 (2020).
52. A. Gasperotti, S. Brameyer, F. Fabiani, K. Jung, Phenotypic heterogeneity of microbial populations under nutrient limitation. *Curr. Opin. Biotechnol.* **62**, 160–167 (2020).

53. T. Julou, L. Zweifel, D. Blank, A. Fiori, E. van Nimwegen, Subpopulations of sensorless bacteria drive fitness in fluctuating environments. *PLoS Biol.* **18**, e3000952 (2020).

54. I. T. Kleijn, L. H. Krah, R. Hermsen, Noise propagation in an integrated model of bacterial gene expression and growth. *PLoS Comput. Biol.* **14**, 1–18 (2018).

55. F. Bertaux, S. Marguerat, V. Shahrezaei, Division rate, cell size and proteome allocation: Impact on gene expression noise and implications for the dynamics of genetic circuits. *R. Soc. Open Sci.* **5**, 172234 (2018).

56. T. Julou, T. Gervais, D. Blank, E. van Nimwegen, Growth rate controls the sensitivity of gene regulatory circuits. bioRxiv **1**, 2022.04.03.486858 (2022).

57. N. Q. Balaban, Persistence: Mechanisms for triggering and enhancing phenotypic variability. *Curr. Opin. Genet. Dev.* **21**, 768–775 (2011).

58. L. Wolf, O. K. Silander, E. van Nimwegen, Expression noise facilitates the evolution of gene regulation. *eLife* **4**, e05856 (2015).