

Optimized bacteria are environmental prediction enginesSarah E. Marzen^{1,*} and James P. Crutchfield^{2,†}¹*Department of Physics, Physics of Living Systems Group, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*²*Complexity Sciences Center and Department of Physics, University of California at Davis, Davis, California 95616, USA*

(Received 8 February 2018; revised manuscript received 17 June 2018; published 16 July 2018)

Experimentalists observe phenotypic variability even in isogenic bacteria populations. We explore the hypothesis that in fluctuating environments this variability is tuned to maximize a bacterium's expected log-growth rate, potentially aided by epigenetic (all inheritable nongenetic) markers that store information about past environments. Crucially, we assume a time delay between sensing and action, so that a past epigenetic marker is used to generate the present phenotypic variability. We show that, in a complex, memoryful environment, the maximal expected log-growth rate is linear in the instantaneous predictive information—the mutual information between a bacterium's epigenetic markers and future environmental states. Hence, under resource constraints, optimal epigenetic markers are causal states—the minimal sufficient statistics for prediction—or lossy approximations thereof. We propose new theoretical investigations into and new experiments on bacteria phenotypic bet-hedging in fluctuating complex environments.

DOI: [10.1103/PhysRevE.98.012408](https://doi.org/10.1103/PhysRevE.98.012408)**I. INTRODUCTION**

Isogenic bacterial populations exhibit phenotypic variability [1–4]. Some variability is unavoidable due to noise in the underlying biological circuits [5]. Such noise is not always detrimental to organism functioning: phenotypic variability can be tuned to maximize population fitness [6,7]. Such optimal phenotypic variability is called *bet hedging* [8,9] and has been implicated in seed germination in annual plants [10,11] and in phenotype switching by bacteriophages [12] and fungi [13–16].

At first blush, it may seem strange that a population of organisms should not simply express the phenotype that grows best in the most probable environment—a deterministic strategy. Imagine, however, that the environment fluctuates somewhat unpredictably (as real environments often do), sometimes reaching a less probable state in which that phenotype does not reproduce. If organisms only express that single phenotype, then eventually the population will go extinct. A population of organisms should, instead, hedge its “bets” about future environmental states, using the unavoidable noise in biological circuits [5] or other mechanisms—e.g., slipped-strand mispairing [2,3]—to express different phenotypes with varying probabilities. Given this, the only question is: how should the population hedge its bets?

The first theoretical analysis of such bet-hedging was provided by Kelly in a classic analysis of gambling; see Refs. [17] and [18, chap. 6]. If one thinks of organisms as money, to draw out the parallel, then gambling and bacterial growth are analogous. Adapting Kelly's setup, only one phenotype can reproduce in any given environmental state. Kelly found in effect that (i) the optimal probability of expressing a

phenotype is the probability of observing the corresponding environmental state and (ii) the maximal expected log-growth rate is linear in the negative entropy of a single environmental state's probability.

Realistically, though, more than one phenotype might reproduce in a particular environment. For example, a bacterial phenotype optimized for growth on a high concentration of lactose can still grow on glucose, albeit with additional energetic expenditure [19]. References [20,21] analyzed bet-hedging in just such a case.

Furthermore, epigenetics provides a mechanism by which organisms can remember the environmental past [22]. This memory acts as *side information* about future environmental states—information that can be used to increase the population's expected log-growth rate [17,18,23].¹ And this suggests in turn that such memory should affect optimal phenotypic variability. In fact, in the context of seed germination, predictive cues about the current environmental state were found to change the optimal germination fraction [25]. We assume that there is a time delay between sensing and action, e.g., as in Ref. [26], so that prior epigenetic memory is used to choose the present phenotype. Hence, one's best guide to the present environmental state is indeed the past environmental state.

Here we solve for optimal phenotypic variability and use this to calculate a population's maximal expected log-growth rate when accounting for (i) nonzero reproduction rates of suboptimal phenotypes, (ii) limited epigenetic memory, and (iii) sensor noise. We find, as one might have expected from Ref. [23], that the instantaneous predictive information—that shared between the organism's present phenotype and future environment states—captures (and not just upper bounds [20,21]) the benefit of epigenetic memory. When combined

*semarzen@mit.edu

†chaos@ucdavis.edu

¹In a different context, this observation about memory was used to improve estimates of the entropy of written English [24].

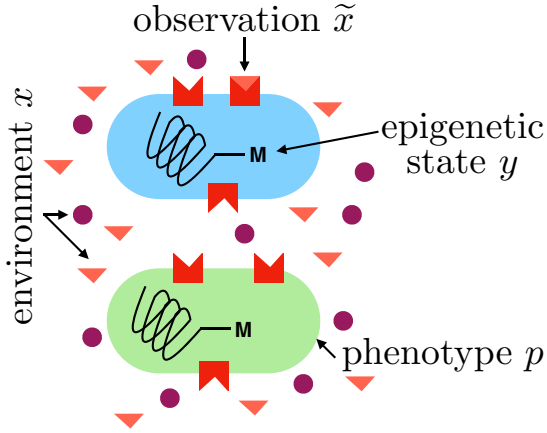


FIG. 1. Population of isogenic bacteria interacting with a fluctuating environment x , given by the concentrations of nutrients (pink triangles and purple circles): Bacteria (i) observe only the concentration \tilde{x} of pink nutrient, (ii) remember aspects of the past environment through their epigenetic state y , which, for instance, could include their genome and any methylations M thereof, and (iii) express a phenotype p (blue or green ovals) that reproduces at different rates depending on the environment. Observations of the environment are assumed to be identical from bacterium to bacterium. When the bacteria experience identical sensations, the epigenetic states of the two bacteria are assumed identical. However, inherent biochemical stochasticity can cause the expression of different phenotypes.

with resource constraints—not considered in Ref. [23]—this predicts that optimal isogenic bacterial populations store epigenetic memories that are causal states or lossy causal states of bacterial observations of the environment, as long as sensors are not too noisy. We conclude with suggestions for testing and extending these results.

II. BACKGROUND

Take the environment to be everything, except the bacterial phenotype, that determines reproductive rates of an individual bacterium. At time t the environment is in a state x_t . What the bacteria observe of the environment at time t is \tilde{x}_t —a noisy subsampling of the full environmental state x_t at time t . For example, the environmental state x_t might consist of a full list of available nutrients, only some of which \tilde{x}_t are sensed by bacteria.

An individual bacterium has a genotype, an *epigenetic state*—all the epigenetic factors such as methylation or the number of proteins that can be inherited above and beyond genetic information—and a phenotype. When we wish to emphasize that the epigenetic state contains information about past environments, we refer to the state as an *epigenetic memory*. We denote the epigenetic state at time t by $y_t \in \mathcal{Y}$, with Y_t its random variable. See Fig. 1.

We assume the environmental time series $x_{-\infty:\infty} = \dots, x_t, x_{t+1}, \dots$ is a realization of a stationary stochastic process. Time increments when the phenotype updates. Bacteria are assumed to stochastically choose a new phenotype every time step based on their previous epigenetic state, where the time delay between perception and action comes from finite biochemical rates. We assume that when a bacterium chooses

its phenotype, it only references its previous epigenetic state and not its previous phenotype.

It is well worth mapping our setup’s assumptions (or, equivalently, those of Ref. [23]) to those previously used to explore the value of information for populations subject to fluctuating environments [27]. We simultaneously relax assumptions A1 (“no information is inherited”) and A3 (“only one phenotype survives”) there, allowing for inheritance only through the epigenetic state, not through the previous phenotype. The last assumption (“epigenetic state selects phenotype”) does not map onto any in Ref. [27]. It constitutes the main insight that allows relaxing both A1 and A3 but still yields closed-form expression for the value of information as the increase in expected log-growth rate arising from storing information about the environment [23].

We, at first, do not allow each bacterium to observe the environment differently; in other words, Ref. [27]’s environmental sensor is the identity map. Later, this, too, is relaxed.

Ultimately, we show that a bacterium should optimally predict its environment (somehow) using the environment’s causal states [28]. Two observed environmental pasts $\tilde{x}_{-\infty:t}$ and $\tilde{x}'_{-\infty:t}$ are considered equivalent, $\tilde{x}_{-\infty:t} \sim_{\epsilon} \tilde{x}'_{-\infty:t}$, if and only if $\Pr(\tilde{X}_t | \tilde{X}_{-\infty:t} = \tilde{x}_{-\infty:t}) = \Pr(\tilde{X}_t | \tilde{X}_{-\infty:t} = \tilde{x}'_{-\infty:t})$. In this, $\tilde{X}_{-\infty:t} = \dots, \tilde{X}_{t-2}, \tilde{X}_{t-1}$ is the chain of random variables representing the observed pasts. The equivalence relation \sim_{ϵ} partitions the set of all pasts into classes called *causal states* $\sigma \in \mathcal{S}$ and induces a rule that maps a past to its causal state: $\sigma = \epsilon(\tilde{x}_{-\infty:t})$. Causal states are the minimal sufficient statistics for predicting the environment, meaning that they constitute the minimal information about the past necessary to predict the future as well as one possibly could given the observations.²

Typically, infinite futures are considered, but in the following, one-step futures are best-suited for the results. The change from infinite futures to one-step futures leads to a slightly modified definition of causal states. The resulting equivalence classes are a coarse-graining of causal states defined as in Ref. [28], but for simplicity we continue to refer to them as causal states.

Let S_t be the random variable corresponding to the causal state at time t . From the probabilities $\Pr(\tilde{X}_t | \tilde{X}_{-\infty:t})$ and the rule $\epsilon(\cdot)$, one obtains a transition dynamic $\Pr(S_{t+1}, \tilde{X}_t | S_t)$ on causal states. The corresponding hidden Markov model is the environment’s minimal, optimal model—its ϵ -machine [28]. It is *unifilar*—that is, given the environment’s current causal state and next observation, its next state is uniquely determined. Of the unifilar hidden Markov models that describe a given environment, the ϵ -machine has the minimal number of states [28,30].

III. RESULTS

First, we find that the instantaneous predictive information defines the quality of an epigenetic state under several assumptions on reproduction rates and environmental statistics. Then, we show that the optimal resource-constrained epigenetic states are the observational causal states. Importantly, this latter

²The reinforcement learning literature has come to call causal states *predictive representations* [29].

result is free from several of the more stringent assumptions required to establish the first result. Finally, we find that when sensors are sufficiently noisy, stochastic switching strategies [7] can beat optimal memoryful bet-hedging strategies.

A. Emergence of instantaneous predictive information

Let n_t be the number of organisms at time t . Let $\Pr(p_t|y_{t-1})$ be a bacterium's *strategy*—the probability that an organism expresses phenotype p_t given epigenetic state y_{t-1} . (Recall that we condition on the previous epigenetic state since there is an effective time delay between perception and action, enforced by finite biochemical kinetic rates.) This conditional probability distribution exists in a *strategy simplex*—the space of valid conditional probability distributions $\Pr(p|y)$. Assume that a bacterium's phenotype at the next time step depends on the previous epigenetic state but is generated independently of its phenotype at the previous time step. Finally, let $f(p_t, x_t)$ be the reproduction rate of phenotype p_t in environment x_t , which might depend on the energetic efficiency of that phenotype in that environment. Let W be the matrix of these reproductive rates.

We measure the fitness of a bacterial population by its expected log-growth rate:

$$r = \left\langle \log \frac{n_{t+1}}{n_t} \right\rangle_{\Pr(x_t, y_{t-1})}. \quad (1)$$

The expected log-growth rate r is a function of epigenetic memories $\Pr(y_t|\tilde{x}_{-\infty:t})$, the phenotypic strategy $\Pr(p_t|y_{t-1})$, and reproductive rates $f(p_t, x_t)$. From Appendix A and Eq. (17) in Ref. [23], one finds a maximal expected log-growth rate:

$$r^* = \max_{\Pr(p_t|y_{t-1})} r = -H[X_t|Y_{t-1}] - \sum_{x_t} \Pr(x_t) \log \sum_{p_t} (W^{-1})_{p_t, x_t} \quad (2)$$

if Eq. (A2) yields an \mathbf{x}_y in the strategy simplex. The first of these two terms ($-H[X_t|Y_{t-1}]$) depends on the scheme that assigns epigenetic states to environmental pasts. The second is independent of such schemes and depends only on environmental statistics and reproduction rates.

Now, recall that the “value of information” in Ref. [27] Δr^* is the increase in maximal expected log-growth rate of a population with epigenetic memory above and beyond that of a population without any epigenetic memory. And so, if Eq. (A2) yields an \mathbf{x}_y in the strategy simplex, then the “value of information” is

$$\begin{aligned} \Delta r^* &= -H[X_t|Y_{t-1}] + H[X_t] \\ &= I[Y_{t-1}; X_t]; \end{aligned} \quad (3)$$

equivalent to Ref. [23]'s Eq. (18), where a cue there is now the previous epigenetic state.

This is the instantaneous predictive information [31]. Hence, epigenetic states with higher instantaneous predictive information are evolutionarily favored.

B. Optimal epigenetic memories are causal states

Now we are ready to analyze optimal epigenetic strategies $\Pr(y_t|\tilde{x}_{-\infty:t})$. Note that y_t has access to information about $\tilde{x}_{-\infty:t}$ but cannot directly access information about x_t . All of y_t 's information about x_t comes through $\tilde{x}_{-\infty:t}$; i.e., $Y_{t-1} \rightarrow \tilde{X}_{-\infty:t} \rightarrow X_{-\infty:t} \rightarrow X_t$. From this, the Data Processing Inequality [18] reveals that

$$\begin{aligned} I[Y_{t-1}; X_t] &\leq I[\tilde{X}_{-\infty:t}; X_t] \\ &\leq I[X_{-\infty:t}; X_t]. \end{aligned} \quad (4)$$

Employing the Data Processing Inequality, we implicitly assume that a bacterium's only guide to the future environment consists of past environmental states. In other words, we assume that an experimentalist, say, does not give the bacterium *additional* side information about the environment. The quantity $I[X_{-\infty:t}; X_t] = H[X_t] - h_\mu$ is also known as the *predicted information rate* or the *total correlation rate* [32,33]. It is largely controlled by the environment's intrinsic randomness or Shannon entropy rate $h_\mu = H[X_t|X_{-\infty:t}]$.

Equation (4) suggests evolution favors populations of organisms which develop epigenetic memories that store as much of the environmental past as possible. However, memory is costly and one should not remember environmental pasts that are not helpful. More specifically, genomes are finite in size and can only support a finite number of epigenetic markers. Hence, the number of possible epigenetic states $|\mathcal{Y}|$ is finite. The balance to strike therefore is to saturate the inequality in Eq. (4) while minimizing a resource cost—the number of possible epigenetic states $|\mathcal{Y}|$. In short, epigenetic memories store the minimal amount of information about the observed environment's past needed to predict the environment's future. They are, therefore, the minimal sufficient statistics of prediction of the future environment with respect to past observations.

How might epigenetic memories store such information? After all, a bacterium cannot directly access the observed environment's past $\dots, \tilde{x}_{t-2}, \tilde{x}_{t-1}$ at time t . However, a bacterium's future epigenetic state y_{t+1} depends on both its previous epigenetic state y_t and the present environmental observation \tilde{x}_{t+1} . In other words, a bacterium's epigenetic state is generated by an input-dependent dynamical system in which input is the environmental observation. If the update rule for how the bacterium's future epigenetic state y_{t+1} depends on the previous epigenetic state y_t and the present environmental observation \tilde{x}_{t+1} are chosen so as to mimic the environment's ϵ -machine transition dynamic, then the bacterium's epigenetic state y_t at time t will be the observed environment's causal state [28]. This is the limit to what is realizable from an input-dependent dynamical system. Hence, a bacterium's optimal *realizable* epigenetic memories and so its epigenetic states are the causal states of the observed environment.

More generally, Eq. (A2) might not give a valid conditional probability distribution or the matrix W there might not be invertible. Even then, maximization of expected log-growth rate combined with resource limitations implies that optimal epigenetic memories are causal states. To show this, we first show that expected log-growth rate is maximized when the epigenetic memories store the entire observed environmental past. Then, we show that this maximum is also achieved when epigenetic memories are minimal sufficient statistics

of prediction of the future environment with respect to past observations. Finally, the aforementioned resource constraints imply that optimal realizable epigenetic memories are causal states.

Let us explain this and so provide a sketch of its proof. As stated, we must first show that expected log-growth rate is maximized when the epigenetic memories store the entire environmental past. To see this, note that any $\Pr(p_t|y_{t-1})$, for any realizable y_{t-1} , can be represented if $y_t = \tilde{x}_{-\infty:t+1}$. Hence,

$$\max_{\Pr(p_t|y_{t-1})} r \leq \max_{\Pr(p_t|\tilde{x}_{-\infty:t+1})} r.$$

Then

$$\max_{\Pr(y_{t-1}|\tilde{x}_{-\infty:t})} \max_{\Pr(p_t|y_{t-1})} r \leq \max_{\Pr(p_t|y_{t-1}); y_{t-1}=\tilde{x}_{-\infty:t}} r.$$

The left-hand side is a maximum over all possible epigenetic states, whereas the right-hand side chooses a particular epigenetic state. And so, the opposite inequality also holds. Therefore, we obtain the equality, as desired:

$$\max_{\Pr(y_{t-1}|\tilde{x}_{-\infty:t})} \max_{\Pr(p_t|y_{t-1})} r = \max_{\Pr(p_t|y_{t-1}); y_{t-1}=\tilde{x}_{-\infty:t}} r.$$

Next, we argue that this maximum is also achieved when epigenetic memories are minimal sufficient statistics of prediction of the future environment with respect to past observations. Note that the expression for r depends only on $\Pr(x_t|y_{t-1})$, averaged over $\Pr(y_{t-1})$. This, in turn, implies that maximal expected log-growth rate can be achieved by any sufficient statistic of prediction. If we prefer sufficient statistics with smaller $|\mathcal{Y}|$, then we find that optimal realizable epigenetic memories are causal states [28], as stated earlier.

C. Detrimental effects of noisy sensors

Natural environments vary widely in terms of structure and predictability. How do evolved bacteria respond? To probe this in a controlled manner, we can design novel experiments in which bacterial observations are described by finite ϵ -machines. There are two reasons for using finite ϵ -machines to probe bacterial behavior. First, when environments are generated by finite ϵ -machines, one can quantitatively predict optimal phenotypic strategies. See Appendix A. Second, even finite ϵ -machines are quite rich, as they can have a wide range of stochasticity and structure. Regardless of the ϵ -machine, results from the previous subsection suggest that bacteria should become predictors of such complex environments.

However, sensors are noisy and this degrades a bacterium's ability to predict and detect regularities, especially in less-predictable or less-structured environments. Bacteria that noisily sense and bet-hedge might be outcompeted by other nonsensing stochastic switchers [7]. Are optimized bacteria still environmental prediction engines?

We study the effects of sensor noise via an illustrative example. Consider an experiment that has been designed so that at any time step, lactose is everywhere at one of two concentrations: 0 and a concentration that we set to be 1 without loss of generality. The experimentalist changes the lactose concentration in time according to the unifilar hidden Markov model shown in Fig. 2 (top).

Bacteria observe the present lactose concentration via transcription of the *lacZ* gene. The *lacZ* gene is more likely to be

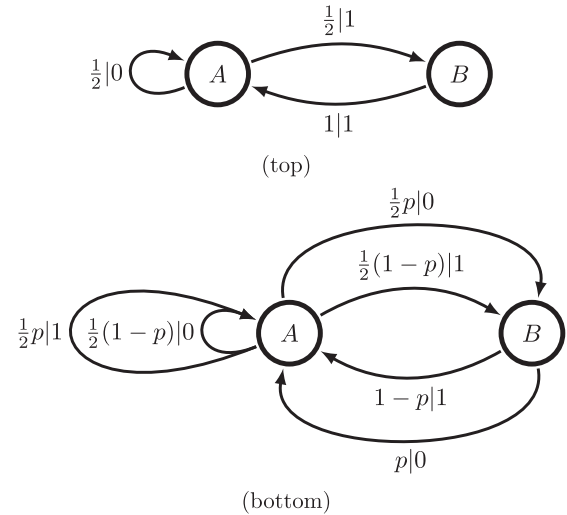


FIG. 2. Hidden Markov models that describe the environment's time evolution x_t (top) and bacterial observations \tilde{x}_t (bottom). Note that Fig. 2 (bottom) is a nonunifilar hidden Markov model; the next state is not uniquely determined by the prior state and the observed symbol because both symbols can be emitted from both states.

transcribed when there is lactose in the environment, as the lactose will bind to a *lac* repressor that inhibits transcription. Such binding is governed by thermal fluctuations. Hence, there is an error probability p that 0 is incorrectly sensed as 1 and that 1 is missed as 0. (In general, these missensing probabilities are not identical, but we assume that they are identical for the sake of simplicity.) This probability p depends on the temperature, on the concentration of lactose when present, and on the free-energy landscape of the sensor and the sensed [34].

Due to this sensing noise, the time series of bacterial observations from the point of view of any particular bacterium is a realization of the process generated by the hidden Markov model in Fig. 2 (bottom). In this model, environmental state transitions can produce either concentrations 0 or 1 with probabilities determined by the error parameter p . Figure 2 (bottom) is built from Fig. 2 (top) by altering state transitions to include the missensing probabilities of the noisy sensor.

What should a bacterium faced with this observed time series do? Ideally, as discussed above, it should estimate whether the environment is in state A or state B in Fig. 2 (bottom) from its observations 0 and 1. When the sensor is noiseless—i.e., in the absence of thermal fluctuations or when $p = 0$ —the bacterium can employ a simple input-dependent dynamical system to correctly estimate whether the environment is in state A or B : that of Fig. 2 (top). In this way, the bacterium reduces the error in its prediction of the environment and so maximizes its fitness.

When bacterial sensors are noisy, a perfect classification of pasts of bacterial observations into state A or B is no longer possible, reducing fitness. To optimize fitness, now bacteria must try to find causal states of a much more complicated stochastic process—that generated by Fig. 2 (bottom)—which has an infinite number of causal states that arise from its (noise-induced) nonunifilarity [35]. (See Appendix B for evidence that the number of causal states is infinite.) Due to resource constraints, though, an optimized bacterium cannot

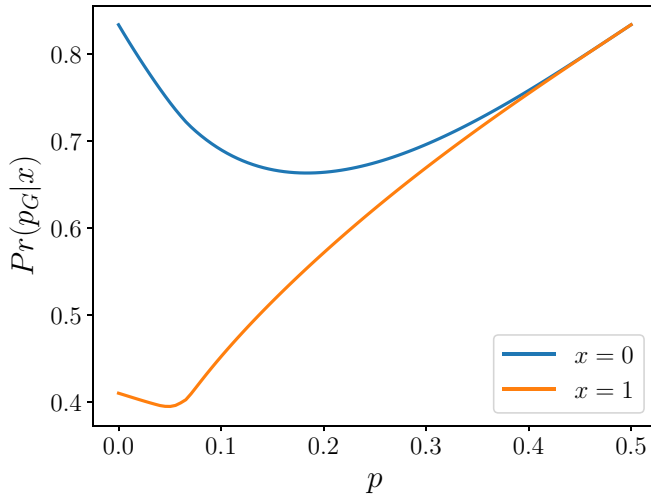


FIG. 3. Optimal phenotypic response depends heavily on sensor noisiness: Horizontal axis shows the sensing error rate p ; the vertical axis shows the probability $\Pr(p_t = p_G|x_t)$ of exhibiting a nonlactose-digesting phenotype given that the environment is in state $x_t \in \{0, 1\}$. When $x_t = 0$, the bacterium optimally expresses the nonlactose-digesting phenotype; when $x_t = 1$, it optimally expresses the lactose-digesting phenotype. As the expressed phenotype can only depend on past bacterial observations, such an optimal response is impossible. And so, as sensor noisiness p increases, the response tracks further from optimal.

internally represent the infinite number of states needed for optimal prediction. Rather, it must now calculate lossy causal states [36]. From Appendix B, the predictive information I_{pred} saturates quickly as a function of the number of causal states $|\mathcal{Y}|$, implying that resource constraints still permit a near-optimal storage of predictive information.

However, our optimized, noisily sensing bacterium perceives different lactose concentrations for the same environmental concentration. This heterogeneity in sensation leads to additional heterogeneity in phenotype expression above and beyond that predicted by the optimal bet-hedging strategy.

To illustrate this, we imagine that there are two possible phenotypes: one p_G that does not have the ability to produce phenotypic machinery to digest lactose and one p_L that does have the ability to produce said machinery. For the associated reproduction rates, suppose that $f(p_G, 0) = 1$, $f(p_G, 1) = 1$, $f(p_L, 0) = 0.75$, and $f(p_L, 1) = 2$. These take into account both that producing the enzymes which digest lactose is expensive and that lactose, when digested, provides the bacterium with energy [19]. If we assume that the epigenetic memory fully stores information about pasts of length 10, then $\Pr(p_G|x)$ depends on the sensor noisiness characterized by error probability p . See Fig. 3. We have assumed, for simplicity, that bacteria use a slightly suboptimal phenotypic strategy: that their phenotypic strategy maximizes expected log growth rate in the case that all bacteria have seen the same bacterial observations. See Appendix B.

Figure 4 shows that the maximal expected log-growth rate r^* decreases with increasing p , implying that growth is negatively affected by sensor noise. Furthermore, from simulations described in Appendix C, simple stochastic switching strategies sans sensing can outcompete optimal phenotypic

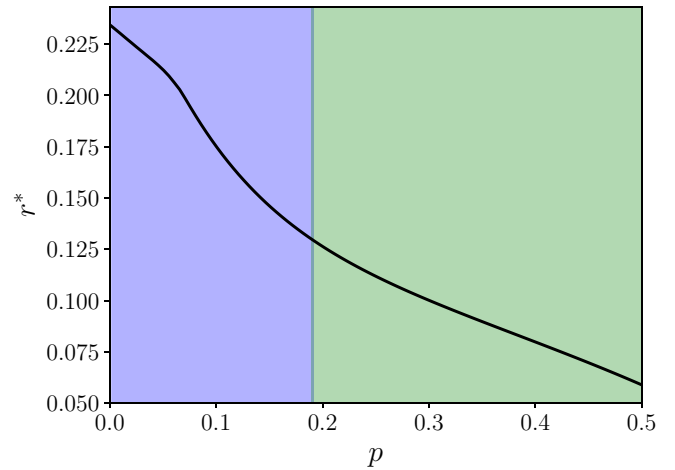


FIG. 4. Maximal expected log-growth rate quickly decreases with increasing sensor noise level: r^* versus noise level p for environments with order-10 Markov memories. From simulations described in Appendix C, simple stochastic switching strategies sans sensing can outcompete optimal phenotypic bet-hedging strategies with noisily sensed epigenetic memory when $r^* < 0.13$ or when $p > 0.19$. The blue region indicates where optimal phenotypic bet-hedging strategies with sensing outcompete stochastic switching strategies sans sensing, and the green region indicates the opposite.

bet-hedging strategies with noisily sensed epigenetic memory when $r^* < 0.13$ or when $p > 0.19$. And these limitations ignore the costs of sensing and memory storage [7]. It is apparently not worthwhile to store expensive epigenetic memory when sensor noise is too large. The critical level of sensor noise depends on the reproductive rates, the temporal statistics of the environment, and the effect of epigenetic memory cost on reproductive rates. For instance, when environments are independent, identically distributed (IID), the stochastic switching strategy will likely outcompete the phenotypic bet-hedging strategy even when there is no sensor noise [26].

IV. CONCLUSIONS

These results lead us to propose that isogenic bacterial populations must predict their environment to maximize their expected log-growth rate, assuming that their sensors are not so noisy that stochastic switching strategies without sensing can outcompete optimal phenotypic bet-hedging strategies with sensing. This conclusion and Eq. (A2) give explicitly testable predictions for new kinds of bacterial-evolution experiments in which populations evolve subject to a fluctuating *memoryful* environment. For instance, one can subject populations to partly random, partly predictable patterns of antibiotics. The prediction is that the bacteria, if sensor noise is small enough, will develop optimal phenotypic bet-hedging behavior in which their probability of exhibiting a particular phenotype implies epigenetic memory, i.e., with phenotypic variability given by Eq. (A2) and with epigenetic memories that correspond to causal states of the environment. (Epigenetic mechanisms might include postreplicative DNA methylation [37], which might occur on timescales as fast as minutes [38].) Although the above analysis focused on bacteria, similar results apply to the phenotype-switching fungi cited earlier.

That said, the setup in Ref. [7] might be more appropriate for interfacing with experiment. As such, we briefly describe an extension of that setup that should yield similar qualitative results to those presented here. Reference [7] studied phenotypic bet-hedging in a continuous-time system and assessed the difference between stochastically switching phenotypes (bet-hedging) and switching to the best phenotype based on sensing. In point of fact, there is a time delay between sensing and action that can be explicitly built into a model of environmental sensing and phenotypic switching, as we have done here. One should then find that *memory* of past environmental states, above and beyond instantaneous sensing of present environmental states, can be used to better select the next phenotype. The environment's inherent stochasticity will also lead such optimally sensing populations to not only utilize memory of past fluctuations but also to stochastically choose phenotypes.

For randomly selected processes, ϵ -machines are usually not finite. Indeed, in our illustrative example, the tiniest bit of sensor noise led to an infinite ϵ -machine. Thus, the resource constraints mentioned earlier become paramount when addressing more naturalistic environments and more realistic setups. It is surprisingly easy to put resource constraints and predictive information on the same footing in this framework based solely on their effect on expected log-growth rate.

Consider Eq. (A3). More stringent constraints on bacterium size will tend to increase reproductive rates $f(p, x)$, since less material is required to generate a new bacterium. Resource constraints therefore will increase the second term in Eq. (A3). However, stronger resource constraints tend to diminish the predictive information captured by the bacterial population, as given by the first term in Eq. (A3). We therefore expect that the input-dependent dynamical system supporting a bacterium's epigenetic states will usually find "lossy causal states" [36] rather than causal states. The degree of trade-off between resource constraints and predictive information will be determined by the environment and the organism's ability to grow in said environment. Lossy causal states can be calculated using the methods of Ref. [36], although here we found suboptimal lossy causal states using the methods of Ref. [35].

The derivation above assumed that the environment was so large that its evolution was independent of bacterial phenotypes. However, bacteria certainly affect their environment, at the very least by secreting molecules and removing nutrients. Ideally, we would not assume that the environment's evolution was independent of the bacteria's actions, thereby closing the sensorimotor loop and allowing for niche construction [39]. We expect relaxing this assumption to yield much more complicated quantifiers of the quality of epigenetic memory, given the difficulty of solution of partially observable Markov decision processes (POMDPs); e.g., as described in Refs. [40–42]. However, we expect causal states to be optimal epigenetic states in any case, since the belief states used in the solution of POMDPs are causal states.

ACKNOWLEDGMENTS

The authors thank J. Horowitz, D. Amor, A. Solon, J. England, C. Ellison, C. Hillar, and S. DeDeo for helpful discussions and the Santa Fe Institute for its hospitality during

visits, where J.P.C. is an External Faculty member. This material is based on work supported by, or in part by, the John Templeton Foundation Grant No. 52095, the Foundational Questions Institute Grant No. FQXi-RFP-1609, and the US Army Research Laboratory and the US Army Research Office under Contract No. W911NF-13-1-0390. S.E.M. was funded by an MIT Physics of Living Systems Fellowship.

APPENDIX A: MAXIMAL EXPECTED LOG-GROWTH RATE

To keep our development self-contained, we derive here the maximal expected log-growth rate and the maximizing distribution over phenotypes. From the setup, we straightforwardly obtain:

$$\begin{aligned} n_{t+1} &= \sum_{p_t} [\Pr(p_t|y_{t-1})n_t]f(p_t, x_t) \\ &= \left[\sum_{p_t} \Pr(p_t|y_{t-1})f(p_t, x_t) \right] n_t, \end{aligned}$$

valid when n_t is sufficiently large. This yields an expected log-growth rate:

$$\begin{aligned} r &= \left\langle \log \frac{n_{t+1}}{n_t} \right\rangle \\ &= \left\langle \log \left[\sum_{p_t} \Pr(p_t|y_{t-1})f(p_t, x_t) \right] \right\rangle \\ &= \sum_{y_{t-1}, x_t} \Pr(y_{t-1}, x_t) \log \left[\sum_{p_t} \Pr(p_t|y_{t-1})f(p_t, x_t) \right]. \quad (\text{A1}) \end{aligned}$$

Note that we focus on the expected log-growth rate as a natural measure of a population's fitness, rather than on an individual's fitness, which might be better measured via expected growth rate. Indeed, the lesson from phenotypic bet-hedging is that what is good for the population is not necessarily good for the individual. To survive, an individual should choose a strategy that survives in all environments, even if it grows slowly in some. However, a population has the luxury of betting some organisms on phenotypes that might not survive in certain environments if they grow much faster in others. Hence, we are interested in what kinds of isogenic bacterial *populations* evolve. However, since these populations are isogenic, we describe the evolved population by describing properties of the individual bacterium.

We seek the bet-hedging strategy $\Pr(p_t|y_{t-1})$ that maximizes expected log-growth rate r . Our derivation closely follows that of Ref. [21], with the key change that we now allow for side-information from epigenetic memory. We maximize r , subject to the constraint that $\sum_{p_t} \Pr(p_t|y_{t-1}) = 1$ for all y_t , via the Lagrangian:

$$\begin{aligned} \mathcal{L} &= \sum_{y_{t-1}, x_t} \Pr(y_{t-1}, x_t) \log \left[\sum_{p_t} \Pr(p_t|y_{t-1})f(p_t, x_t) \right] \\ &\quad + \sum_{y_{t-1}} \lambda_{y_{t-1}} \sum_{p_t} \Pr(p_t|y_{t-1}), \end{aligned}$$

with respect to $\Pr(p_t|y_{t-1})$, where $\lambda_{y_{t-1}}$ is the Lagrange multiplier for each epigenetic state y_{t-1} . Note that if the bacterial

population strongly affected the environment's dynamics, then $\Pr(y_{t-1}, x_t)$ would depend on $\Pr(p_t|y_{t-1})$. Instead, we assume the environment is so large that the bacteria population does not affect it.

To find the strategy $\Pr(p_t|y_{t-1})$ that maximizes r , we take derivatives of the Lagrangian and set them to 0:

$$\begin{aligned} 0 &= \frac{\partial \mathcal{L}}{\partial \Pr(p_t|y_{t-1})} \\ &= \sum_{x_t} \Pr(x_t|y_{t-1}) \frac{f(p_t, x_t)}{\sum_{p_t} \Pr(p_t|y_{t-1}) f(p_t, x_t)} - \lambda_{y_{t-1}}. \end{aligned}$$

And so:

$$\lambda_{y_{t-1}} = \sum_{x_t} \Pr(x_t|y_{t-1}) \frac{f(p_t, x_t)}{\sum_{p_t} \Pr(p_t|y_{t-1}) f(p_t, x_t)}.$$

Let \mathbf{x}_y be the vector of optimal strategies $\Pr(p_t|y_{t-1})$, \mathbf{p}_y the vector of $\Pr(x_t|y_{t-1})$, and W the matrix with elements $f(p_t, x_t)$. Then, the preceding result in matrix form is as follows:

$$\lambda_y \mathbf{1} = W(\mathbf{p}_y \odot [W^\top \mathbf{x}_y]^{\odot -1}),$$

where the $\mathbf{1}$ s vector $\mathbf{1}$ has the length of the number of possible phenotypes and \odot is the Hadamard product, so that \odot represents component-wise multiplication and $[W^\top \mathbf{x}_y]^{\odot -1}$ represents component-wise inversion. If W is invertible, then we solve for \mathbf{x}_y :

$$\mathbf{x}_y = \frac{1}{\lambda_y} (W^\top)^{-1} (\mathbf{p}_y \odot [W^{-1} \mathbf{1}]^{\odot -1}),$$

and, using the normalization condition $\mathbf{1}^\top \mathbf{x}_y = 1$, we fortuitously find that:

$$\mathbf{x}_y = (W^\top)^{-1} (\mathbf{p}_y \odot [W^{-1} \mathbf{1}]^{\odot -1}). \quad (\text{A2})$$

Note that this is the maximizing conditional distribution *if* it is in the strategy simplex and *if* W is invertible. One might relax the condition that W is invertible, if W is square, via the Drazin inverse. In sum, Eq. (A2) determines the optimal strategy for phenotypic variability given a particular epigenetic memory. It is comparable to a manipulation of Eq. (16) in Ref. [23].

The expected log-growth rate r is, from Eq. (A1), a function of epigenetic memories $\Pr(y_t|x_{-\infty:t+1})$, the phenotypic strategy $\Pr(p_t|y_{t-1})$, and reproductive rates $f(p_t|x_t)$. Given the optimal strategy \mathbf{x}_y from Eq. (A2), one finds a maximal expected log-growth rate:

$$\begin{aligned} r^* &= \sum_{y_{t-1}, x_t} \Pr(y_{t-1}, x_t) \log \frac{\Pr(x_t|y_{t-1})}{\sum_{p_t} (W^{-1})_{p_t, x_t}} \\ &= -H[X_t|Y_{t-1}] - \sum_{x_t} \Pr(x_t) \log \sum_{p_t} (W^{-1})_{p_t, x_t}. \quad (\text{A3}) \end{aligned}$$

Equation (A3) is comparable to Eq. (17) in Ref. [23], in which $y(e|c)$ is set to $p(e|c)$ and in which $d(e)$ is calculated in terms of W .

APPENDIX B: SENSOR NOISE

We first calculate $I_{\text{pred}} = I[Y_T; \tilde{X}_T]$. Here, we take Y_t to be a uniform coarse-graining of the mixed state simplex [35,43]. We cannot find the recurrent mixed states, but instead calculate

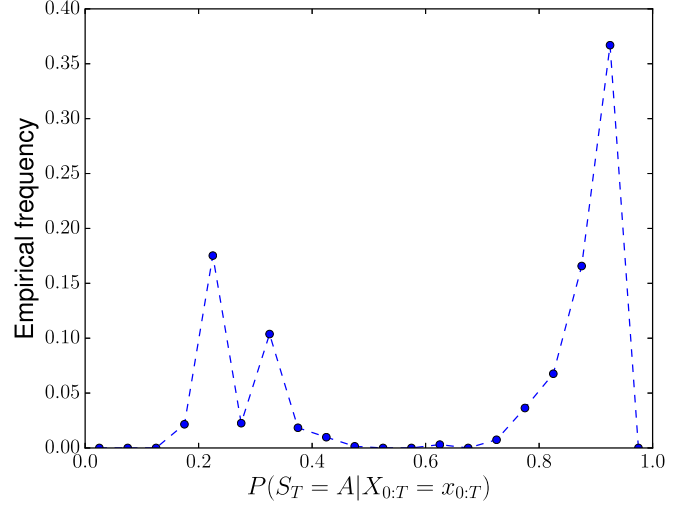


FIG. 5. For the noisy Even Process of Fig. 2 (bottom), mixed states are in a one-dimensional simplex. A mixed state is completely characterized by the mixed state distribution $\Pr(S_T = A | \tilde{X}_{0:T} = \tilde{x}_{0:T})$, where S_T is random variable for the process' two internal states A and B . This distribution is shown as an empirical histogram over the unit interval, here shown with $T = 12$ and 20 bins along the unit interval. This mixed state presentation suggests an infinite number of mixed states, since the empirical frequency distribution does not look to be a sum of delta functions.

transient mixed states for trajectories of length T . Since there are only two states in the underlying hidden Markov model, the mixed-state simplex is one dimensional and characterized completely by the probability that one is in state A given a trajectory $\tilde{x}_{0:T}$. Our calculations require only two labeled-transition matrices:

$$T^{(0)} = \begin{bmatrix} \frac{1}{2}(1-p) & p \\ \frac{1}{2}p & 0 \end{bmatrix} \quad \text{and} \quad T^{(1)} = \begin{bmatrix} \frac{1}{2}p & 1-p \\ \frac{1}{2}(1-p) & 0 \end{bmatrix}.$$

The equilibrium probability μ over states A and B is

$$\begin{aligned} \mu &= \text{eig}_1(T^{(0)} + T^{(1)}) \\ &= \begin{pmatrix} \frac{2}{3} \\ \frac{1}{3} \end{pmatrix}, \end{aligned}$$

where $\text{eig}_1(M)$ is the normalized (i.e., entries sum to 1) eigenvector of eigenvalue 1 of matrix M . The probability of observing any particular trajectory is

$$p(\tilde{x}_{0:T}) = \mathbf{1}^\top T^{(\tilde{x}_{T-1})} \dots T^{(\tilde{x}_0)} \mu$$

and the corresponding mixed state has

$$\begin{bmatrix} \Pr(S_T = A | \tilde{X}_{0:T} = \tilde{x}_{0:T}) \\ \Pr(S_T = B | \tilde{X}_{0:T} = \tilde{x}_{0:T}) \end{bmatrix} = \frac{1}{p(\tilde{x}_{0:T})} T^{(\tilde{x}_{T-1})} \dots T^{(\tilde{x}_0)} \mu.$$

The empirical distribution over mixed states is shown in Fig. 5 for $p = 0.1$.

To coarse-grain, the mixed-state simplex—completely characterized by $\Pr(S_T = A | \tilde{X}_{0:T} = \tilde{x}_{0:T})$ —is divided into $|\mathcal{Y}|$ equal-sized bins, corresponding then to representation y_T . Let $\delta_{|\mathcal{Y}|} : \tilde{x}_{0:T} \rightarrow y$ describe the map:

$$\delta_{|\mathcal{Y}|}(\tilde{x}_{0:T}) = [|\mathcal{Y}| \Pr(S_T = A | \tilde{X}_{0:T} = \tilde{x}_{0:T})].$$

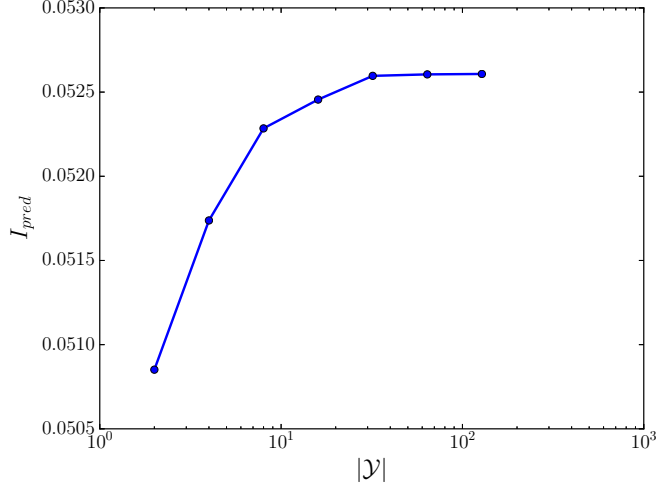


FIG. 6. Predictive information I_{pred} quickly saturates as a function of the number $|\mathcal{Y}|$ of epigenetic states: At $p = 0.1$, with Y_t being a uniform coarse-graining of 2^{15} mixed states ($T = 15$) with $|\mathcal{Y}|$ coarse-grained states. This curve is representative of various p and is typical for such curves [36,44,45].

We would like to calculate $I_{\text{pred}} = I[Y_T; \tilde{X}_T]$, which requires calculating $\Pr(\tilde{X}_T | Y_T)$. We calculate the latter from:

$$\Pr(\tilde{X}_T | Y_T) = \sum_{\sigma} \Pr(\tilde{X}_T | S_T = \sigma) \Pr(S_T = \sigma | Y_T),$$

where

$$\begin{aligned} \Pr(S_T = \sigma | Y_T) &= \sum_{\tilde{x}_{0:T}} \Pr(S_T = \sigma | \tilde{X}_{0:T} = \tilde{x}_{0:T}) \\ &\quad \times \Pr(\tilde{X}_{0:T} = \tilde{x}_{0:T} | Y_T) \\ \Pr(S_T = \sigma | Y_T = y) &= \sum_{\tilde{x}_{0:T}: \delta_{|\mathcal{Y}|}(\tilde{x}_{0:T}) = y} \Pr(S_T = \sigma | \tilde{X}_{0:T} = \tilde{x}_{0:T}) \\ &\quad \times p(\tilde{x}_{0:T}) / \sum_{\tilde{x}_{0:T}: \delta_{|\mathcal{Y}|}(\tilde{x}_{0:T}) = y} p(\tilde{x}_{0:T}) \end{aligned}$$

and where $\Pr(\tilde{X}_T | S_T = \sigma)$ can be calculated from $1^\top T_{\sigma}^{(\tilde{x}_T)}$, so that

$$\Pr(\tilde{X}_T = 0 | S_T = A) = \frac{1}{2} \quad \text{and} \quad \Pr(\tilde{X}_T = 0 | S_T = B) = p,$$

and so on. Figure 6 shows the predictive information I_{pred} so calculated as a function of $|\mathcal{Y}|$ for the particular value $p = 0.1$.

Finally, we would like to compute both optimal phenotypic heterogeneity $\Pr(p | x_t)$ and maximal expected log growth rate r^* as a function of sensor noise p . In the most general case, this turns out to be beyond the present scope, for reasons detailed below. If the bacterial population remains isogenic, however, then we can calculate phenotypic heterogeneity from consideration of the optimal phenotypic strategy given a particular environmental past. Of course, all the bacteria “sees”—all that influences its phenotypic strategy—is the observed environmental past. In other words, we have:

$$\Pr(p | x_{-\infty:t}) = \sum_{\tilde{x}_{-\infty:t}} \Pr(p | \tilde{x}_{-\infty:t}) \Pr(\tilde{x}_{-\infty:t} | x_{-\infty:t}).$$

A tricky part is to correctly calculate $\Pr(\tilde{x}_{-\infty:t} | x_{-\infty:t})$. Naively, it may seem that

$$\Pr(\tilde{x}_{t-L:t} | x_{t-L:t}) = p^{H(\tilde{x}_{t-L:t}, x_{t-L:t})} (1-p)^{L-H(\tilde{x}_{t-L:t}, x_{t-L:t})},$$

where $H(a, b)$ is the Hamming distance based on how we have defined the sensor noise, but this is incorrect. The probability of a cell experiencing a certain observation $\tilde{x}_{t-L:t}$ is affected by the phenotypic strategy pursued by the cell, as that changes the fraction of cells having observed $\tilde{x}_{t-L:t}$, which then changes the fraction of cells observing a particular $\tilde{x}_{t-L:t+1}$. In reality, supposing that there are $n_t(\tilde{x}_{t-L:t})$ cells that experience observations $\tilde{x}_{t-L:t}$ at time t , then

$$\begin{aligned} n_{t+1}(\tilde{x}_{t+1-L:t+1}) &= \sum_{\tilde{x}'_{t-L:t}} n_t(\tilde{x}'_{t-L:t}) \langle f(p, x_t) \rangle_{\Pr(p | \tilde{x}'_{t-L:t})} \\ &\quad \times \Pr(\tilde{x}_{t+1-L:t+1} | \tilde{x}'_{t-L:t}), \end{aligned} \quad (\text{B1})$$

where

$$\Pr(\tilde{x}_{t+1-L:t+1} | \tilde{x}'_{t-L:t}) = \Pr(\tilde{x}_{-\infty:t} | x_t) \delta_{\tilde{x}_{t+1-L:t}, \tilde{x}'_{t+1-L:t}}, \quad (\text{B2})$$

where $\Pr(\tilde{x}_{-\infty:t} | x_t)$ is $1-p$ if they are identical and p otherwise in this particular example. This implies that the expected log growth rate is

$$r = \left\langle \frac{\sum_{\tilde{x}_{t+1-L:t+1}} n_{t+1}(\tilde{x}_{t+1-L:t+1})}{\sum_{\tilde{x}_{t-L:t}} n_t(\tilde{x}_{t-L:t})} \right\rangle_t. \quad (\text{B3})$$

Given a phenotypic strategy $\Pr(p | \tilde{x}_{t-L:t})$, we can calculate the expected log growth rate given the evolution equation above.

Note that Eq. (B3) gives a different equation for expected log growth rate than previously given in Appendix A. There is no longer a closed-form solution for the optimal phenotypic strategy that maximizes expected log growth rate. To make progress, we assume a suboptimal phenotypic strategy: We assume that bacteria experiencing observations $\tilde{x}_{-\infty:t}$ choose a phenotypic strategy that would have maximized expected log growth rate had the entire population observed $\tilde{x}_{-\infty:t}$. This should work well when $\tilde{x}_{-\infty:t}$ is “typical,” which happens with high probability. This will lead to an overestimate of the detrimental effects of sensor noise. In other words, to find an approximate $\Pr(p | \tilde{x}_{-\infty:t})$, we utilize the results of Appendix A and find that

$$\Pr(p | \tilde{x}_{-\infty:t}) = (\Pr(x_t | \tilde{x}_{-\infty:t}) \odot [W^{-1} \mathbf{1}]^{\odot -1}) / W^\top,$$

for

$$W = \begin{pmatrix} 0 & 1 \\ \frac{1}{4} & 2 \end{pmatrix}.$$

When the so-calculated $\Pr(p | \tilde{x}_{-\infty:t})$ is not within the simplex, it is on the simplex faces or edges. To deal with this likely scenario, we simply calculate the expected log growth rate conditioned on observing the past $\tilde{x}_{-\infty:t}$, $\sum_{x_t} \Pr(x_t | \tilde{x}_{-\infty:t}) \log [\sum_p \Pr(p | \tilde{x}_{-\infty:t}) f(p, x_t)]$ and choose the phenotypic strategy at the edge of the simplex that maximizes this conditional expected log growth rate. For the two-phenotype problem considered here, we compare $\sum_{x_t} \Pr(x_t | \tilde{x}_{-\infty:t}) \log [f(p_G, x_t)]$ to $\sum_{x_t} \Pr(x_t | \tilde{x}_{-\infty:t}) \log [f(p_L, x_t)]$ and choose the strategy $\Pr(p | \tilde{x}_{-\infty:t}) = \delta_{p, p_G}$ should the former outweigh the latter and $\Pr(p | \tilde{x}_{-\infty:t}) = \delta_{p, p_L}$ should the latter outweigh the former. In

practice, we replace the infinite pasts $x_{-\infty:t}$ and $\tilde{x}_{-\infty:t}$ with finite-length L pasts, and make L sufficiently large that the answer does not change on increases in L . Finite L is not only computationally necessary, but biologically reasonable given our requirement on resource constraints. We have chosen $L = 10$.

APPENDIX C: STOCHASTIC SWITCHING STRATEGIES

In a stochastic switching strategy, nothing is sensed, but bacteria randomly switch from one phenotype to another. In the simple illustrative example studied in the main text, if n_t^G is the number of glucose-favoring phenotypes at time t and n_t^L is the number of lactose-favoring phenotypes at time t , then

$$\begin{pmatrix} n_{t+1}^G \\ n_{t+1}^L \end{pmatrix} = \begin{pmatrix} 1-a & b \\ a & 1-b \end{pmatrix} \begin{bmatrix} f(p_G, x_t) & 0 \\ 0 & f(p_L, x_t) \end{bmatrix} \begin{pmatrix} n_t^G \\ n_t^L \end{pmatrix}.$$

In this model, bacteria first grow each time step and then stochastically switch. Environments are generated from the hidden Markov model in Fig. 2 (top) for T time steps, and the expected log-growth rate is approximated as

$$\hat{r} = \frac{1}{T} \sum_{t=1}^T \log \frac{n_{t+1}^G + n_{t+1}^L}{n_t^G + n_t^L}.$$

Larger T leads to more accurate estimates of the true expected log-growth rate. Parameter sweeps over a and b were

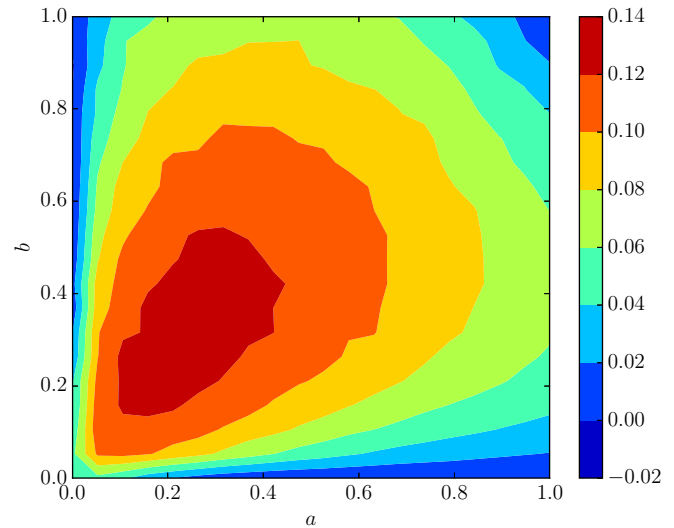


FIG. 7. Estimated expected log-growth rate \hat{r} with $T = 10^5$ as a function of stochastic switching parameters a and b . At maximum, stochastic switching gives an expected log-growth rate of $\hat{r}^* \approx 0.13$.

performed, resulting in the contour plot of expected log-growth rate \hat{r} in Fig. 7. The maximal expected log-growth rate from a stochastic switching strategy is $\hat{r}^* \approx 0.13$.

- [1] H. R. Bonifield and K. T. Hughes, Flagellar phase variation in *Salmonella enterica* is mediated by a posttranscriptional control mechanism, *J. Bacteriol.* **185**, 3567 (2003).
- [2] E. R. Moxon, P. B. Rainey, M. A. Nowak, and R. E. Lenski, Adaptive evolution of highly mutable loci in pathogenic bacteria, *Curr. Biol.* **4**, 24 (1994).
- [3] C. D. Bayliss, D. Field, and E. R. Moxon, The simple sequence contingency loci of *Haemophilus influenzae* and *Neisseria meningitidis*, *J. Clin. Invest.* **107**, 657 (2001).
- [4] D. B. Kearns, F. Chu, R. Rudner, and R. Losick, Genes governing swarming in *Bacillus subtilis* and evidence for a phase variation mechanism controlling surface motility, *Molec. Microbiol.* **52**, 357 (2004).
- [5] M. Kærn, T. C Elston, W. J. Blake, and J. J. Collins, Stochasticity in gene expression: From theories to phenotypes, *Nat. Rev. Gen.* **6**, 451 (2005).
- [6] H. J. E. Beaumont, J. Gallie, C. Kost, G. C. Ferguson, and P. B. Rainey, Experimental evolution of bet hedging, *Nature* **462**, 90 (2009).
- [7] E. Kussell and S. Leibler, Phenotypic diversity, population growth, and information in fluctuating environments, *Science* **309**, 2075 (2005).
- [8] J. Seger and H. J. Brockmann, What is bet-hedging? *Oxf. Surv. Evol. Biol.* **4**, 182 (1987).
- [9] I. G. de Jong, P. Haccou, and O. P. Kuipers, Bet hedging or not? A guide to proper classification of microbial survival strategies, *Bioessays* **33**, 215 (2011).
- [10] D. Cohen, Optimizing reproduction in a randomly varying environment, *J. Theor. Biol.* **12**, 119 (1966).
- [11] M. G. Bulmer, Delayed germination of seeds: Cohen's model revisited, *Theor. Pop. Biol.* **26**, 367 (1984).
- [12] S. Maslov and K. Sneppen, Well-temperate phage: optimal bet-hedging against local environmental collapses, *Sci. Rep.* **5**, 10523 (2015).
- [13] D. R. Soll and B. Kraft, A comparison of high frequency switching in the yeast *Candida albicans* and the slime mold *Dictyostelium discoideum*, *Genes* **9**, 615 (1988).
- [14] N. Jain, A. Guerrero, and B. C. Fries, Phenotypic switching and its implications for the pathogenesis of *Cryptococcus neoformans*, *FEMS Yeast Res.* **6**, 480 (2006).
- [15] A. Guerrero, N. Jain, D. L. Goldman, and B. C. Fries, Phenotypic switching in *Cryptococcus neoformans*, *Microbiology* **152**, 3 (2006).
- [16] J. Pérez-Martín, J. A. Uría, and A. D. Johnson, Phenotypic switching in *Candida albicans* is controlled by a SIR2 gene, *EMBO J.* **18**, 2580 (1999).
- [17] J. L. Kelly, A new interpretation of information rate, in *The Kelly Capital Growth Investment Criterion: Theory and Practice*, (World Scientific, Singapore, 2011), pp. 25–34.
- [18] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley-Interscience, New York, 2006), 2nd ed.
- [19] E. Dekel and U. Alon, Optimality and evolutionary tuning of the expression level of a protein, *Nature* **436**, 588 (2005).
- [20] A. R. Barron and T. M. Cover, A bound on the financial value of information, *IEEE Trans. Info. Th.* **34**, 1097 (1988).
- [21] C. T. Bergstrom and M. Lachmann, Shannon information and biological fitness, in *Information Theory Workshop, 2004* (IEEE, Los Alamitos, CA, 2004), pp. 50–54.

- [22] S. Henikoff and J. M. Gready, Epigenetics, cellular memory and gene regulation, *Curr. Biol.* **26**, R644 (2016).
- [23] M. C. Donaldson-Matasci, C. T. Bergstrom, and M. Lachmann, The fitness value of information, *Oikos* **119**, 219 (2010).
- [24] T. Cover and R. King, A convergent gambling estimate of the entropy of English, *IEEE Trans. Info. Th.* **24**, 413 (1978).
- [25] D. L. Venable and L. Lawlor, Delayed germination and dispersal in desert annuals: Escape in space and time, *Oecologia* **46**, 272 (1980).
- [26] E. Jablonka, B. Oborny, I. Molnar, E. Kisdi, J. Hofbauer, and T. Czaran, The adaptive advantage of phenotypic memory in changing environments, *Phil. Trans. R. Soc. Lond.* **350**, 133 (1995).
- [27] O. Rivoire and S. Leibler, The value of information for populations in varying environments, *J Stat. Phys.* **142**, 1124 (2011).
- [28] C. R. Shalizi and J. P. Crutchfield, Computational mechanics: Pattern and prediction, structure and simplicity, *J. Stat. Phys.* **104**, 817 (2001).
- [29] M. L. Littman, R. S. Sutton, and S. P. Singh, Predictive representations of state, In *NIPS* **14**, 1555 (2001).
- [30] N. Travers and J. P. Crutchfield, Equivalence of history and generator ϵ -machines, [arXiv:1111.4500](https://arxiv.org/abs/1111.4500).
- [31] S. Still, D. A. Sivak, A. J. Bell, and G. E. Crooks, Thermodynamics of Prediction, *Phys. Rev. Lett.* **109**, 120604 (2012).
- [32] R. G. James, C. J. Ellison, and J. P. Crutchfield, Anatomy of a bit: Information in a time series observation, *CHAOS* **21**, 037109 (2011).
- [33] R. G. James, K. Burke, and J. P. Crutchfield, Chaos forgets and remembers: Measuring information creation, destruction, and storage, *Phys. Lett. A* **378**, 2124 (2014).
- [34] R. Phillips, J. Kondev, J. Theriot, and H. Garcia, *Physical Biology of the Cell* (Garland Science, New York, 2012).
- [35] S. E. Marzen and J. P. Crutchfield, Nearly maximally predictive features and their dimensions, *Phys. Rev. E* **95**, 051301(R) (2017).
- [36] S. E. Marzen and J. P. Crutchfield, Predictive rate-distortion for infinite-order Markov processes, *J. Stat. Phys.* **163**, 1312 (2016).
- [37] J. Casadesús and D. Low, Epigenetic gene regulation in the bacterial world, *Microbiol. Molec. Bio. Rev.* **70**, 830 (2006).
- [38] J. W. Kappler, The kinetics of dna methylation in cultures of a mouse adrenal cell line, *J. Cell. Physiol.* **75**, 21 (1970).
- [39] F. J. Odling-Smee, K. N. Laland, and M. W. Feldman, *Niche Construction: The Neglected Process in Evolution* (Princeton University Press, Princeton, NJ, 2003).
- [40] N. Meuleau, K.-E. Kim, L. P. Kaelbling, and A. R. Cassandra, Solving pomdps by searching the space of finite policies, in *Proceedings of the 15th Conference on Uncertainty Artificial Intelligence* (Morgan Kaufmann, New York, 1999), pp. 417–426.
- [41] F. Doshi-Velez, D. Pfau, F. Wood, and N. Roy, Bayesian nonparametric methods for partially-observable reinforcement learning, *IEEE Trans. Patt. Anal. Mach. Intel.* **37**, 394 (2015).
- [42] M. Hausknecht and P. Stone, Deep recurrent q-learning for partially observable MDPs, *CoRR*, [arXiv:1507.06527](https://arxiv.org/abs/1507.06527).
- [43] D. Blackwell, *The Entropy of Functions of Finite-state Markov Chains*, Vol. 28 (Publishing House of the Czechoslovak Academy of Sciences, Prague, 1957).
- [44] S. Still and J. P. Crutchfield, Structure or noise? 2007. Santa Fe Institute Working Paper 2007-08-020; [arXiv:0708.0654](https://arxiv.org/abs/0708.0654).
- [45] S. Still, J. P. Crutchfield, and C. J. Ellison, Optimal causal inference: Estimating stored information and approximating causal architecture, *CHAOS* **20**, 037111 (2010).