

Probabilistic adaptation in changing microbial environments

Yarden Katz^{a,b}, Michael Springer^a

^a*Dept. of Systems Biology, Harvard Medical School, Boston, MA*

^b*Berkman Klein Center for Internet & Society, Harvard University, Cambridge, MA*

22 July 2016

Abstract

Microbes growing in animal host environments face fluctuations that have elements of both randomness and predictability. In the mammalian gut, fluctuations in nutrient levels and other physiological parameters are structured by the animal host's behavior, diet, health and microbiota composition. Microbial cells that are able to anticipate these fluctuations by exploiting this structure would likely gain a fitness advantage, by adapting their internal state in advance. We propose that the problem of adaptive growth in these structured changing environments can be viewed as probabilistic inference. We analyze environments that are “meta-changing”: where there are changes in the way the environment fluctuates, governed by a mechanism unobservable to cells. We develop a dynamic Bayesian model of these environments and show that a real-time inference algorithm (particle filtering) for this model can be used as a microbial growth strategy implementable in molecular circuits. The growth strategy suggested by our model outperforms heuristic strategies, and points to a class of algorithms that could support real-time probabilistic inference in natural or synthetic cellular circuits.

Key words: Epigenetic adaptation, microbiology, Bayesian inference, systems biology, gut microbiota, cellular circuits

Introduction

Outside the laboratory, microbes are faced with rich and changing environments. To improve their chances of survival, single microbial cells must adapt to fluctuations in nutrients and other environmental conditions. The mammalian gut, home to prokaryotic and eukaryotic microbes [1,2], is a striking example of a changing environment with elements of both randomness and order. Nutrients and metabolites may fluctuate stochastically in the gut, but these changes are structured by the host's physiology, diet and behavior [3,4,5]. Cells that exploit this noisy structure and anticipate changes in their environment would likely gain a fitness advantage.

It remains unclear what the information processing capabilities of microbial populations are in such changing environments. To what extent are cells able to learn from their environment's history and use this information to predict future changes? How sophisticated are the resulting computations, and in what environments will they lead to increased fitness? Insight into these questions would shed light on the type of environments cells were selected for and may guide the search for molecular mechanisms that implement adap-

tive computation. This direction could also inform how microbes become pathogenic. The yeast *C. albicans*, for example, can turn from a harmless human gut commensal to a pathogen depending on the host environment and its nutrient composition [6,7,8]. A better understanding of how microbes like *C. albicans* perceive and adapt to their environment may suggest ways of manipulating the environment to control pathogenic growth.

Progress on these questions requires analysis at multiple levels of abstraction, as outlined by David Marr [9] for information-processing in the nervous system. First, the computational task solved by cells has to be specified. For microbial adaptation, this would mean characterizing the space of possible changing environments and identifying the cellular strategies that would result in optimal growth in each environment. Second, the algorithms and representations that cells need to execute the growth strategy would have to be described. Finally, at the implementation level, we have to give an account of how molecular interactions give rise to the algorithm and the necessary representations. A complete account of microbial adaptation would ultimately integrate the three levels.

There has been much work on understanding the molecular and genetic determinants of microbial growth in changing environments (e.g., using experimental evolution [10,11]), but less on defining the abstract computational problem that microbes face when adapting to such envi-

Email addresses: yarden@hms.harvard.edu (Yarden Katz), michael.springer@hms.harvard.edu (Michael Springer).

ronments. In this work, we focus on the computational and algorithmic aspects of adaptive growth in changing environments. We computationally characterize a set of structured dynamic environments, where fluctuations are driven by an unobservable mechanism (“meta-changing” environments), and derive an adaptive strategy for optimal growth in these environments. Our focus is on changing nutrient environments, since nutrient metabolism can serve as a model for microbial information-processing more broadly.

Nutrient metabolism as a system for studying microbial information-processing

A natural context in which to study the microbial response to changing environments is metabolic adaptation to nutrients. Because of its strong effect on growth, the way cells adapt to nutrients is a highly selectable trait, either genetically in long-term changing environments (as shown by experimental evolution studies [12,13]) or epigenetically in environments that change on shorter time scales [14,15].

While the control of nutrient and carbon source metabolism has been studied extensively in yeast and other microbes [16], there is generally no simple mapping between the environment’s nutrient composition and microbial cell state (such as the choice of which metabolic pathway to upregulate, or the rate at which to grow). The elaborate molecular machinery for nutrient sensing and uptake suggests that the mapping may be quite complex.

Some of the complexity arises from the fact that microbes prefer to consume some nutrients over others, and that distinct nutrients require different and sometimes mutually exclusive pathways to be expressed. Glucose is generally the preferred sugar and its presence inhibits the expression of pathways required to metabolize alternative sugars like galactose [17]. In yeast, distinct glucose transporters are upregulated depending on glucose levels in the environment, which are sensed by dedicated glucose sensors *Snf3* and *Rgt2* [18,19,20]. Additionally, many promoters in diverse yeast species were shown to be responsive to the presence of different carbon sources in the environment [21]. It has also recently been shown that in environments containing multiple nutrients, cells might be sensitive to complex functions of nutrient levels. Yeast cells decide to turn on the machinery necessary to metabolize galactose (GAL pathway) based on the ratio of glucose to galactose levels in their environment [22].

In addition to molecular complexity of nutrient signaling, there are also memory effects at play in the nutrient response. For example, prior exposure to galactose in yeast alters the rate at which the GAL pathway will be induced upon subsequent galactose exposures [14], and a similar memory phenotype has also been observed for lactose in glucose-lactose switching environments in bacteria [23]. The environment’s nutrient history can also affect single-cell variation in gene expression. Biggar and Crabtree showed that depending on whether previously grown in glucose or raffinose, populations can exhibit single-cell variation in GAL pathway levels when switched to an environment containing a mixture of glucose and galactose [24]. Other lines of the-

oretical and experimental work showed that single-cell variation can be a form of “bet-hedging” that leads to increased fitness under certain conditions [25,26].

Taken together, the intricate molecular machinery underlying nutrient signaling, the effects of nutrient memory, and single-cell variability in response to fluctuations suggest that microbes process information about their environment [27], and take into account both the environment’s history and their internal cell state in making a decision.

Changing discrete environments and inference-based adaptation

To ask how the environment’s history influences microbial decision-making, a number of theoretical and experimental studies have used changing discrete environments [25,23,14]. A discrete environment is shown in Figure 1A (top), where there are two alternating nutrients. Although natural environments are far more complex, discrete switches have been experimentally useful for uncovering mechanisms of nutrient memory [23,14]. Also, in interacting with a host environment, microbes may sense some fluctuations as effectively discrete. For example, en route from the external environment to the gastrointestinal tract, microbes experience sharp shifts in temperature (Figure 1A). Once in the gastrointestinal tract, microbes can face distinct pH regimes, ranging from acidic environment of the stomach (pH 1.5-5) to the intestinal duodenum (pH 5-7), jejunum (pH 7-9) and ileum (pH 7-8) to the colon (pH 5-7) [1]—these fluctuations are shown schematically in Figure 1A. The gut lumen also contains an oxygen gradient [28], and experimentally induced oxygenation or oxygenation as part of medical procedures (such as ileostomies) result in sharp shifts in oxygen levels that reversibly alter microbiota composition [29,28]. Thus, both discrete and continuous features contribute to the gut environment, and discrete environments are a useful approximation for studying the response to environmental change.

In early theoretical work on changing environments [30], Richard Levins argued that the statistical relations between signals in the fluctuating environment are central to adaptation. While distinct cell populations or strains may have different costs associated with each environmental state—e.g., distinct yeast strains exhibit different “preparation times” when switched from a glucose to a galactose environment [31]—the statistical properties of the environment remain informative for adaptation regardless of these costs.

Here, we develop a computational framework for characterizing the statistical structure of changing discrete environments and the adaptive strategies that support optimal growth in these environments. We focus on environments that are characterized by a blend of randomness and order, of the sort one would expect in rich natural environments like the gut or other microbial ecosystems. We propose that adaptation to changing environments can be framed as inference under uncertainty [32]. Although we illustrate our results in terms of glucose-galactose adaptation in yeast, our framework applies broadly to microbial adaptation in other types of fluctuating environments.

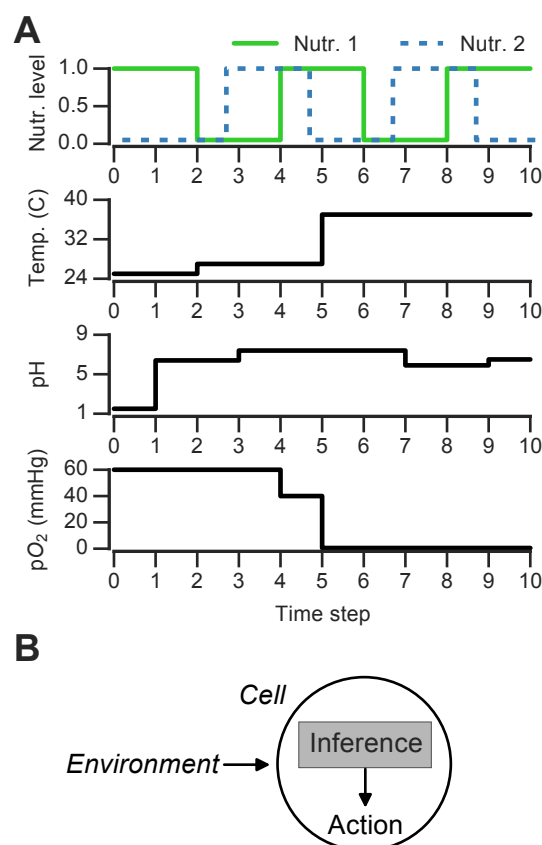


Fig. 1. Discrete changing environments and cellular adaptation to change as inference. (A) Examples of discrete changing environments (time in arbitrary units). Top plot indicates sharp fluctuations in two nutrients. Remaining plots show changes thought to be experienced by microbes when interacting with host gut, which include changes in temperature (25°C in external environment, 27°C on human skin, 37°C in gut), pH and oxygen levels (see main text). (B) Schematic of interaction between cell and changing environment in our framework. Cells sense dynamic environment over time, make inferences about the future state of the environment and use these predictions to take action (e.g., upregulate genes required to metabolize a nutrient).

Results

Growth advantage of using the environment's probabilistic structure

We first asked whether an adaptive strategy, which exploits the probabilistic structure of the environment, would pay off in terms of growth compared with a non-adaptive strategy. We considered a class of *Markov nutrient environments* that fluctuate between two nutrients, glucose and galactose, where nutrient changes follow a Markov model. The probabilistic structure of these environments is determined by two parameters: the probability of transitioning from glucose to glucose ($\theta_{\text{Glu} \rightarrow \text{Glu}}$) and the probability of transitioning from galactose to glucose ($\theta_{\text{Gal} \rightarrow \text{Glu}}$), as shown in Figure 1A. (Equivalently, the model can be parameterized by $\theta_{\text{Glu} \rightarrow \text{Gal}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$, since $\theta_{\text{Glu} \rightarrow \text{Gal}} = 1 - \theta_{\text{Glu} \rightarrow \text{Glu}}$.) Intuitively, the higher $\theta_{\text{Glu} \rightarrow \text{Glu}}$ and the lower $\theta_{\text{Gal} \rightarrow \text{Glu}}$, the

more likely we are to encounter glucose in the environment. Different settings of the transition probabilities can produce qualitatively different environments (Figure 2A). We also assume that one nutrient (in this case, glucose) confers a higher growth rate than alternative nutrients, which is typically the case.

In our framework, the behavior of a cell population is determined by a *policy*: a mapping from the environment's past state to a future action. We compared the performance of two policies in Markov environments: a *posterior predictive* policy, in which cells tune to the most probable nutrient (based on $\theta_{\text{Glu} \rightarrow \text{Glu}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$), and a non-adaptive strategy in which cells are constitutively tuned to the preferred nutrient, glucose. The quantity of interest in the posterior predictive policy is the posterior predictive distribution, which is the probability of a nutrient at the next time step given the previously observed nutrients: $P(C_{t+1} | C_{0:t})$, where C_{t+1} denotes the nutrient at time $t + 1$ and $C_{0:t}$ denotes the environment's nutrient history up until t . The posterior predictive policy chooses the nutrient that maximizes this probability (see Methods for details).

To compare the fitness difference between these policies, we simulated population growth with each policy using a highly simplified growth model, similar to one used in [25]. In this model, we assume that cells respond to changes in the environment with a delay, or "lag". Cells tune to a nutrient at time t and incur a change in growth rate as a consequence of this decision at a later time $t + k$, where k is the lag parameter (we assume here that $k = 1$). For growth kinetics, we assume that: (1) cells grow exponentially when their nutrient state matches the environment's state, (2) there is no switching cost for cells between nutrient states, and (3) when cells are "mismatched" to their environment—i.e., tuned to a nutrient that is not present—their growth rate is zero (this assumption is supported by the observation that cells lacking Gal4, a transcription factor required to activate the GAL pathway, cannot grow in galactose alone [22]). Altogether, our growth assumptions represent extreme conditions, but they serve as a useful starting point for seeing when an adaptive strategy is beneficial to population growth.

Given these growth assumptions, we plotted for each Markov environment the ratio of expected growth rate using the posterior predictive policy to the expected growth rate under the glucose-only policy in Figure 2B (the detailed calculation of these ratios is described in Methods.) To constrain our choice of growth rates, we analyzed growth measurements of 61 yeast strains cultured with different sugars as primary carbon sources from [31] (see Methods for details). As expected, median growth rate in glucose was higher than in other sugars (Figure S1A). Across strains, growth rate in glucose was on average ~ 1.5 fold higher than in galactose and in some strains over 3-fold higher (Figure S1B).

We find that in environments where glucose yields a significantly larger (~ 2 - 4 fold) growth rate than galactose, the posterior predictive policy outperforms the glucose-only policy only in specific regimes of the space of possible Markov environments (red regions in Figure 2B). For a wide range of environments, the non-adaptive policy is equal to

or better than the adaptive policy. As expected, when the difference in growth rate between glucose and galactose gets smaller (Figure 2B, left to right), the payoff from using the posterior predictive policy is greater. However, as our analysis of growth rates in yeast strains showed, glucose typically confers a substantially higher growth rate than galactose.

This idealized calculation shows the importance of the probabilistic structure of the environment in assessing where an adaptive strategy would pay off. This suggests that in order to highlight the advantage of adaptive strategies, specific types of environmental fluctuations will have to be used.

Adapting to meta-changing environments with real-time inference

Our analysis of Markov nutrient environments above doesn't capture several key aspects of adaptive growth in changing environments. First, our environment's changes had a simple "flat" structure describable by only two parameters, whereas natural environment may be generated by far more complex underlying mechanisms. Second, our comparison of the adaptive and glucose-only strategies assumed that the transition probabilities governing the environment, $\theta_{\text{Glu} \rightarrow \text{Glu}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$, are known and can be used by cells. In reality, if this information can be used by cells, it has to be learned from the environment and cannot be assumed as given. Third, such information has to be learned in real-time as cells must respond to the environment while it is changing. We now address each of these aspects in turn.

Changes in complex environments may be governed by dynamic processes that are unobservable to cells. Some environments may oscillate between noisy regimes, where the nutrient switches are less predictable, and stable regimes where the nutrient switches are either rare or more predictable. Such environments can be thought of as "meta-changing" in the sense that there's a change in the way they fluctuate: the probability of being a specific state of the environment (e.g., where a nutrient is available) changes through time, depending on an unobserved condition, such as whether we're in a noisy or stable regime. The transition from noisy to stable regimes might itself be governed by another time-varying mechanism. As an intuitive example of meta-changing environments, consider the eating routine of animals like us. During feeding, bursts of nutrients that are otherwise scarce may become available in the gut. The fluctuations in nutrient levels within a feeding period will depend on what and how much is being consumed. The separation between meals is also subject to randomness, but can still be predictable, depending on how consistent we are in our eating schedule. This high-level structure can be exploited by adaptive systems to anticipate future changes and to separate noisy fluctuations in nutrients from signals of feeding periods.

To understand the adaptive strategies that may be used for effective growth in such environments, we developed a dynamic Bayesian model of meta-changing environments. In our model, we assume a fixed number of hidden "switch states" that correspond to regimes in the environment, and these states are used to generate the fluctuations in nutrients (Figure 3A—see Methods for full model description.)

The switch states and their dynamics are not observable, and therefore have to be learned from the environment. To do this, our model has a prior distribution over the dynamics of the unobservable switch states and the nutrient transitions in the environment (Methods). Through experience with the environment, these priors are updated using Bayesian inference to learn the dynamics that drive nutrient fluctuations. Formally, the goal is to compute the posterior predictive distribution over nutrients, $P(C_{t+1} | C_{0:t})$, which depends on the nutrient history $C_{0:t}$ and on the hidden switch state S_{t+1} :

$$P(C_{t+1} | C_{0:t}) = \sum_{S_{t+1}} P(C_{t+1} | C_t, S_{t+1} = i) P(S_{t+1} | C_{0:t}) \quad (1)$$

For an inference-based strategy to be biologically plausible, it has to be carried out in real-time since cells respond to the environment while it is changing. To compute the posterior predictive distribution (Eq. 1) in real-time, we use a particle filtering algorithm [33] (described in detail in Algorithm S1 and Methods). Rather than naively "memorizing" the environment's history, in particle filtering a representation of the posterior distribution is continuously updated as the environment changes (Methods). The uncertainty of the distribution is represented by a set of "particles" (which can be thought of as samples from a distribution). Each particle corresponds to a configuration of the hidden states of our system. For our model, each particle p is a set $\{s, \mathbf{S}, \mathbf{C}\}$, where s corresponds to the value of the hidden switch state, \mathbf{S} is a transition matrix that tracks the frequencies of transitions between switch states and \mathbf{C} is a transition matrix that tracks the frequency of transitions between nutrients (Methods).

The particle filtering algorithm can be understood by analogy to evolution through mutation and selection. Initially, all particles are weighted equally. Before the environment changes, particles are probabilistically assigned to new configuration based on our model of the environment ("mutation" step). When a new state of the environment is observed, the particles are re-weighted by their fit to this observation and probabilistically resampled using the updated weights ("selection" step). This process repeats as the environment continues to change. Particles that represent more probable states of the environment will get "selected" for through time, while the "mutation" and resampling steps ensure that diversity is brought into the particle population. As each particle p goes through this process, the hidden state and nutrient state configurations in its trajectory are counted by updating the transition matrices \mathbf{S} and \mathbf{C} . Since the particles are propagated through this process in parallel, particle filtering takes what naively would be a serial computation (requiring the complete nutrient history) and converts it to a parallel one that can be performed in real-time. This property may make particle filtering amenable to implementation by molecular circuits, as discussed later.

Signatures of Bayesian adaptation in meta-changing environments

Our model makes a number of predictions about the dynamics of adaptation by systems that represent hidden

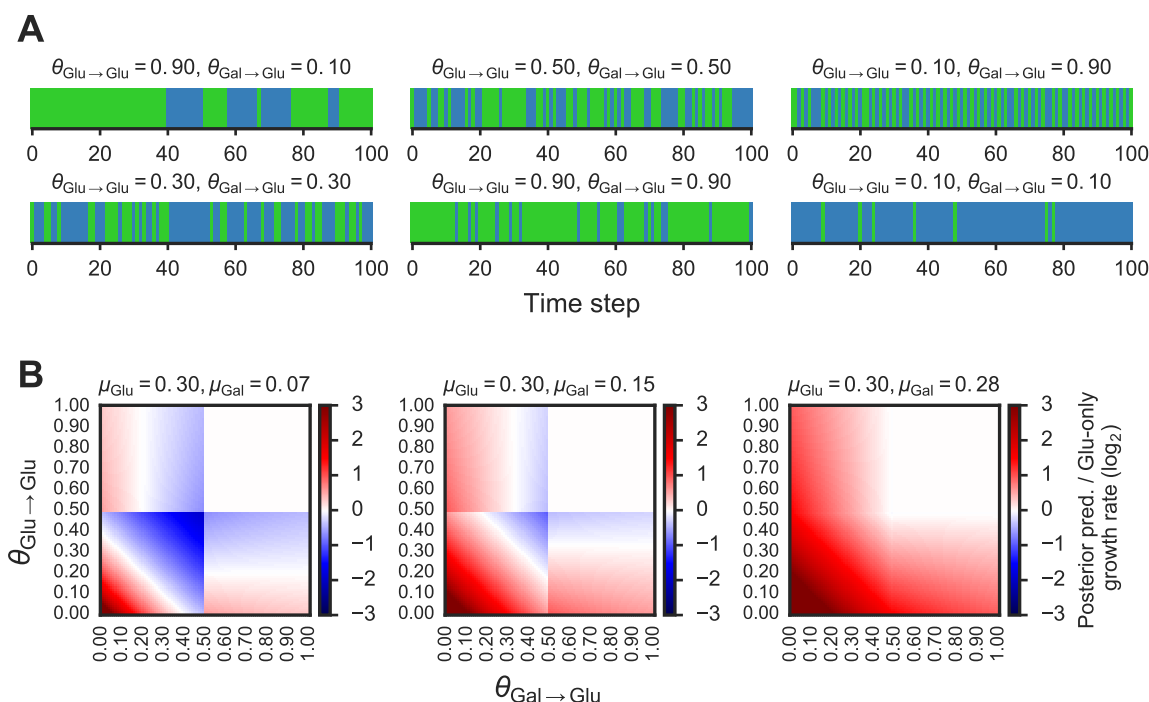


Fig. 2. Fitness benefit of exploiting the probabilistic structure of Markov nutrient environments. (A) Discrete Markov nutrient environments, characterized by two parameters: the transition probability from a glucose state back to glucose state ($\theta_{\text{Glu} \rightarrow \text{Glu}}$) and from galactose state to glucose state ($\theta_{\text{Gal} \rightarrow \text{Glu}}$). Environment assumed to switch from fixed levels of glucose to galactose, visualized as rectangles (collapsing the y-axis from nutrient environment such as one shown in Figure 1A). (B) Comparison of expected growth rates using the posterior predictive compared with the glucose-only strategy. Heat maps show fold-change in expected growth using posterior predictive strategy relative to glucose-only expected growth rate, as a function of the transition probabilities ($\theta_{\text{Glu} \rightarrow \text{Glu}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$) that parametrize the environment (see Methods for detailed calculation). From left to right, increasing galactose growth rate (μ_{Gal}) with fixed glucose growth rate (μ_{Glu}).

environmental states. In Figure 3B, a meta-changing environment is shown that switches between two regimes: one with periodic switches between glucose and galactose, and another where glucose is constant. The posterior predictive probability of glucose in the next time step, $P(C_{t+1} | C_t)$, as it gets updated by real-time inference, is plotted along the environment (Figure 3, top).

The change in the posterior predictive distribution has a number of characteristic features. Starting with a uniform probability over the nutrients, the posterior predictive distribution slowly changes to “learn” the first periodic regime of the environment (Figure 3B, red line). When the environment shifts to the constant regime (at $t = 20$, Figure 3B), the posterior distribution also updates slowly to reflect this. However, when the environment shifts again to the periodic regime (at $t = 40$), the posterior predictive distribution updates faster to reflect the periodicity. Similarly, when the environment shifts to the constant region for the second time (at $t = 60$), the posterior predictive distribution changes even more quickly, since the hidden switch state where the environment is periodic has been seen before. More generally, our model predicts that the more familiar stretches of the environment will be adapted to more quickly. By con-

trast, a Markov model with a “flat” structure that only tracks transition probabilities between glucose and galactose does not show this behavior (Figure 3B, dotted line). This is one of several subtle predictions about the dynamics of adaptation one would expect from an inference-based strategy that uses a representation of hidden environmental states.

Posterior predictive adaptation confers growth advantage in meta-changing environments

We next asked how beneficial the adaptive patterns that result from representing hidden environmental states (of the sort shown in Figure 3) would be for growth. We compared the inference-based growth policy to other policies in meta-changing environments. We considered meta-changing environments that fluctuate between two hidden states: one where there is periodic switching between glucose and galactose, and another where glucose is constant in the environment (Figure 4A). Each hidden state corresponds to a Markov environment (parameterized by a pair of transition probabilities, as described earlier). The switching dynamics between the hidden states are controlled by two transition probabilities, which we labelled p_1 and p_2 in Figure 4A. We compared the population sizes obtained by using the pos-

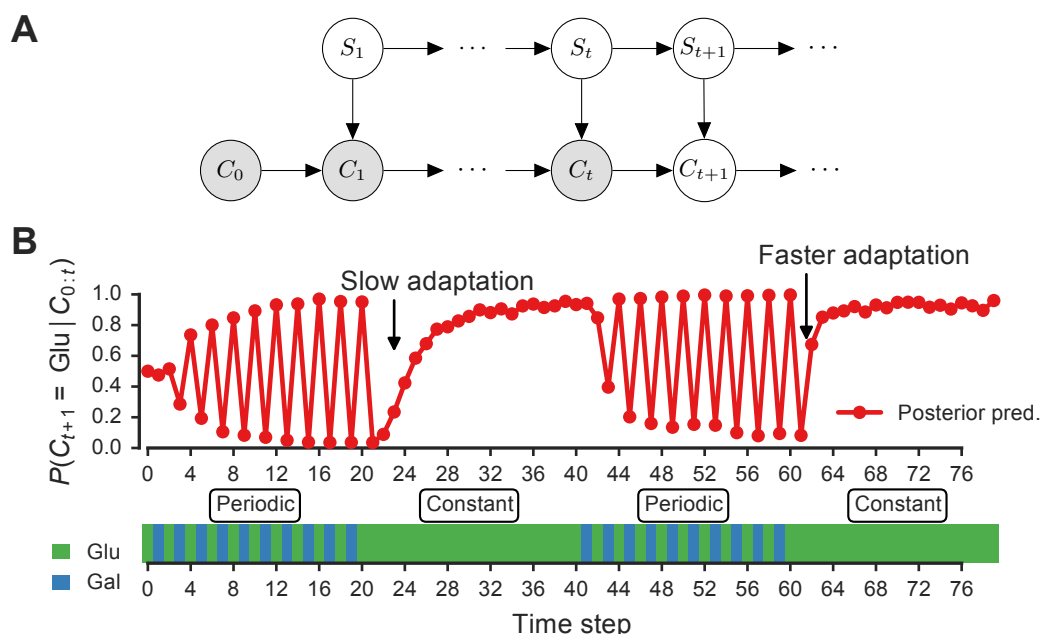


Fig. 3. **Signature of adaptive behavior resulting from using the posterior predictive inference strategy.** (A) Dynamic Bayesian model for meta-changing nutrient environments in graphical model notation [34]: grey nodes represent observed variables, white nodes represent hidden variables. Nutrient value at time t , C_t , depends on the nutrient value at time $t-1$, C_{t-1} , and on the switch state S_t . See Figure S2 for detailed graphical model. (B) Posterior predictive probability, obtained by particle filtering, of glucose using the full model (red), and using a Markov model with “flat” structure and no hidden states (dotted black line). In the first transition from the periodic to constant glucose environment, the posterior predictive probability in the full model updates more slowly compared to the second transition from periodic to constant environment.

terior predictive policy to that obtained by plastic, random and glucose-only policies. In the “plastic” policy (described in [25]), cells tune to the environmental condition that they experienced last, whereas in the random policy the decision is made uniformly at random. We found that across different settings of p_1 and p_2 , the posterior predictive policy generally results in substantially larger population sizes compared with other policies (Figure 4B).

A bet-hedging growth policy based on real-time inference

While all the policies we have considered so far act at the population-level, bet-hedging has been proposed as an adaptive strategy in fluctuating environments [26]. With bet-hedging, different fractions of the population are tuned to different environment conditions, and these cellular states can be inherited through several generations through epigenetic mechanisms [26]. Previously, a bet-hedging policy was shown in simulation to give increased fitness when the bet-hedging proportions matched the frequencies of the environment’s fluctuations [25] (in that study, bet-hedging was referred to as “carry-over”). However, as we discussed above, there is no way for a population of cells to know these probabilities in advance. Our real-time inference scheme lends itself to a bet-hedging policy where the fraction of cells devoted to an environmental state is determined by the posterior predictive probability of this state, which is learned in real-time. Using a growth fitness simulation in

the same meta-changing environments, we found that a bet-hedging posterior predictive policy performs similarly to its population-level counterpart and outperforms a random bet-hedging policy (Figure S3). Real-time probabilistic inference therefore gives us a principled approach for adaptively setting bet-hedging proportions.

Adapting to multi-nutrient meta-changing environments

In natural nutrient environments, unlike in most laboratory conditions, multiple nutrients that can be metabolized by cells may be available. We next asked how distinct growth strategies would do in such multi-nutrient environments.

We analyzed Markov environments with three nutrients—glucose, galactose and maltose—that many yeast strains can grow on as primary carbon sources (as shown in Figure S1). Three qualitatively distinct environments generated by varying the transition probabilities that control these environments are shown in Figure 5A. The first environment has “persistent” stretches of each of the nutrients, the second is a mixture of periodic glucose-galactose switches and persistent maltose states, and the third is one where the presence of glucose signals an upcoming stretch of galactose, while galactose signals an upcoming persistent maltose stretch (Figure 5). Even with only three nutrients, qualitatively rich environments such as these can be constructed.

For fitness comparisons of growth policies, we considered meta-changing environments that switch between the

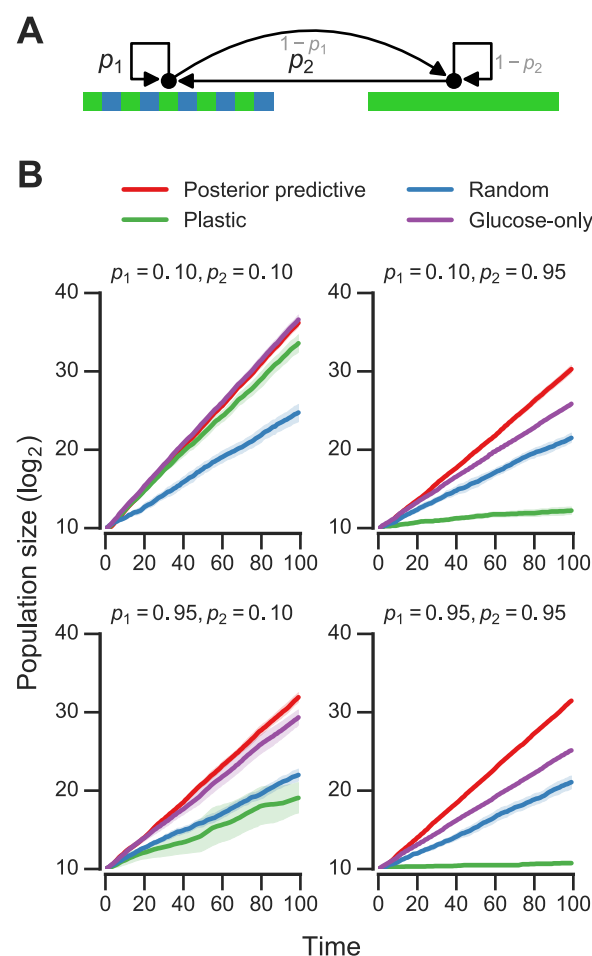


Fig. 4. Fitness of growth policies in meta-changing environments. (A) Meta-changing environment with two hidden states: a periodic environment and a constant environment. The hidden state transitions are parameterized by the probability of transitioning from the periodic environment to itself, p_1 , and from the constant to the periodic environment, p_2 . (B) Population sizes obtained by growth using different policies in meta-changing environments (mean population size from 20 simulations plotted with bootstrap confidence intervals as shaded regions). Four environments shown, each parameterized by different settings of p_1, p_2 (where μ_{Glu} is twice μ_{Gal}). Posterior predictive policy generally outperforms other policies. Mean population size from 20 simulations plotted with bootstrap confidence intervals (shaded regions).

second and third multi-nutrient environments shown in Figure 5A. These switches, as in previous meta-changing environments, are governed by two transition probabilities, p_1 and p_2 (Figure 5B, top). It's straightforward to apply the posterior predictive policy to multi-nutrient environments by running real-time inference using a version of our Bayesian model that assumes three rather than two nutrient states. We compared the growth performance of this posterior predictive policy to that of the plastic, random and glucose-only policies for different settings of p_1 and p_2 (Figure 5B). In all cases, the posterior predictive policy produced larger population sizes than other policies.

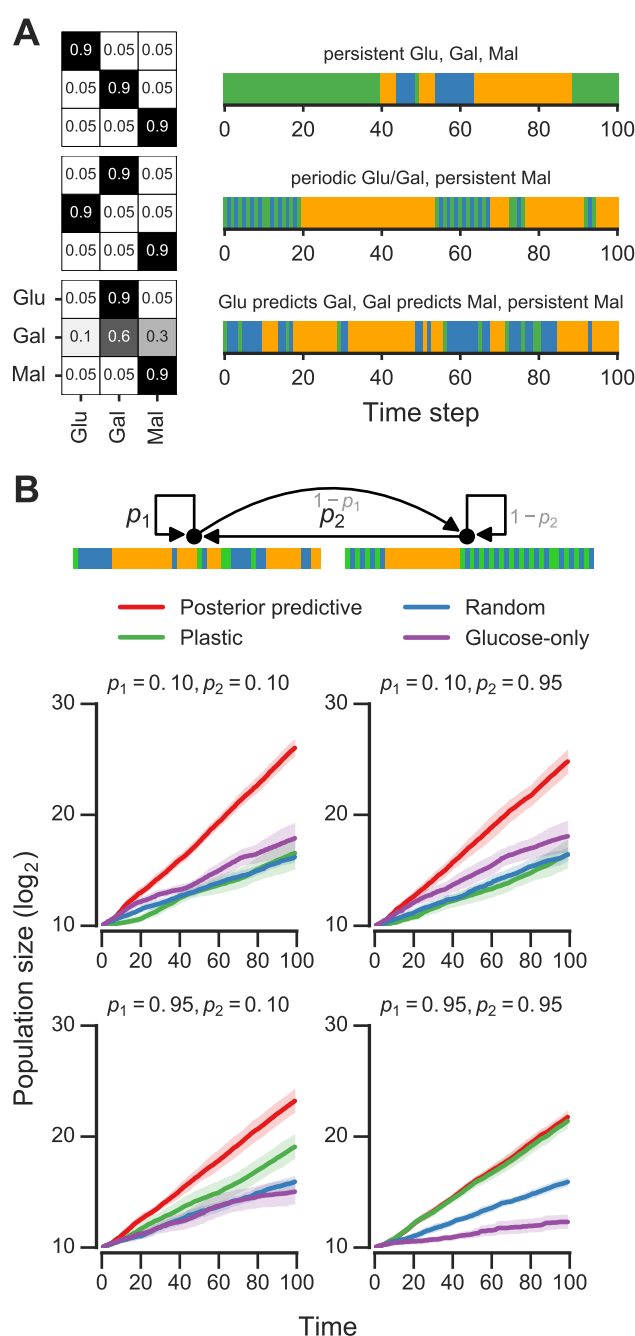


Fig. 5. Fitness of growth policies in multi-nutrient meta-changing environments. (A) Three multi-nutrient Markov environments where glucose, galactose and maltose fluctuate. Transition probability matrices shown as heat maps (left) along with the environments they produce starting with glucose as initial state (right). (B) Top: a meta-changing multi-nutrient environment that switches between the second and third Markov environments shown in panel A. Bottom: population sizes obtained using different policies in meta-changing environment shown in top, shown for four different settings of p_1, p_2 . Growth rate settings used: μ_{Glu} was twice μ_{Gal} and $\mu_{\text{Gal}} = \mu_{\text{Mal}}$. Mean population size from 20 simulations plotted with bootstrap confidence intervals (shaded regions).

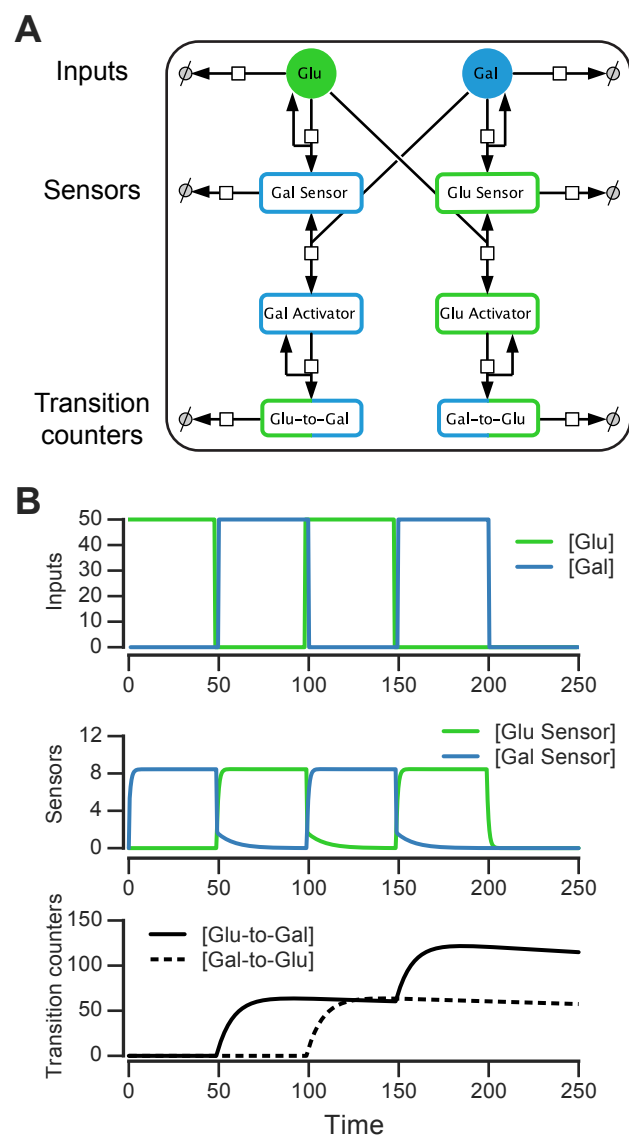


Fig. 6. Molecular circuit implementation of a nutrient transition counter. (A) Network of chemical reactions for implementing nutrient transition counter. \emptyset denotes null species in degradation reactions. Reaction equations represented by the network are listed in Table S1. (B) Simulation of circuit component concentrations (middle, bottom) given the input doses of glucose and galactose (top). (Concentrations are in A.U.)

Molecular circuit design for inference-based adaptation

Our results suggest that cell populations that use a probabilistic inference-based growth strategy can achieve greater fitness, but it's not clear how such a strategy can be realized in molecular circuits. We next asked if the particle filtering algorithm can guide the design of molecular circuits that implement the posterior predictive strategy.

In order to implement the posterior predictive strategy, we need a way to encode the inference algorithm's representation of the environment in molecules. Our particle filtering algorithm uses a minimal representation of (Markovian) en-

vironments that a circuit that performs inference in real-time would have to track. This representation consists of the hidden switch state s , and a pair of transition matrices S and C that track hidden state and nutrient state transitions, respectively (as described above and in Methods). When the environment changes, these representations are updated. At the molecular level, binary state such as the hidden switch state s can be represented easily using phosphorylation, protein dimerization or other known molecular switches. The more challenging part of implementing inference is keeping track of the environment's changes. As the environment changes, our inference algorithm tracks the number of times the environment transitioned from one nutrient state to another, effectively keeping a counter. Recall that the Markov environments we analyzed can be parameterized by the frequency of transitions from glucose to galactose, and from galactose to glucose. Therefore, our circuit would have to distinguish glucose-to-galactose from galactose-to-glucose switches and "count" these switches as the environment fluctuates.

We constructed a molecular circuit that implements such a nutrient transition counter. The eight-component circuit is shown in Figure 6A and its reaction equations are listed in Table S1 (all parameters used in reactions are given in Methods). The circuit consists of sensors that are activated by sugars, activators that are produced downstream of the sensors, and proteins that count glucose-to-galactose and galactose-to-glucose transitions. In the left branch of the circuit, glucose catalytically activates a galactose sensor (Figure 6A). When galactose is present, the galactose sensor reversibly forms a galactose activator ($\text{Gal_Sensor} + \text{Gal} \rightleftharpoons \text{Gal_Activator}$). Similarly, in the right branch of the circuit, galactose catalytically activates a glucose sensor, which in the presence of glucose reversibly forms a glucose activator ($\text{Glu_Sensor} + \text{Glu} \rightleftharpoons \text{Glu_Activator}$). The glucose activator triggers the production of a glucose-to-galactose transition counting molecule, while the galactose activator produces a galactose-to-glucose transition counting molecule. The transition counting molecules are assumed to have a very slow degradation rate, and this rate determines the stability of the counter's "memory".

We simulated the behavior of this circuit in an environment that switches between glucose and galactose (Figure 6B, top). The glucose sensor is active during galactose pulses and the galactose sensor is active during glucose pulses. When the environment first switches from glucose to galactose (at $t = 50$, Figure 6B), the galactose activator is formed and triggers a spike in the glucose-to-galactose counter (Figure 6B, bottom). When the environment switches from galactose to glucose, the galactose-to-glucose counter spikes (at $t = 100$). The counter molecules are highly stable, so their level forms the "memory" of these two transitions. When the environment switches from glucose to galactose for the second time (at $t = 150$), the glucose-to-galactose counter spikes again to a level roughly twice that of the galactose-to-glucose counter, while the galactose-to-glucose counter is unaffected. After all nutrient switches, the circuit retains that 2 glucose-to-galactose and 1 galactose-to-glucose transitions have been observed. This information can be used downstream to implement an

inference-based adaptive strategy, like the posterior predictive strategy.

A key feature of this circuit architecture is that sensors associated with one nutrient (e.g., glucose) get activated by other nutrients (such as galactose). This “crosstalk” between the two arms of the pathway enables the environmental change tracking that is needed for inference. While a true digital counter is unbounded, this molecular counter’s dynamic range and reliability is limited by the degradation rates and dynamic ranges of the molecular components involved (such as the sensors and counting molecules in Figure 6A).

This gives a proof of concept design of the core machinery needed for real-time inference in our probabilistic model, but a full implementation would require integration with the basal glucose/galactose signaling network and careful analysis of the circuit’s robustness and precision.

Discussion

Fluctuations in complex environments can be driven by mechanisms that cells cannot sense directly. Growth strategies that represent these hidden environmental features and use inference to predict future changes give cells a fitness advantage. An important future direction would be to test if signatures of adaptation by inference (such as those in Figure 3B) can be observed experimentally in glucose-galactose switching with yeast [14] or glucose-lactose switching with bacteria [23].

Although we assumed in our fitness simulations that the goal is to maximize population-level fitness, other goals—like minimizing the probability of population extinction [35]—can be more relevant in some environments, especially for small population sizes, and these should be investigated further. Another limitation of our analysis is the assumption that environmental fluctuations follow a Markov process; an assumption violated by many natural time-varying processes. However, dynamic Bayesian models similar to the one presented here have been extended to handle non-Markov environments [36]. It will be fruitful to experiment with these models and compare their assumptions to the statistical properties of natural microbial environments.

Another future challenge is to link the continuous features of the environment (which can be clearly sensed by microbes) to more abstract discrete structure like that of meta-changing environments. Elegant work by Sivak and Thomson derived optimal enzyme induction kinetics for the noisy statistics of an environment with continuously varying nutrients [37]. This suggests that in an ideal adaptive system, principles of optimal inference are at work at multiple layers—from the abstract computational problem of anticipating the next nutrient to the quantitative decision of how much of the relevant enzymes to induce. More work is needed to link abstract computations to these lower mechanistic levels.

To represent the structure of meta-changing environments, our model posited a finite number of hidden states that drive nutrient fluctuations. The number of hidden states was fixed in advance, but nonparametric dynamic Bayesian

models offer a principled alternative [38,39]. In these models, the number of hidden states is learned from observation. Recent work in computational linguistics [40] proposed a particle inference algorithm for a nonparametric dynamic Bayesian model of word segmentation, a task that, like nutrient adaptation, has to be performed in real-time. It would be interesting to investigate whether molecular kinetics can implement such nonparametric Bayesian inference procedures.

While we have focused on glucose-galactose environments, our framework generally applies to environments that change too quickly for mutation and natural selection to take hold. This is distinct from cases where natural selection (e.g., through experimental evolution, as in [12,13]) rewires circuits genetically to better respond to the predictable structure of the environment.

It remains open how inference-based adaptive strategies that apply on short timescales can be implemented at the molecular level, either in natural or engineered cellular circuits. The molecular mechanisms needed to implement these strategies are likely to be epigenetic, based in chromatin [14] or stable protein inheritance [15]. Recent work argued compellingly for developing methods that “compile” abstract computational problems, like probabilistic inference, to molecular descriptions that are physically implementable [41]. In this work, a scheme was proposed for solving exact inference for probabilistic graphical models using chemical reaction networks, with DNA strand displacement as the physical instantiation [41]. This choice of substrate is implausible as a mechanism for cellular computation, however (and striving for exact as opposed to approximate inference may be too restrictive for many computational problems).

In a different approach, an intracellular kinetic scheme that implements a real-time probabilistic decision procedure for a simple continuously changing environment was proposed [42]. An open challenge is to extend these schemes to handle structured environments, such as the meta-changing environments we have considered, and to define the molecular components would be needed to build these circuits *in vivo*. We have outlined an implementation of one critical part of such a circuit, but more work is needed to understand the precision and reliability of such circuits in the presence of gene expression variability and cell division.

Real-time inference algorithms, such as particle filtering, have the potential to guide the construction of synthetic cellular circuits that adapt to rich changing environments. Since particle filtering algorithms rely on noise, these procedures point to areas where biochemical noise (in gene expression or protein interactions) would not only be tolerated, but would be required for inference to work. These algorithmic constraints may inform the design of synthetic circuits that implement probabilistic computation out of noisy molecular parts.

Methods

Growth rate measurements

Growth rates for 61 yeast strains were taken from [31]—cultures were prepared as described therein. Briefly, OD600 growth measurements were log-transformed, fit by

splines and the region with maximal derivative in the spline fit (“exponential phase”) was used to calculate the growth rate (defined as doublings per hour). Each strain was measured in duplicate and the average growth rate was used in Figure S1.

Optimal policies in Markov environments with two nutrients

In Figure 2, the ratio of the expected growth rate obtained by following a posterior predictive policy (where the most probable nutrient under the posterior is chosen in the next step) to expected growth rate using a glucose-only policy is shown. Below, we describe in detail how this ratio was calculated.

To compare the growth rate differences between a glucose-only policy and the posterior predictive policy in two-nutrient Markov environments, we assumed an idealized case where the transition probabilities $\theta_{\text{Glu} \rightarrow \text{Glu}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$ are known. The “optimal” policy relative to an environment is one that maximizes the expected growth rate. The expected growth rate is dependent on the environment’s transition probabilities, the growth rates afforded by each of the nutrients, as well as the cost of being “mismatched” to the environment (i.e. being tuned to a nutrient that isn’t present in the environment.)

In our Markov nutrient environment, there are two environment states (glucose or galactose) and two possible actions for the cell population, each associated with a different growth rate. Notation for these states and parameters is as follows:

- Environment states: the state of the environment at time t is represented by the random variable C_t , which takes on one of two values, $c_1 = \text{Glu}$, $c_2 = \text{Gal}$.
- Transition probabilities: $\theta_{\text{Glu} \rightarrow \text{Glu}}$, $\theta_{\text{Gal} \rightarrow \text{Glu}}$
- Actions: tuning to glucose ($\alpha_1 = \text{Glu}$) or tuning to galactose ($\alpha_2 = \text{Gal}$)
- Growth rates associated with each action and environment state:
 - When tuned to glucose in glucose environment: V_{11}
 - When tuned to glucose in galactose environment: V_{12}
 - When tuned to galactose in galactose environment: V_{22}
 - When tuned to galactose in glucose environment: V_{21}

A policy π is a mapping from an environment’s state to an action. We can now write down the *conditional growth rate* associated with a particular policy, $R(\pi | C_{t-1})$, which is the growth rate given the previous environment state C_{t-1} . Let π_1 and π_2 correspond to policies that constitutively tune to glucose or galactose, respectively. The conditional growth rates for these policies are:

$$R(\pi_1 | C_{t-1}) = V_{11}P(C_t = c_1 | C_{t-1}) + V_{12}P(C_t = c_2 | C_{t-1})$$

$$R(\pi_2 | C_{t-1}) = V_{22}P(C_t = c_2 | C_{t-1}) + V_{21}P(C_t = c_1 | C_{t-1})$$

For simplicity, we assume that the growth rate is zero when the internal state of the cells is mismatched to the environment, i.e. $V_{12} = V_{21} = 0$. The conditional growth rates then simplify to:

$$R(\pi_1 | C_{t-1}) = V_{11}P(C_t = c_1 | C_{t-1})$$

$$R(\pi_2 | C_{t-1}) = V_{22}P(C_t = c_2 | C_{t-1})$$

To get the expected growth rates, we sum over the possible states of the environment at time $t - 1$, yielding:

$$R(\pi_1) = V_{11}[\theta_{\text{Glu} \rightarrow \text{Glu}} + \theta_{\text{Gal} \rightarrow \text{Glu}}]$$

$$R(\pi_2) = V_{22}[(1 - \theta_{\text{Glu} \rightarrow \text{Glu}}) + (1 - \theta_{\text{Gal} \rightarrow \text{Glu}})]$$

The policy π_1 is optimal when $\frac{R(\pi_1)}{R(\pi_2)} > 1$.

Unlike the glucose-only or galactose-only policy, the posterior predictive policy chooses the next action based on the transition probabilities $\theta_{\text{Glu} \rightarrow \text{Glu}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$. This policy, denoted π_3 , chooses the most probable nutrient as a function of the environment’s previous state C_{t-1} :

$$\pi_3(C_{t-1}) = \begin{cases} \begin{cases} \alpha_1 & \text{if } \theta_{\text{Glu} \rightarrow \text{Glu}} > 0.5 \\ \alpha_2 & \text{otherwise} \end{cases} & \text{if } C_{t-1} = c_1, \\ \begin{cases} \alpha_1 & \text{if } \theta_{\text{Glu} \rightarrow \text{Glu}} > 0.5 \\ \alpha_2 & \text{otherwise} \end{cases} & \text{if } C_{t-1} = c_2, \end{cases}$$

We can now write the expected growth rate for the posterior predictive policy by analyzing each of the cases involving the transition probabilities $\theta_{\text{Glu} \rightarrow \text{Glu}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$:

$$R(\alpha_3) = \begin{cases} V_{11}\theta_{\text{Glu} \rightarrow \text{Glu}} + V_{11}\theta_{\text{Gal} \rightarrow \text{Glu}}, & \text{if } \theta_{\text{Glu} \rightarrow \text{Glu}}, \theta_{\text{Gal} \rightarrow \text{Glu}} \geq 0.5 \\ V_{11}\theta_{\text{Glu} \rightarrow \text{Glu}} + V_{22}(1 - \theta_{\text{Gal} \rightarrow \text{Glu}}), & \text{if } \theta_{\text{Glu} \rightarrow \text{Glu}} \geq 0.5 \text{ and } \theta_{\text{Gal} \rightarrow \text{Glu}} < 0.5 \\ V_{22}(1 - \theta_{\text{Glu} \rightarrow \text{Glu}}) + V_{11}\theta_{\text{Gal} \rightarrow \text{Glu}}, & \text{if } \theta_{\text{Glu} \rightarrow \text{Glu}} < 0.5 \text{ and } \theta_{\text{Gal} \rightarrow \text{Glu}} \geq 0.5 \\ V_{22}(1 - \theta_{\text{Glu} \rightarrow \text{Glu}}) + V_{22}(1 - \theta_{\text{Gal} \rightarrow \text{Glu}}), & \text{if } \theta_{\text{Glu} \rightarrow \text{Glu}}, \theta_{\text{Gal} \rightarrow \text{Glu}} < 0.5 \end{cases}$$

The ratio $\frac{R(\pi_3)}{R(\pi_1)}$ is shown in Figure 2 as a function of $\theta_{\text{Glu} \rightarrow \text{Glu}}$ and $\theta_{\text{Gal} \rightarrow \text{Glu}}$ for different values of each nutrient’s growth rate (V_{11} and V_{22}).

Bayesian model for meta-changing nutrient environments

Here we describe in detail the dynamic Bayesian model used to adapt to meta-changing environments. This model is in the family of dynamic probabilistic models, also called “switching state space models”, that have been widely used in artificial intelligence, robotics and machine learning [43,44]. In our model, there are K -many hidden switch states that are used to produce one of J -many “outputs”. The outputs are the different nutrients and the switch states correspond to stretches of the environment such as the “periodic” or “constant” regions shown in Figure 4. Each switch state represents a transition matrix that is used to generate nutrient fluctuations. The nutrient produced at time t , C_t , depends on the value of the switch state S_t and the previous nutrient C_{t-1} . The next switch state S_{t+1} is generated conditioned on S_t based on a separate set of

switch state transition probability, then C_{t+1} is generated conditioned on S_{t+1} so on.

The transition probabilities associated with each switch state, as well as the probabilities of transitioning between switch states, all have to be learned from the environment. We therefore place a prior on these transition probabilities.

The full graphical model including hyperparameters is shown in Figure S2 using plate notation [34]. Formally, a switch state takes on one of $1, \dots, K$ values. The initial switch state S_1 is drawn from a probability distribution on the initial switch state values, π_{s_1} . Since π_{s_1} is unknown, we place a prior on it using the Dirichlet distribution, i.e.:

$$\begin{aligned}\pi_{s_1} &\sim \text{Dirichlet}(\alpha_{s_1}) \\ S_1 | \pi_{s_1} &\sim \text{Multinomial}(\pi_{s_1}, 1)\end{aligned}$$

which means $P(S_1 = i | \pi_{s_1}) = \pi_{s_1}^{(i)}$, where $\pi_{s_1}^{(i)}$ denotes the i th entry in π_{s_1} . For clarity, we will sometimes omit the explicit value assignment for a random variable, and write $P(S_t = i | \pi_{s_1})$ as simply $P(S_1 | \pi_{s_1})$. The switch states at later time points are generated as follows: each switch state $i \in \{1, \dots, K\}$ has an associated probability vector π_i , whose j th entry $\pi_i^{(j)}$ determines the probability of transitioning from switch state i to j . The probability of the switch state at time step $t > 1$ therefore depends on the previous switch state's value S_{t-1} and the switch state transition probabilities:

$$\begin{aligned}\pi_i &\sim \text{Dirichlet}(\alpha_s) \\ S_t | S_{t-1} = i &\sim \text{Multinomial}(\pi_i, 1)\end{aligned}$$

Similarly, each nutrient state can take one value $i \in \{1, \dots, J\}$, and the initial nutrient state C_0 is drawn from a probability distribution, π_{c_0} , which is in turn drawn from a Dirichlet prior distribution:

$$\begin{aligned}\pi_{c_0} &\sim \text{Dirichlet}(\alpha_{c_0}) \\ C_0 | \pi_{c_0} &\sim \text{Multinomial}(\pi_{c_0}, 1)\end{aligned}$$

The probability of a nutrient output at time $t > 0$ depends on the previous nutrient C_{t-1} and on the switch state at time t . The switch state value indexes which transition probability distribution to use for the outputs, and as before, the transition probabilities are drawn from a Dirichlet prior:

$$\begin{aligned}\pi'_{i,j} &\sim \text{Dirichlet}(\alpha_c) \\ C_t | C_{t-1} = i, S_t = j &\sim \text{Multinomial}(\pi'_{i,j}, 1)\end{aligned}$$

The posterior predictive distribution, $P(C_{t+k} | C_{0:t})$, is the main quantity of interest (we assume $k = 1$ as in all of our simulations). This distribution can be calculated recursively in dynamic probabilistic models, a property that we will use in a later section to derive a real-time estimation procedure for this distribution using particle filtering. (For an accessible introduction to recursive estimation of Bayesian posteriors and to particle filtering, see Ch. 1 in [44], Ch. 1-3 in

[45] or [33]). We derive the posterior predictive distribution in steps. First, consider the posterior distribution over the switch state at time t given the environment history up until and including time t , $P(S_t | C_{0:t})$, called the *filtering posterior* distribution, which is obtained by marginalizing out the switch state S_{t-1} :

$$P(S_t | C_{0:t}) \propto P(C_t | S_t) \sum_{S_{t-1}} P(S_t | S_{t-1}) P(S_{t-1} | C_{0:t-1})$$

Note that the third term is the filtering posterior over the switch state at time $t-1$, which can be rewritten as we did above in terms of the filtering posterior at $t-2$, and so on—this shows that the posterior can be computed recursively. The base case of the recursion is given by our prior distributions over the initial nutrient state C_0 and the initial switch state S_1 (as shown in Figure S2).

The posterior predictive distribution $P(C_{t+1} | C_{0:t})$ can then be written as a product that uses the filtering posterior, marginalizing out the hidden switch states:

$$P(C_{t+1} | C_{0:t}) \propto \sum_{S_t} \sum_{S_{t+1}} P(C_{t+1} | C_t, S_{t+1}) P(S_{t+1} | S_t) P(S_t | C_{0:t-1})$$

The distributions that the filtering posterior decomposes to depend on parameters that are unobserved, such as the transition probabilities. Since Dirichlet-Multinomial distributions are conjugate [46], we can analytically integrate out these parameters. For example, the transition probabilities π for the switch states can be integrated out:

$$P(S_{t+1} | S_t = i) = \int P(S_{t+1} | S_t = i, \pi_i) P(\pi_i | \alpha_s) d\pi_i$$

Similarly, the probability of observing a nutrient given the previous nutrient and the switch state, $P(C_{t+1} | C_t, S_{t+1})$, can be computed while integrating out the nutrient transition probabilities $\pi'_{i,j}$:

$$P(C_{t+1} | C_t = j, S_{t+1} = i) = \int P(C_{t+1} | C_t = j, S_{t+1} = i, \pi'_{i,j}) P(\pi'_{i,j} | \alpha_c) d\pi'_{i,j}$$

We discuss in detail how to estimate these distributions from observations in the section on particle filtering.

Values of hyperparameters: In all analyses, we set hyperparameters $\alpha_{s_1}, \alpha_{c_0}, \alpha_c$ to all 1's vector. The hyperparameter on switch state transitions α_s was set such that self-transitions get the hyperparameter value 2, i.e. for the i th switch state, $\alpha_s^{(i)} = 2$, and all other entries in α_s are set to 1. This encodes a weakly “sticky” prior that slightly favors self-transitions for hidden switch states.

Real-time inference using particle filtering

We estimate the posterior predictive distribution above in real-time using particle filtering [33] as outlined in Algorithm S1. Particle filtering for our model is implemented in the [particlefever](#) library (available on Github.) We used 200 particles for all simulations with particle states initialized from the prior.

In particle filtering, a distribution of hidden state values (in our model, the hidden switch states) is represented using a set of particles. Each particle corresponds to a configuration of the hidden states (configurations are typically assigned from the prior distribution at time $t = 0$). The particles are assigned weights that are initialized to be uniform. Starting with time t , the particle filtering algorithm works by first *predicting* hidden state values particles $t + 1$. When a data point at time $t + 1$ is observed, the weight of each particle is *updated* to be proportional to the likelihood of the new data point given the particle’s hidden state configuration. This process is repeated as data points are observed. To prevent particle degeneracy (a case where particles get weights that are too low), a resampling step is used where particles are sampled in proportion to their weights and the weights are reset to be uniform.

We now turn to the representation of our state space that is encoded in each particle. As discussed above, because of the conjugacy of Dirichlet-Multinomial distributions [46], we can analytically integrate out the transition probabilities in our model. This means that these transition probabilities don’t have to be represented in our particles. Instead, the sufficient statistics’ for our model are simply: (1) a matrix of counts \mathbf{S} where $\mathbf{S}^{(i,j)}$ is the number of times hidden switch state i transitioned to hidden state j in the particle’s trajectory, and (2) a three-dimensional array of counts \mathbf{C} where $\mathbf{C}^{(i,j,k)}$ is the number of times nutrient i transitioned to nutrient j under switch state k . The predict/update cycles for a particle $p = \{s_t, \mathbf{S}, \mathbf{C}\}$, where s is the particle’s hidden state at time t , are:

Prediction step: For each particle p we draw a new switch state for $t + 1$, $s_{t+1} \sim P(S_{t+1} | S_t = s_t)$. By conjugacy, the switch state transition probabilities π_{s_t} can be integrated out, yielding the posterior predictive distribution for a Dirichlet-Multinomial:

$$P(S_{t+1} = s_{t+1} | S_t = s_t) \propto \frac{\alpha_s + \mathbf{S}^{(s_t, s_{t+1})}}{\sum_k (\alpha_s + \mathbf{S}^{(s_t, k)})}$$

which can be sampled from.

Updating step: When a nutrient c_{t+1} is observed, we update the weight w of our particle $p = \{s_{t+1}, \mathbf{S}, \mathbf{C}\}$ in proportion to $P(C_{t+1} | C_t, S_{t+1})$:

$$w \propto P(C_{t+1} = c_{t+1} | C_t = c_t, S_{t+1} = s_{t+1})$$

Integrating out the nutrient transition probabilities $\pi'_{c_t, c_{t+1}}$, this also gives the Dirichlet-Multinomial posterior predictive

distribution:

$$w \propto \frac{\alpha_c + \mathbf{C}^{(c_t, c_{t+1}, s_{t+1})}}{\sum_k (\alpha_c + \mathbf{C}^{(c_t, k, s_{t+1})})}$$

See Algorithm S1 for remaining details.

Fitness simulations

We simulated growth with different policies using a simple model of exponential growth. Cells were assumed to grow exponentially with a growth rate determined by the environment’s nutrient state. The initial population size and the time duration of each environment simulated are as described in figure legends. All code for fitness simulations is available in the paper’s [Github repository](#).

Molecular implementation of nutrient transition counter

The nutrient transition counter model was drawn in CellDesigner (version 4.4) [47], serialized as an SBML file, and simulated in Python using [libRoadRunner](#). SBML file for the model is available at [10.6084/m9.figshare.3493994](#). A detailed report of the chemical reactions and rate parameters was generated by SBML2 \LaTeX [48] and is available at [10.6084/m9.figshare.3492185](#).

Acknowledgments

We thank Matt Johnson, Tommi Jaakkola, Bo Hua, Jennifer Chen, Nikolai Slavov, Andrew Bolton, Eric Jonas, Lauren Surface, Ariella Azoulay and Josh Tenenbaum for helpful discussions. We thank Jue Wang for sharing raw data files from published growth rate measurements.

References

- [1] J. Walter and R. Ley. The human gut microbiome: ecology and recent evolutionary changes. *Annu. Rev. Microbiol.*, 65:411–429, 2011.
- [2] L. W. Parfrey, W. A. Walters, and R. Knight. Microbial eukaryotes in the human microbiome: ecology, evolution, and future directions. *Front Microbiol.*, 2:153, 2011.
- [3] C. A. Thaiss, D. Zeevi, M. Levy, G. Zilberman-Schapira, J. Suez, A. C. Tengeler, L. Abramson, M. N. Katz, T. Korem, N. Zmora, Y. Kuperman, I. Biton, S. Gilad, A. Harmelin, H. Shapiro, Z. Halpern, E. Segal, and E. Elinav. Transkingdom control of microbiota diurnal oscillations promotes metabolic homeostasis. *Cell*, 159(3):514–529, Oct 2014.
- [4] L. A. David, A. C. Materna, J. Friedman, M. I. Campos-Baptista, M. C. Blackburn, A. Perrotta, S. E. Erdman, and E. J. Alm. Host lifestyle affects human microbiota on daily timescales. *Genome Biol.*, 15(7):R89, 2014.
- [5] L. A. David, C. F. Maurice, R. N. Carmody, D. B. Gootenberg, J. E. Button, B. E. Wolfe, A. V. Ling, A. S. Devlin, Y. Varma, M. A. Fischbach, S. B. Biddinger, R. J. Dutton, and P. J. Turnbaugh. Diet rapidly and reproducibly alters the human gut microbiome. *Nature*, 505(7484):559–563, Jan 2014.
- [6] A. J. Brown, G. D. Brown, M. G. Netea, and N. A. Gow. Metabolism impacts upon Candida immunogenicity and pathogenicity at multiple levels. *Trends Microbiol.*, 22(11):614–622, Nov 2014.
- [7] C. A. Kumamoto and M. D. Vences. Alternative Candida albicans lifestyles: growth on surfaces. *Annu. Rev. Microbiol.*, 59:113–133, 2005.

- [8] F. Cottier and F. A. Muhlschlegel. Sensing the environment: response of *Candida albicans* to the X factor. *FEMS Microbiol. Lett.*, 295(1):1–9, Jun 2009.
- [9] David Marr. Vision: A computational investigation into the human representation and processing of visual information. *New York: WH Freeman*, 1982.
- [10] F. J. Poelwijk, M. G. de Vos, and S. J. Tans. Tradeoffs and optimality in the evolution of gene regulation. *Cell*, 146(3):462–470, Aug 2011.
- [11] A. M. New, B. Cerulus, S. K. Govers, G. Perez-Samper, B. Zhu, S. Boogmans, J. B. Xavier, and K. J. Verstrepen. Different levels of catabolite repression optimize growth in stable and variable environments. *PLoS Biol.*, 12(1):e1001764, Jan 2014.
- [12] A. Mitchell, G. H. Romano, B. Groisman, A. Yona, E. Dekel, M. Kupiec, O. Dahan, and Y. Pilpel. Adaptive prediction of environmental changes by microorganisms. *Nature*, 460(7252):220–224, Jul 2009.
- [13] I. Tagkopoulos, Y. C. Liu, and S. Tavazoie. Predictive behavior within microbial genetic networks. *Science*, 320(5881):1313–1317, Jun 2008.
- [14] Sarah R Stockwell, Christian R Landry, and Scott A Rifkin. The yeast galactose network as a quantitative model for cellular memory. *Molecular bioSystems*, 11(1):28–37, Jan 2015.
- [15] D. F. Jarosz, A. K. Lancaster, J. C. Brown, and S. Lindquist. An evolutionarily conserved prion-like element converts wild fungi from metabolic specialists to generalists. *Cell*, 158(5):1072–1082, Aug 2014.
- [16] J. R. Broach. Nutritional control of growth and development in yeast. *Genetics*, 192(1):73–105, Sep 2012.
- [17] J. M. Gancedo. Carbon catabolite repression in yeast. *Eur. J. Biochem.*, 206(2):297–313, Jun 1992.
- [18] G. M. Santangelo. Glucose signaling in *Saccharomyces cerevisiae*. *Microbiol. Mol. Biol. Rev.*, 70(1):253–282, Mar 2006.
- [19] S. Busti, P. Coccetti, L. Alberghina, and M. Vanoni. Glucose signaling-mediated coordination of cell growth and cell cycle in *Saccharomyces cerevisiae*. *Sensors (Basel)*, 10(6):6195–6240, 2010.
- [20] S. Ozcan, J. Dover, and M. Johnston. Glucose sensing and signaling by two glucose receptors in the yeast *Saccharomyces cerevisiae*. *EMBO J.*, 17(9):2566–2573, May 1998.
- [21] K. Weinhandl, M. Winkler, A. Glieder, and A. Camattari. Carbon source dependent promoters in yeasts. *Microb. Cell Fact.*, 13:5, 2014.
- [22] R. Escalante-Chong, Y. Savir, S. M. Carroll, J. B. Ingraham, J. Wang, C. J. Marx, and M. Springer. Galactose metabolic genes in yeast respond to a ratio of galactose and glucose. *Proc. Natl. Acad. Sci. U.S.A.*, 112(5):1636–1641, Feb 2015.
- [23] Guillaume Lambert and Edo Kussell. Memory and Fitness Optimization of Bacteria under Fluctuating Environments. *PLoS Genetics*, 10(9):e1004556, 2014.
- [24] Stephen R Biggar and Gerald R Crabtree. Cell signaling can direct either binary or graded transcriptional responses. *The EMBO journal*, 20(12):3167–3176, 2001.
- [25] E. Jablonka, B. Oborny, I. Molnár, E. Kisdi, J. Hofbauer, and T. Czárán. The adaptive advantage of phenotypic memory in changing environments. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 350(1332):133–141, 1995.
- [26] J. W. Veening, W. K. Smits, and O. P. Kuipers. Bistability, epigenetics, and bet-hedging in bacteria. *Annu. Rev. Microbiol.*, 62:193–210, 2008.
- [27] C. G. Bowsher and P. S. Swain. Environmental sensing, information transfer, and cellular decision-making. *Curr. Opin. Biotechnol.*, 28:149–155, Aug 2014.
- [28] L. Albenberg, T. V. Esipova, C. P. Judge, K. Bittinger, J. Chen, A. Laughlin, S. Grunberg, R. N. Baldassano, J. D. Lewis, H. Li, S. R. Thom, F. D. Bushman, S. A. Vinogradov, and G. D. Wu. Correlation between intraluminal oxygen gradient and radial partitioning of intestinal microbiota. *Gastroenterology*, 147(5):1055–1063, Nov 2014.
- [29] A. L. Hartman, D. M. Lough, D. K. Barupal, O. Fiehn, T. Fishbein, M. Zasloff, and J. A. Eisen. Human gut microbiome adopts an alternative state following small bowel transplantation. *Proc. Natl. Acad. Sci. U.S.A.*, 106(40):17187–17192, Oct 2009.
- [30] Richard Levins. *Evolution in Changing Environments: Some Theoretical Explorations*. Princeton University Press, 1968.
- [31] J. Wang, E. Atolia, B. Hua, Y. Savir, R. Escalante-Chong, and M. Springer. Natural variation in preparation for nutrient depletion reveals a cost-benefit tradeoff. *PLoS Biol.*, 13(1):e1002041, Jan 2015.
- [32] E.T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- [33] Olivier Cappé, Simon J. Godsill, and Eric Moulines. An Overview of Existing Methods and Recent Advances in Sequential Monte Carlo. *Proceedings of the IEEE*, 95(5):899–924, May 2007.
- [34] Michael I. Jordan. Graphical models. *Statistical science*, 19(1):140–155, 2004.
- [35] R. C. Lewontin and D. Cohen. On population growth in a randomly varying environment. *Proc. Natl. Acad. Sci. U.S.A.*, 62(4):1056–1060, Apr 1969.
- [36] Shun-zheng Yu. Hidden semi-markov models. *Artificial Intelligence*, 2010.
- [37] D. A. Sivak and M. Thomson. Environmental statistics and optimal regulation. *PLoS Comput. Biol.*, 10(9):e1003826, Sep 2014.
- [38] E. Fox, E.B. Sudderth, M.I. Jordan, and A.S. Willsky. Bayesian nonparametric inference of switching dynamic linear models. *Trans. Sig. Proc.*, 59(4):1569–1585, April 2011.
- [39] Matthew J. Johnson and Alan S. Willsky. Bayesian nonparametric hidden semi-markov models. *Journal of Machine Learning Research*, 14(1):673–701, 2013.
- [40] Benjamin Borschinger and Mark Johnson. A Particle Filter algorithm for Bayesian Word segmentation. In *Proceedings of the Australasian Language Technology Association Workshop 2011*, pages 10–18, Canberra, Australia, December 2011.
- [41] Nils E Napp and Ryan P Adams. Message passing inference with chemical reaction networks. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2247–2255. Curran Associates, Inc., 2013.
- [42] T. J. Kobayashi. Implementation of dynamic Bayesian decision making by intracellular kinetics. *Phys. Rev. Lett.*, 104(22):228104, Jun 2010.
- [43] Zoubin Ghahramani and Geoffrey E. Hinton. Variational learning for switching state-space models. *Neural Computation*, 12:963–996, 1998.
- [44] Kevin Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, UC Berkeley, Computer Science Division, 2002.
- [45] Simo Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, New York, NY, USA, 2013.
- [46] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall, London, 1995.
- [47] A. Funahashi, Y. Matsuoka, A. Jouraku, M. Morohashi, N. Kikuchi, and H. Kitano. CellDesigner 3.5: A versatile modeling tool for biochemical networks. *Proceedings of the IEEE*, 96(8):1254–1265, Aug 2008.

- [48] A. Drager, H. Planatscher, D. Motsou Wouamba, A. Schroder, M. Hucka, L. Endler, M. Golebiewski, W. Muller, and A. Zell. SBML2L(A)T(E)X: conversion of SBML files into human-readable reports. *Bioinformatics*, 25(11):1455–1456, Jun 2009.

Supplementary Information

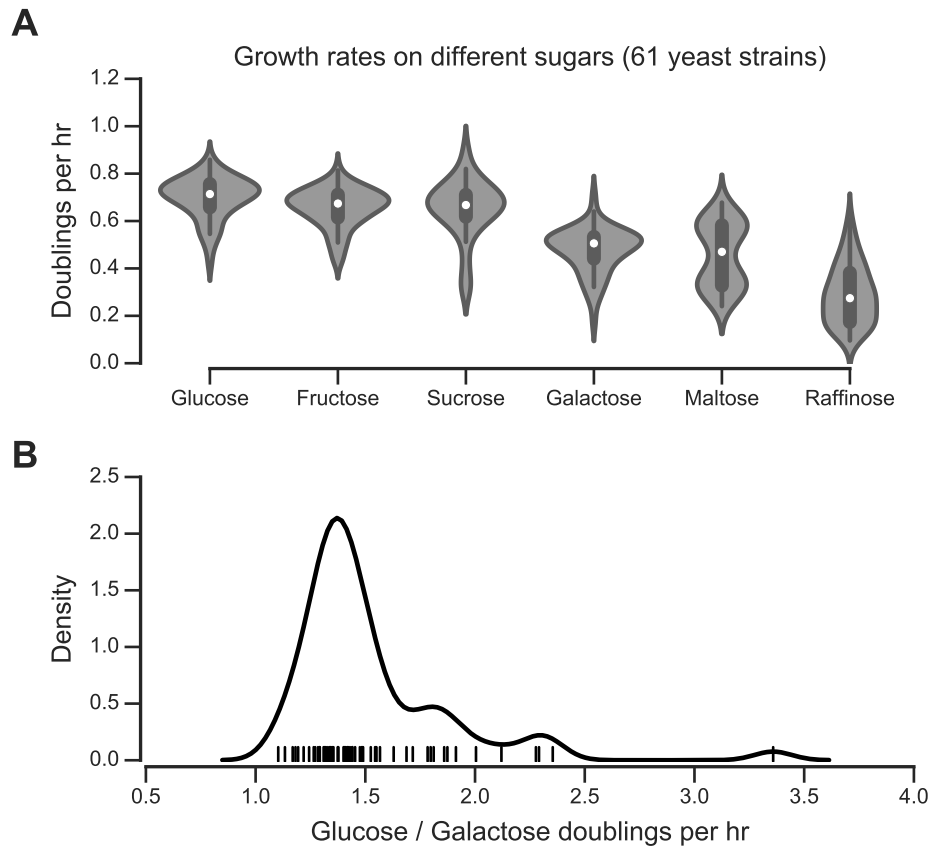


Fig. S1. **Growth rates of 61 yeast strains on different sugar carbon sources.** (A) Data from [31]. Mean growth rates (doublings per hour) from two replicate cultures grown with different sugars as primary carbon source. (B) Distribution of the ratio of glucose to galactose growth rate for 61 yeast strains.

Algorithm S1 Particle filtering algorithm for real-time inference in meta-changing environments.

- 1: initialize N particles $P := p_1, \dots, p_N$ from prior
 - 2: each $p_i = \{s_i, \mathbf{S}_i, \mathbf{C}_i\}$, where s_i is switch state value, \mathbf{S}_i switch transition matrix, \mathbf{C}_i nutrient transition array
 - 3: initialize particle weights $W := w_1, \dots, w_N$ uniformly, $w_i := \frac{1}{N}$
 - 4: **while** next nutrient C_t **do**
 - 5: **for** each particle p_i **do**
 - 6: weigh particle by likelihood of observed nutrient, $w_i := P(C_t | C_{t-1}, S_t = s_i, \mathbf{S}_i, \mathbf{C}_i)$
 - 7: update nutrient transition array \mathbf{C}_i
 - 8: resample particles by weights, $P := \text{RESAMPLE}(P, W)$
 - 9: reset weights, $w_i := \frac{1}{N}$
 - 10: **for** each particle p_i **do**
 - 11: sample new particle state for $t + 1$, $p_i := P(S_{t+1} | S_t, \mathbf{S}_i)$
 - 12: update particle switch transition matrix \mathbf{S}_i
-

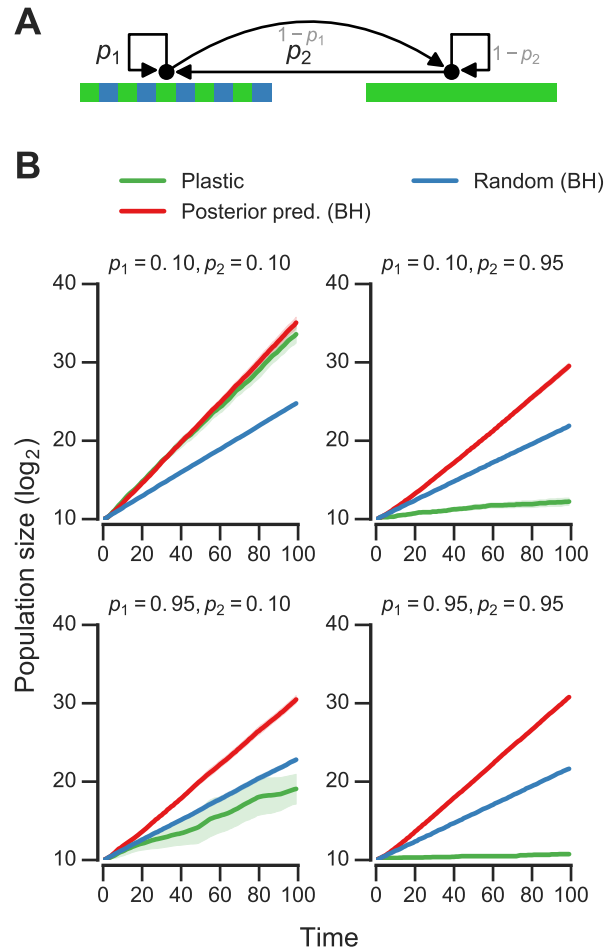


Fig. S3. **Fitness of bet-hedging policies in meta-changing environments.** (A) Meta-changing environment (same as Figure 5A). (B) Population sizes obtained using different growth policies in meta-changing environment shown in (A). “Posterior pred. (BH)” indicates a bet-hedging policy where fraction of population tuned to a nutrient is set by the real-time estimate of the posterior predictive probability of the nutrient, “Random (BH)” indicates a bet-hedging policy where fraction of population tuned to nutrient is set randomly. “Plastic” policy is a non-bet-hedging policy plotted for reference (same as Figure 5B.) Mean population size from 20 simulations plotted with bootstrap confidence intervals (shaded regions).

1	$\text{Glu} \longrightarrow \emptyset$
2	$\text{Gal} \longrightarrow \emptyset$
3	$\text{Glu_Sensor} \longrightarrow \emptyset$
4	$\text{Gal} \longrightarrow \text{Glu_Sensor} + \text{Gal}$
5	$\text{Gal_Sensor} \longrightarrow \emptyset$
6	$\text{Gal_Sensor} + \text{Gal} \rightleftharpoons \text{Gal_Activator}$
7	$\text{Glu} \longrightarrow \text{Gal_Sensor} + \text{Glu}$
8	$\text{Gal_Activator} \longrightarrow \text{Glu_To_Gal} + \text{Gal_Activator}$
9	$\text{Glu_To_Gal} \longrightarrow \emptyset$
10	$\text{Gal_To_Glu} \longrightarrow \emptyset$
11	$\text{Glu_Activator} \longrightarrow \text{Gal_To_Glu} + \text{Glu_Activator}$
12	$\text{Glu_Sensor} + \text{Glu} \rightleftharpoons \text{Glu_Activator}$

Table S1

Reaction equations for molecular implementation of transition counters. Equations with all rate parameters are available in [10.6084/m9.figshare.3492185](https://doi.org/10.6084/m9.figshare.3492185).