Shortsighted Evolution Constrains the Efficacy of Long-Term Bet Hedging

Eric Libby[1],* and William C. Ratcliff[2]

1. Department of Mathematics and Mathematical Statistics and Integrated Science Lab, Umeå University, 907 36 Umeå, Sweden; and Santa Fe Institute, Santa Fe, New Mexico 87501;   2. School of Biological Sciences, Georgia Institute of Technology, Atlanta, Georgia 30332

abstract: To survive unpredictable environmental change, many organisms adopt bet-hedging strategies that are initially costly but provide a long-term fitness benefit. The temporal extent of these deferred fitness benefits determines whether bet-hedging organisms can survive long enough to realize them. In this article, we examine a model of microbial bet hedging in which there are two paths to extinction: unpredictable environmental change and demographic stochasticity. In temporally correlated environments, these drivers of extinction select for different switching strategies. Rapid phenotype switching ensures survival in the face of unpredictable environmental change, while slower-switching organisms become extinct. However, when both switching strategies are present in the same population, then demographic stochasticity—enforced by a limited population size—leads to extinction of the faster-switching organism. As a result, we find a novel form of evolutionary suicide whereby selection in a fluctuating environment can favor bet-hedging strategies that ultimately increase the risk of extinction. Population structures with multiple subpopulations and dispersal can reduce the risk of extinction from unpredictable environmental change and shift the balance so as to facilitate the evolution of slower-switching organisms.

Keywords: bet hedging, extinction, stochastic switching, evolutionary suicide.

Introduction

To survive unpredictable environmental change, some organisms have evolved bet-hedging strategies in which they produce phenotypes maladapted to the current environment in case they may be beneficial at some later stage should the environment change (Cohen 1966; Slatkin 1974; Seger and Brockmann 1987). For example, clonal bacterial lineages have evolved to produce both fast- and slow-growing

phenotypes, where the slower-growing phenotype can better survive lethal antibiotic exposure (Balaban et al. 2004; Keren et al. 2004). At the heart of all bet-hedging strategies is a temporal trade-off in which short-term fitness costs are exchanged for long-term fitness gains (Seger and Brockmann 1987; Kussell et al. 2005; Simons 2009). The time interval over which the long-term benefits of bet hedging are realized depends on many factors, including the frequency and stringency of environmental change (Levins 1962; Bull 1987; Donaldson-Matasci et al. 2008). Environmental change, however, will only rarely be the sole determinant of whether bet hedging is successful. Other processes, such as competition and demographic stochasticity, can act on shorter timescales than environmental fluctuations, causing bet-hedging genotypes to become extinct before their long-term benefits are realized. In this article, we investigate the interplay between two processes—environmental change and demographic stochasticity—and their effects on bet-hedging populations.

Bet-hedging strategies take many forms across a diverse array of organisms (Crump 1981; Tonegawa 1983; Hairston and Munns 1984; Andreadis 1990; Fell 1995; Danforth 1999; Martin 1999; Smits et al. 2006; Ratcliff et al. 2015). For example, desert annuals face considerable across-year variation in spring rainfall that is difficult, if not impossible, to predict accurately. An annual plant that produces seeds that germinate simultaneously risks losing them all should there be a drought. Instead, desert annuals have evolved to delay germination in some seeds as a hedge against across-year variation in rainfall (Cohen 1966; Venable 2007). Another example is the bacterial pathogen *Haemophilus influenzae*, which faces a potentially lethal immune response when infecting a host. Rather than expressing a single antigen and risking elimination on detection, a single clone generates offspring with diverse surface antigens, increasing the probability that some of the population will avoid destruction by the host immune response (Hosking et al. 1999; Bayliss et al. 2001; Moxon and Kussell 2017). Along with the diversity of bet-hedging strategies, organisms have also evolved different molecular

mechanisms to generate phenotypic diversity, including contingency loci, stochastic gene expression, developmental instability, and asymmetric cell division (Moxon et al. 1994; Markow 1995; Kaern et al. 2005; Martins and Locke 2015). These mechanisms work together with reproductive strategies (e.g., asexual/sexual, clutch size) to enact particular forms of bet hedging (Ratcliff et al. 2015). Considering the gamut of bet-hedging strategies is outside the scope of this article. Rather, we focus on microbial bet hedging, which has been well studied and is the subject of a large body of theoretical research (Thattai and van Oudenaarden 2004; Kussell and Leibler 2005; Wolf et al. 2005*a*, 2005*b*; King and Masel 2007; Gaal et al. 2010; Visco et al. 2010).

One way that microbial bet hedging occurs is through stochastic phenotype switching, in which a single genotype produces phenotypic heterogeneity in the absence of an apparent signal or regulatory response (Smits et al. 2006; Beaumont et al. 2009; Libby and Rainey 2011; Rainey et al. 2011; Martins and Locke 2015). Mathematical models of microbial bet hedging typically assume that the organism switches reversibly between at least two distinct phenotypic states and that these distinct phenotypic states are each suited to different possible environmental states. Such models align well with many empirical bet-hedging systems, including $CAP^{+/-}$ phenotypes in *Pseudomonas fluorescens* (Gallie et al. 2015), antigen expression in *Salmonella* (Andrewes 1922; Wildschutte et al. 2004), competence in *Bacillus subtilis* (Maamar et al. 2007), and galactose utilization in engineered populations of *Saccharomyces cerevisiae* (Acar et al. 2008). For the purposes of modeling, phenotypic switching is usually assumed to be random and independent of other factors so that population-level heterogeneity follows a binomial or multinomial distribution depending on the number of phenotypic states.

Since bet hedging is not adaptive in a single environment, all theoretical models feature changing environmental states. However, they differ in their treatment of environmental variation. For example, environmental states can change randomly or after a fixed amount of time (Thattai and van Oudenaarden 2004; Kussell and Leibler 2005; Acar et al. 2008; Salathe et al. 2009; Visco et al. 2010; Müller et al. 2013), and they can be symmetric or asymmetric in terms of how they switch and the selective pressures they exert (Salathe et al. 2009; Gaal et al. 2010). The type of environmental fluctuations determine whether switching is advantageous. If, for example, the environment changes frequently and severely bottlenecks the population—as may occur in the presence of antibiotics—then there is immediate selection against organisms that do not switch (Libby and Rainey 2011). If, instead, the environment changes infrequently or maladapted phenotypes suffer only mild growth-rate reductions, then it may be better not to switch phenotypes (King and Masel 2007; Müller et al. 2013). In

environmental regimes that favor switching, a great deal of work has examined which rates of stochastic switching maximize fitness in terms of the number of descendants. For instance, in exponentially expanding populations, the greatest fitness benefit is afforded by the rate of stochastic switching that maximizes long-term geometric mean fitness by reducing intergenerational variance (Levins 1962; Cohen 1966; Moran 1992; Donaldson-Matasci et al. 2008), a result that can be obtained analytically in certain continuous population models (Kussell and Leibler 2005; Müller et al. 2013).

The role of extinction is a key dividing feature among theoretical models of bet hedging via stochastic phenotype switching. Models in which genotypes do not face the threat of extinction define the fitnesses of different switching rates in terms of long-term population growth rates, where long-term is often equated with the infinite time limit such that arbitrarily many environmental fluctuations have occurred (Kussell and Leibler 2005). Moreover, fitness is typically assessed without regard to transient short-term population dynamics, in which a genotype might have come close to extinction. In contrast, models with extinction may use some way of culling the population, through finite population sizes that limit the total number of organisms or unpredictable environmental catastrophes (King and Masel 2007). Measuring long-term fitness can be challenging because extinction events limit how many genotypes are present in any given specification of long-term; in the infinite time limit, no genotypes may survive. One approach is to link the long-term survival probability to a driver of extinction. For instance, a genotype's probability of surviving an upcoming environmental fluctuation may be proportional to the number of a certain phenotype that it produces beforehand (King and Masel 2007; Libby and Rainey 2011). These models, however, consider only a single driver of extinction, and it is often environmental catastrophe. This leaves an open question: what happens when there are two potential drivers of extinction?

In this article, we examine the temporal dynamics of bet hedging using a model in which genotypes can become extinct in two different ways, via demographic stochasticity and environmental fluctuations. We solve for the rate of stochastic switching that is fittest in terms of survival when each source of extinction acts in isolation and when both are present. Two possible scenarios arise, depending on the model parameters. When environmental risk is uncorrelated or negatively correlated, the rate of switching that optimizes growth rates in the short term also minimizes the risk of extinction in both the short term and long term. Interestingly, when environmental risk is temporally autocorrelated (i.e., once established, the same kind of environmental catastrophes tend to repeat), we find that populations evolve switching rates that minimize short-term extinction risk. More precisely, we find that switch rates

that maximize short-term growth rates leave a population more susceptible to extinction via environmental catastrophes, resulting in an "extinction trap" that limits the long-term efficacy of stochastic bet-hedging strategies.

Methods

Stochastic Disaster Model

We consider a population of genotypes that switch between two phenotypic states, *A* and *B*. The defining characteristic of a genotype is its probability/rate of switching between phenotypes, which we denote as *P*. We assume that the switch occurs stochastically following reproductive events so that each time an *A* or *B* phenotype cell reproduces there is a fixed probability *P* that it yields a cell of the opposite type. The phenotypic states do not differ in any fitness-relevant trait other than susceptibility to risk. This risk manifests in disasters that target either *A* or *B* phenotypes (these are ecological events such as antibiotic exposure, phage exposure, or immune recognition), removing them from the population completely. We simulate the evolution of populations using a discrete-time approach.

In each time step, ecological disaster may strike the population, removing one cellular phenotype. The default for simulations in this article is a disaster probability of 10% (we also explore the effects of varying this probability). If a disaster occurs, we determine which phenotype it targets. A disaster targets the same phenotype as the previous disaster with probability t_c and the alternate phenotype with probability $1 - t_c$ (the first disaster target is random). If $t_c = 0.5$, then the disaster targets a phenotype with no memory of previous targets. For $t_c > 0.5$ there is an increased chance that the targeted phenotype will be the same as before. For $t_c < 0.5$ the opposite is true, and in the extreme case of $t_c = 0$ the targeted phenotype switches every time there is a disaster. In this way, the parameter t_c determines the duration of an environment hostile toward *A* or *B* phenotypes and the temporal correlation of disasters. Thus, the expected number of disasters of the same type occurring in a row is $1/(1 - t_c)$.

Following the potential disaster, there is an opportunity for population turnover whereby a small fraction of the population is randomly chosen for death independent of their *A* or *B* phenotype. We use a default population turnover probability of 10% for most results in this article but also explore the effects of changing this parameter. This random death and the subsequent reproduction step allow standard demographic turnover to have a role in population diversification between disasters.

After the effects of a possible disaster and population turnover, the remaining organisms reproduce until the population is restored to a fixed size, the carrying capacity *N*.

For most results of this article we consider $N = 1,000$. Increasing this parameter increases the duration of competitions, slowing the rate of extinction. Reproduction to carrying capacity occurs through an iterative process whereby organisms are randomly chosen to reproduce according to their relative frequency in the population. The discrete time step ends once the carrying capacity is reached. We simulate the populations until one genotype becomes extinct or a maximum number of time steps occur, here 10^6. Computer simulations were conducted in the programming language Julia and are provided in the supplemental material (available online).[1]

Evolutionary Model

We expand the stochastic disaster model to allow for the evolution of switching probabilities. When each organism reproduces, there is a probability of 10^{-3} that a new mutant will appear with a switching probability chosen from 10^{-5} to 1 according to $10^{\text{unif}(-5,1)}$. Simulations begin with a clonal population with organisms that have a low switch probability, $P = .0001$. The population goes through rounds of death and reproduction in a fixed environment (t_c) for 10^5 iterations or until the population becomes extinct. Computer simulations were conducted in MATLAB, and code is provided in the supplemental material.

Fluctuating Environment Model

We also consider a model of two environmental states E_A and E_B that switch probabilistically. As before, we consider organisms that exhibit one of two phenotypes (*A* or *B*) and switch between phenotypic states as a consequence of reproduction. In each environmental state, there is one phenotype—the favored phenotype—that is more likely to die than reproduce, while the opposite is true for the other phenotype. So the *A* phenotypes have a mean positive population growth rate in one environmental state, say E_A, and a mean negative population growth rate in the other, E_B. Similarly, the *B* phenotypes have a mean positive growth rate in E_B and a mean negative growth rate in E_A. We simulate the model through discrete time steps. During each time step, there is a probability that the environment switches states: either .05 for fast environmental switching or .005 for slow environmental switching. After determining the state of the environment, each organism either reproduces, dies, or does nothing according to some probability. For the simulations in this article, the probability that an organism reproduces is .1. In each environmental state, the favored phenotype dies according to .01 probability while the other

---

1. Code that appears in *The American Naturalist* is provided as a convenience to readers. It has not necessarily been tested as part of peer review.

phenotype dies with probability .25. If the number of re-producing organisms would cause the population to exceed the carrying capacity ($N = 1,000$ in these simulations), then we randomly choose a subset to reproduce so that the population exactly reaches the carrying capacity. Computer simulations were conducted in MATLAB, and code is provided in the supplemental material.

Metapopulation Model

The metapopulation model is an extension of the stochastic disaster model across n subpopulations. Each subpopulation experiences disasters independently from other subpopulations. They all follow the same algorithm as used in the stochastic disaster model with two exceptions. First, there is the possibility of mutation that allows organisms to give rise to mutants with different probabilities of switching between phenotypes. Mutations occur during reproductive events according to a probability u. If a mutation occurs, we randomly choose another probability of switching without bias. For speed of computation, we consider a set of 10 distinct probabilities of switching divided evenly over log space from 10^{-3} to 1. The second difference in the metapopulation model is that we permit migration between subpopulations. At the end of a time step—after every subpopulation has experienced a potential disaster, background turnover, and reproduction—we permit migration. Every organism may migrate to a new subpopulation with probability x. If an organism is chosen to migrate, then we randomly select another subpopulation without bias and add the organism to the population, preserving its phenotypic state. We note that this may temporarily leave some subpopulations slightly above or below carrying capacity. Alternatively, we could have simulated migration after disasters and population turnover, but this would allow a newly migrated organism a chance to reproduce as well, which would make invasion far easier. Instead, by simulating migration at the end of a time step prior to the next round of disasters it makes invasion more difficult and separates migration and reproduction across two time steps. Computer simulations were conducted in MATLAB, and code is provided in the supplemental material.

Results

Extinction via Environmental Challenge

In our model, the probability of switching phenotypes determines whether an organism can survive the challenge of unpredictable ecological disasters. Here we consider genotypes with particular switching probabilities (i.e., switching strategies) in isolation to determine how they fare in surviving recurring disasters. We vary the correlation in disaster type (i.e., targeting either A or B phenotype individuals),

t_c, from 0 to 1 and compute for each value of t_c the number of times out of 1,000 stochastic simulations organisms with a particular probability of switching survive 10^5 rounds of potential disasters, population turnover, and regrowth (see fig. 1). For fast switching probabilities, $P \geq .1$, genotypes rarely become extinct (less than 1% of the time for $P = .1$). In contrast, for slow switching probabilities, $P \leq .025$, all organisms became extinct before the end of 10^5 rounds. For switching probabilities in between these values (e.g., $P = .05$ and $P = .075$), there is a nonmonotonic relationship between t_c and the frequency of extinction. Genotypes are less likely to become extinct at the extremes of t_c (i.e., close to 0 or 1), as compared to intermediate values (e.g., $t_c = 0.65$).

With t_c close to 1, there are long periods of disasters targeting the same phenotype. In the extreme case where $t_c = 1$, disasters never switch the targeted phenotype, so as long as an organism diversified into two phenotypes at least once it would survive. As t_c decreases from 1, the probability of extinction increases because disasters more frequently switch between targeted phenotypes.

At the opposite extreme, with t_c close to 0, disasters frequently switch between targeted phenotypes. While this seems like a challenging environment to survive, it is actually easier than an environment with longer periods of disasters targeting the same phenotype. This is because in our simulations extinction happens when a genotype is phenotypically uniform and there is a disaster targeting that phenotype. For this to occur, a disaster must strike one phenotype, say A, and the surviving phenotype B must fail to
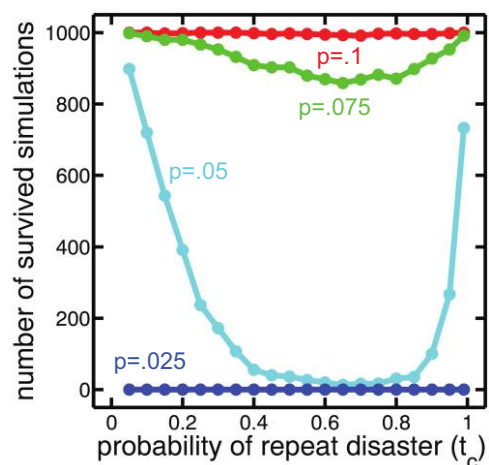

Figure 1: Survival as a function of the probability of switching and the probability of a repeat disaster. The number of simulations (out of 1,000) survived by genotypes with different switch probabilities is shown as a function of the probability that a disaster repeats the phenotype it targets. Fast switchers ($P \geq .1$) rarely become extinct, while slow switchers ($P \leq .025$) always become extinct. Intermediate values of P show a nonmonotonic relationship such that extinction reaches a maximum around $t_c = 0.65$.
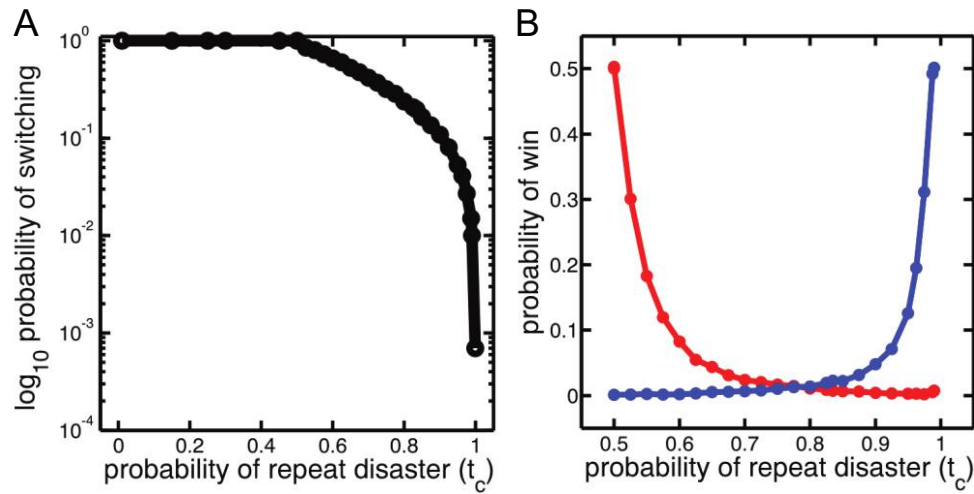
Figure 2: Optimal switching strategy versus the probability that a disaster targets the same phenotype. *A*, The switching probability that beats all others in mixed pair populations is shown as a function of the probability of repeat disasters t_c for $0 \leq t_c \leq 0.99$. The switching probability decreases with increasing values of t_c. *B*, A switch probability of $P = 1$ (red) or $P = .01$ (blue) is paired against the optimal switch probability for a range of t_c values (on the horizontal axis). As measured by the number of wins, both strategies quickly drop in performance by a factor of more than five with a 0.1 change in t_c.

produce an *A* by the time the next disaster strikes. The number of opportunities for a *B* to produce an *A* is determined by two factors: (1) the population growth back to carrying capacity and (2) the population turnover at carrying capacity while waiting for the next disaster. When disasters occur frequently and the background rate of cellular death is low, the dominant force in determining whether the *B* phenotypes will diversify is their distance from carrying capacity—which is determined by the number of *A* annihilated in the previous disaster. When disaster type is correlated in time so that disasters repeatedly target the same phenotype, there is little opportunity for reproduction-driven diversification between disasters, and with each disaster, the frequency of the minority phenotype should decline.[2] The poorly diversified genotype is thus highly susceptible to extinction if the next disaster targets the opposite phenotype (delivering a one-two punch). If, instead, disasters switch targets every time, as is the case with $t_c = 0$, then they are continually targeting the majority of the population and providing significant opportunity for diversification.

Extinction via Demographic Stochasticity

If avoiding extinction from environmental challenges were the only fitness-relevant measure, then the best switching strategy would be to switch as quickly as possible. However,

---

2. This argument relies on the fact that if a population at carrying capacity is hit by a disaster and grows back to carrying capacity, then the phenotype targeted by the disaster always decreases in relative frequency, regardless of the switch rate.

other population processes may have the potential to change the dynamics of selection. For example, a carrying capacity constrains the population size and induces a type of zero-sum competition. A switching strategy that produces phenotypes destroyed by recurring disasters will lose spots in the population to a switching strategy that invests more in phenotypes not currently targeted by disasters. We now examine which switching strategies survive longer in mixed populations constrained by a carrying capacity. We vary the probability of repeat disasters (t_c) from 0 to 1, and for each value we consider pairs of microbial strains with different probabilities of switching sampled between 0 and 1. For each t_c we compute the optimal switching strategy—that is, the one that drives others to extinction more often than it itself becomes extinct—by using pairwise competition experiments where each competition consists of 1,000 simulations. We find that the more likely it is that disasters target the same phenotype, the lower the optimal probability of switching (see fig. 2*A*). If disasters frequently switch between target phenotypes, such that $t_c \leq 0.5$, then the best strategy in paired populations is to rapidly diversify and switch phenotypes often. Thus, in this regime the optimal switching probability is $P = 1.0$.

If, instead, disasters seldom switch the phenotype they target ($t_c \gg 0.5$), then rapid diversification is costly. Consider the case in which a disaster has removed all of the *A* phenotypes. As the *B* phenotypes reproduce to reach the carrying capacity, any *A* types they produce will likely be lost to the next disaster. On the other hand, failing to diversify at all, $P = 0$, will lead to the genotype becoming extinct should the disaster switch the phenotype it targets. When risk is cor-

related in time, the optimal switching strategy must strike a balance between diversifying too much (and producing offspring with a low probability of surviving the next disaster) and failing to diversify and risking extinction. As a point of reference from figure 2*A*, if $t_c = 0.99$, then the optimal switch probability that strikes this balance is $P = .01$. These results in which it is best to switch less often in more stable environments echo earlier studies of bet-hedging populations in the absence of a carrying capacity (Leigh 1970; Seger and Brockmann 1987; Kussell et al. 2005).

Although the optimal switching strategy changes with the temporal correlation of disaster type (t_c), it is unclear how a suboptimal switching strategy fares. To assess this, we picked the best switching strategies for $t_c = 0.5$ ($P = 1.0$) and $t_c = 0.99$ ($P = .01$) and competed them against the optimal switching strategies for a range of t_c values (see fig. 2*B*). The performance quickly drops off such that if either strategy is competed against the optimal strategy at a t_c different by 0.1, it wins less than 10% of the time. Furthermore, $P = 1$ competes as unsuccessfully against $P = .01$ at $t_c = 0.99$ as the reverse when $P = .01$ competes against $P = 1$ at $t_c = 0.5$—they each win less than 1% of the time.

Due to the stochastic nature of these simulations, suboptimal strategies can occasionally beat optimal strategies. To understand what happens in these circumstances, we investigate the pairing of $P = 1$ and $P = .01$ switching strategies when $t_c = 0.99$. On the rare occasions that $P = 1$ wins, the dynamics of disasters show frequent change in the targeted phenotype, mimicking an environment with a lower value of t_c in which $P = 1$ is more adaptive (see fig. 3*A*). In contrast, the more typical scenario is that disasters infrequently switch the target phenotype and thereby penalize strategies that adopt rapid phenotypic diversification (see fig. 3*B*). The trajectory of this extinction shows that each disaster gives an incremental numerical advantage to the slower switching strategy. This acts as a steady drain that ultimately leads the $P = 1$ genotype to extinction.

The different trajectories in panels *A* and *B* of figure 3 demonstrate the two ways in which organisms can become extinct in our model. First, the slow switcher may diversify poorly, allowing the fast switcher to win when a disaster occurs (see fig. 3*C*). Alternatively, repeated disasters of the same type can gradually whittle down the fast switcher's numbers (fig. 3*D*). This is due to the fact that when $t_c > 0.5$, offspring of the opposite phenotype are likely to be killed in the next disaster. The eventual extinction of the fast-switching strain occurs during the background death and replacement phase of our simulations, through random demographic fluctuations. If each organism has a probability α to be chosen for death and replacement, then the probability that a population of m organisms becomes extinct in a single round of replacement is α^m. For this form of extinction to be realized, a genotype must be rare, that is, m must be small.

Although genotypes could become rare randomly through a set of unfortunate replacement events, it usually occurs because of a particular sequence of disasters. For instance, there could be a sequence of sudden switches in the phenotype targeted for disaster that would leave genotypes that switch infrequently as the rare type. Alternatively, when disasters repeatedly target a single phenotype, then the genotype that switches frequently can become rare. The latter case befalls the genotype with $P = 1$ in an environment with $t_c = 0.99$. To illustrate how this happens, consider two genotypes with switching probabilities P_1 and P_2 that just experienced a disaster eliminating all m of one phenotype. In the growth back to carrying capacity, we assume for simplicity that they evenly split the remaining spots in the population, that is, $m/2$ reproductive events are allotted to each genotype after the first disaster. Over k disasters and no replacement other than growth back to carrying capacity, the genotype with a switch probability of P_1 will gain an amount shown in equation (1):

$$(1 - P_1)\frac{m}{2} + (1 - P_1)\frac{m}{2}\frac{(P_1 + P_2)}{2} + \dots$$

$$+ (1 - P_1)\frac{m}{2}\frac{(P_1 + P_2)^k}{2^k} \tag{1}$$

$$= (m/2)(1 - P_1)\frac{1 - [(P_1 + P_2)/2]^k}{1 - [(P_1 + P_2)/2]}.$$

Thus, the genotype with P_1 will have $(1 - P_1)/(1 - P_2)$ as much of the m pool as the genotype with P_2. If one genotype switches with $P = 1$, then the other will eventually get the entire pool of m. This route to rarity is particularly effective if m is close to the carrying capacity N. Once a strain is rare, then background population turnover can lead it to extinction.

Now we weigh the competing pressures of environmental change and demographic forces. In an environment of $t_c = 0.99$, a strategy of $P = .01$ always becomes extinct when considered in isolation (fig. 1, all $P < .025$ become extinct), and yet it is the best strategy in mixed populations (fig. 2). The reason for this seeming contradiction is a significant difference in timescales (see fig. 4*A*). The time it takes $P = .01$ to become extinct from environmental challenges is at least 10-fold greater than the time it takes to drive another organism to extinction in a population of limited size. Thus, while we expect $P = .01$ to become extinct eventually from an environmental change, it has enough time to outlast faster switchers with $P = 1$ who are lost to demographic forces in mixed populations.

The different timescales for environmental challenges and demographic forces can be adjusted by changing parameters such as the rate of background population turnover (see fig. 4*B*) or disaster probability (see supplemental material). By decreasing the probability of background population turnover, we can increase the time it takes for demographic forces to
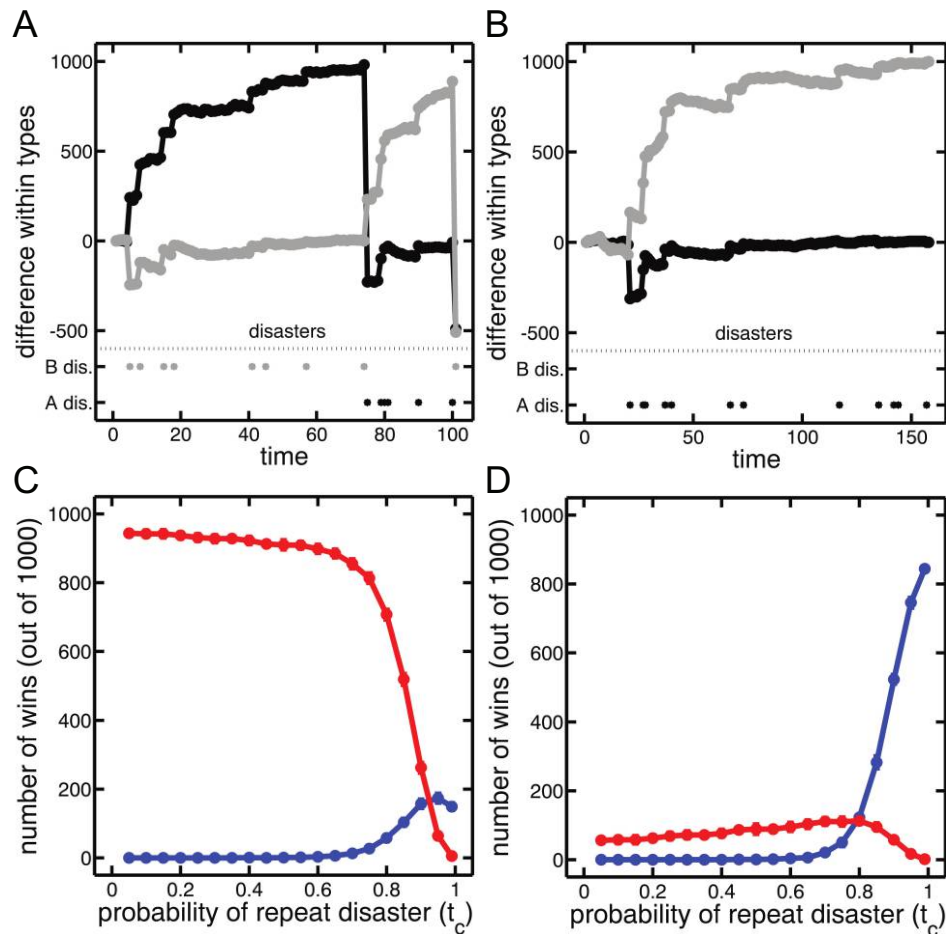
Figure 3: Characteristic manner in which different switching strategies win. *A,* A sample simulation of a mixed population of slow ($P = .01$) and fast ($P = 1$) switching strategies over the course of many disasters in a $t_c = 0.99$ environment. The differences between *A* types (black; i.e., $A_{slow} - A_{fast}$) and *B* types (gray; $B_{slow} - B_{fast}$) are plotted. The phenotype targeted by the disaster is shown at the bottom. The faster-switching phenotype wins because the disasters switch targets and mimic an environment with a lower t_c value. *B,* Another sample simulation similar to *A* shows an example where the slow-switching phenotype wins because disasters repeatedly target the same phenotype and slowly diminish the fast-switching population. *C,* The number of wins (out of 1,000) that occur because a disaster switches targets is shown as a function of the probability of repeat disasters. When the fast switcher ($P = 1$; red) wins, it more often occurs in this manner. *D,* Similar to *C* except the number of wins that occur through background population turnover in between disasters is shown. In comparison to *C,* the slow switcher ($P = .01$; blue) wins more than 80% of its victories in this manner. Thus, fast switchers tend to win when disasters target both *A* and *B* phenotypes in rapid succession, while slow switchers tend to win via a slow draining process.

play out. In addition, the decreased rate of background population turnover results in less diversification, hastening extinction. The net effect is that when risk is highly correlated in time (high t_c) and background population turnover is low, the time frame for extinction by environmental challenge can be shorter than extinction by demographic forces. When this occurs, the slow switcher loses to the fast switcher (see fig. 4*C*).

Evolution of Switching Strategies

So far we have considered populations of only two genotypes in an environment with fixed t_c. To see whether diverse pop-

ulations containing many genotypes can evolve to respond to the selective pressure imposed by the value of t_c, we implement an evolutionary simulation in which genotypes can mutate to give rise to new genotypes with new characteristic switch rates. We begin each simulation with a clonal population with a low probability of switching, $P = .0001$. The population goes through rounds of death and reproduction in a fixed environment (t_c) for 10^5 iterations or until the population becomes extinct. Figure 5*A* shows evolution within a $t_c = 0.99$ environment. The average probability of switching from 1,000 simulations evolves to a switching probability of $P = .01$, which was the optimal switching strategy
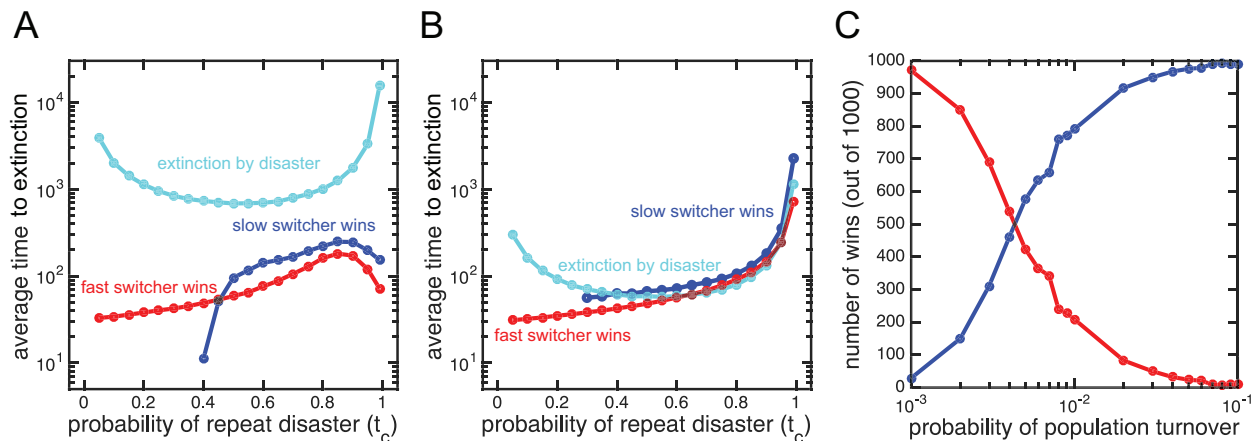
Figure 4: Timescale separation for environmental challenges and demographic processes. *A,* The timescales over which organisms become extinct through environmental challenges or demographic forces are shown as a function of the probability of repeat disasters. Each point is an average of 10,000 simulations. The time for the slow switching strategy ($P = .01$) to become extinct because of disasters (cyan) is at least 10 times longer than it takes either $P = .01$ (blue) or $P = 1$ (red) to win in a mixed population. The fast switching strategy did not become extinct and so is not plotted. *B,* The same as in *A* except that the rate of background population turnover is 100 times lower ($\alpha = 0.001$). The timescale for demographic forces at $t_c = 0.99$ is now longer than the timescale for environmental challenges. This means that the $P = .01$ strategy will become extinct from a disaster before it can win in a mixed population. *C,* The number of wins out of 1,000 for $P = .01$ (blue) or $P = 1$ (red) switchers is shown as a function of background population turnover when $t_c = 0.99$. The lower value of background turnover in *B* is where the fast switchers win more often. As background turnover increases to the value in *A,* the timescale for environmental challenges becomes longer, allowing demographic forces to dominate and the slower switcher to win more frequently.

in paired populations. However, all populations become extinct prior to the end of 10^5 rounds. This sharply contrasts with evolution in a $t_c = 0.5$ environment in which only six of 1,000 simulations become extinct (fig. 5*B*). All other simulations last the entire duration with an average switching probability close to $P = 1$, which was the optimal switching strategy in paired populations.

The speed of adaptation and general results are further confirmed in panels *C* and *D* of figure 5 when populations experience an environmental shift. Figure 5*C* shows the evolution of a population that survived $t_c = 0.99$ (a very rare event) and is subsequently transferred to a different environment with $t_c = 0.5$. In $t_c = 0.99$, the population evolves to the switch probability: $P = .01$. As the environment changes to $t_c = 0.5$, the population adapts by evolving to a switch rate close to $P = 1$. The population remains at a high probability of switching and survives for the rest of the simulation. The reverse environmental fluctuation is shown in figure 5*D*: a population evolving in $t_c = 0.5$ is transferred to an environment with $t_c = 0.99$. In this simulation, the population evolves to $P > .7$ in $t_c = 0.5$. When the environment shifts to $t_c = 0.99$, the population evolves to $P = .01$ and fluctuates before ultimately becoming extinct. In this scenario, the evolutionary response to the mixed population left the winning genotypes far more vulnerable to extinction from unpredictable environmental change.

Alternate Realization of Model

In many previous models of bet hedging, the adaptive benefit of switching phenotypes is determined by fluctuations between discrete environmental states—particularly the frequency at which environmental states switch. Such models often feature two environmental states and organisms that can switch between two phenotypic states, each better suited to a different environmental state (Thattai and van Oudenaarden 2004; Kussell and Leibler 2005; Wolf et al. 2005*a*, 2005*b*; Gaal et al. 2010; Visco et al. 2010). We can frame our model in a similar way with only a few alterations and observe the main findings of our article.

Our model imposes selection for switching through recurring disasters that target either the *A* or the *B* phenotype. In a sense, a set of disasters targeting the same phenotype can be viewed as an environmental state. For example, if a string of disasters target the *B* phenotype, then this could correspond to an environmental state E_A in which the *A* phenotype is favored or protected. The duration of environmental states would then be determined in part by the temporal correlation of disasters, such that larger values of t_c would correspond to longer durations. The key feature of our model that sets it apart from previous models is the dual threat of extinction: from environmental change and from demographic stochasticity. We can incorporate these routes to extinction into a model
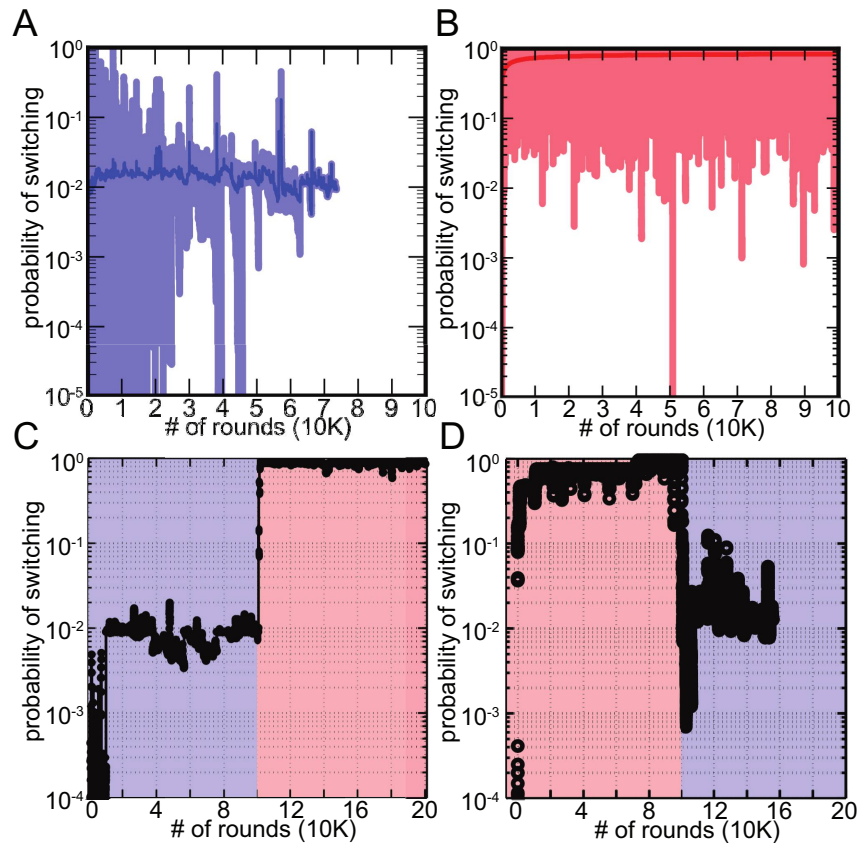
Figure 5: Evolution of switch rates in environments with different probabilities of repeat disasters. *A*, The average probability of switching over 1,000 simulations of evolving populations (dark blue) is plotted over time in an environment with $t_c = 0.99$. The lighter area shows simulated populations removing the top and bottom 5%. Populations quickly evolve to $P = .01$ and then become extinct. *B*, Similar to *A* but with $t_c = 0.5$ and a red scale for coloring. Populations evolve a probability of switching close to $P = 1$, and all but six out of 1,000 survive the duration of the simulation. *C*, Similar to *A* and *B* but environments switch from $t_c = 0.99$ (blue) to $t_c = 0.5$ (red). It took many simulations to find a population that survived $t_c = 0.99$, but once it did, transfer to an environment with $t_c = 0.5$ saw the evolution of higher probabilities of switching close to $P = 1$ (average switch probability shown in black). *D*, Same as *C* but in reverse order. The population average (black) evolves to the optimal switch probability in mixed pair populations $P = 1$ in $t_c = 0.5$ (red) and $P = .01$ in $t_c = 0.99$ (blue) but ultimately becomes extinct in $t_c = 0.99$.

of fluctuating environmental states by making two changes. First, we assume that in each environmental state one phenotype has a net positive growth rate such that on average it reproduces more often than it dies (e.g., *A* in E_A and *B* in E_B) while the other phenotype has a net negative growth rate (*B* in E_A and *A* in E_B). Second, we impose a carrying capacity that limits the total population size and prevents it from potentially escaping the threat of extinction. By making these two changes, we have a model of two fluctuating environmental states that exhibits dynamical behavior similar to our model of recurring disasters.

Figure 6*A* shows a simulation of our two environmental state models (for a description of the implementation, see "Methods"). As the environment changes between states from E_A to E_B, the *B* phenotype increases and the *A* pheno-

type decreases, owing to their respective positive and negative growth rates. We compete organisms with two different switch rates, one fast ($P = 1$) and one slow ($P = .01$), in models with different switch rates between environmental states. If the environmental states switch relatively slowly, with probability .005, then the slow switcher wins more often. If, instead, the environmental states switch frequently, with probability $P = .05$, then the fast switcher wins more often (see fig. 6*B*). As with our model of recurring disasters, the slow switcher is prone to extinction even if it outlasts the fast switcher in a mixed population. Figure 6*C* shows the average duration of each phenotype when it wins in a mixed population. The fast switcher survives the duration of the simulation in 1,049 out of 1,065 instances or 98.5%. In contrast, the slow switcher never survived the duration of the simulation.
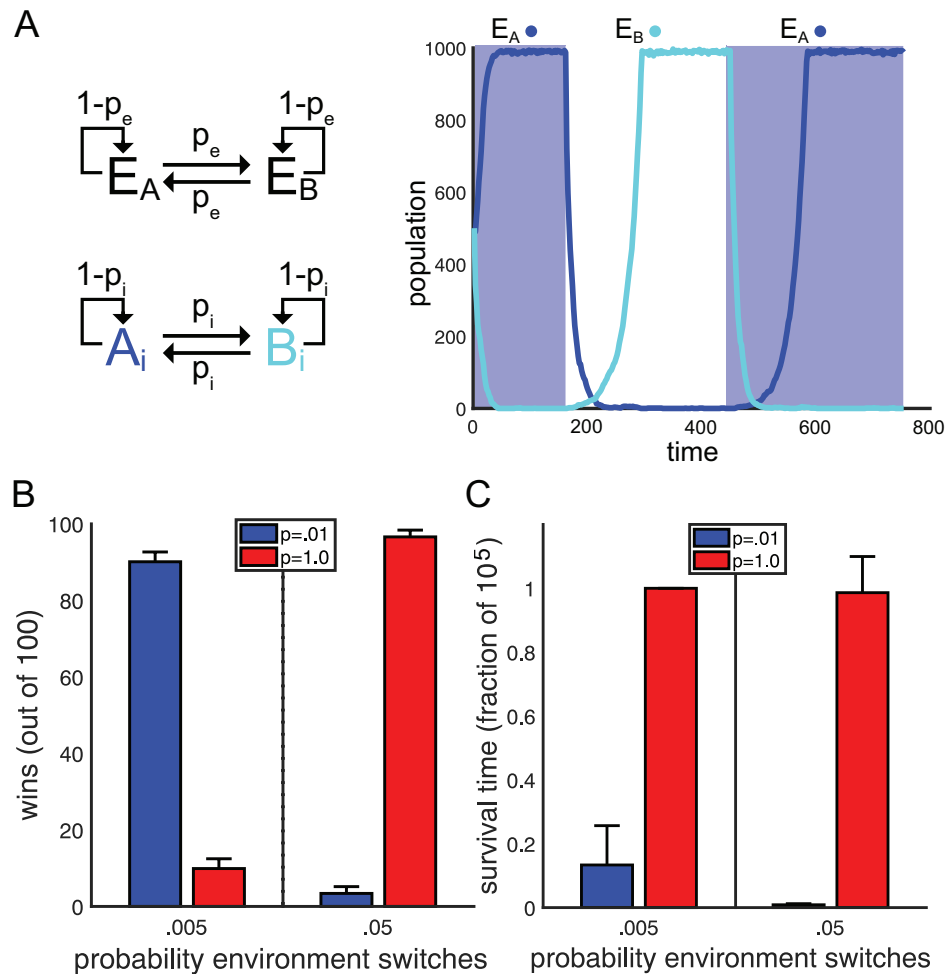
Figure 6: Model with fluctuating environments. *A*, A schematic depicts a model in which both the organisms and the environment switch between two discrete states. The organisms switch between *A* and *B* phenotypes, while the environment switches between E_A and E_B states. Each phenotype has a positive population growth rate in one environmental state (e.g., *A* in E_A) and a negative population growth rate in the other (e.g., *A* in E_B). The simulation shows that when the environment is in the E_A state the *A* phenotypes rise in prevalence while the *B* phenotypes plummet to close to zero. The opposite is true when the environmental state switches to E_B. *B*, The bar graphs show the number of wins of slow ($P = .01$) and fast ($P = 1$) switchers in simulations where environmental states switch with a probability of either .005 or .05. As expected, the slow switchers win more often when the environmental state switches infrequently, while the fast switchers win more often when the environmental state switches frequently. *C*, The bar graphs show the average survival time of switching strategies when they won in *B*. When the slow switcher wins, it becomes extinct before the end of the simulation. In contrast, when the fast switcher wins, it survives the duration of the simulation.

Population Rescue

A $t_c = 0.99$ environment selects for the evolution of switching strategies (e.g., $P = .01$) that ultimately lead the population to extinction. However, if there are many distinct populations with migration between them, it is possible that extinct populations could continually be restored by other surviving populations. This would mean that we could observe the same switching strategy across several populations even though each is heading toward extinction.

To test this hypothesis, we use a metapopulation model composed of *n* subpopulations, each seeded with $P = .01$ switching organisms. For now, we ignore the possibility of mutations. The dynamics of subpopulations evolve independently with the exception of migration of organisms between them (see fig. 7*A* and description in "Methods"). We simulated the metapopulation model for $n = 1, 2, 5$, and 10 subpopulations and migration probabilities of 10^{-3}, 10^{-4}, and 10^{-5}. Similar to our previous analyses, we find that when $n = 1$ a population of $P = .01$ switching strategies always be-
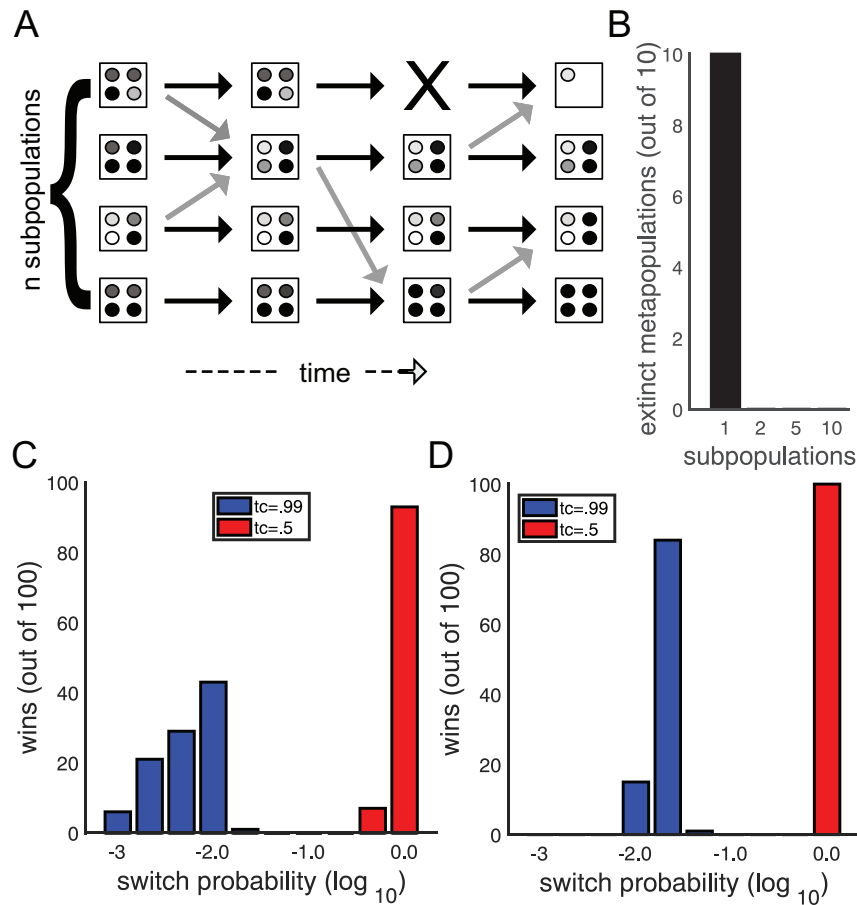
Figure 7: A metapopulation model rescues the slow switcher. *A*, A schematic shows the structure of the metapopulation model. There are *n* subpopulations that each contain distinct populations. At discrete time intervals, organisms can migrate (represented by gray arrows) to new subpopulations. The black arrows indicate the evolution of a subpopulation in which the frequencies of genotypes change and mutations give rise to new genotypes. Occasionally subpopulations can become extinct (represented by the X) but through migration can be seeded by another subpopulation. *B*, The bar graph shows the number of times a metapopulation containing only $P = .01$ switching genotypes becomes extinct. Here the migration probability is 10^{-5} and the mutation probability is 0. All $n = 1$ populations become extinct, but for $n \geq 2$ the metapopulations are able to survive the duration of the simulation (i.e., 10^5 time steps). *C*, The bar graph shows the number of wins for different switching strategies out of 100 runs of the metapopulation model with $n = 10$ and a migration probability of 10^{-3}. For a switching strategy to win, it must be the most abundant switching strategy in all subpopulations or the most abundant in the majority of subpopulations when the simulation reaches 10^5 time steps. The blue bars correspond to a $t_c = 0.99$ environment, and the switching strategy that wins most often is $P = 10^{-2}$, but the majority of winning strategies are less than 10^{-2}. The red bars correspond to a $t_c = 0.5$ environment, and the strategy that wins most often is $P = 1$. *D*, Same as *C* but the migration probability is 10^{-5}. In the $t_c = 0.5$ environment, the strategy that wins most often is still $P = 1$. However, in the $t_c = 0.99$ environment, $P = 10^{-1.66}$ is the strategy that wins most often, and no switching strategy below 10^{-2} wins.

comes extinct. If there are at least two subpopulations ($n \geq 2$) with migration, then it is enough to keep the metapopulation alive for all migration probabilities tested (for migration probability 10^{-5}, see fig. 7*B*).

If we allow subpopulations to evolve via mutation as well as migrate, we find that the migration probability can play a role in determining what switching strategies evolve. Panels *C* and *D* of figure 7 show the most abundant switching strategy for 100 simulations, each composed of 10 subpopulations

for two different migration probabilities (10^{-3} and 10^{-5}). We find that if disasters target the same phenotype according to $t_c = 0.5$, then in both cases of migration probabilities the $P = 1$ switching strategy evolves most often. In a $t_c = 0.99$ environment, we find that the higher migration probability permits the evolution of slower-switching organisms. This is because the frequent migration allows slower-switching organisms to survive in the metapopulation as a whole even though they drive individual subpopulations to extinction.

Discussion

Stochastic phenotype switching is a canonical microbial bet-hedging strategy that can increase fitness and long-term survival in unpredictable environments (Kussell et al. 2005; Smits et al. 2006; Beaumont et al. 2009; Libby and Rainey 2011; Rainey et al. 2011; Martins and Locke 2015). In this article, we describe an evolutionary conflict between the short- and long-term fitness effects of this trait. We study a mathematical model in which unpredictable environmental disasters select for phenotypic diversification. In environments where disaster type is correlated in time (the same cellular phenotype is repeatedly targeted), selection favors diversification strategies that switch infrequently, which allows them to outlast other switching strategies over the short term but ultimately results in their extinction in the long term. This trade-off disappears when the disaster type is uncorrelated in time—here selection favors a switch rate that can outlast other switching genotypes in the short term and survive environmental challenges in the long term.

The stochastic model presented in this article has four key elements: genotypes that switch between phenotypes, correlated disasters that annihilate a specific phenotype, population turnover caused by nondisaster death, and regrowth back to a carrying capacity. These elements are quite general and may appear in many real biological systems. Ecologically, one can think of our model as being a description of free-living bacteria that risk exposure to lytic phages capable of infecting only one of the two phenotypes (Zaleski et al. 2005; Labrie et al. 2010) or bacteria living in a larger organism that risk detection by its immune system. These disasters occur stochastically, but the type of disaster (favoring either *A* or *B* cells) may be correlated in time. Furthermore, population expansion may be limited by the available space and resources found in the environment or host. Interestingly, in the bacteria-host scenario, if the host continually mounts a response against the most abundant bacterial phenotype, then this would be similar to an environment $t_c = 0$ when "disasters" (i.e., immune responses) continually switch between targeted phenotypes. In this case, the survival probability of a switching organism actually increases when compared to environments with intermediate values of t_c.

The evolution of slow switch rates is an example of a phenomenon known as evolutionary suicide (or deterioration) in which adaptation leads a population to extinction (Parvinen 2005; Rankin and López-Sepulcre 2005; Dieckmann and Ferrière 2009). Evolutionary suicide has been studied in a broad range of circumstances, including the evolution of dispersal rates within a structured metapopulation (Gyllenberg and Parvinen 2001; Gyllenberg et al. 2002), increased body size and smaller carrying capacity (Matsuda and Abrams 1994), common goods in a chemostat (Parvinen 2005), altruistic traits (Le Galliard et al. 2003), and aggregation phenomena (Fiegna

and Velicer 2003; Nonaka et al. 2013). Our model describes a novel mechanism through which evolutionary suicide may arise. Interestingly, the route to evolutionary suicide in our model differs considerably from most other examples. In classic models of evolutionary suicide, shrinking population size makes a population susceptible to extinction from environmental or demographic stochasticity. In contrast, reduced population size does not play a role in the evolutionary suicide within our models. In our stochastic model, the population size is restored to the carrying capacity at the end of each time step, so the population size is fixed.

Many prior models studying stochastic switching in fluctuating environments consider continuous populations that grow without limits (Levins 1962; Cohen 1966; Moran 1992; Kussell and Leibler 2005; Donaldson-Matasci et al. 2008). This has two major effects relevant to the study of bet hedging: first, it eliminates the risk of extinction caused by environmental fluctuations, and second, it decouples competing lineages so that the behavior of one lineage has no effect on the fitness of competitors within the same population. We relax both of these constraints and find that imposing a carrying capacity on the population radically changes the ability for natural selection to favor the bet-hedging strategy that maximizes long-term survival. By enforcing a lower limit on the probability of extinction, a carrying capacity allows extinction to play a powerful demographic role in shaping life-history evolution. Without a carrying capacity, populations could expand to the point that extinction is no longer a threat. Carrying capacities and limits to population size in general are a phenomenon we expect to see in many biological systems because they arise whenever there are limited resources.

In addition to affecting survival, carrying capacities are also instrumental in demographic competition. By limiting opportunities for reproduction, a carrying capacity couples the fitness consequences of one strain's switching strategy to its competitors. Specifically, the effects of employing one strategy determines the number of available reproductive events for the other strategy—either in the same round of growth or in a future round. As a result, strains with higher immediate fitness can displace competitors even though they have a lower probability of long-term survival (see fig. 5). These competitive effects should become more influential if one phenotype reproduces more slowly than another. For example, in the case of bacterial persistence (Balaban et al. 2004), successful switching strategies balance a trade-off between a fast-growing/antibiotic-susceptible phenotype and a dormant/antibiotic-resistant phenotype (Kussell et al. 2005). Imposing a carrying capacity on a population of microbes hedging against antibiotic exposure via persistence would create an additional trade-off in which production of dormant cells would limit the number of other cells that could be produced. These trade-offs also exist in systems with regular, predictable environmental change as might be found in experimental populations

of *Pseudomonas fluorescens* (Libby and Rainey 2013). The only requirements are a limit to population growth and organisms that can produce more than one phenotype with different reproductive rates.

Throughout the article, we adopt the view that the bet-hedging strategy consists only of a switch rate that acts independently of other cues or information. However, within our system there are many sources of information that may be useful if incorporated into a bet-hedging strategy. For example, the population composition surviving a disaster provides useful information that could permit the evolution of more adaptive bet-hedging strategies (Leimar 2009; Dall et al. 2015). Moreover, if there were cues preceding an environmental change—perhaps through sensing population density or some other signal—then there is theoretical work showing that this permits the evolution of more regulated mechanisms of phenotype switching (Kussell and Leibler 2005; Donaldson-Matasci et al. 2010; Wolinsky and Libby 2016). Although the consideration of how information can be incorporated into bet-hedging strategies lies outside the scope of this article, we imagine that the evolution of more complex bet-hedging strategies may be able to avoid evolutionary suicide.

Finally, the stochastic phenotype switching considered in this article fits within the hierarchical model of bet hedging advanced by Andrew Simons (Simons 2002): stochastic switching itself is a primary bet-hedging trait that effectively improves fitness in unpredictable, fluctuating environments, while rapid switch rates can be a second-order bet-hedging trait beneficial only over long time periods. With this view, organisms that do not switch phenotypes are quickly driven to extinction by environmental fluctuations. Long-term selection clearly favors switch rates rapid enough to avoid extinction, but this can work against short-term selection for slow switching imposed by competition. We find that depending on the interplay and timescales between these different drivers of selection, isolated populations can become extinct before ever evolving second-order bet-hedging traits. In addition, larger metapopulations can restore extinct subpopulations through migration and further limit the evolution of rapid switching. Our results suggest that by investigating more complex interactions between different drivers of selection and population structures, future studies may uncover a vast array of factors governing the evolution of bet-hedging strategies.

Acknowledgments

Literature Cited

Acar, M., J. T. Mettetal, and A. van Oudenaarden. 2008. Stochastic switching as a survival strategy in fluctuating environments. Nature Genetics 40:471–475. doi:10.1038/ng.110.

Andreadis, T. G. 1990. Observations on installment egg hatching in the brown salt-marsh mosquito, *Aedes cantator*. Journal of the American Mosquito Control Association 6:727–729.

Andrewes, F. W. 1922. Studies in group-agglutination. I. The salmonella group and its antigenic structure. Journal of Pathology and Bacteriology 25:505–521. doi:10.1002/path.1700250411.

Balaban, N. Q., J. Merrin, R. Chait, L. Kowalik, and S. Leibler. 2004. Bacterial persistence as a phenotypic switch. Science 305:1622–1625. doi:10.1126/science.1099390.

Bayliss, C. D., D. Field, and E. R. Moxon. 2001. The simple sequence contingency loci of *Haemophilus influenzae* and *Neisseria meningitidis*. Journal of Clinical Investigation 107:657–666. doi:10.1172/JCI12557.

Beaumont, H. J. E., J. Gallie, C. Kost, G. Ferguson, and P. B. Rainey. 2009. Experimental evolution of bet-hedging. Nature 462:90–93. doi:10.1038/nature08504.

Bull, J. J. 1987. Evolution of phenotypic variance. Evolution 41:303–315. doi:10.2307/2409140.

Cohen, D. 1966. Optimizing reproduction in a randomly varying environment. Journal of Theoretical Biology 12:119–129. doi:10.1016/0022-5193(66)90188-3.

Crump, M. L. 1981. Variation in propagule size as a function of environmental uncertainty for tree frogs. American Naturalist 117:724–737.

Dall, S. R. X., J. M. McNamara, and O. Leimar. 2015. Genes as cues: phenotypic integration of genetic and epigenetic information from a Darwinian perspective. Trends in Ecology and Evolution 30:327–333.

Danforth, B. N. 1999. Emergence dynamics and bet hedging in a desert bee, *Perdita portalis*. Proceedings of the National Academy of Sciences of the USA 266:1985–1994.

Dieckmann, U., and R. Ferrière. 2009. Adaptive dynamics and evolving biodiversity. Pages 188–224 *in* R. Ferrière, U. Dieckmann, and D. Couvet, eds. Evolutionary conservation biology. Cambridge University Press, Cambridge.

Donaldson-Matasci, M. C., C. T. Bergstrom, and M. Lachmann. 2010. The fitness value of information. Oikos 119:219–230.

Donaldson-Matasci, M. C., M. Lachmann, and C. T. Bergstrom. 2008. Phenotypic diversity as an adaptation to environmental uncertainty. Evolutionary Ecology Research 10:493–515.

Fell, P. E. 1995. Deep diapause and the influence of low temperature on the hatching of the gemmules of *Spongilla lacustris* (L.) and *Eunapius fragilis* (Leidy). Invertebrate Biology 114:3–8.

Fiegna, F., and G. J. Velicer. 2003. Competitive fates of bacterial social parasites: persistence and self-induced extinction of *Myxococcus xanthus* cheaters. Proceedings of the Royal Society B 270:1527–1534.

Gaal, B., J. W. Pitchford, and A. J. Wood. 2010. Exact results for the evolution of stochastic switching in variable asymmetric environments. Genetics 184:1113–1119. doi:10.1534/genetics.109.113431.

Gallie, J., E. Libby, F. Bertels, P. Remigi, C. B. Jendresen, G. C. Ferguson, N. Desprat, et al. 2015. Bistability in a metabolic network underpins

the de novo evolution of colony switching in *Pseudomonas fluorescens*. PLoS Biology 13:e1002109. doi:10.1371/journal.pbio.1002109.

Gyllenberg, M., and K. Parvinen. 2001. Necessary and sufficient conditions for evolutionary suicide. Bulletin of Mathematical Biology 63:981–993.

Gyllenberg, M., K. Parvinen, and U. Dieckmann. 2002. Evolutionary suicide and evolution of dispersal in structured metapopulations. Journal of Mathematical Biology 45:79–105.

Hairston, N. G., and W. R. Munns. 1984. The timing of copepod diapause as an evolutionary stable strategy. American Naturalist 123:733–751.

Hosking, S. L., J. E. Craig, and N. J. High. 1999. Phase variation of *lic1A*, *lic2A* and *lic3A* in colonization of the nasopharynx, bloodstream and cerebrospinal fluid by *Haemophilus influenzae* type b. Microbiology 145:3005–3011.

Kaern, M., T. C. Elston, W. J. Blake, and J. J. Collins. 2005. Stochasticity in gene expression: from theories to phenotypes. Nature Reviews Genetics 6:451–464. doi:10.1038/nrg1615.

Keren, I., N. Kaldalu, A. Spoering, Y. Wang, and K. Lewis. 2004. Persister cells and tolerance to antimicrobials. FEMS Microbiology Letters 230:13–18. doi:10.1016/S0378-1097(03)00856-5.

King, O. D., and J. Masel. 2007. The evolution of bet-hedging adaptations to rare scenarios. Theoretical Population Biology 72:560–575. doi:10.1016/j.tpb.2007.08.006.

Kussell, E., R. Kishony, N. Q. Balaban, and S. Leibler. 2005. Bacterial persistence: a model of survival in changing environments. Genetics 169:1807–1814. doi:10.1534/genetics.104.035352.

Kussell, E., and S. Leibler. 2005. Phenotypic diversity, population growth, and information in fluctuating environments. Science 309:2075–2078. doi:10.1126/science.1114383.

Labrie, S. J., J. E. Samson, and S. Moineau. 2010. Bacteriophage resistance mechanisms. Nature Reviews Microbiology 8:317–327.

Le Galliard, J. F., R. Ferrière, and U. Dieckmann. 2003. The adaptive dynamics of altruism in spatially heterogeneous populations. Evolution 57:1–17.

Leigh, E. G., Jr. 1970. Natural selection and mutability. American Naturalist 104:301–305.

Leimar, O. 2009. Environmental and genetic cues in the evolution of phenotypic polymorphism. Evolutionary Ecology 23:125–135.

Levins, R. 1962. Theory of fitness in a heterogeneous environment. I. The fitness set and the adaptive function. American Naturalist 96:361–373. doi:10.1086/282245.

Libby, E., and P. B. Rainey. 2011. Exclusion rules, bottlenecks and the evolution of stochastic phenotype switching. Proceedings of the Royal Society B 278:3574–3583. doi:10.1098/rspb.2011.0146.

———. 2013. Eco-evolutionary feedback and the tuning of proto-developmental life cycles. PLoS ONE 8:e82274.

Maamar, H., A. Raj, and D. Dubnau. 2007. Noise in gene expression determines cell fate in *Bacillus subtilis*. Science 317:526–529. doi:10.1126/science.1140818.

Markow, T. A. 1995. Evolutionary ecology and developmental instability. Annual Review of Entomology 40:105–120.

Martin, K. L. M. 1999. Ready and waiting: delayed hatching and extended incubation of anamniotic vertebrate terrestrial eggs. American Zoologist 39:279–288.

Martins, B. M., and J. C. Locke. 2015. Microbial individuality: how single-cell heterogeneity enables population level strategies. Current Opinion in Microbiology 24:104–112. doi:10.1016/j.mib.2015.01.003.

Matsuda, H., and P. A. Abrams. 1994. Runaway evolution to self-extinction under asymmetrical competition. Evolution 48:1764–1772.

Moran, N. A. 1992. The evolutionary maintenance of alternative phenotypes. American Naturalist 139:971–989.

Moxon, E. R., P. B. Rainey, M. A. Nowak, and R. E. Lenski. 1994. Adaptive evolution of highly mutable loci in pathogenic bacteria. Current Biology 4:24–33. doi:10.1016/S0960-9822(00)00005-1.

Moxon, R., and E. Kussell. 2017. The impact of bottlenecks on microbial survival, adaptation, and phenotypic switching in host-pathogen interactions. Evolution 71:2803–2816.

Müller, J., B. A. Hense, T. M. Fuchs, M. Utz, and C. Pötzsche. 2013. Bet-hedging in stochastically switching environments. Journal of Theoretical Biology 336:144–157. doi:10.1016/j.jtbi.2013.07.017.

Nonaka, E., K. Parvinen, and Å. Brännström. 2013. Evolutionary suicide as a consequence of runaway selection for greater aggregation tendency. Journal of Theoretical Biology 317:96–104.

Parvinen, K. 2005. Evolutionary suicide. Acta Biotheoretica 53:241–264.

Rainey, P. B., H. J. Beaumont, G. C. Ferguson, J. Gallie, C. Kost, E. Libby, and X. X. Zhang. 2011. The evolutionary emergence of stochastic phenotype switching in bacteria. Microbial Cell Factories 10(suppl. 1):S14. doi:10.1186/1475-2859-10-S1-S14.

Rankin, D. J., and A. López-Sepulcre. 2005. Can adaptation lead to extinction? Oikos 111:616–619.

Ratcliff, W. C., P. Hawthorne, and E. Libby. 2015. Courting disaster: how diversification rate affects fitness under risk. Evolution 69:126–135. doi:10.1111/evo.12568.

Salathe, M., J. Van Cleve, and M. W. Feldman. 2009. Evolution of stochastic switching rates in asymmetric fitness landscapes. Genetics 182:1159–1164. doi:10.1534/genetics.109.103333.

Seger, J. B., and H. Brockmann. 1987. What is bet hedging? Pages 182–211 *in* P. Harvey and L. Partridge, eds. Oxford surveys in evolutionary biology. Vol. 4. Oxford University Press, Oxford.

Simons, A. M. 2002. The continuity of microevolution and macroevolution. Journal of Evolutionary Biology 15:688–701.

———. 2009. Fluctuating natural selection accounts for the evolution of diversification bet hedging. Proceedings of the Royal Society B 276:1987–1992.

Slatkin, M. 1974. Hedging one's evolutionary bets. Nature 250:704–705. doi:10.1038/250704b0.

Smits, W. K., O. P. Kuipers, and J.-W. Veening. 2006. Phenotypic variation in bacteria: the role of feedback regulation. Nature Reviews Microbiology 4:259–271. doi:10.1038/nrmicro1381.

Thattai, M., and A. van Oudenaarden. 2004. Stochastic gene expression in fluctuating environments. Genetics 167:523–530. doi:10.1534/genetics.167.1.523.

Tonegawa, S. 1983. Somatic generation of antibody diversity. Nature 302:575–581.

Venable, D. L. 2007. Bet hedging in a guild of desert annuals. Ecology 88:1086–1090.

Visco, P., R. J. Alled, S. N. Majumdar, and M. R. Evans. 2010. Switching and growth for microbial populations in catastrophic responsive environments. Biophysical Journal 98:1099–1108. doi:10.1016/j.bpj.2009.11.049.

Wildschutte, H., D. M. Wolfe, A. Tamewitz, and J. G. Lawrence. 2004. Protozoan predation, diversifying selection, and the evolution of antigenic diversity in *Salmonella*. Proceedings of the National Academy of Sciences of the USA 101:10644–10649. doi:10.1073/pnas.0404028101.

Wolf, D. M., V. V. Vazirani, and A. P. Arkin. 2005a. Diversity in times of adversity: probabilistic strategies in microbial survival games. Journal of Theoretical Biology 234:227–253. doi:10.1016/j.jtbi.2004.11.020.

———. 2005*b*. A microbial modified prisoner's dilemma game: how frequency-dependent selection can lead to random phase variation. Journal of Theoretical Biology 234:255–262. doi:10.1016/j.jtbi.2004.11.021.

Wolinsky, E., and E. Libby. 2016. Evolution of regulated phenotypic expression during a transition to multicellularity. Evolutionary Ecology 30:235–250.

Zaleski, P., M. Wojciechowski, and A. Piekarowicz. 2005. The role of Dam methylation in phase variation of *Haemophilus influenzae* genes involved in defense against phage infection. Microbiology 151:3361–3369.
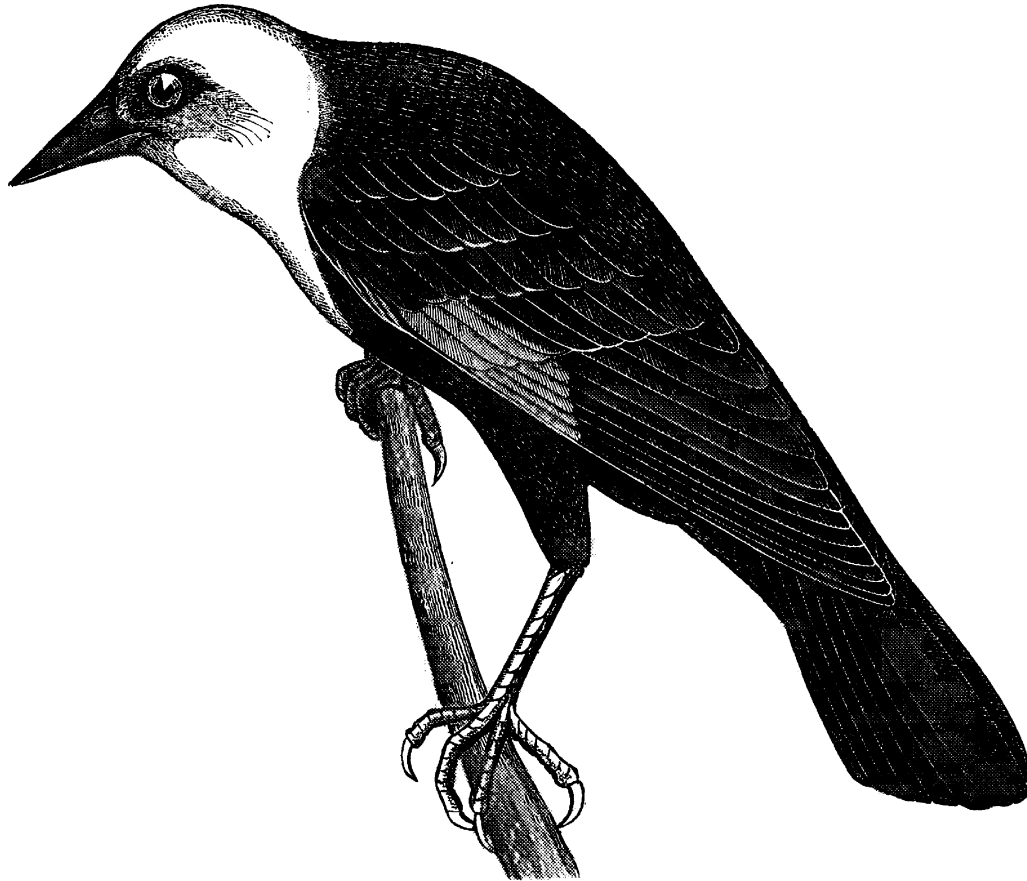
"The *Xanthocephalus icterocephalus*, as this species is called, is one of our handsomest blackbirds, the body being glossy black, the head, neck and breast rich yellow, and the wings having a white spot. The female is smaller than the male, and plain dark brown, with the yellow of the head restricted and clouded over. Together with the Lark Finch (*Chondestes grammaca*) and Prairie Hen, it is one of the earliest indications that the westward bound traveller has of approaching a different ornithological region." From "The Yellow-Headed Blackbird" by Elliott Coues (*The American Naturalist*, 1871, 5:195–200).