Phenotype Switching and Mutations in Random Environments

Drew Fudenberg & Lorens A. Imhof

Bulletin of Mathematical Biology
A Journal Devoted to Research at the Junction of Computational, Theoretical and Experimental Biology Official Journal of The Society for Mathematical Biology

ISSN 0092-8240
Volume 74
Number 2

Bull Math Biol (2012) 74:399-421
DOI 10.1007/s11538-011-9687-8

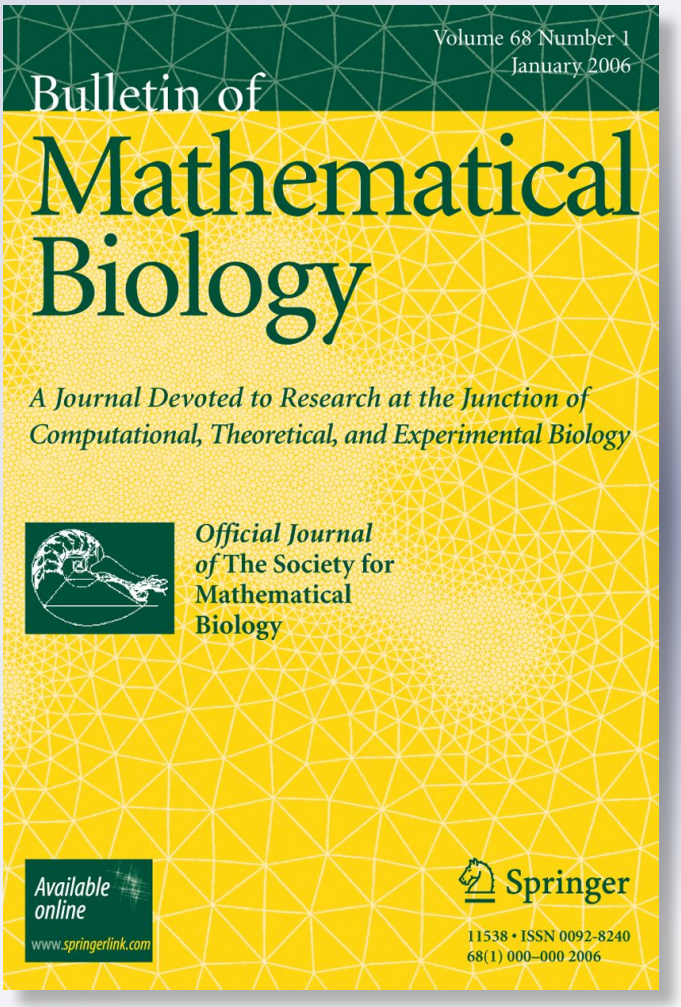

Springer

Your article is protected by copyright and all rights are held exclusively by Society for Mathematical Biology. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

Author's personal copy

Bull Math Biol (2012) 74:399–421
DOI 10.1007/s11538-011-9687-8

ORIGINAL ARTICLE

Phenotype Switching and Mutations in Random Environments

Drew Fudenberg · Lorens A. Imhof

Received: 12 October 2010 / Accepted: 4 August 2011 / Published online: 7 September 2011
© Society for Mathematical Biology 2011

Abstract Cell populations can benefit from changing phenotype when the environment changes. One mechanism for generating these changes is stochastic phenotype switching, whereby cells switch stochastically from one phenotype to another according to genetically determined rates, irrespective of the current environment, with the matching of phenotype to environment then determined by selective pressure. This mechanism has been observed in numerous contexts, but identifying the precise connection between switching rates and environmental changes remains an open problem. Here, we introduce a simple model to study the evolution of phenotype switching in a finite population subject to random environmental shocks. We compare the successes of competing genotypes with different switching rates, and analyze how the optimal switching rates depend on the frequency of environmental changes. If environmental changes are as rare as mutations, then the optimal switching rates mimic the rates of environmental changes. If the environment changes more frequently, then the optimal genotype either maximally favors fitness in the more common environment or has the maximal switching rate to each phenotype. Our results also explain why the optimum is relatively insensitive to fitness in each environment.

Keywords Bet-hedging · Finite population · Fluctuating environment · Optimal switching rate · Stochastic evolution

D. Fudenberg
Department of Economics, Harvard University, Cambridge, MA 02138, USA

L.A. Imhof (✉)
Statistische Abteilung, Hausdorff Zentrum für Mathematik, Universität Bonn, 53113 Bonn, Germany
e-mail: limhof@uni-bonn.de

Author's personal copy

1 Introduction

A major challenge in evolutionary biology is to understand how population-level response to environmental fluctuations evolves (Levins 1968; Tuljapurkar 1990). This is relatively straightforward for a population of organisms that can sense environmental changes and induce an appropriate response in their own phenotype or that of their offspring, but such sensing has costs (DeWitt et al. 1998) and is not always feasible. An alternative mechanism is for the population to maintain phenotypic variation ("bet-hedging"), so that there are always some individuals who are adapted to other environments (Cohen 1966; Gillespie 1974; Slatkin 1974; Philippi and Seger 1989). Another mechanism to induce change is adaptive mutation, where genetic variation occurs as a reaction to a new environment (Rosenberg 2001).

Recent studies direct interest to situations where adaptation by mutation would be too slow. A particularly interesting mechanism to accelerate adaptation is stochastic phenotype switching (Thattai and van Oudenaarden 2004; Avery 2006; Acar et al. 2008): Each cell may switch randomly from one phenotype to another according to genetically encoded switching rates. E.g., genetically identical *E. coli* cells subject to fluctuating antibiotic treatment switch stochastically between an antibiotic-sensitive type with normal growth and a resistant type with reduced growth (Balaban et al. 2004; Kussell et al. 2005). The cells protect themselves by suspending growth. Further examples where a similar strategy has been observed include various bacteria (Maamar et al. 2007; Süel et al. 2007), the yeast prion (True and Lindquist 2000) and viruses (Stumpf et al. 2002).

Here, we introduce a simple finite-population model to study evolutionary selection of competing genotypes with different phenotypic switching rates. The phases where a genotype or phenotype is nearly or completely extinct are crucial for the evolution of the system (Kussell et al. 2005), as here the influence of mutation and switching is most pronounced. Due to these bottlenecks, finite-population effects can be decisive even though the mean size of the population is large. In our finite population model, reproduction, genetic mutations, phenotype switches, and environmental changes are all stochastic. This distinguishes our framework from most work on evolution in changing environments where infinite populations are considered with deterministic cycles (Kimura 1967; Leigh 1970; Ishii et al. 1989; Ben-Porath et al. 1993; Lachmann and Jablonka 1996; Kussell and Leibler 2005; Bürger et al. 2006; Acar et al. 2008) or stochastic cycles (Jablonka et al. 1995; Thattai and van Oudenaarden 2004). In particular, to determine optimal switching rates, standard arguments for infinite populations based on geometric mean fitness (Stumpf et al. 2002) are not suitable for our model (King and Masel 2007). The natural conjecture is that optimal switching rates should closely mimic the rates of environmental changes, but this is not necessarily true. Our results help to identify and understand when this conjecture is correct.

Our analysis rests on stochastic techniques from evolutionary game theory (Foster and Young 1990; Fudenberg and Harris 1992; Kandori et al. 1993; Young 1993; Fudenberg and Imhof 2006), which in turn are inspired by the general study of random perturbations of dynamical systems (Freidlin and Wentzell 1998). These techniques provide more precise predictions in settings where the unperturbed system

Author's personal copy

has multiple local attractors, so that its long-run behavior depends on the initial conditions. Specifically, the perturbations lead to an ergodic system that fluctuates between various attractors, and it is often possible to compare the fraction of the time the system spends at each attractor (Samuelson 1997; Fudenberg and Levine 1998; Imhof et al. 2005; Sandholm 2009). In this paper, we extend the perturbation approach to study phenotype switching. The approach enables us to consider a fairly general specification of the underlying evolutionary process, instead of committing to, e.g., the Moran process, or to particular forms of the fitness function. Moreover, our approach allows for competition between phenotypes and between genotypes in the same population; the importance of this aspect was pointed out by Thattai and van Oudenaarden (2004) and Donaldson-Matasci et al. (2008).

To highlight the effect of environmental changes and switching probabilities, we consider symmetric situations, where the relative advantage of phenotype A in environment E is the same as that of phenotype B in environment F. We also suppose that only the phenotype has a direct effect on fitness and that there are lower and upper bounds on the rate of phenotype switching. The lower bound reflects costs of detecting and correcting errors (Thomas et al. 1998; Ibba and Söll 1999). The upper bound means that phenotype switching cannot be arbitrarily fast compared to the frequency of environmental changes. Throughout, we assume that mutations are rarer than phenotypic switches.

We first consider the case where changes of the environment occur more frequently than mutations. As reported in Corollary 1, we find that if the environment stays on average a long time in one state, say E, and a short time in F, the optimal switching rate to the phenotype superior in E is as high and the optimal switching rate for the other direction is as low as possible. If the environment spends about the same amount of time in each state or remains sufficiently long in each state, the optimal switching rates are maximal for both directions.

For an intuitive explanation, suppose the resident phenotype is well adapted to the environment and that then the environment changes. A high switching rate in the right direction generates a quick adaptation to the new environment, which reduces the risk of being invaded and taken over during the critical phase when the resident phenotype is inferior. The advantage becomes less relevant if the environment only stays in the new state for a short time after a successful adaptation. On the other hand, a high switching rate may cause many detrimental switches, and the disadvantage becomes more relevant the longer the environment remains in the new state. In the case at hand, the critical phases following environmental changes play a decisive role, and it turns out that the positive effect of a high switching rate tends to be stronger than the negative effect. These observations explain why if the environment remains on average a long time in one state, the optimal switching rate to the phenotype superior in that state is maximal, irrespective of how long the environment stays in the other state. They also explain why the optimal switching rates are both maximal when the environment spends about the same amount of time in each state. However, if the environment remains only for a short time in one state, say F, and for a long time in E, the optimal switching rate to the phenotype superior in F is low, because the positive effect of a high switching rate has little time to be felt, but the negative effect is felt for a long time.

Author's personal copy

We find that the optimal switching rates are relatively insensitive to the details of the dynamics and exact values of the fitness parameters for each environment. In our model, the effect of fitness on persistence of phenotypes in a fixed environment without mutations is captured by fixation probabilities. The fixation probability of a given phenotype is the probability that a single individual having this phenotype takes over a population where otherwise only the other phenotype is present. Our analysis shows that the optimal switching rates are constant over wide ranges of the sizes of the fixation probabilities. This result helps to understand why this phenomenon was observed in the simulation results of Kussell et al. (2005) for a related model.

We then consider the case where environmental changes are about as rare as mutations. Here, the optimal switching rate to the phenotype that is superior in the more common environment is again as high as possible. In contrast to the previous case, the phases immediately after environmental changes, when the resident phenotype has not yet adapted, are now negligible compared to the entire periods between changes of the environment. Thus, the optimal switching rate for the other direction is, in general, not as large as possible, but the ratio of the optimal switching rates mimics the ratio of the rates of environmental changes, as reported in Corollary 2. Again, the optimal rates are relatively insensitive to fitness parameters.

In summary, then, we find that in our model there are three kinds of optimal genotypes. The genotype that maximally favors fitness in the more common environment, the genotype that maximizes switching in each direction, and the genotype whose switching rates mirror the rates of environmental changes. The third type is optimal if environmental changes are as rare as mutations, one of the first two types is optimal if the environment changes more frequently.

The paper is organized as follows. The model is described in Sect. 2. In Sect. 3, we first determine for any two competing genotypes which one is favored in the sense of being resident more than half the time. Building on these pairwise comparisons we derive optimal switching rates when rates may vary in an arbitrary interval. In Sect. 4, we present some numerical examples and we discuss related literature and possible extensions. Proofs are given in the Appendix.

2 The Model

We consider an asexual population of fixed size N. The population evolves in an environment that can be in two possible states, E and F. There are four types of individuals, A_1, B_1, A_2, B_2; A and B are phenotypes and the index indicates the genotype. The phenotype of an individual affects its fitness. The fitness of a given individual depends on its own phenotype, on the current frequencies of both phenotypes, and on the environment, but not on genotypes, as only the phenotypes have an immediate influence on reproductive success. The influence of genotypes on reproductive success is only indirect, through realized phenotypes. We are interested in the connection between small switching probabilities and the long-run evolutionary success of corresponding genotypes when the environment is randomly changing.

We first describe the behavior of our model in the absence of genetic mutations, phenotypic switches and environmental changes. Rather than specifying a particular

Author's personal copy

evolutionary process, we assume that for each environment there is a discrete-time Markov chain that describes the evolution of the population without mutation or phenotype switching when the environment is fixed. The common state space of the two chains is the set of possible population compositions. We denote the transition matrices of the chains by Π_E and Π_F. We assume that under Π_E and under Π_F, absent types remain absent, and that starting from any composition of the population, each type initially present has a positive chance of taking over the population. Moreover, because the transition probabilities depend on an underlying fitness function that is independent of genotype, evolution is governed by random drift if a homogeneous population is invaded by a single individual with the same phenotype but the other genotype; in this case, the invader takes over with probability $1/N$. When the invader has a different phenotype, however, evolution need not be governed by random drift because the relative fitness of the phenotypes plays a role. Let ϕ_{AB}^E denote the probability that eventually the whole population has phenotype B given that the environment is E and initially every individual except one has phenotype A and one has phenotype B. As suggested by the notation, we assume that this probability depends only on the phenotypes and not on the genotypes. We define ϕ_{BA}^E, ϕ_{AB}^F and ϕ_{BA}^F analogously. At this point we have imposed few restrictions on the evolutionary process; the key assumptions are that the fitness of each type depends only on its phenotype and the environment, but not on the genotype, and that every offspring inherits both the genotype and the phenotype of its parent. Phenotype here refers to a viable form of the overall organism. Errors in DNA replication can lower fitness and even lead to a nonviable organism, and processes that correct these errors can also have a fitness cost. This corresponds to a link between genotype and fitness that we have abstracted away in this model. Below we represent these links with upper and lower bounds on the rate of phenotype switching.

We now use these two chains to construct a process with environmental changes as well as phenotypic switches and genetic mutations. The state of the system is the current composition of the population and the current environment. We assume that at every time step, a transition takes place according to the relevant matrix Π_E or Π_F, and then a random shock may occur. This shock has the following three possible forms:

(i) *Environmental change:* If the environment is in state E, it changes to F with probability $\epsilon\eta_{EF}$. If the environment is in state F, it changes to E with probability $\epsilon\eta_{FE}$.
(ii) *Phenotypic switching:* An individual is chosen at random, each being equally likely. If it is of type A_i, it switches to B_i with probability $\epsilon\beta_i$; if it is of type B_i, it switches to A_i with probability $\epsilon\alpha_i$.
(iii) *Genotypic mutation:* With probability $2\epsilon\mu$ a randomly chosen individual is replaced by an individual with the other genotype, each phenotype occurring with probability $\epsilon\mu$.

We assume that $\eta_{EF}, \eta_{FE}, \alpha_i, \beta_i$, and μ are positive and that the shocks are independent across time and type of shock.

The process so obtained is a Markov chain with two time scales, where $\epsilon > 0$ is a small parameter that quantifies the separation of the time scales. On the fast time

Author's personal copy

scale the evolution of the population is given by the transition matrices Π_E and Π_F. As long as the environment is in state E, the population evolves according to Π_E and after a geometrically distributed amount of time with expectation of order $1/\epsilon$, a shock occurs: Either the environment changes, or an individual switches its phenotype, or a mutation takes place. The relative frequencies of these rare events are determined by the parameters $\eta_{EF}, \eta_{FE}, \alpha_i, \beta_i, \mu$; where the parameters α_i and β_i control the probabilities of switches by genotype i to phenotypes A and B, respectively. For now, the relation between these parameters is arbitrary, but in our results we will assume that μ is much smaller than α_i and β_i so that genotype mutations are rare compared to phenotype switches.

Note that while the environment can influence fitness, and thus the population, the population does not influence the environment. The lengths of time between consecutive environmental changes are independent and have a geometric distribution. The mean duration is $1/(\epsilon\eta_{EF})$ in environment E and $1/(\epsilon\eta_{FE})$ in F.

Invariant Distribution Genetic mutations and environmental changes ensure that the evolutionary process is irreducible and so has a unique invariant distribution for each $\epsilon > 0$. This distribution describes the long-run behavior of the process; it specifies for each state the proportion of time that the process spends there. A direct analysis of the invariant distribution seems impractical in most cases of interest. Even if Π_E and Π_F are sparse matrices, the complexity of a numerical computation of the invariant distribution seems prohibitive unless N is small, because the size of the state space is of order N^3 (King and Masel 2007). However, we are mostly interested in the behavior of the system when ϵ becomes small, so that we can take advantage of the separation of time scales; this is the essence of the perturbation approach of Freidlin and Wentzell (1998).

Let π denote the limit of the invariant distribution as $\epsilon \to 0$. When ϵ is small, the process spends nearly all the time at states where all individuals are of the same type. We denote these states by $(ALL\, A_1, E), \ldots, (ALL\, B_2, F)$; e.g. $(ALL\, A_1, E)$ is the state where the whole population is of type A_1 and the environment is in state E. By extending the arguments of Fudenberg and Imhof (2006), we show that π concentrates on these eight states and that the probabilities $\pi(ALL\, A_1, E), \ldots, \pi(ALL\, B_2, F)$ can be determined by computing the unique invariant probability vector of a reduced process that moves only between these states. Moreover, the transition probabilities for this reduced process correspond to the idea that transitions between the eight states arise only when a shock occurs; these transition probabilities are computed conditional on some shock occurring, so they are scaled up by $1/\epsilon$. The reduced process shifts from one environment to the other with probabilities η_{EF} and η_{FE}. For the phenotype or genotype to change, there must be a mutation that attains fixation before the next shock; for example the probability that the reduced process assigns to a transition from $(ALL\, A_1, E)$ to $(ALL\, B_1, E)$ is the conditional probability β_1 of a mutation from A_1 to B_1 multiplied by the probability ϕ_{AB}^E that the B-phenotype achieves fixation before the next random event. (Note that the absorption probability ϕ_{AB}^E is for a Markov chain whose state space is only of size $N+1$ as opposed to order N^3. When the chains Π_E and Π_F are birth–death processes, as with Moran, the absorption probabilities are given by a simple recurrence formula.)

Author's personal copy

The same heuristic underlies the definition of the other elements of the transition matrix M of the 8-state process. Our results show that if $\lambda = (\lambda_1, \ldots, \lambda_8)$ is the unique vector such that $\lambda M = \lambda$ and $\lambda_1 + \cdots + \lambda_8 = 1$, then

$$\pi(ALL\,A_1, E) = \lambda_1, \quad \pi(ALL\,B_1, E) = \lambda_2, \quad \ldots, \quad \pi(ALL\,B_2, F) = \lambda_8.$$

The formal description of M and proof of this and our other results are given in the Appendix.

An immediate consequence of the representation of π as an eigenvector of M is that the long-run success of each type depends on the given transition matrices Π_E and Π_F only through the fixation probabilities ϕ but is otherwise unaffected by the details of the dynamics in each environment. This result explains in particular a corresponding invariance property observed in simulations of a closely related model with deterministic environmental changes (Kussell et al. 2005). The representation is also a key ingredient for our analysis of optimal switching rates, where it turns out that the optimal rates are constant over wide ranges of the sizes of the fixation probabilities.

3 Results

To compare the evolutionary success of the genotypes let $ALL\,G_i$ denote the set of states where all individuals have genotype i. We say that genotype 1 is favored under rare genotype mutations if $\pi(ALL\,G_1) > \pi(ALL\,G_2)$ for all $\mu > 0$ sufficiently small. Note that $\pi(ALL\,G_1) + \pi(ALL\,G_2) = 1$ and

$$\pi(ALL\,G_i) = \pi(ALL\,A_i, E) + \pi(ALL\,B_i, E) + \pi(ALL\,A_i, F) + \pi(ALL\,B_i, F),$$

showing that the favored genotype can be determined from the invariant probability vector λ of M.

Note also that the system spends most of its time in states where only one genotype is present; the favored one is prevalent more than half of the time. Finally, note that this is a qualitative condition, and does not pin down the share of the favored genotype; we present some numerical examples in the discussion.

Throughout this section, we consider symmetric environments, where

$$\phi_{AB}^E = \phi_{BA}^F, \qquad \phi_{BA}^E = \phi_{AB}^F. \tag{1}$$

We assume symmetry in the fitness functions and environments to highlight the role of the probabilities η of environmental change, but our technical apparatus easily applies to asymmetric situations.

We also assume that in environment E phenotype A has an advantage over B in the sense that

$$\Phi_{BA}^E := \lim_{N\to\infty} \phi_{BA}^E > 0, \qquad \lim_{N\to\infty} \phi_{AB}^E = 0, \tag{2}$$

where N is the population size. This condition is quite weak and is satisfied by many processes, as shown by the approximation results of Fudenberg and Imhof (2008). For

Author's personal copy

example, these conditions are satisfied by the frequency-independent Moran process if in environment E, phenotype A has fitness a and B has fitness b with $a > b$, and in environment F, A has fitness b and B has fitness a; here $\Phi_{BA}^{E} = 1 - (b/a)$. The conditions are also satisfied by the frequency-dependent Moran process (Fudenberg et al. 2006) obtained from a 2×2 hawk-dove game. Moreover, all our results hold under the weaker condition that $\phi_{AB}^{E} = o((\phi_{BA}^{E})^3)$ as $N \to \infty$, even if ϕ_{BA}^{E} is not bounded away from zero. For any given game and imitation process, this condition is easily checked using the asymptotic results of Fudenberg and Imhof (2008). By "imitation process" we mean a process under which absent types remain absent when there are no mutations or switches. Some further technical conditions must be satisfied for the asymptotic results to hold.

3.1 Environmental Changes of the Order of Phenotypic Switches

Our first result is for the case where the frequency of environmental changes is of the same order as the frequency of phenotypic switches, and both are much more common than genotype mutations. That is, we keep the other parameters fixed and examine the behavior of the limit distribution π when μ is small. We simplify the analysis by considering large populations. We expect, however, that our conclusions will typically hold for small or medium populations as well, because the fixation probabilities, on which our asymptotic arguments are based, usually converge exponentially quickly (Fudenberg and Imhof 2008). Numerical simulations support this conjecture; we discuss some of them in the discussion.

Theorem 1 *Consider competing genotypes $i = 1, 2$ with switching parameters α_i, β_i. Suppose the fixation probabilities ϕ satisfy (1) and (2). Suppose the probabilities of environmental changes are $\epsilon\eta_{EF}$ and $\epsilon\eta_{FE}$, so that environmental changes are of the order of phenotypic switches.*

If

$$\left(\frac{1}{\alpha_2} + \frac{1}{\beta_2} - \frac{1}{\alpha_1} - \frac{1}{\beta_1}\right)\Phi_{BA}^{E} > (\eta_{EF} - \eta_{FE})\left(\frac{1}{\alpha_2\beta_1} - \frac{1}{\alpha_1\beta_2}\right), \tag{3}$$

then genotype 1 is favored under rare genetic mutations, provided the population is sufficiently large. That is, there exist $N_0 < \infty$ and $\mu_0(N) > 0$ such that $\pi(ALL\,G_1) > \pi(ALL\,G_2)$ if $N \geq N_0$ and $\mu < \mu_0(N)$.

If the reverse of inequality (3) holds, then genotype 2 is favored under rare genetic mutations, provided the population is sufficiently large.

Theorem 1 allows us to determine the optimal switching behavior in a given set of competing genotypes. It is costly to have too high a rate of switching; it is also costly to have too low a rate of switching because detecting and correcting errors is costly. We therefore consider switching parameters α and β in a given interval $[s, S]$ with $0 < s < S$. We say that $(\alpha_1, \beta_1) \in [s, S]^2$ is optimal if the corresponding genotype is favored over every competitor, that is, if inequality (3) holds for all $(\alpha_2, \beta_2) \in [s, S]^2$ with $(\alpha_2, \beta_2) \neq (\alpha_1, \beta_1)$.

If the environment is usually in state E, where phenotype A is better than B, one may expect that a successful genotype should have a high switching rate from B to

Author's personal copy

A and a low one from A to B. It turns out that the first part of the guess is correct and that the second part holds under some restrictions. A simple computation based on Theorem 1 leads to the following classification of optimal switching parameters for the case where genetic mutations are rare compared to environmental changes and phenotypic switches.

Corollary 1 *Suppose* (1) *and* (2) *hold. Suppose the probabilities of environmental changes are* $\epsilon\eta_{EF}$ *and* $\epsilon\eta_{FE}$, *so that environmental changes are of the order of phenotypic switches. Then the optimal pair of switching parameters in* $[s, S]$ *is given as follows*:

(a) *If* $\eta_{EF} - \eta_{FE} < -S\Phi_{BA}^{E}$, *then* (S, s) *is optimal.*
(b) *If* $|\eta_{EF} - \eta_{FE}| < S\Phi_{BA}^{E}$, *then* (S, S) *is optimal.*
(c) *If* $\eta_{EF} - \eta_{FE} > S\Phi_{BA}^{E}$, *then* (s, S) *is optimal.*

Case (a), and by symmetry, case (c) agree with the conclusion from the heuristic argument. Recall that η_{EF} and η_{FE} determine the probabilities that the state of the environment changes from E to F or from F to E, respectively, and that Φ_{BA}^{E} can be regarded as a measure of the advantage of the better adapted phenotype. Thus, if $\eta_{EF} - \eta_{FE} < -S\Phi_{BA}^{E}$, the environment remains on average longer in state E than in F, and the optimal switching parameter for switches from phenotype B to A is indeed maximal and the other one is minimal.

Note, however, that the classification depends on the difference $\eta_{EF} - \eta_{FE}$, not on the ratio η_{EF}/η_{FE}, that is, the ratio of the mean lengths of stays in the two environments. If both η_{EF} and η_{FE} are small compared to $S\Phi_{BA}^{E}$, the inequality in (b) is satisfied irrespective of the ratio η_{EF}/η_{FE}. To understand why in this case both optimal switching parameters are maximal even if the environment stays on average much longer in one state than in the other, note that the long-run success of a genotype depends critically on its ability to resist invasion by the competing genotype. To see how the risk of such an invasion can be minimized suppose genotype 1 has reached fixation. As we consider the case where μ, the parameter determining the mutation rate, tends to 0 and as η_{EF} and η_{FE} are small, we can assume that before a mutant having genotype 2 occurs, there will be many environmental changes and between any two of them there can be many phenotypic switches. If the environment is in state E, phenotype A is superior and, therefore, A_1 will soon become the resident type. Then there can be occasional switches to B_1, but type B_1 is very likely to disappear quickly as long as the environment remains in state E. When the environment changes to F, the resident phenotype is still A, which is now inferior. By phenotypic switching, an individual of type B_1 will occur and this type takes over. Again, the new resident type B_1 is nearly proof against invasion by A_1 as long as the environment stays in state F. Thus, the population evolves typically in cycles as follows: the resident phenotype is the one that is best in the current environment. When the environment changes, there is a phase when the resident phenotype is inferior; then the superior phenotype takes over and persists until the environment changes again. Only during the phases immediately after the environment changes does the genotype have a high risk of being replaced; the expected lengths of these phases are about $1/(\epsilon\alpha_1)$

Author's personal copy

and $1/(\epsilon\beta_1)$, that is, they are inversely proportional to the switching parameters α_1 and β_1. This shows explicitly the way in which the ability to adapt quickly to a new environment contributes to success. The optimal genotype in case (b) minimizes the risk of being invaded and replaced by minimizing the lengths of the risky periods. This means of protection seems reasonable independent of the lengths of the ensuing less risky phases and so independent of the relative amount of time spent in each environment. On the other hand, large switching parameters have a detrimental effect after the initial risky periods, but the effect is small because an occasionally occurring inferior phenotype is likely to die out quickly. Moreover, a second factor that could have an influence on the long-run success of genotype 1, the ability to invade and take over a population with resident genotype 2, is hardly affected by α_1 and β_1.

A look at the risky and less risky phases is also helpful in understanding cases (a) and (c). In case (a), $\eta_{FE} > S\Phi^E_{BA} > 0$ while η_{EF} can be arbitrarily small. The mean of the time the environment remains in state F is at most $1/(\epsilon S\Phi^E_{BA})$. Even for a genotype whose switching parameter β is maximal, the length of the risky phase after an environmental change to F would still be comparable to the length of the whole stay in F; the low-risk phase in F would be relatively short or would not be reached at all. For the optimal genotype, β is minimal, which suggests that the small loss from not minimizing the length of the risky phase in F is more than compensated for by the gain in environment E.

3.2 Environmental Changes of the Order of Genotype Mutations

We now turn to the case where the frequency of environmental changes is of the same order as the frequency of genetic mutations, with both less likely than phenotypic switches. We assume that the other probabilities are as described before, while the probabilities of the environmental changes are now $\epsilon\mu\eta^*_{EF}$ and $\epsilon\mu\eta^*_{FE}$, where η^*_{EF} and η^*_{FE} are fixed positive numbers, and as before μ will become arbitrarily small. That is, we assume that η_{EF} and η_{FE} are coupled to μ according to

$$\eta_{EF} = \mu\eta^*_{EF}, \qquad \eta_{FE} = \mu\eta^*_{FE}, \tag{4}$$

so that the ratio $\eta_{EF}/\eta_{FE} = \eta^*_{EF}/\eta^*_{FE}$ is fixed. In this case, we say that genotype 1 is favored under rare environmental changes and rare genetic mutations if $\pi(ALL\,G_1) > \pi(ALL\,G_2)$ for all μ sufficiently small.

Theorem 2 *Consider competing genotypes $i = 1, 2$ with switching parameters α_i, β_i. Suppose the fixation probabilities ϕ satisfy (1) and (2). Suppose (4) holds, so that environmental changes are of the order of genotype mutations.*

(a) *If*

$$\left(\frac{\alpha_1}{\beta_1} - \frac{\alpha_2}{\beta_2}\right)\left(\frac{\eta^*_{FE}}{\eta^*_{EF}} - \frac{\alpha_1}{\beta_1}\frac{\alpha_2}{\beta_2}\right) > 0, \tag{5}$$

then genotype 1 is favored under rare environmental changes and rare genetic mutations, provided the population is sufficiently large. That is, there exist $N_0 < \infty$ and $\mu_0(N) > 0$ such that $\pi(ALL\,G_1) > \pi(ALL\,G_2)$ if $N \geq N_0$ and $\mu < \mu_0(N)$.

Author's personal copy

If the reverse of inequality (5) *holds, then genotype* 2 *is favored under rare environmental changes and rare genetic mutations, provided the population is sufficiently large.*

(b) *If* $\alpha_2 = t\alpha_1$ *and* $\beta_2 = t\beta_1$ *with* $0 < t < 1$, *then genotype* 1 *is favored under rare environmental changes and rare genetic mutations, provided the population is sufficiently large.*

To derive from Theorem 2 the optimal switching behavior for rare environmental changes consider for a fixed pair of switching parameters α_1, β_1, the interval

$$I = \left[\min\left(\frac{\alpha_1}{\beta_1}, \frac{\eta^*_{FE}}{\eta^*_{EF}}\frac{\beta_1}{\alpha_1}\right), \max\left(\frac{\alpha_1}{\beta_1}, \frac{\eta^*_{FE}}{\eta^*_{EF}}\frac{\beta_1}{\alpha_1}\right)\right].$$

Then, in the sense of Theorem 2, genotype 1 is favored over every competing genotype 2 with switching parameter ratio $\alpha_2/\beta_2 \notin I$, whereas genotype 2 is favored if α_2/β_2 is an interior point of I. If $\alpha_1/\beta_1 = \sqrt{\eta^*_{FE}/\eta^*_{EF}}$, then interval I degenerates and has no interior point, so that genotype 1 is favored over every competitor with $\alpha_2/\beta_2 \neq \alpha_1/\beta_1$. Thus, if environmental changes are rare, then the ratio $\alpha_1/\beta_1 = \sqrt{\eta^*_{FE}/\eta^*_{EF}}$ is optimal. By Theorem 2(b), of two genotypes with this ratio, the one with larger switching rates is favored. We have thus obtained the following characterization of the genotype that is optimally adapted to rare environmental changes.

Corollary 2 *Suppose* (1), (2), *and* (4) *hold, so that environmental changes are of the order of genotype mutations. Then the optimal pair of switching parameters in* $(0, S]$ *is*

$$\left(S, S\sqrt{\frac{\eta^*_{EF}}{\eta^*_{FE}}}\right) \quad \text{if } \eta^*_{FE} > \eta^*_{EF},$$

$$\left(S\sqrt{\frac{\eta^*_{FE}}{\eta^*_{EF}}}, S\right) \quad \text{if } \eta^*_{FE} \leq \eta^*_{EF}.$$

It is interesting to compare this result with the corresponding classification of optimal genotypes for less rare environmental changes. Again, for the optimal genotype the switching parameter controlling switches to the phenotype more efficient in the environment visited longer is always maximal. The other switching parameter now depends on the ratio η^*_{EF}/η^*_{FE} and approaches the maximal value when $\eta^*_{EF}/\eta^*_{FE} \to 1$. The ratio of the optimal switching parameters simply reflects the ratio of the times spent in the environments. To understand the difference from the classification derived earlier consider the situation where a genotype has reached fixation and has to resist invasion by a competitor. When the frequency of environmental changes is of the order of the frequency of phenotypic switches, the initial phases when the resident phenotype has not yet adapted to a new environment constitute a significant part of the overall times between changes in the environment. Controlling the risk of being invaded and replaced during these phases is crucial to the success of

Author's personal copy

Table 1 Shares of genotypes. The dynamics are given by a Moran process in symmetric environments, where in environment E, phenotype A has fitness a and B has fitness b. In case I, environmental changes are on the order of phenotypic switches ($\eta_{EF} = 1.0$, $\eta_{FE} = 0.1$). In case II, environmental changes are on the order of genetic mutations ($\eta_{EF} = \mu$, $\eta_{FE} = 0.1\mu$). The results show that the fitness ratio has a relatively small impact on the shares of the genotypes and that our large population results remain valid for populations of moderate size

N	a/b	μ	Case I $\pi(ALL\, G_1)$	Case II $\pi(ALL\, G_1)$
100	2	10^{-3}	0.6161	0.4083
		10^{-4}	0.6155	0.4832
		10^{-5}	0.6155	0.4982
100	100	10^{-3}	0.6165	0.4004
		10^{-4}	0.6158	0.4830
		10^{-5}	0.6158	0.4982
1000	2	10^{-3}	0.6327	0.2552
		10^{-4}	0.6320	0.3844
		10^{-5}	0.6319	0.4826
1000	100	10^{-3}	0.6252	0.2434
		10^{-4}	0.6245	0.3834
		10^{-5}	0.6244	0.4826

a genotype, the risk during the remaining time plays only a small role. In the present setting where the environment changes much more slowly, the lengths of the initial risky phases are negligible compared to the overall times between changes in the environment. Now it is the proportion of time spent in each environment that determines the best switching parameters.

4 Discussion

4.1 Numerical Examples

The general results are qualitative and apply to the limit of small μ. Table 1 presents numerical calculations of the limit distribution π for the frequency-independent Moran process with $(\alpha_1, \beta_1) = (0.01, 1.0)$ and $(\alpha_2, \beta_2) = (0.1, 0.1)$. In case I, $\eta_{EF} = 1.0$ and $\eta_{FE} = 0.1$. Here, by Theorem 1, genotype 1 is favored under rare genotypic mutations when the population is sufficiently large. In case II, $\eta^*_{EF} = 1.0$ and $\eta^*_{FE} = 0.1$ with $\eta = \eta^*\mu$ in each environment. Here, by Theorem 2, genotype 2 is favored under rare environmental changes. In case I, the share of the favored genotype is between 0.6 and 0.65, even when the fitness advantage a/b of matching the environment is large. In case II, the share of the favored genotype ranges from slightly above 0.5 to 0.76. The share is again relatively insensitive to the fitness advantage a/b, but in case II, the genetic mutation rate and the population size play a larger role than in case I. The observed insensitivity to the value of a/b can be explained by the facts that the shares are determined by the fixation probabilities ϕ and that here ϕ^E_{AB} converges to 0 exponentially quickly (Fudenberg and Imhof 2008), so that ϕ^E_{AB} is negligible compared to $1/N$ and ϕ^E_{BA} unless a/b is close to 1.

To visualize the optimality result obtained from Theorem 1, consider genotypes with switching parameters in $[s, S] = [e^{-8}, 1]$. Under the Moran process with $a/b = 2$ and $\eta_{EF} = 1$, $\eta_{FE} = 0.01$, the optimal parameters are $\alpha_1 = e^{-8}$, $\beta_1 = 1$. Figure 1

Author's personal copy

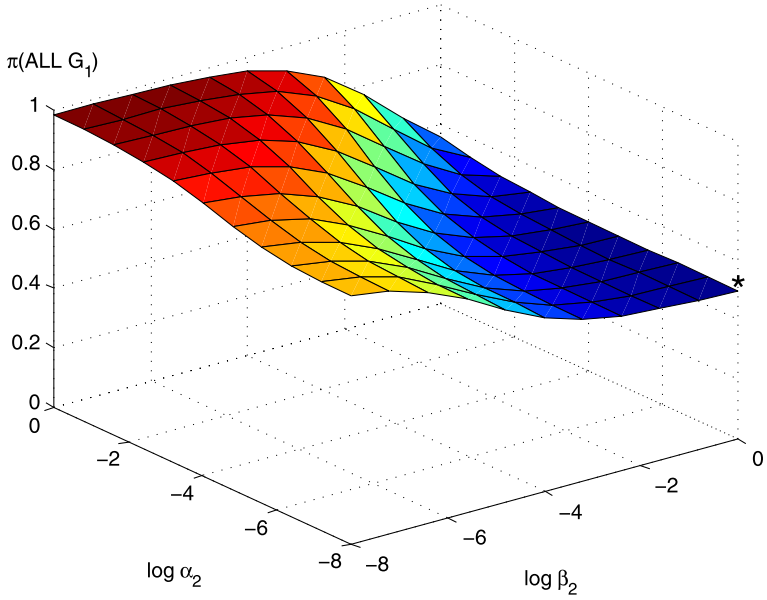

Fig. 1 Share of the optimal genotype. The dynamics are given by a Moran process in symmetric environments, where in environment E, phenotype A has fitness $a = 2$ and B has fitness $b = 1$. The frequencies of environmental changes are determined by $\eta_{EF} = 1$ and $\eta_{FE} = 0.01$, so that the environment is most of the time in state F, where phenotype B is superior. The optimal genotype among all genotypes with switching parameters in $[s, S] = [e^{-8}, 1]$ is given by $\alpha_1 = e^{-8}$ and $\beta_1 = 1$, which corresponds to switching slowly from B to A and relatively quickly from A to B. The graph of the share $\pi(ALL\, G_1)$ shows the extent to which the optimal genotype outperforms all the competitors with switching parameters in $[s, S]$. The asterisk indicates the position of the optimal genotype. The population size is $N = 1000$ and the genotypic mutation rate is determined by $\mu = 0.01$

displays the share of the optimal genotype against every possible competitor. The share is nearly 1 for ill-adapted genotypes.

Our analysis considers the limit π of the invariant distributions as ϵ goes to 0. Table 2 compares π to the invariant distribution ν_ϵ in an example for particular values of ϵ. In this example, the approximation by π seems to be fairly accurate for $\epsilon \leq 10^{-4}$. A theoretical investigation of the quality of this type of approximation in a related model has recently been carried out by Wu et al. (2011).

4.2 Related Literature

Most work (Kimura 1967; Leigh 1970; Ishii et al. 1989; Ben-Porath et al. 1993; Lachmann and Jablonka 1996; Kussell and Leibler 2005; Bürger et al. 2006; Acar et al. 2008) on evolution in changing environments considers infinite populations and deterministic cycles; Jablonka et al. (1995) and Thattai and van Oudenaarden (2004) consider an infinite population and both deterministic and stochastic cycles. King and Masel (2007) consider finite populations governed by a Moran process when phenotype B has zero fitness in environment E and environment F is absorbing; they obtain exact results under the additional assumption that in E, any new-born

Author's personal copy

Table 2 Invariant distribution ν_ϵ and limit distribution π. The dynamics are given by a Moran process in symmetric environments, where in environment E, phenotype A has fitness $a = 1.5$ and B has fitness $b = 1$. The population size is $N = 20$. The invariant distribution ν_ϵ describes the long-run behavior of a Markov chain that has, in this example, about 3,500 states. For the homogeneous states ξ, the limit distribution $\pi(\xi) = \lim_{\epsilon \to 0} \nu_\epsilon(\xi)$ can be calculated as an eigenvector of an 8×8 matrix, see the lemma in the Appendix. Parameter values are $\eta_{EF} = 1.0$, $\eta_{FE} = 0.5$, $\alpha_1 = 1.0$, $\beta_1 = 0.5$, $\alpha_2 = 0.4$, $\beta_2 = 0.8$, and $\mu = 0.1$

State ξ	$\nu_\epsilon(\xi)$			$\pi(\xi)$
	$\epsilon = 10^{-3}$	$\epsilon = 10^{-4}$	$\epsilon = 10^{-5}$	
$(ALL\,A_1, E)$	0.0829	0.0904	0.0911	0.0912
$(ALL\,B_1, E)$	0.0651	0.0669	0.0670	0.0671
$(ALL\,A_2, E)$	0.0455	0.0505	0.0511	0.0511
$(ALL\,B_2, E)$	0.1217	0.1237	0.1239	0.1239
$(ALL\,A_1, F)$	0.1194	0.1288	0.1298	0.1299
$(ALL\,B_1, F)$	0.1762	0.1821	0.1827	0.1827
$(ALL\,A_2, F)$	0.0582	0.0637	0.0643	0.0643
$(ALL\,B_2, F)$	0.2835	0.2891	0.2896	0.2897
$\sum$	0.9526	0.9952	0.9995	1.0000

B's are immediately replaced. By building on the results of Fudenberg and Imhof (2006, 2008), we obtain analytic results without these simplifications. Kussell et al. (2005) report finite-population simulations of a model that is roughly similar to ours. Perhaps the chief difference is that instead of a fixed population size, the population is rescaled according to a Poisson distribution whenever it exceeds a fixed threshold. If the total population size exceeds the limit when there is exactly one mutant wild-type cell, this cell dies in the next step with probability greater than 0.36; with the fixed population rescaling of this paper, the cell would die only with probability of order $1/N$, so that invaders are more likely to be successful. A second difference is that in Kussell et al. (2005) the time scale of the process depends on the fitness of the prevailing phenotypes, while in ours it does not.

Livnat et al. (2005) present an evolutionary analysis for the trade-off between offspring number and quality. The environment is fixed. They compare spenders who lay large clutches and investors who lay smaller clutches but invest more per offspring. Thus, investors sacrifice current success to invest in future generations and so correspond, in our model, to a genotype with high switching rates. Livnat et al. show for an infinite population that the optimal strategy depends on the mutation rate and that investors are more abundant for low mutation rates. This is in rough agreement with our findings. Consider the symmetric case where $\eta_{EF} = \eta_{FE}$. Then in the settings of Corollaries 1 and 2, individuals that invest in future generations by choosing large switching rates are more abundant, provided the mutation rate is small and the population is large.

4.3 Possible Extensions

Our assumption that every mutant has either phenotype with the same probability can be relaxed. For example, instead of assumption (iii) on mutation probabilities we

Author's personal copy

may assume that with probability $\epsilon(\mu+\mu')$ a randomly chosen individual is replaced by an individual with the other genotype, where the phenotype remains unchanged with probability $\epsilon\mu$ and is changed with probability $\epsilon\mu'$, where $\mu > 0$ and $\mu' > 0$. Theorem 1 holds under this assumption as well, provided the condition $\mu < \mu_0(N)$ is replaced by $\max(\mu, \mu') < \mu_0(N)$, see Remark A.1(c) in the Appendix.

In ongoing research, we show that the conclusions of this paper apply more generally. In particular it appears that Theorem 2(a) holds if $\eta_{EF} = \mu^k \eta^*_{EF}$ and $\eta_{FE} = \mu^k \eta^*_{FE}$ for some $k > 1$; this allows for genotypic mutations to be more frequent than environmental changes but less frequent than phenotypic switches. We can also relate phenotypic switching to Parrondo's paradox (Harmer and Abbott 2002): switching between two losing gambles (phenotypes) can lead to a winning genotype.

In Corollaries 1 and 2, we assumed constraints for the switching parameters to ensure existence of an optimum. One might think that an upper bound would not be required because too high a switching parameter would lead to too many switches in the wrong direction and so would rule out optimality. To see why this intuition is misleading consider the symmetric case where $\eta_{EF} = \eta_{FE} = 1$ and consider genotype 1 with $\alpha_1 = \beta_1 = 1$ and genotype 2 with $\alpha_2 = \beta_2$ very large. Suppose A_2 has reached fixation and is well-adapted to the current environment. Often, an individual will switch to the inferior type B_2, but in each instance B_2 is unlikely to take over. If B_2 does become the resident type, a switch to A_2 will occur soon, and A_2 is then likely to take over again. Thus, even if $\alpha_2 = \beta_2$ is very large, this does not mean that the fraction of time where genotype 2 is vulnerable to invasion is close to 1. It only means that there will be many transitions from $ALL\, A_2$ to $ALL\, B_2$ and back. Moreover, since our results apply to the limit distribution obtained as ϵ tends to 0, switching will be a rare event compared to reproduction, even for very high switching parameters. It is an open problem to obtain optimal switching rates in our finite-population stochastic setting without assuming constraints.

Acknowledgements We thank the anonymous referees, Carl Bergstrom, Ken Wachter, and conference participants at GRIPS, Harvard University, and the Max Planck Institute for Evolutionary Biology in Plön for helpful comments and the NSF grant 0646816 for financial support.

Appendix

Here, we present proofs of the claims in the main text.

The common state space of the family of Markov chains we study is given by

$$\mathcal{S} = \big\{(x_1, x_2, x_3, x_4, u) \in \{0, 1, \ldots, N\}^4 \times \{E, F\} : x_1 + \cdots + x_4 = N\big\},$$

where $x_1, \ldots, x_4$ denote the number of individuals of type A_1, B_1, A_2, B_2, respectively, and u is the state of the environment. For example, in terms of this notation, the state $(ALL\, A_1, E)$ is denoted by $(N, 0, 0, 0, E)$. We denote the transition probabilities by $p_\epsilon(\xi, \xi')$, $\epsilon \geq 0$, $\xi, \xi' \in \mathcal{S}$. If $\epsilon = 0$, the environment does not change and mutations and switches do not occur. For $(x, u), (x', u') \in \mathcal{S}$,

$$p_0\big((x, u), (x', u')\big) = \begin{cases} \Pi_u(x, x') & \text{if } u = u', \\ 0 & \text{if } u \neq u', \end{cases}$$

Author's personal copy

where $\Pi_u(x, x')$ are the entries of the given transition matrices Π_u, $u = E, F$.

For $\epsilon > 0$, the set of recurrent states is irreducible and the chain has therefore a unique invariant distribution ν_ϵ, say. The following lemma shows that π, the weak limit of ν_ϵ as $\epsilon \to 0$, can be determined by computing the invariant probability vector of the block matrix

$$M = \begin{bmatrix} M_E & \eta_{EF} I_4 \\ \eta_{FE} I_4 & M_F \end{bmatrix},$$

where I_4 is the 4×4 unit matrix,

$$M_u = \begin{pmatrix} m^u_{11} & \beta_1 \phi^u_{AB} & \mu/N & \mu\phi^u_{AB} \\ \alpha_1 \phi^u_{BA} & m^u_{22} & \mu\phi^u_{BA} & \mu/N \\ \mu/N & \mu\phi^u_{AB} & m^u_{33} & \beta_2 \phi^u_{AB} \\ \mu\phi^u_{BA} & \mu/N & \alpha_2 \phi^u_{BA} & m^u_{44} \end{pmatrix}, \quad u = E, F,$$

and the diagonal entries of M_E and M_F are such that the row sums of the block matrix M are 1. Note that M is irreducible, so that there is unique vector $\lambda = (\lambda_1, \ldots, \lambda_8)$ such that $\lambda M = \lambda$ and $\lambda_1 + \cdots + \lambda_8 = 1$. We call λ the invariant probability vector of M, even though M is not necessarily a stochastic matrix since the diagonal elements can be negative. However, rescaling ϵ if necessary we can assume without loss of generality that M is a stochastic matrix. In this case, λ represents the invariant distribution of the reduced process with state space $\mathcal{H} = \{(ALL\, A_1, E), \ldots, (ALL\, B_2, F)\}$, the set of states where the population is homogeneous.

Lemma *For every $\xi \in \mathcal{S}$, the limit $\pi(\xi) = \lim_{\epsilon \to 0} \nu_\epsilon(\xi)$ exists. For every $\xi \in \mathcal{S} \setminus \mathcal{H}$, $\pi(\xi) = 0$, and if λ is the invariant probability vector of M, then*

$$\begin{aligned} \pi(ALL\, A_1, E) &= \lambda_1, & \pi(ALL\, B_1, E) &= \lambda_2, \\ \pi(ALL\, A_2, E) &= \lambda_3, & \pi(ALL\, B_2, E) &= \lambda_4, \\ \pi(ALL\, A_1, F) &= \lambda_5, & \pi(ALL\, B_1, F) &= \lambda_6, \\ \pi(ALL\, A_2, F) &= \lambda_7, & \pi(ALL\, B_2, F) &= \lambda_8. \end{aligned}$$

Proof Let $\tilde{\mathcal{S}}$ denote the set of states that are recurrent under p_ϵ if $\epsilon > 0$. The set $\tilde{\mathcal{S}}$ is the same for all $\epsilon > 0$ and $\mathcal{H} \subset \tilde{\mathcal{S}}$. Moreover, if $\xi \in \mathcal{S} \setminus \tilde{\mathcal{S}}$, then $\nu_\epsilon(\xi) = 0$, so that $\pi(\xi) = \lim_{\epsilon \to 0} \nu_\epsilon(\xi) = 0$.

To determine $\pi(\xi)$ for the remaining states, we apply Theorem 2 of Fudenberg and Imhof (2006). We consider in the rest of the proof only the restricted state space $\tilde{\mathcal{S}}$. Note that for each $\epsilon > 0$, the chain restricted to $\tilde{\mathcal{S}}$ is irreducible and its invariant distribution is given by $\{\nu_\epsilon(s) : s \in \tilde{\mathcal{S}}\}$. We now verify Assumptions 6 to 9 of Fudenberg and Imhof (2006). Under p_0, the states in $\mathcal{H}$ are absorbing and all other states in $\tilde{\mathcal{S}}$ are transient. Thus, Assumptions 6 and 7 are met. Assumption 8 requires that the limit $\lim_{\epsilon \to 0} p_\epsilon(\xi, \xi')/\epsilon$ exists for all $\xi \in \mathcal{H}$ and all $\xi' \in \tilde{\mathcal{S}}$ with $\xi' \neq \xi$. We verify this only for $\xi = (ALL\, A_1, E)$, the argument for the other elements of $\mathcal{H}$ is similar. If the current state is $\xi = (ALL\, A_1, E)$, then the set of states that can be reached in a single step with a probability of order ϵ is given by

$$\mathcal{S}_1 = \big\{(N-1, 1, 0, 0, E), (N-1, 0, 1, 0, E), (N-1, 0, 0, 1, E), (N, 0, 0, 0, F)\big\}.$$

Author's personal copy

Specifically, for a step to $(N-1,1,0,0,E)$, one individual must switch from A_1 to B_1, which has probability $\epsilon\beta_1$, and there must not occur a genetic mutation nor an environmental change. Hence,

$$\lim_{\epsilon\to 0}\frac{p_\epsilon((ALL\,A_1,E),(N-1,1,0,0,E))}{\epsilon}=\beta_1$$

and similarly,

$$\lim_{\epsilon\to 0}\frac{p_\epsilon((ALL\,A_1,E),(N-1,0,1,0,E))}{\epsilon}=\mu,$$
$$\lim_{\epsilon\to 0}\frac{p_\epsilon((ALL\,A_1,E),(N-1,0,0,1,E))}{\epsilon}=\mu,$$
$$\lim_{\epsilon\to 0}\frac{p_\epsilon((ALL\,A_1,E),(N,0,0,0,F))}{\epsilon}=\eta_{EF}.$$

For $\xi=(ALL\,A_1,E)$ and all $\xi'\in\tilde{\mathcal{S}}\setminus[\mathcal{S}_1\cup\{\xi\}]$, $p_\epsilon(\xi,\xi')=o(\epsilon)$. Assumption 8 is thus seen to hold.

To see that matrix M corresponds to the transition matrix Λ defined in Sect. 4 of Fudenberg and Imhof (2006) note that for each state in $\mathcal{S}_1$ there is, under p_0, exactly one state in $\mathcal{H}\setminus\{(ALL\,A_1,E)\}$ that can be reached from it, namely

$$(ALL\,B_1,E),(ALL\,A_2,E),(ALL\,B_2,E),(ALL\,A_1,F),$$

respectively, and the corresponding absorption probabilities are ϕ_{AB}^E, $1/N$, ϕ_{AB}^E and 1. This shows that the first row of M agrees with the definition of Λ and the other rows can dealt with in much the same way. Our assumptions that η_{EF}, η_{FE}, μ and all the fixation probabilities ϕ are positive ensure that M is irreducible, so that Assumption 9 is satisfied as well. The remaining part of the assertion now follows from Theorem 2 in Fudenberg and Imhof (2006). □

From now on, we consider symmetric environments satisfying (1) and we suppose that

$$\phi_{AB}^E=o\big(\big(\phi_{BA}^E\big)^2\big)\quad(N\to\infty).\tag{6}$$

This condition is slightly weaker than the corresponding condition in the main text where we require that $\phi_{AB}^E=o((\phi_{BA}^E)^3)$. The stronger condition will only be needed in Theorem A.2(b).

Genotypes are characterized by the parameters α_i and β_i. In all the pairwise comparisons of the long-run successes of competing genotypes, we keep these parameters fixed and study $\pi(ALL\,G_i)$ when μ becomes small. This means that in our results genetic mutations are rare compared to phenotypic switches. The following theorem treats the case where the frequency of environmental changes is of the same order as the frequency of phenotypic switches. Here, the parameters η_{EF} and η_{FE} are fixed positive numbers.

Author's personal copy

Theorem A.1 *If*

$$\liminf_{N\to\infty}\left(\frac{1}{\alpha_2}+\frac{1}{\beta_2}-\frac{1}{\alpha_1}-\frac{1}{\beta_1}\right)\phi^E_{BA} > (\eta_{EF}-\eta_{FE})\left(\frac{1}{\alpha_2\beta_1}-\frac{1}{\alpha_1\beta_2}\right), \tag{7}$$

then there exists $N_0<\infty$ *and for every* $N\geq N_0$ *there exists* $\mu_0(N)>0$ *such that for* $N\geq N_0$ *and* $0<\mu\leq\mu_0(N)$,

$$\pi(ALL\,G_1) > \pi(ALL\,G_2).$$

If

$$\limsup_{N\to\infty}\left(\frac{1}{\alpha_2}+\frac{1}{\beta_2}-\frac{1}{\alpha_1}-\frac{1}{\beta_1}\right)\phi^E_{BA} < (\eta_{EF}-\eta_{FE})\left(\frac{1}{\alpha_2\beta_1}-\frac{1}{\alpha_1\beta_2}\right),$$

then there exists $N_0<\infty$ *and for every* $N\geq N_0$ *there exists* $\mu_0(N)>0$ *such that for* $N\geq N_0$ *and* $0<\mu\leq\mu_0(N)$,

$$\pi(ALL\,G_1) < \pi(ALL\,G_2).$$

Proof In view of (1), matrix M from the lemma is here given by

$$M=\begin{pmatrix}
m^E_{11} & \beta_1\phi^E_{AB} & \frac{\mu}{N} & \mu\phi^E_{AB} & \eta_{EF} & 0 & 0 & 0\\
\alpha_1\phi^E_{BA} & m^E_{22} & \mu\phi^E_{BA} & \frac{\mu}{N} & 0 & \eta_{EF} & 0 & 0\\
\frac{\mu}{N} & \mu\phi^E_{AB} & m^E_{33} & \beta_2\phi^E_{AB} & 0 & 0 & \eta_{EF} & 0\\
\mu\phi^E_{BA} & \frac{\mu}{N} & \alpha_2\phi^E_{BA} & m^E_{44} & 0 & 0 & 0 & \eta_{EF}\\
\eta_{FE} & 0 & 0 & 0 & m^F_{11} & \beta_1\phi^E_{BA} & \frac{\mu}{N} & \mu\phi^E_{BA}\\
0 & \eta_{FE} & 0 & 0 & \alpha_1\phi^E_{AB} & m^F_{22} & \mu\phi^E_{AB} & \frac{\mu}{N}\\
0 & 0 & \eta_{FE} & 0 & \frac{\mu}{N} & \mu\phi^E_{BA} & m^F_{33} & \beta_2\phi^E_{BA}\\
0 & 0 & 0 & \eta_{FE} & \mu\phi^E_{AB} & \frac{\mu}{N} & \alpha_2\phi^E_{AB} & m^F_{44}
\end{pmatrix},$$

where the diagonal entries are such that the row sums of M are 1. Let δ_i denote the cofactor of the entry in row i and column 8 of $M-I_8$. Then, as M is irreducible, $(\sum_i\delta_i)^{-1}(\delta_1,\ldots,\delta_8)$ is the unique invariant probability vector of M and $\delta_1<0,\ldots,\delta_8<0$; see, e.g., Fiedler (1986), pp. 78–80. Hence, by the lemma,

$$\pi(ALL\,G_1)=\frac{\delta_1+\delta_2+\delta_5+\delta_6}{\delta_1+\cdots+\delta_8},\qquad \pi(ALL\,G_2)=\frac{\delta_3+\delta_4+\delta_7+\delta_8}{\delta_1+\cdots+\delta_8}.$$

Consequently,

$$\pi(ALL\,G_1)>\pi(ALL\,G_2)\iff \Delta:=-\delta_1-\delta_2+\delta_3+\delta_4-\delta_5-\delta_6+\delta_7+\delta_8>0.$$

One may verify that Δ can be written as

$$\Delta=\mu\big(\phi^E_{BA}-\phi^E_{AB}\big)\big[\big(\eta_{EF}+\eta_{FE}\big)\eta_{EF}\eta_{FE}\big(\phi^E_{BA}\big)^2\big\{(\eta_{EF}-\eta_{FE})(\beta_1\alpha_2-\alpha_1\beta_2)$$
$$+(\alpha_1\beta_1\beta_2+\alpha_1\alpha_2\beta_1-\alpha_2\beta_1\beta_2-\alpha_1\alpha_2\beta_2)\phi^E_{BA}\big\}+\phi^E_{AB}\psi_1+\mu\psi_2\big],$$

Author's personal copy

where ψ_1 and ψ_2 are polynomials in the ϕ's, α's, β's, η's, and in μ and $1/N$. Hence,

$$\sup_{\substack{0<\mu\le 1,\\ N=2,3,\ldots}} |\psi_1| + |\psi_2| < \infty.$$

To determine the sign of Δ note first that by (6), there is $N_1 < \infty$ so that $\phi^E_{BA} - \phi^E_{AB} > 0$ for all $N \ge N_1$. Suppose that inequality (7) holds. Then the limit inferior, as $N \to \infty$, of the term in braces is positive. In view of (6), it follows that there exists $N_2 < \infty$ and for every $N \ge N_2$ there exists $\mu_0(N) > 0$ such that the term in brackets is positive if $N \ge N_2$ and $0 < \mu \le \mu_0(N)$. Thus, if $N \ge N_0 := \max(N_1, N_2)$ and $0 < \mu \le \mu_0(N)$, then $\Delta > 0$, and so $\pi(ALL\, G_1) > \pi(ALL\, G_2)$. The proof of the remaining assertion is similar. □

Remark A.1 (a) If the limit $\Phi^E_{BA} = \lim_{N\to\infty} \phi^E_{BA}$ exists, condition (7) becomes

$$\left(\frac{1}{\alpha_2} + \frac{1}{\beta_2} - \frac{1}{\alpha_1} - \frac{1}{\beta_1}\right)\Phi^E_{BA} > (\eta_{EF} - \eta_{FE})\left(\frac{1}{\alpha_2\beta_1} - \frac{1}{\alpha_1\beta_2}\right), \tag{8}$$

which is the condition given in Theorem 1 of the main text.

(b) Under the additional assumption that $\liminf_{N\to\infty} \phi^E_{BA} > 0$, the assertion of Theorem A.1 holds with some μ_0 that does not depend on N. To see this note that under the additional condition, the product of $(\phi^E_{BA})^2$ and the term in braces discussed in the proof of Theorem A.1 is bounded away from 0 if N is sufficiently large.

(c) To prove the claim in the discussion that Theorem 1 holds with a relaxed version of assumption (iii) on mutation probabilities assume, instead, that at every time step, with probability $\epsilon(\mu + \mu')$ a randomly chosen individual is replaced by an individual with the other genotype, where the phenotype remains unchanged with probability $\epsilon\mu$ and is changed with probability $\epsilon\mu'$. Then, in matrix M, in the entries at positions $(1,4)$, $(2,3)$, $(3,2)$, $(4,1)$, $(5,8)$, $(6,7)$, $(7,6)$, and $(8,5)$ the factor μ has to be replaced by μ'. It can now be shown that Δ, obtained analogously from the modified M, can be written as

$$\begin{aligned}\Delta = \mu'\big(\phi^E_{BA} - \phi^E_{AB}\big)\big[&(\eta_{EF} + \eta_{FE})\eta_{EF}\eta_{FE}\big(\phi^E_{BA}\big)^2\big\{(\eta_{EF} - \eta_{FE})(\beta_1\alpha_2 - \alpha_1\beta_2)\\ &+ (\alpha_1\beta_1\beta_2 + \alpha_1\alpha_2\beta_1 - \alpha_2\beta_1\beta_2 - \alpha_1\alpha_2\beta_2)\phi^E_{BA}\big\} + \phi^E_{AB}\psi_1 + \mu\psi_2 + \mu'\psi_3\big],\end{aligned}$$

where ψ_1, ψ_2, ψ_3 are polynomials in the ϕ's, α's, β's, η's, μ, μ', and $1/N$. It follows as before that if (7) holds, then there exists N_0 and for each $N \ge N_0$ there exists $\mu_0(N) > 0$ such that $\pi(ALL\, G_1) > \pi(ALL\, G_2)$, provided $N \ge N_0$ and $\max(\mu, \mu') < \mu_0(N)$.

We next derive optimal switching parameters for the case where environmental changes are of the order of phenotypic switches. We consider switching parameters in a given interval $[s, S]$ with $0 < s < S$ and we assume for simplicity that $\Phi^E_{BA} = \lim_{N\to\infty} \phi^E_{BA}$ exists. In view of Theorem A.1, we say that $(\alpha_1, \beta_1) \in [s, S]^2$ is optimal if inequality (8) holds for all $(\alpha_2, \beta_2) \in [s, S]^2$ with $(\alpha_2, \beta_2) \ne (\alpha_1, \beta_1)$. That is, an optimal genotype is favored over every competitor if mutations are rare and the population is large.

Author's personal copy

Corollary A.1 *Let* $\Phi_{BA}^{E} = \lim_{N\to\infty} \phi_{BA}^{E}$.

(a) *If* $\eta_{EF} - \eta_{FE} < -S\Phi_{BA}^{E}$, *then* (S, s) *is optimal.*
(b) *If* $|\eta_{EF} - \eta_{FE}| < S\Phi_{BA}^{E}$, *then* (S, S) *is optimal.*
(c) *If* $\eta_{EF} - \eta_{FE} > S\Phi_{BA}^{E}$, *then* (s, S) *is optimal.*

Proof (a) Suppose $\eta_{EF} - \eta_{FE} < -S\Phi_{BA}^{E}$. Then for all $\alpha_2, \beta_2 \in [s, S]$,

$$\Phi_{BA}^{E} - \frac{\eta_{EF} - \eta_{FE}}{\beta_2} > 0, \qquad \Phi_{BA}^{E} + \frac{\eta_{EF} - \eta_{FE}}{\alpha_2} < 0.$$

Therefore, if $(\alpha_1, \beta_1) = (S, s)$ and $(\alpha_2, \beta_2) \in [s, S]^2 \setminus \{(\alpha_1, \beta_1)\}$, then

$$\begin{aligned}\frac{\Phi_{BA}^{E}}{\alpha_2} + \frac{\Phi_{BA}^{E}}{\beta_2} &= \left(\Phi_{BA}^{E} - \frac{\eta_{EF} - \eta_{FE}}{\beta_2}\right)\frac{1}{\alpha_2} + \left(\Phi_{BA}^{E} + \frac{\eta_{EF} - \eta_{FE}}{\alpha_2}\right)\frac{1}{\beta_2} \\ &> \left(\Phi_{BA}^{E} - \frac{\eta_{EF} - \eta_{FE}}{\beta_2}\right)\frac{1}{\alpha_1} + \left(\Phi_{BA}^{E} + \frac{\eta_{EF} - \eta_{FE}}{\alpha_2}\right)\frac{1}{\beta_1},\end{aligned}$$

showing that inequality (8) holds.

(b) Suppose that $|\eta_{EF} - \eta_{FE}| < S\Phi_{BA}^{E}$. Then for all $\alpha_2, \beta_2 \in [s, S]$,

$$\begin{aligned}\left|(\eta_{EF} - \eta_{FE})\left(\frac{1}{\alpha_2 S} - \frac{1}{S\beta_2}\right)\right| &\le \left|\frac{1}{\alpha_2} - \frac{1}{\beta_2}\right|\Phi_{BA}^{E} \\ &= \left(\frac{1}{\min(\alpha_2, \beta_2)} - \frac{1}{\max(\alpha_2, \beta_2)}\right)\Phi_{BA}^{E} \\ &\le \left(\frac{1}{\alpha_2} + \frac{1}{\beta_2} - \frac{2}{S}\right)\Phi_{BA}^{E},\end{aligned}$$

where the first inequality is strict unless $\alpha_2 = \beta_2$ and the second inequality is strict unless $\max(\alpha_2, \beta_2) = S$. It now follows that inequality (8) holds if $(\alpha_1, \beta_1) = (S, S)$ and $(\alpha_2, \beta_2) \in [s, S]^2 \setminus \{(\alpha_1, \beta_1)\}$.

The proof of (c) is similar to that of (a). □

We next turn to the pairwise comparison of genotypes in the case where the frequency of environmental changes is of the same order as the frequency of genetic mutations. We formalize this relation by assuming that η_{EF} and η_{FE} are coupled to μ by

$$\eta_{EF} = \mu\eta_{EF}^{*}, \qquad \eta_{FE} = \mu\eta_{FE}^{*}, \tag{9}$$

where η_{EF}^{*} and η_{FE}^{*} are fixed positive numbers. In the following theorem, α_i and β_i are again fixed and $\pi(ALL\, G_i)$ is studied when μ and so, by (9), η_{EF} and η_{FE} are small. That is, genetic mutations and environmental changes are rare compared to phenotypic switches.

Author's personal copy

Theorem A.2 *Assume* (9).

(a) *If*

$$\left(\frac{\alpha_1}{\beta_1}-\frac{\alpha_2}{\beta_2}\right)\left(\frac{\eta^*_{FE}}{\eta^*_{EF}}-\frac{\alpha_1}{\beta_1}\frac{\alpha_2}{\beta_2}\right)>0, \tag{10}$$

then there exists $N_0<\infty$ *and for every* $N\geq N_0$ *there exists* $\mu_0(N)>0$ *such that for* $N\geq N_0$ *and* $0<\mu\leq\mu_0(N)$,

$$\pi(ALL\,G_1)>\pi(ALL\,G_2). \tag{11}$$

The implication remains true if "$>$*" is replaced by "*$<$*" in* (10) *and* (11).

(b) *Suppose that*

$$\phi^E_{AB}=o\big(\big(\phi^E_{BA}\big)^3\big)\quad(N\to\infty). \tag{12}$$

If $\alpha_2=t\alpha_1$ *and* $\beta_2=t\beta_1$ *with* $0<t<1$, *then there exists* $N_0<\infty$ *and for every* $N\geq N_0$ *there exists* $\mu_0(N)>0$ *such that inequality* (11) *holds for all* $N\geq N_0$ *and* $0<\mu\leq\mu_0(N)$.

Proof In view of (1) and (9), matrix M is now given by

$$M=\begin{pmatrix}
m^E_{11} & \beta_1\phi^E_{AB} & \frac{\mu}{N} & \mu\phi^E_{AB} & \mu\eta^*_{EF} & 0 & 0 & 0\\
\alpha_1\phi^E_{BA} & m^E_{22} & \mu\phi^E_{BA} & \frac{\mu}{N} & 0 & \mu\eta^*_{EF} & 0 & 0\\
\frac{\mu}{N} & \mu\phi^E_{AB} & m^E_{33} & \beta_2\phi^E_{AB} & 0 & 0 & \mu\eta^*_{EF} & 0\\
\mu\phi^E_{BA} & \frac{\mu}{N} & \alpha_2\phi^E_{BA} & m^E_{44} & 0 & 0 & 0 & \mu\eta^*_{EF}\\
\mu\eta^*_{FE} & 0 & 0 & 0 & m^F_{11} & \beta_1\phi^E_{BA} & \frac{\mu}{N} & \mu\phi^E_{BA}\\
0 & \mu\eta^*_{FE} & 0 & 0 & \alpha_1\phi^E_{AB} & m^F_{22} & \mu\phi^E_{AB} & \frac{\mu}{N}\\
0 & 0 & \mu\eta^*_{FE} & 0 & \frac{\mu}{N} & \mu\phi^E_{BA} & m^F_{33} & \beta_2\phi^E_{BA}\\
0 & 0 & 0 & \mu\eta^*_{FE} & \mu\phi^E_{AB} & \frac{\mu}{N} & \alpha_2\phi^E_{AB} & m^F_{44}
\end{pmatrix},$$

where the diagonal entries are such that the row sums of M are 1. Let δ_i denote the cofactor of the entry in row i and column 8 of $M-I_8$ and let $\Delta=-\delta_1-\delta_2+\delta_3+\delta_4-\delta_5-\delta_6+\delta_7+\delta_8$. As in the proof of Theorem 1, it follows from the lemma that $\pi(ALL\,G_1)>\pi(ALL\,G_2)$ if and only if $\Delta>0$.

To prove (a) suppose that inequality (10) holds. One may verify that

$$\Delta=\mu^3\big(\phi^E_{BA}-\phi^E_{AB}\big)\Bigg[\phi^E_{AB}\phi^E_{BA}(\alpha_1\beta_2-\alpha_2\beta_1)$$
$$\times\left\{\big(\phi^E_{BA}\big)^2\left(\eta^*_{EF}+\eta^*_{FE}+\frac{2}{N}\right)(\beta_1\beta_2\eta^*_{FE}-\alpha_1\alpha_2\eta^*_{EF})+\phi^E_{AB}\psi_1\right\}+\mu\psi_2\Bigg],$$

where

$$\sup_{\substack{0<\mu\leq 1,\\ N=2,3,\ldots}}|\psi_1|+|\psi_2|<\infty.$$

Author's personal copy

By (10), the differences $\alpha_1\beta_2 - \alpha_2\beta_1$ and $\beta_1\beta_2\eta^*_{FE} - \alpha_1\alpha_2\eta^*_{EF}$ are either both positive or both negative. In view of (6), there exists $N_0 < \infty$ such that for all $N \geq N_0$, $\phi^E_{BA} - \phi^E_{AB} > 0$ and the term in braces has the same sign as $\alpha_1\beta_2 - \alpha_2\beta_1$. Thus, if $N \geq N_0$ there is $\mu_0(N) > 0$ so that for all $0 < \mu \leq \mu_0(N)$, $\Delta > 0$ and so $\pi(ALL\, G_1) > \pi(ALL\, G_2)$. The proof for the case where the inequalities in (10) and (11) are reversed is similar.

To prove (b) suppose that $0 < t < 1$, $\alpha_2 = t\alpha_1$ and $\beta_2 = t\beta_1$. Then

$$\Delta = \mu^4\big(\phi^E_{BA} - \phi^E_{AB}\big)(1-t) \times \left[t\big(\phi^E_{BA}\big)^3\left(\eta^*_{EF} + \eta^*_{FE} + \frac{2}{N}\right)\eta^*_{EF}\eta^*_{FE}\alpha_1\beta_1(\alpha_1+\beta_1) + O\big(\phi^E_{AB}\big) + O(\mu)\right]$$

and the assertion follows from (12). □

Remark A.2 Under the additional assumption that $\liminf_{N\to\infty} \phi^E_{BA} > 0$, the assertion of Theorem A.2(b) holds with some μ_0 that does not depend on N. This is obvious from the representation of Δ given in the proof of Theorem A.2(b).

References

Acar, M., Mettetal, J. T., & van Oudenaarden, A. (2008). Stochastic switching as a survival strategy in fluctuating environments. *Nat. Genet.*, *40*, 471–475.

Avery, S. V. (2006). Microbial cell individuality and the underlying sources of heterogeneity. *Nat. Rev., Microbiol.*, *4*, 577–587.

Balaban, N. Q., Merrin, J., Chait, R., Kowalik, L., & Leibler, S. (2004). Bacterial persistence as a phenotypic switch. *Science*, *305*, 1622–1625.

Ben-Porath, E., Dekel, E., & Rustichini, A. (1993). On the relationship between mutation rates and growth rates in a changing environment. *Games Econ. Behav.*, *5*, 576–603.

Bürger, R., Willensdorfer, M., & Nowak, M. A. (2006). Why are phenotypic mutation rates much higher than genotypic mutation rates? *Genetics*, *172*, 197–206.

Cohen, D. (1966). Optimizing reproduction in a randomly varying environment. *J. Theor. Biol.*, *12*, 119–129.

DeWitt, T. J., Sih, A., & Wilson, D. S. (1998). Costs and limits of phenotypic plasticity. *Trends Ecol. Evol.*, *13*, 77–81.

Donaldson-Matasci, M. C., Lachmann, M., & Bergstrom, C. T. (2008). Phenotypic diversity as an adaptation to environmental uncertainty. *Evol. Ecol. Res.*, *10*, 493–515.

Fiedler, M. (1986). *Special matrices and their applications in numerical mathematics*. Dordrecht: Martinus Nijhoff.

Foster, D., & Young, P. (1990). Stochastic evolutionary game dynamics. *Theor. Popul. Biol.*, *38*, 219–232.

Freidlin, M. I., & Wentzell, A. D. (1998). *Random perturbations of dynamical systems* (2nd ed.). New York: Springer.

Fudenberg, D., & Harris, C. (1992). Evolutionary dynamics with aggregate shocks. *J. Econ. Theory*, *57*, 420–441.

Fudenberg, D., & Imhof, L. A. (2006). Imitation processes with small mutations. *J. Econ. Theory*, *131*, 251–262.

Fudenberg, D., & Imhof, L. A. (2008). Monotone imitation dynamics in large populations. *J. Econ. Theory*, *140*, 229–245.

Fudenberg, D., & Levine, D. K. (1998). *The theory of learning in games*. Cambridge: MIT Press.

Fudenberg, D., Nowak, M. A., Taylor, C., & Imhof, L. A. (2006). Evolutionary game dynamics in finite populations with strong selection and weak mutation. *Theor. Popul. Biol.*, *70*, 352–363.

Gillespie, J. H. (1974). Natural selection for within-generation variance in offspring number. *Genetics*, *76*, 601–606.

Author's personal copy

Harmer, G. P., & Abbott, D. (2002). A review of Parrondo's paradox. *Fluct. Noise Lett.*, *2*, R71–R107.
Ibba, M., & Söll, D. (1999). Quality control mechanisms during translation. *Science*, *286*, 1893–1897.
Imhof, L. A., Fudenberg, D., & Nowak, M. A. (2005). Evolutionary cycles of cooperation and defection. *Proc. Natl. Acad. Sci. USA*, *102*, 10797–10800.
Ishii, K., Matsuda, H., Iwasa, Y., & Sasaki, A. (1989). Evolutionarily stable mutation rate in a periodically changing environment. *Genetics*, *121*, 163–174.
Jablonka, E., Oborny, B., Molnár, I., Kisdi, É., Hofbauer, J., & Czárán, T. (1995). The adaptive advantage of phenotypic memory in changing environments. *Philos. Trans. R. Soc. Lond. Ser. B*, *350*, 133–141.
Kandori, M., Mailath, G. J., & Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, *61*, 29–56.
Kimura, M. (1967). On the evolutionary adjustment of spontaneous mutation rates. *Genet. Res.*, *9*, 23–34.
King, O. D., & Masel, J. (2007). The evolution of bet-hedging adaptations to rare scenarios. *Theor. Popul. Biol.*, *72*, 560–575.
Kussell, E., & Leibler, S. (2005). Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, *309*, 2075–2078.
Kussell, E., Kishony, R., Balaban, N. Q., & Leibler, S. (2005). Bacterial persistence: A model of survival in changing environments. *Genetics*, *169*, 1807–1814.
Lachmann, M., & Jablonka, E. (1996). The inheritance of phenotypes: An adaptation to fluctuating environments. *J. Theor. Biol.*, *181*, 1–9.
Leigh, E. (1970). Natural selection and mutability. *Am. Nat.*, *104*, 301–305.
Levins, R. (1968). *Evolution in changing environments*. Princeton: Princeton University Press.
Livnat, A., Pacala, S. W., & Levin, S. A. (2005). The evolution of intergenerational discounting in offspring quality. *Am. Nat.*, *165*, 311–321.
Maamar, H., Raj, A., & Dubnau, D. (2007). Noise in gene expression determines cell fate in Bacillus subtilis. *Science*, *317*, 526–529.
Philippi, T., & Seger, J. (1989). Hedging one's evolutionary bets, revisited. *Trends Ecol. Evol.*, *4*, 41–44.
Rosenberg, S. M. (2001). Evolving responsively: Adaptive mutation. *Nat. Rev. Genet.*, *2*, 504–515.
Samuelson, L. (1997). *Evolutionary games and equilibrium selection*. Cambridge: MIT Press.
Sandholm, W. (2009). Evolutionary game theory. In R. A. Meyers (Ed.), *Encyclopedia of complexity and system science* (pp. 3176–3205). Berlin: Springer.
Slatkin, M. (1974). Hedging one's evolutionary bets. *Nature*, *250*, 704–705.
Stumpf, M. P. H., Laidlaw, Z., & Jansen, V. A. A. (2002). Herpes viruses hedge their bets. *Proc. Natl. Acad. Sci. USA*, *99*, 15234–15237.
Süel, G. M., Kulkarni, R. P., Dworkin, J., Garcia-Ojalvo, J., & Elowitz, M. B. (2007). Tunability and noise dependence in differentiation dynamics. *Science*, *315*, 1716–1719.
Thattai, M., & van Oudenaarden, A. (2004). Stochastic gene expression in fluctuating environments. *Genetics*, *167*, 523–530.
Thomas, M. J., Platas, A. A., & Hawley, D. K. (1998). Transcriptional fidelity and proofreading by RNA polymerase II. *Cell*, *93*, 627–637.
True, H. L., & Lindquist, S. L. (2000). A yeast prion provides a mechanism for genetic variation and phenotypic diversity. *Nature*, *407*, 477–483.
Tuljapurkar, S. (1990). *Population dynamics in variable environments*. New York: Springer.
Wu, B., Gokhale, C. S., Wang, L., & Traulsen, A. (2011). How small are small mutation rates? *J. Math. Biol.*, to appear.
Young, H. P. (1993). The evolution of conventions. *Econometrica*, *61*, 57–84.