

On the Relationship between Mutation Rates and Growth Rates in a Changing Environment*

ELCHANAN BEN-PORATH

Tel Aviv University and MEDS KGSM Northwestern University

EDDIE DEKEL

University of California at Berkeley and Tel Aviv University

AND

ALDO RUSTICHINI

New York University and CORE

Received August 8, 1992

The paper examines the relationship between the mutation rate and the rate of growth of the population when the environment is changing. It is shown that while the global maximum of the growth rate is attained at a non-zero mutation rate, zero is always a local maximum. This suggests that a population with an initially low mutation rate will evolve to a zero mutation rate. *Journal of Economic Literature* Classification Number: C72. © 1993 Academic Press, Inc.

1. INTRODUCTION

Recent papers have examined dynamic models of learning and of evolution.¹ The dynamic processes in these models usually include two ingredi-

* We are grateful to Dilip Abreu, Bob Anderson, Vince Crawford, Drew Fudenberg, George Mailath, Eric Maskin, Matthew Rabin, Jeroen Swinkels, two referees, and seminar audiences at the University of Wisconsin in Madison, Harvard/MIT, Boston University, Princeton, and Northwestern for many helpful comments. We are very thankful to Gadi Barlevy and Debbie Minehart, both of whom provided extremely useful research assistance. We are also pleased to acknowledge financial support from the NSF, the Sloan Foundation, and Tel Aviv University.

¹ These include Canning (1989), Foster and Young (1990), Fudenberg and Harris (1992), Fudenberg and Kreps (1988), Kandori *et al.* (1993), Kandori (1991), Kandori and Rob (1992), Samuelson (1991), Young (1993), and a large number of papers based on Maynard Smith's (1982) evolutionarily motivated solution concept, ESS.

ents. First, the process incorporates some aspects of rationality or optimality, such as the assumption in evolutionary models that better strategies become more prevalent. Second, it is often assumed that there is an exogenous random force, such as mutation, leading players to change their strategies. Usually the rate of change is assumed to be small, and almost all of the research considers only the case where the rate of change approaches zero.

In this paper we use an evolutionary model to examine how the outcome of the dynamic process depends on the rate of change. In particular we examine how the long-run growth rate of the population depends on the mutation rate. Insofar as evolutionary forces select among mutation rates, our analysis sheds light on what rates of change might evolve.²

It is important to clarify that we use the notion of evolution in two distinct ways. First, we explicitly model an evolutionary process in which more successful strategies become more prevalent. Second, we informally argue that evolution would select mutation rates that lead to a higher growth rate. While we do not explicitly model this latter selection process, we do discuss in Section 6 the extent to which an explicit model would support our conclusions.

Our primary goal is to understand the forces that determine whether one rate of mutation does better—in terms of yielding a higher growth rate for the population—than another. It is intuitive that a small rate of mutation does better than a larger mutation rate in stationary environments. This is because any positive rate of mutation, no matter how small, will introduce the optimal strategy; and after a long enough time the dynamics of evolution will lead to almost all of the population playing this best strategy. Then, in the long run, mutations will only have costs and no benefits: a higher rate of mutation will just mean that there are more mutations away from the optimal strategy. In changing environments, however, a higher mutation rate can be beneficial, even in the long run: mutations can give a jump start to a strategy that suddenly becomes good due to the environment changing. More precisely, mutations are beneficial after periods in which few people played the best strategy; so they can be helpful after each time that the environment changes.

The preceding discussion suggests that higher rates of mutation are better when the environment changes rapidly. This idea can be explored in several ways. One might ask whether the rate of mutation that maximizes the population's growth rate is increasing in the rate of fluctuation of the environment. More specifically, one could ask if the mutation rate

² Fudenberg and Levine (1990) endogenize the rate of experimentation in a model where players optimally choose how much to experiment and what to play. By contrast, most of the research that is concerned with how play evolves over time has used models where players are not fully rational optimizers. Our work suggests what mutation rates will arise endogenously in a *non-optimizing* framework.

that maximizes the population's growth rate goes to zero as the rate of change of the environment goes to zero. Both these questions focus on global changes: assuming that evolutionary forces select the globally optimum mutation rate, will the mutation rate decrease when the environment fluctuates less? Will it disappear when the environment fluctuates arbitrarily slowly? We show that the answer to this global type of question is yes: quite generally the mutation rate that maximizes the growth rate converges to zero when the rate of change of the environment converges to zero; and in a special case we can prove that the optimal mutation rate is increasing in the rate of environmental change.³

Even if evolutionary forces select among mutation rates, they need not succeed in selecting the global optimum. Therefore we also explore local selection pressures, and address two questions. First, are there local maxima which are not global maxima? Second, if for a given rate of change of the environment and a given mutation rate it is beneficial to increase the mutation rate slightly, is it therefore beneficial to increase the mutation rate when the environment changes more rapidly? The answer to these local-selection questions is not so clear cut. In one special case the intuition above is validated: if it is beneficial to slightly increase the mutation rate for a given rate of environmental change, then a fortiori it is beneficial when the environment changes more rapidly. Similarly, in this special case, the only local maximum is the global maximum. However, in general, the intuition is incomplete. Loosely speaking, if the mutation rate is so small that the "best" strategies do not become prevalent soon enough before the environment changes again, then it is best to do without the mutation altogether. More precisely, when the mutation rate is small then a change to a strategy that is currently optimal will last so long that there will be another change in the environment and the strategy will no longer be optimal.⁴ Indeed, our main result is that, in general, for *any* rate of change of the environment, if the mutation rate is small enough, then it is better to decrease it further. That is, zero mutations is a local maximizer of the growth rate of the population. As is well known, local maximizers may well be selected by evolutionary forces, so this identifies a force that will favor zero mutation rates.⁵

³ In the general case we think the second result is still true, but we have not proven this.

⁴ A related intuition is that without mutation evolutionary forces would select a strategy that is best in the "average" environment. So for small mutation rates, one should focus on the average environment; and since the average environment is (by definition) not changing, the previous intuition that smaller mutation rates are better in fixed environments is applicable.

⁵ The fact that there is a local maximum at zero can also be used to show that for given rates of mutation and of environmental change it may be good to slightly increase the mutation rate, while for a more rapidly changing environment it is best to slightly decrease the mutation rate.

1.1. *A Sketch of the Model and Results*

In our model each player faces a single-person choice problem, so the fitness of a strategy does not depend on other players' choices. But we assume a changing environment which means that the strategies' fitness, i.e., the structure of payoffs, is changing over time. Excluding the effects of mutation, the proportion of players choosing a particular strategy evolves according to the replicator dynamics. This means that the number of players choosing a particular strategy in the next period is the number today multiplied by that strategy's fitness/payoff today. However, before playing, each player might mutate and switch to another strategy. For a given specification of payoffs and how they change over time, and a given rate of mutation, this dynamic process generates a long-run rate of growth for the population. We examine the long-run growth rate as a function of the mutation rate and the rate of change of the environment. We consider two ways in which the environment changes over time: in one there is a deterministic cycle between two environments; in the second the two environments follow a Markov process.⁶

The stochastic model raises an interesting issue concerning how to evaluate when one mutation rate is better than another. One might expect that the mutation rate which will be selected is the one with the highest growth rate of the expected population. Alternatively, in comparing two mutation rates, perhaps the one which is more likely to have higher growth will be selected. This second approach seems likely to suffer from the drawback that there might not be one mutation rate that is most likely to have a higher growth rate when compared to any other mutation rate. It turns out that this is not a problem and we can find a mutation rate that dominates any other mutation rate with probability 1. To contrast the two measures, assume that we want to compare two different mutation rates, both of which are represented in one population. Since the individuals with different mutation rates will go through the same environmental changes, if one mutation rate does better than the other with probability one, then almost surely the proportion of the population with the better mutation rate will converge to one. Thus, if we were to explicitly model selection of mutation rates, the mutation rate that has probability one of dominating would be selected. Therefore, the probability-of-domination comparison seems more interesting. The growth rate of the expected population size seems relevant only if there are many different islands with both mutation rates represented on each, and in each island the environment is changing independently. Then, the proportion of the population that has the mutation rate that yields higher growth of the expected

⁶ We discuss in the paper, and show in the Appendix, that for the deterministic model our results hold for any finite number of strategies and environments.

population will, with probability one, converge to 1 when we add up the populations on all the islands.

The paper is organized as follows. In Section 2 we introduce the model by considering a fixed environment, and we obtain the natural preliminary result that in such a model the growth rate is decreasing in the mutation rate.⁷ We then present the more interesting model where the environment is changing. In Section 3 we consider the knife-edge case where non-best strategies are wiped out immediately since this case is easier to analyze and helps demonstrate the issues that come up in the stochastic model. We show that in this case the mutation rate that maximizes the growth rate is strictly greater than zero and is larger when the rate of change of the environment increases. More precisely, in the deterministic-cycles model the optimal mutation rate is $1/n$ where n is the length of a cycle. The same result obtains in the stochastic model when we compare mutation rates according to the probability of domination: with probability 1 the mutation rate $1/n$ has a higher growth rate than any other mutation rate, where n is the expected length of time in an environment. In both models it is also the case that there are no local maxima other than the global maxima. The growth rate of the expected population size, however, turns out to be a decreasing function of the mutation rate. Thus comparing mutation rates using the probability-of-domination measure seems more attractive a priori, as argued above, and yields the same results as the deterministic cyclic model.

The main part of the paper, Section 4, contains the analysis of a changing environment when non-best strategies can survive. For the deterministic model we first show that the globally maximizing mutation rate is strictly greater than zero and it converges to zero as the cycle length converges to infinity. Our main result is that there is always a local maximum at zero. Though we have been unable to analytically confirm these results for the stochastic model (when we compare mutation rates according to the probability of one having a higher growth rate than another), we have confirmed them using simulations.⁸

In Section 5 we discuss several issues concerning the robustness and the interpretation of our model.

Naturally, there is a literature in biology that relates to this paper. Kimura (1960, 1967) was one of the first to examine which mutation rate would be optimal when the environment is changing. He did not explicitly

⁷ See also Robson (1993, Theorem 2) who has shown that zero mutations is a global maximizer when the environment is constant.

⁸ We can show that in the stochastic model, if we compare mutation rates according to the growth rate of the expected population, then—as in the previous model—zero is a locally maximizing growth rate. While this result agrees with the deterministic model it still seems less interesting on a priori grounds, hence it is not presented in the paper.

consider a mutator gene, and in this respect our paper is similar to his work; the literature since Kimura's work has examined the implications of explicitly modeling genes that control the mutation rate. In asexual models of reproduction Leigh (1970) and Maynard Smith (1978) explicitly introduced mutator genes and their conclusions are the same as Kimura's. These three papers conclude that in asexual populations the optimal mutation rate will be selected and will equal $1/n$, the rate of change of the environment. Thus our question, model, and some of our results are similar to their work. Nevertheless, our main result—that in general zero mutations are a local maximizer of the growth rate, and hence may be selected—has not been noted. In addition the significant distinction between the growth rate of the expected population and the probability of dominance has not been discussed. Most of the literature after Kimura seems to have focused on sexual models, and has identified a "reduction principle."⁹ Roughly speaking, this literature argues that when the mutator gene is separate from the strategy gene, the mutation rate that will evolve is zero. We provide a qualitatively different justification for a reduction principle: since a zero mutation rate is a local maximizer, it could be the outcome of evolutionary forces.

2. THE MODEL

We begin with the case of a fixed environment. This case serves as a useful benchmark, and helps develop some of the intuition which is used below in the discussion of the more interesting case where the environment is not fixed. The exogenous parameters of the model are

$$\begin{array}{r} E_1 \\ s_1 \quad a \\ s_2 \quad b, \end{array}$$

where E_1 is the environment, $a > b > 0$ are the payoffs, and s_i are the strategies.¹⁰ The dynamic process is based on the replicator dynamics: Each person is characterized by a strategy choice; the number of children of each person is given by that person's payoffs (often called fitness

⁹ See Twomey and Feldman's (1990) introduction for a survey of such work, including Karlin and McGregor (1974), Holsinger and Feldman (1983), Liberman and Feldman (1986), and other papers that reach similar conclusions for rates of recombination and of migration.

¹⁰ For simplicity we assume that there are only two strategies, and it will be clear that all the results for the case of a fixed environment hold for any finite number of strategies; we discuss the case of infinitely many strategies in Footnote 13.

payoffs); all non-mutant children follow their parent's strategy, while a certain proportion of children mutate and switch to the other strategy. We consider a population all of whose individuals have the same mutation rate, $\mu \in [0, \frac{1}{2}]$, and ask how the long-run rate of growth depends on μ .¹¹ Let $x_i(t)$ denote the size of the population playing s_i in period t . Then the replicator dynamics are

$$\begin{aligned} x(t+1) &\equiv (x_1(t+1), x_2(t+1)) \\ &= ((1-\mu)ax_1(t) + \mu bx_2(t), \mu ax_1(t) + (1-\mu)bx_2(t)). \end{aligned}$$

Let $\bar{x}(t)$ be the total population in period t , $\bar{x}(t) \equiv x_1(t) + x_2(t)$. The rate of growth in period t is denoted by $\gamma_t \equiv \bar{x}(t+1)/\bar{x}(t)$, and the long-run rate of growth is then $\lim_{t \rightarrow \infty} \gamma_t$.¹² (Henceforth we write $\lim_t \gamma_t$ for $\lim_{t \rightarrow \infty} \gamma_t$. Of course, γ_t depends on the starting point $x(0)$ and the mutation rate μ ; we do not include these arguments for notational simplicity.) We focus on how the long-run growth rate depends on μ , and write $\gamma(\mu) \equiv \lim_t \gamma_t$.

LEMMA 2.1. *For all $0 < \mu < \frac{1}{2}$, $\lim_t \gamma_t$ exists and is independent of the initial condition $x(0)$. Therefore $\gamma: (0, \frac{1}{2}) \rightarrow \mathbb{R}$ is a well-defined function.*

Proof. This and all other proofs that are not presented in the body of the paper are in the Appendix.

In the knife-edge case where $\mu = 0$, the long-run growth rate depends on $x(0)$. In particular, $\lim_t \gamma_t = a$ if $x_1(0) > 0$, and $\lim_t \gamma_t = b$ if $x_1(0) = 0$. To see this note that if no one is playing s_1 at the beginning and if the mutation rate is zero, then no one will ever play s_1 . Therefore, in any period t , the per-period growth rate, γ_t , is b . On the other hand, if anyone is playing s_1 at time zero, then the proportion playing s_1 will converge over time to 1. This is because the portion of the population playing s_1 is growing at the rate a , which is larger than the rate of growth of those playing s_2 . Since in the limit the proportion playing s_1 is 1, the rate of growth in the limit is a . We define $\gamma(0) \equiv a$, and then $\gamma: [0, \frac{1}{2}) \rightarrow \mathbb{R}$ is well-defined.

LEMMA 2.2 *The growth rate $\gamma(\mu)$ is an analytic function.*

It is helpful to understand when a higher mutation rate leads to a higher growth rate. One intuitive argument is that increasing μ should help after a period in which the best strategy is played by a small portion of the

¹¹ Mutation rates larger than $\frac{1}{2}$ are unnatural, and in any case the interesting results concern small mutation rates.

¹² One might alternatively want to define the long-run growth rate by considering the limit of the average growth rate: $\lim_{t \rightarrow \infty} [\bar{x}(t)/\bar{x}(0)]^{1/t}$. Clearly when the limits exist they coincide.

population, and it should hurt after a period in which the best strategy was played by many people. To see this, consider a fictitious stage after reproduction has occurred but before mutations occur. If at that stage the majority of people play the best strategy, then mutations will cause more people to stop playing the best strategy than it will cause people to begin playing it. Thus, once most of the population is playing the best strategy, increased mutation decreases growth. This suggests that $\gamma(\mu)$ is strictly decreasing.¹³

PROPOSITION 2.3. *For all $0 \leq \mu < \frac{1}{2}$, $\gamma'(\mu) < 0$.*

As discussed in the Introduction, the nature of the dependence of the growth rate on the mutation rate tells us how the mutation rate might evolve, if the mutation rate itself were subjected to evolutionary forces. Proposition 2.3 implies that if the environment were fixed then the mutation rate would tend to zero. This intuitive idea has also arisen in the biological literature regarding the evolution of sexual reproduction. An early theory explaining why sexual reproduction evolved was that, relative to asexual reproduction, it increases variability. But, in line with Proposition 2.3, more recently it has been noted that the variability caused by sexual reproduction only "scrambles up [our] perfectly good genome."¹⁴ Biologists now explain that variability helps by keeping the species one step ahead of parasites that are constantly changing: "Pathogens . . . can reproduce . . . in seconds and mutate many times . . . [G]enetic variability gives us hosts at least a fighting chance. . . ."¹⁵ This change in the environment that is faced by hosts is endogenous; it is a consequence of changes by the parasites who are also subject to evolutionary forces. While we do not model such endogenous changing, our analysis below can be seen as a preliminary step toward a complete game-theoretic treatment of this interaction, since it examines how different mutation rates fare against a given rate of change of the environment. We now present a model appro-

¹³ If there were an infinite number of strategies then one might expect that even in a fixed environment there would not be a maximum at zero mutations. (Therefore, as we see below, one might expect that in the case of a changing environment, there would not be a local maximum at zero.) The intuition is that it might be best to keep on searching for a better strategy. However, this is not the case, and in fact the maximum at zero is robust. First, if there exists a best strategy, it is not hard to see that there will be a maximum at zero. Similarly, if the reproduction rates (payoffs) are bounded and continuous, then the growth rate will be decreasing in the mutation rate. (For any pair of mutation rates, there will be a point in time at which most of the population will be playing a strategy which is almost best. If the almost best is close enough to the best, then the potential gain from mutations into better strategies is small, while the cost of shifting into worse strategies is large.) Finally, if the payoffs are not bounded then it is not clear how to define the growth rate.

¹⁴ Richard Michod, quoted in Gutin (1992, p. 36).

¹⁵ Gutin (1992, p. 38).

appropriate for examining the idea that variability, or a high rate of mutation, is good when the environment, and hence the best response to the environment, is changing.

For simplicity, the model and results are presented assuming two environments and two strategies. All our results extend to any finite number of strategies and environments; the general case is presented in the Appendix.

	E_1	E_2
s_1	a	c
s_2	b	d

The two environments are denoted by E_1 and E_2 , and the two strategies are s_1 and s_2 . For the problem to be distinct from the fixed-environment case we assume $a > b > 0$ and $d > c > 0$. So, for $i = 1, 2$, s_i is best in E_i . We consider the generic case where $ac \neq bd$ and assume w.l.o.g. $ac > bd$. (The relevance of this assumption becomes clear later.) When the environment is E_1 , the population evolves as before; and in E_2 the dynamics are $(x_1(t+1), x_2(t+1)) = ((1-\mu)cx_1(t) + \mu dx_2(t), \mu cx_1(t) + (1-\mu)dx_2(t))$.

The remaining feature of the model that needs to be described is how the environment is determined. We assume a deterministic cycle of length $2n$: Starting in E_1 the environment switches back and forth every n periods between E_1 and E_2 .

We can define the per-period growth rate $\gamma_t = \bar{x}(t+1)/\bar{x}(t)$ as before, and of course it depends on μ and on n . However, during a cycle we would expect different growth rates, e.g., if the population adapts to new environments slowly then the growth immediately after the environment shifts may be less than the growth immediately before the shift, so γ_t need not converge. The long-run growth rate is defined instead by taking the per-period geometric average of the growth rate over a cycle. That is, we consider the limit of how much the population grows during $2n$ periods and take the $2n$ root of this limit.

$$\gamma(\mu, n) \equiv \lim_{k \rightarrow \infty} [\bar{x}(2n(k+1))/\bar{x}(2nk)]^{1/2n}.$$

LEMMA 2.4. $\gamma(\mu, n)$ is well defined for $\frac{1}{2} > \mu > 0$, that is, the limit exists and it is independent of the initial state $x(0)$. Moreover, for any n , $\gamma(\cdot, n)$ is analytic w.r.t. μ .

LEMMA 2.5. The long-run growth rate is the same when averaging over any $2n$ consecutive periods: $\gamma(\mu, n) = \lim_{k \rightarrow \infty} [\bar{x}(2n(k+1) + i)/\bar{x}(2nk + i)]^{1/2n}$ for all $i = 0, 1, \dots, 2n - 1$. Moreover, the long-run growth

rate equals the limit of the average growth rate: $\gamma(\mu, n) = \lim_{t \rightarrow \infty} [\bar{x}(t)/\bar{x}(0)]^{1/t}$.

To clarify the calculation of the long-run growth rate consider the simple knife-edge case where $\mu = 0$ and $x_1(0) = 0$. Then $\lim_{k \rightarrow \infty} [\bar{x}(2n(k+1) + i)/\bar{x}(2nk + i)]^{1/2n} = (bd)^{1/2}$, since all the population will be playing s_2 ; and out of the $2n$ periods of a cycle there will be n periods with growth rate b (when the environment is E_1) and n periods with growth rate d (during E_2).

3. CHANGING ENVIRONMENTS: RESULTS WHEN NON-OPTIMAL STRATEGIES ARE EXTINGUISHED

Before analyzing the general model described above, we start with a special case where $b = c = 0$. In this case, a non-optimal strategy is extinguished immediately. This knife-edge case is easier to solve analytically, and therefore serves as a useful step toward the general model. In this case the intuition that mutations are better when the environment fluctuates more rapidly is confirmed quite generally.

PROPOSITION 3.1. *The long-run growth rate $\gamma(\mu, n)$ is strictly quasi-concave in μ for any n , and the unique maximum is at $\mu^* = 1/n$.*

Proof. The proof is a simple calculation. Starting from $(x_1(0), x_2(0))$, at the beginning of period n , the population will be $(x_1(n), x_2(n)) = (a^n(1 - \mu)^n x_1(0), a^n(1 - \mu)^{n-1} \mu x_1(0))$. (Note that $x_2(0)$ is irrelevant since individuals choosing s_2 do not reproduce in any period where the environment is E_1 .) The environment then changes and n periods later, at the beginning of period $2n$, the environment will change again. At that point we have $x_1(2n) = a^n d^n (1 - \mu)^{2n-2} \mu^2 x_1(0)$. Thus the per-period growth rate is given by $\gamma(\mu, n) = (ad)^{1/2} (1 - \mu)^{1-1/n} \mu^{1/n}$, and then $\gamma(\cdot, n)$ for any n is increasing (decreasing) in μ if $\mu < (>) 1/n$. QED

Thus, if the mutation rate is selected by any hill-climbing process, it will reach the optimal rate; and the optimal rate is increasing at the rate of change of the environment.¹⁶

A natural question to raise is whether this conclusion also holds in a stochastic model. Instead of going through n periods of each environment, it seems more plausible that the environment is determined randomly, but with a high likelihood of persistence. The most interesting analog to the

¹⁶ While the models are different, this conclusion is essentially the same as the result of Maynard Smith (1978), which he attributes to Leigh (1970), who in turn followed Kimura (1960, 1967).

cycling model seems to be a Markov chain, where the expected length of time in a particular environment is n , so that in each period t there is a probability $1/n$ of switching to a new environment. That is $P(E^{t+1} = E_i | E^t = E_j) = 1/n$, $i = 1, 2$; $i \neq j$. For this model we restrict attention further to the case where $a = d$.

It is less clear how we should compare different mutation rates in this stochastic model. That is, what is the analog to the comparison of mutation rates according to population growth rates, as above? One might presume that the mutation rate which will be selected is the one with the highest growth rate of the expected population.¹⁷ We first consider this possibility.

In order to state our results for the randomly changing environment it is helpful to introduce some notation. Let h_t denote a t period history specifying the environment in each period, $h_t \in \{E_1, E_2\}^t$. Denote by $c(h_t)$ the number of times that the environment changed in the history h_t . If we start in environment E_i , with an initial population $(x_1(0), x_2(0))$ of individuals with mutation rate μ , then after a history h_t the total population (after reproduction) will be $a^t(1 - \mu)^{t-c(h_t)}\mu^{c(h_t)}x_i(0)$. (In any period t where the environment is E_i , each individual who plays s_i will have a children. Of these how many will bear children? If the environment in $t + 1$ is again E_i then a proportion $1 - \mu$ will reproduce so that the effective growth rate, so long as there are no changes in the environment, is $a(1 - \mu)$. If the environment in period $t + 1$ changed, then only the mutants born to the t period players will reproduce. Thus, after a change in the environment the effective growth rate is only $a\mu$.) The per-period growth rate after history h_t and starting from environment E_i is then $[a^t(1 - \mu)^{t-c(h_t)}\mu^{c(h_t)}x_i(0)/(x_1(0) + x_2(0))]^{1/t}$. This leads to the following result.

PROPOSITION 3.2. *The growth rate of the expected population size at time t is decreasing in μ , for $1/n \in (0, \frac{1}{2})$ and for any t .*

Proof. If there are k switches of environment in the first t periods, and the initial environment was E_i then the total population after t periods of reproduction is $a^t(1 - \mu)^{t-k}\mu^k x_i(0)$. This events occurs with probability $\binom{t}{k} (1 - 1/n)^{t-k}(1/n)^k$. Thus the expected population at time t is

$$\sum_{k=1}^t \binom{t}{k} (1 - 1/n)^{t-k}(1/n)^k a^t (1 - \mu)^{t-k} \mu^k x_i(0) \\ = [(1 - 1/n)(1 - \mu) + (1/n)\mu]^t x_i(0)$$

which, for $1/n < \frac{1}{2}$, is decreasing in μ .

QED

¹⁷ The only explicit model of a random environment of which we are aware does in fact focus on the growth rate of the expected population, see Gillespie (1981).

Thus in this case $\mu^* = 0$ is a global maximizer. This surprising conclusion, which conflicts with the deterministic cycling model, is due to the fact that we used expectations to measure the success of a mutation rate: an extremely high growth rate in some very unlikely history yields high growth of the expected population. In fact, a population with mutation rate 0 will almost surely be extinct in finite time.¹⁸ This raises doubts about comparing mutation rates according to the growth rate of the expected population. Moreover, as we argued in the Introduction, comparing mutation rates by calculating expectations is relevant only if we think that there are simultaneously many islands in each of which the environment is determined independently *and* if we want to know which is the mutation rate that is most common after a long enough time when we aggregate the populations on all the islands.

Instead, we may be interested in asking which mutation rate is most common in most islands, or if there is only one location we would ask which mutation rate is most likely to be the dominant proportion of the population. Thus, instead of looking at the growth rate of the expected population, it seems more relevant to consider the probability of domination. That is, in comparing two mutation rates to see which will be selected, we should ask which is more likely to yield a larger population. We now demonstrate that this will lead to a result exactly analogous to Proposition 3.1 from the deterministic cycling model.

Denote by $\bar{x}(\mu, h_t)$ the total population after history h_t and with mutation rate μ . For each k , where k is the number of switches in history h_t , mutation rate μ_1 has less growth than μ_2 if and only if $a^k(1 - \mu_2)^{t-k}\mu_2^k \geq a^k(1 - \mu_1)^{t-k}\mu_1^k$, which is equivalent to $(1 - \mu_2)^{1-r}\mu_2^r \geq (1 - \mu_1)^{1-r}\mu_1^r$, where r is the proportion of times in which a switch occurs.

By the weak law of large numbers the probability that $|r - 1/n|$ is greater than ε is converging to 0 as t goes to infinity. So if we have

$$(1 - \mu_2)^{1-1/n}\mu_2^{1/n} > (1 - \mu_1)^{1-1/n}\mu_1^{1/n} \quad (1)$$

then the probability that the population with μ_2 at time t is greater than the population with μ_1 at time t , is converging to 1 as t goes to infinity.

Thus we compare pairs of mutation rates according to Eq. (1). The function $f(\mu) = (1 - \mu)^{1-1/n}\mu^{1/n}$, is strictly quasi-concave and is maximized at $\mu^* = 1/n$. Thus the probability that a population with mutation rate μ^*

¹⁸ That is, it will be extinct at time t with probability approaching 1 as t goes to infinity. This is because after one period only individuals playing one of the strategies remain, and once the environment changes all these individuals will be extinct. The probability of the environment changing within the first t periods is, of course, converging to 1 as t converges to infinity.

will be larger than any population with a different mutation rate converges to 1. Formally, we have proven the following results.

PROPOSITION 3.3. For any pair μ_1, μ_2 ,

$$P(\lim_{t \rightarrow \infty} ([\bar{x}(\mu_1, h_t)]^{1/t} - [\bar{x}(\mu_2, h_t)]^{1/t}) > 0) = 1 \Leftrightarrow f(\mu_1) > f(\mu_2).$$

COROLLARY 3.4. $P(\lim_{t \rightarrow \infty} ([\bar{x}(1/n, h_t)]^{1/t} - [\bar{x}(\mu, h_t)]^{1/t}) > 0) = 1$ if $\mu \neq 1/n$.

This section has shown that, for the special case where non-optimal strategies are immediately extinguished ($b = c = 0$) and optimal strategies yield the same payoffs ($a = d = 0$), the mutation rate that yields the highest growth rate is $1/n$ where n is the (expected) length of time in an environment. We have also argued that the correct method of evaluating mutation rates when environments are random is *not* the growth rate of the expected population but rather the likelihood of having a higher growth rate, and we have shown that using this method the conclusions of the deterministic and stochastic models coincide. We now return to the more general model.

4. CHANGING ENVIRONMENTS: RESULTS WITH GENERAL PAYOFFS

This section focuses on the deterministic cycling model with general payoffs and addresses the main question of this paper: what does $\gamma(\mu, n)$ look like? A natural conjecture might be that¹⁹ for n large enough we get something like two back-to-back fixed-environment models. To state this more precisely, let $\gamma^i(\mu)$ be the long-run growth rate if the environment is E_i forever.²⁰ If n is large, then during the E_1 phase we get to a point where the per-period growth rate is close to $\gamma^1(\mu)$ and stay there long enough so that the average growth in this phase is close to $\gamma^1(\mu)$. The case is similar in the E_2 phase. This intuition is drawn in Fig. 1, and it leads to the proposition below.

PROPOSITION 4.1. $\lim_{n \rightarrow \infty} \gamma(\mu, n) = [\gamma^1(\mu)\gamma^2(\mu)]^{1/2}$, for $\mu \in (0, \frac{1}{2})$, where $\gamma^i(\mu)$, $i = 1, 2$, denotes the long-run growth rate when the environment is E_i forever (that is, when the environment is not changing).

¹⁹ We ignore for the moment knife-edge cases where one strategy is not represented in the population.

²⁰ Thus $\gamma^1(\mu)$ is the same as $\gamma(\mu)$ which was defined preceding Lemma 2.1 where we dealt with the unchanging environment case.




FIGURE 1

COROLLARY 4.2. $\lim_{\mu \rightarrow 0}(\lim_{n \rightarrow \infty} \gamma(\mu, n)) = (ad)^{1/2}$.

However, the conjecture (as well as the proposition) does not tell the whole story. In particular, $\gamma(0, n) = (ac)^{1/2}$ if $x_1(0) > 0$.²¹ To see why this is true, recall that we assumed that $ac > bd$, so those playing s_1 will reproduce more during a cycle than those playing s_2 . Hence the proportion of people playing s_1 converges to 1 (see Fig. 2). We define $\gamma(0, n) \equiv (ac)^{1/2}$; with this definition $\gamma(\cdot, n)$ is analytic w.r.t. μ .

The previous paragraph says that, for every n and almost every starting point, if $\mu = 0$, then the population converges to a state where the proportion playing a best reply to the average environment is 1. By average we mean an environment in which s_1 has payoffs $(ab)^{1/2}$ and s_2 yields $(cd)^{1/2}$. Continuity of the growth rate w.r.t. the mutation rate implies that, given any n , there is a neighborhood of zero, such that for every mutation rate in that neighborhood the following holds: (i) the limit proportion of the population that is playing a best reply to the average environment is close




FIGURE 2

²¹ Thus, the order of limits in the corollary is important.




FIGURE 3

to 1, and (ii) the long-run rate of growth is close to $(ac)^{1/2}$, rather than $(ad)^{1/2}$.²²

In contrast to the fixed-environment case, here the rate of mutation that maximizes growth is, for large enough n , different from zero. This demonstrates the benefit of mutations when the environment is changing: by providing a “kick” in the right direction, mutations enable the population to adapt to the current environment. To clarify the contrast with the fixed environment case, recall that the only role of mutations in a fixed environment is to introduce a strategy that is not present in the population. We saw in Proposition 2.3 that once the strategy is present mutations are harmful. We see now that in the changing environment case, mutations serve a significant additional role: for any n , if μ is too small then the presence of s_2 is irrelevant; despite being good in E_2 , s_2 is driven to extinction. Mutations must be large enough to give a “kick” in the right direction.

So far we have discovered two properties of the function $\gamma(\mu, n)$. First, for all n , $\gamma(0, n) = (ac)^{1/2}$. Second, if we let $\mu^*(n)$ denote the mutation rate that maximizes the growth rate, then Propositions 2.3 and 4.1 imply that for n large, $\gamma(\mu^*(n), n)$ is approximately $(ad)^{1/2}$. We might guess that the function is single peaked at $\mu^*(n)$, (so that the growth rate as a function of the mutation rate, for given n , is quasiconcave). However, simulations suggest that the function is as in Fig. 3. Not only is it not single peaked; it is decreasing in a neighborhood of zero. That this somewhat surprising feature is general is stated in the proposition below.

PROPOSITION 4.3. *For every n there exists a number $\delta(n) > 0$ such that $\gamma(\mu, n)$ is decreasing in the interval $[0, \delta(n)]$.*

²² Continuity of the growth rate (except at knife-edge cases) follows from continuity of the eigenvalues of the matrix representing the growth during a $2n$ cycle; see the proofs of Lemmas 2.1 and 2.4.

That is, near zero, an increase in the mutation rate is costly. The rough intuition for this result is that, when μ is small relative to $1/n$, then the “kick” that is given by μ doesn’t help enough. The problem is that in this case the “kick” doesn’t only cause more people to play s_2 during E_2 , but rather it results in more people playing s_2 during *both* E_2 and E_1 . Alternatively put, if μ is small then the strategy selected is one which maximizes against the average environment (not one that tries to maximize in each environment separately), and when maximizing against the unchanging average environment, as in the case of a fixed environment, less mutation is better.

For a more precise intuition, consider a $2n$ cycle starting in E_2 . When the mutation rate is small we know that most of the population is playing s_1 , which is a best reply to the average environment. So mutations from s_2 to s_1 are negligible all along the cycle. Therefore, we can focus on mutations from s_1 into s_2 only. Consider then the growth rate, during a cycle lasting $2n$ periods, of a lineage starting from a group who mutated from s_1 to s_2 . (The mutation rate must be small relative to n in order for us to be able to claim that *throughout* the cycle we can ignore mutations back out of this lineage into s_1). Instead of reproducing $a^n c^n$ these mutants and their descendants reproduce only $b^n d^n$. So the mutation is costly.

Proposition 4.3 immediately implies that for n , there is a local maximum at zero.²³ This suggests a reason why, even in a changing environment, the population might evolve to zero mutations.

A natural question is whether this result in the deterministic-cycling model holds in a stochastic version of the model. While we showed that the results for the $b = c = 0$ and $a = d$ case coincide in the deterministic-cycling model and in the stochastic model, we have been unable to extend the result that there is a local maximum of $\gamma(\mu, n)$ in the general cycling model to a stochastic model. We conjecture that the following analog to Proposition 3.3 does hold.

$$\text{Conjecture. } \lim_{\mu \rightarrow 0} P(\lim_{t \rightarrow \infty} ([\bar{x}(0, h_t)]^{1/t} - [\bar{x}(\mu, h_t)]^{1/t}) > 0) = 1.$$

That is, the probability that a mutation rate of zero does better than another strictly positive mutation rate goes to one as the mutation rate becomes small. We have been able to confirm this result in simulations, but we have been unable to prove it analytically.

In terms of understanding the form of $\gamma(\mu, n)$, it is interesting to compare two different n 's, a slowly changing environment and a quickly changing environment. When is an increase in mutation rates more valuable? It might seem that mutations are more valuable in a quickly changing environ-

²³ Note that the local maximum is robust in that it does not depend on the parameters at all.




FIGURE 4

ment. In particular we might expect the following: If $n < n'$, then $\gamma'(\mu, n') > 0 \Rightarrow \gamma'(\mu, n) > 0$. This is because, if an increase is good with n' , then when the environment changes more quickly there are less periods in which the best strategy is chosen; therefore increasing μ should help even more. This intuition is satisfied in Fig. 4, which is obtained by simulations, on the intervals $(0, \mu')$ and $(\mu, \frac{1}{2})$. However, this intuition is violated in the interval (μ', μ) . Moreover, this violation is general (and not only a feature of the simulation): Given any n , there exists an \bar{n} such that for any $n' > \bar{n}$ there is an interval (μ', μ) on which $\gamma'(\mu, n') > 0$ but $\gamma'(\mu, n) < 0$. (This immediately follows from Proposition 4.3 and the fact that $\lim_{n \rightarrow \infty} \mu^*(n) = 0$, which in turn follows immediately from Propositions 2.3 and 4.1.)

A related view of $\gamma(\mu, n)$ can be seen by asking when is more mutation better? Not necessarily in a more quickly changing environment. If the environment changes quickly enough relative to μ , then it might be better (locally) to decrease the rate of mutation. This we saw in Proposition 4.3. On the other hand, if the environment is changing slowly relative to μ , then simulations suggest that, again, a decrease in μ increases growth. And this is partially confirmed by Proposition 4.1, which implies that a smaller growth rate is better than a larger one if n is large enough, that is, for every μ_1, μ_2 with $\frac{1}{2} > \mu_2 > \mu_1 \geq 0$, there exists $N(\mu_1, \mu_2)$ such that for $n \geq N$, $\gamma(\mu_1, n) > \gamma(\mu_2, n)$. However, if the environment is changing, but not too quickly nor too slowly, then an increase in μ is beneficial, since it will enable the population to adapt better in each environment. And the increased adaptation of the population is the dominant feature when n is neither too small nor too large.

5. DISCUSSION

In this section we discuss several issues concerning the robustness and interpretation of our model. One issue was already touched upon in the

Introduction: in what sense will those mutation rates that maximize population growth be selected? A second issue is whether the results seem plausible when alternative dynamic processes are used. Both of these are related to the final point we address, which is whether our results might be relevant for economics.

Throughout this paper we have been comparing populations with different mutation rates, and we have been asking which will have a larger growth rate in *isolation*. This seems like what biologists term group selection. If the populations interacted would not that affect our results? As argued by Maynard Smith (1978), in this type of model the answer, at least in the case of the deterministic cycling model, is clearly no. If all the populations were together then in our model there would still be no interaction among the individuals with different mutation rates and therefore the conclusions would not be affected. More specifically, consider a population with individuals who were characterized by two traits: a strategy choice and a mutation rate. Assume that mutation rates are inherited perfectly and only the strategy choice can mutate. Our results imply that the proportion of individuals who have the mutation rate that yields a higher growth rate would converge to 1.²⁴

The argument above does not imply that the mutation rate that globally maximizes growth would be selected; it only shows that among those mutation rates represented in the population, the one that maximizes population growth would dominate. If, for example, only a small range of mutation rates were represented, and only small changes of mutation rates were possible, then one could end up at a local maximum.

In the random model the group-selection issue is less clear. We have shown that the set of histories on which $\mu = 1/n$ has a larger population in period t than *any* other μ has probability converging to 1 when t goes to infinity. However, it might be that for different sets of possible histories, where each set has probability going to zero, there is some mutation rate that does better than $1/n$ (where the mutation rate may vary from set to set). It may be that the probability of the union of these sets is not converging to zero. Then, if we had a population consisting of subpopulations with all different possible mutation rates, the probability that at least one of these subpopulations would grow more than the subpopulation with mutation rate $1/n$ can be large.

One question of robustness concerns the fact that the mutation rate was assumed throughout not to depend on any exogenous features. Allowing the mutation rate to be time-dependent would dramatically

²⁴ Even if one believes that there is a capacity constraint that limits growth, so long as the death caused by this capacity constraint is random and not a function of mutation rates, the conclusion that growth-maximizing mutation rates would dominate the population would continue to hold.

change our results. In a model where the environment changes every n periods growth would be greater if the mutation rate could be large immediately after the environment changed, and small while the environment is constant. However, while the cycles we model in this paper are simple and easy to identify, cycles in reality are more complex and harder to identify. Therefore the simplifying assumption that mutation rates cannot depend on the cycle or on time seems reasonable.²⁵ Moreover, in a stochastic environment it would not be enough for mutation rates to be time-dependent; they would have to depend on some aspect which seems hard to identify, such as the environment itself, the proportion of players choosing particular strategies, or the growth rate of the population. The idea that the mutation rate should not be constant is interesting, and very similar to models of simulated annealing and other adaptive algorithms used to solve for optima. While examining what rules governing mutation rates might evolve seems worthwhile—especially in the context of learning—it is beyond the scope of this current paper.

Another question that is especially important if one wants to evaluate the relevance of this to economics concerns the robustness of our results to using alternative dynamic processes. The interior maximum in the case of changing environments relies on a feature of the replicator dynamics.²⁶ In this dynamic process, if there is almost no one paying a particular strategy then the process is slow. That is, if a strategy in some period is played by very few agents, then even if that strategy is very good, only a small proportion of people will be playing it in the next round (since only “children” of those playing this strategy can “inherit” it). Therefore, it can be very beneficial for mutations to give a “jump start” to strategies that become good after a change in environment. If, however, everyone can observe everything then more plausible dynamic processes—especially if one has in mind a learning process—might have very fast movement as soon as any one player adopted a good strategy. Thus, the conclusion that in a changing environment there is an interior maximum seems more likely to extend to dynamic processes where players learn only from some individuals, possibly as in local-learning models (e.g., Ellison (1992), Blume (1993)), than to dynamics where good strategies are immediately adopted by everyone.

We have argued so far that evolutionary forces would lead most of the population to exhibit mutation rates that are global or local maximizers of the population’s growth rate. The most direct interpretation of these

²⁵ This is not to say that the observed proportion of individuals with different mutation rates will be constant. Immediately after a change in environment those individuals with higher mutation rates will do better. Our results suggest however that the cycles will dampen and converge to most of the population having a growth-maximizing mutation rate.

²⁶ We are grateful to Jeroen Swinkels for raising this issue.

results seems to be as a biological model. Similar to other models of evolution in the literature there are some leaps of faith that are needed to justify giving these results an economic interpretation. First, we must assume that individuals "inherit" from others what strategy to play and how often to "mutate." So, we need to justify the dynamic process of inheritance embodied in the replicator dynamics and to consider the extent to which mutations and experiments are similar.

The replicator dynamics are motivated using payoffs that are given in terms of biological fitness, not utility. Moreover, there has only been limited success at justifying the replicator dynamics in a learning context (see, e.g., Nachbar (1990) and Cabrales (1993)). However, other than the issue discussed in the preceding paragraph, the intuition underlying our main results on an interior global maximum and a local maximum at zero do not seem to depend on the precise process. It is, nevertheless, crucial that we believe that individuals decide both what to do and how often to try something new based on the success of these modes of behavior in their immediate surroundings. The interpretation of mutations as experiments is standard in the literature; however, this relies on the premise that a player who tries something new and the people who follow that player (his/her "children") need to be locked into the experiment (see, e.g., Kandori and Rob (1992) for justifications of this type of experiment).

Despite the strong assumptions needed for the economic interpretation, we think the model and its conclusions are interesting. First, they are interesting even if only as results about evolution. In addition, it seems intuitive that new strategies may be tried too rarely because *slightly* more experimentation will lead a player to lag behind the ever changing environment and others will then not mimic those who do try new things a little more often. It will seem better to stick to the tried and true which is best on average.

7. APPENDIX

We begin with some known results that will be used in the proofs.

THEOREM 7.1 *Let D be a strictly positive, real-valued, n -by- n matrix. There is a unique eigenvalue of D with largest modulus. This eigenvalue is positive, real, and its associated eigenvector can be taken to be positive.*

Proof. See, e.g., Theorem 1, p. 288, in Bellman (1974).

QED

THEOREM 7.2. *Let D be a real-valued n -by- n diagonal matrix whose maximal element on the diagonal is unique. There is a unique eigenvalue of D with largest modulus. This eigenvalue is positive, real, and its associated eigenvector has zeros in all places except for the location associated with the maximal element on the diagonal.*

Proof. The proof is immediate. QED

Henceforth we denote the largest eigenvalue of a matrix D by $\lambda(D)$.

THEOREM 7.3. *For D a strictly positive, n -by- n matrix and $x(0) = (x_j(0))_{j=1}^n$ a positive vector, let $x(t) \equiv x(0)D^t$. Then $\lim_{t \rightarrow \infty} x(t) / \sum_{j=1}^n x_j(t)$ exists, is independent of $x(0)$, and is equal to the eigenvector associated with $\lambda(D)$. Moreover, $\lim_{t \rightarrow \infty} \sum_{j=1}^n x_j(t) / \sum_{j=1}^n x_j(0)$ exists and is equal to $[\lambda(D)]^t$.*

Proof. See, e.g., Theorem 4, p. 292, in Bellman (1974). QED

THEOREM 7.4 *Let D be a real-valued, n -by- n , diagonal matrix whose maximal element on the diagonal is unique, and let $x(0)$ be a strictly positive vector. Then $\lim_{t \rightarrow \infty} x(t) / \sum_{j=1}^n x_j(t)$ exists, is independent of $x(0)$, and is equal to the eigenvector associated with $\lambda(D)$. Moreover, $\lim_{t \rightarrow \infty} \sum_{j=1}^n x_j(t) / \sum_{j=1}^n x_j(0)$ exists and is equal to $[\lambda(D)]^{it}$.*

Proof. The proof is immediate. QED

THEOREM 7.5. *Let $D: [0, \frac{1}{2}] \rightarrow \mathfrak{R}^{n^2}$ define a n -by- n matrix whose entries are polynomial functions of a real-valued parameter, say z . Assume that for any z there is a unique maximal eigenvalue $\lambda(z) \equiv \lambda(D(z))$. Then $\lambda(z)$ is analytic in z .*

Proof. See, e.g., Kato (1982, pp. 74, 83–84, or 148). Note that it is easy to show that $\lambda(z)$ is C^∞ using the implicit function theorem and the fact that $\lambda(z)$ is an eigenvalue with multiplicity 1. The proof of analyticity is not so immediate. QED

THEOREM 7.6. *Consider a two-by-two, strictly-positive matrix $D(z) = \begin{bmatrix} d_{11}(z) & d_{12}(z) \\ d_{21}(z) & d_{22}(z) \end{bmatrix}$. Assume that the elements are converging in z as follows: $\lim_{z \rightarrow 0} d_{12}(z) = \lim_{z \rightarrow 0} d_{21}(z) = 0$, and $\lim_{z \rightarrow 0} d_{11}(z) > \lim_{z \rightarrow 0} d_{22}(z) \geq 0$. Then the eigenvector associated with $\lambda(D(z))$ is continuous in z at 0 and can be taken to converge to $(1, 0)$.*

Proof. This is a straightforward calculation. QED

In the proofs of our results we use the following additional notation. Let $p(t) \equiv x_i(t) / \bar{x}(t)$ be the proportion of the population that plays strategy s_i , at time t starting from some strictly positive population vector $x(0)$. The long-run proportion is $p(\mu) = \lim_{t \rightarrow \infty} p(t)$; the existence of $p(\mu)$ follows immediately from Theorem 7.3 for the case of $\mu > 0$ and from Theorem 7.4 when $\mu = 0$, if one defines D as in the proof of Lemma 2.1 below.

LEMMA 2.1. For all $0 < \mu < \frac{1}{2}$, $\lim_t \gamma_t$ exists and does not depend on the initial condition $x(0)$, so $\gamma: (0, \frac{1}{2}) \rightarrow \mathfrak{R}$ is a well-defined function.

Proof of Lemma 2.1. Let $D \equiv \begin{bmatrix} (1-\mu)a & \mu a \\ \mu b & (1-\mu)b \end{bmatrix}$. Then by Theorem 7.3 $\gamma(\mu)$ is given by the maximal eigenvalue of D and by Theorem 7.1 this eigenvalue is real and positive.

LEMMA 2.2. The growth rate $\gamma(\mu): (0, \frac{1}{2}) \rightarrow \mathfrak{R}$ is an analytic function.

Proof of Lemma 2.2. By Theorems 7.3 and 7.4, $\gamma(\mu)$ is the maximal eigenvalue and by Theorems 7.1 and 7.2 it is unique. Then by Theorem 7.5 it is an analytic function.

PROPOSITION 2.3. For all $0 \leq \mu < \frac{1}{2}$, $\gamma'(\mu) < 0$.

Proof of Proposition 2.3. It is easy to see that $p(t+1) = (1-\mu)ap(t) + \mu b(1-p(t)) / (ap(t) + b(1-p(t)))$. Therefore, taking the limit, we obtain

$$p(\mu) = \frac{(1-\mu)ap(\mu) + \mu b(1-p(\mu))}{ap(\mu) + b(1-p(\mu))}.$$

Consider the functions $F(p, \mu) = ((1 - \mu)ap + \mu b(1 - p))/(ap + b(1 - p)) - p$, and $G(p, \mu) = F(p, \mu)[ap + b(1 - p)]$, where p and μ are in $[0, 1]$. Since $F(p(\mu), \mu) = 0$, and since $G(p, \mu) = 0 \Leftrightarrow F(p, \mu) = 0$ then, applying the implicit function theorem to G , it is easy to see that

$$p'(\mu) = -\frac{-ap(\mu) + b(1 - p(\mu))}{(1 - \mu)a - \mu b - 2ap(\mu) - b(1 - p(\mu)) + bp(\mu)}.$$

Since the denominator equals $(a - b)(1 - 2p(\mu)) - \mu(a + b)$, we have that, if $p(\mu) > \frac{1}{2}$, then $p'(\mu) < 0$.

We now show that indeed $p(\mu) > \frac{1}{2}$ when $\mu < \frac{1}{2}$. First, it is easy to see that if $0 < \mu$, then $F(0, \mu) > 0$; and if, in addition, $0 < \mu < \frac{1}{2}$, then $F(\frac{1}{2}, \mu) > 0$. Also, $F(p, \mu) > 0 \Leftrightarrow G(p, \mu) > 0$. Taking partial derivatives of G as below, we see that G is concave in p .

$$\frac{\partial G(p, \mu)}{\partial p} = (a - b)(1 - 2p) - \mu(a + b); \quad \frac{\partial^2 G(p, \mu)}{\partial p^2} = -2(a - b) < 0.$$

Since G is concave in p , and since $G(\frac{1}{2}, \mu) > 0$ and $G(0, \mu) > 0$, it follows that $G(p, \mu) > 0$ for every $p \in [0, \frac{1}{2}]$.

Since $G(p(\mu), \mu) = 0$ we conclude that $p(\mu) > \frac{1}{2}$. It follows that $p'(\mu) < 0$. Since $\gamma(\mu) = ap(\mu) + b(1 - p(\mu))$ we obtain $\gamma'(\mu) = (a - b)p'(\mu) < 0$.

LEMMA 2.4. $\gamma(\mu, n)$ is well defined for all $\frac{1}{2} > \mu > 0$, that is, the limit exists and it is independent of the initial state $x(0)$. Moreover, for any n , $\gamma(\cdot, n)$ is analytic w.r.t. μ .

Proof of Lemma 2.4. The structure of the proof and the theorems used are the same as in the proof of Lemma 2.1.

For environment E_1 let $A(1) \equiv \begin{bmatrix} (1-\mu)^a & \mu a \\ \mu b & (1-\mu)^b \end{bmatrix}$ and for E_2 let $A(2) \equiv \begin{bmatrix} (1-\mu)^c & \mu c \\ \mu d & (1-\mu)^d \end{bmatrix}$. Now define $D \equiv [A(1)]^n \cdot [A(2)]^n$, and proceed as in the proofs of Lemmas 2.1 and 2.2.

LEMMA 2.5. The long-run growth rate is the same when averaging over any $2n$ consecutive periods, that is, $\gamma(\mu, n) = \lim_{k \rightarrow \infty} [\bar{x}(2n(k+1) + i)/\bar{x}(2nk + i)]^{1/2n}$ for all $i = 0, 1, \dots, 2n - 1$.

Proof of Lemma 2.5. We prove the claim for $i = n$, i.e., for the case where the cycle starts just after the system has shifted to the environment E_2 . The proof for any other i is similar. Let $\gamma_t \equiv \bar{x}(t + 2n)/\bar{x}(t)$ denote the growth rate over a $2n$ cycle starting in period t . A simple modification of the proof of Lemma 2.4 shows that there is a limit growth rate for cycles of length $2n$ that start in E_2 , that is, $\lim_{k \rightarrow \infty} \gamma_{(2k+1)n}$ exists. Let $\bar{\gamma}(\mu, n)$ denote the $2n$ th root of this limit. Let $t = 2kn$ be some period in which the system moves back to the environment E_1 . Let m be some natural number, and let $t' = 2(k + m)n$ be a period m cycles later. Consider the growth of the population from time t to time t' , that is $\bar{x}(t')/\bar{x}(t)$. If k is large enough, then $\gamma_{2l(k+l)n}$ and $\gamma_{(2(k+l)n-n)}$, $l = 0, \dots, m - 1$, are close to their limits. So we have $\bar{x}(t')/\bar{x}(t) = [\gamma(\mu, n)]^{2nm} + R_1(k)$ where $\lim_{k \rightarrow \infty} R_1(k) = 0$. On the other hand, we have $\bar{x}(t')/\bar{x}(t) = (\bar{x}(t+n)/\bar{x}(t))([\bar{\gamma}(\mu, n)]^{2n(m-1)})(\bar{x}(t')/\bar{x}(t' - n)) + R_2(k)$, where $\lim_{k \rightarrow \infty} R_2(k) = 0$. Rearranging terms we obtain that

$$\left[\frac{\gamma(\mu, n)}{\bar{\gamma}(\mu, n)} \right]^{2nm} = \frac{(\bar{x}(t+n)/\bar{x}(t))(\bar{x}(t')/\bar{x}(t' - n))}{[\bar{\gamma}(\mu, n)]^{2n}} + R_3(k), \quad \text{where } \lim_{k \rightarrow \infty} R_3(k) = 0.$$

If $\gamma(\mu, n) > \bar{\gamma}(\mu, n)$ then the left-hand side tends to infinity when $m \rightarrow \infty$. If $\gamma(\mu, n) < \bar{\gamma}(\mu, n)$ then the left-hand side tends to zero when $m \rightarrow \infty$. Since the right-hand side is

bounded from above and it is bounded from below above zero, it follows that $\gamma(\mu, n) = \bar{\gamma}(\mu, n)$.

PROPOSITION 4.1. $\lim_{n \rightarrow \infty} \gamma(\mu, n) = (\gamma^1(\mu)\gamma^2(\mu))^{1/2}$ for $0 < \mu < \frac{1}{2}$.

Proof of Proposition 4.1. Let $\mu \in (0, \frac{1}{2})$, and ε be given. There exists $\delta > 0$ such that

$$\begin{aligned} [\gamma^1(\mu)\gamma^2(\mu)]^{1/2} + \varepsilon/2 &> [(\gamma^1(\mu) + \delta)(\gamma^2(\mu) + \delta)]^{1/2} \\ &> [(\gamma^1(\mu) - \delta)(\gamma^2(\mu) - \delta)]^{1/2} \\ &> [(\gamma^1(\mu)\gamma^2(\mu))^{1/2} - \varepsilon/2]. \end{aligned}$$

It follows from Lemma 2.1 that there exists a number T such that for every initial state of the population, $x(0)$, and for every $t > T$, the rate of growth at time t in environment E_i , which we denote by γ_t^i , satisfies $|\gamma_t^i - \gamma^i(\mu)| < \delta$. That is, if the rate of change of the environment is slow enough, then in each environment there will be a point after which the per-period growth rate is almost equal to the long-run growth rate if the environment were fixed.

It is easy to see that there exists a number $N > T$, such that for every $n > N$,

$$([\gamma^1(\mu) - \delta]^{n-T} b^T [\gamma^2(\mu) - \delta]^{n-T} c^T)^{1/2n} \geq [(\gamma^1(\mu) - \delta)(\gamma^2(\mu) - \delta)]^{1/2} - \varepsilon/2 > [\gamma^1(\mu)\gamma^2(\mu)]^{1/2} - \varepsilon,$$

and

$$\begin{aligned} ([\gamma^1(\mu) + \delta]^{n-T} a^T [\gamma^2(\mu) + \delta]^{n-T} d^T)^{1/2n} &\leq [(\gamma^1(\mu) + \delta)(\gamma^2(\mu) + \delta)]^{1/2} + \varepsilon/2 \\ &< [\gamma^1(\mu)\gamma^2(\mu)]^{1/2} + \varepsilon. \end{aligned}$$

Since for every t , $a \geq \gamma^1(t) \geq b$ and $d \geq \gamma^2(t) \geq c$, it follows that for every n ,

$$\left| \left[\frac{\bar{x}(2n)}{\bar{x}(0)} \right]^{1/2n} - [\gamma^1(\mu)\gamma^2(\mu)]^{1/2} \right| < \varepsilon.$$

COROLLARY 4.2. $\lim_{\mu \rightarrow 0} (\lim_{n \rightarrow \infty} \gamma(\mu, n)) = (ad)^{1/2}$.

Corollary 4.2 follows immediately from Proposition 4.1.

PROPOSITION 4.3. For every n there exists a number $\delta(n) > 0$, such that $\gamma(\mu, n)$ is decreasing in the interval $[0, \delta(n)]$.

Proof of Proposition 4.3. We show that for every n there exists $\delta(n)$ such that for every $\mu \in (0, \delta(n))$, $\gamma(\mu, n) < \gamma(0, n)$. We then show that this implies the proposition. Consider a $2n$ cycle which starts in environment E_2 . Let $\bar{p}(\mu)$ denote the limit of the proportion that plays the strategy s_1 at the beginning of such a cycle. It is easy to see that the matrix $[A(2)]^n \cdot [A(1)]^n$, where $A(1)$ and $A(2)$ are as defined in Lemma 2.4, satisfies the provisions of Theorem 7.6. It follows that the function $\bar{p}(\mu)$ is continuous at 0 and therefore $\lim_{\mu \rightarrow 0} \bar{p}(\mu) = \bar{p}(0) = 1$.

Let $P_i, i = 0, \dots, 2n - 1$, be the set of individuals in the population who mutate at stage i . Since $\lim_{\mu \rightarrow 0} \bar{p}(\mu) = 1$, for every n there exists $\delta(n)$ such that for every $\mu \in (0, \delta(n))$ we can choose the sets P_i so that the following two properties are satisfied:

- (1) $P_i \cap P_j = \emptyset$ for $i \neq j$.
- (2) In each P_i the proportion that mutates from s_1 to s_2 is greater than $\frac{1}{2}$.

Property 1 can be satisfied when μ is small, by relabeling people if necessary. Property 2 can be satisfied because when $\bar{p}(\mu)$ is close to 1 and μ is small the proportion of the population that plays s_1 is close to 1 throughout the cycle.

Consider a cycle of length $2n$, that begins in period 0 in environment E_2 , and assume that

the initial proportion playing s_j is $\bar{p}(\mu)$. We now consider the effect of allowing only $i < 2n$ periods of mutation, so that mutations occur at rate μ up to but not including stage i , and no mutations occur afterward. We show that the growth rate in such cycles is decreasing in i , that is, allowing more periods of mutation, decreases the growth over this cycle. Let $\bar{\gamma}_i$ denote the rate of growth during a cycle of such a dynamic process (i.e., the ratio between the size of the population at the end and its size at the beginning of the cycle, when the population mutates up to but not including stage i and then stops mutating). With this definition $\bar{\gamma}_{2n-1} = \gamma(\mu, n)$ because a mutation at the last stage of the cycle does not affect the size of the population at the end of the cycle.

Claim. For every $i = 0, 1, \dots, 2n - 2$, $\bar{\gamma}_{i+1} < \bar{\gamma}_i$.

Proof of Claim. It follows from Property 2 that if there is a mutation at stage i then the measure of the set that changes from s_1 to s_2 is greater than the measure that changes from s_2 to s_1 . We can analyze the system in terms of the net transformations, so we can view the mutants born at the end of period i as a transformation in which a set of individuals with a positive measure changes from s_1 to s_2 (and none change from s_2 to s_1). Let \bar{P}_i denote this set of players (who were born in period i and will play and reproduce in period $i + 1$), and let m_i denote its measure. At the end of the cycle, the measure of the population that originated from \bar{P}_i is $m_i d^{n-(i-1)} b^n$ if $i < n$ and it is $m_i b^{2n-(i+1)}$ if $i > n$. However, if we allowed mutations only up to but not including stage $i - 1$, then the measure of the population that originates from \bar{P}_i at the end of the cycle is $c^{n-(i-1)} a^n$ if $i < n$ and it is $a^{2n-(i+1)}$ if $i \geq n$. We have $a^k > b^k$ for every k and $c^k a^n > d^k b^n$ for every $0 \leq k \leq n$ (because by assumption $c^n a^n > d^n b^n$ and $d > c$). It follows that the mutation at stage i decreases the growth of the lineage following \bar{P}_i and therefore the growth of the whole population.

The claim implies that $\bar{\gamma}_0 > \bar{\gamma}_{2n-1}$. This means that when the proportion of the population that plays the strategy s_1 in the beginning of the cycle is $\bar{p}(\mu)$, then the growth in the cycle is higher if there is no mutation. To complete the first part of the proof we note that if the proportion of the population that plays s_1 is 1 and $\mu = 0$ then the growth is higher than $\bar{\gamma}_0$. So we have $\gamma(0, n) > \bar{\gamma}_0 > \bar{\gamma}_{2n-1} = \gamma(\mu, n)$.

Thus, we have shown that for sufficiently small μ , $\gamma(0, n) > \gamma(\mu, n)$. To complete the proof of the proposition, we show that there exists an $\epsilon > 0$ such that $\gamma'(\mu) < 0$ for $\mu \in (0, \epsilon)$. It is easy to see that the result that was established above implies that there is no open interval $(0, \delta)$ where $\gamma' > 0$. Next assume by contradiction that there is an open interval on which $\gamma' \geq 0$ while there is no open interval on which $\gamma' > 0$. Then there is a sequence $\{\mu_k\}_{k=1}^\infty$ such that $\mu_k \rightarrow 0$ and $\gamma'(\mu_k) = 0$. By theorems 7.1, 7.2, and 7.5 γ' is analytic it, follows that $\gamma' = 0$ but this implies that γ is constant which is false. QED

The rest of the Appendix deals with the extension of our results to the general case of m environments and n strategies.

Let E_1, \dots, E_m be a finite set of environments and s_1, \dots, s_n be a finite set of actions. We assume a deterministic cycle of length m , so that the environment at time t is $E_{\lfloor t \bmod m \rfloor}$. We let $a_{ij} > 0$ denote the payoff when the environment is E_j and the action is s_i . Let $\mu > 0$ be the mutation rate.

We now define the matrix that determines the dynamic process in a way that is similar to the procedures in Lemmas 2.1, 2.2, and 2.4. First, for mutation rate μ and environment j , define an n -by- n matrix $A(j, \mu)$ as follows:

$$[A(j, \mu)]_{p,q} \equiv \begin{cases} (1 - \mu)a_{p,j} & p = q \\ \frac{\mu}{(n-1)} a_{p,j} & p \neq q. \end{cases}$$

The matrix $A(j, \mu)$ enables us to define how the population changes from one period to the

next: if $x_i(t)$ is the size of the population playing s_i in period t then the replicator dynamics are $x(t+1) \equiv x(t)A(t \bmod m, \mu)$.

Let $\bar{x}(t) \equiv \sum_{i=1}^n x_i(t)$ be the total population at time t , so $[\bar{x}(t+m)/\bar{x}(t)]^{1/m}$ is the per-period growth over a cycle through the m environments starting in period t . Let $\gamma(\mu) \equiv \lim_{k \rightarrow \infty} [\bar{x}(km)/\bar{x}((k-1)m)]^{1/m}$. Define $D(\mu) \equiv \prod_{j=1}^m A(j, \mu)$. Again using Theorems 7.1, 7.2, 7.3, 7.4, and 7.5, $\gamma(\mu)$ is well defined and analytic. An argument similar to the proof of Lemma 2.5 shows that $\gamma(\mu) = \lim_{k \rightarrow \infty} [\bar{x}(km+l)/\bar{x}((k-1)m+l)]^{1/m}$ for all l ; that is, the limit per-period growth rate over a cycle does not depend on where we start to measure the cycle.

Note that the proofs of Proposition 4.1 and of Corollary 4.2 can be adapted immediately to the general case with n strategies and m environments.

We now show that when one strategy is best on average, $\gamma(\mu)$ is decreasing in a neighborhood of 0.

THEOREM 7.7. *Given the matrix of payoffs $(a_{i,j})$, if $\prod_{j=1}^m a_{i,j} \neq \prod_{j=1}^m a_{i',j}$, for all $i \neq i'$, then there is a number $\delta > 0$ such that $\gamma(\mu)$ is decreasing in the interval $[0, \delta]$.*

Proof of Theorem 7.7. We assume w.l.o.g. that if $i' > i$ then $\prod_{j=1}^m a_{i,j} > \prod_{j=1}^m a_{i',j}$. In particular, s_1 is the strategy that maximizes the growth over a cycle. Let $a = \prod_{j=1}^m a_{1,j}$ and $b = \prod_{j=1}^m a_{2,j}$.

Let $p_i(\mu)$ denote the limit of the proportion of the population playing s_i . For $\mu > 0$, Theorem 7.3 implies that this limit exists; for $\mu = 0$ Theorem 7.4 implies that the limit exists and equals 1 for $i = 1$ and zero otherwise.

Claim. $\lim_{\mu \rightarrow 0} p_1(\mu) = 1$.

Proof of Claims. Assume by contradiction that there exists an $\varepsilon > 0$ and a sequence μ_k such that $\lim_{k \rightarrow \infty} \mu_k = 0$ and $p_1(\mu_k) \leq 1 - \varepsilon$. Then $[1 - p_1(\mu_k)]/p_1(\mu_k) \geq 1 - p_1(\mu_k) \geq \varepsilon$.

Let $\bar{p}(\mu) \equiv [p(\mu) \cdot D(0)] / \sum_{i=1}^n [p(\mu) \cdot D(0)]_i$ be the vector of proportions that is obtained by running the process for one cycle with mutation rate zero starting from initial population proportions of $p(\mu)$. Then we have

$$\frac{1}{\bar{p}_1(\mu)} \leq \frac{p_1(\mu) \cdot a + (1 - p_1(\mu)) \cdot b}{p_1(\mu) \cdot a} = 1 + \frac{1 - p_1(\mu)}{p_1(\mu)} \cdot \frac{b}{a},$$

where the inequality follows because the expression in the middle corresponds to a situation where the fraction of the population that does not play s_1 plays the second best, namely s_2 .

Since $1/p_1(\mu) = 1 + [1 - p_1(\mu)]/p_1(\mu)$ and $1 + [1 - p_1(\mu_k)]/p_1(\mu_k) \geq \varepsilon$, we conclude that for every μ_k

$$\frac{1}{p_1(\mu_k)} \geq \frac{1}{\bar{p}_1(\mu_k)} + \varepsilon \left(1 - \frac{b}{a}\right). \quad (2)$$

But since $\lim_{k \rightarrow \infty} D(\mu_k) \rightarrow D(0)$ and $p(\mu_k) \cdot D(\mu_k) = p(\mu_k)$ it follows that $\lim_{k \rightarrow \infty} 1/p_1(\mu_k) = 1/\bar{p}_1(\mu_k)$, which contradicts Eq. (2). QED

This completes the proof of the claim.

We return now to the proof the Theorem 7.7.

Given a mutation rate μ and a sequence of periods consisting of N consecutive cycles, we define a t -process to be one where the initial proportions of the strategies are $p(\mu)$ and where there are mutations at rate μ up to, but not including, period t , and from period t to the end there are no mutations. We denote by γ_t , $1 \leq t \leq Nm$ the growth rate along the segment in a t -process. Formally define $A^t(\mu) \equiv \prod_{j=1}^{t-1} A(j \bmod m, \mu) \cdot \prod_{j=1}^{Nm-t} A(j \bmod m,$

0), and $\gamma_t \equiv \gamma_t(\mu) \equiv \sum_{i=1}^n [p(\mu) \cdot A^i(\mu)]_i$. Note that, for convenience, we have defined γ_t to be the rate of growth along the segment and not the per-period growth rate. With this definition γ_{Nm} is the growth along the segment with mutation rate μ while γ_1 is the growth along the segment with no mutations.

Now choose $\bar{\mu}$ so that for every $0 < \mu < \bar{\mu}$ the proportion of the population that plays s_1 is close to 1 all along the cycle. (How close is determined later.) It is possible to pick $\bar{\mu}$ in such a way because $\lim_{\mu \rightarrow 0} p_1(\mu) = 1$, by the claim, and because the cycle is of finite length m . Since $p(\mu) \cdot D(\mu) = p(\mu)$ it follows that the proportion that plays s_1 is close to 1 all along the segment.

We now show that given such a $\bar{\mu}$ we can pick N so that

$$\frac{\gamma_1}{\gamma_{Nm}} > 1. \tag{3}$$

We prove Eq. (3) by establishing a lower bound on γ_t/γ_{t+1} for every t such that $Nm - t$ is sufficiently large.

Let $L_t(t_2)$ denote the size of the population at the end of period t_2 in a t_1 -process. Let $M = Nm$. Since $L_t(t) = L_{t-1}(t)$ we have

$$\frac{\gamma_t}{\gamma_{t+1}} = \frac{L_t(M)}{L_{t+1}(M)} = \frac{L_t(M)}{L_t(t)} \bigg/ \frac{L_{t+1}(M)}{L_{t+1}(t)}.$$

Let p be the proportion of the population that plays s_1 in the period t after replication but before mutation. Let $\bar{t} = t + m - t \pmod{m}$ and let $t' = (M - \bar{t})/m$. Let N_1 be a number such that for every $1 \leq i < i' \leq n$ and for every $1 \leq l \leq m$,

$$\prod_{j=i}^m a_{i,j} \left[\prod_{k=1}^m a_{i,k} \right]^{N_1} \geq \prod_{j=i}^m a_{i',j} \left[\prod_{k=1}^m a_{i',k} \right]^{N_1}.$$

Assume that $t' \geq N_1$. It follows that

$$\frac{L_t(M)}{L_t(t)} \geq pc_1 a^{t'} + (1-p)c_2 d^{t'}, \tag{4}$$

where $c_1 \equiv \prod_{j=t+1 \pmod{m}}^m a_{1,j}$ and $c_2 \equiv \prod_{j=t+1 \pmod{m}}^m a_{n,j}$ and $d \equiv \prod_{j=1}^m a_{n,j} = \min_i \prod_{j=1}^m a_{i,j}$.

The RHS of Eq. (4) is a lower bound for the LHS because it estimates the growth of the population when the fraction that does not play s_1 plays the worst strategy, s_n .

$$\frac{L_{t+1}(M)}{L_{t+1}(t)} \leq [p(1-\mu) + (1-p)\mu]c_1 a^{t'} + (1-[p(1-\mu) + (1-p)\mu])c_3 b^{t'}, \tag{5}$$

where $c_3 \equiv \prod_{j=t+1 \pmod{m}}^m a_{2,j}$.

The RHS of Eq. (5) is an upper bound for the LHS because it measures the growth of the population when the fraction that does not play s_1 plays the second best strategy s_2 .

Since $a > b \geq d$ it follows by simple algebra that if p is sufficiently close to 1 and t' is sufficiently high then the ratio between the RHS of Eq. (4) and the RHS of Eq. (5) is larger than $1/(1 - \mu/2)$.

So let K be a number such that $\gamma_t/\gamma_{t+1} > 1/(1 - \mu/2)$ for $t \leq M - K$. Since $M - K$ is finite and since $a_{i,j} \neq 0$ there exists some number $k > 0$ such that $\gamma_{M-K+1}/\gamma_M > k$. We now have

$$\frac{\gamma_1}{\gamma_M} = \frac{\gamma_1}{\gamma_2} \cdot \frac{\gamma_2}{\gamma_3} \cdots \frac{\gamma_{M-k}}{\gamma_{M-k+1}} \cdot \frac{\gamma_{M-k+1}}{\gamma_M}.$$

Clearly if M is large enough then $\prod_{i=1}^{M-k} \gamma_i/\gamma_{i+1} \geq (1/1 - \mu/2)^{M-k} > 1/k$. It follows that $\gamma_M < \gamma_1$. Let $\gamma(p, \mu)$ denote the growth rate on the segment consisting of N cycles, when the initial proportions of the strategies are p and the mutation rate is μ . Thus $\gamma_M = \gamma(p(\mu), \mu)$ and $\gamma_1 = \gamma(p(\mu), 0)$. It is easy to see that $\gamma(e_1, 0) > \gamma(p(\mu), 0)$ and therefore we have $\gamma(0) = [\gamma(e_1, 0)]^{1/M} > [\gamma(p(\mu), 0)]^{1/M} > [\gamma(p(\mu), \mu)]^{1/M} = \gamma(\mu)$. We have shown that zero is a strict local maximum of $\gamma(\cdot)$. An argument that is similar to the one given in the proof of Proposition 4.3 establishes that $\gamma(\cdot)$ is decreasing in a neighborhood of zero. QED

REFERENCES

- BELLMAN, R. (1974). *Introduction to Matrix Analysis*. New York: McGraw-Hill.
- BLUME, L. E. (1993). "The Statistical Mechanics of Strategic Interaction," mimeo, Cornell University.
- CABRALES, A. (1993). "Stochastic Replicator Dynamics," mimeo, University of California, San Diego.
- CANNING, D. (1989). "Social Equilibrium," mimeo, Pembroke College, Cambridge University.
- ELLISON, G. (1992). "Learning, Local Interaction and Coordination," Mimeo, MIT.
- FOSTER, D., AND YOUNG, P. (1990). "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology* **38**, 219-239.
- FUDENBERG, D., AND HARRIS, C. (1992). "Evolutionary Theory with Aggregate Shocks," *J. Econ. Theory* **57**, 420-441.
- FUDENBERG, D., AND KREPS, D. (1988). "A Theory of Learning, Experimentation, and Equilibrium in Games," mimeo, MIT and Stanford.
- FUDENBERG, D., AND LEVINE, D. (1990). "Steady State Learning and Self-Confirming Equilibrium," unpublished manuscript.
- GILLESPIE, J. H. (1981). "Mutation Modification in a Random Environment," *Evolution* **35**, 468-476.
- GUTIN (1992). "Why Bother," *Discovery*, June.
- HOLSINGER, K. E., AND FELDMAN, M. W. (1983). "Modifiers of Mutation Rate: Evolutionary Optimum with Complete Selfing," *Proc. Natl. Acad. Sci.* **78**, 6772-6784.
- KANDORI, M. (1991). "Independent Mutations and Long Run Equilibria in Games: A Continuous Time Formulation," mimeo, Princeton University.
- KANDORI, M., MAILATH, G., AND ROB. R. (1993). "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* **61**, 29-56.
- KANDORI, M., AND ROB. R. (1992). "Evolution of Equilibria in the Long Run: A General Theory and Applications," CARESS WP 92-06R, University of Pennsylvania.
- KARLIN, S., AND MCGREGOR, J. (1974). "Towards a Theory of the Evolution of Modifier genes," *Theoretical Population Biology* **5**, 59-103.
- KATO, T. (1982). *A Short Introduction to Perturbation Theory for Linear Operators*. New York: Springer-Verlag.
- KIMURA, M. (1960). "Optimum Mutation Rate and Degree of Dominance as Determined by the Principle of Minimum Genetic Load," *J. Genet.* **57**, 21-34.

- KIMURA, M. (1967). "On the Evolutionary Adjustment of Spontaneous Mutation Rates," *Genet. Res. Cambridge* **73**, 1-18.
- LEIGH, E. (1970). "Natural Selection and Mutability," *Amer. Naturalist* **104**, 301-305.
- LIBERMAN, U., AND FELDMAN, M. W. (1986). "Modifiers of Mutation Rate: A General Reduction Principle," *Theoretical Population Biology* **30**, 125-142.
- MAYNARD SMITH, JOHN (1978). "The Evolution of Sex." Cambridge: Cambridge University Press.
- MAYNARD SMITH, JOHN (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- NACHBAR, J. (1990). "'Evolutionary' Selection Dynamics in Games: Dynamics and Limit Properties," *Int. J. Game Theory* **19**, 59-89.
- ROBSON, A. (1993). "A Biological Basis For Expected and Non-Expected Utility," mimeo, University of Western Ontario.
- SAMUELSON, L. (1991). "How to Tremble if You Must." SSRI WP 9122, University of Wisconsin.
- TWOMEY, M. J., AND FELDMAN, M. W. (1990). "Mutation Modification with Multiplicative Fertility Selection," *Theoretical Population Biology* **37**, 320-342.
- YOUNG, H. P. (1993). "The Evolution of Conventions," *Econometrica* **61**, 57-84.