

Stochastic Evolutionary Game Dynamics

Dean Foster

and

Peyton Young

Reprinted from THEORETICAL POPULATION BIOLOGY

Vol. 38, No. 2, October 1990

All Rights Reserved by Academic Press, New York and London

Printed by Catherine Press, Ltd., Tempelhof 41, B-8000 Brugge, Belgium

Copyright 1990 by Academic Press, Inc.

The concept of an evolutionary stable strategy (ESS) is a useful tool for studying the dynamics of natural selection. One of its limitations, however, is that it does not capture the notion of long-run stability when the system is subjected to stochastic effects. We define the concept of stability in a stochastic dynamical system, and show that it differs from both the traditional ESS and the concept of an attractor in a dynamical system. The stochastically stable set may be computed analytically using recent advances in potential theory.

1. Introduction

The concept of an evolutionary stable strategy (ESS) has become one of the principal tools for analyzing the dynamics of natural selection. The essential idea is that a population is “stable” if it cannot be invaded by a small number of individuals playing a different strategy (Maynard Smith and Price, 1973).¹ Since small perturbations due to mutations and other chance events are almost inevitable, stability requires that any small deviation from the equilibrium state be self-correcting. The system should, over time, evolve back to the equilibrium situation.

As various authors have pointed out, however, this definition does not capture the full meaning of dynamical stability. We need to view the evolution of population frequencies as a dynamical system. In this framework in ESS represents only *one* type of stable outcome, other asymptotic regimes such as limit cycles and strange attractors must also be considered (Taylor and Jonker, 1978; Zeeman, 1979; Schuster and Sigmund, 1983, 1986, Sigmund, 1987).

In this paper we shall argue that *neither* evolutionarily stable strategies nor the more general idea of attractors is quite the right concept of dynamical stability in a biological context. What is missing is an adequate account of stochastic effects. This is not to say that stochastic elements are ignored in these models, only that they are not fully taken into account. For example, the ESS criterion says that any small perturbation to the system will eventually die out. The limitation of this definition is that it treats each perturbation as if it were an isolated event. In reality, a system is *continually* being subjected to small perturbations that arise through mutations, as well as ordinary chance events that affect the reproductive success of individuals. This distinction between continual and isolated stochastic influences is fundamental. If

¹ For excellent reviews of ESS see Hines (1987) and Hofbauer and Sigmund (1988).

each successive perturbation dies out (or almost dies out) before the next one is felt, then obviously the system can never move very far from its “equilibrium” position, and the ESS definition of stability remains valid.

But this situation must be considered quite exceptional. If stochastic effects operate continually on the system, then even if they are arbitrarily small they may qualitatively change the long-run outcome. The reason is that, over the long run, it is likely that some succession of perturbations will accumulate and kick the system out of any immediate locus of an ESS. How soon the system eventually returns to this particular ESS depends on the global structure of the dynamical process. In a stochastic setting, then, the ESS condition is insufficient as a criterion of stability because it is only a *local* criterion.

The notion of an attractor in dynamical systems theory comes closer to a full description of dynamical stability. Yet even this model is not sufficient for two reasons. First, the limiting behavior of the system may depend on the assumed starting point. Second, only certain states in the attractor set may be stochastically stable. Which states are “selected” depends on the specific structure of the stochastic process.

In this paper we propose a general definition of stability in a stochastic dynamical system.² Roughly speaking, a state P is a *stochastically stable equilibrium* (SSE) if, in the long run, it is nearly certain that the system lies within every small neighborhood of P as the noise tends slowly to zero.³ More generally the *stochastically stable set* (SSS) is the set of states S such that, in the long run, it is nearly certain that the system

² A different approach is taken by Hines (1980, 1982) who studies stability issues created by mutation.

³ If one had an exact estimate of the noise term (say from biological data) then there would be no need to consider the limit as the noise tends to zero: we could study the behavior of the limiting distribution as

lies within every open set containing S as the noise tends slowly to zero.

The stochastically stable set is always nonempty. But it is by no means equivalent to the set of evolutionary stable strategies even when the latter exist. Often only a subset of the evolutionarily stable strategies are stochastically stable. In some cases *none* of them may be stochastically stable.

Our major point is to show that the introduction of stochastic effects may qualitatively change the asymptotic behavior of an evolutionary system. Second, we give a precise definition of the stochastically stable set, and show that it minimizes a suitably defined potential function. The potential function can be used to explicitly compute the stochastically stable set, as we show by example.

2. Stochastic Evolutionary Game Dynamics

We begin with the following model due to Taylor and Jonker (1978). Let a population consist of finitely many distinct phenotypes $i = 1, 2, \dots, n$. Each phenotype is associated with a fixed “strategy.” Let $\mathbf{p}(t) = [\rho_1(t), \dots, \rho_n(t)]^T$ denote the relative proportions of the n phenotypes in the total population at time t . $\mathbf{p}(t)$ is a *state* of the system. The model of selection is as follows. Individuals continually meet each other in random encounters. Whenever an individual of type- i meets an individual of type- j , the “payoff” to i is a_{ij} where the payoff is measured in terms of the change in i 's Darwinian Fitness: after an encounter between i and j , i 's reproductive rate is increased by the amount a_{ij} .

time goes to infinity. In the absence of information on the variance, however, it makes sense to consider the limiting behavior of the system when the noise is very small but nonvanishing.

Given the proportions $\mathbf{p}(t)$, and assuming a large number of (instantaneous) random encounters between types, the current rate of increase for i can be written

$r_i(t) = \sum_{j=1}^n \mathbf{a}_{ij} \rho_j$. The *average* rate of increase of the whole population is

$r_i(t) = \sum_{i=1}^n \sum_{j=1}^n \rho_i(t) \mathbf{a}_{ij} \rho_j(t)$. The relative rate of increase in the frequency of i is given by the replicator equation

$$\dot{\rho}_i(t) / \rho_i(t) = (\mathbf{A}\mathbf{p}(t))_i - \mathbf{p}(t)^T \mathbf{A}\mathbf{p}(t), \quad (1)$$

where A is the $n \times n$ matrix of payoffs (\mathbf{a}_{ij}).

The biological model on which the replicator equation is based is inherently stochastic in nature. Obviously, not every encounter between a type- i individual and a type- j individual results in *exactly* the same change in fitness. In reality, the change in reproductive fitness resulting from chance encounters is a random variable, and \mathbf{a}_{ij} is its expectation. If the population of each type is very large, and encounters are frequent in every brief time period, then the variance of \mathbf{a}_{ij} will be very small, so Eq. (1) will be a good approximation of the path that the system follows at each time t . Nevertheless, it is only an approximation. Furthermore, a real biological system is subject to all sorts of other random effects (in addition to changes in fitness resulting from encounters) such as variability in mating success, foraging success, infant mortality, and so forth.

While most authors acknowledge the existence of such stochastic influences, they have generally assumed that their effect on the outcome is so small as to be negligible. For example, Peck and Feldman (1988) assert that “the population size and the number of offspring produced per female are sufficiently large to allow us to

ignore stochastic effects.” While this assumption may be reasonable in some situations, we shall demonstrate that, in general, even if the stochastic effects are *arbitrarily* small, they may qualitatively change the asymptotic behavior of the system. Hence they cannot be ignored.

The stochastic terms arise from a variety of factors. First, there is natural variability in the payoff rate a_{ij} that results from environmental influences. Second, there is variability in the number of individuals of type i that actually do meet individuals of type j in any given time period. The relative probability that an i meets a j is of course $\rho_i\rho_j$, but the number of such interactions fluctuates: If, for example, there is perfect mixing of the types, then the number of individuals of each type in the next generation has a multinomial distribution. Third, there is background mutation, and possibly immigration of individuals from other gene pools.⁴

Assuming that the populations and the number of interactions per period are large, these sources of variability can be well-approximated (on the interior of the simplex $\sum \rho_i(t) = 1$) by a continuous-time, continuous space Wiener process. The boundary behavior of this process is quite complex, owing to the fact that the variability in the number of i - j interactions tends to zero as either ρ_i or ρ_j tends to zero. Indeed, with finite populations and no mutations, it is certain that the process will eventually hit the boundary because some phenotype dies out. Once this happens the process can never leave the boundary again.

This assumption seems biologically unrealistic, however. In practice, background mutation and immigration, however minute they may be, keep each ρ_j bounded

⁴ For consideration of these effects from a population genetics perspective see Crow and Kimura (1970).

away from zero. Thus it makes sense to study the behavior of the process in some interior envelope of the state space: $\mathcal{S}_\Delta = \{\sum p_i(t) = 1, p_i(t) \geq \Delta > 0 \text{ for all } i\}$, for some small $\Delta > 0$. On this subspace, the noise is well-approximated by a Wiener process of the form

$$dp_i(t) = p_i(t) \left((A\mathbf{p}(t))_i dt - \mathbf{p}(t)^T A\mathbf{p}(t) dt + \sigma (\Gamma(\mathbf{p}) d\mathbf{W}(t))_i \right). \quad (2)$$

Here $W(t)$ is a continuous, white-noise process with mean zero and unit rate variance covariance matrix. $\Gamma(\mathbf{p})$ is continuous in \mathbf{p} and has the property that $\mathbf{p}^T \Gamma(\mathbf{p}) = [0, 0, \dots, 0]^T$. Mutation and immigration keep the process within the space \mathcal{S}_Δ almost all of the time, and the boundary acts, to a good approximation, as if it were reflecting.⁵

Our goal is to study the asymptotic behavior of (2) as σ converges to zero. In particular we are interested in the connection between the limiting behavior of (2) and the ESS of the system (if any exist). We shall find that the introduction of a vanishingly small noise term often causes the system to “select” among the ESS: some of the ESS may be stochastically stable, while others are not. Furthermore, in some systems, none of the ESS are stochastically stable.

A precise definition and characterization of the stochastically stable set is deferred to the next section. First we illustrate the ideas by several elementary examples.

EXAMPLE 1. Consider the following symmetric payoff matrix:

⁵ Roughly speaking, the boundary is “reflecting” if the process bounces off of it according to the equal angles law. This is how it would be simulated in discrete time for example. See Karlin and Taylor (1975) for a more precise statement.

$$A = \begin{bmatrix} 4 & 0 \\ 0 & 8 \end{bmatrix}.$$

This game describes a situation where being like everyone else is desirable. If most of the population is playing strategy 1 then it is better to play 1, whereas if most of the population is playing 2, then it is better to also play 2. The population on the whole grows faster, however, if everyone plays 2. The ESS of this game are the states all-1 and all-2, as the reader may easily verify. These are also the Nash equilibria of the two-person game A and the attractors of the corresponding dynamical system (1).⁶ (6) [6 – If both players play strategy one with probability 2/3, then this is a Nash equilibrium, but it is unstable.]

Now let us examine the asymptotic behavior of the system when a small noise term is introduced. Let $p(t)$ be the proportion of the population playing 1 and $1 - p(t)$ the proportion playing 2. Pick $\Delta > 0$, and consider the state space $p \in [\Delta, 1 - \Delta]$. Assume for simplicity that the noise is uniform over the state space and over time. Thus we are considering the equation

$$dp(t) = p(t) \left(4p(t) - 4p^2(t) - 8(1-p(t))^2 \right) dt + \sigma dW(t), \quad (3)$$

where $W(t)$ is normal $(0, t)$. We shall study the evolution of the system as the variance σ^2 becomes arbitrarily small.

⁶ If both players play strategy one with probability $\frac{2}{3}$, then this is a Nash equilibrium, but it is unstable.




FIGURE 1

Figure 1 shows a simulated sample path for i between zero and 100 with the process starting at $\rho_0 = \frac{1}{2}$ and $\sigma = 0.6$. As t increases there are frequent transitions from a population playing nearly all-1 to a population playing nearly all 2. The system spends more time near all-2 than it does near all-1 but it does not appear to *converge* to a Δ -neighborhood of the all-2 state.

This is confirmed when we examine the limiting distribution of $p(t)$. As $t \rightarrow \infty$, the probability of finding the population near any given value $\rho \in (\Delta, 1 - \Delta)$ is positive,

and the limiting distribution $f_{\sigma}(\mathbf{p})$ is independent of the starting point of the paths.⁷ (This is a general feature of many stochastic models, and one of their great analytical advantages: the asymptotic behavior of the system does not depend in any essential way on the initial state.) Figure 2 shows the distribution $f_{\sigma}(\mathbf{p})$ for the σ -values 0.6, 0.4, and 0.2. As σ goes to zero, the probability becomes increasingly concentrated near $\rho = \Delta$. In other words, the state $\rho = \Delta$ (all type-2 except for the new mutants) is a stochastically stable equilibrium.




FIGURE 2

⁷ The existence of the limiting distribution is assured because $\Gamma(\mathbf{p})$ is bounded away from zero for all

Why is the mostly all-2 state selected over the mostly all-1 state, even though both are ESS and both are stable attractors? Intuitively the explanation is the following. Suppose that most of the population is playing strategy 1 (i.e., $p \cong 1 - \Delta$). Against this population an organism playing 1 will get a payoff of four (on average) whereas an organism playing strategy 2 will get a payoff of zero. Since most organisms are playing 1, the average payoff for everyone is about four. Thus, there is selection pressure against any organism playing strategy 2. In order for the population to switch to playing strategy 2, the number of type-2 players must grow (due to random drift) in spite of being disadvantaged. Once the number playing 2 is more than one-third of the population, selection will tend to drive the population towards type-2.

A similar story holds for the transition from mostly type-2 players to mostly type-1. But this transition is harder because the selection pressure against it is stronger. Hence it is less likely. In summary, the population will fluctuate between being mostly type-1 and being mostly type-2. But when most of the population is playing 2, it takes longer to switch to 1 than vice versa. Furthermore, the difference in transition times becomes larger as the noise term becomes weaker. So, as $\sigma \rightarrow 0$, the probability that almost everyone is playing strategy 2 approaches unity.




FIGURE 3

Example 2. Consider the matrix

$$A = \begin{bmatrix} 4 & 0 & 8 \\ 0 & 8 & 0 \\ 0 & 7.5 & 0 \end{bmatrix}$$

This is like Example 1 except that a third strategy has been added. This new strategy is dominated by the second one: playing strategy 2 gives at least as high a payoff as playing strategy 3 no matter what strategy the opponent plays. While this inferior strategy is not itself selected, it has the effect of destabilizing the previous equilibrium and shifting the outcome from nearly all-2 to nearly all-1. The reason is this: the presence of a small number of 3's gives an extra boost to the 1's, but not to themselves or to the 2's. If the 2's start to grow, then they help boost the number of 3's, which in turn boosts the number of 1's. Thus a chance movement away from all-1, say towards 2, triggers the "catalytic" sequence $2 \rightarrow 3 \rightarrow 1$, and the system is pushed back toward the all-1 state. The basins of attraction are such that it is harder for the system to drift away from the mostly all-1 state than from the mostly all-2 state.

This example shows that the SSE is not necessarily the ESS having the highest.

growth rate. Stochastic stability does not imply that the reproductive rate is maximized.

Example 3. In this example we show that there may exist no SSE. Rather, there is a set of points that is asymptotically stable. Consider the matrix

$$A = \begin{bmatrix} 3 & 1 & 4 \\ 4 & 3 & 1 \\ 1 & 4 & 3 \end{bmatrix}$$

The unique Nash equilibrium of this game is $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) = \mathcal{N}$. This is not an ESS,

however, because from any interior point other than \mathcal{N} , the deterministic dynamical system spirals outward and converges to a limit cycle (see Fig. 3). The specific shape of this cycle depends, of course, on the value of the mutation rate Δ . As Δ becomes smaller, it approximates the boundary of \mathcal{S}_Δ more and more closely.

Now fix $\Delta > 0$ and introduce a small level of noise. The limiting density of the process is concentrated in a small envelope containing this limit cycle, and the limit cycle is the asymptotically stable set of the stochastic process. However, as Δ becomes smaller the limiting density is increasingly concentrated in the “corners” of the limit cycle. In other words, if both Δ and σ are sufficiently small, the probability is almost one that the process is within ε of one of the three corners. The dynamics of the process are as follows. Typically it is near a corner and lingers there for a long time. Occasionally, it skitters alongside the boundary toward the next corner, where it loiters for another long period of time.

EXAMPLE 4. In this example we show that even if there is a unique ESS, it is not

necessarily an SSE. Add a fourth phenotype to Example 3 to obtain the following matrix:

$$A = \begin{bmatrix} 3 & 1 & 4 & -1 \\ 4 & 3 & 1 & -1 \\ 1 & 4 & 3 & -1 \\ -1 & -1 & -1 & 0 \end{bmatrix}$$

This new type imposes a negative payoff on each of the others (and *vice versa*) but it is harmless against itself. If type 4 is sufficiently numerous, then it has the highest relative reproduction rate, and the process converges to all-4. Therefore it is an ESS, and in fact it is the unique one. If, on the other hand, type 4 falls below some critical proportion of the population, then they are disadvantaged relative to the others and the process converges to the limit cycle described in the previous example. Furthermore this second basin of attraction is the dominant one for the stochastic process. As the noise goes to zero, the process is almost certain to be in the limit-cycle regime, and type 4 will almost have died out. Thus the process is typically nowhere near the unique ESS of the system.

3. THE STOCHASTICALLY STABLE SET

The preceding examples show that stochastic stability is a more refined idea than either attractors or evolutionarily stable strategies. We now give a formal definition.

Let $f_\sigma(\cdot)$ be the limiting density of $\mathbf{p}(t)$ as $t \rightarrow \infty$. This limit always exists because we assume that $\Gamma(\mathbf{p})$ is bounded away from zero for all $\mathbf{p} \in \mathcal{S}_j$. If the variance σ^2 could

be estimated precisely, then it would suffice, of course, to study the distribution f_σ . In this case the stochastically stable set could be defined, somewhat loosely, as the smallest set of states that is 95 % or 99 % probable. In the absence of such information, it makes sense to study the behavior of f_σ as $\sigma \rightarrow 0$. Unfortunately, this limit may not exist for general dynamical systems. (An open question is whether it exists in dynamical systems of the specific form (2).) Hence we shall consider the limit superior $\overline{\lim}_{\sigma \rightarrow 0} f_\sigma$, which always does exist.

DEFINITION. The population vector \mathbf{p}^* is *stochastically stable* if, as $\sigma \rightarrow 0$, the limiting density assigns positive probability to every small neighborhood of \mathbf{p}^* ; more precisely,

$$\forall \varepsilon > 0 \overline{\lim}_{\sigma \rightarrow 0} \int_{\mathcal{N}_\varepsilon(\mathbf{p}^*)} f_\sigma(\mathbf{p}) d\mathbf{p} > 0,$$

where

$$\mathcal{N}_\varepsilon(\mathbf{p}^*) = \{\mathbf{p} : |\mathbf{p} - \mathbf{p}^*| < \varepsilon\}.$$

A stochastically stable population vector always exists, as we shall presently show. Furthermore, the *set* of stochastically stable vectors has the following characterization: it is the smallest closed set S such that, as $\sigma \rightarrow 0$ and $t \rightarrow \infty$, it is highly likely that the system is arbitrarily close to S independently of the initial state.

THEOREM 1. The set of stochastically stable vectors is nonempty, and it is the unique minimal closed set S satisfying

$$\lim_{\varepsilon \rightarrow 0} \lim_{\sigma \rightarrow 0} \int_{\mathcal{N}_\varepsilon(S)} f_\sigma(\mathbf{p}) d\mathbf{p} = 1,$$

(4)

where

$$\mathcal{N}_\varepsilon(\mathcal{S}) = \{\mathbf{p} : \exists \mathbf{x} \in \mathcal{S}, |\mathbf{x} - \mathbf{p}| < \varepsilon\}$$

The proof of this result is given in the Appendix.

4. COMPUTATION OF THE STOCHASTICALLY STABLE SET

We next turn to the question of how to compute the SSS. This can be done by finding the minimum of a suitably defined potential function.

In general, suppose that $\mathbf{p}(t)$ satisfies the stochastic differential equation

$$d\mathbf{p}(t) = \mu(\mathbf{p}(t))dt + \sigma\Gamma(\mathbf{p}(t))dW(t),$$

where $W(t)$ is a Gaussian noise, $\Gamma(\cdot)$ is bounded away from 0, and both $\mu(\cdot)$ and $\Gamma(\cdot)$ are continuous. Then there exists a function $\mathcal{K}(\sigma)$ such that the following limit exists for all points \mathbf{p} (see Freidlin and Wentzell, p. 190, 1984):

$$\lim_{\sigma \rightarrow 0} (-2\sigma^2 \log f_\sigma(\mathbf{p}) - \mathcal{K}(\sigma)) = \mathcal{U}(\mathbf{p}). \quad (5)$$

$\mathcal{U}(\mathbf{p})$ can be viewed as a potential function. In particular it has the property that- the stochastic process evolves to states having minimum potential.

THEOREM 2. *The stochastically stable set is contained in the set of minimum potential.*

The proof of this result is given in the Appendix.

For the one dimensional case (i.e., two phenotypes), $U(p)$ can be explicitly computed from the formula

$$\mathcal{U}(p) = -\int_0^p \mu(x) / \Gamma(x)^2 dx.$$

For the first example considered in this paper, μ and Γ are determined by Eq. (3). Thus $U(p)$ can be computed as follows:

$$\mathcal{U}(p) = -\int_0^p x(4x - 4x^2 - 8(1-x)^2) dx = 4p^2 - (20/3)p^3 + 3p^4.$$

$U(p)$ has a unique minimum at $p = 0$, so the state of nearly all type-2 phenotypes is the SSE.

5. CONCLUSION

In this paper we have argued that evolutionary stability means long-run viability when the dynamical system is subjected to small but non-vanishing perturbations. We have seen that stability in this sense does not necessarily coincide with the ESS or (in some cases) With *any* of the ESS. It would be interesting to examine the implications of this model when the noise term is derived from specific data on mutation rates and the variability of reproductive success.

APPENDIX

Here we prove Theorem 1. Fix the state space to be

$\mathcal{S}_\Delta \equiv \left\{ \sum \rho_i(t) = 1, \rho_i(t) \geq \Delta \text{ for all } i \right\}$, for some small $\Delta > 0$. First we establish the following.

LEMMA 1. *The set S of all stochastically stable vectors satisfies*

$$\lim_{\varepsilon \rightarrow 0} \lim_{\sigma \rightarrow 0} \int_{\mathcal{N}_\varepsilon(S)} f_\sigma(\mathbf{p}) d\mathbf{p} = 1. \quad (4)$$

Proof. Pick $\alpha > 0$, and let W be the complement of the α -neighborhood of the stochastically stable set. W is compact. Consider the collection of sets

$\left\{ \mathcal{N}_{\varepsilon(\mathbf{q})}(\mathbf{q}) \right\}_{\mathbf{q} \in W}$ where $\varepsilon(\mathbf{q})$ satisfies

$$\overline{\lim}_{\sigma \rightarrow 0} \int_{\mathcal{N}_\varepsilon(S)} f_\sigma(\mathbf{p}) d\mathbf{p} = 1.$$

This collection of open sets covers W and hence a finite subcovering exists.

Call this subcovering $\left\{ \mathcal{N}_i \right\}_{i=1, n}$. Consider

$$1 \leq \int_{\mathcal{N}_\alpha(\mathbf{q})} f_\sigma(\mathbf{p}) d\mathbf{p} = 0.$$

Rearranging terms and taking the limit superior we get

$$1 - \lim_{\sigma \rightarrow 0} \int_{\mathcal{N}_\alpha(S)} f_\sigma(\mathbf{p}) d\mathbf{p} \leq \sum_{i=1}^n \int_{\mathcal{N}_i} f_\sigma(\mathbf{p}) d\mathbf{p} = 0.$$

But

$$\overline{\lim}_{\sigma \rightarrow 0} \sum_{i=1}^n \int_{\mathcal{N}_i} f_{\sigma}(\rho) d\rho \leq \sum_{i=1}^n \overline{\lim}_{\sigma \rightarrow 0} \int_{\mathcal{N}_i} f_{\sigma}(\mathbf{p}) d\mathbf{p} = 0.$$

Thus,

$$\overline{\lim}_{\sigma \rightarrow 0} \int_{\mathcal{N}_{\alpha}(S)} f_{\sigma}(\mathbf{p}) d\mathbf{p} \geq 1.$$

Since the limit superior must be less than or equal to one, we see that the limit exists and equals one.

From (4) it follows that the stochastically stable set S is nonempty. Next we shall show that S is the unique smallest closed subset of \mathcal{S}_{Δ} satisfying (4).

It is left to the reader to check that S is closed. Next we shall show that no closed set that is properly contained in S does satisfy (4). Let $\mathcal{K} \subset \mathcal{S}$ where \mathcal{K} is closed. Consider a point \mathbf{p}^* which is in the stochastically stable set but not in \mathcal{K} . Because \mathcal{K} is closed, there exists an ε -neighborhood around \mathbf{p}^* such that this neighborhood does not intersect \mathcal{K} . Because \mathbf{p}^* is stochastically stable, we know that

$$\overline{\lim}_{\sigma \rightarrow 0} \int_{\mathcal{N}_{\varepsilon/2}(\mathbf{p}^*)} f_{\sigma}(\mathbf{p}) d\mathbf{p} > 0.$$

$\mathcal{N}_{\varepsilon/2}(\mathcal{K})$ is contained in the set $|\mathbf{p} - \mathbf{p}^*| \geq \varepsilon/2$, so

$$\lim_{\sigma \rightarrow 0} \int_{\mathcal{N}_{\varepsilon/2}(K)} f_{\sigma}(\mathbf{p}) d\mathbf{p} \leq \lim_{\sigma \rightarrow 0} \int_{|\mathbf{p}-\mathbf{p}^*| \geq \varepsilon/2} f_{\sigma}(\mathbf{p}) d\mathbf{p} < 1.$$

Thus the set K does not satisfy (4).

Finally, there can be no other minimal closed set satisfying (4), because if two sets satisfy (4) then so does their intersection. This concludes the proof of Theorem 1.

THEOREM 2. *The stochastically stable set is contained in the set of minimum potential.*

Proof. We will prove that if \mathbf{p} is not in the minimum set, then \mathbf{p} is not in the stochastically stable set. Without loss of generality, assume that the minimum value of $\mathcal{U}(\cdot)$ is zero. Consider a point \mathbf{p} not contained in the minimum set of $\mathcal{U}(\cdot)$. By continuity of $\mathcal{U}(\cdot)$ there exists an ε -neighborhood around \mathbf{p} such that $\mathcal{U}(\cdot)$ is bounded below by a positive α . Pick a positive β smaller than α . There exists a point \mathbf{q} and a positive δ such that $\mathcal{U}(\cdot)$ is less than β in the δ -neighborhood of \mathbf{q} . In other words, the following holds:

$$\inf_{\mathcal{N}_{\varepsilon}(\mathbf{p})} \mathcal{U}(x) \geq \alpha \geq \beta \geq \sup_{\mathcal{N}_{\delta}(\mathbf{q})} \mathcal{U}(x) \geq \inf_{S_{\Delta}} \mathcal{U}(x) \equiv 0.$$

Choose α' and β' such that $\alpha > \alpha' > \beta' > \beta$. By (5), for all sufficiently small σ ,

$$-2\sigma^2 \log f_{\sigma}(x) - \mathcal{K}(\sigma) \geq \alpha' \quad \text{for all } x \in \mathcal{N}_{\varepsilon}(\mathbf{p}) \tag{6}$$

Therefore,

$$f_{\sigma}(\mathbf{x}) < \exp\left(-\frac{K(\sigma) + \alpha'}{2\sigma^2}\right) \quad \text{for all } x \in \mathcal{N}_{\varepsilon}(\mathbf{p}). \quad (7)$$

Similarly,

$$f_{\sigma}(\mathbf{x}) < \exp\left(-\frac{K(\sigma) + \beta'}{2\sigma^2}\right) \quad \text{for all } x \in \mathcal{N}_{\delta}(\mathbf{q}). \quad (8)$$

Now integrate both sides of (8) over $\mathcal{N}_{\delta}(\mathbf{q})$ and deduce that

$$1 \geq \int_{\mathcal{N}_{\delta}(\mathbf{q})} f_{\sigma}(x) dx \geq \exp\left(-\frac{K(\sigma) + \alpha'}{2\sigma^2}\right) V\delta^n,$$

where $V\delta^n$ is the volume of the δ -neighborhood around \mathbf{q} . Therefore

$$K(\sigma) \geq 2\sigma^2 \log(V\delta^n) - \alpha'.$$

Combining this lower bound on $K(\sigma)$ with Eq. (7) we deduce that

$$\begin{aligned} \int_{\mathcal{N}_{\varepsilon}(\mathbf{p})} f_{\sigma}(x) dx &\leq V\varepsilon^n \exp\left(-\frac{K(\sigma) + \alpha'}{2\sigma^2}\right) \\ &\leq V\varepsilon^n \exp\left(-\log(V\delta^n) - (\alpha' - \beta')/(2\sigma^2)\right) \\ &= (\varepsilon/\delta)^n \exp\left(-(\alpha' - \beta')/(2\sigma^2)\right). \end{aligned}$$

Because $\alpha' - \beta'$ is positive, we see that the right hand side converges to zero as σ converges to zero. Thus, \mathbf{p} is not in the stochastically stable set.

REFERENCES

- Crow, J., and Kimura, M. 1970. "An Introduction to Population Genetics Theory," Harper & Row, New York.
- Freidlin, M. I., and Wentzell, A.D. 1994. "Random Perturbations of Dynamical Systems," Springer-Verlag. Berlin/New York.
- Hines, W. G. S. 1980. Strategy stability in complex populations, *J. Appl. Probab.* **17**, 600-610.
- Hines, W. G. S. 1982. Mutation, perturbations and evolutionarily stable states, *J. Appl Probab.* **19**, 204-209.
- Hines, W. G. S. 1987. Evolutionary stable strategies: A review of basic theory, *Theor. Pop. Biol.* **31**, 195-272.
- Hofbauer, J., and Sigmund, K. 1988. "The Theory of Evolution and Dynamical Systems," Cambridge Univ. Press.
- Karlin, S., and Taylor, H.M. 1975. "A First Course in Stochastic Processes," Academic Press, New York.
- Maynard Smith, J., and Price, G.R. 1973. The logic of animal conflict, *Nature (London)* **246**, 15-18.
- Peck, J.R. and Feldman, M.W. 1988. Kin selection and the evolution of monogamy, *Science*, **240**, 1672-1674.

Schuster, P., and Sigmund, K. 1983. Replicator dynamics, *J. Theor. Biol.* **100**, 535-538.

Schuster, P., and Sigmund, K. 1986. Evolutionary game dynamics, *Mondes en Developpement*, **54-55**, 229-236.

Sigmund, K. 1987. Game dynamics, mixed strategies and gradient systems, *Theor. Pop. Biol.* **32**, 114-126.

Taylor, P.D., and Jonker, L.B. 1978. Evolutionarily stable strategies and game dynamics, *Math. Biosci.* **40**, 145-156.

Zeeman, E. C. 1979. Population dynamics from Same theory, in "Proc. Int. Conf. Global Theory of Dynamical Systems," pp. 491-497. Northwestern, Evanston, IL