Simple Formulas for Stationary Distributions and Stochastically Stable States

William H. Sandholm*
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706
whs@ssc.wisc.edu
http://www.ssc.wisc.edu/~whs

April 11, 2006

* I thank Larry Samuelson, an anonymous referee, and an anonymous Associate Editor for helpful comments. Financial support from NSF Grant SES-0092145 is gratefully acknowledged.

Abstract

We derive some simple formulas for limiting stationary distributions for models of stochastic evolution in two-strategy population games. As an application of these formulas, we investigate the robustness of equilibrium selection results to the assumption that the level of noise in agents' choice rules is vanishingly small.

Keywords: evolutionary game theory, stochastic stability, equilibrium selection.

1. Introduction

In this paper, we revisit stochastic evolution and equilibrium selection in binary choice games, a topic studied by Kandori, Mailath, and Rob (1993), Blume (1993, 2003), and Binmore and Samuelson (1997), among many others. Compared to the existing literature, this paper offers two innovations. First, instead of restricting attention to equilibrium selection in random matching games with two stable monomorphic Nash equilibria, one corresponding to each pure strategy, we allow nonlinear "playing the field" games possessing any finite number of Nash equilibria. Second, rather than just determining the stochastically stable state, we derive a formula that describes the entire limiting stationary distribution, and show that this formula takes an exceptionally simple form under the two best-known specifications of the agents' choice rules: the mutation model of Kandori, Mailath, and Rob (1993), and the logit model of Blume (1993).

Our stochastic stability analysis follows the approach of Binmore and Samuelson (1997), Young (1998, Sec. 4.5) and Benaim and Weibull (2003): we fix the noise level in the agents' choice rules, and consider the limit of the stationary distributions as the population size grows large. In doing so, we adopt Binmore and Samuelson's (1997) view that economic agents make mistakes at nonnegligible rates.

One reason for the common focus on vanishingly small mutation rates in evolutionary models is tractability: for instance, by taking noise levels to zero, one can take advantage of Freidlin and Wentzell's (1998) methods for computing limiting stationary distributions. A possible defense of the vanishing noise assumption is that limiting stationary distributions will be close to the stationary distributions that obtain at small noise levels; if this is true, then equilibrium selection results established in the limit are robust. Robustness of this sort is also important because it serves to mitigate the well-known waiting time critique: since the time spans necessary for the equilibrium selection results to be meaningful increase sharply as the noise level falls, results that are robust to higher noise levels are more likely to have economic relevance.

The approach to stochastic stability we take in this paper enables us to address this question of robustness in a direct fashion, by comparing stochastic stability results obtained at fixed noise levels to those obtained by taking these levels to zero. As a simple application of our formulas, we obtain a nonrobustness result: for every positive noise level ε, one can construct a game in which a monomorphic equilibrium is selected at noise levels below ε, while an interior equilibrium is selected at noise levels above ε. But we also derive two positive results. First, we prove that the "opposite" of the previous example cannot be constructed: if an interior equilibrium x^* is selected for small values of ε, then no state that is more asymmetric than x^* can be selected for larger values of ε. Second, we show that whenever the only candidates for stability are the monomorphic states, analyses of the small noise limit are robust to the introduction of nonnegligible noise levels.

2. The Model

2.1 Two-Strategy Population Games

Each agent in a large population chooses between two strategies, 0 and 1. The population state $x \in [0, 1]$ represents the proportion of agents choosing strategy 1. Payoffs are denoted by F: $[0, 1] \to \mathbf{R}^2$, where $F_i(x)$ is the payoff to strategy $i \in \{0, 1\}$ at population state x. We assume that each payoff function is piecewise continuous, and that $F_1 - F_0$ changes sign at most finitely many times on [0, 1].[1] (For some simple examples, see Section 5.)

2.2 Choice Rules

Suppose that a population of N agents recurrently plays the game F described above. At each time $t \in T^N = \{0, \frac{1}{N}, \frac{2}{N}, \dots\}$, one agent is drawn at random from the population and given the opportunity to switch strategies.[2] When an agent receives a revision opportunity, he chooses a new strategy by employing a *choice rule* $\sigma = (\sigma_{01}, \sigma_{10})$, where σ_{ij}: $[0, 1] \times \mathbf{R}^2 \to [0, 1]$. Here, $\sigma_{ij}(x, \pi)$ represents the probability that an agent playing strategy i who receives a revision opportunity opts to switch to strategy j; we assume that this probability is a continuous function of the current state x and current payoff vector π, and that σ_{01} and σ_{10} are bounded away from zero.

2.3 The Evolutionary Process

A game F, a population size N, and a choice rule σ together induce a Markov chain $\{X_t^N\}_{t \in T^N}$ on the discrete state space $U^N = \{0, \frac{1}{N}, \dots, 1\} \subset [0, 1]$. Since only one agent is given the opportunity to switch strategies in each period, $\{X_t^N\}$ is a *birth and death chain*:

---

[1] The Nash equilibria of F are (i) the interior states x at which $F_1(x) = F_2(x)$; (ii) state 0, if $F_0(0) \geq F_1(0)$; and (iii) state 1, if $F_1(1) \geq F_0(1)$. Imagine a continuous-time deterministic evolutionary dynamic on [0, 1] that respects payoffs in F. If payoffs are continuous, the stable states of the dynamic are those Nash equilibria where F_1 crosses F_0 from above, where we interpret this statement in an appropriate way at the monomorphic states. If payoffs are discontinuous, we must add to this list those states $\hat{x}$ at which both $\lim_{x \uparrow \hat{x}} (F_1(x) - F_0(x)) > 0$ (assuming $\hat{x} \neq 0$) and $\lim_{x \downarrow \hat{x}} (F_1(x) - F_0(x)) < 0$ (assuming $\hat{x} \neq 1$). In either case, the assumption that the sign of $F_1 - F_0$ changes finitely often ensures that both the number of Nash equilibria and the number of stable states are finite.

[2] Having the period length equal the inverse of the population size is the natural time scale, since it fixes each agent's expected rate of revision at one opportunity per time interval of unit length. However, since our interest in this paper is in stationary distributions, this assumption plays no role.

if the state at time t is x, then the state at time $t + \frac{1}{N}$ must be $x - \frac{1}{N}$, x, or $x + \frac{1}{N}$.

The transition probabilities of $\{X_t^N\}$ are easy to derive. For the state to increase by one increment, the randomly drawn agent must initially play strategy 0, and must decide to switch to strategy 1; hence

$$p_x \equiv P\left(X_{t+\frac{1}{N}} = x + \tfrac{1}{N} \,\middle|\, X_t = x\right) = (1 - x)\ \sigma_{01}(x, F(x)).$$

Similarly, the state decreases by one increment if the randomly drawn agent initially plays strategy 1 and switches to strategy 0:

$$q_x \equiv P\left(X_{t+\frac{1}{N}} = x - \tfrac{1}{N} \,\middle|\, X_t = x\right) = x\ \sigma_{10}(x, F(x)).$$

In the next two sections, we show that by keeping separate the terms in p_x and q_x corresponding to arrivals of revision opportunities and to choices of strategies, we can derive simple formulas for the limiting stationary distribution of $\{X_t^N\}$.

3. Limiting Stationary Distributions

Given our lower bound on the choice probabilities σ_{01} and σ_{10}, the Markov chain $\{X_t^N\}$ is aperiodic and irreducible, and so has a unique stationary distribution μ^N. This distribution describes the long run behavior of $\{X_t^N\}$ in two distinct ways: it is the limiting distribution of this process, and it also describes the limiting empirical distribution of the process with probability one (see, e.g., Durrett (2005)).

Our main result characterizes the limiting stationary distribution of $\{X_t^N\}$ in terms of the payoff function F and choice rule σ. To state this result, we define the *entropy function* h: $[0, 1] \to \mathbf{R}$ to be the continuous concave function

$$h(x) = \begin{cases} -\left(x \log x + (1-x)\log(1-x)\right) & \text{if } x \in (0,1); \\ 0 & \text{if } x = 0 \text{ or } x = 1. \end{cases}$$

We then define the function m: $[0, 1] \to \mathbf{R}$ by

$$m(x) = \int_0^x \log\left(\frac{\sigma_{01}(y, F(y))}{\sigma_{10}(y, F(y))}\right) dy \ + h(x).$$

Theorem 1 shows that for each state x, the value of $m(x)$ describes the exponential growth rate of μ_x^N relative to that of μ_0^N.

Theorem 1: *Suppose $x \in [0, 1]$ is rational. Then considering only population sizes N for which Nx is an integer, we have that*

$$\lim_{N\to\infty} \frac{1}{N} \log\left(\frac{\mu_x^N}{\mu_0^N} \right) = m(x).$$

In other words, when N is large, the ratio μ_x^N / μ_0^N *is of order* $\exp(Nm(x))$.

Proof: Fix a population size N and a population state $x \in U^N - \{0, 1\}$. (The result is trivial when $x = 0$, and a minor modification of the argument to follow proves the result when $x = 1$.) It is well known that the stationary distribution of the birth and death chain $\{X_t^N\}$ is of the form

$$\mu_x^N = \mu_0^N \prod_{k=1}^{Nx} \frac{p_{\frac{k-1}{N}}}{q_{\frac{k}{N}}}.$$

(see Durrett (2005), p. 297). Therefore,

$$\log\left(\frac{\mu_x^N}{\mu_0^N} \right) = \log\left(\prod_{k=1}^{Nx} \frac{p_{\frac{k-1}{N}}}{q_{\frac{k}{N}}} \right) = \sum_{k=1}^{Nx} \log\left(\frac{p_{\frac{k-1}{N}}}{q_{\frac{k}{N}}} \right) = \sum_{k=1}^{Nx} \log\left(\frac{p_{\frac{k}{N}}}{q_{\frac{k}{N}}} \right) + \log\left(\frac{p_0}{p_x} \right).$$

Since $x < 1$ is fixed and σ_{01} is bounded away from zero,

$$\log\left(\frac{p_0}{p_x} \right) = \log\left(\frac{\sigma_{01}(0, F(0))}{(1-x)\sigma_{01}(x, F(x))} \right)$$

is a finite constant, allowing us to conclude that

$$\begin{aligned}
\lim_{N\to\infty} \frac{1}{N} \log\left(\frac{\mu_x^N}{\mu_0^N} \right) &= \lim_{N\to\infty} \frac{1}{N} \sum_{k=1}^{Nx} \log\left(\frac{p_{\frac{k}{N}}}{q_{\frac{k}{N}}} \right) \\
&= \int_0^x \log\left(\frac{(1-y)}{y} \cdot \frac{\sigma_{01}(y, F(y))}{\sigma_{10}(y, F(y))} \right) dy \\
&= \int_0^x \left(\log\left(\frac{\sigma_{01}(y, F(y))}{\sigma_{10}(y, F(y))} \right) + \log(1-y) - \log y \right) dy \\
&= \int_0^x \log\left(\frac{\sigma_{01}(y, F(y))}{\sigma_{10}(y, F(y))} \right) dy + h(x). \quad \blacksquare
\end{aligned}$$

The idea of converting the usual formula for the stationary distribution into a Riemann sum, and then using this sum to describe the limiting stationary distribution in terms of an integral, is not new: both Binmore and Samuelson (1997) and Blume (2003) employ this technique. But by distinguishing arrivals of revision opportunities from choices of strategies in the transition probabilities p_x and q_x, we are able to express the stationary distribution as the sum of two terms: a term that depends on payoffs and the choice rule, and an entropy term. This separation is very useful in performing explicit computations, as the examples in the next section show.

To understand the role of the entropy function in the limiting stationary distribution, consider a situation in which switching probabilities are constant and are equal in each direction: $\sigma_{01} \equiv \sigma_{10} \equiv s \in (0, 1]$. In this case, evolution is driven entirely by the random assignment of revision opportunities, creating a tendency for the Markov chain $\{X_t^N\}$ to move toward the center of the unit interval. It is easy to verify that in this case, the stationary distribution μ^N is simply the binomial(N, $\frac{1}{N}$) distribution, scaled down by a factor of N so as to place its mass on the set $U^N = \{0, \frac{1}{N}, \ldots, 1\}$. According to a well-known result from large deviations theory (see Durrett (2005, p. 74)), this distribution satisfies

$$\lim_{N\to\infty} \tfrac{1}{N} \log \mu^N([x,1]) = -\big(x\log(2x) + (1-x)\log(2(1-x)\big) = h(x) - \log 2$$

whenever $x \in [\frac{1}{2}, 1)$, with virtually all of the mass in the interval $[x, 1]$ accruing from states very close to x.[3] This limit result accords exactly with Theorem 1 above.

4. Examples

We now use Theorem 1 to describe the limiting stationary distributions generated by the two best-known specifications of agents' choice rules: those of Kandori, Mailath, and Rob (1993) and Blume (1993). Each of these choice rules is parameterized by a *noise level* $\varepsilon > 0$, representing the degree to which the rule deviates from exact optimization.

Example 1: The mutation model. Kandori, Mailath, and Rob (1993) study stochastic evolution under the choice rule

$$\sigma_{ij}(x,\pi) = \begin{cases} 1-\varepsilon & \text{if } \pi_j > \pi_i, \\ \varepsilon & \text{if } \pi_j < \pi_i, \\ \text{arbitrary} & \text{if } \pi_j = \pi_i, \end{cases}$$

where $\varepsilon \in (0, 1)$. Notice that under this rule,

$$\log\left(\frac{\sigma_{01}(x,F(x))}{\sigma_{10}(x,F(x))}\right) = \log\left(\frac{1-\varepsilon}{\varepsilon}\right)\operatorname{sgn}\big(F_1(x) - F_0(x)\big).$$

whenever $F_0(x)$ and $F_1(x)$ differ. Therefore,

---

[3] Since $h(\frac{1}{2}) = \log 2$, the shift by log 2 in the previous formula ensures that $\lim_{N\to\infty} \frac{1}{N} \log \mu^N([\frac{1}{2},1]) = 0$.

$$\lim_{N\to\infty} \frac{1}{N} \log\left(\frac{\mu_x^N}{\mu_0^N}\right) = \int_0^x \log\left(\frac{1-\varepsilon}{\varepsilon}\right) \operatorname{sgn}\big(F_1(y) - F_0(y)\big)\, dy \; + h(x)$$

$$= \log\left(\frac{1-\varepsilon}{\varepsilon}\right) l(x) + h(x).$$

Here, l is the *best response potential function* for F, defined by

$$l(x) = \lambda\big([0,x] \cap B_1\big) - \lambda\big([0,x] \cap B_0\big) = 2\lambda\big([0,x] \cap B_1\big) - x,$$

where λ is Lebesgue measure and $B_i = \{y \in [0, 1]: F_i(y) > F_j(y)\}$ is the set of states at which strategy i is optimal.

The value of $l(x)$ is the difference between the measures of two sets: the set of states in the interval $[0, x]$ where strategy 1 is optimal, and the set of states in $[0, x]$ where strategy 0 is optimal.[4] Thus, the function l provides a generalization of the "mutation counting" arguments introduced by Kandori, Mailath, and Rob (1993) for the case of 2 x 2 coordination games: in essence, the function l keeps a running total of *net* number of mutations to move from state x to state 0 (cf Corrolary 2 below). The mutation rate ε determines the relative weights placed on $l(x)$ and on the entropy $h(x)$ in the limiting stationary distribution. §

Example 2: The logit model. Our other basic choice rule is the *logit choice rule* of Blume (1993):

$$\sigma_{ij}(x, \pi) = \frac{\exp(\varepsilon^{-1}\pi_j)}{\sum_{k\in\{0,1\}} \exp(\varepsilon^{-1}\pi_k)},$$

where $\varepsilon \in (0, \infty)$. Under this rule,

$$\log\left(\frac{\sigma_{01}(x, F(x))}{\sigma_{10}(x, F(x))}\right) = \log\left(\frac{\exp(\varepsilon^{-1}F_1(x))}{\exp(\varepsilon^{-1}F_0(x))}\right) = \varepsilon^{-1}\big(F_1(x) - F_0(x)\big),$$

so

$$\lim_{N\to\infty} \frac{1}{N} \log\left(\frac{\mu_x^N}{\mu_0^N}\right) = \int_0^x \varepsilon^{-1}\big(F_1(y) - F_0(y)\big)\, dy \; + h(x)$$

$$= \varepsilon^{-1} f(x) + h(x).$$

[4] To graph l, place the point of a pencil at the origin; as the point proceeds to the right, move it upward at rate 1 whenever strategy 1 is optimal, and move it downward at rate 1 whenever strategy 0 is optimal.

In this last expression, the function

$$f(x) = \int_0^x \big(F_1(y) - F_0(y)\big)\,dy$$

is the *potential function* for the game F (see Sandholm (2001)). As in the previous example, the relative weights placed on potential $f(x)$ and entropy $h(x)$ in the limiting stationary distribution are determined by the noise level ε.

Interestingly, the function $\varepsilon^{-1} f(x) + h(x)$ is also quite useful for analyzing logit evolution in potential games with more than two strategies—see Hofbauer and Sandholm (2005) and Benaim and Sandholm (2005). §

5. Equilibrium Selection with Fixed and Vanishing Noise

5.1 Stochastic Stability and Limit Stochastic Stability

We now use Theorem 1 and Examples 1 and 2 to derive equilibrium selection results. To begin, we follow Binmore and Samuelson (1997), Young (1998, Sec. 4.5), and Benaim and Weibull (2003) by fixing the noise level ε and considering the limit of the stationary distributions μ^N as the population size N grows large. We call state $x^* \in [0, 1]$ (*uniquely*) *stochastically stable* if

$$\lim_{N \to \infty} \mu^N = \delta_{\{x^*\}},$$

where the limit refers to convergence in distribution, and where $\delta_{\{x^*\}}$ represents a point mass at state x^*. Put differently, x^* is stochastically stable if nearly all the mass in the measure μ^N lies in an arbitrarily small neighborhood of x^* once N is large enough.

With this definition in hand, the corollary below follows directly from Theorem 1.

Corollary 1: *If the function m has a unique maximizer x^*, then x^* is stochastically stable.*

Like Binmore and Samuelson (1997), we believe that in most situations, errant choices are made with sufficient frequency to be best modeled using a fixed positive noise level. But for reasons noted in the Introduction, it is common in the literature on stochastic evolution to focus on the limiting case of vanishing noise. With this in mind, we call the state x^* *limit stochastically stable* (*LSS*) if it is the limit of stochastically stable states as the amount of noise in agents' decisions goes to zero. Formally, x^* is LSS if

$$\lim_{\varepsilon \to 0} \lim_{N \to \infty} \mu^{N,\varepsilon} = \delta_{\{x^*\}}.$$

The next corollary is an immediate consequence of our previous results.

Corollary 2: (*i*) *If the best response potential function l has a unique maximizer x*, then x* is LSS in the mutation model.*

(*ii*) *If the potential function f has a unique maximizer x*, then x* is LSS in the logit model.*

We now apply Corollary 2 to two examples. The first is the canonical one.

Example 3: Random matching in normal form coordination games and risk dominance. Suppose that agents are randomly matched to play a symmetric normal form game with payoff matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Then $F_0(x) = (1 - x)a + xb$ and $F_1(x) = (1 - x)c + xd$, so the potential function for F is $f(x) = \frac{1}{2}(a-b-c+d)x^2 - (a-c)x$. If we assume that $a > c$ and $b > d$, then F is a coordination game with Nash equilibria at states 0, 1, and $x^* = (a - c)/(a - b - c + d)$. Since the best response regions for strategies 0 and 1 are separated by state x^*, the best response potential function for F is $l(x) = |x - x^*| - x^*$.

In a symmetric 2 x 2 coordination game, strategy 1 is said to be *risk dominant* if $x^* < \frac{1}{2}$, or, equivalently, if $c + d > a + b$. Since l is convex with $l(0) = 0$ and $l(1) = 1 - 2x^*$, state 1 is LSS in the mutation model if and only if it is risk dominant. Moreover, since f is convex with $f(0) = 0$ and $f(1) = \frac{1}{2}(c+d-a-b)$, state 1 is LSS in the logit model if and only if it is risk dominant. Thus, in coordination games based on random matching, mutation LSS and logit LSS agree. §

The other example is borrowed from Blume (2003). Once again, our formulas make the analysis trivial.

Example 4: A nonlinear coordination game. Consider a nonlinear coordination game with payoff functions $F_0(x) = 1$ and $F_1(x) = ax^2$, where $a > 1$. This game has potential function $f(x) = \frac{1}{3}ax^3 - x$, Nash equilibria at states 0, 1, and $x^* = \frac{1}{\sqrt{a}}$, and best response potential function $l(x) = |x - x^*| - x^*$. State 1 is LSS in the mutation model if $x^* < \frac{1}{2}$, which is true whenever $a > 4$. On the other hand, state 1 is LSS in the logit model whenever this state maximizes f on [0, 1]; since f is convex, this is true whenever $\frac{1}{3}a - 1 = f(1) > f(0) = 0$, or, equivalently, whenever $a > 3$.

The reason for this discrepancy is easy to see. In the mutation model, only the signs of payoff differences matter; in the logit model, the magnitudes of these differences are important as well. Thus, since $|F_1 - F_0|$ tends to be larger when strategy 1 is optimal, state 1 is selected more readily under logit choice than under mutations. §

5.2 On the Robustness of the Small Noise Limit

While the last two examples show that our formulas ease the computation of LSS states, we can also apply them to more subtle questions. In this final section, we use our formulas to investigate the robustness of limit stochastic stability to nonvanishing noise levels.

Example 5: A jump in the stochastically stable state at a low noise level. Consider a game with payoff functions $F_0(x) = 0$ and $F_1(x) = 81x^2 - 108x + \frac{203}{6}$. This game has Nash equilibria at states $x^* \approx .5031$, $y^* \approx .8302$, and 1. The potential function for F is $f(x) = 27x^3 - 54x^2 + \frac{203}{6}x$, and the local maximizers of f are states x^* and 1. Since $-1.2082 \approx f(x^*) < f(1) \approx -1.1667$, state 1 is limit stochastically stable in the logit model.

On the other hand, if we fix a positive noise level ε, then the stochastically stable state is the maximizer of $m(x)$, or, equivalently, the maximizer of $\varepsilon m(x) = f(x) + \varepsilon h(x)$. If $\varepsilon = .1$, then the local maximizers of εm occur at $x_\varepsilon \approx .5031 \approx x^*$ and $z_\varepsilon \approx 1$; since $-1.1389 = \varepsilon m(x_\varepsilon) > \varepsilon m(z_\varepsilon) \approx -1.1667$, x_ε is logit stochastically stable at noise level .1. In fact, the switch from selecting state x_ε to selecting state z_ε occurs at a noise level of approximately .05993. §

This example is easily extended to a general result on the sensitivity of stochastically stable states to low noise levels, which we state here without proof.

Corollary 3: *Consider either the mutation model or the logit model of choice. For all $\varepsilon > 0$, there is a game F in which state $\frac{1}{2}$ is stochastically stable at all noise levels above ε, and a state close to* 1 *is stochastically stable at all noise levels below ε.*

Corollary 3 tells us that any strictly positive noise level may be enough to disrupt an LSS state. Thus, without specific information about payoffs in the game in question, the robustness of the LSS selection to variations in the noise level is uncertain.

In Example 5, lowering the noise level caused the stochastically stable state to jump from a point in the interior of the unit interval to a point near the boundary. The direction of this jump is no accident. Recall from Examples 1 and 2 that by lowering the noise level, we reduce the weight on entropy in the function m. Since the entropy function is concave and symmetric about $\frac{1}{2}$, lowering the noise level favors states that are closer to the boundary.

Two positive results immediately follow.

Corollary 4: *Consider either the mutation model or the logit model of choice.*

(*i*) *Suppose that state* x^* *is LSS. Then no state outside the interval from* $1 - x^*$ *to* x^* *is stochastically stable at any positive noise level.*

(*ii*) *Suppose that state* $x_\varepsilon \approx 1$ *is stochastically stable at noise level* ε*. Then states near* 1 *are stochastically stable at all smaller noise levels, and states near* 0 *are not stochastically stable at any noise level.*

Corollary 4(*i*) says that if some interior state x^* is limit stochastically stable, then no state that is more asymmetric than x^* is stochastically stable at any positive noise level. Thus, if all candidate stable states besides x^* are closer to the endpoints of the unit interval than x^*, a limiting analysis that selects x^* also rules out the selection of the other candidate stable states at all positive noise levels. Corollary 4(*ii*) tells us that when the only candidate stable states are the monomorphic states, limit stochastic stability is robust: if one monomorphic state is stochastically stable at some noise level, the other monomorphic state is not stochastically stable at any noise level. Corollary 4(*ii*) does not tell us whether the stochastically stable state will jump to an interior position at some sufficiently high noise level, but this can always be determined through a direct application of the formulas from Section 4.

References

Benaim, M., Weibull, J., 2003. Deterministic approximation of stochastic evolution in games. Econometrica 71, 873-903.

Benaim, M., Sandholm W. H., 2005. Logit evolution and large deviations in potential games. Unpublished manuscript, Université de Neuchâtel and University of Wisconsin.

Binmore, K., Samuelson, L., 1997. Muddling through: noisy equilibrium selection. J. Econ. Theory 74, 235-265.

Blume, L. E., 1993. The statistical mechanics of strategic interaction. Games Econ. Behav. 5, 387-424.

Blume, L. E., 2003. How noise matters. Games Econ. Behav. 44, 251-271.

Durrett, R., 2005. Probability: Theory and Examples, 3rd ed. Brooks/Cole, Belmont, CA.

Freidlin, M. I., Wentzell, A. D., 1998. Random Perturbations of Dynamical Systems, 2nd ed. Springer, Berlin.

Hofbauer, J., Sandholm W. H., 2005. Evolution in games with randomly disturbed

payoffs. J. Econ. Theory, forthcoming.
Kandori, M., Mailath, G. J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. Econometrica 61, 29-56.
Sandholm, W. H., 2001. Potential games with continuous player sets. J. Econ. Theory 97, 81-108.
Young, H. P., 1998. Individual Strategy and Social Structure. Princeton University Press, Princeton.