DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

STOCHASTIC EVOLUTION WITH SLOW LEARNING

Alan Beggs

Number 33

October 2000

Manor Road Building, Oxford OX1 3UQ

Stochastic Evolution with Slow Learning*

Alan Beggs
Wadham College
Oxford OX1 3PN
UK

Current Version: August 2000

Abstract

This paper studies the extent to which diffusion approximations provide a reliable guide to equilibrium selection results in finite games. It is shown that they do for a class of finite games with weak learning provided that limits are taken in a certain order. The paper also shows that making mutation rates small does not in general select a unique equilibrium but making selection strong does.

Keywords and Phrases: Equilibrium selection, Diffusion approximation, Evolutionary game theory, Risk dominance.

JEL Classifications Numbers: C73, C72

1 Introduction

This paper considers stochastic models of equilibrium selection in games. There has been substantial interest in evolutionary models of equilibrium selection since the work of Kandori, Mailath and Rob (1993) and Young (1993). This work has explored the idea that models of evolution with random shocks may help predict which equilibrium will be played in games with multiple equilibria. Most of this work has considered models of selection in discrete time. Models of equilibrium selection in continuous time (see for example Foster and Young (1990) and Fudenberg and Harris (1992)) have sometimes given rather different answers and it is therefore of some interest to understand the relationship between the two approaches.

The typical approach in the papers inspired by Kandori, Mailath and Rob (1993) and Young (1993) is to take some model of dynamic adjustment and perturb it (say by introducing mutations) so that the resulting Markov process has a unique stationary distribution. The perturbation is then allowed to tend to zero and the question is which of the equilibria of the unperturbed process does the stationary distribution converge to. This equilibrium is then taken to be the one selected by evolution. A notable result of Kandori, Mailath and Rob is that the equilibrium selected in symmetric 2×2 co-ordination games is the risk-dominant equilibrium in the sense of Harsanyi and Selten (1988), under minimal assumptions on the underlying deterministic dynamic. Ellison (1995) generalises this result. Bergin and Lipman (1996) discuss its limitations.

The work above considers games with a fixed finite population. Many commonly studied dynamics (for example the replicator dynamic) are often studied using differential equations, that is with a continuum of players. Foster and Young (1990) and Fudenberg and Harris (1992) consider perturbations of differential equations by small stochastic noise. In their context, the equilibrium selected depends rather closely on the chosen dynamic. This contrast with the finite models is perhaps disturbing. This paper aims to investigate the connection between finite population and differential equation models.

Some previous papers have considered finite population models when the number of players is large. If the source of randomness in the model is disturbances at the individual level and the strength of selection remains fixed regardless of sample size, then, by the law of large numbers, one would expect uncertainty to average out at the aggregate level. The sample paths for a large population should therefore be well approximated by the solutions to a deterministic differential equation. This is indeed the case if one considers behaviour over a fixed finite interval of time. Binmore, Samuelson and Vaughan (1994) and Sandholm (1999), for example, provide proofs in an economic context. Ethier and Kurtz (1986) Chapter 11 have a general treatment. Nevertheless this approximation may not be helpful if one is interested in the stationary distribution, which involves an indefinite time interval. In particular, the finite process may have a unique stationary distribution

for each population size but the differential equation may have multiple stationary points.

This paper pursues a different approach. In the papers considered above, the learning process at the aggregate level is almost deterministic when the population size is large. For some models the force of selection or learning may be weak or slow at the individual level, even though it is important in the aggregate. In these cases there may be significant randomness even in a large population. To capture this, this paper considers taking the limit where the strength of selection at the individual level becomes small as the population size becomes large. The force of this assumption is not that this is literally so but that selection is weak at the individual level in comparison to the size of the population. One might draw a loose analogy with the Poisson approximation to the Binomial, which is applied when the Binomial probabilities are small in comparison to the sample size.

Under these assumptions the paper shows that a suitably scaled version of the finite process converges to stochastic differential equation. Moreover, the stationary distribution of the finite model converges to that of the limiting diffusion process, and so one can use the latter to study the former. In particular, the equilibrium selection results derived from the limiting models carry over approximately to the corresponding finite models in large populations. One can therefore regard stochastic differential equation or diffusion models as providing information about finite models in a large population, albeit finite models of a rather different kind to those usually considered.

This kind of approximation has been much studied in genetics where diffusion models are commonly used to obtain information about otherwise intractable finite state models (see for example Ewens (1979) or Karlin and Taylor (1981)). The assumption that learning is slow may mean that the approximations are more useful in biological games, in which it is plausible that the force of selection is weak and random variation is important.

The paper considers two simple models of selection with finite populations in a 2×2 symmetric game. The models differ only in that in the first only one agent changes strategy at a time (according to a Poisson process) while in the second all agents change strategy simultaneously. The process is similar to that studied in Binmore and Samuelson (1997): agents become discontented with a certain probability and imitate the choice of another agent. In addition strategies may change randomly with a small probability. One can also interpret the model in terms of offspring. It is shown that as the population size becomes large the model converges to a diffusion process.

In terms of equilibrium selection in a 2×2 symmetric coordination game, one can obtain two kinds of results. If one simply lets the mutation probability tend to zero then the stationary distribution does not converge to a unique equilibrium of the game. Rather, the equilibrium distribution puts weight on both pure strategy equilibria with more weight on the equilibrium with greater 'fitness'. In other words, if the mutation rate is very small, the system spends most of its time in the

2

neighbourhood of one or other of the equilibria but occasionally switches between them. Intuitively, this is because the system even in the absence of mutation is still stochastic on account of individual-level randomness. If mutation allows the system to escape from one equilibrium, it is therefore relatively easy for the system to move against the determinstic forces of selection and reach the other equilibrium.

On the other hand, if one fixes the mutation rate and the force of selection becomes large (though not so large that the drift is large) then the force of selection dominates and the equilibrium distribution will in the limit put all its weight on the equilibrium with greater fitness. Here even if mutation allows the system to escape from one equilibrium, the deterministic force of selection is so strong that it is pushed back to the equilibrium with greater mean fitness. As in the work of Binmore and Samuelson (1997) fitness need not coincide with risk dominance.

It is, however, shown that the results are sensitive to the order in which limits are taken. If one takes the limit as population size becomes large and then considers the limit as the force of selection becomes strong, then one obtains the result above. On the other hand if the force of selection becomes strong and then one allows the population to become large, the risk-dominant equilibrium may be selected. Some care is therefore needed in the application of these approximations.

The closest paper in the literature to this one is Corradi and Sarin (1996). They also construct a diffusion approximation from a finite population model. The main differences are that they consider a different model and that they only consider the case without mutation. In this case the pure strategy equilibria are absorbing states of the diffusion model, so this does not have a unique invariant distribution, which is the focus of this paper.[1]

Foster and Young (1990) consider adding perturbations to the replicator dynamic but do not motivate these particularly well. Fudenberg and Harris (1992) derive a stochastic model by considering shocks at the aggregate level, in contrast to the approach here which considers individual-level randomness. Cabrales (1996) generalises Fudenberg and Harris' approach to games with more than two strategies. Vaughan (1996) gives a simple treatment of some selection models with diffusion. None of the papers in this paragraph deals with the relationship between finite and infinite games.

The paper is structured as follows. Section 2 gives an intuitive account of the ideas behind the diffusion approximation. Section 3 outlines the model. Section 4 gives formal results on convergence to a diffusion process. Section 5 discusses equilibrium selection. Section 6 considers the order of limits and other variations on the basic model. Section 7 briefly concludes.

---

[1]Their paper was written earlier than this one, although I only learned of it while writing up these results.

2   Informal Sketch of the Model

This Section gives an intuitive outline of the main ideas behind the diffusion approximation. Consider a population of size N playing a symmetric 2×2 game, such as that shown in Figure 1. Suppose that they are randomly matched against one another. The payoff to playing strategy 1 will depend on the fraction of the population currently playing it, x. It will be assumed that rather than always playing a best response players occasionally revise their strategies. Suppose that individual players switch from playing strategy 1 to playing strategy 2 according to a Poisson process with rate $u(x)$ and from 2 to 1 with rate $v(x)$. The details of this process will be specified in the next Section. The total number of players playing stategy 1 is Nx and so the aggregate rate of switching from 1 to 2 will be $Nxu(x)$ or $Nl(x)$, where $l(x) = xu(x)$. Similarly players switch from 2 to 1 at rate $Nr(x)$, where $r(x) = (1-x)v(x)$.

Consider an interval of time of length h. By the properties of Poisson processes, if h is small then at most one player is likely to change strategy in that time interval. With probablity about $Nl(x)h$ one player switches from 1 to 2 and so the number playing stategy 1 falls by 1, or equivalently the fraction of the population playing 1 falls by $1/N$. Similarly with probability about $Nr(x)h$ the fraction of the population playing 1 falls by $1/N$.

It follows that the expected change in the fraction of players playing strategy 1 is about

$$((r(x) - l(x))\, h \tag{1}$$

and the variance of the change is about

$$(l(x) + r(x))\, \frac{1}{N} h \tag{2}$$

where in (2) it has been assumed that h^2 is small compared to h and so the mean squared contributes little to the variance.

As N becomes large, the variance of the change becomes negligible and so one might expected the paths of the process to be well approximated by a deterministic differential equation with slope $r(x) - l(x)$. Binmore, Samuelson and Vaughan (1994) show that over a fixed finite interval of time this is indeed the case. Ethier and Kurtz (1986) Chapter 11 provides a general proof. Nevertheless this result is not entirely satisfactory if one is interested in equilibrium selection. The underlying finite process may have a unique stationary distribution but the differential equation may have multiple stationary points. If the differential equation starts in a neighbourbood of a (stable) stationary point it will remain in it but the finite process will always escape from it if enough time elapses. It follows that the differential equation cannot be used to analyse the long-run behavior, and so the stationary distribution, of the finite process.

The differential equation approximation is analogous to the Law of Large Numbers. By analogy with the Central Limit Theorem one might hope from (2) that

if one multiplies the deviations of the finite sample process from the deterministic path by $\sqrt{N}$, these would converge to a well-defined process and hence the approximation can be improved. This can be done (see Ethier and Kurtz (1986) Chapter 11 for a treatment) but it is well-known that this is not accurate enough for equilibrium selection. Binmore, Samuelson and Vaughan (1995) provide an example of this in an economic context. Kushner (1982) provides an example in another context. Sandholm (1999) suggests, however, that it may be useful for studying the stability of equilibrium points.

This paper pursues another route. Suppose that players switch very slowly between strategies, that is learning or selection is slow. More precisely suppose that

$$l(x) = r(x) + \frac{a(x)}{N} \tag{3}$$

(1) and (2) then become

$$\frac{a(x)}{N}h \tag{4}$$

and, ignoring terms of higher order than $1/N$

$$2\frac{l(x)}{N}h \tag{5}$$

Now both the mean and variance go to zero at the same rate. For large N the change over any finite interval becomes almost unnoticeable but this can be cured by looking at the changes over increasingly large time intervals. Putting $h = N\tau$ then the mean and variance converge to

$$a(x)\tau \tag{6}$$

and

$$2l(x)\tau \tag{7}$$

Replacing h by τ can be thought of as redefining time so that 1 unit of time in the new time scale corresponds to N units of time in the old.

Now intuitively, a diffusion process, X, is a continuous-time process with continuous sample paths such that

$$\mathbb{E}(X_{t+h} - X_t | I_t) = b(X_t)h + o(h) \tag{8}$$

$$\mathbb{E}\big((X_{t+h} - X_t)^2 | I_t\big) = \sigma^2(X_t)h + o(h) \tag{9}$$

I_t denotes the history of the process up to time t. In other words, $b(x)$ measures the infinitesimal drift of the process and $\sigma(x)$ its infinitesimal standard deviation.[2]

---

[2]The notation $o(h)$ means that the term is of smaller order than h, that is when divided by h it tends to zero as h tends to zero.

The diffusion processes of interest here can be generated as the solutions of a stochastic differential equation of the following form

$$dX = b(X)dt + \sigma(X)dW \qquad (10)$$

where W is standard Brownian motion. Karlin and Taylor (1981) or Durrett (1996) are excellent sources for further information on such processes.[3]

Comparison of (8) and (9) with (6) and (7) strongly suggests that under the assumptions above the finite-population process under consideration converges to a diffusion process as the population becomes large. Moreover the variance of this process is non-zero and this will guarantee that, unlike the determinstic approximation first considered, it has a unique stationary distribution and so this can be used to approximate the stationary distribution of the finite process. Precise results to this effect will be presented in Section 4.

The assumption that selection effects become weak as population size becomes large, as set out in (3), may seem peculiar but one can draw a comparison with the Poisson approximation to the Binomial distribution. To derive the approximation there one considers the effect of making the probability of success, P, go to zero as population size becomes large, holding NP constant. Nevertheless in applying it one does not imagine that P literally depends on population size. Rather it is assumed that the approximation is likely to be applicable when N is large and P is small. Similarly here taken literally, the assumption means that the net rate of switching between strategies at the individual level goes to zero as N becomes large, holding the aggregate number of switches constant. Again in applying the approximation it is not imagined that this is literally so. Rather the implication is that the approximation is likely to be useful when the population size is large and selection or learning effects are weak at the individual level.

The assumption of re-scaling time may also seem unusual. By itself it is innocuous. The need for it, however, reflects the fact the process moves very slowly as size of the population becomes large. In particular one may worry that if the system is not at its stationary distribution initially it may take so long for it to converge to it that it has little practical relevance — an issue of concern in the literature on evolutionary selection (see for example Kandori, Mailath and Rob (1993) and Binmore and Samuelson (1997)). For this reason the results may perhaps be more relevant in biological games where the evolutionary time-scale may be very large.

A mathematically equivalent interpretation, which may be more palatable in an economic context, is that rather than time being rescaled events take place more rapidly. This has the advantage of more rapid convergence times, though it implies that the number of events per unit time becomes unbounded as N becomes large.

---

[3]Note that the fact that (10) is driven by Brownian motion, which has Gaussian increments, does not imply in general that X itself has Gaussian increments. This is approximately true in a short time interval but the approximation is not in general exact. In particular, the diffusion processes considered here will remain in the unit interval and so must have bounded increments.

6

The same ideas can be applied if time is discrete. Suppose that if a fraction x of the population is playing strategy 1, the probability a player plays strategy 1 next period is $p(x)$. The number of players playing strategy 1 next period, if players choose independently, then follows a Binomial distribution with parameters N and $p(x)$. The expected change in x is therefore

$$p(x) - x \tag{11}$$

and the variance is

$$\frac{p(x)\,(1 - p(x))}{N} \tag{12}$$

If one assumes that

$$p(x) = x + \frac{a(x)}{N} \tag{13}$$

then to order $1/N$ the mean change in x becomes

$$\frac{a(x)}{N} \tag{14}$$

and the variance becomes

$$\frac{x(1 - x)}{N} \tag{15}$$

One would therefore expect that, if one rescales time as above, the process will converge to a diffusion process with drift $a(x)$ and variance $x(1 - x)$.

The ideas in this Section are made precise in Sections 3 and 4.

3   The Model

This section outlines the basic models of adjustment: one with adjustment by one player at a time, the other with simultaneous adjustment. The reason for presenting two separate models is that they are of some intrinsic interest and also serve to illustrate the fact that, as shown in Section 4, different models may have rather similar diffusion approximations.

The framework common to both is that there is a symmetric 2×2 game with strategies 1 and 2 (see Figure 1) and a finite population of N players. Players are randomly matched against each other and it is assumed for simplicity that their payoff depends simply on the proportion of players x playing strategy 1.[4] Strictly x for each player should be the proportion of other players playing 1

---

[4]Thus implicitly players are matched against each other infinitely rapidly. Kandori, Mailath and Rob (amongst others) also make this assumption

(excluding himself) but this does not affect the asymptotics[5] and to save notation this distinction is disregarded. Hence if a proportion x of players play strategy 1 the expected payoff to 1 is

$$\pi_1(x) = ax + b(1-x) \tag{16}$$

and to 2 is

$$\pi_2(x) = cx + d(1-x) \tag{17}$$

There are three cases to be distinguished. If $a > c$ and $d > b$ then the game is a co-ordination game and there are three Nash equilibria: all players play strategy 1, all play strategy 2 and a mixed-strategy equilibrium.[6] If (say) $a > d$ then the Pareto-dominant equilibrium is all playing strategy 1, but if $a - c < d - b$ it is commonly suggested that all playing 2 is the more likely outcome. In the example of Figure 2, if everyone plays strategy 1 then they receive a payoff of 9, which is better than the payoff of 7 achieved by all playing 2. On the other hand, playing 1 is quite risky: since the payoff to playing 1 against someone playing 2 is 0, it only requires a small probability of deviation by others (whether by mistake or whatever) to make switching to 2 attractive. 2 is a rather safer strategy and requires a larger probability of deviation to induce people to switch to 1. It can be argued, therefore, that in this case the risk-dominant equilibrium, all playing 2, is more plausible than the Pareto-dominant equilibrium, although Harsanyi and Selten (1988) themselves do not take this view: 'The solution that . . . results from the application of our general concept to this class [2×2 games] gives absolute priority to payoff dominance.'[7] Much of the recent work in evolutionary games has been devoted to seeing whether the prediction that the risk-dominant equilibrium will be played is borne out in explicit models of evolution.

The remaining cases are less interesting. If $a < c$ and $d < b$ then there is a unique mixed strategy equilibrium. Regarded as a 2×2 game there are also two asymmetric pure strategy equilibria but these are not relevant here as players are assumed to follow the same strategy regardless of whether they are chosen to play row or column. Finally if $a > c$ and $b > d$ then strategy 1 is the dominant-strategy equilibrium, while if $a < c$ and $b < d$ strategy 2 is the dominant-strategy equilibrium.

3.1 Overlapping Generations

This model is similar to that developed by Binmore and Samuelson (1997). In a biological context it can be traced back to Moran (1958). Time is continuous.

---

[5]At least so long as Assumption 3 in Section 3 is satisfied.

[6]In common with other papers, it is assumed that each player player plays a pure strategy. A mixed-strategy equilibrium corresponds to different players playing different strategies.

[7]Harsanyi and Selten (1988) p. 90. I am grateful to an anonymous referee for drawing my attention to this quotation.

It is assumed that players randomly become dissatisfied with their current strategies (or in a biological interpretation die). The rate of dissatisfaction depends on the payoff of the current strategy. More precisely the rate of at which players switch from strategy 1 is given by a Poisson process with rate $\lambda_1^N(x)$, from 2 with rate $\lambda_2^N(x)$. The superscripts on λ (and μ below) reflect possible dependence on population size, to be discussed below, and are not exponents. Note that from (16) and (17), x summarises the current payoffs of the two strategies. A player who becomes dissatisfied randomly chooses another player to imitate and plays the strategy the latter is playing with probability $1 - \mu^N$ and plays the opposite strategy with probability μ^N. μ^N reflects the idea of random error or 'mutation'.

The probability that the number of players playing strategy 1 increases by 1 in an interval of length h if currently j $(j < N)$ players are playing strategy 1 is therefore, letting $x = j/N$,

$$Nr(x) + o(h) = N(1 - x)\lambda_2^N(x)\, h\left[x(1 - \mu^N) + (1 - x)\mu^N\right] + o(h) \qquad (18)$$

which is of the form assumed in Section 2. This expression simply reflects the fact that the number of players playing strategy 1 can only increase if one of the players playing strategy 2 becomes discontented, which happens at rate $N(1 - x)\lambda_2^N(x)$ on account of the Poisson assumption (since each individual player becomes discontented according to a Poisson process with rate $\lambda_2^N(x)$, the total rate of discontentment is also a Poisson process with rate $N - j$ times that). The number of players playing strategy 1 will increase if either this player samples a player playing strategy 1 and obeys his recommendation (probability $x(1 - \mu^N)$) or he samples a player playing strategy 2 and disobeys his recommendation (probability $(1 - x)\mu^N$). Strictly this assumes that a player can sample himself but ruling this out does not affect the asymptotics and simply makes the notation messier. The final term of smaller order than h simply reflects the fact that under the Poisson process more than one change in state is very unlikely if the time interval is short.

Similarly the probability that the number of players playing strategy 1 decreases if j players $(j > 0)$ in an interval of length h is given by

$$Nl(x) + o(h) = Nx\,\lambda_1^N(x)\, h\left[(1 - x)(1 - \mu^N) + x\mu^N\right] + o(h) \qquad (19)$$

The probability of the number of players playing strategy 1 changing by more than 1 is of smaller order than h , since events follow a Poisson process, and so is negligible if h is small.

Under the above assumptions, the model has a unique stationary distribution[8] and in the spirit of Kandori, Mailath and Rob (1993) it is of interest to determine the behaviour of the equilibrium when μ^N is small. Note that when there is no mutation both $x = 0$ and $x = 1$ are absorbing states.

This structure is similar to that employed by Binmore and Samuelson (1997)[9]. It reflects a simple model of satisficing behaviour: players become discontented

---

[8]See for example Asmussen (1987) Theorems 4.3 and 4.6 in Chapter II (p. 40–41).

[9]Unlike theirs, this paper works directly in continuous time

and experiment by imitation. Note that players do not compare their own payoff to the payoff enjoyed by the person they are imitating and in this sense are rather naïve. Schlag (1996) considers more sophisticated imitation rules. One could make the imitation probability reflect payoff differences by making μ^N a function of x but this is not pursued here. It seems more interesting to see whether fairly simple behaviour can yield the equilibrium behaviour predicted by sophisticated models of play.

The biological interpretation is fairly straightforward. Every so often a member of the population dies and one of members is selected randomly to breed to replace him. The offspring has the same type as his parent but mutates with probability μ^N. The assumption of fixed population is perhaps artificial but may reflect environmental constraints on total population size.

This model could be analysed directly, by using standard techniques for birth-death chains (see for example Binmore and Samuelson (1997) or Amir and Berninghaus (1996)). It is, however, of interest to compare it to the next model which cannot be solved directly but has very similar limiting properties.

3.2   Non-overlapping Generations

The second model considered is perhaps easier to motivate in biological terms rather than learning. It is essentially the well-known Fisher-Wright model in genetics (see for example Ewens (1979) for a detailed discussion). In contrast to the previous model, all agents change their strategies at once. This model is fairly intractable and the diffusion approximation allows a vast simplification.

In this model time is discrete. Again there is a finite number of agents N. At the end of each period, each player produces a large number of offspring. The number of offspring depends on the current payoff to the strategy being played. The relative number produced by an agent of type 1 compared to one of type 2 is denoted $s^N(x)$. If strategy 1 is yielding a higher payoff than strategy 2 then $s^N(x)$ exceeds 1. The offspring are of the same type as their parent but there is a probability μ^N that they mutate into the opposite type. It follows that the probability that a randomly selected offspring is of type 1 is

$$p^N(x) = \frac{(1 - \mu^N)s^N(x)x + \mu^N(1 - x)}{s^N(x)x + (1 - x)} \tag{20}$$

Random variation occurs because there is only room for a finite number N of these offspring in the next generation. The next generation is selected by random sampling from the offspring, so the probability that there are j players of type 1 in the next generation given that the current fraction is x is

$$p_j^N = \binom{N}{j} p^N(x)^j (1 - p^N(x))^{N-j} \tag{21}$$

As with the previous model this has a unique stationary distribution[10] but in this case it impossible to characterise it directly.

One could give this model a learning interpretation if one assumed that players learn by imitation and a player is relatively more likely to be chosen for consulation (by a factor $s^N(x)$) if the strategy he is employing the one currently more successful. The information mechanism underlying this is perhaps less plausible than a simple biological model.

4 Diffusion Approximations

In this section, the ideas of Section 2 are applied to the models of Section 3. In the case of overlapping generations, it is assumed that selection and mutation are weak in the following sense:

Assumption 1 $\lambda_1^N(x) = 1 + \frac{\lambda_1(x)}{N} + o\left(\frac{1}{N}\right)$ and $\lambda_2^N(x) = 1 + \frac{\lambda_2(x)}{N} + o\left(\frac{1}{N}\right)$ and $\mu^N = \frac{\mu}{N} + o\left(\frac{1}{N}\right)$.

and for the non-overlapping generations model

Assumption 2 $s^N(x) = 1 + \frac{s(x)}{N} + o\left(\frac{1}{N}\right)$ and $\mu^N = \frac{\mu}{N} + o\left(\frac{1}{N}\right)$.

The force of these assumptions is that the selective advantages of the two strategies are small at the individual level, so as a result random variation from the environment is not completely dominated by selective forces. The terms of order smaller than $1/N$ (if any) in the selection and mutation terms do not affect the asymptotics. In finite samples, of course, their form may matter but evaluating the adequacy of the asymptotic approximations in finite samples is left to future research.

Let Z^N be the number of players playing strategy 1. In the case of overlapping generations, one considers $X^N(t) = Z^N(Nt)/N$. In the case of non-overlapping generations, $X^N(t) = Z^N([Nt])/N$, where $[\ldots]$ denotes the integer part. In other words, X^N is the mean number of players playing strategy 1 and time is measured so that one unit of time in the new process corresponds to N generations in the old process. (In the case of overlapping generations, the rate of events per individual is approximately 1 from Assumption 1, so it takes roughly 1 unit of time in the old process for a generation to revise their strategies). Note that in the case of non-overlapping generations, the process is embedded in continuous time by defining X^N to be constant between shifts in strategy.

If one uses Assumption 1 for the overlapping generations process and works to order $1/N$ then it is easy to check, using (18) and (19), that (6) becomes

$$x(1-x)(\lambda_2(x) - \lambda_1(x))\tau + \mu(1-x)\tau - \mu x\tau \qquad (22)$$

---

[10]See for example Corollaries 3.5 and 3.6 in Chapter I (p.15–16) in Asmussen (1987).

and (7) is given by

$$2x(1-x)\tau \tag{23}$$

It is therefore plausible that the process converges to a diffusion process with mean and variance given by the above equations (divided by τ).

Similarly, in the case of non-overlapping generations, it follows from Assumption 2 and (20) that (14) becomes

$$(s(x)x(1-x) + \mu(1-x) - \mu x)\frac{1}{N} \tag{24}$$

and (15) is as before. Again it is therefore to be expected that the process converges to a diffusion process with mean given by (24) divided by $1/N$ and variance $x(1-x)$.

Formally, one can regard the models as defining a stochastic process with sample paths in $D_I[0, \infty)$, the space of functions from $[0, \infty)$ to $I = [0, 1]$ which are right-continuous and have left limits. Each model defines a probability measure over all such possible paths, denote this by Δ^N. The content of the result below is that the probability measures converge weakly to that yielded by equation (10), denote this by Δ. Weak convergence is denoted by $\Longrightarrow$.[11]

For the following result one technical assumption is required. It guarantees, in conjunction with the form of the variance term, that the limiting diffusion has a unique solution.[12] Recall that a function g satisfies a Lipschitz condition if there exists a constant K such that $|g(x) - g(y)| \leq K|x - y|$ for all x,y. If g is continuously differentiable, then this will be true provided the derivative is uniformly bounded (as it must be since the unit interval is compact).

Assumption 3 $\lambda_1(x)$ and $\lambda_2(x)$, and hence $\phi(x) = \lambda_2(x) - \lambda_1(x)$, and $s(x)$ satisfy a Lipschitz condition.

One then has,

Theorem 1 *Under Assumptions 1 to 3, provided $X^N(0) \Longrightarrow X(0)$, then $\Delta^N \Longrightarrow \Delta$, where Δ is the probability measure on $D_I[0, \infty)$ generated by (10) with initial condition $X(0)$, where $b(x) = (\lambda_2(x) - \lambda_1(x)) x(1-x) + \mu(1-x) - \mu x$ and $\sigma(x) = \sqrt{2x(1-x)}$ in the case of overlapping generations, and $b(x) = s(x)x(1-x) + \mu(1-x) - \mu x$ and $\sigma(x) = \sqrt{x(1-x)}$ in the case of non-overlapping generations.*

Proof: *See Appendix.*

In other words, provided the initial condition in the discrete model converges to that in the continuous model, the probability distribution of the sample paths also does.

---

[11]To be precise, one needs to specify a topology on the space of sample paths for weak convergence to be a well-defined notion. The appropriate topology is the Skorokhod topology — see Ethier and Kurtz (1986) Chapter 3 for a definition and a general treatment of weak convergence.

[12]See for example Durrett (1996) Chapter 5, Theorem 3.3.

The diffusion process constructed above has the property that it always remains in the unit interval, as it must if it is to approximate the finite process. Intuitively, this follows from the fact that the drift points inward at the boundaries and the variance vanishes sufficiently rapidly as the process approaches the boundaries (0 or 1). In fact for large μ, the process cannot hit the boundaries but if it starts there it can leave them for the interior of the interval. If μ is small, however, then the boundaries can be reached.[13] Nevertheless, the specification above ensures that the process cannot leave the unit interval[14] and in fact spends measure zero time at the boundaries.[15]

For current purposes, the most interesting fact is not the convergence of the sample paths but the fact that the stationary distributions of the discrete model converge to that of the diffusion process.

Under Assumption 3, (10) has a unique stationary distribution on $[0, 1]$ with density function (with respect to Lebesgue measure) given by

$$f(x) = C \frac{1}{\sigma^2(x)} \exp \left(\int_\xi^x \frac{2b(s)}{\sigma^2(s)} \, ds \right) \tag{25}$$

where $b(x)$ and $\sigma(x)$ have the form given in Theorem 1,[16] C is a constant chosen so that the density integrates to 1 and ξ is an arbitrary point (its value only affects that of C). One has

Theorem 2 *The stationary distribution of the mean number of players playing strategy 1 converges weakly to that given by (25), where $b(x)$ and $\sigma(x)$ have the form given in Theorem 1.*

Proof: *See Appendix.*

This result will be used in the next Section to study the invariant distributions of the process with small noise. Theorem 1 holds even if μ is zero, but in this case the boundaries are absorbing and so there is no stationary distribution of the limiting process. Hence Theorem 2 fails to apply. One can however show that the absorption probabilities of the boundary points converge to those for the limiting process (see for example Ethier (1979)) but this result will not be needed here.

It is straightforward to extend Theorem 1 to general $n \times n$ games and show that the limiting distribution is the limit of the finite stationary distributions. Except

---

[13]To be more precise the boundaries can be reached when $\mu < 1$ with overlapping generations and $\mu < 1/2$ with non-overlapping generations. See Durrett (1996) Chapter 6.5, or Karlin and Taylor (1981) Chapter 15, Sections 5 and 6 for details.

[14]See for example, Ethier (1976)

[15]See for example Shiga (1981) Lemma 3.2 for a proof. In other words, the boundary points are instantaneously reflecting (see for example Revuz and Yor (1991) p.285). Norman (1980) has a detailed discussion of the implied boundary conditions for the limiting model with these kind of approximations.

[16]See for example Karlin and Taylor (1981) Chapter 15, Sections 5 and 6 or Rogers and Williams (1987) Chapter V.52 and V.53.

in a few special cases, however, explicit solutions for the stationary distribution of a multi-dimensional diffusion process are not known, so the results are less useful. Attention is therefore restricted to 2×2 games.

It is perhaps briefly worth comparing the limiting model to those considered in the literature previously. Foster and Young (1990) consider a model with $b(x) = (\alpha + \beta x)x(1 - x)$ and $\sigma(x) = \epsilon$, in the spirit of Freidlin and Wentzell (1984), where α and β are constants. The constant variance term implies that the process may leave the unit interval, so they impose the condition that the process be instantaneously reflected at the boundary to prevent this but this assumption does not seem well-motivated. Fudenberg and Harris (1992) assume aggregate shocks to payoffs and deterministic mutation and obtain a model of the form $b(x) = (\alpha + \beta x)x(1-x) + \mu x - \mu(1-x)$ and $\sigma(x) = x(1-x)$. This is similar to the model obtained above with linear $s(x)$ except that the variance term in their model is the square of the term here. This has the consequence that the boundaries are unattainable in their model so long as μ is strictly positive. Here the boundaries are attainable for small μ but as noted above this does not cause difficulties. Corradi and Sarin (1996) consider a model of the form $b(x) = (\alpha + \beta x)x(1 - x)$ and $\sigma(x) = \sqrt{x(1-x)}$, so that it has the same form as the current one with linear $s(x)$ and no mutation, so the boundaries are absorbing.

5   Equilibrium Selection

This section applies the results obtained in the previous section to equilibrium selection. The natural procedure, by analogy with previous work, is to let the mutation rate tend to zero. It is shown in Section 5.1 that here this does not lead to unique equilbrium selection. Instead when mutation rates become small the process spends most of its time in the neighbourhood one or other of the equilibria but occasionally switches between them. Intuitively this is because even in the absence of mutation there is still considerable randomness in the system due to randomness in the process of strategy choice. Once mutation has allowed the system to escape from one of the boundaries, this randomness in the selection process makes it relatively easy to move counter to the forces of the deterministic selection dynamic and reach the other boundary.

Section 5.2 considers the effects of eliminating this randomness by making selection strong. In this case unique selection is achieved. Now, even if mutation allows the system to escape from one of the boundaries, the strength of the deterministic component of selection will force it back to the equilibrium with greater mean fitness.

It follows from Theorem 2 that the stationary distribution of the limiting process in the case of overlapping generations has a density of the form

$$f(x) = Cx^{\mu-1}(1-x)^{\mu-1} \exp \left(\int_{\xi}^{x} \phi(y) \, dy \right) \tag{26}$$

14

where C is a normalising constant (and is a function of the parameters). In the case of non-overlapping generations, the above becomes

$$f(x) = C x^{2\mu-1}(1-x)^{2\mu-1} \exp\left(\int_\xi^x 2s(y)\,dy\right) \qquad (27)$$

The extra factor 2 in (27) does not affect the asymptotics and for definiteness, attention will focus on (26). To save notation, let $H(x) = \exp\left(\int_\xi^x \phi(y)\,dy\right)$ and $K(x) = \exp\left(\int_\xi^x 2s(y)\,dy\right)$.

5.1   Weak Mutation

It is easy to see from (26) that as μ tends to zero most of the density of the process, and so the probability mass concentrates on the end-points, 0 and 1. Intuitively, as μ becomes smaller and smaller it becomes harder and harder to escape from the boundaries and so the process spends most of its time there. The relative magnitudes of the densities near 0 and 1 are determined by $H(0)$ and $H(1)$ and so the following is plausible:

Theorem 3 *As μ tends to zero the stationary distribution given by (26) (respectively (27)) converges weakly to a discrete distribution that puts positive probability mass only on 0 and 1, and that in the ratio $H(0)$ to $H(1)$ (respectively $K(0)$ to $K(1)$).*

For a proof see the Appendix. Referring to the definition of $H(x)$ one sees

Corollary 1 *The limiting distribution of (26) (respectively (27)) puts greater weight on the point 1 if and only if $\int_0^1 \phi(x)\,dx > 0$ (respectively $\int_0^1 s(x)\,dx > 0$).*

$\phi(x)$, and $s(x)$, represent the difference in reproductive fitness or difference in rates at which players become dissatisfied when a proportion x of players is playing strategy 1. Hence the result can be interpreted as saying that more weight will be put on strategy 1 if it has higher mean fitness.

If one denotes the stationary distribution in a finite model by ν_μ^N, then the implication of the results above and in the previous section is therefore

Corollary 2 $\lim_{\mu\to 0} \lim_{N\to\infty} \nu_\mu^N$ *(in the sense of weak convergence) is given by the discrete distribution in Theorem 3.*

In other words provided one chooses μ sufficiently small then for sufficiently large N (which may depend on μ), the ratio of probability mass on 0 and 1 will be approximately $H(0)/H(1)$. The result therefore gives information about the corresponding discrete models provided limits are taken in the order indicated.

This result contrasts with the usual results in the literature (for example Kandori, Mailath and Rob (1993)), which predict that all the players will end up

playing the same strategy. The reason is not too hard to see. Although the mutation rates tend to zero, there is still aggregate randomness in the system in the limit from randomness in individual choice (or selection). As a result, once mutation allows escape from (say) $x = 0$, aggregate randomness means it possible to move relatively quickly to the neighbourhood of $x = 1$. The system then spends a long time there until mutation allows it to escape again. As a result the system spends almost all its time at the end-points (in the ratio $H(0)$ to $H(1)$, see Asmussen (1987) Chapter V) but occasionally switches between them.[17]

So far no restrictions have been put on the nature of $\phi(x)$ and $s(x)$. The following seems reasonable.

Assumption 4 $\mathrm{sgn}(\phi(x)) = \mathrm{sgn}(\pi_1(x) - \pi_2(x))$ *and similarly for* $s(x)$.

In other words, players playing strategy 1 reproduce to a greater extent or become dissatisfied more slowly, if strategy 1 currently yields a higher payoff. In the special case when $\phi(x)$ is simply $\alpha(\pi_1(x) - \pi_2(x))$, where α is a constant, then strategy 1 has higher mean fitness if and only if it risk-dominates strategy 2. This can be seen in Figure 3. The difference in payoffs is shown by the dotted line. The total area below between the line and the x-axis, which is proportional to the mean difference in payoffs is positive only if $1 - x^*$ exceeds x^*. The point x^* is the mixed-strategy equilibrium. The case shown corresponds to there being a co-ordination game. Here if there are deterministic best-reply dynamics (everyone switches to the best-reply when switching) the population converges to everyone playing strategy 1 if the initial proportion players playing 1 exceeds x^*. The arrows on the x-axis show the direction of motion under best-reply dynamics. This region is therefore referred to as the basin of attraction of strategy 1. Strategy 1 risk-dominates strategy 2 precisely if it has a larger basin of attraction.

In fact, strategy 1 has greater mean fitness than strategy 2 provided $\phi(x) = f(\pi_1(x) - \pi_2(x))$ and f is an odd function (that is $f(-y) = -f(y)$), in other words fitness solely depends on the payoff difference between the two strategies. This can be seen from the diagram: an equal distance from x^* gives the same (absolute) difference in payoffs between the two strategies, but the risk-dominant strategy has a longer basin to integrate over. This yields

Corollary 3 *In the co-ordination case, the limiting distribution puts more weight on strategy 1 when $\phi(x)$ (or $s(x)$) is an odd function of the payoff difference between 1 and 2 if and only if 1 risk-dominates strategy 2.*

On the other hand, this clearly need not be the case if fitness is not an odd function of the payoff difference between the strategies. In the case of the model of players becoming dissatisfied with their current strategies, the spirit is that rather

---

[17]The model of Binmore and Samuelson (1997) which is quite similar gives unique selection. The difference is that here the selective difference between the strategies tends to zero as the population size grows so that the system is still stochastic in the limit.

than comparing strategies they only look at the payoff of their current strategy before deciding whether to switch. It would therefore be more natural to suppose that $\phi(x) = f(\pi_1(x)) - f(\pi_2(x))$, which does not have the form above unless f is linear. Similarly in the case of selective differences in reproduction one might argue that the selective difference should be the ratio of functions of the current payoffs and this would not in general make $s(x)$ a function of the payoff difference. The support for the prediction that the risk-dominant strategy will be selected is therefore rather limited.

As a description of co-ordination games the model is not unattractive: the system spends most of its time in the neighbourhood of the pure-strategy equilibria but occasionally switches between them. The observant reader will have noticed, however, that the system will still spend a positive amount of time in the neighbourhood of $x = 0$ even if strategy 1 strictly dominates strategy 2. The reason is not hard to understand: since randomness does not vanish in the limit it is still possible for the inferior strategy to dominate by chance. This is, perhaps, unsatisfactory and the case of strong selection, considered in the next sub-section, will not have this feature.

5.2 Strong Selection

The previous sub-section considered what happened as mutation becomes weak but even in the limit there is considerable randomness in selection. This sub-section considers what happens as the deterministic component of selection becomes strong and the randomness is eliminated. In this case unique equilibrium selection is achieved. For even if the system escapes from the boundaries, the deterministic component of the selection dynamic will tend to push the system back to the equilibrium with greater mean fitness.

To be more precise, consider replacing $\phi(x)$ by $k\phi(x)$ where k measures the strength of selection. The equation for the evolution of the limiting model becomes

$$dx = k\phi(x)x(1-x) + \mu(1-x) - \mu x + \sqrt{2x(1-x)}dW \qquad (28)$$

As k becomes large, both the effects of mutation and random variation become small relative to the deterministic component and so one might expect unique selection.[18] Note that from (25), this has the same effect on the stationary distribution as holding ϕ fixed but replacing μ by μ/k and $2x(1-x)$ by $2x(1-x)/k$, that is, allowing the mutation terms and variance to tend to zero. Although formally equivalent, this is perhaps harder to interpret in terms of the underlying finite models.

Assumptions 1 and 2 become

Assumption 5 $\lambda_1^N(x) = \max\{1 + k\frac{\lambda_1(x)}{N} + o\left(\frac{1}{N}\right), 0\}$ *and* $\lambda_2^N(x) = \max\{1 + k\frac{\lambda_2(x)}{N}, 0\} + o\left(\frac{1}{N}\right)$ *and* $\mu^N = \frac{\mu}{N} + o\left(\frac{1}{N}\right)$.

---

[18] For an example of this kind of exercise in a biological context, see Gillespie (1983).

and

Assumption 6 $s^N(x) = \max\{1 + k\frac{s(x)}{N} + o\left(\frac{1}{N}\right), 0\}$ and $\mu^N = \frac{\mu}{N} + o\left(\frac{1}{N}\right)$.

This simply ensures that all the relevant parameters are non-negative (as was implicitly assumed before). Since k is held fixed as N tends to infinity, for large N these restrictions are irrelevant and so the previous limiting results apply. Using the notation of the previous section, it follows from (26) that the limiting density for the case of overlapping generations is

$$f(x) = Cx^{\mu-1}(1-x)^{\mu-1}\exp\left(k\int_\xi^x \phi(y)\,dy\right) \tag{29}$$

As k becomes large, one would expect almost all weight to go to points for which $\int_\xi^x \phi(y)$ is largest. Under Assumption 4, these points can only be Nash equilibrium points. Hence the following is plausible (proof in Appendix)

Theorem 4 *As k tends to infinity, the stationary distribution of the limiting model converges weakly*
(a) in the co-ordination case to a point mass at $x = 1$ if $\int_0^1 \phi(x)\,dx > 0$ ($\int_0^1 s(x)\,dx > 0$ with non-overlapping generations) or to a point mass at $x = 0$ if $\int_0^1 \phi(x)\,dx < 0$ (respectively $\int_0^1 s(x)\,dx < 0$).
(b) to a point mass at $x = 1$ if strategy 1 is dominant or to a point mass at $x = 0$ if strategy 2 is dominant
(c) in the mixed strategy case to a point mass at x^*, the mixed-strategy equilibrium point.

As before, this provides information on the finite models. If one denotes the stationary distribution of a finite model by ν_k^N, then the previous result and the results of the last section imply

Corollary 4 $\lim_{k\to\infty}\lim_{N\to\infty}\nu_k^N$ *(in the sense of weak convergence) is given by the appropriate discrete distribution of Theorem 4.*

So in the co-ordination case if k is large enough, then almost all the probability mass will be put on the equilibrium with greater fitness if N is large enough.

As in the last sub-section, the equilibrium with greatest mean fitness need not be the risk-dominant equilibrium.

6 Discussion

6.1 Order of Limits

The last section showed that when selection becomes strong, the finite model puts almost all weight on the equilibrium with greater mean fitness when the population

size is large, if one takes limits in the order: first population size goes to infinity and then selection becomes strong. One might ask what happens if one reverses the order of limits. For the moment, attention is restricted to the model with non-overlapping generations. If one fixes N and lets k become large, the finite model is still ergodic for all finite k and for $k = \infty$ as μ^N is still positive. One has

Lemma 1 $\lim_{k\to\infty} \nu_k^N = \nu_\infty^N$ *(in the sense of weak convergence).*

This is proven in the Appendix but is easily seen since any limit of ν_k^N is a stationary distribution for the model when $k = \infty$ and since the model is ergodic, this must be ν_∞^N. In any case, it follows that it is enough to consider the model when $k = \infty$ in evaluating the reverse order limit.

From Assumption 5 and (20), it follows that when $k = \infty$, the probability that a member of the next generation plays strategy 1 if the current proportion playing 1 is x is given by[19]

$$p^N(x) = \begin{cases} 1 - \mu^N & \text{if 1 is a best reply} \\ \mu^N & \text{otherwise} \end{cases} \tag{30}$$

In other words, the process is now best-reply dynamics perturbed by an error probability μ^N. As N becomes large, μ^N becomes small and so the results of Kandori, Mailath and Rob (1993) suggest that in the co-ordination case the risk-dominant equilibrium will be selected. This not immediate as they consider the case of small error probability and fixed population size but the following is shown in the Appendix:[20]

Theorem 5 *In the co-ordination case, ν_∞^N tends weakly to a discrete distribution which puts all probability mass on the risk-dominant equilibrium.*

Putting this together with Lemma 1 yields

Corollary 5 $\lim_{N\to\infty} \lim_{k\to\infty} \nu_k^N$ *is a discrete distribution placing all probability mass on the risk-dominant equilibrium.*

This contrasts with Corollary 4 which shows that if limits are taken in the reverse order, the risk-dominant equilibrium need not be selected. Intuitively, making k large first makes all dynamics like perturbed best-reply dynamics. The work of Kandori, Mailath and Rob (1993) therefore suggests that if k becomes large faster than N, the risk-dominant equilibrium will be selected. When N

---

[19]For convenience assume that the mixed strategy equilibrium point is irrational, so it is not an element of the state space for finite N. This is not essential.

[20]Kandori, Mailath and Rob (1993) conjecture (p. 47) that if the error probability is held fixed and N tends to infinity, probability mass will concentrate on the risk-dominant equilibrium. The argument in the Appendix can easily be adapted to prove this as well.

becomes larger faster than k, however, the detailed interaction of the dynamics with the random terms is important, so the form of the dynamic matters.

The result for the case of overlapping generations is slightly different. To simplify the analysis, it is convenient to replace Assumption 5 with the particular assumption

Assumption 7 $\lambda_1^N(x) = \exp\left(k\frac{\lambda_1(x)}{N}\right)$ and $\lambda_2^N(x) = \exp\left(k\frac{\lambda_2(x)}{N}\right)$ and $\mu^N = \frac{\mu}{N}$.

Note that since $\exp\left(k\frac{\lambda_1(x)}{N}\right) = 1 + k\frac{\lambda_1(x)}{N} + o\left(\frac{1}{N}\right)$, this satisfies Assumption 5. This specific form is tractable and has the feature, under Assumption 4, that as k becomes large the relative frequency of death (or dissatisfaction) of the less fit strategy tends to infinity.

In this case, the limiting model as k tends to infinity is not ergodic.[21] Here as k becomes large, the less fit strategy dies off much faster than the fit one and so it becomes in the limit impossible for less fit strategy to grow, since the only way it can do so is by mutations after deaths of the fitter strategy, which are increasingly rare. This contrasts with the case of non-overlapping generations, where all strategies die at the same rate and mutation in the reproduction process can still lead to entry of less fit strategies even though they reproduce at a much lesser rate than the fitter ones. In any event, one does not obtain perturbed best-reply dynamics in the limit as k tends to infinity here and one cannot appeal to Lemma 1. One can however show (proof in Appendix):

Theorem 6 *With overlapping generations,* $\lim_{N\to\infty}\lim_{k\to\infty}\nu_k^N = \lim_{k\to\infty}\lim_{N\to\infty}\nu_k^N$ *in the co-ordination case.*

In other words, the limiting distribution puts all weight on the equilibrium with greater mean fitness in the co-ordination case. This contrasts with the result for non-overlapping generations and is perhaps more reassuring. The contrast can perhaps be understood in terms of the remarks before the Theorem.[22] [23]

Nevertheless, the results with non-overlapping generations raise the question of the order of limits. Which limit is more appropriate is open to debate. Loosely, one might think that taking N to become large faster than k is appropriate when one is interested in the effect of varying the strength of selection in a large population, whereas the reverse order is more appropriate when one is interested in the effect of population size when selection is strong. The results of the previous section are perhaps more appealing but some care is clearly required in applications.

---

[21]Note that only relative rates of death matter for the stationary distribution, so one can always normalise the model so the rates are finite.

[22]Note that this result is not a question of non-overlapping versus overlapping generations: one can construct overlapping generations models which have the features of the non-overlapping process and also yield the risk-dominant equilibrium

[23]Note that result in Theorem 6 depends on the exponential form. In general one will not obtain equality of the double limits even with overlapping generations.

6.2 Asymmetric Mutation

It is straightforward to allow mutation rates to be asymmetric. Suppose that the probability that some-one mutates from playing strategy 1 is $\mu_1^N = \mu_1/N$ and from strategy 2 is $\mu_2^N = \mu_2/N$.[24] Then one obtains in the limiting model

$$dx = \phi(x)x(1-x) + \mu_2(1-x) - \mu_1 x + \sqrt{2x(1-x)}dW \qquad (31)$$

In the case of strong selection, considered in Section 4.2 the results are unaffected. In the case of weak mutation, the answer depends on how fast μ_1 and μ_2 tend to zero. It is straightforward to show that if $\mu_1/\mu_2 \to 0$, that is the probability of mutation or escape from $x = 0$ goes to zero much faster than from $x = 1$, then $x = 0$ will be selected in the limit. Since mutation is conceived of as arising from random error, this does not seem a very appealing resolution of the selection issue in the absence of convincing story as to why mutation rates should behave like this.

7   Conclusion

This paper has investigated the extent to which diffusion approximations give a guide to equilibrium selection results in finite models. It showed that for a class of models where selection is weak at the individual level and randomness is important, they do so provided that limits are taken in a certain order. They need not always do so, however, and this suggests some caution as to their use.

The paper also showed that the usual procedure of letting mutation rates go to zero need not result in unique equilibrium selection. In these models, there is still considerable randomness even in the absence of mutation and once one has escaped from one equilibrium, it is relatively easy to move between equilibria. The system therefore spends more time in the neighbourhood of the equilibrium with greater mean fitness but moves in between the equilibria periodically. By contrast, if selection becomes strong, the deterministic component dominates and so it becomes very difficult to escape from the neighbourhood of the equilbrium with greater mean fitness. Unique equilibrium selection therefore results.

The results show that the prediction of the diffusion approximations that the risk-dominant equilibrium need not always be selected is borne out in finite models with a large population and weak selection. In contexts where these are appropriate, therefore, the risk-dominant equilibrium need not always be observed.

---

[24]In the case of random sampling, this assumes that the probability of mutation depends on the strategy sampled not the original strategy.

Appendix

Proof of Theorem 1

In the case of non-overlapping generations, Theorem 1.1 of Chapter 10 (p. 415) of Ethier and Kurtz (1986) proves the result for the case when $s(x)$ is a constant. The argument in the general case is exactly the same. The assumption that $s(x)$ satisfies a Lipschitz condition means that the conditions of Theorem 2.1 in Chapter 8 (p. 371) are satisfied.

In the case of overlapping generations, the argument is as above except one appeals to Corollary 8.7 in Chapter 4 instead of Corollary 8.9 and Theorem 6.1 in Chapter 1 instead of Theorem 6.5.

Proof of Theorem 2

In the case of non-overlapping generations, the result follows directly from Theorem 1 of Norman (1975). In the case of overlapping generations, the argument indicated there on p. 573 yields the result.

Proof of Theorem 3

Denote the stationary distribution for μ positive by ν^μ. The set of probability measures on $[0, 1]$ is compact in the weak topology (see Billingsley (1968) p. 37) and so it is enough to prove that any convergent subsequence of ν^μ has the indicated limit. Further by Billingsley (1968) p. 2 it is enough to show that for any $0 < a < 1$, $\nu^\mu[0, a]/\nu^\mu(a, 1]$ converges to $H(0)/H(1)$ and $\nu^\mu(a, b]/\nu^\mu[0, a]$ converges to 0 if $0 < b < 1$. Now for any δ

$$\frac{\nu^\mu[0, a]}{\nu^\mu(a, 1]} = \frac{\nu^\mu[0, \delta] + \nu^\mu(\delta, a]}{\nu^\mu(a, 1 - \delta] + \nu^\mu(1 - \delta, 1]} \tag{32}$$

Since H is continuous given $\epsilon > 0$ one can find δ (with $\delta < a < 1 - \delta$) such that $H(0) - \epsilon < H(x) < H(0) + \epsilon$ for $x \leq \delta$ and $H(1) - \epsilon < H(x) < H(1) + \epsilon$ for $x \geq 1 - \delta$. Now it follows from (26) that if $\mu < 1$

$$\nu^\mu[0, \delta] \geq C(H(0) - \epsilon)\frac{\delta^\mu}{\mu} \tag{33}$$

and

$$\nu^\mu[0, \delta] \leq C(H(0) + \epsilon)(1 - \delta)^{\mu-1}\frac{\delta^\mu}{\mu} \tag{34}$$

Similarly if $\mu < 1$

$$\nu^\mu(1 - \delta, 1] \geq C(H(1) - \epsilon)\frac{\delta^\mu}{\mu} \tag{35}$$

and

$$\nu^\mu(1 - \delta, 1] \leq C(H(1) + \epsilon)(1 - \delta)^{\mu-1}\frac{\delta^\mu}{\mu} \tag{36}$$

Also if $\mu < 1$

$$\nu^\mu(\delta, a] \leq C\bar{H}\delta^{\mu-1}(1-a)^{\mu-1}a \tag{37}$$

where $\bar{H}$ is an upper bound for $H(x)$ on $[0,1]$. Similarly

$$\nu^\mu(a, 1-\delta] \leq C\bar{H}a^{\mu-1}(1-\delta)^{\mu-1}(1-a) \tag{38}$$

It follows that

$$\liminf_{\mu \to 0} \frac{\nu^\mu[0, a]}{\nu^\mu(a, 1)} \geq \frac{H(0) - \epsilon}{H(1) + \epsilon}(1 - \delta) \tag{39}$$

and

$$\limsup_{\mu \to 0} \frac{\nu^\mu[0, a]}{\nu^\mu(a, 1]} \leq \frac{H(0) + \epsilon}{H(1) + \epsilon}\frac{1}{(1 - \delta)} \tag{40}$$

Since ϵ and δ can be made arbitrarily small, it follows that $\nu^\mu[0, a]/\nu^\mu(a, 1]$ converges to $H(0)/H(1)$. That $\nu^\mu(a, b]/\nu^\mu[0, a]$ converges to 0 follows easily from (33) and (37) and (38).

Proof of Theorem 4

It follows from Assumption 4 that in $R(x) = \int_\xi^x \phi(x)\, dx$ achieves its maximum at a Nash equilibrium. The proof is given for the co-ordination case. The proofs for the other cases are similar.

Suppose that $\int_0^1 \phi(x)\, dx < 0$, so that $R(x)$ achieves a unique maximum at $x = 0$. Denote the stationary distribution for finite k by ν^k. Since the set of probability measures on $[0, 1]$ is compact (see Billingsley (1968) p. 37) it suffices to show that for each $a > 0$, $\nu^k[0, a]/\nu^k(a, 1]$ tends to infinity as k tends to infinity (see Billingsley (1968) p. 2). Since R is continuous and attains a unique maximum at $x = 0$ one can find $0 < \delta < a$ such that $s' \equiv \inf_{x \in [0, \delta)} R(x) > \sup_{y \in (a, 1]} R(y) \equiv \bar{s}$.

Now

$$
\begin{aligned}
\frac{v^k[0, a]}{\nu[a, 1]} &= \frac{\int_0^\delta x^{\mu-1}(1-x)^{\mu-1}\exp\left(kR(x)\right)\, dx + \int_\delta^a x^{\mu-1}(1-x)^{\mu-1}\exp\left(kR(x)\right)\, dx}{\int_a^1 x^{\mu-1}(1-x)^{\mu-1}\exp\left(kR(x)\right)\, dx} \\
&\geq \frac{\exp(ks')\int_0^\delta x^{\mu-1}(1-x)^{\mu-1} + \int_\delta^a x^{\mu-1}(1-x)^{\mu-1}\exp\left(kR(x)\right)\, dx}{\exp(k\bar{s})\int_a^1 x^{\mu-1}(1-x)^{\mu-1}\, dx}
\end{aligned}
\tag{41}
$$

Since the second integral in the numerator is bounded below by zero and $s' > \bar{s}$, the expression after the inequality tends to infinity as k tends to infinity, which proves the required result.

23

Proof of Lemma 1

Consider an ergodic finite Markov chain with transition probabilities p_{ij}^k, where k is a parameter and i and j states. A stationary distribution for the chain is the unique probability vector π^k with $\sum_i \pi_i^k p_{ij}^k = \pi_j^k$ for all j. Suppose p_{ij}^k converges to p_{ij}^∞ for each i,j as k tends to infinity and the resulting chain is ergodic with stationary distribution π^∞. Since the set of probability vectors on a finite state space is compact to show that π^k converges to π^∞ (since the state space is finite, weak convergence is equivalent of each component of the vector) it suffices to show that any convergent subsequence of π^k converges to π^∞. Taking limits of the equation for a stationary distribution one finds that the limit of any convergent subsequence is a stationary distribution for p_{ij}^∞, which since an ergodic chain has a unique stationary distribution must be π^∞.

Proof of Theorem 5

Let x^* denote the mixed-strategy equilibrium point. Arguing as in the proof of Theorem 4 it is enough to show that $\nu_\infty^N[a,b]/\nu_\infty^N[a,b]^c$ (where c denotes complement) tends to zero for any $0 < a < x^* < b < 1$ and $\nu_\infty^N[0,x^*)/\nu_\infty^N(x^*,1]$ tends to zero if $x = 1$ is the risk-dominant equilibrium.

To show the first part, note that for a Markov Chain with stationary distribution π_i, for any set A,

$$\sum_{i \in A} \sum_{j \notin A} \pi_i p_{ij} = \sum_{j \notin A} \sum_{i \in A} \pi_j p_{ji} \tag{42}$$

In other words, the net probability flux out of A must be zero. Now for fixed N, the process forms a Markov Chain with state space $\{0, \frac{1}{N}, \ldots, 1\}$. Hence one can apply (42) to $A = [a,b] \bigcap \{0, \frac{1}{N}, \ldots, 1\}$. Now for $i \in A$ there exists $j \notin A$ such that $p_{ij} \geq \left(1 - \frac{\mu}{N}\right)^N$ (one can leap to 0 or 1 by taking a best response — as the mixed strategy point does not belong to the state space). On the other hand for $j \notin A$, $p_{ji} \leq \frac{\mu}{N}$ for all i in A, as it requires at least one mistake to enter A. It follows from (42) that $\frac{\nu_\infty^N[a,b]}{\nu^N[a,b]^c}$ converges to zero as N tends to ∞, as was to be shown.

On the other hand, note that for any point i in $[0,x^*)$, the probability of leaping to a point j in $(x^*,1]$ is the same: namely the probability of making j mistakes. Similarly the probability of leaping to a point in $[0,x^*)$ is independent of the starting point in $(x^*,1]$. It follows from (42) that

$$\nu_\infty^N[0,x^*)B(N, \frac{\mu}{N}, Nx^*) = \nu_\infty^N(1-x^*,1]B(N, \frac{\mu}{N}, N(1-x^*)) \tag{43}$$

where $B(N,p,r)$ denotes the probability of making at least r mistakes when mistakes are drawn from a Binomial distribution with mistake probability p and population size N. (To save notation Nx^* is written rather than the integer part.) Let $B^*(N,p,r)$ denote the probability of making exactly r mistakes. Now it is

straightforward to check that if $r \geq Np$

$$B^*(N, p, r) \leq B(N, p, r) \leq \frac{1-p}{1 - \frac{(N+1)p}{k+1}} B^*(N, p, r) \tag{44}$$

(The first inequality is obvious, the second follows by considering the ratios of the Binomial probabilities — see for example Shorack and Wellner (1986) p. 482.) Applying this to (43), one finds

$$\frac{\nu_\infty^N[0, x^*]}{\nu_\infty^N(x^*, 1]} \leq \frac{B^*(N, \frac{\mu}{N}, Nx^*)}{B^*(N, \frac{\mu}{N}, N(1 - x^*))} \frac{1 - \frac{\mu}{N}}{1 - \frac{(N+1)\frac{\mu}{N}}{Nx^*+1}}$$

$$= \left(\frac{\mu}{N}\right)^{N(2x^*-1)} (1 - \frac{\mu}{N})^{N(1-2x^*)} \frac{1 - \frac{\mu}{N}}{1 - \frac{(N+1)\mu/N}{Nx^*+1}} \tag{45}$$

Letting N tend to infinity, it follows that if $x^* < 1/2$, so that $x = 1$ is the risk-dominant equilibrium, $\nu_\infty^N[0, x^*]/\nu_\infty^N(x^*, 1]$ tends to zero, as was to be shown.

Proof of Theorem 6

For fixed N, the process forms a birth-death process with state-space $\{0, \frac{1}{N}, \ldots, 1\}$. For notational ease the states will be referred to as 0 to N. For more detailed references on birth-death chains see Amir and Berninghaus (1996) or Binmore and Samuelson (1997). Let π_i be the stationary probability that the process has state i. Let θ_i be the transition rate from state i to $i + 1$ and γ_i the transition rate from i to $i - 1$. Then the stationary distribution must satisfy the following detailed balance equation

$$\pi_i \theta_i = \pi_{i+1} \gamma_{i+1} \tag{46}$$

From (19) and (18) it follows that here

$$\theta_i = (N - i) \exp\left(k\frac{\lambda_2\left(\frac{i}{N}\right)}{N}\right) \left[\frac{i}{N}(1 - \mu^N) + \frac{N-i}{N}\mu^N\right] \tag{47}$$

$$\gamma_i = i \exp\left(k\frac{\lambda_1\left(\frac{i}{N}\right)}{N}\right) \left[\frac{N-i}{N}(1 - \mu^N) + \frac{i}{N}\mu^N\right] \tag{48}$$

Using (46) and the above expressions for the transitions rates, it follows that for fixed N as k tends to infinity the probability measure converges to one with all weight on the end points, with all weight on 1 if

$$\frac{1}{N}\sum_{i=0}^{N-1} \lambda_2\left(\frac{i}{N}\right) - \lambda_1\left(\frac{i+1}{N}\right) > 0 \tag{49}$$

If N is large enough this will be true if and only if

$$\int_0^1 (\lambda_2(x) - \lambda_1(x))\, dx > 0 \tag{50}$$

This is precisely the criterion of Theorem 4.

25

References

[1] Amir, M. and S. Berninghaus (1996) 'Another Approach to Learning and Mutation in Games', *Games and Economic Behavior*, **14**, 19–43

[2] Asmussen, S. (1987) *Applied Probability and the Theory of Queues*, John Wiley and Sons, Chichester

[3] Bergin, J. and B. Lipman (1996) 'Evolution with State-Dependent Mutations', *Econometrica*, **64**, 943–956

[4] Billingsley, P. (1968) *The Convergence of Probability Measures*, John Wiley and Sons, New York

[5] Binmore, K., and L. Samuelson (1997) 'Muddling Through: Noisy Equilibrium Selection', *Journal of Economic Theory*, **74**, 235–65

[6] Binmore, K., Samuelson, L. and R. Vaughan (1994) 'Musical Chairs: Modelling Noisy Evolution', mimeo, University College, London

[7] Corradi, V. and R. Sarin (1996) 'Continuous Approximations of Stochastic Evolutionary Game Dynamics', mimeo, Texas A&M University

[8] Durrett, R. (1996) *Stochastic Calculus: A Practical Introduction*, CRC Press, Boca Raton

[9] Ellison, G. (1995) 'Basins of Attraction, Long Run Equilibria and the Speed of Step-by-Step Evolution', mimeo, MIT

[10] Ethier, S. (1976) 'A Class of Degenerate Diffusion Processes Occurring in Population Genetics', *Communications in Pure and Applied Mathematics*, **29**, 483–493

[11] Ethier, S. (1979) 'Limit Theorems for Absorption Times of Genetic Models', *Annals of Probability*, **7**, 622–638

[12] Ethier, S. and T. Kurtz (1986) *Markov Processes: Characterization and Convergence*, John Wiley and Sons, New York

[13] Ewens, W. (1979) *Mathematical Population Genetics*, Springer Verlag, Berlin

[14] Foster, D. and P. Young (1990) 'Stochastic Evolutionary Game Dynamics', *Journal of Theoretical Biology*, **38**, 219–232

[15] Freidlin, M. and A. Wentzell (1984) *Random Perturbations of Dynamical Systems*, Springer Verlag, New York

[16] Fudenberg, D. and C. Harris (1992) 'Evolutionary Dynamics with Aggregate Shocks', *Journal of Economic Theory*, **57**, 420–441

[17] Gillespie, J. (1983) 'Some Properties of Finite Populations Experiencing Strong Selection and Weak Mutation', *The American Naturalist*, **121**, 691–708

[18] Harsanyi, J. and R. Selten (1988) *A General Theory of Equilibrium Selection in Games*, Cambridge, MA

[19] Kandori, M., Mailath, G. and R. Rob 'Learning, Mutation and Long-Run Equilibria in Games', *Econometrica*, **61**, 29–56

[20] Karlin, S. and H. Taylor (1981) *A Second Course on Stochastic Processes*, Academic Press, New York

[21] Kushner, H. (1982) 'A Cautionary Note on the Use of Singular Perturbation Methods for "Small Noise" Models', *Stochastics*, **6**, 117–120

[22] Moran, P. (1958) 'Random Processes in Genetics', *Proceedings of the Cambridge Philosophical Society*, **54**, 60–71

[23] Norman, M.F. (1975) 'Limit Theorems for Stationary Distributions', *Advances in Applied Probability*, **7**, 561–575

[24] Norman, M.F. (1980) 'A "Psychological" Proof that Certain Markov Semigroups Preserve Differentiability', *SIAM-AMS Proceedings*, **13**, 197–211

[25] Revuz, D. and M. Yor (1991) *Continuous Martingales and Brownian Motion*, Springer Verlag, Berlin

[26] Rogers, L. and D. Williams (1987) *Diffusions, Markov Processes, and Martingales*, Volume 2, John Wiley, Chichester

[27] Sandholm, W. (1999) 'Evolution and Equilibrium under Inexact Information', mimeo, University of Wisconsin

[28] Schlag, K. (1996) 'Why Imitate, and if so, How? A Bounded Rational Approach to Multi-Armed Bandits', mimeo, University of Bonn

[29] Shiga, T. (1981) 'Diffusion Processes in Population Genetics', *Journal of Mathematics of Kyoto University*, **21**, 133–151

[30] Shorack, G. and J. Wellner (1986) *Empirical Processes with Applications to Statistics*, John Wiley, New York

[31] Vaughan, R. (1996) 'Evolutive Equilibrium Selection I: Symmetric Two Player Binary Choice Games', mimeo, University College, London

[32] Young, P. (1993) 'The Evolution of Conventions', *Econometrica*, **61**, 57–84

Figures

	1	2
1	a,a	b,c
2	c,b	d,d

Figure 1

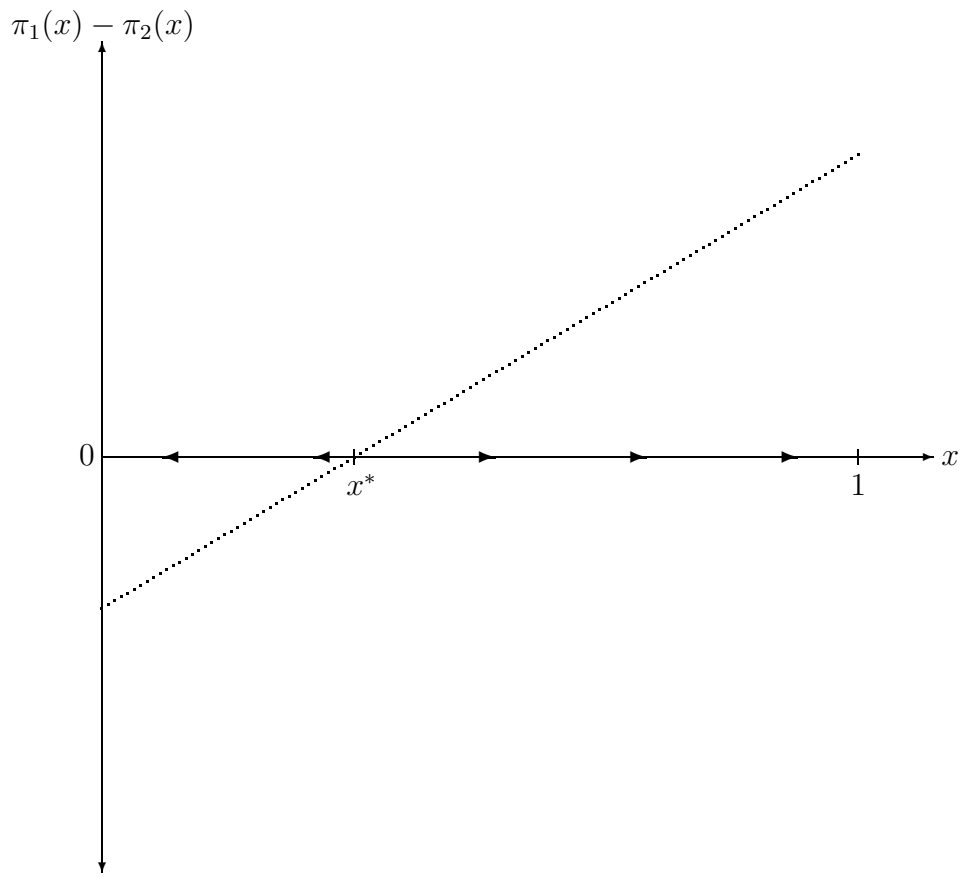|	1	2
1	9,9	0,8
2	8,0	7,7

Figure 2

Figure 3