Evolutionary Dynamics with Large Aggregate Shocks

Daniel B. Neill

October 2003
CMU-CS-03-197

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

This paper examines a class of evolutionary models in which large shocks cause frequent movement between short-term "stable" equilibria. Mutations are rare in our model, but their effects are magnified by a "spread process" which causes a finite proportion of the population to initially adopt the entering strategy before the short-term selection dynamics takes effect. We examine the long run invariant distribution for a variety of games, under several different spread processes: most interestingly, we find that cooperative strategies prevail *in the long run* in the Finitely Repeated Prisoner's Dilemma game, contrary to the backward induction solution. We also study equilibrium selection in 2x2 and NxN coordination games, establishing conditions under which the risk-dominant equilibrium is selected, and demonstrate rapid convergence to the long run invariant distribution.

Keywords: game theory, evolutionary games, large aggregate shocks

1 Introduction

The concept of Nash equilibrium (Nash, 1950) and its various refinements (ex. Selten, 1975) serve as the foundations of non-cooperative game theory, and have been widely applied in economics and many other fields. These equilibrium concepts attempt to answer the question of what actions a perfectly rational player should choose in a given decision situation, under the strong assumptions that players have common knowledge of rationality and unlimited computational powers. These assumptions often do not apply to economic interactions in the real-world, and thus their results sometimes conflict with experimental observations. Moreover, notions of equilibrium fail to answer two basic questions: when does play converge to an equilibrium, and (in games with multiple equilibria) which equilibrium is chosen?

Thus much attention has been drawn to the study of "evolutionary games," which assume a population of players interacting over time, and provide an explicit dynamic process of how players adapt their behavior. These processes include not only evolution in the Darwinian sense, but also individual processes including imitation, adaptation, and learning. Players are not assumed to have common knowledge or perfect rationality, though often some degree of (myopic or bounded) rationality is assumed. The assumption common to all of these models is that of "fitness-dependent selection": strategies which earn higher payoffs have higher fitness and are able to survive and reproduce, while less successful strategies die off. This may be because biological or economic forces select against unsuccessful behaviors, or because players tend to adopt or imitate more successful behaviors. For a more in-depth discussion of evolutionary game theory, the reader is directed to the recent texts of Weibull (1995), Fudenberg & Levine (1998), Vega-Redondo (1996), and Hofbauer & Sigmund (1998), as well as the review articles of Samuelson (2002) and Mailath (1998).

We focus on a class of models which combine two processes: *selection* and *mutation*. Selection is any process by which the proportion of "better" strategies in the population increases, and the proportion of "worse" strategies in the population decreases, where a strategy's fitness is measured by its payoff in the current population. As discussed above, these include both biological processes (fitter strategies are more likely to survive and have offspring), as well as individual processes of adaptation (individuals are more likely to adopt strategies which are more successful). We assume that the selection process does not introduce new strategies into the population; it only changes the proportions of strategies present in the population, possibly causing some strategies present in the population to die off. New strategies are introduced to the population by the process of mutation: we assume that mutations are relatively rare, and that the new mutant strategy is chosen at random. Again, a variety of mutation processes are possible, ranging from genetic mutation, to invasion by a new species, to individual experimentation or discovery of new ideas.

For a useful viewpoint on models combining selection and mutation, we turn to the biological literature on population dynamics (Eshel et al, 1998; Eshel, 1996). Eshel et al describe evolution as a process of "trial and error," in which the population is carried from one state to another as the result of new mutations being introduced into the system. They refer to the change in population shares due to selection as "short-term evolution"; a stable state of the selection dynamics may be a stable (monomorphic or polymorphic) population, or a cyclic or even chaotic fluctuation in the population shares. "Long-term evolution," on the other hand, is "the process whereby successful mutations invade the population, renewing the process of short-term evolution towards a new stable equilibrium, cycle, or state of chaos" (Eshel et al, 1998). Essentially, we can think of the selection-mutation model as a series of time steps, where each time step consists of a mutation (introduction of a new strategy into the population) followed by convergence of the selection dynamics to a stable state (which may be identical to, or substantially different from, the previous stable state). Eshel (1996) notes that it is impossible to precisely predict the long-term behavior of an evolutionary system, because "it is impossible to tell when and in which order the various rare, potentially advantageous mutations will appear in the population." However, in some cases it is possible to find strategies which are "stable" in the long-term, under certain assumptions on the *frequency*, *type*, and *impact* of mutations.

The most influential concept of long-term stability is the "evolutionarily stable strategy" (ESS), as defined by Maynard Smith & Price (1973) and Maynard Smith (1982). Maynard Smith's model assumes a large homogeneous population of a strategy X, which is visited periodically by a small number of mutants playing some alternative strategy A. If players of the common strategy obtain a higher average payoff than players of the alternative strategy, then X resists the invasion, and natural selection eliminates the mutants. Otherwise, the mutants will be able to invade the population, and possibly to replace the common strategy. If strategy X can resist invasion by any alternative strategy A, we say that X is *evolutionarily stable*. The ESS concept makes three important assumptions: first, mutations must be rare enough that the short-term dynamics can evolve to a stable state; in other words, an ESS must be allowed sufficient time to fight off each invasion before a new mutation occurs. Second, only a single type of mutant is allowed to invade on any given time step; a simultaneous invasion by two or more types of mutant is not considered.[1] Third (and most importantly to our discussion), the number of mutants is assumed to be extremely small in relation to the population size. In fact, the initial population share of the mutant strategy is assumed to be infinitesimal: thus each strategy's "average payoff" is simply its payoff against the common strategy, and the strategies' payoffs against the alternative strategy are only relevant in the case of a tie. As we will discuss, however, in many real-world interactions the alternative strategy will be initially adopted by a finite (and possibly large) proportion of the population. In these cases, the self-payoff of the alternative strategy, as well as the payoff of the common strategy against the alternative strategy, become important to the evolutionary dynamics.

Another influential long-term stability concept is that of "stochastic stability," pioneered by Foster & Young (1990), Kandori et al (1993), and Young (1993). Foster and Young were the first to consider games with continual small stochastic shocks; these shocks, accumulated over time, can move the system between the basins of attraction of different equilibria. They argued that, in games with multiple strict Nash equilibria (and thus multiple evolutionarily stable strategies), some equilibria are more likely to emerge than others in the presence of shocks; the "stochastically stable equilibria" are defined as the set of equilibria which occur with finite probability as the amount of noise goes to zero. Kandori et al and Young extend

[1] A stronger notion of stability, due to Boyd and Lorberbaum (1987) requires strategies to resist simultaneous invasion by multiple mutant types; however, this requirement is sufficiently strong that no stable strategies exist for many games.

this concept by presenting methods for calculating which equilibria are stochastically stable in 2x2 and NxN coordination games respectively; we examine their work, as well as several other papers on stochastic stability, in our discussion on equilibrium selection in Section 4. For now, we note that the concept of stochastic stability resembles evolutionary stability in its assumptions on the type and impact of mutations. The difference is in the frequency of mutations: shocks to the population are assumed to be continual rather than extremely rare, and thus a different set of techniques must be used to compute which strategies are stable.

Thus we have considered "evolutionarily stable" strategies, which are robust to isolated small shocks, and "stochastically stable" strategies, which are robust to continuous small shocks. Neither of these concepts consider cases where the impact of a shock is significant: when a mutation introduces a non-negligible proportion of a new strategy into the population. In this paper we consider isolated large shocks: mutations are rare, but each mutation causes a finite (and possibly large) proportion of the population to adopt the entering strategy before the short-term selection dynamics takes effect. If the impact of the shocks is large enough, many games will have no stable equilibrium; in such games, any population can be invaded by a sufficiently large proportion of some other strategy. Thus, we focus on the *invariant distribution* (also called the limiting distribution): the proportion of time, in the long run, that the population spends in each state. Strategies which are more evolutionarily stable (resistant to invasion by other strategies) and more evolutionarily potent (able to invade other strategies) will occur with higher probability in the long run.

Our model relies heavily on the notion of a "spread process," a process that causes the entering strategy to spread to some proportion of the population before the short-term dynamics takes effect. We consider a number of scenarios (including models of evolution, learning, and adaptation) and the type of spread process that might result in each. The spread process can be represented mathematically by a probability density function on the interval [0,1], representing the proportion of the population to which the entering strategy spreads: we call this the "spread function." The long run invariant distribution for a given game is dependent on the spread function; however, under relatively loose constraints, the invariant distribution will be unique (and thus, not dependent on initial conditions) for a given game and spread function.

Our model has a number of interesting features, which distinguish it from other models of evolutionary dynamics. First, it is very general, representing a variety of evolutionary scenarios with large aggregate shocks. Under certain constraints on the spread function, we can even examine a model similar to the stochastic stability models of Kandori et al (1993) and Young (1993), where large shocks can result from the (extremely rare) occurrence of multiple simultaneous mutations. Second, frequent large shocks result in faster convergence to the long-run invariant distribution; thus we can focus on this distribution rather than the variety of "quasi-stable" states which may (depending on the initial conditions) persist for long periods of time in models with small shocks. Third, for coordination games and other games with no stable polymorphisms, calculation of the invariant distribution is generally simpler than finding stochastically stable strategies, since transition probabilities are based only on a notion of pairwise risk dominance. Fourth, the system may spend a large proportion of time far from Nash equilibrium: strategies with high average risk dominance may occur frequently in the long run, even if these strategies are (strictly or weakly) dominated. In fact, we find that cooperative strategies prevail *in the long run* in the Finitely Repeated Prisoner's Dilemma game, contrary to the backward induction solution. This is very different than other evolutionary models, where (partial) cooperation may be quasi-stable, but defection eventually prevails. Finally, a model with large aggregate shocks is less likely to be "stuck" at any given equilibrium for long periods of time; we conjecture that more experimentation allows a population to adapt more quickly to a changing environment, resulting in greater probability of survival in the long run.

In section 2, we present our model of evolutionary dynamics with large aggregate shocks, assuming a spread process represented by an arbitrary spread function. We prove various properties of the model (ex. existence of a unique invariant distribution) under certain constraints on the game and the spread function. In section 3, we justify the assumption of large aggregate shocks by presenting a variety of (evolutionary and adaptive) scenarios where spread processes occur. In sections 4-5, we present results of our model. Section 4 discusses equilibrium selection in 2x2 and NxN coordination games, and several other games with no stable polymorphisms; we examine the conditions under which a risk-dominant equilibrium is selected, and the speed of convergence to the invariant distribution. Section 5 discusses our results for the Finitely Repeated Prisoner's Dilemma game. Section 6 extends the model to games with stable polymorphisms, and presents a variety of examples. In section 7 we consider extensions and applications of the model, and section 8 concludes the paper.

2 The model

We assume a symmetric, two player game in strategic form; let $S = s_1 \ldots s_N$ be the (finite) set of pure strategies available to each player. The game can be represented by an NxN payoff table W, where w_{ij} is the payoff that a player of strategy s_i receives against a player of strategy s_j. We assume an infinite population of players, playing the game represented by W: we assume further that the population is well-mixed, so that a player's payoff is equal to the average of his payoffs against every player in the population. At a given point in time, the population can be represented by a N-dimensional vector $\vec{x}$, where x_i is the proportion of the population playing strategy s_i, and $\sum_{i=1}^{N} x_i = 1$. Then the payoff to a player of strategy s_i can be computed as $u_i = \sum_{j=1}^{N} x_j w_{ij}$. We initially assume that the game has the *no stable polymorphisms* (NSP) property: for any pair of strategies s_i and s_j, either $w_{ii} > w_{ji}$, or $w_{jj} > w_{ij}$, or both. Games with stable polymorphisms will be treated in section 6. We now discuss our evolutionary model in detail. We first consider the short-term dynamics (i.e. the effects of selection alone), then the simple long-term dynamics (i.e. the effects of selection plus isolated small shocks), and finally the long-term dynamics with spread (i.e. the effects of selection plus isolated large shocks).

2.1 Short-term dynamics

As discussed above, the short-term dynamics represents a process of selection: over time, the proportion of "better" strategies in the population increases, and the proportion of "worse" strategies in the population decreases, where a strategy's fitness is measured by its payoff in the current population. In our model, we assume that the short-term dynamics are payoff-monotone and continuous in time (Nachbar, 1990; Friedman, 1991): strategies with higher payoffs u_i have higher growth rates $\frac{\dot{x}_i}{x_i}$, where the dot represents the derivative with respect to time. One such process is the well-known "replicator equation" (Taylor & Jonker, 1978), in which the growth of a strategy's population share is proportional to the difference between the strategy's payoff and the average payoff of the population: $\dot{x}_i = x_i(u_i - \bar{u})$, where $\bar{u} = \sum_{j=1}^{N} x_j u_j$. Though we assume (for concreteness) that the short-term dynamics evolves according to the replicator equation, all payoff-monotone, deterministic dynamics give identical results for two-strategy interactions. On the other hand, the result of a multi-strategy evolutionary interaction may be dependent on different strategies' rates of adjustment; we consider this possibility in more detail in Section 6.

Note that we are considering only "pure selection" dynamics rather than "innovation" dynamics: this means that strategies not currently present in the population will not appear until they are reintroduced by random mutations. This rules out models such as the best response dynamics (Gilboa & Matsui, 1991), in which individuals switch strategies in the direction of the best reply to the current population. Our use of a pure selection dynamics for short-term evolution differs from the work of Kandori et al (1993) and Young (1993), who assume the best response dynamics, but several other papers on stochastic stability (ex. Foster & Young, 1990; Fudenberg & Harris, 1992) assume forms of the replicator dynamics. Because of our assumption of a pure selection dynamics, our model is probably not the best description of a "fully rational" population with complete information and unlimited computational powers; if players were fully rational, we would expect them to be able to compute and play the best response strategy even if that strategy is not currently present in the population. On the other hand, the assumption of a pure selection dynamics makes sense when considering models of imitation (a strategy that is not present cannot be imitated) or Darwinian evolution; it also applies to adaptive players with bounded rationality (or limited knowledge of the strategy space) who cannot simply compute and play a best response.

We are most interested in the stable states of the short-term dynamics: for any combination of strategies, the dynamics will eventually evolve to a stable population, or one of several other limiting distributions (e.g. a limit cycle or chaotic attractor). For two-strategy interactions, the limiting distributions are particularly simple, since limit cycles and chaotic attractors cannot occur. For games with no stable polymorphisms, the limiting distributions are even simpler, since evolution always results in a monomorphic population (a population where all individuals play the same strategy). To see this, we consider the definition of an NSP game: given a pair of strategies s_i and s_j, we have either $w_{ii} > w_{ji}$, or $w_{jj} > w_{ij}$, or both.

If only the first inequality holds, players of strategy s_i always score higher than players of strategy s_j in any population consisting of a mixture of the two strategies. Thus by the selection dynamics, the proportion of s_i increases, and the proportion of s_j decreases, until the population consists entirely of strategy s_i. In this case, any initial proportion of strategy s_i can take over, and completely wipe out, strategy s_j. We say that s_i *evolutionarily dominates* s_j, and write $s_j \to s_i$. Conversely, if only the second inequality holds, s_j will take over any population consisting of a mixture of the two strategies (eliminating strategy s_i). Thus we say that s_j evolutionarily dominates s_i, and write $s_i \to s_j$.

If both inequalities hold, s_i and s_j are in *bistable equilibrium*: either strategy s_i or strategy

s_j will take over the population depending on the initial proportions of the two strategies. Assuming that the initial proportion of strategy s_i is m, and the initial proportion of strategy s_j is $1-m$, we calculate $u_i = mw_{ii} + (1-m)w_{ij}$ and $u_j = mw_{ji} + (1-m)w_{jj}$. Thus strategy s_i will take over if its initial proportion is higher than:

$$m = \frac{w_{jj} - w_{ij}}{w_{ii} - w_{ji} + w_{jj} - w_{ij}}.$$

In this case we write $s_i \overset{m}{\Leftrightarrow} s_j$, or equivalently, $s_j \overset{1-m}{\Leftrightarrow} s_i$.

If neither inequality holds, a *stable polymorphism* is possible: both strategies may survive in the combined population. We consider games with stable polymorphisms in Section 6, and extend our model to deal with these games.

2.2 Long-term dynamics with small shocks

We begin by considering a simple model of selection and mutation: we assume that a large (initially homogeneous) population is visited periodically by a small number of mutants playing an alternative strategy. We first select the initial strategy s randomly (assuming some, possibly uniform, prior distribution) from $s_1 \ldots s_N$. Then a mutant strain s' is selected randomly from our space of strategies, and added to the population; the combined population evolves according to the short-term selection dynamics until it reaches a stable state. Because we have assumed that the game is NSP, the mutant strain s' either takes over the population (replacing strategy s) or dies off, and the population remains homogeneous. This process of random mutation, followed by selection, is repeated indefinitely. This model of invasion dynamics has been used by Nowak & Sigmund (1993) to study the repeated Prisoner's Dilemma game: Nowak and Sigmund perform a computer simulation, and examine which strategies most frequently dominate the population over a large number of time steps. Rather than simulating this random process, however, we prefer to apply probabilistic techniques in order to obtain exact results for the limiting distribution in the long run.

Since the population remains homogeneous, and there are N possible strategies, we can treat the process as a Markov chain with N states, where the state i is determined by which strategy s_i is currently dominating the population. Letting one "time step" consist of introduction of a single mutant strategy, followed by convergence of the selection dynamics to a homogeneous population, we can compute the probability of transitioning from each strategy to each other strategy. The probability p_{ij} of transitioning from s_i to s_j is the product of two probabilities: the prior probability that strategy s_j will attempt to invade, and the probability q_{ij} that the invasion attempt will succeed. We generally assume that mutations are completely random (undirected). Thus the prior probability of selecting any strategy s_j is equal to $\frac{1}{N}$, and $p_{ij} = \frac{q_{ij}}{N} \leq \frac{1}{N}$ for all strategies s_i and s_j, $i \neq j$. The probability of a "self-transition" p_{ii} is the probability that either an invasion will fail, or that strategy s_i will be selected to invade itself: $p_{ii} = 1 - \sum_{j \neq i} p_{ij} \geq \frac{1}{N}$.

Once we have computed p_{ij} for each pair of states i and j, we can compute the invariant distribution $\vec{y}$ of the Markov chain by solving the matrix equation $\vec{y}P = \vec{y}$, subject to the constraint that $\sum_{i=1}^{N} y_i = 1$. However, the Markov chain may or may not have a single ergodic set, and thus the long run invariant distribution may or may not be unique. If the

Markov chain has a single ergodic set, then the invariant distribution $\vec{y}$ is unique: regardless of the initial conditions, we would expect each strategy s_i to dominate the population a proportion of the time that converges in the limit to y_i. Thus y_i can be thought of as a measure of the evolutionary performance of a strategy s_i, for a given game and a given spread function. If the Markov chain has multiple ergodic sets, on the other hand, then the invariant distribution is dependent on the initial conditions. As we prove below, the invariant distribution will be unique whenever the spread function meets certain (relatively loose) conditions.

Now we consider how the invasion success probabilities q_{ij} are computed. If strategy s_j evolutionarily dominates strategy s_i, then s_j will always succeed in invading s_i, and s_i will never succeed in invading s_j: thus $q_{ij} = 1$ and $q_{ji} = 0$. This is why we write $s_i \rightarrow s_j$ for the evolutionary dominance relation: in this case the evolutionary process can only move from i to j, and not in the reverse direction. For a bistable equilibrium $s_i \overset{m}{\Leftrightarrow} s_j$, the probability that an invasion is successful depends on our assumptions about the impact of a mutation (the initial proportion of the entering strategy). A simple model of small shocks, based on Maynard Smith's invasion criteria (1982), would assume that the population of the invader is vanishingly small compared to the population being invaded. Thus s_j invades s_i if $w_{ji} > w_{ii}$, or $w_{ji} = w_{ii}$ and $w_{jj} > w_{ij}$. For this "simple invasion model," if s_i and s_j are in bistable equilibrium, neither strategy will succeed in invading the other: the proportion of mutants always remains below the threshold needed to invade. Thus for strategies in bistable equilibrium, the simple invasion model gives $q_{ij} = q_{ji} = 0$. Note that the resulting Markov chain often has multiple ergodic sets, and in these cases the invariant distribution is dependent on initial conditions. For example, for any 2x2 coordination game, the two strategies are in bistable equilibrium, so the simple invasion model converges to one equilibrium or the other, depending on the proportions of strategies in the initial population.

2.3 Long-term dynamics with large shocks

Now consider a large, initially homogenous population of a strategy s_j, and an attempted invasion by a strategy s_i, where $s_i \overset{m}{\Leftrightarrow} s_j$ and $m \approx 0$. According to the simple invasion model, strategy s_i cannot invade strategy s_j, and $q_{ji} = q_{ij} = 0$. However, in this case it is clear that strategy s_i can take over strategy s_j even if only a very small (but finite) proportion of the population plays s_i. This finite proportion of s_i mutants can arise in a number of ways. If the population is finite, and we assume that a player faces all members of the population including himself, a single s_i mutation results in a finite proportion of s_i. If the population is finite, and a player faces all members of the population but himself, multiple mutations are necessary for the self-payoff of the invader to be relevant.[2] Even in an infinite population, various "spread processes" may cause a finite, and possibly large, proportion of the population to adopt the entering strategy; we consider these processes in detail in the

[2]In this case, we must also take the finite population correction of Schaffer (1988) into account. For very large but technically finite populations, this correction turns out to be irrelevant for games with no stable polymorphisms, but may be relevant in other games; Neill (2003) discusses this issue in more detail, and proposes criteria for a *large population ESS*, as distinct from Maynard Smith's *infinite population ESS* and Schaffer's *finite population ESS*.

next section. Thus the assumption that the proportion of invaders is zero, and thus that $q_{ij} = q_{ji} = 0$ even for very small m, does not necessarily make sense for many evolutionary scenarios.

We define, for each common strategy s_i and invading strategy s_j, a real-valued random variable X_{ij}, restricted to the closed interval [0,1]. Assuming that strategy s_j attempts to invade strategy s_i, X_{ij} determines the proportion of the population to which the invading strategy spreads. We initially assume that the spread of the invading strategy is *independent* of the types of the common or invading strategies: thus we can define a spread process by a single real-valued random variable X, restricted to [0,1]. The distribution of the random variable X can be characterized by the cumulative distribution function $F(x) = \Pr(X \leq x)$, which we call the "cumulative spread function." The cumulative spread function is monotone increasing, and is continuous from the right. We assume that the spread process always results in a population that contains some non-zero proportion of both the common and invading strategies: thus $\lim_{x \downarrow 0} F(x) = F(0) = 0$, and $\lim_{x \uparrow 1} F(x) = F(1) = 1$. We call this the "large aggregate shocks" (LASH) property. For simplicity, we also assume that $F(x)$ is absolutely continuous, and thus we can define the function $f(x) = \frac{dF(x)}{dx}$. Then $f(x)$ is the probability density function for the random variable X: $\int_a^b f(x)\,dx = \Pr(a \leq X \leq b)$. Thus we call $f(x)$ the "spread function." We have $f(x) \geq 0$ for all $x \in [0,1]$, and $\int_0^1 f(x)\,dx = 1$. It will also often be useful to work with the quantity $\Pr(X \geq x) = 1 - F(x)$. We denote this by the function $\bar{F}(x)$.

Now we consider the invasion success probability q_{ij} for a given spread function $f(x)$. As for the simple invasion model, if $s_i \to s_j$, we have $q_{ij} = 1$ and $q_{ji} = 0$. This is because the spread function always results in a mixture of the two strategies; then the selection dynamics will result in s_j taking over the combined population, and wiping out s_i. For bistable equilibria, on the other hand, we no longer have $q_{ij} = q_{ji} = 0$. For $s_i \overset{m}{\Leftrightarrow} s_j$, we know that s_i will invade s_j if the proportion of mutants is at least m: thus $q_{ji} = \Pr(X \geq m) = \int_m^1 f(x)\,dx = \bar{F}(m)$. Similarly, s_j will invade s_i if the proportion of mutants is at least $1-m$: thus $q_{ij} = \Pr(X \geq 1-m) = \int_{1-m}^1 f(x)\,dx = \bar{F}(1-m)$.

We now consider various common types of spread function, computing the transition probabilities q_{ji} and q_{ij} for each when $s_i \overset{m}{\Leftrightarrow} s_j$. The simplest spread function is the "uniform spread function," $f(x) = 1$ for $x \in [0,1]$. In this case, $F(x) = x$, so $q_{ji} = 1-m$, and $q_{ij} = m$. We can also consider an "interval-uniform spread function," $f(x) = \frac{1}{b-a}$ for $x \in [a,b]$, and $f(x) = 0$ otherwise. In this case, $q_{ji} = 1$ if $m \leq a$, $q_{ji} = 0$ if $m \geq b$, and $q_{ji} = \frac{b-m}{b-a}$ if $a \leq m \leq b$. Similarly, $q_{ij} = 0$ if $m \leq 1-b$, $q_{ij} = 1$ if $m \geq 1-a$, and $q_{ij} = \frac{b-(1-m)}{b-a}$ if $1-b \leq m \leq 1-a$. We denote the uniform spread function by U, and the interval-uniform spread functions by $U(a,b)$.

Several "degenerate" spread functions are also important to our discussion: a degenerate spread function is the limit of a sequence of spread functions, but is technically a distribution (generalized function) rather than a function in the strict sense. For example, the "delta spread function" is important for evolutionary scenarios where the population always consists of the same proportion of mutants. This can occur when the probability of an individual agent adopting the entering strategy is identical for all agents, and choices are made independently: since the variance of the proportion of mutants decreases with the population size, the aggregate shock becomes deterministic as the population size goes to infinity. Of course, many other spread functions are possible, since choices are not always independent or identically distributed: for instance, if individuals make choices based on some characteristic (such as aspiration levels) which is assumed to be a uniformly distributed random variable, then this would result in a uniform (or interval-uniform) spread function. The delta spread function $\delta(k)$, for $0 < k < 1$, is defined as the limit, for small ϵ, of the interval-uniform functions centered at k with width 2ϵ: $\delta(k) = \lim_{\epsilon \to 0} U(k-\epsilon, k+\epsilon)$. Thus all of the probability mass of $f(x)$ is concentrated at the point $x = k$. The cumulative spread function $F(x)$ is equal to 0 if $x < k$, 1 if $x > k$, and $\frac{1}{2}$ if $x = k$. Thus for $s_i \overset{m}{\Leftrightarrow} s_j$, $q_{ji} = \bar{F}(m) = 1$ if $m < k$, 0 if $m > k$, and $\frac{1}{2}$ if $m = k$. Similarly, $q_{ij} = 0$ if $m < 1-k$, 1 if $m > 1-k$, and $\frac{1}{2}$ if $m = 1-k$.

Using degenerate spread functions, we can also connect our model to various models in the evolutionary games literature. These include models based on evolutionary stability (and thus, assuming isolated small shocks) as well as models based on stochastic stability (and thus, assuming continuous small shocks). We find that the latter type of model, but not the former, obeys the LASH property.

First, we consider the limit of the interval-uniform spread functions $U(0, \epsilon)$ as $\epsilon \to 0$. In this case, all of the probability mass of $f(x)$ is concentrated at the point $x = 0$. The cumulative spread function $F(x) = 0$ if $x = 0$, and $F(x) = 1$ otherwise; thus for $s_i \overset{m}{\Leftrightarrow} s_j$, we have $q_{ij} = q_{ji} = 0$. Since the invasion success probabilities are identical to those given by Maynard Smith's invasion criteria, we denote this spread function by MS. Note that MS does not satisfy the LASH property given above, since $\lim_{x \downarrow 0} F(x) \neq 0$; thus any results that we prove for LASH spread functions may not hold for this degenerate spread function.

Next, we consider a finite but large population N, with discrete generations and a probability ϵ of mutation. Then let Y be a random variable corresponding to the number of mutations in a given generation: Y is binomially distributed with parameters (N, ϵ). For $\epsilon \approx 0$, the probability that $Y \geq xN$ decreases exponentially as $\epsilon^{xN} = e^{-kx}$, where $k = -N \log \epsilon \gg 0$. Thus we consider cumulative spread functions of the form $F(x) = \frac{1-e^{-kx}}{Z}$, where k is very large, and $Z = 1 - e^{-k} \approx 1$ is a normalizing factor which ensures that $F(1) = 1$. The associated spread function is $f(x) = \frac{dF(x)}{dx} = \frac{ke^{-kx}}{Z}$. For large but finite k, almost all of the probability mass of $f(x)$ is concentrated near $x = 0$, but we still have $f(x) > 0$ for all $x \in [0, 1]$. Now consider two transition probabilities q_{ji} and q'_{ji}, where $s_i \overset{m}{\Leftrightarrow} s_j$ and $s'_i \overset{m'}{\Leftrightarrow} s'_j$. We calculate $q_{ji} = \bar{F}(m) = \frac{e^{-km} - e^{-k}}{1-e^{-k}} \approx e^{-km}$. Similarly, $q'_{ji} = \bar{F}(m') \approx e^{-km'}$. Thus, if $m > m'$, we find $\frac{q_{ji}}{q'_{ji}} \approx e^{-k(m-m')} \approx 0$ for k sufficiently large. Using similar reasoning, we obtain $\frac{q'_{ij}}{q_{ij}} \approx 0$ for large k. This is very similar to the Kandori-Mailath-Rob model: any number of mutations are possible, but since mutations are very unlikely, a transition which requires more mutations is much less likely than a transition that requires less mutations. Hence we call this the KMR spread function; of course, as discussed above, our model is not identical to that of Kandori et al (1993) since it assumes the replicator dynamics rather than the best-response dynamics. We also note that the KMR spread function obeys the LASH property: we have $\lim_{x \downarrow 0} F(x) = F(0) = 0$ and $\lim_{x \uparrow 1} F(x) = F(1) = 1$. Thus any properties that we prove for LASH functions also hold for KMR; this would not be true if we made the assumption of an infinite (rather than finite but very large) value of k.

This brings us to an important question: under what circumstances does a unique invariant distribution exist for a given game W and a given spread function $f(x)$? We examine this question in more detail in Section 4, but a partial answer is given by the following theorem:

Theorem 2.1 *For a given spread function f with associated cumulative spread function F, there exists a unique invariant distribution (independent of initial conditions) for all NSP games if and only if $F(\frac{1}{2}) < 1$.*

Proof Assume $F(\frac{1}{2}) < 1$. Then for each pair of strategies s_i and s_j, either $q_{ij} > 0$ or $q_{ji} > 0$ or both. To see this, we first note that, if $s_i \rightarrow s_j$, $q_{ij} = 1 > 0$, and if $s_j \rightarrow s_i$, $q_{ji} = 1 > 0$. If $s_i \overset{m}{\Leftrightarrow} s_j$, then $q_{ij} = 1 - F(1-m)$, and $q_{ji} = 1 - F(m)$. In this case, at least one of m and $1-m$ is less than or equal to $\frac{1}{2}$. Thus, since F is non-decreasing, $F(m) \le F(\frac{1}{2})$ or $F(1-m) \le F(\frac{1}{2})$. Then since $F(\frac{1}{2}) < 1$, we know that $F(m) < 1$ or $F(1-m) < 1$, and thus $q_{ji} > 0$ or $q_{ij} > 0$. Thus if $F(\frac{1}{2}) < 1$, we know that for each pair of states i and j in our Markov chain, either i is accessible from j, or j is accessible from i, or both. This implies that the Markov chain has a single ergodic set (though the entire chain may not be ergodic, since some states may be transient). To see that there is a single ergodic set, consider the possibility that multiple ergodic sets exist. Then pick states i and j from two separate ergodic sets; either i is accessible from j or j is accessible from i, which contradicts the assumption that both sets are irreducible. Thus we have a Markov chain with a single ergodic set, and possibly some other transient states. We also note that the chain is aperiodic, since every state has a positive probability of self-transition. Thus a unique invariant distribution $\vec{y}$ exists, such that all states i in the ergodic set have positive probabilities y_i, and all transient states j have probabilities $y_j = 0$ (Kemeny & Snell, 1960, p.117). To prove the converse, assume $F(\frac{1}{2}) = 1$. Then the 2x2 coordination game defined by $w_{ii} = w_{jj} = 1$ and $w_{ij} = w_{ji} = 0$ does not have a unique invariant distribution, since both states are absorbing. ∎

This implies that the uniform spread function U always has a unique invariant, as does any interval-uniform spread function $U(a,b)$ where $b > \frac{1}{2}$. Similarly, the delta spread function $\delta(k)$ always has a unique invariant for any $k \ge \frac{1}{2}$. Likewise, the KMR spread function has a unique invariant. The MS spread function, on the other hand, does not always have a unique invariant, and as noted above, the invariant distribution is not unique for any 2x2 coordination game.

Finally, we consider one other property of spread functions: whether the current strategy is given a benefit or a penalty for being established. Most of the time, the spread function will be "conservative": for all $m \in (0,1)$, given a pair of strategies $s_i \overset{m}{\Leftrightarrow} s_j$, the sum of the probabilities q_{ij} and q_{ji} is at most one. This means that, given an opposing strategy s_j, a strategy s_i is at least as likely to take over the combined population when it is the common strategy as when it is the invading strategy. In other words, strategies are not penalized for being established. A special case occurs when the spread function is "fair": for all $m \in (0,1)$, given a pair of strategies $s_i \overset{m}{\Leftrightarrow} s_j$, we have $q_{ij} + q_{ji} = 1$. In this case, neither the common strategy nor the invading strategy is given an advantage, and the probability of s_i taking over the combined population will be identical whether s_i invades or is invaded by s_j. In the next section, we consider several spread processes which result in fair spread functions. We note that the uniform spread function U is fair, as is the $\delta(\frac{1}{2})$ spread function, and any interval-uniform spread function $U(a,b)$ with $\frac{a+b}{2} = \frac{1}{2}$. In fact, a necessary and sufficient condition for fairness is that the spread function is symmetric about $x = \frac{1}{2}$. Conservative but unfair spread functions include $\delta(k)$ for $k < \frac{1}{2}$, as well as the MS and KMR spread functions, and any interval-uniform spread function $U(a,b)$ with $\frac{a+b}{2} < \frac{1}{2}$. It is also possible for the spread function to be "progressive," where the common strategy is penalized for being established, but we conjecture that evolutionary scenarios with this sort of spread function would be rare.

3 Models with large aggregate shocks

We now examine a variety of evolutionary and adaptive scenarios which result in large aggregate shocks, and consider which spread functions would be appropriate for each model. We have already discussed one such scenario: simultaneous multiple mutations in a finite population, as in the stochastic stability model of Kandori, Mailath, and Rob (1993). In this case, we have shown that the spread function is very conservative; $f(x)$ decreases exponentially with the proportion of mutants, and almost all of the probability mass is concentrated near $x = 0$. In the following subsections, we consider a variety of other spread processes which would result in large aggregate shocks even in very large or infinite populations. We divide these into three groups: models of Darwinian evolution, models of learning and adaptation, and single player "imaginary play" models.

We also note that, in all of these models, we continue to make the assumption that the population is well mixed: every individual's fitness is measured by his average payoff against the entire population. If we assume instead that the game is spatially localized, and individuals play only their immediate neighbors, it is possible that the invaders will form spatially clustered groups. This would increase the importance of the invader's self-payoff, since each invader would face a proportion of invaders that is significantly higher than the actual proportion of invaders in the population. As a result, the "effective size" of shocks may be increased. We neglect this effect in the following models, but note that it is another reason why large shocks, rather than only negligible ones, must be considered.

3.1 Models of Darwinian evolution

The first evolutionary scenario that we consider is the invasion of a population by the *migration* of individuals from outside the population. This is very different than invasions resulting from mutations within the population: while we would expect mutations to be isolated (or at least, confined to a kin group), migrants could invade a population in large numbers, possibly even greater than the size of the initial population. Nevertheless, we typically assume that the number of migrants is, on average, smaller than the population being invaded, resulting in a conservative spread function. For example, the simulations of Nowak & Sigmund (1993) assume that the proportion of the invading strategy is only 1% of the combined population, equivalent to a $\delta(.01)$ spread function. A variant of this scenario occurs when "invasions" result from the combination of populations that were formerly spatially isolated. Assume a large number of such populations; at random intervals, two populations meet, combine, and then the combined population evolves according to the short-term selection dynamic. If we focus on a single initial population, this is equivalent to a model where the population is occasionally invaded by a very large number of migrants. On average, we would expect the size of the "invading" population to be equal to the size of the common population, resulting in a fair spread function. If every isolated population is exactly the same size, we would have a $\delta(\frac{1}{2})$ spread function; more likely, the size of each population is dependent on other unmodeled factors, and thus can be treated as a random variable with some specified distribution. For example, if each population is chosen uniformly from the interval $(0, N)$, this would result in a spread function that is fair and approximately uniform. It is also possible that the size of a population is a function of its self-payoff, in which case the spread function becomes dependent on the common and invading strategies.

A second evolutionary scenario is the *temporary shock* model. We assume that the proportion of the mutant strategy in the population is initially small, but a large (but temporary) payoff shock allows the proportion of mutants to grow. This could model major but temporary environmental changes such as natural disasters, disease epidemics, or seasonal fluctuations. This scenario was inspired by the stochastic evolutionary models of Foster & Young (1990), Fudenberg & Harris (1992), and Cabrales (2000), which assume small but continual shocks in the replicator equation (Foster & Young) or the payoffs (Fudenberg & Harris, Cabrales) respectively. The difference is that the stochastic models assume small fluctuations, and thus, that the direction of a shock can change from instant to instant: we assume a large shock which either a) gives the common strategy an advantage, or b) gives the mutant strategy an advantage, for an extended period of time. In the former case, we assume that the small initial proportion of the mutant strategy dies out: this case can be ignored from the perspective of our long-term evolutionary model. In the latter case, the mutant strategy spreads through the population, with any spread function possible depending on the length of the shock and the speed of adjustment. We assume that the duration of a shock is not long enough for the common strategy to be wiped out, and thus the spread function obeys the LASH property. Some types of shock might favor the invading strategy over the common strategy: for instance, if the invading strategy s_j and common strategy s_i are antagonistic, and the invader has the advantage of surprise, this might result in large values of w_{ji} and small values of w_{ij}, until the common strategy can develop defenses against the invader's attacks. On the other hand, natural disasters, diseases, or seasonal changes would result in a random shock direction, either in favor of the invaders or in favor of the common strategy.

Another evolutionary scenario, closely related to the temporary shock model, is *neutral evolution*: an initially neutral mutation spreads through the population by drift, with its effects (i.e. differing payoffs from the common strategy) triggered by later environmental changes or other mutations. Kimura (1983) has argued that such neutral mutations play a significant role in determining the course of evolution. If drift is completely random, and if the boundary states are absorbing, either the common strategy or the mutant strategy will eventually take over the population; if the initial proportion of the mutant is small, the mutant will be wiped out with high probability. If, on the other hand, a constant flux of mutants or migrants causes the boundary states to be reflecting, a uniform spread function will result. Binmore & Samuelson (1999) instead consider directed drift, where the drift results from unmodeled factors whose impact is small enough that they can be neglected except when the selective pressure is weak. In this case, directed drift in favor of an initially neutral mutation would cause that mutation to spread through the population, as in the random shocks model. This directed drift can result from a number of factors: for example, genetic linkage could cause the neutral mutation to be positively correlated with another (evolutionarily beneficial) trait.

The final evolutionary scenario we consider is *Darwinian evolution in a communicating population*. It is well known that the spread of ideas through a communicating population can occur at much shorter time scales than those required for natural selection. In this scenario, a single mutant conceives of an "idea" which is contrary to the established idea and thus prescribes a pattern of behavior different from the common strategy. This idea is then spread (by direct communication or by imitation of behavior) to various members of the population, who in turn spread the idea to others, until the idea has been distributed through the entire population. Some proportion of the population will adopt the invading idea, changing their behavior as a result, and pass the changed behavior on to their offspring by teaching or imitation. We assume that every individual in the original population independently chooses whether or not to adopt the new idea, based on some measure of its "attractiveness." The attractiveness of an idea is not necessarily related to its long-term payoff: consider the rapid spread of cigarette smoking before its long-term health risks were known, and its continued prevalence even in the face of medical warnings and legal restrictions. It is, of course, possible that the attractiveness of a strategy may be positively correlated to its payoff, depending on whether the benefits and costs are immediately apparent; this would result in a strategy-dependent spread function. For simplicity, we assume that attractiveness is strategy-independent: for example, if attractiveness is uniformly distributed, and the population is large enough to ignore variance, this could result in a uniform spread function. Once some proportion of the population has adopted the entering strategy, natural selection can take effect, and the population evolves according to the replicator equation.

3.2 Models of learning, imitation, and adaptation

We now consider a variety of scenarios which are based not on evolution in the Darwinian sense, but on individual adaptation through processes such as learning and imitation. One simple adaptive process is proportional imitation (Schlag, 1998): each instant, some fraction of the population randomly samples another agent's strategy and switches (if the other agent's realized payoff is higher) with probability proportional to how much better the other agent performed. Schlag derives the proportional imitation rule through an axiomatic approach; others have justified proportional imitation by uniformly distributed costs of switching strategies, or uniformly distributed "aspiration levels" (Gale et al, 1995), where an agent becomes active only when his payoff falls below his aspiration level. Schlag has shown that, for a large population, the trajectories of the proportional imitation process converge to the replicator dynamics; moreover, a much broader class of imitation processes converges to some payoff-monotone selection dynamics.

In adaptive models, large aggregate shocks can result from a variety of phenomena. First, if a minority of the population makes their choices based on some criterion other than expected payoff maximization, this can result in non-negligible spread. For example, if a small proportion k of the population are "progressives" who always initially adopt an entering strategy, this will result in a $\delta(k)$ spread function. Alternatively, consider a game with two pure strategies, where most players are "adaptive" (readjusting their strategy at the end of each interaction, according to its relative success) but a small proportion of the population are "fanatics" who maintain their original strategy regardless of its relative success. This model, used in the paper of Bicchieri & Rovelli (1995), would also result in a delta spread function. Yet another possibility is that some proportion k of the population is "Kantian": they choose the strategy they would wish to become a general rule (i.e. the strategy with higher self-payoff). This results in a strategy-dependent spread function: $\delta(k)$ if the self-payoff of the entering strategy is higher, and MS otherwise. In all of these cases, we expect the proportion of the "minority" to be relatively small, and thus we have a conservative spread function. We also note that a distinct "minority" population is unnecessary: if each person has independent probability k of initial adoption, and the population is sufficiently large, this becomes equivalent to a model where a minority of k initially adopts the entering strategy.

Another interesting adaptive scenario results when players have *foresight*: they are not myopic optimizers, but have the ability to choose strategies which will be more successful in the post-invasion environment. This is related to the concept of robustness to equilibrium entrants (Swinkels, 1992), a weaker form of evolutionary stability which requires entering strategies to be a best response to some mixture of the common and entering strategies. However, we assume that players can only choose between the common and entering strategies, rather than "innovating" and playing a best response. Thus entering strategies need only be a *better* reply than the common strategy against some mixture of the two strategies. We imagine a "coalition" model, where a single individual conceives of a new idea, and convinces a group of other players to coordinate decisions to play that strategy. This agreement must be stable against possible deviations from the deviation: but since the only possible deviation we consider is playing the common strategy instead of the alternative strategy, the agreement is stable whenever the invading strategy's self-payoff is greater than the payoff of the common strategy against the invading strategy. Then any spread function is possible, depending on the size of the coalition formed. However, if players are sufficiently rational (and this is common knowledge), we might also want to place another condition on the agreement: players will only agree to play a deviation which will (in the long run) increase their average payoff, and thus the self-payoff of the entering strategy must also be higher than the self-payoff of the common strategy. In this case, we again have a strategy-dependent spread function: any spread function (depending on the distribution of coalition sizes) if the entering strategy has higher self-payoff, and the MS spread function if the common strategy has higher self-payoff.

Our final group of adaptive scenarios are *uncertainty models*, where the spread function results from individuals' initial uncertainty about the relative benefits or harms of the entering strategy. We assume that all individuals are informed about their possible choices (either the common or entering strategy) and cannot "innovate" and choose a different strategy. Individuals receive noisy statistical information about the expected payoffs of each strategy in the population; at random intervals, a small proportion of the population is given the option of switching strategies, and each individual chooses the strategy that he believes to have higher payoff. This is similar but not identical to the directed imitation model of Bjornerstedt & Weibull (1995) and Weibull (1995), in which a player receives noisy information about his own expected payoff and the expected payoff of a randomly sampled player, switching if the other strategy's payoff is higher. We consider several cases in which players are initially uncertain about the payoffs of the entering strategy, resulting in large aggregate shocks.

The first case we consider is when individuals initially believe (incorrectly) that they know the relative performance of the entering strategy. In this case, we assume that the statistical information that players receive is so noisy that each individual's decision is independent of the true relative performance. Let k equal the probability that an individual believes that the entering strategy's payoff is higher; we assume that this probability is equal for all individuals in the population. Then the proportion of the entering strategy converges toward k, and if the duration of payoff uncertainty is long enough, this results in a $\delta(k)$ spread function. Once the amount of uncertainty is reduced to more normal levels, selective pressures will push the population in the direction of the strategy with higher actual expected payoff, and the population will evolve according to a payoff-monotone dynamics.

The second case we consider is when the relative performance of the entering strategy is unknown, but the individuals *know* that they do not know its relative performance. In this case, players may be completely indifferent until the uncertainty is resolved. If players occasionally switch strategies by imitating another player at random, or (rarely) choosing one of the two strategies at random, this results in an undirected random walk with reflecting boundaries, and hence a uniform spread function. Alternatively, choices whether to switch strategies may be controlled by parameters such as "desire for novelty" and the cost of switching strategies, and the spread function is determined by the distribution of these parameters.

The third and final case of large aggregate shocks in uncertainty models is when imitation is driven by dissatisfaction (Bjornerstedt, 1995). In Bjornerstedt's model, agents with less successful strategies switch strategies (by imitating an individual chosen at random from the population) at a higher rate than individuals with more successful strategies. Bjornerstedt shows that this converges to a payoff-monotone dynamics, and if the review rate is linearly decreasing in current payoff, it converges to the replicator dynamics. We assume that, initially, individuals switch only if and only if they are dissatisfied with the old strategy; since they do not know the payoffs of the new strategy yet, they cannot be dissatisfied with it. The proportion dissatisfied may be dependent on the self-payoff of the common strategy, in which case the spread function is strategy-dependent, or dependent on other unmodeled factors, in which case the spread function may be strategy-independent. In all of these cases, once the uncertainty of relative payoffs is resolved, the population will evolve according to a payoff-monotonic imitation dynamics, whether this results from proportional imitation, directed imitation, or imitation due to dissatisfaction.

3.3 Single player, "imaginary play" models

Our final class of scenarios is very different than typical models of evolution or of adaptation. Instead, we consider a single, boundedly rational player, who is attempting to decide which strategy to play in an NSP game. However, the player does not know all of his possible strategy choices, nor does he have the computational ability to decide which strategy is a best reply to a given strategy. Thus his method, rather than optimizing, is one of "brainstorming": he maintains a current "favorite" strategy and continually compares a randomly selected ("discovered") strategy to his current favorite, switching strategies if the new strategy is "better." Since play is imaginary, there is no cost to switching strategies. Thus the current strategy is given no bonus for being established; strategies are compared on their own merits. If the new strategy evolutionarily dominates the current strategy, the player will switch strategies; conversely, if the current strategy evolutionarily dominates the new strategy, the player will not switch. If the two strategies are bistable, there are several possibilities. One option is to always keep the strategy that "risk-dominates" the other. Strategy s_i risk-dominates strategy s_j if $w_{ii} - w_{ji} > w_{jj} - w_{ij}$. Thus, if $s_i \overset{m}{\Leftrightarrow} s_j$, s_i risk-dominates s_j when $m < \frac{1}{2}$. Thus, if the player chooses always to keep the risk-dominant strategy, his choice of strategy evolves as in a LASH model with $\delta(\frac{1}{2})$ spread function. Alternatively, rather than a deterministic approach, the player can take a stochastic approach, where he keeps the risk-dominant strategy with probability proportional to its level of risk-dominance; this has the advantage of distinguishing between stronger and weaker risk-dominance relations. Thus, if $s_i \overset{m}{\Leftrightarrow} s_j$, he would choose s_i with probability $1 - m$. In this case, his choice of strategy evolves as in a LASH model with uniform spread function. Many other spread functions are possible; we assume that the common strategy is given no bonus or penalty for being established, and hence the resulting spread function is fair. Thus, in addition to a variety of evolutionary and adaptive scenarios, our model of large aggregate shocks also describes the decision-making process of a rational player with limited knowledge of the strategy space: a player who cannot simply compute and choose the optimal strategy, but must instead continually explore and test new alternatives.

4 Equilibrium selection

One of the most widely discussed topics in game theory is equilibrium selection in games with multiple Nash equilibria. There exists an extensive literature on refinements of Nash equilibrium, and procedures for convergence to equilibrium; for a general reference, see Harsanyi and Selten (1988), and for connections to evolutionary games, see Samuelson (1997). For the purposes of our discussion, we focus on two major concepts, risk-dominance and p-dominance. Harsanyi and Selten (1988) define *risk-dominance* for symmetric NxN games: a strategy s risk-dominates a strategy s' if s is a better response than s' against a mixture $\frac{1}{2}s + \frac{1}{2}s'$. They then define a *risk-dominant equilibrium* as a Nash equilibrium that risk-dominates every other Nash equilibrium; not every game has a risk-dominant equilibrium, but if one exists, it is unique. For NSP games, we note that s risk-dominates s' if $s' \to s$, or if $s \overset{m}{\Leftrightarrow} s'$ and $m < \frac{1}{2}$. Morris, Rob, and Shin (1995) define the stronger notion of p-dominance: a symmetric Nash equilibrium (s,s) is p-dominant if s is a strict best response to any mixed strategy placing probability at least p on s. Thus $\frac{1}{2}$-dominance is a refinement of risk-dominance: every $\frac{1}{2}$-dominant strategy is risk-dominant, but not the reverse.

4.1 Simple 2x2 games

Our main focus for equilibrium selection, as is common in the literature, is coordination games, or games with two or more strict Nash equilibria. But as a preliminary exercise (and demonstration of the properties of our evolutionary model), we first consider two simple symmetric 2x2 games: the Prisoner's Dilemma (with payoffs given in Table 1), and the joint venture game (with payoffs given in Table 2).

In the Prisoner's Dilemma, there is a single strict Nash equilibrium, (Y,Y). Since $X \rightarrow Y$, strategy Y takes over any polymorphic population containing both X and Y. Since we assume a LASH spread function, any mutation results in a population consisting of a mixture of X and Y, and then the short-term selection dynamics converge to a state containing only Y. Thus the long run invariant distribution places all of the probability mass on Y for any LASH spread function.

In the joint venture game, there are two Nash equilibria: the strict equilibrium (X,X) and the weak equilibrium (Y,Y). For this game we calculate $Y \rightarrow X$. Again, since the spread function is LASH (and thus, the initial population always contains some proportion of the strict Nash equilibrium strategy), the short-term dynamics always converges back to the strict Nash equilibrium. Thus the invariant distribution places all of the probability mass on X for any LASH spread function.

Thus, in both of these examples, the size of the mutation is not important: as long as there is some non-zero proportion of the dominant strategy, the short-term dynamics will always converge to that strategy. For games with multiple strict equilibria, on the other hand, the short-term dynamics can converge to either equilibrium, depending on the basin of attraction in which the short-term process begins. Hence significant shocks are needed to move the long-term process between strict equilibria, whether resulting from the aggregation of small shocks (as in the stochastic stability literature) or resulting from isolated but large shocks as in the present model.

4.2 Coordination games

As discussed above, equilibrium selection in games with multiple strict equilibria has been a major focus of the evolutionary games literature, especially in the recent literature on models of stochastic evolutionary dynamics. Here we consider two 2x2 coordination games: in each game, both strategy X and strategy Y are strict Nash equilibria. In the game given in Table 3, strategy Y risk-dominates strategy X; the two strategies are in bistable equilibrium with $X \overset{4/5}{\Leftrightarrow} Y$. Also, Y Pareto-dominates X, giving higher payoffs to both players. Thus this is a simple coordination game, where the risk-dominant and Pareto-dominant equilibria coincide. In this case, nearly all of the stochastic evolutionary models in the literature will choose the equilibrium that is both risk-dominant and Pareto-dominant. A more interesting situation occurs in the game given in Table 4: strategy Y again risk-dominates strategy X, with $X \overset{4/5}{\Leftrightarrow} Y$, but in this case X is Pareto-dominant. In this case, the question of which equilibrium to select is more difficult; while most evolutionary models choose the risk-dominant equilibrium, some choose the Pareto-dominant equilibrium instead.

Our model is closest in spirit to the models of Kandori et al (1993) and Young (1993), since in these models the long run invariant distribution depends on the probability of large jumps across basins of attraction, as opposed to depending on the accumulation of smaller shocks over time (as in Foster & Young, 1990, and Fudenberg & Harris, 1992). Kandori et al consider a discrete time, finite population best-response dynamics, where each period each player mutates with a small probability ϵ. Young considers sampled fictitious play models, in which players with finite memory choose best responses to a sample of the history, but can make mistakes. In both of these models, movement between equilibria takes a large number of simultaneous mutations, and thus it can be represented by an exponentially decreasing spread function. They then look at the limit as the probability of mutation (or error) goes to 0, resulting in the (degenerate) KMR spread function given above. In a 2x2 coordination game, both dynamics converge to the risk-dominant equilibrium.

Other stochastic stability models, where small stochastic shocks accumulate over time, cannot be easily represented by a LASH model. Foster and Young (1990) add a Brownian motion noise term to the replicator equation, giving a system of stochastic differential equations. Because of their reliance on gradual local movements, their equilibrium selection results depend not only on the payoffs, but also on the details of the selection and mutation processes. Similarly, Fudenberg & Harris (1992) assume stochastic payoff shocks. Their dynamics does not result in a unique invariant distribution; it converges to one of the equilibria with probabilities depending on the initial conditions. This is because the perturbation in payoffs does not have a significant effect on population shares when the population is nearly homogeneous. They also consider a further modification of the replicator dynamics assuming a deterministic flow of mutations as well as payoff shocks. This makes the system ergodic, and results in convergence to the risk-dominant equilibrium.

Thus most of the models above, when they give unique equilibrium selection results, select the risk-dominant equilibrium over the Pareto-dominant equilibrium. However, certain models give the opposite results. Robson & Vega-Redondo (1996) consider random matching of players (rather than the assumption of "expected payoffs" or round-robin matching). Because transitions can occur when sufficiently many players of a Pareto-dominant strategy are randomly matched together (causing the rest of the population to switch to this strategy), their model results in convergence to the Pareto-dominant equilibrium in 2x2 symmetric coordination games. The "noisy selection" models of Binmore et al (1995) and Binmore & Samuelson (1997), representing a "musical chairs" model of evolution and a "muddling" model of learning respectively, give still different equilibrium selection results: their dynamics converge to either the Pareto-dominant or risk-dominant equilibrium depending on the payoffs and certain parameters of the model. The difference is that the noisy selection models depend on the absolute payoffs as compared to the background "death rate" or revision probability. In the models of Kandori et al and Young, on the other hand, fitnesses depend only on the relative payoffs $w_{ii} - w_{ji}$ and $w_{jj} - w_{ij}$, resulting in selection according to risk-dominance rather than payoff-dominance.

Our model, like the models of Kandori et al and Young, considers only the relative payoffs, and hence we would expect it to select the risk-dominant strategy. In fact, this is the case: for symmetric 2x2 coordination games, the long run invariant distribution of our model places more probability mass on the risk-dominant strategy, regardless of whether or not this strategy is Pareto-dominant. We note that $X \overset{4/5}{\Leftrightarrow} Y$ for both of the coordination games presented above; hence the transition probabilities, and our evolutionary results, are identical for both games. However, the exact results depend on the spread function: we calculate $q_{XY} = \bar{F}(\frac{1}{5})$ and $q_{YX} = \bar{F}(\frac{4}{5})$. Thus, if $F(\frac{1}{5}) < 1$, we have $q_{XY} > 0$, so the system has a unique invariant distribution. In this case, the invariant probability mass of X is $\frac{\bar{F}(\frac{4}{5})}{\bar{F}(\frac{1}{5})+\bar{F}(\frac{4}{5})}$, and the invariant probability mass of Y is $\frac{\bar{F}(\frac{1}{5})}{\bar{F}(\frac{1}{5})+\bar{F}(\frac{4}{5})}$. Since $\bar{F}(\frac{1}{5}) \geq \bar{F}(\frac{4}{5})$, the system places at least half the mass on the risk-dominant equilibrium Y; this inequality is strict for all spread functions with probability mass in the interval $(\frac{1}{5}, \frac{4}{5})$. Also, note that if the system has a unique invariant distribution and $F(\frac{4}{5}) = 1$, then all of the probability mass is placed on the risk-dominant equilibrium. For example, the delta spread function $\delta(k)$, for any $\frac{1}{5} \leq k < \frac{4}{5}$, places all of the probability mass on the risk-dominant equilibrium. For the uniform spread function U, we have $F(x) = x$, so we calculate $q_{YX} = \frac{1}{5}$ and $q_{XY} = \frac{4}{5}$. Thus the risk-dominant equilibrium gets 80% of the probability mass for the uniform spread function. For the KMR spread function, we have $\frac{q_{YX}}{q_{XY}} \approx 0$, so the risk-dominant equilibrium gets almost all of the probability mass. Finally, as noted above, the MS spread function does not have a unique invariant distribution: it converges to either X or Y depending on the initial conditions.

4.3 NxN coordination games

We now consider equilibrium selection for NxN coordination games, focusing on an example game given in Young (1993). For this game, given in Table 5, we calculate $X \overset{7/8}{\Leftrightarrow} Y$, $X \overset{4/7}{\Leftrightarrow} Z$, and $Y \overset{3/5}{\Leftrightarrow} Z$. Thus Z risk-dominates X and Y, and Y risk-dominates X. Since Z is the risk-dominant equilibrium in this game, we might expect the models of Kandori et al (1993) and Young (1993) to converge to this equilibrium. However, as Young shows, both models converge to strategy Y instead, and thus do not always choose the risk-dominant strategy. We note that, though strategy Z is risk-dominant, it is not $\frac{1}{2}$-dominant, since Y (not Z) is the best reply to a population consisting of $\frac{1}{2}X + \frac{1}{2}Z$. Maruta (1997) and Ellison (2000) show that the Kandori et al and Young models choose the $\frac{1}{2}$-dominant strategy, if one exists, in an NxN coordination game.

We now consider the results of our model for this game. We note that state Z is absorbing whenever $F(\frac{4}{7}) = 1$, state Y is absorbing whenever $F(\frac{2}{5}) = 1$, and state X is absorbing whenever $F(\frac{1}{8}) = 1$. Thus, any spread function with $F(\frac{2}{5}) < 1$ and $F(\frac{4}{7}) = 1$ will converge to an invariant distribution which places all of the probability mass on Z. This includes the $\delta(\frac{1}{2})$ spread function, which as we show below, always converges to the risk-dominant equilibrium if one exists.

Theorem 4.1 *For any symmetric coordination game W with a risk-dominant equilibrium (s_i, s_i), the LASH model with $\delta(\frac{1}{2})$ spread function converges to an invariant distribution which places all of the probability mass on s_i.*

Proof For each other strategy s_j, we know $s_i \overset{m}{\Leftrightarrow} s_j$ with $m < \frac{1}{2}$. Thus $F(m) = 0$ and $F(1-m) = 1$ for the $\delta(\frac{1}{2})$ spread function. This implies that $q_{ji} = \bar{F}(m) = 1$ and $q_{ij} = \bar{F}(1-m) = 0$. Hence s_i is absorbing, and accessible from every other strategy, so it is the unique absorbing state. ∎

There are two other possibilities for the spread function. If $F(\frac{2}{5}) = 1$, then both Y and Z are absorbing. In this case, the invariant distribution is not unique, and is dependent on initial conditions. If, on the other hand, $F(\frac{4}{7}) < 1$, then none of the equilibria are absorbing; as a result, the Markov chain is ergodic, and all states will have non-zero probability mass in the invariant distribution. Nevertheless, the KMR spread function will place almost all of the probability mass on the risk-dominant equilibrium Z for any given (large) rate of exponential decrease k. As k goes to infinity, the invariant distribution converges to one which places all probability mass on Z. In fact, the KMR spread function always selects the risk-dominant equilibrium if one exists:

Theorem 4.2 *For any symmetric coordination game W with a risk-dominant equilibrium (s_i, s_i), the LASH model with KMR spread function converges to an invariant distribution which places all of the probability mass on s_i.*

Proof For each other strategy s_j, we know $s_i \overset{m}{\Leftrightarrow} s_j$ with $m < \frac{1}{2}$. For the KMR spread function $F(x) = \frac{1-e^{-kx}}{Z}$, the transition probability $q_{ji} = \bar{F}(m) = \frac{e^{-km}-e^{-k}}{Z}$. Let $m_0 < \frac{1}{2}$ be the highest value of m such that $s_i \overset{m}{\Leftrightarrow} s_j$. Then the probability p_{in} of entering state i, if originally in a different state, is at least $\frac{e^{-km_0}-e^{-k}}{ZN} \approx \frac{e^{-km_0}}{N}$. Similarly, the probability p_{out} of leaving state i is at most $\frac{(N-1)(e^{-k(1-m_0)}-e^{-k})}{ZN} \approx \frac{(N-1)(e^{-k(1-m_0)})}{N}$. Thus the ratio $\frac{p_{in}}{p_{out}}$ is at least $\frac{e^{-km_0}}{(N-1)(e^{-k(1-m_0)})} = \frac{1}{N-1}e^{-k(2m_0-1)}$. Since N and m_0 are constants, and $m_0 < \frac{1}{2}$, we know that $\frac{p_{in}}{p_{out}}$ goes to infinity as k goes to infinity, and hence the invariant distribution places all of the probability mass on s_i. ∎

We note that this result is very different than the results of Young (1993), even though the spread processes are essentially identical (i.e. multiple rare mutations in a large population). In Young's 3x3 game, our model chooses the risk-dominant equilibrium Z while his model chooses equilibrium Y. This difference results from the fact that Young's model assumes the best-response dynamics rather than the replicator dynamics. As a result, when computing the probability of a strategy s_j taking over a strategy s_i, he calculates the smallest proportion of mutants of *any sort*, not just the smallest proportion of s_j mutants, needed to go from s_i to s_j. Thus the "shortest path" from s_i to s_j might actually be through invasion by a third strategy s_k, or even a combination of strategies s_j and s_k. If s_j is a strict best response to the combined population, it will take over even if it is not currently present in the population. For example, in Young's game, strategy Y can take over strategy Z if the population is invaded by at least a $\frac{3}{8}$ proportion of X mutants. Moreover, the alternative strategy must be a *best response* to take over the population, not simply a better response than the common strategy: thus, a $\frac{3}{8}$ proportion of X mutants cannot take over Z, because Y (not X) is the best response. These complications make finding the stochastically stable equilibrium difficult, since a best possible spanning tree must be found.[3] Our model, based only on pairwise risk-dominance comparisons, is much simpler to compute.

Young notes that risk-dominance and stochastic stability differ in two respects. First, risk-dominance considers only the subgame consisting of the two strategies, while stochastic stability examines all transitions between the two strategies, including those that go

[3]For more details on this process, see Young (1993), Kandori et al (1993), and Ellison (2000).

through other strategies. Second, a risk-dominant strategy does not always exist, while a best spanning tree (and thus, a set of stochastically stable equilibria) can always be found. Our method keeps the simplicity of risk-dominance calculations, while avoiding Young's criticisms. The key is that pairwise risk dominance is used to determine transition probabilities in our model: thus, even though the short-term dynamics only considers subgames consisting of two strategies, indirect transitions are essential to the long-term dynamics. Moreover, for a large class of spread functions, a unique invariant distribution can be found, even in games with no risk-dominant strategy.

Finally, we compute the results for Young's 3x3 game, with the uniform spread function U. In this case, we find the invariant distribution $[X\ Y\ Z] = [.1601\ .4231\ .4168]$. Thus Y and Z have nearly equal probability mass, but Y has slightly greater probability even though Z is risk-dominant. This result occurs because Z transitions to X far more often than Y does: $p_{ZX} = \frac{1}{7}$ and $p_{YX} = \frac{1}{24}$. Thus Y outperforms Z even though $p_{YZ} = \frac{1}{5}$ is greater than $p_{ZY} = \frac{2}{15}$. Though Young's model also selects equilibrium Y, this occurs for a very different reason: the stability of Z is weakened because of an indirect short-term transition from Z to Y through X, which is impossible in our model. Also, Young's equilibrium selection result is strong (all probability mass is placed on Y), while our equilibrium selection result places almost as much mass on Z. We believe that this is a reasonable solution to Young's game, where there are good reasons to select either Y or Z; compare this to the results for the 2x2 coordination games above, where we select the risk-dominant equilibrium much more strongly (with probability mass 0.8 in the long run invariant distribution). Thus the uniform spread function allows greater discrimination between the relative strengths of equilibria in different games. We also note two other advantages, over the delta and KMR spread functions respectively. First, the uniform spread function can distinguish between equilibria based on amount of risk-dominance, rather than simply a binary decision as to which strategy risk-dominates the other; this is also true for KMR, or any other spread function where $f(x) > 0$ for $x \in (0,1)$, but is not true for delta spread functions. Second, spread functions such as KMR (where almost all of the probability mass of $f(x)$ is near $x = 0$) result in very long waiting times before the distribution of states converges to the long run invariant distribution. Functions with larger spread, such as the uniform or $\delta(\frac{1}{2})$ spread functions, do not share this problem; we consider this issue in more detail below.

4.4 Other NSP games

Next we consider two other examples of NSP games that are not coordination games: games in which both evolutionary dominance relations and bistable equilibria occur. First, consider the game in Table 6, initially assuming $\mu = 1$. In this game we have a cycle of evolutionary dominance relations: $X \to Y \to Z \to W \to X$. Imagining the four strategies as corners of a square, we expect movement around the square in a single direction, from X to Y to Z to W, back to X. Also, since $X \overset{1/2}{\Leftrightarrow} Z$ and $Y \overset{1/2}{\Leftrightarrow} W$, we can also have transitions across the diagonal of the square if $F(\frac{1}{2}) < 1$. In any case, since the system is both ergodic and symmetric, we know that the invariant distribution places 25% of the probability mass on each of the four pure strategies. Now let us consider the same payoff table with $\mu = 1.5$. X is still dominated by Y and still dominates W, but now $X \overset{2/5}{\Leftrightarrow} Z$. Thus, assuming that

$F(\frac{3}{5}) > F(\frac{2}{5})$, we would expect X to have the most, and Z to have the least, probability mass in the invariant distribution. For a uniform spread function, we find the invariant distribution $[X\ Y\ Z\ W] = [.27\ .26\ .23\ .24]$. For the $\delta(\frac{1}{2})$ spread function, we find $[X\ Y\ Z\ W] = [.35\ .30\ .15\ .20]$. Thus when the symmetry is broken by increasing the self-payoff of one strategy, its probability mass increases as expected. However, for the KMR spread function, we find convergence to $[\frac{1}{4}\ \frac{1}{4}\ \frac{1}{4}\ \frac{1}{4}]$. This is because, in the limit as $k \to \infty$, the bistable transitions become extremely rare and hence irrelevant. Similarly, for a $\delta(k)$ spread function with $k < \frac{2}{5}$, the bistable transitions never occur, so we again find convergence to $[\frac{1}{4}\ \frac{1}{4}\ \frac{1}{4}\ \frac{1}{4}]$. For this game, the KMR and some delta spread functions fail to pick out the obvious focal point equilibrium.

Next, we consider the game in Table 7. Note that we have a 3-cycle of evolutionary dominance relations: $X \to Y \to Z \to X$. Also, W is bistable with each of X, Y, and Z; we calculate $X, Y, Z \overset{\frac{\mu}{1+\mu}}{\Leftrightarrow} W$. Thus the probability mass on W is $\frac{\bar{F}(\frac{1}{1+\mu})}{\bar{F}(\frac{1}{1+\mu})+3\bar{F}(\frac{\mu}{1+\mu})}$, and the total probability mass on X, Y, and Z is $\frac{3\bar{F}(\frac{\mu}{1+\mu})}{\bar{F}(\frac{1}{1+\mu})+3\bar{F}(\frac{\mu}{1+\mu})}$; by symmetry, we know the probability is split evenly between the three strategies, and each has probability $\frac{\bar{F}(\frac{\mu}{1+\mu})}{\bar{F}(\frac{1}{1+\mu})+3\bar{F}(\frac{\mu}{1+\mu})}$. For a uniform spread function this results in the invariant distribution $[X\ Y\ Z\ W] = [\frac{1}{\mu+3}\ \frac{1}{\mu+3}\ \frac{1}{\mu+3}\ \frac{\mu}{\mu+3}]$. For example, if $\mu = 2$, we have $[X\ Y\ Z\ W] = [\frac{1}{5}\ \frac{1}{5}\ \frac{1}{5}\ \frac{2}{5}]$. For the $\delta(\frac{1}{2})$ and KMR spread functions, three different distributions are possible depending on the value of μ. If $\mu > 1$, W gets all of the probability mass: since it risk-dominates each of the other strategies, it is an absorbing state. If $\mu < 1$, the three strategies X, Y, and Z evenly split the probability mass: since W is risk-dominated by each of the other strategies, it is a transient state and gets no probability mass. If $\mu = 1$, then $F(\frac{1}{1+\mu}) = F(\frac{\mu}{1+\mu}) = F(\frac{1}{2})$, so all four strategies get $\frac{1}{4}$ of the probability mass by symmetry.

4.5 General results for NSP games

We now prove some general results for NSP games, based on a notion of *pairwise risk-dominance*. As in Neill (2001), we define the pairwise risk-dominance $\mathrm{dom}(s_i \mid s_j) = 1$ if $s_j \to s_i$, 0 if $s_i \to s_j$, and $1 - m$ if $s_i \overset{m}{\Leftrightarrow} s_j$. Then we note the following three properties of pairwise risk-dominance. First, $\mathrm{dom}(s_i \mid s_j) + \mathrm{dom}(s_j \mid s_i) = 1$. Second, $q_{ij} = \bar{F}(\mathrm{dom}(s_i \mid s_j))$. Third, $q_{ji} = \bar{F}(\mathrm{dom}(s_j \mid s_i))$. Now, let a *path* from strategy s_i to strategy s_j be a sequence of strategies $s_{(1)} = s_i, s_{(2)}, \ldots, s_{(k)} = s_j$. Let the *shock threshold* st of a path be the maximum value of $\mathrm{dom}(s_{(i)} \mid s_{(i+1)})$ for $i = 1 \ldots k-1$. It is clear that s_j is accessible from s_i through the path $s_{(1)} \ldots s_{(k)}$ if and only if $F(\mathrm{st}) < 1$. Let the *minimum shock threshold* $\mathrm{mst}(i \to j)$ be the minimum value of st over all paths from s_i to s_j. It is clear that s_j is accessible from s_i if and only if $F(\mathrm{mst}(i \to j)) < 1$. We also note that $\mathrm{mst}(i \to j) \leq \mathrm{dom}(s_i \mid s_j)$ for all i and j, since the direct path from s_i to s_j has shock threshold $\mathrm{dom}(s_i \mid s_j)$.

Now, we define $\mathrm{mst}(i \leftrightarrow j)$ to be the minimum of $\mathrm{mst}(i \to j)$ and $\mathrm{mst}(j \to i)$. It is clear that either s_j is accessible from s_i or s_i is accessible from s_j, if and only if $F(\mathrm{mst}(i \leftrightarrow j)) < 1$. We also note that $\mathrm{mst}(i \leftrightarrow j) \leq \min(\mathrm{dom}(s_i \mid s_j), \mathrm{dom}(s_j \mid s_i)) \leq \frac{1}{2}$ for all i and j. Now define the *maximum pathwise bistability* (mptb) of an NSP game as the maximum value of $\mathrm{mst}(i \leftrightarrow j)$ for all strategies s_i and s_j, $i \neq j$. Also, we define the *maximum pairwise* *bistability* (mprb) of an NSP game as the maximum value of $\min(\mathrm{dom}(s_i \mid s_j), \mathrm{dom}(s_j \mid s_i))$ for all strategies s_i and s_j, $i \neq j$. It is clear that mptb $\leq$ mprb $\leq \frac{1}{2}$ for every NSP game. We now prove the following theorem:

Theorem 4.3 *For a given NSP game with mptb* $= k$*, and for a given spread function* $f(x)$ *with associated cumulative spread function* $F(x)$*, the long run invariant distribution is unique if and only if* $F(k) < 1$*.*

Proof As we argue above, either s_i is accessible from s_j or s_j is accessible from s_i, if and only if $F(\mathrm{mst}(i \leftrightarrow j)) < 1$. Also, for all i and j, $\mathrm{mst}(i \leftrightarrow j) \leq k$, so $F(\mathrm{mst}(i \leftrightarrow j)) \leq F(k)$. Then if $F(k) < 1$, we know $F(\mathrm{mst}(i \leftrightarrow j)) < 1$, and thus either s_i is accessible from s_j or s_j is accessible from s_i, for all strategies s_i and s_j, $i \neq j$. Hence the Markov chain has a single ergodic set, and a unique invariant distribution, as per Theorem 2.1. If $F(k) = 1$, we know $F(\mathrm{mst}(i \leftrightarrow j)) = 1$, and thus neither of s_i and s_j is accessible from the other, for some pair of strategies s_i and s_j. Hence the Markov chain has multiple ergodic sets, and the invariant distribution is not unique, as per Theorem 2.1. ∎

This immediately implies the following useful corollary:

Corollary 4.4 *For a given NSP game with mprb* $= k$*, and for a given spread function* $f(x)$ *with associated cumulative spread function* $F(x)$ *such that* $F(k) < 1$*, the long run invariant distribution is unique.*

This is useful because the maximum pairwise bistability is much easier to compute than the maximum pathwise bistability: mprb requires only calculation over all pairs of strategies, while mptb requires calculation over all possible paths between strategies. Moreover, Theorem 2.1 follows trivially since mprb $\leq \frac{1}{2}$:

Corollary 4.5 *For a given NSP game, and for a given spread function* $f(x)$ *with associated cumulative spread function* $F(x)$ *such that* $F(\frac{1}{2}) < 1$*, the long run invariant distribution is unique.*

If we limit ourselves to coordination games, we can prove several other useful results. For a coordination game, we know $0 < \mathrm{dom}(s_i \mid s_j) < 1$ for all strategies s_i and s_j. This implies that $0 < \mathrm{mst}(i \to j) < 1$ for all i and j. Thus, if $F(x) < 1$ for all $x \in (0, 1)$, every strategy is accessible from every other strategy. This implies that the system is ergodic, and hence there exists a unique invariant distribution where *all* strategies have non-zero probability mass. This is the case for the uniform and KMR spread functions; however, for the KMR spread function, the probability masses may converge to 0 as the rate of exponential decrease k goes to infinity. Also, if $F(x) = 1$ for all $x \in (0, 1)$, no strategy is accessible from any other strategy, and hence the invariant distribution is not unique. This is the case for the MS spread function.

Finally, we consider some sufficient conditions for the model to give *strong* equilibrium selection results, i.e. when all of the probability mass is concentrated at a single strategy. Let the *minimum dominance score* of a strategy s_i be defined as $d_i = \min_{j \neq i} \mathrm{dom}(s_i \mid s_j)$.

Then state i is absorbing if and only if $F(d_i) = 1$. Now assume that the game and the spread function are such that the invariant distribution is unique, that is, $F(k) < 1$ for $k \leq \frac{1}{2}$ equal to the maximum pathwise bistability of the game. In this case, the model gives strong equilibrium selection results if exactly one equilibrium is absorbing. Let D denote the maximum d_i over all strategies s_i, and let d denote the second highest d_i. Then the model gives strong equilibrium selection results if $F(d) < 1$ and $F(D) = 1$. This includes any spread function $\delta(k)$ where $d \leq k < D$. Thus for any game with a unique maximum d_i, there exists a spread function $f(x)$ such that the invariant distribution places all probability mass on s_i.

We also note that $d + D \leq 1$, and thus if $D > \frac{1}{2}$, we have $d < \frac{1}{2} < D$. This is the case when there exists a risk-dominant strategy s_i, and in this case $D = d_i$. Thus the spread function $\delta(\frac{1}{2})$ always puts all probability mass on the risk-dominant strategy if one exists. Even if no risk-dominant strategy exists, we can still have strong equilibrium selection for some spread function $\delta(k)$ for $d \leq k < D \leq \frac{1}{2}$. For example, consider a game similar to Young's 3x3 game, but with $X \overset{7/8}{\Leftrightarrow} Y$, $Y \overset{3/5}{\Leftrightarrow} Z$, and $X \overset{3/7}{\Leftrightarrow} Z$, i.e. the relationship between X and Z has been reversed. In this case, X is absorbing if $F(\frac{1}{8}) = 1$, Y is absorbing if $F(\frac{2}{5}) = 1$, and Z is absorbing if $F(\frac{3}{7}) = 1$. Hence we will obtain strong equilibrium selection results for any $\delta(k)$ where $\frac{2}{5} \leq k < \frac{3}{7}$; in these cases strategy Z will be selected with probability 1, even though it is not risk-dominant.

4.6 A note on time to reach equilibrium

Though the stochastic stability models of Kandori et al (1993) and Young (1993) have been extremely influential in the evolutionary games literature, certain aspects of the models have caused some researchers to question the applicability of their results. In particular, both models depend on the occurrence of a large number of simultaneous mutations to move between the basins of attraction of different equilibria. Since the models assume that mutations are very rare ($\epsilon \approx 0$), and the necessary number of mutations k may be large, movements between equilibria occur with a very small probability ϵ^k. Hence, though the models *eventually* converge to the risk-dominant equilibrium in 2x2 symmetric coordination games, they may take a huge number of time steps to reach this equilibrium. As Ellison (1993) and Binmore & Samuelson (1997, 1999) have argued, for plausible values of the payoffs and the population size, the Kandori et al and Young models change basins so infrequently that *for all practical purposes* the result is determined by the initial condition. At least three solutions to this problem have been proposed. Ellison (1993) shows that convergence to the risk-dominant equilibrium is much faster when interaction is "local," i.e. players interact with, and choose best replies to, only a few close "neighbors" rather than the population as a whole. However, this solution is not applicable when the population is well mixed, and each player interacts with a large sample of the population, not only his close neighbors. Binmore, Samuelson, and Vaughan (1995) and Binmore & Samuelson (1997) show that incorporating noise into the selection process itself increases the speed of convergence to the invariant distribution, by offering the system more opportunities to climb out of basins of attraction that are not selected in the long run. They give an example where their model takes 5000 periods to converge, while the Kandori et al model requires an incredible 10^{72}

periods. Nevertheless, several thousand periods is still a long time until convergence to the invariant, and the waiting time still grows explosively as population size increases. Thus Binmore & Samuelson (1999) focus on the history-dependent convergence of the short-term dynamics (which they call the "long run") as opposed to the limiting distribution of the long-term dynamics (which they call the "ultralong run"). Since, they argue, the time needed to reach the "ultralong run" makes this irrelevant in many contexts, the "long run" results will be more relevant in applied work. We cannot dispute this criticism of the Kandori et al and Young models: if a model requires a huge number of periods to converge to the invariant distribution, this distribution may not be as relevant on the time scales under consideration. However, the "long run" results of Binmore & Samuelson (1999) have the disadvantage of being strongly dependent on initial conditions: thus even a very poor equilibrium may be selected if the system happens to start in that equilibrium's small basin of attraction.

Our solution is to allow large aggregate shocks, as opposed to models such as Kandori et al, where the impact of shocks is negligible. We find that models with large shocks have the advantage of converging much more quickly to the invariant distribution: thus the results of these models are both independent of initial conditions, and applicable on "reasonable" time scales, assuming that the shocks are sufficiently large. A simple way to think of this is that large shocks enable the system to move much more frequently between equilibria, resulting in rapid mixing of the Markov chain. To quantify this claim, we compute bounds on the speed of convergence for three games, and thus demonstrate that the uniform and $\delta(\frac{1}{2})$ spread functions (where large shocks are frequent) converge much more quickly than the KMR spread function (where large shocks are very infrequent). To do so, we rely on a well-known "minorization" result on the speed of convergence of Markov chains (Rosenthal, 1995), which we present here without proof.

Assume that the Markov chain Γ is finite (N states), aperiodic, and has a single ergodic set. Let $P = \{p_{ij}\}$ be the matrix of transition probabilities from state i to state j. Then for any initial probability distribution $\vec{x}_0$, we know that the sequence $\vec{x}_k = \vec{x}_0 P^k$ will converge to the invariant distribution $\vec{y}$ as k goes to infinity. To determine the speed of convergence, we can measure the *total variation distance* $V_k = \|\vec{x}_k - \vec{y}\| = \frac{1}{2}\sum_{i=1\ldots N}|x_{ki} - y_i|$. We know that V_k goes to zero as k goes to infinity, and the relative speeds of convergence of two chains can be measured by how fast V_k goes to zero for each. In particular, we choose an arbitrary constant $\epsilon = .005$, and compute the minimum number of periods k such that $V_k \leq \epsilon$ for each chain. To do so, we note that since the chain has a single ergodic set, there exists a state j_0 and a positive integer k_0 such that j_0 is accessible from every state i in k_0 steps. Then let $\beta = \sum_{j=1\ldots N} \min_{i=1\ldots N} p_{ij}^{k_0}$. Since $p_{ij_0}^{k_0} > 0$ for each state i, we know $\beta > 0$. Given these values of k_0 and β, Rosenthal (1995) proves that the total variation distance is bounded by $V_k \leq (1-\beta)^{\lfloor k/k_0 \rfloor}$.

For example, for the 2x2 coordination games given in Tables 3-4, we have $X \overset{4/5}{\Leftrightarrow} Y$, and thus $p_{XY} = \frac{\bar{F}(\frac{1}{5})}{2}$ and $p_{YX} = \frac{\bar{F}(\frac{4}{5})}{2}$. Also, for any spread function where the invariant distribution is unique, we know that Y is accessible from X in 1 step. Then $\beta = p_{XY} + p_{YX} = \frac{\bar{F}(\frac{1}{5})+\bar{F}(\frac{4}{5})}{2}$, and $V_k \leq (1-\beta)^k$. We now consider the uniform spread function, the $\delta(\frac{1}{2})$ spread function, and the KMR spread function with $\bar{F}(x) = \frac{e^{-100x} - e^{-100}}{1-e^{-100}} \approx e^{-100x}$. Then $\beta = \frac{1}{2}$ for the uniform and $\delta(\frac{1}{2})$ spread functions, while β for the KMR spread function is approximately

$\frac{e^{-20}}{2} \approx 10^{-9}$. Thus we know that the uniform and $\delta(\frac{1}{2})$ spread functions take 8 periods to guarantee $V_k \leq .005$, while the KMR spread function takes approximately 5×10^9 periods.

Next we consider the 3x3 coordination game of Young (1993), with payoffs given in Table 5. For this game, recall that $X \overset{7/8}{\Leftrightarrow} Y$, $X \overset{4/7}{\Leftrightarrow} Z$, and $Y \overset{3/5}{\Leftrightarrow} Z$. For simplicity, we assume a spread function with $F(\frac{3}{7}) < 1$; in this case, Z is accessible from X and Y in 1 step. Then $\beta = p_{YX} + p_{XZ} + p_{ZY} = \frac{\bar{F}(\frac{7}{8}) + \bar{F}(\frac{3}{7}) + \bar{F}(\frac{3}{5})}{3}$, and $V_k \leq (1-\beta)^k$. For the uniform spread function U, we compute $\beta = \frac{307}{840}$, and thus the total variation distance V_k is guaranteed to be less than .005 in 12 periods. For the $\delta(\frac{1}{2})$ spread function, $\beta = \frac{1}{3}$, and thus $V_k < .005$ in 14 periods. For the KMR spread function given above, we compute $\beta \approx \frac{e^{-300/7}}{3} \approx 8 \times 10^{-20}$, and thus $V_k < .005$ in approximately 7×10^{19} periods.

Finally, we consider the NSP game given in Table 6, assuming $\mu = 1$. For this game, we have $X \to Y \to Z \to W \to X$, $X \overset{1/2}{\Leftrightarrow} Z$, and $Y \overset{1/2}{\Leftrightarrow} W$. We first consider the case where $F(\frac{1}{2}) < 1$, allowing transitions between X and Z, and between Y and W. In this case, any strategy is accessible from any other strategy in 2 periods: for example, strategy X is not accessible from strategy Y in 1 period, but is accessible from strategy Y in two periods via either Z or W. The minimum transition probability for each of the four strategies is $2 \times \frac{\bar{F}(0)}{4} \times \frac{\bar{F}(\frac{1}{2})}{4} = \frac{\bar{F}(\frac{1}{2})}{8}$. Thus we calculate $\beta = \frac{\bar{F}(\frac{1}{2})}{2}$, and $V_k \leq (1-\beta)^{\lfloor k/2 \rfloor}$. For the uniform and $\delta(\frac{1}{2})$ spread functions, we have $\beta = \frac{1}{4}$, and thus $V_k < .005$ in 37 periods. For the KMR spread function, we have $\beta = \frac{e^{-50}}{2} \approx 10^{-22}$, and thus the total variation distance is guaranteed to be less than .005 in approximately 10^{23} periods. However, we can also derive a tighter bound by noting that every strategy is accessible from any other strategy in three periods, with minimum probability $\left(\frac{\bar{F}(0)}{4}\right)^3 = \frac{1}{64}$, regardless of the spread function. In other words, we can follow the cycle of dominance relations, $X \to Y \to Z \to W \to X$, until reaching any of the four strategies. Thus $\beta = \frac{4}{64} = \frac{1}{16}$, and $V_k \leq (1-\beta)^{\lfloor k/3 \rfloor}$. This gives us $V_k < .005$ in 247 periods.

Thus, for the game in Table 6, the cycle of dominance relations allows the KMR spread function to converge to the invariant distribution in only ≈ 250 periods. For the other two examples, however, the KMR function takes a very long time to converge, on the order of 5 billion periods for the 2x2 coordination game, and 70 billion billion periods for Young's 3x3 coordination game. This was assuming that the spread function $f(x)$ decreases as e^{-100x}; this value approximates a population of 22 individuals with a 1% chance of mutation. For a larger population, or smaller probability of mutation, the invariant distribution takes even longer to converge. For the uniform or $\delta(\frac{1}{2})$ spread functions (where large shocks are frequent), on the other hand, convergence was extremely fast in all examples, ranging from 8 to 37 periods. This demonstrates that models with large shocks converge quickly enough for the invariant distribution to be an accurate reflection of the behavior of the model on reasonable time scales.

5 The Finitely Repeated Prisoner's Dilemma

We now consider the Prisoner's Dilemma game, as given in Table 1, in more detail. The Prisoner's Dilemma is a widely used model of interactions between individuals with partially conflicting goals, where each player must choose between *cooperation* for mutual benefit, and *defection* for individual benefit (harming the other player). If both players cooperate, each receives a high payoff R as a *reward* for mutual cooperation. If both defect, each receives a low payoff P as a *punishment* for mutual defection. If one player defects while the other cooperates, the defector receives a very high payoff T as a *temptation* to defect, and the cooperator receives a very low *sucker* payoff S. For the Prisoner's Dilemma game in Table 1, strategy X corresponds to cooperation and strategy Y corresponds to defection; we assume the payoffs $T = 3$, $R = 2$, $P = 0$, and $S = -1$. Note that the Pareto-optimal result (X,X) is dominated by the strict Nash equilibrium (Y,Y). Thus, despite the fact that mutual cooperation is preferred to mutual defection, each player scores higher if he defects regardless of the opponent's choice, and hence mutual defection is the only rational outcome.

In an "Iterated Prisoner's Dilemma" (IPD), a Prisoner's Dilemma interaction is repeated between the same two players over a number of rounds. Based on past results, reciprocal altruism can develop, enabling mutual cooperation to become a rational option. The IPD game has sparked various theories of the evolution of cooperation based on reciprocity, and these models have been applied to fields ranging from economics to biology (Axelrod & Hamilton, 1981; Maynard Smith, 1982; Axelrod, 1984).

The Finitely Repeated Prisoner's Dilemma (FRPD) is a variant of the Iterated Prisoner's Dilemma in which the number of rounds M is given and known in advance by both players. We might expect that two rational strategies would be able to achieve cooperation in the FRPD; however, under the assumption of common knowledge of rationality, we can use an argument of backward induction to show that each player defects continually. This argument proceeds as follows: on round M, both players will defect, as in the one-shot Prisoner's Dilemma. Then on round $M - 1$, both players know that, whatever they do now, both will defect on the next round in any case. Since the player's move on round $M - 1$ will not influence the opponent to reciprocate cooperation on future moves, it is rational to defect on round $M - 1$. This argument can be applied to round $M - 2$ and each preceding move, in turn, and we reach the conclusion that two rational players will defect on every round. This argument was first put forth by Luce & Raiffa (1957), and Aumann (1995) proved that common knowledge of rationality implies the backward induction solution. But human players tend to treat the FRPD as a Prisoner's Dilemma interaction of indeterminate length until nearly the end of the game, reciprocating the other player's cooperation until the final few rounds. Thus we have a "paradox of rational choice" in which perfectly rational agents perform significantly worse than agents with imperfect rationality.

Traditional "solutions" to the Finitely Repeated Prisoner's Dilemma attempt to justify cooperation by weakening the assumptions of common knowledge of (perfect) rationality or of the structure of the game. The best known result along these lines is the "Gang of Four" paper (Kreps et al, 1982): Kreps et al explain cooperation by incomplete information about one or both players' options, motivation, or behavior. They prove an upper bound on the number of defections in the FRPD if players believe that one player may possibly (with probability ϵ) be committed to a "Tit for Tat" (reciprocating) strategy. Also, if both players are uncertain about the other's payoff function, they prove that both cooperative and non-cooperative sequential equilibria are possible. Similarly, Fudenberg & Maskin (1986) assume that one player may possibly be "crazy" (punishing defections that would otherwise be too near the end of the game to be deterred by credible threats), resulting in cooperative equilibria. Neyman (1999) shows that cooperation can be achieved in the FRPD if the number of stages is not common knowledge; even an exponentially small deviation is sufficient for cooperation. Neyman (1985) considers finite automata playing the FRPD, and shows that bounded complexity results in cooperation. Finally, Radner (1986) shows that if players "satisfice" rather than strictly optimizing (i.e. each player is satisfied if he gets ϵ-close to the Nash equilibrium payoff), then cooperation is possible.

Evolutionary models have also been applied to the Finitely Repeated Prisoner's Dilemma, but most of the results in the literature tend to reinforce the backward induction argument rather than suggesting how cooperation can emerge. In fact, assuming a simple invasion model based on the criteria of Maynard Smith (1982), we observe an evolutionary equivalent of the backward induction argument. First assume that we have a population of continual cooperators; this strategy can be invaded, and taken over, by a strategy that cooperates until move $M-1$ and defects on move M. The strategy that cooperates until move $M-1$, in turn, can be taken over by a strategy that cooperates until move $M-2$, and defects thereafter. This process continues, with each strategy being invaded by a strategy that starts defecting one round earlier, until the population consists entirely of continual defectors. Continual defection is evolutionarily stable, and thus the evolutionary process ends with a population of defectors.

In other evolutionary models, cooperation may be quasi-stable and persist for long periods of time, but in the long run, behavior eventually converges to defection. Nachbar (1992) assumes a polymorphic population (with all strategies initially present), and allows this population to evolve according to a deterministic, payoff-monotone selection dynamics. Again we observe a form of "evolutionary backward induction": each pure strategy in turn has a time interval where it achieves a large population share, then its population share decreases as it is replaced by the strategy which starts defecting one round earlier. This continues until the population consists entirely of strategies which defect on every round. Interestingly, Nachbar's simulations demonstrate that partial cooperation can flourish for extended periods of time. Since defectors perform poorly in early (mostly cooperative) stages, their population shares are small, and it takes a long time before they can recover a significant population share under the replicator dynamics. He gives an example of a 6-round FRPD where the strategy "cooperate in the first round and defect afterward" maintains a 99% population share for 300 periods of simulation, before defectors take over. Nevertheless, in the long run, defection is permanently triumphant. Cressman (1996) extends Nachbar's results, proving that "any initially polymorphic population evolves under the replicator dynamics to one where all strategy types remaining defect against each other at every stage."

However, some researchers have observed that a defecting population, though immune to isolated small mutations, might be invaded by a finite proportion of a cooperative strategy. Axelrod (1984) observes that, though defection is evolutionarily stable, the invasion of a small proportion of conditional cooperators might topple the defecting regime. Schuessler (1989) concludes from this that "a cyclical rise and decline of cooperativeness has to be expected if random perturbation occurs ... the gradual decline of cooperation holds true as a general tendency, but not for all periods of the evolutionary process." In fact, the conclusion we draw from this is much stronger: shocks resulting in a finite proportion of mutants will enable *cooperators* to dominate the population in the long run. This is very different from the results of Nachbar and Cressman, where cooperation is "quasi-stable" but defectors eventually triumph. In the next subsection, we prove that, for sufficiently long FRPD games, the population is dominated by a cooperative strategy *almost all of the time* in the long run invariant distribution. Moreover, this is true not only for models with very large shocks, but holds for any model that satisfies the LASH property (i.e. assumes a finite proportion of invaders).

5.1 Finitely Repeated Prisoner's Dilemma: results

To simplify our calculations, and present the paradox in its most common and essential form, we assume that once cooperation has failed (i.e. either player has defected), it cannot be resumed, and the rest of the game will consist of mutual defection. Since we have assumed $P = 0$, this is equivalent to an FPRD game that stops upon either player's defection. Alternatively, we can think of this as a standard no-noise FRPD, with all strategies chosen from the strategy space s_i, $i = 1 \ldots (M+1)$. A strategy s_i will cooperate on turn k if and only if both of the following are true: $k < i$, and its opponent has cooperated on turns 1 through $k-1$. Thus s_1 always defects, and s_{M+1} always reciprocates cooperation with cooperation.

We now calculate the total payoff w_{ij} to strategy s_i against strategy s_j, in an M-round FRPD game. There are three cases to consider. First, if $i = j$, both strategies will cooperate until round $i-1$ and defect thereafter, receiving total payoff $2(i-1)$. Second, if $i > j$, strategy s_i receives $j-1$ reward payoffs, followed by a sucker payoff. Thus its total payoff is $2(j-1)-1$. Third, if $i < j$, strategy s_i receives $i-1$ reward payoffs, followed by a temptation payoff. Thus its total payoff is $2(i-1)+3$. We can summarize this as $w_{ij} = 2\min(i-1, j-1) + 3(i < j) - (i > j)$, where ($i$ op j) equals 1 if the expression is true and 0 otherwise. Now we compute the evolutionary relationship between strategies s_i and s_j, assuming without loss of generality that $i > j$. If $i - j = 1$, we find $w_{jj} - w_{ij} = 1$ and $w_{ji} - w_{ii} = 1$, and thus $s_i \to s_j$. If $i-j > 1$, we find $w_{jj} - w_{ij} = 1$ and $w_{ii} - w_{ji} = 2(i-j-1)-1$, and thus $s_i \overset{m}{\Leftrightarrow} s_j$, where $m = \frac{1}{2(i-j-1)}$. Note that if $i \gg j$, we have $m \approx 0$, and thus mutants of type s_i can invade a population of type s_j even if they make up only a very small (but finite) proportion of the population. On the other hand, if shocks are negligible (as in the MS spread function), then s_i cannot invade. We also note that, for any spread function, strategy s_1 is accessible from every other strategy. This implies that there is a single ergodic set, and thus the long run invariant distribution is unique.

We now examine the invariant distribution $\vec{y}$, first for models with small shocks (the MS spread function) and then for models with large shocks (all other spread functions we have considered). For the MS spread function, the only possible transitions are from s_j to s_{j-1}, for all $j > 1$. Hence the number of rounds of cooperation decreases monotonically, until s_1 (the strategy that defects continually) dominates the population. s_1 is the unique absorbing state, and thus it receives all of the probability mass in the invariant distribution. For models with large shocks, on the other hand, we expect to observe two phenomena in the FRPD: "evolutionary backward induction" (s_j invaded by s_{j-1}) and "evolutionary forward progression" (s_j invaded by s_i, $i \gg j$). For a large number of rounds M, and for most strategies s_j, there are a large number of strategies s_i with $i \gg j$. Thus we would expect evolutionary forward progression to be much more common than evolutionary backward induction. As a result, we would expect the evolutionary process to be "pushed" toward continual cooperation rather than continual defection, resulting in domination of the population by strategies that cooperate until nearly the end of the game.

We now calculate the long run invariant distribution $\vec{y}_M$ for various game lengths M, initially assuming a uniform spread function. For each M, we record three values: the best strategy s_i (the strategy with the highest y_{Mi}), the proportion of the time the population is dominated by "cooperative" strategies (strategies that cooperate at least 90% of the game), and the proportion of the time the population is dominated by "uncooperative" strategies (strategies that cooperate at most 10% of the game). These results are given in Table 8.

From these results, we observe that defectors dominate the population for very short games, but for longer games, the population is dominated by strategies that cooperate until near the end of the game. For $M = 20$, less than 1% of dominating strategies were uncooperative, and more than 50% of dominating strategies were cooperative. For $M = 800$, less than 0.01% of dominating strategies were uncooperative, and 99% were cooperative. s_1, the strategy which always defects, was the most successful for $M \leq 3$, but for all $M > 3$, the most successful strategy was s_M, the strategy that cooperates through the next-to-last round.

Our empirical results suggest that the proportion of "cooperative" strategies goes to 1 as M goes to infinity. In fact, we can prove a much stronger result than this: for any constant k, $0 < k < 1$, the proportion of strategies cooperating through turn kM goes to 1 as $M \to \infty$. In other words, if we had chosen 99.999% cooperation instead of 90% as the threshold for a "cooperative" strategy, the population would still be dominated by cooperative strategies for large enough M. Moreover, this is true not only for the uniform spread function, but for any LASH spread function (i.e. all of the spread functions we have considered, except for MS). To show this, we first prove the following lemma:

Lemma 5.1 *For the Finitely Repeated Prisoner's Dilemma game,* $\lim_{k\to\infty} q_{i(i+k)} = 1$ *and* $\lim_{k\to\infty} q_{(i+k)i} = 0$ *for any spread function satisfying the LASH property.*

Proof For $k > 1$, we know $s_i \overset{m}{\Leftrightarrow} s_{i+k}$, where $m = 1 - \frac{1}{2(k-1)}$. Thus $\lim_{k\to\infty} q_{i(i+k)}$ $= \lim_{k\to\infty} \left(1 - F\left(\frac{1}{2(k-1)}\right)\right) = 1 - \lim_{k\to 0} F(k)$. Since the spread function is LASH, we know $\lim_{k\to 0} F(k) = 0$, and thus $\lim_{k\to\infty} q_{i(i+k)} = 1$. Similarly, $\lim_{k\to\infty} q_{(i+k)i}$ is equal to $\lim_{k\to\infty} \left(1 - F\left(1 - \frac{1}{2(k-1)}\right)\right) = 1 - \lim_{k\to 1} F(k)$. Since the spread function is LASH, we know $\lim_{k\to 1} F(k) = 1$, and thus $\lim_{k\to\infty} q_{(i+k)i} = 0$. ∎

Now consider an M-stage Finitely Repeated Prisoner's Dilemma game, with stage payoffs given in Table 1 and the strategy space defined as above. Let $\vec{y}_M$ be the unique invariant distribution for this game. Let $f(x)$ denote the spread function for this game, with associated cumulative spread function $F(x)$. We then state the main theorem as follows:

Theorem 5.2 *If the spread function $f(x)$ satisfies the LASH property, then for any constant k, $0 < k < 1$, $\sum_{i>kM} y_{Mi} \to 1$ and $\sum_{i<kM} y_{Mi} \to 0$ as the number of stages M goes to infinity.*

Proof For simplicity of notation, we assume that kM is not an integer. We call a strategy s_i "cooperative" if $kM < i \leq M+1$, and "uncooperative" if $1 \leq i < kM$. Let y_{high} denote the proportion of cooperative strategies in the invariant distribution, and y_{low} denote the proportion of uncooperative strategies. We will show that, for any fixed $\epsilon \approx 0$, $\lim_{M\to\infty} y_{low} \leq \epsilon$. It then follows trivially that $\lim_{M\to\infty} y_{low} = 0$, and thus $\lim_{M\to\infty} y_{high} = 1$.

First, let p_{low} denote the probability of switching to an uncooperative strategy, given that the current strategy is cooperative:

$$p_{low} = \frac{\sum_{i=\lceil kM\rceil}^{M+1} y_{Mi} \sum_{j=1}^{\lfloor kM\rfloor} p_{ij}}{\sum_{i=\lceil kM\rceil}^{M+1} y_{Mi}}$$

Similarly, let p_{high} denote the probability of switching to a cooperative strategy, given that the current strategy is uncooperative:

$$p_{high} = \frac{\sum_{i=1}^{\lfloor kM\rfloor} y_{Mi} \sum_{j=\lceil kM\rceil}^{M+1} p_{ij}}{\sum_{i=1}^{\lfloor kM\rfloor} y_{Mi}}$$

We know that $y_{low} = \frac{p_{low}}{p_{low}+p_{high}}$. From this, we can find an upper bound for y_{low} by computing an upper bound for p_{low} and a lower bound for p_{high}. To do this, we rely on the following observations: for $j > i$, p_{ij} is decreasing with i and increasing with j, and p_{ji} is decreasing with j and increasing with i. Thus p_{low}, the weighted average of $\sum_{j=1}^{\lfloor kM\rfloor} p_{ij}$ for all $i = \lceil kM\rceil \ldots (M+1)$, is bounded by:

$$p_{low} \leq \sum_{j=1}^{\lfloor kM\rfloor} p_{\lceil kM\rceil j} = \sum_{j=1}^{\lfloor kM\rfloor} \frac{q_{\lceil kM\rceil j}}{M+1}$$

Similarly, p_{high}, the weighted average of $\sum_{j=\lceil kM\rceil}^{M+1} p_{ij}$ for all $i = 1 \ldots \lfloor kM\rfloor$, is bounded by:

$$p_{high} \geq \sum_{j=\lceil kM\rceil}^{M+1} p_{\lfloor kM\rfloor j} = \sum_{j=\lceil kM\rceil}^{M+1} \frac{q_{\lfloor kM\rfloor j}}{M+1}$$

We now define $Z = \epsilon(1-k)M$. Then we can compute the following bounds for p_{low} and p_{high}:

$$p_{low} \leq \sum_{j=1}^{\lfloor kM\rfloor - \lceil Z\rceil} \frac{q_{\lceil kM\rceil j}}{M+1} + \sum_{j=\lfloor kM\rfloor - \lfloor Z\rfloor}^{\lfloor kM\rfloor} \frac{q_{\lceil kM\rceil j}}{M+1} \leq \sum_{j=1}^{\lfloor kM\rfloor - \lceil Z\rceil} \frac{q_{\lceil kM\rceil (\lfloor kM\rfloor - \lceil Z\rceil)}}{M+1} + \sum_{j=\lfloor kM\rfloor - \lfloor Z\rfloor}^{\lfloor kM\rfloor} \frac{q_{\lceil kM\rceil \lfloor kM\rfloor}}{M+1}$$

$$= (\lfloor kM\rfloor - \lceil Z\rceil)\frac{q_{\lceil kM\rceil (\lfloor kM\rfloor - \lceil Z\rceil)}}{M+1} + \frac{\lceil Z\rceil}{M+1}$$

$$= (\lfloor kM\rfloor - \lceil \epsilon(1-k)M\rceil)\frac{q_{\lceil kM\rceil (\lfloor kM\rfloor - \lceil \epsilon(1-k)M\rceil)}}{M+1} + \frac{\lceil \epsilon(1-k)M\rceil}{M+1}.$$

$$p_{high} \geq \sum_{j=\lceil kM \rceil}^{\lceil kM \rceil + \lfloor Z \rfloor} \frac{q_{\lfloor kM \rfloor j}}{M+1} + \sum_{j=\lceil kM \rceil + \lceil Z \rceil}^{M+1} \frac{q_{\lfloor kM \rfloor j}}{M+1} \geq \sum_{j=\lceil kM \rceil}^{\lceil kM \rceil + \lfloor Z \rfloor} \frac{q_{\lfloor kM \rfloor \lceil kM \rceil}}{M+1} + \sum_{j=\lceil kM \rceil + \lceil Z \rceil}^{M+1} \frac{q_{\lfloor kM \rfloor (\lceil kM \rceil + \lceil Z \rceil)}}{M+1}$$

$$= (M - \lfloor kM \rfloor - \lfloor Z \rfloor) \frac{q_{\lfloor kM \rfloor (\lceil kM \rceil + \lceil Z \rceil)}}{M+1}$$

$$= (M - \lfloor kM \rfloor - \lfloor \epsilon(1-k)M \rfloor) \frac{q_{\lfloor kM \rfloor (\lceil kM \rceil + \lceil \epsilon(1-k)M \rceil)}}{M+1}.$$

We now apply Lemma 5.1, giving us $q_{\lceil kM \rceil (\lfloor kM \rfloor - \lceil \epsilon(1-k)M \rceil)} \to 0$ and $q_{\lfloor kM \rfloor (\lceil kM \rceil + \lceil \epsilon(1-k)M \rceil)} \to 1$ as $M \to \infty$. Thus we have:

$$\lim_{M \to \infty} p_{low} \leq \lim_{M \to \infty} \frac{\lceil \epsilon(1-k)M \rceil}{M+1} = \epsilon(1-k).$$

$$\lim_{M \to \infty} p_{high} \geq \lim_{M \to \infty} \frac{M - \lfloor kM \rfloor - \lfloor \epsilon(1-k)M \rfloor}{M+1} = (1-\epsilon)(1-k).$$

$$\lim_{M \to \infty} y_{low} = \lim_{M \to \infty} \frac{p_{low}}{p_{low} + p_{high}} \leq \frac{\epsilon(1-k)}{\epsilon(1-k) + (1-\epsilon)(1-k)} = \epsilon.$$

Thus for any ϵ, $\lim_{M \to \infty} y_{low} \leq \epsilon$ and $\lim_{M \to \infty} y_{high} = 1 - y_{low} \geq 1 - \epsilon$. Therefore y_{low} goes to 0, and y_{high} goes to 1, as M goes to infinity. ∎

Thus, if the game length is sufficiently long, cooperative strategies dominate the population for any LASH spread function. This does not include the MS spread function, of course, but it does include the KMR spread function for any *fixed* rate of decrease k. Note that this result for KMR is dependent on the order of limits: if we take $k \to \infty$ before $M \to \infty$, then the KMR spread function does not obey the LASH property, and gives results identical to MS (i.e. convergence to continual defection). It should also be noted that, for models where $F(x) = 1$ except for $x \approx 0$, very large M may be required for cooperation to succeed. To demonstrate this, we compute (for a variety of spread functions) the minimum M such that an invasion of s_1 by s_{M+1} is successful with probability $q_{1(M+1)} > \frac{1}{2}$. Since $s_1 \overset{m}{\Leftrightarrow} s_{M+1}$ with $m = 1 - \frac{1}{2(M-1)}$, we must find the minimum M such that $\bar{F}\left(\frac{1}{2(M-1)}\right) > \frac{1}{2}$. For the uniform or $\delta(\frac{1}{2})$ spread functions, $\bar{F}(x) > \frac{1}{2}$ for $x < \frac{1}{2}$, and thus we have $M > 2$. On the other hand, consider the interval-uniform spread function $U(0, \epsilon)$ for a small, fixed $\epsilon \approx 0$. In this case, we have $\bar{F}(x) > \frac{1}{2}$ for $x < \frac{\epsilon}{2}$, and thus we have $M > 1 + \frac{1}{\epsilon}$. For example, for $\epsilon = 0.001$, we obtain $M > 1001$. Similarly, for the delta spread function $\delta(\epsilon)$, we obtain $M > 1 + \frac{1}{2\epsilon}$, and thus $M > 501$ for $\epsilon = 0.001$. Finally, for the KMR spread function $\bar{F}(x) \approx e^{-kx}$, we have $M > 1 + \frac{k}{2 \ln 2}$; for $k = 1000$, we have $M > 722$. Thus for some spread functions, the game length must be very long for cooperation to be rational. Nevertheless, since cooperation is rational for long enough games, the evolutionary FRPD paradox is resolved.

6 Games with stable polymorphisms

Up to this point, we have considered only games which satisfy the "no stable polymorphisms" (NSP) property: for any pair of strategies s_i and s_j, either $w_{ii} > w_{ji}$ or $w_{jj} > w_{ij}$, or both. Games with this property are "nice" in the sense that, given an initially polymorphic population consisting of two strategies, the population will evolve under the replicator dynamics (or any other payoff-monotone selection dynamics) to a state where only a single strategy is present. Thus, under the assumptions that mutations are isolated (i.e. the population is given sufficient time to evolve to a short-term stable state between mutations) and that only a single type of mutant can invade at a time, we can consider the long-term dynamics as a series of transitions between monomorphic populations (populations consisting only of a single strategy type). This enables a simple Markov chain representation, where the number of states is equal to the number of pure strategies, and transition probabilities are independent of the underlying short-term dynamics.

For games which do not satisfy the NSP property, however, it may be possible for multiple strategies to coexist in a stable polymorphic population. Consider an interaction between two strategies s_i and s_j, for which the NSP property is *not* satisfied, i.e. $w_{ij} \geq w_{jj}$ and $w_{ji} \geq w_{ii}$. There are four cases to consider. First, if $w_{ij} > w_{jj}$ and $w_{ji} > w_{ii}$, then no matter what the initial proportions of strategies s_i and s_j, the two strategies reach a *stable equilibrium* where the proportion of strategy s_i is given by:

$$p = \frac{w_{ij} - w_{jj}}{w_{ij} - w_{jj} + w_{ji} - w_{ii}}$$

In this case, we write $s_i \overset{p}{\leftrightarrow} s_j$, or equivalently $s_j \overset{1-p}{\leftrightarrow} s_i$.

If $w_{ij} > w_{jj}$ and $w_{ji} = w_{ii}$, any initial proportion of strategy s_i will take over the combined population, reducing the population share of strategy s_j to zero. However, we are not guaranteed that strategy s_j will be completely wiped out. As Neill (2003) shows, a small number of s_j strategists may be able to survive, under the more realistic assumptions that the population is very large but technically finite, and that players face all members of the population *excluding* themselves. This can make a huge difference if the population is then invaded by a third strategy, possibly allowing s_j to reemerge. According to Neill's *large population ESS*, strategy s_i will wipe out strategy s_j if $w_{ij} > w_{ji}$. In this case, s_i dominates s_j, and we write $s_j \rightarrow s_i$ as above. If $w_{ij} \leq w_{ji}$, a finite number $M = \frac{w_{ji}-w_{jj}}{w_{ij}-w_{jj}}$ of s_j strategists can survive. In this case, we say that s_i "quasi-dominates" s_j, and write $s_j \overset{-}{\rightarrow} s_i$. Under the simpler assumption that players face all members of the population *including* themselves, s_i wipes out s_j in any case, and we can treat domination and quasi-domination as identical. Conversely, if $w_{ji} > w_{ii}$ and $w_{ij} = w_{jj}$, then s_j dominates s_i if $w_{ji} > w_{ij}$, and quasi-dominates s_i otherwise.

Finally, if $w_{ij} = w_{jj}$ and $w_{ji} = w_{ii}$, then (assuming that the population size is infinite) both strategies achieve identical payoffs in the combined population.[4] We denote these interactions by $s_i \overset{-}{\leftrightarrow} s_j$. In this case, under a payoff-monotone, deterministic selection dynamics, the population shares of strategies s_i and s_j remain constant (i.e. the proportion

[4]Neill (2003) deals with the case of a large but finite population, which we do not discuss here.

of invaders remains at its initial value as determined by the spread function). However, we generally assume that *drift* occurs: the population shares of s_i and s_j are affected by shocks which are vanishingly small compared to the selective pressures. These tiny perturbations can usually be ignored, but must be considered in cases when there is no selective pressure. As discussed below, we make certain assumptions on the type of drift which allow us to simplify our calculations; see Binmore & Samuelson (1999) for a more detailed discussion of the impact of drift on the short-term selection dynamics.

In all of these cases (except for dominance), the result of an evolutionary interaction between two strategies may be a polymorphic population in which the two strategies coexist. The proportions of the two strategies in the population may be fixed (as in stable equilibrium) or subject to drift and other model-dependent parameters. A further complication is introduced on the next time step of the long-term dynamics, when another mutant strategy attempts to invade the polymorphic population. The evolutionary result of an interaction between three or more strategies is much more complex to analyze: in addition to stable (monomorphic or polymorphic) states, a variety of other phenomena can occur, including neutral oscillations, heteroclinic cycles, and chaotic attractors (Nowak & Sigmund, 1989). Moreover, the result may be dependent on the initial proportions of each strategy, and since the population shares may exhibit significant (cyclic or chaotic) fluctuations in multi-strategy interactions, these conditions become impossible to predict or even approximate. Additionally, the number of possible "stable states" of the short-term dynamics grows exponentially with the number of strategies considered, making the model intractable for all but the simplest non-NSP games.

Thus, it is clear that some simplifying assumptions must be made if we wish to apply this model to more general classes of games. In particular, we can imagine some process by which a single strategy type is chosen from a polymorphic population, and assume that this process is carried out whenever the short-term dynamics evolves to a polymorphic stable state. To examine some reasonable processes of "single strategy selection," we return to the "imaginary play" scenario discussed above. We assume a boundedly rational player with very limited resources, who can only maintain a single strategy (his current "favorite") in memory. Thus he changes strategies by a process of "discovery," selecting a new strategy randomly from the strategy space, and choosing between the new strategy and his current favorite. As above, since switching is costless, the current favorite is given no bonus (or penalty) for being established. In the case of a stable polymorphism, one simple solution would be for the player to always choose the risk-dominant strategy. In this case, s_i is selected when $w_{ij} - w_{jj} > w_{ji} - w_{ii}$, or in other words, $s_i \overset{p}{\leftrightarrow} s_j$ with $p > \frac{1}{2}$. This deterministic choice of the risk-dominant strategy has parallels in evolutionary and adaptive scenarios: it corresponds to populations whose choices are determined by *majority rule*, whether by voting or by any conflict where the majority is able to triumph. Alternatively, the boundedly rational player can follow a stochastic decision rule, where he chooses the risk-dominant strategy with probability equal to the level of risk-dominance: thus if $s_i \overset{p}{\leftrightarrow} s_j$, strategy s_i would be chosen with probability p. The stochastic choice method also has parallels in evolutionary and adaptive scenarios: it corresponds to picking a single individual out of the combined population at random, and carrying his strategy on to the next time step. This may occur when the choice of a population is determined by the rule of a dictator, or if the new population results from the offspring of a small number of migrant individuals (i.e. the well-known "founder effect" in evolutionary biology). In either of these cases, we assume for simplicity that a dominant or quasi-dominant strategy is always selected, and in the case of $s_i \stackrel{-}{\leftrightarrow} s_j$, that each of the two strategies is selected with probability $\frac{1}{2}$. The first assumption is equivalent to assuming that the random fluctuations due to drift will eventually wipe out any small (measure zero) proportion of mutants. The second assumption is equivalent to assuming that the drift has a directed component favoring one strategy or the other, and that the direction of drift is determined randomly.

More generally, we can define a monotonically increasing function $G(x)$ on the interval [0,1], such that $G(0) = 0$, $G(1) = 1$, and $G(1 - x) = 1 - G(x)$. This function gives the probability that a strategy is selected (at the end of a time step) if it makes up proportion x of a stable polymorphic population. Note that, since a two-strategy stable equilibrium does not depend on the initial proportions of the two strategies, the function $G(x)$ is always "fair," giving no advantage to the previously established strategy. The deterministic rule above corresponds to a function $G(x)$ which equals 0 for $x < \frac{1}{2}$, $\frac{1}{2}$ for $x = \frac{1}{2}$, and 1 for $x > \frac{1}{2}$. The stochastic rule above corresponds to a function $G(x) = x$. Under our simplifying assumption of single strategy selection, we can extend many of the general results in Section 4. For example, a system with the $\delta(\frac{1}{2})$ spread function and deterministic single strategy selection will always strongly select the risk-dominant strategy if one exists. The notion of maximum pathwise bistability can also be extended to non-NSP games: in this case, we calculate shock thresholds for all paths, excluding paths that are "blocked" by having some $s_{(i)} \stackrel{p}{\leftrightarrow} s_{(i+1)}$ such that $G(p) = 1$. Given the maximum pathwise bistability defined in this way, Theorem 4.3 holds for all games; moreover, we still have mptb $\leq \frac{1}{2}$, and thus Corollary 4.5 holds for all games.

6.1 Games with stable polymorphisms: results

We now apply our model to examine various extensive-form games common in the literature on forward and backward induction; typically these games do not satisfy the NSP property. A great deal of work has been done on applying evolutionary and adaptive models to extensive-form games: Fudenberg & Levine (1998, Ch. 6-7) present an excellent overview. We focus here on the large-population models of Nöldeke & Samuelson (1993), Gale et al (1995), Cressman & Schlag (1998), and Hart (2002). These approaches tend to treat extensive-form games differently than their strategic-form equivalents, assuming a distinct population of players at each node of the extensive-form game. Moreover, Nöldeke and Samuelson's model, as well as the fictitious play models discussed by Fudenberg and Levine, consider players characterized both by actions and by conjectures about the composition of the population at (reached and unreached) nodes. Interestingly, most of these models do not lend complete support to the backward induction argument: except for special classes of games, we may have evolutionarily stable equilibria other than the ones selected by backward induction. For Hart's model, on the other hand, the backward induction equilibrium is the unique evolutionarily stable outcome, assuming large populations and a low mutation rate.

Rather than assuming a distinct population at each decision node, we take the approach originated by Selten (1980), considering the "symmetrized normal form" in which players condition their strategy on their role (for example, "play A as player 1, or b as player 2"). In this model, players' expected payoffs are computed by assuming equal probabilities of being assigned to either role. For example, consider the simple "2-legged centipede" game shown in Figure 1; we compare the relative performance of the four role-dependent strategies Cc, Cd, Dc, and Dd. The centipede game can be thought of as a "gift-giving" game similar to the FRPD, except that players alternate turns rather than choosing simultaneously. The (asymmetric) normal form is given in Table 9, and the symmetrized normal form in Table 10.

For this game, we compute $Dc \to Cc \to Cd \to Dd$. We also have $Dc \stackrel{-}{\to} Cd$, $Dc \stackrel{-}{\leftrightarrow} Dd$, and $Cc \stackrel{1/2}{\Leftrightarrow} Dd$. Assuming a uniform or $\delta(\frac{1}{2})$ spread function, we obtain $[Cc\ Cd\ Dc\ Dd] = [\frac{7}{35}\ \frac{10}{35}\ \frac{3}{35}\ \frac{15}{35}]$, and thus mutual defection is selected in the 2-legged centipede game. In this case, our model agrees with the backward induction solution. For longer centipede games, we would expect cooperative strategies to be more successful, as in the Finitely Repeated Prisoner's Dilemma. For a 4-legged centipede game, as given in Figure 2, we find that the population is dominated 81% of the time by strategies which cooperate at least once in both roles, assuming the uniform spread function. Assuming the $\delta(\frac{1}{2})$ spread function, this fraction increases to 87%. For a 40-legged centipede game, we find that the population is dominated by cooperative strategies almost all of the time: for the $\delta(\frac{1}{2})$ spread function, strategies that cooperated at least 18 rounds (out of 20) dominated the population 97% of the time. Similarly, for the uniform spread function, strategies that cooperated at least 18 rounds dominated the population 80% of the time. Thus our results for the centipede game confirm our FRPD results: defection is dominant for very short game lengths, but cooperation dominates for longer games.

Next we consider a simplified version of the "ultimatum" game (Gale et al, 1995). In this game, the first player (the "proposer") must choose whether to make the "fair" offer A or the "unfair" offer B. Then the second player (the "responder") must choose whether or not to accept the offer; if the offer is accepted (choice y), both players profit, but if the offer is rejected (choice n), neither player profits. We make the further assumption that the second player always accepts a fair offer (and this is common knowledge), resulting in the game in Figure 3. The unique backward induction equilibrium is By (proposer makes the unfair offer, and responder accepts the offer); however, An (proposer makes the fair offer, responder rejects unfair offers) is also a (weak) Nash equilibrium of the game. Gale et al (1995) assume separate populations of proposers and responders; they show that which equilibrium is selected depends on the relative rates at which these populations evolve. If proposers learn sufficiently fast compared to responders, play can converge to the weak Nash equilibrium. Though this equilibrium is not asymptotically stable, Binmore & Samuelson (1999) have shown that it can become asymptotically stable in games with drift. We again consider the symmetrized normal form, and compare the relative performance of the four role-dependent strategies Ay, An, By, and Bn. For this game, we calculate $Bn \to An$, $Bn \to By$, and $Ay \to By$. We also have $Ay \stackrel{-}{\leftrightarrow} An$, $Ay \stackrel{3/4}{\leftrightarrow} Bn$, and $An \stackrel{1/2}{\Leftrightarrow} By$. Based on these relations, we would expect By to have a higher proportion of the probability mass than An, since By dominates Ay, and An does not. For a fair spread function with stochastic single strategy selection, we find the invariant distribution $[Ay\ An\ By\ Bn] = [.1000\ .3364\ .5545\ .0091]$. Similarly, for a fair spread function with deterministic single strategy selection, we find $[Ay\ An\ By\ Bn] = [\frac{1}{9}\ \frac{1}{3}\ \frac{5}{9}\ 0]$. Thus our model selects the backward induction equilibrium in the ultimatum game, giving it approximately 55% of the probability mass; the weak Nash equilibrium also dominates the population a significant fraction of the time, while the other two strategies are rarely dominant.

Finally, we examine a more complicated game: the "Dalek" game of Kohlberg & Mertens (1986). This game is given in Figure 4; we initially assume $\mu = 3$. In this game, if player 1 chooses T (the "outside option"), then both players receive 2 points; if he chooses M or B then player 2 must guess which letter he has chosen, and both players receive 3 points or 1 point if he correctly guesses m or b respectively. Both Mm and Tb are Nash equilibria, but Kohlberg & Mertens use an argument of forward induction to conclude that Mm is the only "strategically stable" equilibrium. This argument proceeds as follows: player 1 will never choose B, since no matter what player 2 chooses, player 1 would have been better off choosing T. This implies that player 2 should always guess m, and thus player 1 should always choose M rather than T. We again consider the symmetrized normal form, and compare the performance of the six role-dependent strategies. We do not list all of the 15 evolutionary relations between these strategies, but we note that the forward induction equilibrium is risk-dominant, and it thus receives all of the probability mass for the $\delta(\frac{1}{2})$ spread function with deterministic single strategy selection. It also receives the majority (57%) of the probability mass for the uniform spread function (with stochastic single strategy selection), while the outside option receives 37%. On the other hand, consider the same game with $\mu = \frac{1}{2}$: player 2 receives only half a point for correctly guessing m. In this case, the forward induction solution is no longer risk-dominant, since $Tb \overset{3/7}{\Leftrightarrow} Mm$. Thus we find that the outside option receives 80% of the probability mass, and Mm only 20%, for the $\delta(\frac{1}{2})$ spread function; similarly, the outside option receives 54% of the probability mass, and Mm only 39%, for the uniform spread function. This supports the observation of Binmore & Samuelson (1999) that the forward induction criterion has little predictive power: in games where the forward induction solution Mm is risk-dominated by the outside option Tb, we would not expect Mm to be selected.

6.2 More games with stable polymorphisms

At this point, we revisit the simplifying assumption of "single strategy selection." To what extent is this assumption necessary, or reasonable? As we have argued above, simplifications are necessary in the general case, since multi-strategy interactions can be both extremely complex and unpredictable. On the other hand, for some simple non-NSP games, the evolutionary results can be easily computed without our assumption. For example, consider the symmetric "hawk-dove" game given in Table 11. For this game, we calculate $X \overset{1/2}{\leftrightarrow} Y$. Thus, given any mixture of strategies X and Y, the short-term selection dynamics always converges to a stable state where the population consists of $\frac{1}{2}X + \frac{1}{2}Y$. Hence the interaction is sufficiently simple that single strategy selection is unnecessary; moreover, through the invariant converges to $[X\ Y] = [\frac{1}{2}\ \frac{1}{2}]$ in any case, it makes more sense to think of the population as always being polymorphic rather than being dominated by each strategy half the time, and thus our analysis is more accurate without the simplifying assumption.

This brings up the second half of our initial question: to what extent is single strategy selection *reasonable*? For cases of drift or quasi-dominance, it is very reasonable to assume that (given sufficient time) a monomorphic population will result; this can be guaranteed with some simple assumptions on the type of drift, as discussed above. For stable equilibria, on the other hand, selective pressures prevent the population shares from varying significantly from their equilibrium levels, and hence it is very unlikely that a monomorphic population will result. Moreover, each strategy in the polymorphic population scores higher than if it was the only strategy type present: in other words, single strategy selection *weakens* the surviving type, making it more susceptible to future invasions. This may significantly reduce the long-term success of the strategies with stable equilibrium relations, and hence benefit other strategies which do not form stable equilibria. For an extreme example, consider the "2-stable" game in Table 12: for this game, we calculate $X \overset{2/3}{\Leftrightarrow} Z$, $Y \overset{2/3}{\Leftrightarrow} Z$, and $X \overset{1/2}{\leftrightarrow} Y$. Thus under the assumption of single strategy selection, it is clear that Z receives the majority of the probability mass: for a uniform spread function, we calculate $[X\ Y\ Z] = [\frac{1}{4}\ \frac{1}{4}\ \frac{1}{2}]$. However, we find that the stable combination of X and Y is very resistant to invasion by Z: since X and Y each score 500 against $\frac{1}{2}X + \frac{1}{2}Y$, we find $\frac{1}{2}X + \frac{1}{2}Y \overset{1/251}{\Leftrightarrow} Z$. Thus, for a more accurate model of this game, we can consider a four-state Markov chain, where the states are X, Y, Z, and $(X+Y)$. Whenever X is selected to invade Y or vice-versa, the state transitions to $(X+Y)$. State Z cannot transition to $(X+Y)$, since this would require multiple mutations. Finally, the only transition from $(X+Y)$ is to Z, but this occurs rarely (with probability $\frac{1}{753}$). Given this revised transition matrix, we calculate $[X\ Y\ Z\ (X+Y)]$ = [.0020 .0020 .0098 .9862] for the uniform spread function: thus the population is almost always dominated by the polymorphic combination of X and Y, and not by strategy Z.

Thus the simplifying assumption of single strategy selection, when applied to certain games containing stable equilibria, may dramatically alter the long run invariant distribution of strategies. Because of this, a better idea is to treat the stable equilibrium as a separate population state and to include this in the transition matrix. For some games, the complexities of multi-strategy interaction make this difficult or impossible, but for other games this can be done relatively easily. For example, we consider the asymmetric version of the hawk-dove game, comparing the four role-dependent choices Xx, Xy, Yx, and Yy. For this game, we calculate $Xx, Yy \rightarrow Xy, Yx$. Also, we have $Xx \overset{1/2}{\leftrightarrow} Yy$ and $Xy \overset{1/2}{\Leftrightarrow} Yx$. Thus, with the simplifying assumption of single strategy selection, we obtain $[Xx\ Xy\ Yx\ Yy] = [0\ \frac{1}{2}\ \frac{1}{2}\ 0]$. In other words, the two unconditional strategies "play X" and "play Y" are dominated by the two strategies which choose X or Y conditioned on their role. What happens if we do not assume single strategy selection? In this case, we must consider the stable equilibrium of $\frac{1}{2}Xx + \frac{1}{2}Yy$, and compute its evolutionary relations to Xy and Yx. We find that Xy achieves a payoff of $\frac{1}{2}$ against a 50/50 mixture of Xx and Yy, identical to the payoffs of Xx and Yy against this mixture. However, Xy achieves a self-payoff of 1, while Xx and Yy only score $\frac{1}{2}$ against Xy. Thus Xy dominates the mixture of Xx and Yy; using identical reasoning, Yx also dominates the mixture of Xx and Yy. Hence the role-dependent strategies Xy and Yx still receive all of the probability mass, and the invariant distribution is unchanged. In this case, we obtain identical results with or without single strategy selection, and thus our assumption is both reasonable and useful.

Finally, we revisit the ultimatum game (Figure 3, above), without using the assumption of single strategy selection. From the above analysis, we know $Bn \rightarrow An$, $Bn \rightarrow By$, and $Ay \rightarrow By$. We also have $Ay \overset{-}{\leftrightarrow} An$, $Ay \overset{3/4}{\leftrightarrow} Bn$, and $An \overset{1/2}{\Leftrightarrow} By$. Thus we have one stable equilibrium to consider: the combination of $\frac{3}{4}Ay + \frac{1}{4}Bn$. We first consider an invasion of the polymorphic population by the backward induction equilibrium strategy By. We find that By has a self-payoff of 2, and also scores 2 against the mixture of Ay and Bn. Also, Ay and Bn each score $\frac{3}{2}$ against By, and $\frac{15}{8}$ against the mixture of Ay and Bn. Thus it is clear that Bn dominates the polymorphic population. Next, we consider an invasion of the polymorphic population by the weak Nash equilibrium An. In this case, strategy Ay scores higher against the An mutants than strategy Bn does, so the relative proportions of Ay and Bn will not remain constant as they were in the previous case; we must consider each strategy's payoff against all three strategy types in the combined population. Strategy Ay scores 2, 2, and 1.5 points against strategies Ay, An, and Bn respectively; it weakly dominates An, which scores 2, 2, and 1 against the three strategies, but does not weakly dominate Bn, which scores 2.5, 1, and 0. For an initial population of $\frac{3}{4}(1-k)Ay + \frac{1}{4}(1-k)Bn + (k)An$, we find that Ay outscores both Bn and An, increasing its population share toward 75%, while Bn and/or An decrease toward zero. If An reaches zero first, the population shares of Ay and Bn restabilize to $\frac{3}{4}$ and $\frac{1}{4}$ respectively, and the polymorphic population remains. If Bn reaches zero first, we have a mixture of Ay and An; thus drift can occur, and we assume that one strategy will die off, and one survive, with equal probabilities. However, whether An or Bn reaches zero first is strongly dependent on the short-term selection dynamics; different dynamics may give very different results from an identical starting population.

To see this, let x_1, x_2 and x_3 be the population shares of Ay, An, and Bn respectively. Then, assuming that the dynamics is payoff-monotone, we know that An has a higher growth rate than Bn when $x_1 < \frac{2}{3}$, and a lower growth rate when $x_1 > \frac{2}{3}$. Thus, in the best case for Bn, x_2 and x_3 decrease at approximately the same rate if $x_1 \leq \frac{2}{3}$, and otherwise, x_2 decreases rapidly while x_3 grows at approximately the same rate as x_1. For this case, we find that x_2 always reaches zero before x_3, and thus the polymorphic population resists the invasion by An. In the best case for An, x_2 and x_3 decrease at approximately the same rate when $x_1 \geq \frac{2}{3}$, and otherwise, x_3 decreases rapidly while x_2 grows at approximately the same rate as x_1. For this case, x_3 reaches zero first when the initial value of x_2 is at least half the initial value of x_1, which will occur when the initial proportion of An is at least $\frac{3}{11}$. Thus the probability that An successfully invades the polymorphic population ranges from 0 to $\frac{1}{4} \times \frac{1}{2} \times \bar{F}(\frac{3}{11})$, depending on the short-term selection dynamics. For the replicator dynamics, $\dot{x}_i = x_i(u_i - \bar{u})$, and we have a system of three nonlinear differential equations for the population shares:

$$\dot{x}_1 = x_1 \left(\frac{3}{2} - 2x_1 \right) x_3$$

$$\dot{x}_2 = x_2 \left(1 - 2x_1\right) x_3$$

$$\dot{x}_3 = x_3 \left(\frac{1}{2}x_1 - x_2 - 2x_1x_3 \right)$$

We can, without changing the paths for $x_3 > 0$, rescale the time axis by $2x_3^{-1}$, giving us:

$$\dot{x}_1 = x_1(3 - 4x_1)$$

$$\dot{x}_2 = x_2(2 - 4x_1)$$

$$\dot{x}_3 = x_1 - 2x_2 - 4x_1x_3$$

Then we can solve the first differential equation for x_1, then substitute the result into the second equation and solve for x_2. Finally, we know $x_3 = 1 - x_1 - x_2$. Assuming that the initial proportion of An mutants is k, this gives us:

$$x_1(t) = \frac{\frac{3}{4}(1-k)}{ke^{-3t} + (1-k)}$$

$$x_2(t) = \frac{ke^{-t}}{ke^{-3t} + (1-k)}$$

$$x_3(t) = \frac{\frac{1}{4}(1-k) + k(e^{-3t} - e^{-t})}{ke^{-3t} + (1-k)}$$

Since the minimum value of $e^{-3t} - e^{-t}$ is $-\frac{2\sqrt{3}}{9}$, we find that x_3 remains positive when $\frac{1}{4}(1-k) - \frac{2k\sqrt{3}}{9} > 0$. This equation is satisfied when $k < \frac{24\sqrt{3}-27}{37} \approx .394$. Thus An invades the polymorphic population when $k > .394$, and the probability of invasion is $\frac{1}{4} \times \frac{1}{2} \times \bar{F}(.394)$. Given this result, we can compute transition probabilities for a five-state Markov chain with states Ay, An, By, Bn, and $(Ay + Bn)$. However, Bn is a transient state, so we can exclude it from our computation. Assuming a uniform spread function and the replicator dynamics, we obtain the transition matrix given in Table 13. From this, we calculate the invariant distribution $[Ay\ An\ By\ (Ay + Bn)] = [.0709\ .3275\ .5575\ .0441]$. We note that this is very similar to the invariant distributions obtained under the assumption of single strategy selection. The backward induction equilibrium By dominates the population approximately $\frac{5}{9}$ of the time, the weak Nash equilibrium An dominates the population approximately $\frac{1}{3}$ of the time, and the other $\frac{1}{9}$ of the time, the population consists of Ay and/or Bn. Thus the assumption of single strategy selection is reasonable for the ultimatum game, simplifying our calculations without significantly affecting our results. If this assumption is not made, the invariant distribution is both difficult to compute, and dependent on the short-term dynamics. Nevertheless, the final result is unchanged: our model selects the backward induction equilibrium in the ultimatum game.

7 Discussion

We now consider several other extensions of our evolutionary model. As discussed above, we typically assume that the spread function $f(x)$ is independent of the types of the common or entering strategies; however, the model can easily be extended to allow for state-dependent spread functions. In the most general case, we can define $N(N-1)$ spread functions $f_{ij}(x)$, $i \neq j$, corresponding to the proportion of s_j mutants invading a population of strategy s_i. Then, assuming $s_i \overset{m}{\Leftrightarrow} s_j$, we calculate the invasion success probability $q_{ij} = \bar{F}_{ij}(1-m) = \int_{1-m}^{1} f_{ij}(x)\,dx$. For any other evolutionary relation between the two strategies, the spread function is not relevant. State-dependent spread functions can arise in a number of ways: for example, strategies may differ in their degree of "attractiveness," and thus the proportion of individuals who choose to switch to them. Similarly, in models of imitation resulting from dissatisfaction, different strategies may be more or less inherently "disagreeable," affecting the proportion of individuals who choose to switch from them. Dissatisfaction may be negatively correlated with the self-payoff of the common strategy, while the attractiveness of a new strategy might be positively correlated with its payoff in the current population or (if players have sufficient foresight) its self-payoff. These correlations may not be perfect, however: as discussed above, some behaviors (ex. use of a harmful but addictive substance) may be attractive despite their low payoffs in the long term, while other behaviors (ex. strict regimens of diet and exercise) may be unattractive despite high long-term payoffs. A variety of other factors may influence spread of a strategy. For example, some strategies may be easier to learn, and hence more likely to be adopted; additionally, individuals may be more likely to adopt strategies that are more similar to the strategy that they are currently using.

In addition to this flexibility in defining how new strategies spread through a population, we may also be more flexible in defining how new strategies arise. We typically assume that mutations are random, and every strategy (including the current strategy) has a $\frac{1}{N}$ chance of being selected to invade. Alternatively, mutations may be "state-dependent" as in Bergin & Lipman (1996): strategies may have different probabilities of being selected to invade, and these probabilities may depend on the current strategy as well. In general, we can define N^2 prior values r_{ij}, such that $\sum_{j=1}^{N} r_{ij} = 1$ for all i: then r_{ij} is the probability that strategy s_j will be selected to invade strategy s_i on a given time step. Thus our transition probabilities p_{ij} can be defined as $p_{ij} = r_{ij}q_{ij}$, i.e. the probability that s_j is selected to invade s_i, and that the invasion is successful. As in the case of spread functions, these priors r_{ij} may be based on payoffs (e.g. if individuals are more likely to "discover" strategies which are better replies), complexity (e.g. if simpler strategies are more likely to arise via mutation), or may be arbitrary (depending on the specifics of the evolutionary or adaptive model). We also note that the *rate* of mutations from a particular strategy s_i may be controlled by choosing the value of r_{ii}, then scaling the other priors accordingly.

Though this flexibility in the model parameters can be useful, enabling the model to generalize to a larger set of real-world interactions, it should also be noted that the choice of priors and spread functions can have a significant impact on our evolutionary results. Typically the spread functions have smaller impact, since they only affect transition probabilities between strategies in bistable relationships. Nevertheless, it is clear that any strict equilibrium can be made absorbing (and accessible) by some state-dependent spread function, and thus any equilibrium can be selected in an NxN coordination game. The choice of prior invasion probabilities can affect all transitions, with the exception of those where $q_{ij} = 0$. Thus any state i can be made absorbing by setting $r_{ii} = 1$; however, if this state is not accessible (for example, because s_i is dominated by every other strategy) it will not be selected unless the system begins in that state. In general, we only recommend the use of strategy-dependent priors or spread functions when specific aspects of the evolutionary or adaptive model require them. For a general measure of the performance of a strategy in a non-specific evolutionary model, a single spread function and uniform prior is sufficient; or if we want to simulate a best-response-type dynamic, the prior probabilities can be dependent on the payoffs of each potential mutant against the current population.

Another interesting extension of the model would be to allow the model parameters (ex. spread functions and prior invasion probabilities) to evolve over time rather than being exogeneously specified. One possible application of this type would be to compare the performance of subpopulations with high and low degrees of spread respectively. Imagine a population with two distinct subpopulations, one which communicates new ideas rapidly through the population (and thus evolves according to a spread function with large aggregate shocks) and one which does not communicate (and thus evolves according to Maynard Smith's invasion criteria). Let the population shares of various strategies within a subpopulation evolve as in our long-term model above, but also assume that the total size of a subpopulation varies as an increasing function of the average payoff to members of that subpopulation. Based on our results for the Finitely Repeated Prisoner's Dilemma and centipede games, we conjecture that communicating subpopulations are dominated more often by cooperative strategies, resulting in higher average payoffs. In other words, the communicating subpopulation does not get stuck in poor but evolutionarily stable strategies ("local maxima" of the evolutionary space), resulting in better average performance in the long run. This may allow a communicating subpopulation to out-compete another subpopulation which does not communicate among themselves: thus we conjecture that natural selection will select for the evolution of communication.

Next we consider to what extent our evolutionary model can be applied as a *rational decision process*, or a method which an individual rational player can use to decide which strategy to play against an unknown (and not necessarily rational) opponent. First, we note that there is no reason to expect a priori that a Nash equilibrium strategy will be played, since neither player is assumed to know the other's strategy choice. Thus the rational player must decide what a "reasonable" distribution of opposing strategies would be, and make the choice that maximizes expected utility given this distribution. Three possible alternatives are to assume that the opponent has *evolved* via a process of selection and mutation, has *learned* a strategy from previous interactions, or has *decided* on a strategy based on some method of imaginary play. For the first two cases, we can assume that the opponent has been picked (randomly) from a population resulting from some evolutionary or adaptive model respectively; for the third case, we can assume that the opponent's thought process (leading to his selection of a strategy) simulates such a model. Thus one reasonable assumption would be that the opposing strategy is *sampled from the invariant distribution* of some unknown model.

Now we must consider what parameters are reasonable for this model: in particular, whether mutations are *random* or *directed*, and whether mutations *spread* to a finite proportion of the population. For evolutionary models, mutations result from errors in the reproductive process, and thus are random. For learning models, mutations can be either directed (resulting from a process of myopic optimization) or random (resulting from a process of "discovery" of new ideas). We prefer to assume that mutations are random: the relative performance of a new mutant strategy determines not whether the new behavior is "discovered" by some individual in the population, but whether it is adopted by the population as a whole, or is abandoned and fades from the collective memory. In either case, the population is likely to be finite, and thus the proportion of mutants will be finite; additionally, various processes can facilitate rapid spread of a new behavior through a population, leading to a substantial proportion of the new mutant strategy. Thus it is reasonable to assume that the model satisfies the LASH property.

Similarly, for a model of decision-making via imaginary play, we note that the decision-maker has complete control over the parameters of the model. In particular, he is not required to always choose a best reply to his current "favorite" strategy, but can use randomization to his benefit. If he treats the decision-making process as a problem of *search* over a large strategy space, then including some degree of randomness in the search procedure allows him to search over a larger region of the space, reducing the probability of being trapped in local maxima that are not globally optimal. Additionally, if the decision-maker's computational resources are limited and the strategy space is large, it may be infeasible for him to compute a best reply to the current favorite strategy, and thus a "generate and test" model of decision-making is more applicable than a model of strict and myopic optimization.

Therefore, one reasonable method of decision-making (given no prior information about the opponent) would be to choose a LASH model with some arbitrary but general parameters (e.g. a uniform spread function and uniform prior), compute the invariant distribution, and choose the best reply against this distribution. We note that this method will often result in selection of a Nash equilibrium strategy, as was the case for almost all of the examples we considered. However, our model does not choose a Nash equilibrium strategy in certain cases, including long repeated games such as the Finitely Repeated Prisoner's Dilemma and the centipede game. In these games, the Nash equilibrium solutions (via arguments of iterated dominance in the normal form, or backward induction in the extensive form) are counterintuitive, and the validity of these arguments has been widely questioned in the literature. Our model does not reject the backward induction argument for the FRPD and centipede games, but for sufficiently long games, the effects of backward induction are negligible compared to the effects of "forward progression" (where the current strategy is replaced by one which cooperates much more often), and thus the system spends nearly all of the time far from equilibrium.

More generally, our model differs from Nash equilibrium because it selects strategies based on a measure of "average case" performance, rather than "worst case" performance, against the strategy space. A strategy is not a Nash equilibrium if either player can improve his performance by switching strategies: thus a strategy will not be an equilibrium if it is evolutionarily dominated by even a single alternative strategy. In our model, on the other hand, poor performance against a single strategy may be outweighed by high risk-dominance against a large proportion of the strategy space. For example, if strategy s_i evolutionarily dominates strategy s_j, we know that s_j is not "evolutionarily stable" (and not a Nash equilibrium) since it can be invaded by s_i mutants. However, assuming a uniform prior and a large strategy space, s_i will rarely be selected to invade s_j, and thus this transition will occur infrequently. Moreover, even if s_i is a strict Nash equilibrium (and hence evolutionarily stable), it may be invadable by a small proportion of mutants of many different types, and assuming a LASH spread function, these small proportions may frequently occur. Thus the "stable" strategy s_i may have very low probability in the invariant distribution. On the other hand, s_j may strongly risk-dominate every strategy except s_i, and thus have very high probability in the invariant distribution, despite being "unstable." Thus our model gives results different from Nash equilibrium, not because it makes non-equilibrium strategies any more stable, but because it *weakens* the stability of some equilibrium strategies; in certain cases, this may result in equilibrium strategies being outperformed by strategies which are not equilibria but do better in the "average case."

8 Conclusions

We have presented a class of evolutionary models consisting of selection and mutation. These models are similar to the prior models in the literature in their assumptions on the *type* and *frequency* of shocks. However, we assume that the *impact* of shocks is large: a new mutant strategy can spread to a finite proportion of the population before the short-term selection dynamics takes effect. We consider a variety of processes through which these "large aggregate shocks" can occur, and present a very general model of evolutionary dynamics with large shocks, which includes evolutionary stability and stochastic stability models as special cases. The assumption of large aggregate shocks impacts the results of the model in three main ways. First, it allows us to select between Nash equilibria in games with multiple equilibria, selecting the risk-dominant strategy in a wider class of games than stochastic stability models. Second, large shocks result in much faster convergence to the invariant distribution than in stochastic stability models, allowing this long-term distribution to be relevant on the time scales under consideration. Finally, large shocks weaken the stability of certain poor but "evolutionarily stable" equilibria. As a result of this, we find that cooperative strategies prevail *in the long run* in games such as the Finitely Repeated Prisoner's Dilemma, if the game length is sufficiently long. Thus our evolutionary model gives results which more closely correspond to solutions which are considered "reasonable" in real-world interactions, even when this contradicts the backward induction solution.

9 References

- Robert Aumann (1995). Backward induction and common knowledge of rationality. *Games and Economic Behavior* **8**, 6-19.

- Robert Axelrod (1984). *The Evolution of Cooperation.* New York: Basic Books.

- Robert Axelrod & W.D. Hamilton (1981). The evolution of cooperation. *Science* **211**, 1390-1396.

- J. Bergin and B. Lipman (1996). Evolution with state-dependent mutations. *Econometrica* **64**, 943-956.

- Cristina Bicchieri and C. Rovelli (1995). Evolution and revolution: the dynamics of corruption. *Rationality and Society* **7**(2), 201-224.

- Kenneth G. Binmore and Larry Samuelson (1997). Muddling through: noisy equilibrium selection. *Journal of Economic Theory* **74**, 235-265.

- Kenneth G. Binmore and Larry Samuelson (1999). Evolutionary drift and equilibrium selection. *Review of Economic Studies* **66**(2), 363-393.

- Kenneth G. Binmore, Larry Samuelson, and Richard Vaughan (1995). Musical chairs: modeling noisy evolution. *Games and Economic Behavior* **11**(1), 1-35.

- Jonas Bjornerstedt (1995). Experimentation, imitation, and evolutionary dynamics. Mimeo, Stockholm University.

- Jonas Bjornerstedt and Jurgen Weibull (1995). Nash equilibrium and evolution by imitation. In *The Rational Foundations of Economic Behavior*, Kenneth Arrow et al, ed. London: Macmillan.

- R. Boyd and J. Lorberbaum (1987). No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game. *Nature* **327**, 58-59.

- Antonio Cabrales (2000). Stochastic replicator dynamics. *International Economics Review* **41**(2), 451-481.

- R. Cressman (1996). Evolutionary stability in the finitely repeated prisoner's dilemma game. *Journal of Economic Theory* **68**, 234-248.

- R. Cressman and Karl H. Schlag (1998). The dynamic (in)stability of backwards induction. *Journal of Economic Theory* **83**, 260-285.

- Glenn Ellison (1993). Learning, local interaction, and coordination. *Econometrica* **61**, 1047-1071.

- Glenn Ellison (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Review of Economic Studies* **67**(1), 17-45.

- Ilan Eshel (1996). On the changing concept of evolutionary population stability as a reflection of a changing point of view in the quantitative theory of evolution. *Journal of Mathematical Biology* **34**, 485-510.

- Ilan Eshel, Marcus W. Feldman, and Aviv Bergman (1998). Long term evolution, short term evolution, and population genetic theory. *Journal of Theoretical Biology* **191**(4), 391-396.

- Dean Foster and H. Peyton Young (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology* **38**, 219-232.

- D. Friedman (1991). Evolutionary games in economics. *Econometrica* **59**, 637-666.

- Drew Fudenberg and C. Harris (1992). Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory* **57**, 420-441.
- Drew Fudenberg and David K. Levine (1998). *The Theory of Learning in Games.* Cambridge: MIT Press.
- Drew Fudenberg & Eric Maskin (1986). The folk theorem for repeated games with discounting and incomplete information. *Econometrica* **54**, 533-554.
- John Gale, Kenneth G. Binmore, and Larry Samuelson (1995). Learning to be imperfect: the ultimatum game. *Games and Economic Behavior* **8**, 56-90.
- I. Gilboa and A. Matsui (1991). Social stability and equilibrium. *Econometrica* **58**, 859-867.
- John C. Harsanyi and Reinhard Selten (1988). *A General Theory of Equilibrium Selection in Games.* Cambridge: MIT Press.
- Sergiu Hart (2002). Evolutionary dynamics and backward induction. *Games and Economic Behavior* **41**, 227-264.
- Josef Hofbauer and Karl Sigmund (1998). *Evolutionary Games and Population Dynamics.* Cambridge: Cambridge UP.
- Michihiro Kandori, George J. Mailath, and Rafael Rob (1993). Learning, mutation, and long run equilibria in games. *Econometrica* **61**(1), 29-56.
- John G. Kemeny and J. Laurie Snell (1960). *Finite Markov Chains.* Princeton: D. Van Nostrand.
- Motoo Kimura (1983). *The Neutral Theory of Molecular Evolution.* Cambridge: Cambridge UP.
- Elon Kohlberg and Jean-Francois Mertens (1986). On the strategic stability of equilibria. *Econometrica* **54**, 1003-1037.
- David M. Kreps, Paul Milgrom, John Roberts, and Robert Wilson (1982). Rational cooperation in the finitely repeated prisoner's dilemma. *Journal of Economic Theory* **27**, 245-252.
- R.D. Luce and H. Raiffa (1957). *Games and Decisions.* New York: Wiley.
- George J. Mailath (1998). Do people play Nash equilibrium? Lessons from evolutionary game theory. *Journal of Economic Literature* **36**(3), 1347-1374.
- T. Maruta (1997). On the relationship between risk-dominance and stochastic stability. *Games and Economic Behavior* **19**, 221-234.
- John Maynard Smith (1982). *Evolution and the Theory of Games.* Cambridge: Cambridge UP.

- John Maynard Smith and G.R. Price (1973). The logic of animal conflict. *Nature* **246**, 15-18.

- S. Morris, R. Rob, and H.S. Shin (1995). p-dominance and belief potential. *Econometrica* **63**, 145-157.

- John Nachbar (1990). Evolutionary selection dynamics in games: convergence and limit properties. *International Journal of Game Theory* **19**, 59-89.

- John Nachbar (1992). Evolution in the finitely repeated prisoner's dilemma. *Journal of Economic Behavior and Organization* **19**, 307-326.

- John F. Nash (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences* **36**, 48-49.

- Daniel B. Neill (2001). Optimality under noise: higher memory strategies for the alternating prisoner's dilemma. *Journal of Theoretical Biology* **211**, 159-180.

- Daniel B. Neill (2003). Evolutionary stability for large populations. Mimeo., Carnegie Mellon University.

- Abraham Neyman (1985). Bounded complexity justifies cooperation in the finitely repeated prisoner's dilemma. *Economic Letters* **19**, 227-229.

- Abraham Neyman (1999). Cooperation in repeated games when the number of stages is not commonly known. *Econometrica* **67**, 45-64.

- G. Nöldeke and Larry Samuelson (1993). An evolutionary analysis of backward and forward induction. *Games and Economic Behavior* **5**, 425-454.

- Martin A. Nowak and Karl Sigmund (1989). Game dynamical aspects of the prisoner's dilemma. *Applied Mathematics and Computation* **30**, 191-213.

- Martin A. Nowak and Karl Sigmund (1993). A strategy of win-stay, lose-shift that outperforms tit for tat in the prisoner's dilemma game. *Nature* **364**, 56-58.

- R. Radner (1986). Can bounded rationality resolve the prisoner's dilemma? In *Contributions to Mathematical Economics in honor of Gerard Debreu*, W. Hildenbrand and A. Mas-Collel, eds.

- Arthur J. Robson and Fernando Vega-Redondo (1996). Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory* **70**, 65-92.

- Jeffrey S. Rosenthal (1995). Convergence rates of Markov chains. *SIAM Review* **37**, 387-405.

- Larry Samuelson (1997). *Evolutionary Games and Equilibrium Selection.* Cambridge: MIT Press.

- Larry Samuelson (2002). Evolution and game theory. *Journal of Economic Perspectives* **16**(2), 47-66.
- Mark E. Schaffer (1988). Evolutionary stable strategies for a finite population and variable contest size. *Journal of Theoretical Biology* **132**, 469-478.
- Karl Schlag (1998). Why imitate, and if so, how? A bounded rational approach to multi-armed bandits. *Journal of Economic Theory* **78**, 130-156.
- R. Schuessler (1989). The gradual decline of cooperation: endgame effects in evolutionary game theory. *Theory and Decision* **26**, 133-155.
- Reinhard Selten (1975). Reexamination of the perfectness concept for equilibrium points in extensive-form games. *International Journal of Game Theory* **4**, 25-55.
- Reinhard Selten (1980). A note on evolutionary stable strategies in asymmetric animal conflicts. *Journal of Theoretical Biology* **84**, 93-101.
- Jeroen M. Swinkels (1992). Evolutionary stability with equilibrium entrants. *Journal of Economic Theory* **57**, 306-332.
- Peter D. Taylor and Leo B. Jonker (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* **40**, 145-156.
- Fernando Vega-Redondo (1996). *Evolution, Games, and Economic Behavior*. Oxford: Oxford UP.
- Jurgen Weibull (1995). *Evolutionary Game Theory*. Cambridge: MIT Press.
- H. Peyton Young (1993). The evolution of conventions. *Econometrica* **61**(1), 57-84.

		P2 X	P2 Y
P1	X	2/2	-1/3
	Y	3/-1	0/0

Table 1: Prisoner's Dilemma game

		P2 X	P2 Y
P1	X	2/2	0/0
	Y	0/0	0/0

Table 2: Joint venture game

		P2 X	P2 Y
P1	X	1/1	0/0
	Y	0/0	4/4

Table 3: Simple coordination game

		P2 X	P2 Y
P1	X	9/9	0/7
	Y	7/0	8/8

Table 4: A coordination game where risk- and Pareto- dominant equilibria differ.

		P2 X	P2 Y	P2 Z
P1	X	6/6	0/5	0/0
	Y	5/0	7/7	5/5
	Z	0/0	5/5	8/8

Table 5: Young's 3x3 coordination game

		P2			
		X	Y	Z	W
P1	X	μ/μ	-1/2	0/0	2/-1
	Y	2/-1	1/1	-1/2	0/0
	Z	0/0	2/-1	1/1	-1/2
	W	-1/2	0/0	2/-1	1/1

Table 6: A 4x4 NSP game

		P2			
		X	Y	Z	W
P1	X	1/1	-1/2	2/-1	0/0
	Y	2/-1	1/1	-1/2	0/0
	Z	-1/2	2/-1	1/1	0/0
	W	0/0	0/0	0/0	μ/μ

Table 7: Another 4x4 NSP game

M	invariant distribution	best s_i	cooperative	uncooperative
1	[1 0]	1	0	1
2	[.6 .2 .2]	1	.2	.6
3	[.32 .23 .24 .20]	1	.2	.32
4	[.18 .17 .21 .24 .20]	4	.2	.18
5	[.10 .11 .15 .21 .23 .20]	5	.2	.10
6	[.06 .07 .10 .15 .20 .22 .19]	6	.19	.06
8	[.03 .03 .04 .07 .10 .14 .19 .22 .19]	8	.19	.03
10	[.01 .02 .02 .03 .04 ... .14 .18 .21 .18]	10	.39	.03
20	[.002 .002 .003 ... .09 .13 .17 .19 .17]	20	.53	.007
30	[.001 .001 .001 ... .09 .12 .16 .18 .17]	30	.63	.004
40	[4×10^{-4} ... .05 .06 .09 .12 .16 .18 .16]	40	.71	.002
50	[3×10^{-4} ... .05 .06 .09 .12 .15 .17 .16]	50	.76	.002
100	[6×10^{-5} ... .05 .06 .09 .11 .14 .16 .15]	100	.88	7×10^{-4}
200	[1×10^{-5} ... .05 .06 .08 .11 .14 .15 .14]	200	.94	3×10^{-4}
400	[3×10^{-6} ... .05 .06 .08 .10 .13 .14 .14]	400	.97	2×10^{-4}
800	[8×10^{-7} ... .05 .06 .08 .10 .12 .14 .13]	800	.99	8×10^{-5}

Table 8: FRPD results for uniform spread function

		P2 c	d
P1	C	2/2	-1/3
	D	0/0	0/0

Table 9: 2-legged centipede game, normal form

		P2 Cc	Cd	Dc	Dd
P1	Cc	2/2	0.5/2.5	1/1	-0.5/1.5
	Cd	2.5/0.5	1/1	1/1	-0.5/1.5
	Dc	1/1	1/1	0/0	0/0
	Dd	1.5/-0.5	1.5/-0.5	0/0	0/0

Table 10: 2-legged centipede game, symmetrized form

		P2 X	Y
P1	X	1/1	0/2
	Y	2/0	-1/-1

Table 11: Hawk-dove game

		P2 X	Y	Z
P1	X	1/1	999/999	0/0
	Y	999/999	1/1	0/0
	Z	0/0	0/0	2/2

Table 12: 2-stable game

	Ay	An	By	$(Ay + Bn)$
Ay	.375	.125	.25	.25
An	.125	.75	.125	0
By	0	.125	.875	0
$(Ay + Bn)$	.076	.076	.25	.598

Table 13: Transition matrix for ultimatum game

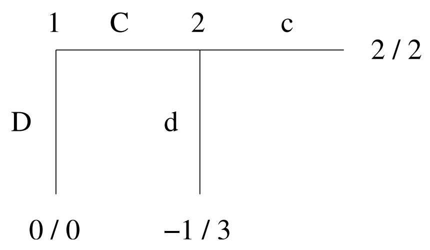

Figure 1: 2-legged centipede game

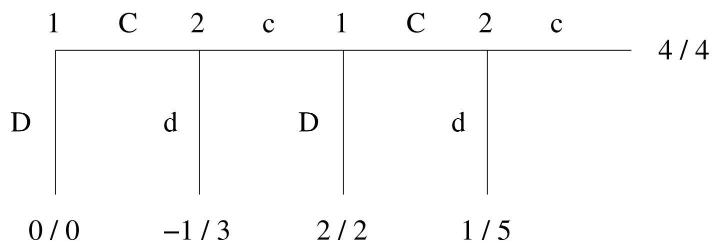

Figure 2: 4-legged centipede game

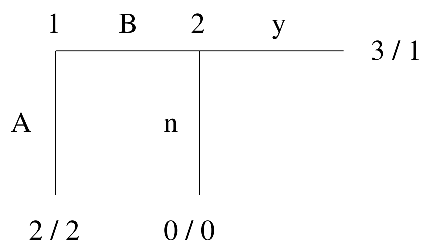

Figure 3: Ultimatum game

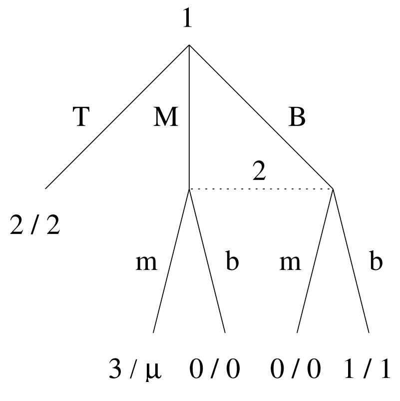

Figure 4: Dalek game