POPULATION GAMES

Lawrence E. Blume

December 1995

This paper was prepared for the Santa Fe Institute Workshop on the Economy as a Complex Evolving System II. The author's research is supported by the National Science Foundation. The author is also grateful for support from the Department of Economics at Tel Aviv University.

Correspondent:

Professor Lawrence Blume
Department of Economics
Uris Hall
Cornell University
Ithaca, NY 14853
LB19@CORNELL.EDU

1. Introduction

In most economists' view, aggregate economic activity results from the interactions of many small economic agents pursuing diverse interests. To some, this statement is a normative proposition. It immediately brings to mind Adam Smith's famous metaphor of the invisible hand, an expression of the normative superiority of markets as a resource allocation device. But even among those economists who hold a more cautious view of the superiority of markets, methodological individualism, the emphasis on the primacy of individual agents, has had an enormous impact on research into aggregate economic phenomena. For instance, it has led to the wholesale scrapping of those parts of macroeconomics lacking adequate "microfoundations", and to the acceptance of the "new classical macroeconomics", search theories of macroeconomic activity and other micro-based models. Nowadays a model of aggregate economic activity is only acceptable if it can be derived by aggregation from acceptable models of the behavior of individual agents.

Methodological individualism has reached a peculiar extreme in modern economic analysis. The great achievement of general equilibrium theory is that, given perfectly competitive behavior, equilibrium price levels can be determined without reference to the institutional details of market organization. Similarly, the inductive view of non-cooperative game theory abstracts from all details of how players interact. In this view, equilibrium describes mutually consistent beliefs among players. The nature of this consistency is the subject of the common-knowledge literature. The paradigmatic strategic interaction has two undergraduates sitting in separate closed rooms, their only link a computer screen, and so forth. The analysis puts enormous emphasis on the decision problem, but employs a model of player interaction which is irrelevant to the study of social organization.

One significant consequence of the lack of specificity in the modelling of how agents interact with one another is the indeterminacy of equilibrium behavior. In general equilibrium analysis this indeterminacy is the subject of a celebrated result, the Debreu-Mantel-Mas Colell-Sonnenschein Theorem, which states that in an exchange economy with more people that goods, any closed set of positive price vectors can be an equilibrium price set for some specification of preferences. Indeterminacy is even more problematic in general equilibrium theory with incomplete markets. Here given preference and endowment data the models can predict a continuum of equilibria. And in game theory, practitioners know that most interesting games come with an embarrassingly large set of equilibria, even after all of one's favorite refinements have been exercised.

Recent developments in so-called evolutionary game theory offer a new methodological approach to understanding the relationship between individual and aggregate behavior. These models, drawing inspiration from similar models in population genetics, utilize only vague descriptions of individual behavior, but bring into sharp focus the aggregate population behavior. Attention shifts from the fine points of individual-level decision theory to dynamics of agent interaction. Consequently the analysis is fairly robust to descriptions of individual behavior, but depends crucially on the underlying institutions and norms that govern interactions. Collectively, these models have come to be known as *population games.*

The idea of population games is as old as Nash equilibrium itself. Literally. In his unpublished Ph.D. dissertation, John Nash writes:[1]

> It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal.
>
> To be more detailed, we assume that there is a population (in the sense of statistics) of participants for each position of the game. Let us also assume that the 'average playing' of the game involves n participants selected at random from the n populations, and that there is a stable average frequency frequency with which each pure strategy is employed by the 'average member' of the appropriate population.
>
> . . .
>
> Thus the assumption we made in this 'mass-action' interpretation lead to the conclusion that the mixed strategies representing the average behavior in each of the populations form an equilibrium point.
>
> . . .
>
> Actually, of course, we can only expect some sort of approximate equilibrium, since the information, its utilization, and the stability of the average frequencies will be imperfect.

Nash's remarks suggest that he saw the mass-action justification as providing a rationale for mixed strategy equilibrium. Only recently have scholars attended to those processes which could serve as the source of the mass-action interpretation. Research has demonstrated the expected result, that the relationship of population games to Nash equilibrium is more complex than Nash's comments suggest.

This paper will discuss that research which has attended to the dynamics of systems continually subject to random perturbations, or noise. Dean Foster and H. Peyton Young (1990) introduced the study of these systems to evolutionary game theory, and also the concept of a *stochastically stable state.* A stochastically stable state is a set of states minimal with respect to the property that the fraction of time the system spends in the set is bounded away from zero as the random perturbations become vanishingly small. Both the Foster and Young paper and Drew Fudenberg and Christopher Harris (1992) introduce noise to the aggregate process. David Canning (1992), Lawrence Blume (1993), Michihiro Kandori, George Mailath and Raphael Rob (1993, hereinafter referred to as KMR) and Young (1993, hereinafter referred to as Young) introduce noise directly

---

[1] Nash (1950), pp. 21–23, quoted in Weibull (1994).

into individual decisionmaking.

The early papers that actually identify stochastically stable sets get their strongest results in two-by-two games. They demonstrate, under differing assumptions, the stochastic stability of the risk dominant equilibria in coordination games. Kandori and Rob (1995) extend the KMR analysis to K-by-K coordination games and supermodular games. Young (1993b) applies the techniques of his (1993) paper to a "divide the dollar" bargaining game. He shows that for a particular class of learning processes, the unique stochastically outcome is a generalization of the Nash bargaining solution. Georg Noldeke and Larry Samuelson (1994) apply related techniques to the problem of two populations of players from which players are matched to play a signalling game.

Many different models of player interaction can be studied with population game techniques. The forerunner of the entire population games literature is Thomas Schellings's (1971) analysis of segregation in housing markets. This paper presents a "local interaction" model; that each player directly competes with only a few players in the population. Blume (1993) developed the asymptotic analysis and selection results for local interaction models, and Glenn Ellison (1993) compared the sample path behavior of local and global interaction models applied to two-by-two coordination games. Blume (1993b) studied the behavior of local interaction in particular classes of $K \times K$ games with pure best-response dynamics; that is, without noise.

More recently, other topics have been studied, including local interaction with endogenous partner choice, forward-lookingbehavior, and variable and endogenous population size. It is almostcertainly too early for a comprehensive survey of the literature, and none will beattempted here. Instead, this paper surveys some of the important questionsaddressed in the literature, with the particular view of pushing towards economic applications.

Left out of this survey is the related literature on deterministic dynamics in a continuum population of players. These models are frequently justified as some kind of large-numbers limit of stochastic population games. But this connection, discussed in section 4.1, is ill-founded.

2. The Basic Model

This section sets up a basic population model. Later sections will specialize this model in different ways. The motivation for the story comes from KMR and Young, but here a modelling strategy is introduced which makes the usual limit analysis transparent.

2.1. Description

Begin with an K-player normal-form game, where player k has I_k pure strategies. Let S denote the set of pure strategy profiles, and let $G_k : S \to \mathbf{R}$ denote player k's payoff function. Suppose that there is a population consisting of N players each of K different types, each type corresponding to a player in the original game.

A *strategy revision process* $\{p_t\}_{t\geq 0}$ tracks the number of players choosing each strategy. The state space for the process is

$$\mathcal{N} = \prod_{k=1}^{K} \left\{ (M_1, \ldots, M_{I_k}) : \sum_{i=1}^{I_k} m_i = N \right\}$$

The vector $n = (n_{11}, \ldots, n_{KI_K})$ lists the number of players playing each strategy. Each state encodes the distribution of strategies played by each type. A Nash equilibrium state, for instance, is a state in which the distribution of strategies played by the types is a (mixed strategy) Nash equilibrium.

At random times τ_{ni} a *strategy revision opportunity* arrives for player n. The times $\tau_{ni} - \tau_{n\,i-1}$ between her ith and $i-1$st strategy revision opportunities are independently distributed, exponentially with parameter σ. In between times, she may not revise her choice. So when player n makes a choice at time τ_{ni}, she is *locked in* to that choice for some random period until the next revision opportunity arrives.

When a strategy revision opportunity arrives, the lucky player will choose according to the immediate expected return given her expectations about play. It is through the process of expectation generation that players' behavior are coupled together. For instance, in the KMR model and throughout this paper except where stated, expectations are just the empirical distribution of current play. Young has a more complicated random expectation model. The underlying assumption in all these models is the *myopia* of players. Whatever their source, expectations are applied to the payoff matrix to compute immediate expected payoffs. That is, decisions are based on immediate returns and not on other, perhaps more long-run desiderata. When there are multiple best responses, the player is equally likely to choose any one of them.

Myopic best-response and lock-in do not completely describe the evolution of choice. The third component of the model is *noisy behavior.* We assume that at each strategy revision opportunity, with some probability the player fails to realize her best response. A parameter ϵ will describe the probability that the player best responds. That is, at a strategy revision opportunity a player of type k will realize a choice drawn from the distribution $b_k(n, \epsilon)$, where $\lim_{\epsilon \to 0} b_k(n, \epsilon) = \mathrm{br}_k(n)$, the uniform distribution on the best responses of the player given her expectations in state N, and $\lim_{\epsilon \to \infty} b_k(n, \epsilon)$ is a given completely mixed distrbution on the I_k strategies of type k. The choice distribution $b_k(n, \epsilon)$ is assumed to be continuous in ϵ for each n. Two such rules are

$$b_k(n, \epsilon) = (1 - \epsilon)\,\mathrm{br}_k(n) + \epsilon q \qquad 0 \le \epsilon \le 1 \tag{2.1}$$

$$b_k(n, \epsilon) = \frac{1}{Z}\left(\exp \epsilon^{-1} \sum_{s_{-k} \in I_{-k}} G_k(s_{-k}, s_k)\; \mathrm{pr}_n(s_{-k}) \right)_{s_k=1}^{I_k} \tag{2.2}$$

where Z is a normalizing constant and pr_n is the empirical distribution of joint strategies of players of types other than k. The first rule is that studied by KMR and, with important modifications, Young. I shall refer to this as the “mistakes model” of choice, because players make mistakes with a fixed probability ϵ. The second rule, the log-linear choice rule, was introduced in Blume (1993) for the study of local interaction models. This model of stochastic choice has a long history in psychology and economics, dating back to Thurstone (1927). Several arguments have been put forth to justify the assumption of random choice: Random experimentation, random preferences and, most important, unmodelled bounded rationality.

Because the probability that two or more players have a strategy revision opportunity at the same instant is zero, the state can advance or decline by atmost 1 unit at any step. The process $\{p_t\}_{t\geq 0}$ is a K-type birth-death process, whose long-run behavior is easy to study. Birth-death processes are characterized by a set of parameters which describe the rate at which the the process switches from one state to the next. For single-type birth-death processes these parameters are called birth- and death- rates. Let λ_{nm} denote the transition rate from state n to state m. For all positive ϵ the process is irreducible, hence ergodic, and the invariant distribution ρ^ϵ is characterized by the *balance conditions*

$$\sum_m \lambda_{nm}\rho^\epsilon(n) - \lambda_{mn}\rho^\epsilon(m) = 0 \tag{2.3}$$

For all single-type birth-death processes and for a special class of multitype birth-death processes known as *reversible* birth-death processes, the invariant distribution is characterized by the *detailed balance conditions*

$$\lambda_{nm}\rho^\epsilon(n) - \lambda_{mn}\rho^\epsilon(m) = 0 \tag{2.4}$$

which directly determine the odds-ratios of the invariant distribution.

Most of the λ_{nm} are 0. They are non-zero only for n-m pairs such that m can be reached from n by a player of some type, say k, switching from some strategy l to another strategy l'. In these cases the λ_{nm} can be conputed as follows: Strategy revision opportunities arrive to each player at rate σ, so they arrive collectively to the group of type k players playing l at rate σn_{kl}. The probability that such a player will switch to l' is $b_k(n,\epsilon)(l')$. Consequently,

$$\lambda_{nm} = \sigma n_{kl} b_k(n,\epsilon)(l')$$

It is hard to solve for ρ^ϵ analytically from the balance conditions (2.3), but ρ^ϵ is accessible by computation. The detailed balance conditions (2.4) can easily be solved for ρ^ϵ in those models where they apply.

Another version of this model has only one type, and the players in the single population play a symmetric game against each other. In this version, as before, the state space is the number of players playing each srategy. When players have only 2 choices, knowledge of the number of players playing one strategy is enough to infer the number of players playing the other, and so the state can be summarized by the number of players playing strategy 1. The resulting process will be a single-type birth death process, and the invariant distribution is characterized by the detailed balance conditions (2.4).

2.2. Best Response Dynamics

When $\epsilon = 0$, a transition from state n to state m where a k player switches from l to l' is possible only if l' is a best response for players of type k in state n. The $\epsilon = 0$ process is stochastic because of the random order of strategy revision opportunities, but it need not be irreducible. In particular, if a game has two or more strict Nash equilibria the process cannot be irreducible. In this case the state space can be partitioned into a finite number of sets $Q_0, Q_1, \ldots, Q_L$ such that Q_0 is the set of transient states, and the Q_l, $l \geq 1$ are the *communication classes* of the processes. That is, for any state in Q_0 the set of paths that remain in Q_0 has probability 0. For $l > 0$ there is a path with positive probability between any two states in each Q_l and there are no paths out of any Q_l. The communication classes provide an ergodic decomposition of the state space $\mathcal{N}$. The process $\{p_t\}_{t \geq 0}$ with initial condition in any communication class Q_l is irreducible on Q_l, and hence will have a unique invariant distribution ρ_l with support Q_l. The set of invariant distributions is the convex hull of the ρ_l. It follows immediately from the balance conditions (2.3) that if ρ is an invariant distribution for the $\epsilon = 0$ process and $\rho(n) > 0$, then $\rho(m) > 0$ for all other states m in the same communication class as n.

The important questions at this stage concern the relationship between the supports of the communication classes and the payoff structure of the game. In the following game, $K = 2$. This game has two Nash equilibria: The strict equilibrium (D, L) and the mixed equilibrium where the players play D and L, respectively, with probability 0, and put equal weight on the remaining two strategies.

	R	C	L
U	$1,3$	$2,2$	$0,0$
M	$2,2$	$1,3$	$0,0$
D	$0,0$	$0,0$	$1,1$

Example 1.

There are two communication classes. One is a singleton; the state in which all row players play D and all column players play L. The other is the set of states in which no row player plays D and no column player plays L. All states not in one of these two communication classes are transient. Notice that the first communication class describes a Nash equilibrium. This illustrates a connection between singleton communication classes and Nash equilibria. A *strict Nash equilibrium* is a Nash equilibrium in pure strategies wherein each player has a unique best response.

Proposition 2.1: A state n is a singleton communication class if and only if it is a strict Nash equilibrium state.

Proof of 2.1: Suppose that n is a singleton communication class. All players of each type must be best-responding to the play of the other types; otherwise there is another state m with $\lambda_{nm} > 0$ and the communication class is not a singleton. Similarly, each type must have a unique best response. If players of a given type k are all choosing a given action i and action j is also a best response, there is a path from n to any other state in which the distribution of play for types other than k is that of n, some players of type k are choosing i and the remaining players are choosing j.

Conversely, suppose that $(i_1, \ldots, i_K)$ is a strict Nash equilibrium, and consider the state n such that $n_{ki_k} = N$. Any change of strategy for any player will leave that player worse off, so $\lambda_{nm} = 0$ for all $m \neq n$. Thus $\{n\}$ is a singleton communication class. $\square$

It is possible for all communication classes to be singletons. The *best-reply graph* of a game G is a directed graph whose vertices are the pure strategy profiles. There is an arrow from s to s' if there is one player i such that $s'_i \neq s_i$ is a best-response for player i to s_{-i} and $s'_j = s_j$ for all other j. A *sink* of the graph is a state from which no arrow exits. Clearly a sink is a strict Nash equilibrium. The graph is *acyclic* if there are no paths that exit and later return to any state s. Then every path from any state s leads to a sink. The graph is *weakly acyclic* if from every state there is a path leading to a sink. Weakly acyclic games may have best response cycles, but there is an alternative best response from some state that leads out of the any cycle.

Theorem 2.1: (Young) If the best-reply graph of the game is weakly acyclic, then each sink is a singleton communication class, and these are the only communication classes.

Many games, including coordination games and ordinal potential games with only strict equilibria, are weakly acyclic.

2.3. Stochastic Stability

When $\epsilon > 0$, the random deviations from best response choice make every state reachable from every other state, and so the process $\{p_t\}_{t\geq 0}$ is irreducible and recurrent. Consequently it will have a unique invariant distribution, and the invariant distribution will have full support. None theless, as ϵ becomes small distinctive behavior emerges. States that are transient when $\epsilon = 0$ can only be reached by stochastic perturbations when ϵ is positive, and these will be very unlikely. It will also be easier to reach some communication classes then others, and harder to leave them. This can be studied by looking at the limiting behavior of the invariant distributions as ϵ becomes small, and this limiting behavior is quite regular. Again let ρ^ϵ denote the invariant distribution for the process $\{p_t\}_{t\geq 0}$ with parameter value ϵ.

Theorem 2.2: If the $b_k(n, \epsilon)$ are continuous in ϵ, then $\rho^* = \lim_{\epsilon \to 0} \rho$ exists and is a probability distribution on $\mathcal{N}$. The distribution ρ^* is an invariant distribution for the process $\{p_t\}_{t\geq 0}$ with parameter $\epsilon = 0$.

The proof of this follows immediately upon noticing that the solution set to the equation system (2.3) varies upper-hemicontinuously with ϵ. $\square$

Those states in the limit distribution ρ^* are stable in the sense that they are likely to arise when small random perturbations of best-reply choice are introduced.

Definition 2.1: A state $n \in \mathcal{N}$ is *stochastically stable* if $\rho^*(n) > 0$.

Stochastic stability is a property of communication classes. If a state n in communication class Q_l is stochastically stable, then all states m in Q_l are stochastically stable.

Stochastic stability is a means of identifying communication classes that are easier to reach and those that are harder. To the extent that it sorts among communication classes, stochastic stability serves the same role in evolutionary game theory that refinements do in eductive game theory. But the motivations are quite different. Where the trembles in the conventional refinements literature exist only in the heads of the players, the trembles of evolutionary game theory are real shocks to a dynamic system, which have an effect on the system's ability to reach particular equilibria.

KMR, Young, Ellison (1995) and Evans (1993) provide powerful techniques that apply to very general stochastic perturbations of discrete time, countable state dynamical systems. One advantage of the birth-death formalism is that its special structure makes the identification of stochastically stable states easier. Identifying the stochastically stable states is straightforward when the detailed balance condition (2.4) is satisfied for all ϵ, and the invariant distributions can be computed.

When only the balance conditions (2.3) are satisfied, stable states can be identified by a technique which is related to Ellison's and Evans'. The following fact is easy to prove. Let $\{x_v\}_{v=0}^{\infty}$ be a discrete time, irreducible Markov process on a countable state space, and let $\{y_u\}_{u=0}^{\infty}$ be the process which results from recording the value of the x_v process every time it hits either state 0 or state 1. It is easily seen that the $\{y_u\}_{u=0}^{\infty}$ process is Markov and irreducible. Let ρ_x and ρ_y denote the invariant distributions for the two processes. Then

$$\frac{\rho_x(1)}{\rho_x(0)} = \frac{\rho_y(1)}{\rho_y(0)}$$

The second process is much easier to study. The transition matrix of the $\{y_u\}_{u=0}^{\infty}$ process is

$$\begin{pmatrix} 1 - f_{01} & f_{10} \\ f_{01} & 1 - f_{10} \end{pmatrix}$$

where f_{ij} is the probability of transiting from i to j without passing again through i (a so-called *taboo probability*). Thus

$$\frac{\rho_x(1)}{\rho_x(0)} = \frac{f_{01}}{f_{10}}$$

The connection to strategy revision processes $\{p_t\}_{t\geq 0}$ goes as follows. Sample $\{p_t\}_{t\geq 0}$ at every strategy revision opportunity. Let x_v denote the value of $\{p_t\}_{t\geq 0}$ at the time of arrival of the vth strategy revision opportunity. The process $\{x_v\}_{v=1}^{\infty}$ is a discrete time Markov chain, called the *embedded chain* of the process $\{p_t\}_{t\geq 0}$. Moreover, because the arrival rate of strategy revision opportunities is unaffected by the value of $\{p_t\}_{t\geq 0}$, the invariant distributions of ρ_x are precisely those of $\{p_t\}_{t\geq 0}$.

The problem of identifying stable states now comes down to a minimization problem. Supose there is a path that goes from 0 to 1 with h trembles, and that at least $h + k$ trembles are required to return from 1 to 0. If the probability of a tremble is ϵ, the ratio f_{01}/f_{10} will be of order ϵ^{-k}, which converges to ∞ as ϵ becomes small. Thus conclude that 1 is more stable than 0. Maximal elements of this "more stable than" order are the stochastically stable states. This method extends in a straightforward manner to sets. It should be clear that if 1 and 0 are in the same communication class and 1 is more stable than 2, then 0 is more stable than 2. Thus we can talk of the stability of communication classes.

An example will illustrate how this works. Consider the game

	a	b	c
A	3 , 2	2 , 3	0 , 0
B	2 , 3	3 , 2	0 , 0
C	0 , 0	0 , 0	x , x

Example 2.

There are two populations of size N playing against each other, and x is positive. There are two communication classes, $Q_1 = \{(0,0,N),(0,0,N)\}$ with support $\{C\} \times \{c\}$ and $Q_2 = \{(L,N-L,0),(M,N-M,0) : 0 \leq L, M \leq N\}$ with support $\{A,B\} \times \{a,b\}$. Dynamics are of the KMR variety: each player best responds with probability $1-\epsilon$ and trembles with probability ϵ to a given completely mixed distribution.

First, find the path from Q_1 to Q_2 with the minimum number of trembles. If enough column players deviate to b, then it will pay the row players to deviate to B. Once enough of them do so, it will be in the interests of the remaining column players to switch to b. A calculation shows that B becomes a best response when fraction $x/(3+x)$ or more of the population plays b, so the number of trembles required is, up to an integer, $Nx/(3+x)$. Clearly no path from Q_1 to Q_2 requires less trembles. On the other hand, the easiest way to move from Q_2 to Q_1 is to start from a state where all row players play B and all column players play a. (This state can be reached from any other state in the communication class using only best responses.) If enough row players deviate to C, then it pays the rest to do so. Once enough row players have switched to C, the column players will want to switch to c. Again, the number of trembles required is approximately $2N/(2+x)$.

The singleton set Q_1 will be stochastically stable when the number of trembles needed to leave exceeds the number needed to return. A computation shows that, for large enough N that the integer problem does not interfere, the critical value of x is $\sqrt{6}$. If $x > \sqrt{6}$, then for N large enough only Q_2 is stochastically stable. If $x < \sqrt{6}$, then only Q_1 is stochastically stable.

In a number of papers, stochastic stability appears to be a kind of equilibrium refinement, selecting, for instance, the risk-dominant equilibrium in a two-by-two coordination game. Samuelson (1994) demonstrates that stochastically stable sets can include states where all players use weakly dominated strategies. And as the previous example demonstrates, stochastically stable sets may contain states which correspond to no Nash equilibria. The singleton set Q_2 contains a Nash equilibrium state. Assuming N is even, the state $\big((N/2, N/2, 0), (N/2, N/2, 0)\big) \in Q_1$ is a Nash equilibrium state. But no other state in Q_1 corresponds to a Nash equilibrium.

3. Examples

The canonical examples of this emerging literature are the results on the stochastic stability of risk-dominant play in coordination games, due to KMR and Young. This section works out a birth-death version of the KMR model, a selection model somewhat in the spirit of Young, a general analysis for potential games and a local interaction model.

3.1. The KMR Model

In the KMR model, expectations held by each player are given my the empirical distribution of play at the moment of the strategy revision opportunity and players' choice follows the mistakes model (2.1). That is, if $p_t > p^*$ the player will choose a, if $p_t < p^*$ the player will choose b, and if $p_t = p^*$ the player will randomize in some given fashion. It summarizes the effects of a presumed *interaction model.* At discrete random moments pairs of players are briefly matched, and the payoff to each player from the match is determined by the two players' choices and the payoff matrix. All matches are equally likely, and players do not know with whom their next match will be.

Let M^* denote the largest integer m such that $M^*/N < p^*$. For all $M \le M^*$ a player's best response is to choose a. It will soon be apparent that without any essential loss of generality we can assume that $M^*/N < p^* < M^* + 1/N$, so that the possibility of indifference never arises. For all $M > M^*$ a player's best response is to choose a. For the KMR model, the birth and death rates are

$$\lambda_M = \begin{cases} \sigma(N-M)\epsilon q_a & \text{if } M < M^*, \\ \sigma(N-M)\big((1-\epsilon) + \epsilon q_a\big) & \text{if } M \ge M^* \end{cases}$$

$$\mu_M = \begin{cases} \sigma M\big((1-\epsilon) + \epsilon q_b\big) & \text{if } M \le M^*, \\ \sigma M \epsilon q_b & \text{if } M > M^*. \end{cases}$$

To see this, consider the birth rate for $M < M^*$. A birth happens when a b player switches to a. There are $N - M$ such players, so the arrival rate of strategy revision opportunities to the collection of b players is $\sigma(N - M)$. When a player gets a strategy revision opportunity in state M, b is her best response. Consequently she will switch to a only by error, which happens with probability ϵq_a. Thus the birth rate $\sigma(N - M)\epsilon q_a$. The other cases are argued similarly.

The invariant distribution is

$$\log \frac{\rho^\epsilon(M)}{\rho^\epsilon(0)} = \begin{cases} \log \binom{N}{M} - M \log \dfrac{1-\epsilon+\epsilon q_b}{\epsilon q_a} & \text{if } M \le M^*, \\ \log \binom{N}{M} + (M - 2M^*) \log \dfrac{1-\epsilon+\epsilon q_b}{\epsilon q_a} & \text{if } M > M^*. \end{cases}$$

When $\epsilon = 1$, the dynamics are governed entirely by the random term. On average, fraction q_a of the population will be choosing a at any point in time. Furthermore, the least likely states are the full coordination states.

As ϵ decreases, the invariant distribution changes in a systematic way. The hump in the middle begins to shrink and shift one way or the other, and the probabilities at the two ends begin to grow. At some point they become local modes of the distribution, and finally they are the only local modes. The probability function is U-shaped. I call this process *equilibrium emergence*.

Equilibrium selection describes the behavior of the invariant distribution for very small ϵ. As ϵ continues to shrink one of the local modes grows at the expense of the other, and as ϵ converges to 0 the invariant distribution converges to point mass at one of the full coordination states. In the birth-death version of the KMR model, the odds ratio of the two states is

$$\log \frac{\rho^\epsilon(N)}{\rho^\epsilon(0)} \approx 2\big(\frac{1}{2} - p^*\big) \log \frac{1 - \epsilon + \epsilon q_b}{\epsilon q_a}$$

which converges to $+\infty$ as $\epsilon \to 0$ since a is risk-dominant. Notice that the selection criterion is independent of the error distribution q. No matter how large q_b, so long as it is less than one full coordination at a is selected for. At $q_b = 1$ there is a dramatic change in behavior. When $q_b = 1$ the invariant distribution converges to point mass at state 0 as ϵ shrinks, because all states greater than p^* are transient.

Equilibrium emergence and selection is ilustrated in the four graphs below. In this game, a is risk dominant and the threshold probability is $p^* = 0.4$

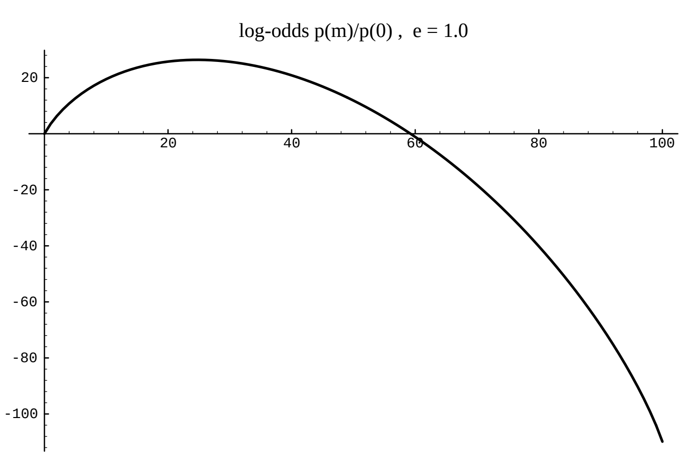

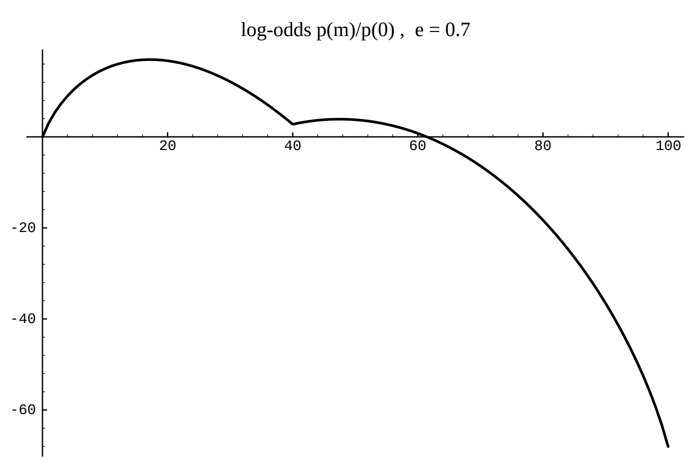

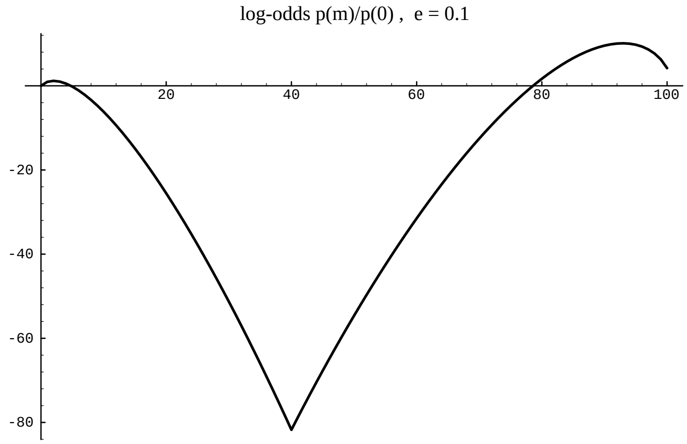

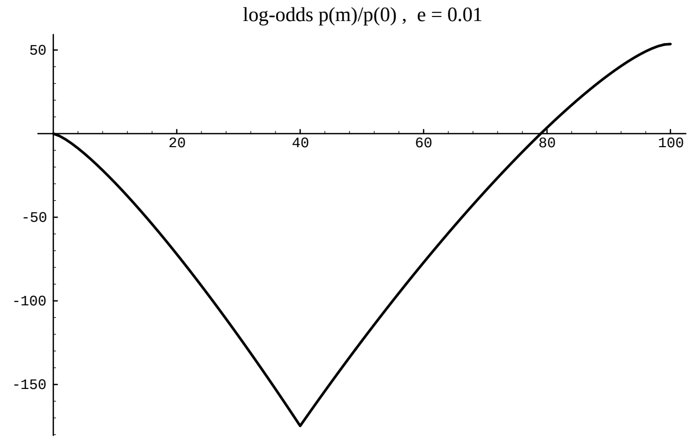

When the tremble probability ϵ is 1, the invariant distribution is the binomial distribution with sample size n and success probability q_a. As ϵ falls, the valley with bottom at the threshold state p^*n quickly appears. When the tremble probability falls to 17% the right and left hump are approximately the same height. By the time the tremble probability reaches 10% the mode on the right, at 93, is $20,000$ times more likely than the mode on the left, at 2. By the time the tremble probability is 1%, the local modes are 0 and 100, and the former occurs with infinitesimal probability.

3.2. Noisy Selection Models

Young's model has a more sophisticated expectations generation mechanism. Each player has a memory of decisions from the immediate past. Each player samples that memory, and best responds to the empirical distribution of the sample. Thus expectations are stochastic, since different samples can be drawn from the same memory. Binmore Samuelson and Vaugn (1993) have called such models "noisy selection models". In the KMR model, in any state there is one action which is a best response, and the other action can arise only because of a tremble. In Young's model and other noisy selection models, any action can be a best response except in a very few absorbing states.

Developing that model here would take too much space. A simpler story in similar spirit goes as follows. A player receiving a decision opportunity when the state of the process is M draws a random sample of size K from a binomial distribution with probability M/N. If the number of successes exceeds p^*K she chooses a, otherwise b. Thus expectations are random, and the mean expectation is M/N. When K is large this model behaves like the KMR model. But when K is small the noise in expectations has a considerable effect on the long-run dynamics. This can be seen in the series of pictures below. On the left $\epsilon = 0.1$ while on the right $\epsilon = 0.01$. The game and all other parameters are the same as before.

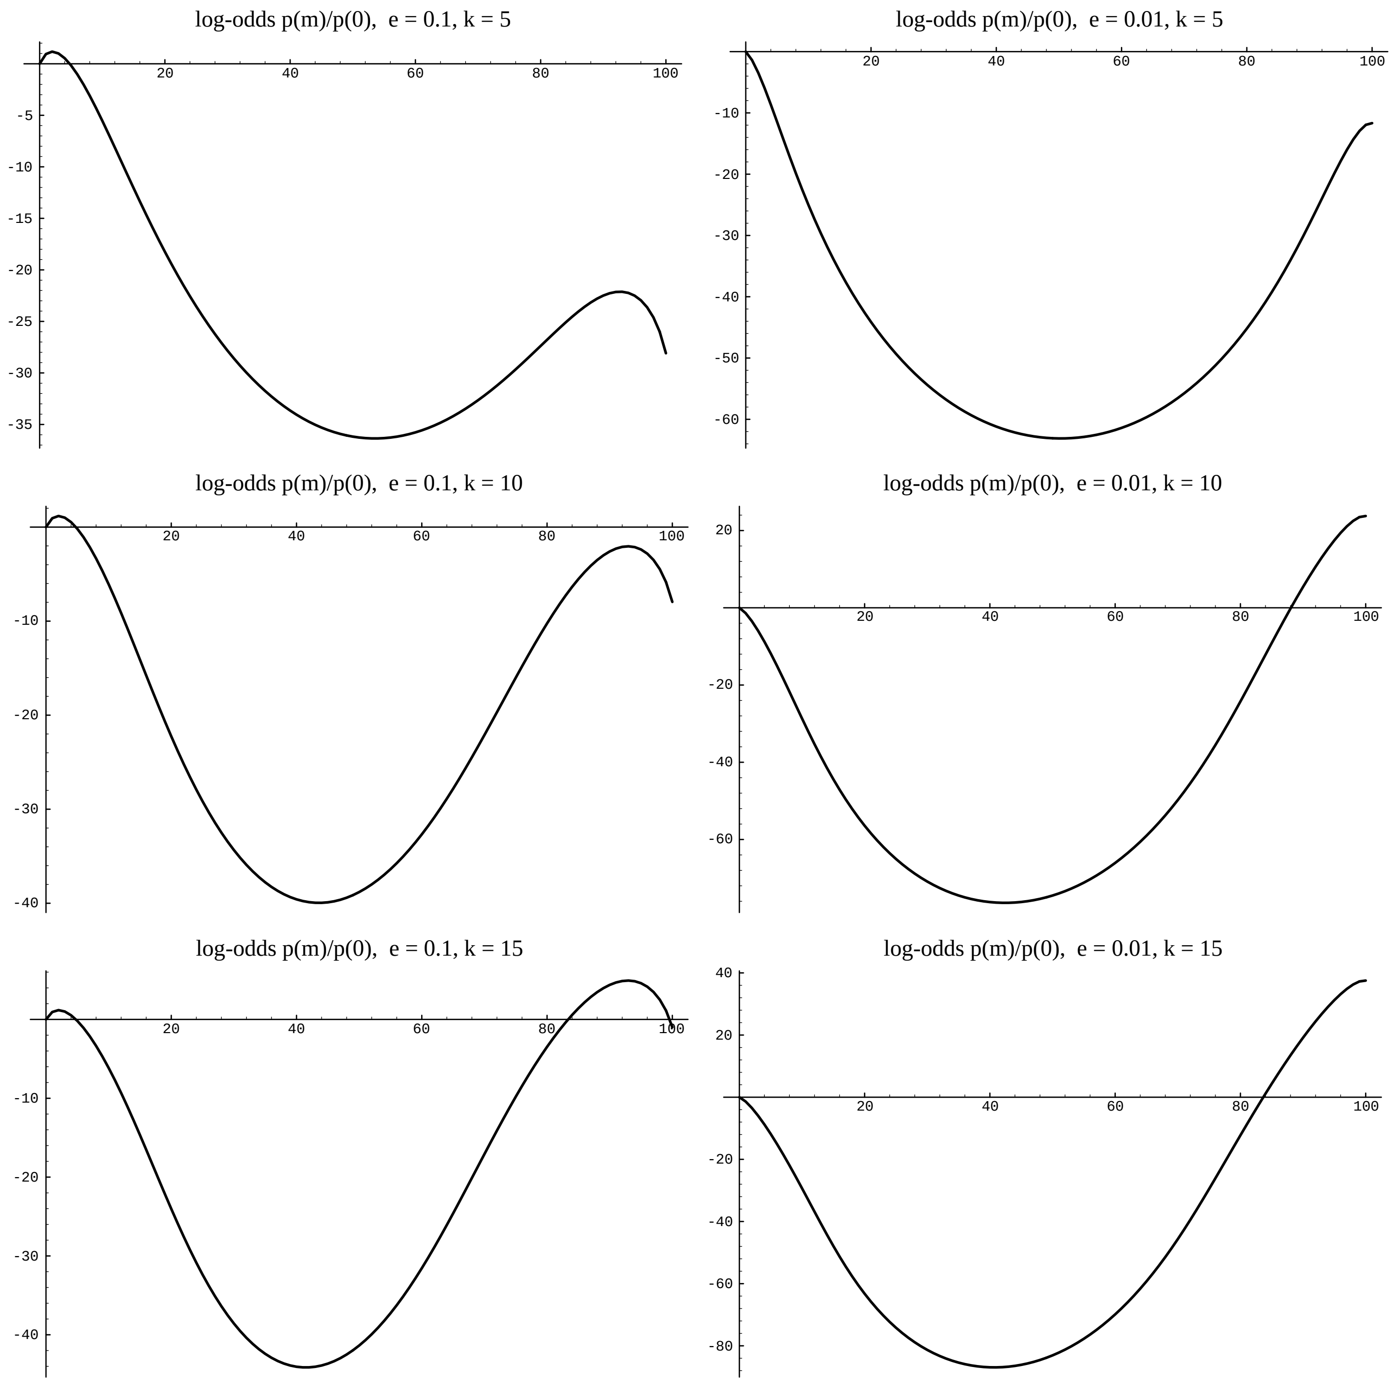

3.3. Potential Games

Potential games are games in which all payoff differences to all players can be described by first differences of a single real-valued function. Nash equilibria appear as local maxima of this potential function, and consequently the dynamics of adjustment to equilibrium are easier to study for potential games than they are for general games. Special as this structure might seem, it turns out that many important economic models have potentials. The analysis of potential games is due to Monderer and Shapley (1993).

Consider a K-player normal-form game with a set S of pure strategy profiles and payoff functions $G_k : S \to \mathbf{R}$.

Definition 3.1: A *potential* for a K-player game is a function $P : S \to \mathbf{R}$ such that for each player k, strategy profiles s and strategies s'_k for player k,

$$G_k(s) - G_k(s'_k, s_{-k}) = P(s) - P(s'_k, s_{-k})$$

Potential games have many special properties for particular dynamics. For instance, fictitious play converges for potential functions. Here we will see that the strategy revision process with the log-linear choice rule satisfies the detailed balance equations if and only if the game has a potential.

If the detailed balance conditions (2.4) are met, the ratios of the transition rates define odds ratios. That is, if the transition rates from n to m and m to n are positive,

$$\frac{\rho(m)}{\rho(n)} = \frac{\lambda_{nm}}{\lambda_{mn}}$$

Imagine now a closed path of states $n, m_1, \ldots, m_J, n$, with positive transition rates between adjacent states in both directions. Since

$$\frac{\rho(m_1)}{\rho(n)} \frac{\rho(m_2)}{\rho(m_1)} \cdots \frac{\rho(n)}{\rho(m_J)} = 1$$

It follows that if the detailed balance conditions are satisfied,

$$\frac{\lambda_{nm_1}}{\lambda_{m_1 n}} \frac{\lambda_{m_1 m_2}}{\lambda_{m_2 m_1}} \cdots \frac{\lambda_{m_J n}}{\lambda_{n m_J}} = 1$$

Conversely, it can be shown that if this condition, known as Kolmogorov's condition, is satisfied, then (2.4) has a solution.

Checking Kolmogorov's condition for all closed paths would be impossible. But in fact it suffices to check it only for four-cycles, because all closed paths can be put together out of four-cycles. let n denote a given state. A four-cycle is the following kind of closed path: (i) a player of type k switches from s_k to s'_k; (ii) a player of type l switches from s_l to $s_{l'}$; (iii) a type k player switches from s'_k back to s_k; (iv) a type l player switches from s'_l back to s_l. (Reversing the order of the return trip breaks the four-cycle up into two two-cycles, and the conditions are always satisfied for two-cycles.) For a four-cycle Kolmogorov's condition is

$$\frac{b_k(n,\epsilon)(s'_k)}{b_k(n - e_{ks_k} + e_{ks'_k}, \epsilon)(s_k)} \frac{b_l(n - e_{ks_k} + e_{ks'_k}, \epsilon)(s'_l)}{b_l(n - e_{ks_k} + e_{ks'_k} - e_{ls_l} + e_{ls'_l}, \epsilon)(s_l)} \frac{b_k(n - e_{ks_k} + e_{ks'_k} - e_{ls_l} + e_{ls'_l}, \epsilon)(s'_k)}{b_k(n - e_{ls_l} + e_{ls'_l}, \epsilon)(s_k)} \frac{b_l(n - e_{ls_l} + e_{ls'_l}, \epsilon)(s_l)}{b_l(n,\epsilon)(s_l)} = 1$$

If players' choice behavior is described by the log-linear choice rule (2.2), a calculation shows that the left-hand side of the preceding equation is

$$\begin{aligned}\exp \frac{1}{\epsilon}\{P(n - e_{ks_k} + e_{ks'_k}) - P(n) + P(n - e_{ks_k} + e_{ks'_k} - e_{ls_l} + e_{ls'_l}) - \\ P(n - e_{ks_k} + e_{ks'_k}) + P(n - e_{ls_l} + e_{ls'_l}) - \\ P(n - e_{ks_k} + e_{ks'_k} - e_{ls_l} + e_{ls'_l}) + P(n) - P(n - e_{ls_l} + e_{ls'_l})\} = 1\end{aligned}$$

so the condition is satisfied.

Since the invariant distribution is characterized by the detailed balance condition (2.4), it is straightforward to compute odds ratios of various states simply by chaining together the detailed balance conditons along a path. The general relationship is

$$\frac{\rho^\epsilon(m)}{\rho^\epsilon(n)} = C_{nm} \exp \frac{1}{\epsilon}\{P(m) - P(n)\}$$

where C_{nm} is a positive constant depending only on n and m, and not on ϵ. Fix a reference state n_0. The invariant distribution is of the form

$$\rho^\epsilon(n) = \frac{C_{n_0 n} \exp \frac{1}{\epsilon} P(n)}{\sum_m C_{n_0 m} \exp \frac{1}{\epsilon} P(m)}$$

There is a general equilibrium selection result for potential games, corresponding to the risk-dominant equilibrium selection result for two-by-two coordination games. As ϵ becomes small, the mass of the invariant distribution concentrates on the set of global maxima of the potential function;

$$\lim_{\epsilon \to 0} \rho^\epsilon\{\operatorname{argmax} P(n)\} = 1.$$

All symmetric two-by-two games are potential games. For coordination games, the potential maximum is at the state where the players coordinate on the risk-dominant equilibrium. Thus there is a connection between this result and those of KMR and Young. KMR and Young introduce trembles in the same way but use different best-response rules. This analysis of potential games uses the same best-response rule as KMR but introduces trembles in a different fashion. Nonetheless all three analyses find the same selection Theorem, that the unique stochastically stable set is the state in which all players play the risk-dominant strategy.

3.4. A Local Interaction Model

The previous models have each player interacting with the entire population. Another important class of models has each player interacting only with a subset of the population — her neighbors. Although each player interacts directly with only a few others, a chain of direct interactions connects each player indirectly with every other player in the population. Imagine a population of players, each interacting with her neighbors in a coordination game. Initially, different geographic regions of the populace are playing different strategies. At the edge of these regions the different strategies compete against one another. Which strategies can survive in the long run? In the long run can distinct regions play in distinct ways, or is heterogeneity not sustainable?

Suppose that players are interacting in a $K \times K$ symmetric game with payoff matrix G. Let I denote the set of K pure strategies. Given is a finite graph. Vertices represent players, and the edges connect neighbors. Let V denote the set of players (vertices), and V_v the set of neighbors of player v. A configuration of the population is a map $\phi : V \to I$ that describes the current choice of each player. If the play of the population is described by the configuration ϕ, then the instantaneous payoff to player v from adopting strategy s is $\sum_{u \in V_v} G\bigl(s, \phi(u)\bigr)$.

A strategy revision process is a respresentation of the stochastic evolution of strategic choice, a process $\{\phi_t\}_{t \geq 0}$. The details of construction of the process are similar to those of the global interaction models. At random times τ_{v_i} player v has a strategy revision opportunity, and $\tau_{vi} - \tau_{vi-1}$, the ith interarrival time, is distributed exponentially with parameter σ and is independent of all other interarrival times. At a strategy revision opportunity, players will adopt a new strategy according to rules such as (2.1) or (2.2), but here best responses are only to the play of v's neighbors.

One would think that the behavior of stochastic strategy revision processes for local interaction models would be profoundly affected by the geometry of the graph. As we shall see, this is certainly the case for some choice rules. But when the game G has a potential, the basic result of the previous section still applies. If G has a potential $P(s, s')$, then the "lattice game" has a potential

$$\tilde{P}(\phi) = \frac{1}{2} \sum_{v} \sum_{u \in V_v} P\bigl(\phi(v), \phi(u)\bigr) .$$

To see that $\tilde{P}$ is indeed a potential for the game on the lattice, compute the potential difference between two configurations which differ only in the play of one player. Fix a configuration ϕ. Let ϕ_v^s denote the new configuration created by assigning s to v and letting everyone else play as specified by ϕ. Then

$$\begin{aligned}
\tilde{P}(\phi_v^{s'}) - \tilde{P}(\phi_v^{s}) &= \frac{1}{2} \sum_{u \in V_v} P\bigl(s', \phi(u)\bigr) - P\bigl(s, \phi(u)\bigr) + P\bigl(\phi(u), s'\bigr) - P\bigl(\phi(u), s\bigr) \\
&= \sum_{u \in V_v} P\bigl(s', \phi(u)\bigr) - P\bigl(s, \phi(u)\bigr) \\
&= \sum_{u \in V_v} G\bigl(\phi(u), s'\bigr) - G\bigl(\phi(u), s\bigr)
\end{aligned}$$

which is the payoff difference for v from switching from s to s'. The middle inequality follows from the symmetry of the two-person game.

If choice is described by the log-linear strategy rule (2.2), then as the parameter ϵ becomes small, the invariant distribution puts increasing weight on the configuration that maximizes the potential. The invariant distribution is

$$\rho^{\epsilon}(\phi) = \frac{\exp -\frac{1}{\epsilon}\tilde{P}(\phi)}{Z(\epsilon)}$$

where $Z(\epsilon)$ is a normalizing constant.

Two-person coordination games are potential games, and the strategy profile that maximizes the potential has both players playing the risk-dominant strategy. Consequently the configuration that maximizes $\tilde{P}$ has each player playing the risk-dominant strategy. More generally, for any symmetric K-by-K potential game, if there is a strategy t such that $\max_{u,v} P(u,v) = P(t,t)$, then the configuration $\phi(v) \equiv t$ maximizes $\tilde{P}$, and so strategy t is selected for as ϵ becomes small. This is true without regard to the geometry of the graph.

This independence result depends upon the game G. A two-by-two symmetric "anti-coordination game" has a potential, but it is maximized only when the two players are playing different strategies. Now the set of maxima of $\tilde{P}$ depends on the shape of the graph. Consider the following two graphs:

and suppose coordinating pays off 0 and failing to coordinate pays off 1. In the first graph $\tilde{P}$ is maximized by a configuration which had the four corners playing identically, but different from that of the player in the center. In the second graph $\tilde{P}$ is maximized only by configurations which have the three corners alternating their play.

4. Analytical Aspects of the Models

Both the short-run and the long-run behavior of population models are important objects of study. Long-run behavior is characterized by the process' ergodic distribution. When the population is large, the short run dynamicsare nearly deterministic, and can be described by the differential equation which describes the evolution of mean behavior. A fundamental feature of these models, but all to often ignored in applications, is that the asymptotic behavior of theshort-run deterministic approximation need have no connection to theasymptotic behavior of the stochastic population process. A classic example of this arises in the KMR model applied to a two-by-two coordination game.

4.1. The Short-run Deterministic Approximation

Begin with the KMR model described above, and rescale time so that when the population is of size N strategy revision opportunities arise at rate σ/N. The mean direction of motion of $\{p_t\}_{t\geq 0}$ depends upon the state, Suppose $p_t < p^*$. The process advances when a b player trembles in favor of a, and it declines when an a player either best responds or trembles in favor of b. Thus the probability of an advance by $1/N$ is $(1-p)\epsilon q_a$, and that of a decline is $p(1-\epsilon+\epsilon q_b)$, and so the expected motion is $-p_t+\epsilon q_a$. Similar reasoning on the other side of p^* gives an expected motion of $1-p_t+\epsilon q_b$. When N is large, this suggests the following differential equation for mean motion of the system.

$$\dot{p}_t = \begin{cases} -p_t + \epsilon q_a & \text{if } p_t < p^*, \\ 1 - p_t - \epsilon q_b & \text{if } p_t > p^*. \end{cases}$$

A simple strong law argument suggests that this should be a good local approximation to the behavior of sample paths for large N. In fact it can be made uniformly good over bounded time intervals. The following Theorem is well-known in population genetics (Ethier and Kurtz, 1986). An elementary proof can be found in Binmore, Samuelson and Vaughn (1993). Fix $p_0 = p(0)$ and let $p(t)$ denote the solution to the differential equation for the given boundary value.

Theorem 4.1: For all $T < \infty$ and $\delta > 0$, if N is large enough then $\sup_{0\leq t\leq T} \|p(t) - p_t\| < \delta$ almost surely.

When $\epsilon = 1$ the equation has a unique globally stable steady state at $p = q_a$. But when ϵ is small p^* divides the domain into basins of attraction for two stable steady states, at ϵq_a and $1-\epsilon q_b$.

Differential equations like this arise in the literature on deterministic population models with a continuum player population. Similar equations are derived for equilibrium search and sorting models. (See, for example, Diamond (1982).) But the problem with using these equations to study long-run behavior can be seen in this example. The asymptotic analysis of the differential equation does notapproximate the long-run behavior of the stochastic process. For small ϵ theprocess spends most of its time near only one equilibrium state — thatwhere players coordinate on the risk-dominant equilibrium. In general itis true that the support of the limit invariant distribution taken as$\epsilon \to 0$ is a subset of the set of stable steady states, but the twosets are not identical.

Other stochastic approximations have also been studied. Binmore, Samuelson and Vaughn (1993) construct a degenerate diffusion approximation by the discrete-time version of expanding the action of the birth-death generator on a function into a Taylor series and dropping the terms of order greater than 2. The resulting approximation is degenerate because as N grows the diffusion coefficient converges to 0. For all finite N the diffusion approximation has a unique invariant distribution which converges to a point mass in the limit. Examples show, however, that this point mass can be concentrated on the wrong equilibrium state; that is, coordination on the risk-dominated equilibrium. They fail to predict the asymptotic behavior of the stochastic population process for the same reason the differential equations fail. Nonetheless they can be useful for making probabilistic statements about the magnitude of the error in the short-run differential approximation.

4.2. How Long is the Long Run?

The equilibrium selection theorems for population games are all asymptotic results. They describe properties of the long-run distribution of states. For economic analysis the long run may be too long. Thus it is also important to have an idea of how long it might take to reach a particular state. First passage time distributions are useful for describing the relevant aspects of sample path behavior.

Consider first one-dimensional models such as the KMR model or the noisy selection model in a two-by-two coordination game with a single population. Suppose strategy a is risk-dominant. The equilibrium selection result says that when ϵ is small, the invariant distribution puts most of its mass on the state N in which all players use a. The first passage time from state M to 1 is $\tau(M) = \inf_t\{p_t = N;\, p_0 = M\}$. The first moment $\nu^1(M)$ of the first passage time distributions is given by the following difference equation:

$$\nu^1(M) = \frac{1 + \lambda(M)\nu^1(M+1) + \mu(M)\nu^1(M-1)}{\lambda(M) + \nu(M)}$$

For a small population of 25, the log of the expected first passage time to state 1 from initial state M is plotted in the following graphs for $\epsilon = 0.1$ and $\epsilon = 0.01$. For these plots, $p^* = 0.4$ so the boundary between the two basins of attraction is at $M = 10$.

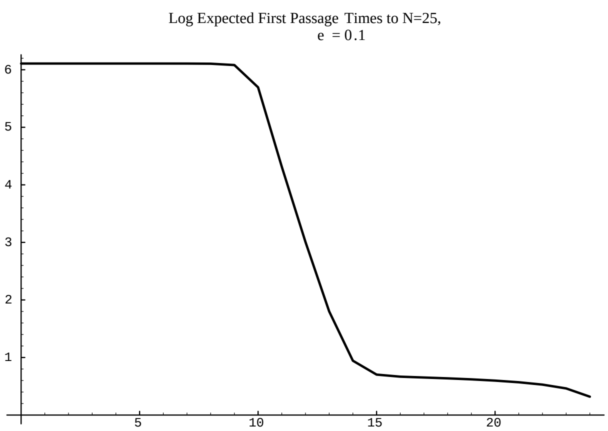

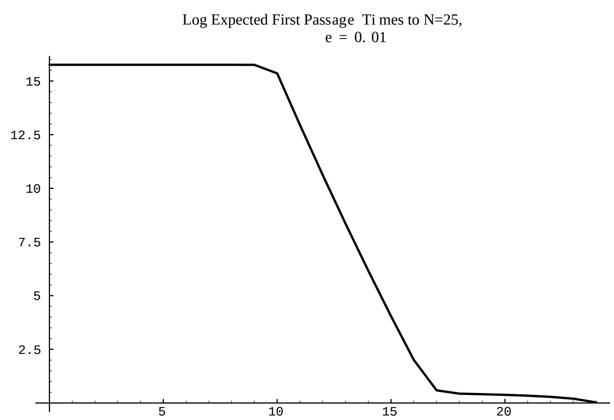

The graphs show that once the process enters the basin of attraction of a given equilibria, it will move to that equilibrium rather quickly. But should it leave that basin of attraction, a long time will pass before the process returns. A better way of scaling time is by the number of strategy revision opportunities required for the transit. Strategy revision opportunities come at rate 1 for each individual — at rate 25 for the entire population. Thus when $\epsilon = 0.1$, for the process to move from state 0 to state 25 requires, on average, on the order of 10^6 strategy revision opportunities for each individual in the population, or 10^7 revision opportunities for the entire population. These orders increase to 10^{15} and 10^{16}, respectively, when $\epsilon = 0.01$.

The usual complaint about stochastic perturbations is that when perturbations are small, the expected first passage time from one equilibrium state to the other is so small that the asymptotic analysis, including the selection results, is irrelevant. The following graph plots expected first passage times from 0 to 25 as a function of the tremble probability ϵ.

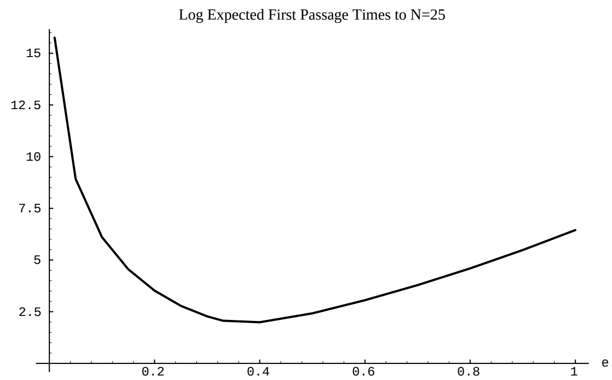

The striking fact illustrated in this graph is that, to the extent that the waiting times are viewed as being long, this is a function of the number of intermediate states that need to be traversed and not a consequence of the improbability of trembles. In this example $\epsilon = 1$ represents pure random drift; at a strategy revision opportunity a move in either direction is equally likely. The expected waiting time for $\epsilon = 1$ is comparable to that of $\epsilon = 0.1$. For $N = 25$ and $\epsilon = 0.1$ the mode of the invariant distribution is at $p = 0.84$, and $\rho(25)/\rho(0) \approx 2400$, so at an ϵ value for which the expected first passage time is comparable to that for random drift, the shape of the invariant distribution already reflects the properties of the two equilibria.

Notice that the expected first passage time initially decreases as ϵ falls. As ϵ becomes smaller, it becomes harder to move between basins of attraction, but easier to move within a basin of attraction to its sink. Initially, the first effect dominates, and so the expected first passage times fall. Subsequently the second effect dominates, and so the expected first passage times increase.

The expected first passage time alone says little about the sample path behavior of the population process. What we really want to know about is the likelihood of fast and slow transitions. A simple application of Markov's inequality gives

$$\text{Prob}\{\tau(M) \geq 2k\nu^1(0)\} \leq 2^{-k}$$

But since $\nu^1(0)$ is so large for even a small population with moderate values of ϵ and p^*, this bound is not very helpful. So far nothing is known about the shape of the lower tail of the first-passage-time distribution; the probability that the passage is quick.

Even if the probability of a fast transition is small, the asymptotic analysis tells us meaningful facts about the behavior of population process. Suppose a stationary population process is observed at some random time τ, uncorrelated with the evolution of the process. Then the distribution of p_τ is described by the ergodic distribution. In other words, an econometrician studying a naturally occurring population process with small ϵ would have to assume he or she is observing a stochastically stable state. This has clear implications for statistical inference concerning the values of payoffs, the likelihood of trembles, and other important parameters of the model.

4.3. Robustness to the Perturbation Process

Population games are intentionally vague about the details of noisy choice. This is an acceptable modelling strategy only if these details do not significantly effect the results. This question has been investigated in two recent papers, by Bergin and Lipman (1994) and Blume (1994).

The equilibrium selection results of KMR and Young are shown by their respective authors to satisfy some invariance to the specification of noise. If the updating rule is the mistakes model, then both KMR and Young demonstrate risk-dominant selection regardless of the distribution q from which choice is determined in the event of a tremble, so long as q places positive weight on all strategies. The log-linear model also leads to risk-dominant equilibrium selection, although it is distinctly different from the mistakes model. But Bergin and Lipman (1994) show that it is possible to introduce noise to choice in a way that leads to risk-dominated equilibrium selection. Their analysis is really much more general than this, but the following example will illustrate their point. Consider again the KMR model, but with the following change: When b is a best response, b is chosen with probability $1-\epsilon^2$, while when a is a best response, a is chosen with probability $1-\epsilon$. So trembling away from b is ϵ times as likely as trembling away from a. A calculation like that of section 3.1 shows that

$$\begin{aligned}\frac{\rho(N)}{\rho(0)} &= C\left(\frac{\epsilon^2}{1-\epsilon^2}\right)^{M^*}\left(\frac{1-\epsilon}{\epsilon}\right)^{N-M^*}\\ &= O(\epsilon^{3M^*-N})\end{aligned}$$

If strategy a is risk dominant, then $2M^* < N$. But in this model, selecting strategy a requires that $3M^* < N$, that $p^* < 1/3$. In fact, by making trembling away from b ϵ^k times as likely as trembling away from a for large k, the threshold probability p^* necessary to guarantee a-selection can be made arbitrarily close to 0.

Since it is possible to change the stochastically stable states by changing the noise process, it becomes important to indentify the class of noise processes which give rise to the same stochastically stable states as do the mistakes model and the log-linear choice model. This question is addressed in Blume (1994) for symmetric two-by-two games. Consider the class of choice rules where the log-odds of choosing a over b depend only on the payoff difference between the two states and on a multiplicative parameter β:

$$\log\frac{\mathrm{pr}\{a\,|\,p\}}{\mathrm{pr}\{b\,|\,p\}} = \beta g\bigl(\Delta(p)\bigr)$$

where $\Delta(p)$ is the payoff difference between a and b when fraction p of the population plays a. This includes the log-linear model, where $g(x) = x$, and the mistakes model with $q_a = q_b = 1/2$, where $g(x) = \mathrm{sgn}(x)$. (Take $\beta = \log\bigl((1-\epsilon)/\epsilon\bigr)$.)

Suppose the choice rule has the property that only payoff differences matter. That is, the odds of choosing a over b when the payoff difference between a and b is Δ are the same as the odds of choosing b over a when the payoff difference between b and a is Δ. Then g will be skew-symmetric. All skew-symmetric g give rise to the same stochastically stable states. Furthermore, if a given g is not skew-symmetric, then there is a game for which this g gives stochastically stable states different from that given by any skew-symmetric g.

The skew-symmetry says that names of strategies do not matter — only their payoff advantage. Although it is a natural property within this framework, this framework does not encompass models such as the mistakes model where $q_a \neq q_b$. However the results extend to the class of choice rules

$$\log \frac{\mathrm{pr}\{a \mid p\}}{\mathrm{pr}\{b \mid p\}} = \beta g\big(\Delta(p)\big) + \beta r\big(\Delta(p), \beta\big)$$

where g is skew-symmetric but r is not, so long as $\limsup_\beta \beta r\big(\Delta(p), \beta\big) < \infty$. This larger class of models includes all mistakes models. (In fact, the only requirement on r is that $\beta \|r(\Delta, \beta)\|$ diverge slower than $\beta g(\Delta)$ as β grows.) The lesson of these two papers is that biases in the effect of deviations from best-response on the identity of stochastically stable states matter, but only if the biases are unbounded as choice converges to pure best-response.

These invariance results for global interaction models are encouraging. They suggest that, within broad limits, the actual specification of noise in choice does not matter. Unfortunately the situation is not so good for local interaction models. Consider the graph on the left:

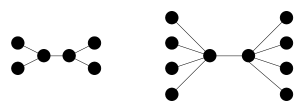

Take any two-by-two coordination game with a unique risk-dominant strategy and $p^* > 1/3$, and consider the two configurations which have all players playing identically. With either the mistakes model (2.1) or the log-linear model (2.2) no other configuration can be stochastically stable (or part of a stochastically stable set. With the mistakes model both fully-coordinated configurations are stochastically stable (singleton) sets. It takes two mutations to move from all play a to all play b, and two to move back. In the log-linear model only the configuration in which everyone plays the risk-dominant strategy is stochastically stable, as it alone maximizes the value of the potential function $\tilde{P}(\phi)$. This is essentially an integer problem, that only a majority of ones' neighbors playing the risk dominant strategy can make the risk dominant strategy a best response. Links can be added to this graph to preserve the phenomenon for an arbitrarily large population of players. For the graph on the right, suppose that $p^* > 2/5$. Then again it takes two mutations to move from one coordinated configuration to the other, so both are stochastically stable. Notice that there is only one way of moving from the dominated state to the dominant state, but 8 ways of moving from the dominated state to the dominant state. Thus the limit distribution will put more weight on "all play the dominant strategy" than on "all play the dominated strategy", but both will be present. Again the potential function $\tilde{P}(\phi)$ is maximized only by the configuration "all play the dominant strategy". How special these examples are is unknown, so the adequacy of not modelling noise is not yet determined. Notice too that these examples demonstrate that the invariance of the selected configuration to the geometry of the graph is not a universal property of nice selection rules. In the graphs above the "all play the dominated strategy" configuration is stochastically stable under the mistakes model, but if more edges were added to the graphs this would no longer be the case.

5. Economic Applications

There are four barriers to the application of population game techniques to serious economic models.

The treatment of time is inadequate, different interaction models are needed, heterogoeneity among agents must be considered, and the population game paradigm must be extended beyond two-player games. In this section I discuss some recent attempts to bring population game models closer to economic problems.

5.1. Time

In the models discussed so far, decisionmaking is essentially static. Players form expectations about the current state of the process, and respond to those expectations, subject to noise in choice, so as to maximize their instantaneous return. This separation between choice and dynamics makes the analysis of population games models particularly simple. But economic decisionmakers are typically concerned about the future as well as the present. Evaluating the future requires that each trader try to forecast the strategy revision process, and take account of these forecasts when searching for the best response at a strategy revision opportunity. If there is any connection between the forecasts and the actual behavior of the strategy revision process, such as the hypothesis that expectations are rational, then the dynamic behavior of the strategy revision process cannot be simply computed from the choice rule.

The substantive effects of allowing decisionmakers to consider the future are important. Some have argued that in coordination games where the risk-dominant strategy is not payoff-dominant, risk-dominant selection is an artifact of myopia. That is, players who put sufficient emphasis on the future will tend to coordinate on the payoff-dominant choice when choice is not too noisy. This question was addressed by Blume (1995), who found quite the opposite. If players are sufficiently patient, then when noise is sufficiently small the only equilibrium is to select the risk-dominant strategy at every outcome, regardless of the state of the process. The equilibrium concept employed here is called a *rational population equilibrium*, where all players correctly identify the stochastic process that describes the evolution of their opponents' play. Given this knowledge, each player's choice problem is a dynamic program.

Consider again the basic model of section 2 where, at random intervals, strategy revision opportunities arrive. That model simply specified a rule of behavior for each player; a choice of action in each state . Now we need to derive such a rule. In order to do so, we need to be specific about how players interact. The story — told but not formally modelled — in most population games papers is the "billiard ball model" of player interaction. At random intervals, two players (billiard balls) meet (collide), and receive payoffs depending on the strategies they are currently using as described by the payoff matrix G. Formally this can be described by assigning to each pair of players i, j a sequence of times $\{\tau_{ijn}\}_{n=0}^{\infty}$ such that $\tau_{ij0} = 0$ and the random variables $\tau_{ij\,n} - \tau_{ij\,n-1}$ are distributed independent of each other and of the strategy revision opportunity interarrival times, exponentially with parameter 1. The time τ_{ijn} is the time of the nth meeting of players i and j.

We shall suppose that the game is symmetric; that is, there is only one player type. Let S denote the set of K pure strategies for the game. Player's actions depend upon what their opponents are doing. Let $\Delta_{N-1} = \{M_1, \ldots, M_K : M_k \geq 0 \quad \text{and} \quad \sum_k M_k = N-1\}$. This is the state space for each player's decision problem. A *policy* is a map $\pi : \Delta_{N-1} \to P(S)$, the set of all probability distributions on S. If all opponents are employing policy π, their collective play, the state of the decision problem, evolves as a multitype birth-death process. Suppose the decision maker is currently employing strategy k. The transition rate from state m to state $m' = m - e_j + e_l$, requiring that a j player switches to l, is

$$\lambda^k_{mm'} = \sigma m_j \big((1-\epsilon)\pi_l(m - \epsilon_j + \epsilon_k) + \epsilon q_l\big) .$$

Since m_j opponents are playing j, σm_j is the arrival rate of strategy revision opportunities to this group. If our player is playing k, then the state of one of her opponents playing j is $m - e_j + e_k$. Thus $\pi_l(m - e_j + e_k)$ is the probability that the opponent playing j will want to choose l, and $(1-\epsilon)\pi_l(m - \epsilon_j + \epsilon_k) + \epsilon q_l$ is the probability that the opponent will end up playing l. Let $\lambda^k_m = \sum_{m'} \lambda^k_{mm'}$.

Suppose that players discount the future at rate r. If all opponents are playing according to policy π, then the optimal choice for a player whose opponent process is in state m is the solution to a dynamic programming problem. That problem has the Bellman equation

$$\begin{aligned} V(k,m) &= E\Big\{e^{-r\tilde\tau} \sum_{m' \neq m} \frac{\lambda^k_{mm'}}{\lambda^k_m + \sigma + 1} V(k,m') + \frac{1}{\lambda^k_m + \sigma + 1}\big(v(k,m) + V(k,m)\big) + \\ &\qquad \frac{\sigma}{\lambda^k_a + \sigma + 1} \max_{\pi \in P(S)} \sum_l \big((1-\epsilon)\pi_l + \epsilon q_l\big) V(l,m)\Big\} \\ &= \sum_{m' \neq m} \frac{\lambda^k_{mm'}}{r + \lambda^k_m + \sigma} V(k,m') + \frac{1}{r + \lambda^k_m + \sigma} v(k,m) + \\ &\qquad \frac{\sigma}{r + \lambda^k_m + \sigma} \max_{\pi \in P(S)} \sum_l \big((1-\epsilon)\pi_l + \epsilon q_l\big) V(l,m) \end{aligned} \tag{5.1}$$

where $v(k,m) = (N-1)^{-1} \sum_{l \in S} m_l G(k,l)$ is the expected value of a match with a randomly chosen opponent. A *rational population equilibrium* is a policy π such that $\pi_k(m) > 0$ implies $k \in \operatorname{argmax} V(k,m)$ when all opponents are using policy π. Among other things, Blume (1995) shows that when r is large, the optimal policy is to maximize the expected payoff flow from the current state — the KMR model. More surprising, in a two-by two coordination game with unique risk-dominant strategy a, if r and ϵ are sufficiently small, the only rational population equilibrium is $\pi_a(m) \equiv 1$.

A virtue of the birth-death formalism is that, for a given conjectured equilibrium, the dynamic program is easily solved and so the equilibrium can be checked. This is really the main point of the paper, that intertemporal equilibrium analysis can be conducted within the population games framework.

5.2. Interaction Models

A criticism frequently levelled at the existing population games literature is that the "billiard ball" model of player interaction is not useful for modelling any interesting economic phenomena.[2] Fortunately, other kinds of player interaction are equally amenable to analysis. Strictly within the population game framework, Blume (1994) also considers a model where players receive a continuous payoff flow from their current match. At random moments, players are matched with new opponents, and so forth. At any moment of time, each player is matched with an opponent. Consequently, the myopic behavior of this model is not the KMR model. Instead, the myopic player maximizes the return against the play of her current opponent. This leads to different asymptotic behavior. In a two-by-two coordination game there are two stochastically stable sets, each one a singleton consisting of a state in which all players coordinate. Patient play in this interaction model mimics that of billiard-ball interaction.

Still other models are within easy reach of contemporary technique. For instance, consider a model where from time to time players are matched, and later these matches break up. New matches can only be made from among the unmatched population. This matching model can be set up by assigning to each player pair two Poisson processes. When the first one advances, the player pair forms a match if both are unmatched; otherwise nothing happens. When the second one advances, if the players are matched they break up; otherwise nothing happens. Now different experiments can be performed with strategy revision opportunities by assuming that strategy revision opportunities arrive at different rates for the paired and unpaired players. Suppose they arrive at rate 0 for the unpaired players. For myopic players this hypothesis generates continuous payoff flow dynamics. On the other hand, suppose they arrive at rate 0 for paired players. This leads to KMR dynamics. Most interesting (and not yet done) is to study the intermediate regimes in order to understand the transition between voter model and KMR dynamics. I conjecture that in two-by-two coordination games all intermediate regimes have only one stochastically stable set, consisting only of the state in which all players coordinate on the risk-dominant strategy.

5.3. Beyond Two-Player Games

More crucial then how players meet is the assumption that the interaction, however it occurs, takes the form of a single two-person game. This section will exhibit two models which depart from the conventional paradigm in different ways. The first model introduces player heterogeneity in with two-person interactions. In the second model, players choose which of two groups to join, and payoffs are determined by the size of the groups.

5.3.1 Matsuyama's Fashion Model

The model, due to Kiminori Matsuyama (1992), provides an elegant study of the evolution of fashion. Refer to the original publication for a discussion of the motivations, modelling strategy and implications. Matsuyama presented what can be interpreted as the large-numbers limit analysis. Here the model will simply be recast as a population model. The simple story is that there are two types of clothes, red and blue, and two types of people. Types 1s are conformists and type 2s are

---

[2] The Menger-Kiotaki-Wright model of monetary exchange (Menger, 1892 and Kiyotaki and Wright, 1989) and some versions of Diamond's search equilibrium model (Diamond, 1982) are exceptions.

non-conformists. Conformists want to look like the rest of the population, and therefore will match the dress of the majority of the population. Non-conformists want to look different, and so will match the dress of the minority of the population. The number of type i people wearing red is m_i. The total population size is n, and the size of the subpopulation of type i is n_i.

This leads to the following specification of dynamics. Transition rates are:

$$m_1 \to m_1 + 1 \quad \text{at rate} \quad \sigma(n_1 - m_1)\begin{cases} \epsilon q_r & \text{if } m_1 + m_2 < n/2 \\ 1 - \epsilon + \epsilon q_r & \text{if } m_1 + m_2 > n/2 \end{cases}$$

$$m_1 \to m_1 - 1 \quad \text{at rate} \quad \sigma m_1 \begin{cases} 1 - \epsilon + \epsilon q_b & \text{if } m_1 + m_2 < n/2 \\ \epsilon q_b & \text{if } m_1 + m_2 > n/2 \end{cases}$$

$$m_2 \to m_2 + 1 \quad \text{at rate} \quad \sigma(n_2 - m_2)\begin{cases} 1 - \epsilon + \epsilon q_r & \text{if } m_1 + m_2 < n/2 \\ \epsilon q_r & \text{if } m_1 + m_2 > n/2 \end{cases}$$

$$m_2 \to m_2 - 1 \quad \text{at rate} \quad \sigma m_2 \begin{cases} \epsilon q_b & \text{if } m_1 + m_2 < n/2 \\ 1 - \epsilon + \epsilon q_b & \text{if } m_1 + m_2 > n/2 \end{cases}$$

Take $\sigma = 1/n$, and let θ denote the fraction of conformists. The differential equation system describing the short run dynamics is

$$\begin{aligned} \dot{p}_1 &= -\theta(p_1 - \epsilon q_r) & &\text{if } p_1\theta + p_2(1-\theta) < 1/2, \\ \dot{p}_2 &= (1-\theta)(1 - p_2 - \epsilon q_b) & & \end{aligned}$$

$$\begin{aligned} \dot{p}_1 &= \theta(1 - p_1 - \epsilon q_b) & &\text{if } p_1\theta + p_2(1-\theta) > 1/2. \\ \dot{p}_2 &= -(1-\theta)(p_2 - \epsilon q_r) & & \end{aligned}$$

where p_i is the fraction of red-dressers in the type-i population. Above the boundary line $p_1\theta + p_2(1-\theta) = 1/2$ the conformists switch to red and the non-conformists to blue. Below the boundary line they switch the other way.

When $\theta < 1/2$, evaluation of the differential equation system shows that for small ϵ the system dives into the boundary line. The slope of the two solutions at the boundary are such that the escape from the boundary cannot occur. When $\theta > 1/2$, the natural sinks near $(0, 1)$ for the regime where conformists switch to blue and the nonconformists to red, and $(0, 1)$ for the other regime, are below and above the line, respectively. (They are only "near" because of the tremble.) Thus curves never hit the line, and the flow is towards the sinks. The two cases are illustrated in the following pictures.

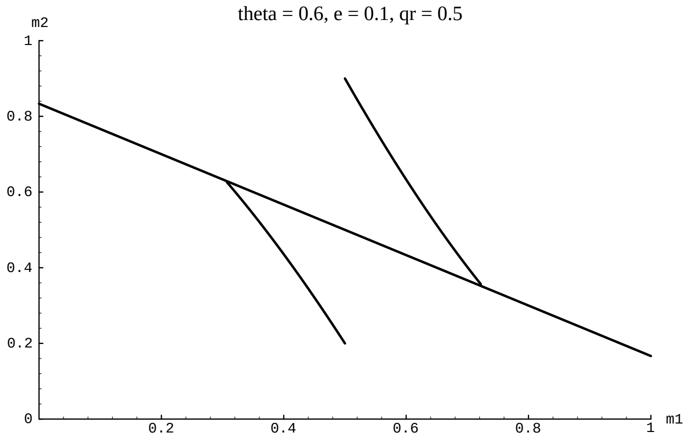

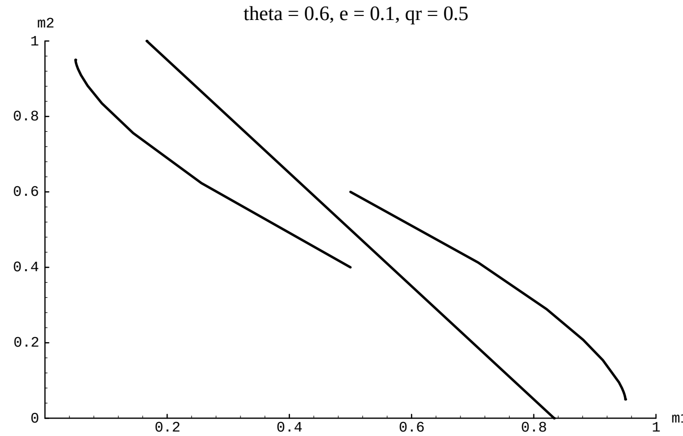

In both cases the curves in the lower part of the boxes are moving up and to the left, and those in the upper boxes are moving down and to the right.

The asypmtotic behavior is interesting, and hinted at by the analysis of the short run behavior. Suppose first that $\theta < 1/2$. The sinks near $(n_1, 0)$ and $(0, n_2)$ are on the "wrong" side of the line. When $\epsilon = 0$ it is easy to see that states (n_1, n_2) where $n_2 > 1/2(1-\theta)$ and $n_2 < (1-\theta)^{-1}(1/2-\theta)$ are transient, and all other states are recurrent. The following picture on the left shows the invariant distribution when $n1 = 8$ and $n2 = 24$. The probability of the mode is about 0.11.

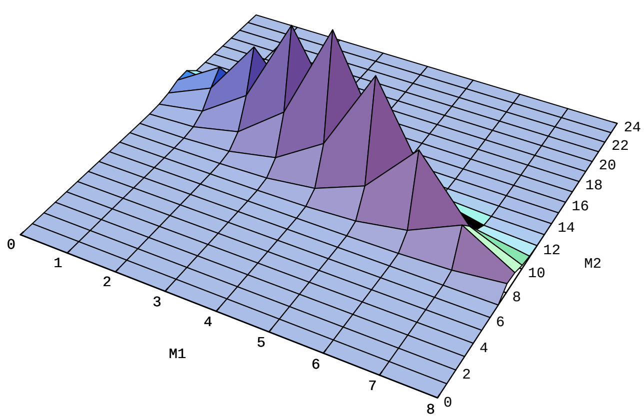

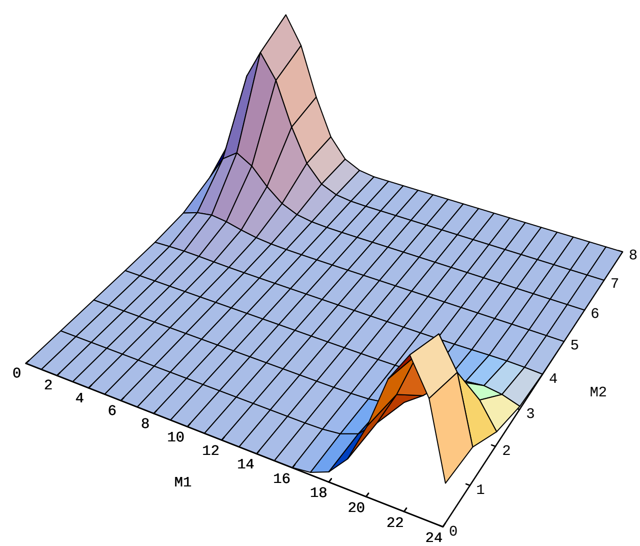

When $\theta > 1/2$ the two sinks are absorbing states for the $\epsilon = 0$ process, and neither is favored by selection. To see this, note that the minimum number of trembles necessary to move from $(n_1, 0)$ to $(0, n_2)$ is the same as the number necessary to move from $(0, n_2)$ back to $(n_1, 0)$. This is illustrated in the picture on the right above, for $\epsilon = 0.1$. It is easy to introduce asymmetries that will lead to a selection result simply by putting in a bias in favor of one color or the other. This has the effect of shifting the line to the left or the right, increasing the number of trembles needed to move in one direction and decreasing the number of trembles needed to move in the other.

Matsuyama has also varied the interaction technology by allowing players to meet other players of the same and the other type at different rates. This increases the possibilities for interesting short-run dynamics. For some specifications of the model, fashion cycles arise. That is, the differential equations describing the large-numbers short-run behavior of the system have cycles as solutions. This illustrates once again how details of the interaction technology effect the qualitative performance of the system.

5.3.2 Keynsian Coordination Failure

The study of strategic complementarities has become important throughout economics. Macro, industrial organization, development economics are only some of the fields where problems posed in terms of strategic complementarities are drawing significant attention. Arthur's (1988) contribution to the first Santa Fe volume on the economy as a complex system surveys some of the dynamic phenomena that can arise from systems exhibiting strategic complementarities.

In the macroeconomic literature on Keynsian coordination failure, the existence of strategic complementarities has been claimed as a source of cyclic economic activity. The argument, crudely summarized from a number of papers, is that for one reason or another a given model will have multiple equilibria, and therefore cyclic economic activity can be understood as a drift from one equilibrium to the next. Cooper and John (1988) were the first to recognize the common analytic theme in this literature is that the existence of a strategic complementarity creates a coordination problem. Different actions are self-reinforcing, so it is possible for players in equilibrium to "coordinate" on one action or another. Coordinating on a suboptimal action is a "coordination failure".

In fact one principle lesson of modern economic analysis is that there is a big leap between static equilibrium and the dynamic process whose long-run behavior the static equilibrium is supposed to represent. The mere existence of multiple equilibria due to strategic complementarities does not guarantee the existence of interesting economic dynamics. So Blume (1994b) investigates a population game with a strategic complementarity to see if it can generate cyclic investment behavior. The model is in the spirit of Diamond's (1982) search equilibrium paper, where the strategic complementarity is due to an externality from market participation. When more traders participate in a market, the returns to market participation are higher.

Imagine, then, a population of N potential investors. Those who choose not to invest receive a fixed payoff flow. Those who invest in a project receive an expected payoff flow whose value depends upon the number of other investment projects currently active. The net return function is $W(M)$. Details of the market organization determine $W(M)$, but one can write down a number of models where $W(M)$ is initially increasing in M due to strategic complementarities in market participation. If there is an M^* such that $W(M) < 0$ for $1 \leq M < M^*$ and $W(M^*) > 0$, there will be at least two Nash equilibria of the market participation game, one where no one participates and one where some number $M \geq M^*$ invest. Decision opportunities arrive for each player at the jump times of a Poisson process, and the process for each player is independent of all other processes. When a non-participant has a decision opportunity she may choose to enter the market, while an investor with a decision opportunity may choose to leave the market. For the moment, suppose that the scrap value of an investment project is 0, and that there are no fixed costs to investment. Finally, suppose that decisions are based on the current return to investment; that is, investors' discount rates are sufficiently high relative to the arrival rate of investment decisions that their investment decision is determined by the return to investment given the current size of the market. Randomness in decisionmaking is a consequence of shocks to expectations. I assume that the resultant random utility choice model has a logit representation. In other words, the probability of investing or not (disinvesting or continuing to invest) at a decision opportunity is given by the log-linear choice rule (2.2).

This model tells a very Keynsian story of the trade cycle. Keynes viewed investment as perhaps interest-inelastic, but very responsive to expectations, and made much of the "instability" of the marginal efficiency of investment on this account. One consequence of this expectations-driven investment analysis is that different kinds of expectations, optimistic and pessimistic, can support different levels of investment of which only one corresponds to full employment. In the present model shocks to expectations may cumulate to drive investment from one basin of attraction to the next.

The investment model departs from the previous two-player game models in that the net gain to participation is not linear in the number of market participants. This nonlinearity has no consequences for the analytical paradigm outlined above. The process describing the number of active investors is a birth-death process, and its properties can be examined by studying the birth and death rates. In particular, the asymptotic behavior of the economy is characterized by a potential function. Let $P(0) = 0$ and $P(M) = P(M-1) + W(M)$. A computation from the detailed balance conditions (2.4) shows that the invariant distribution satisfies the relationship

$$\log \frac{\rho(M)}{\rho(0)} = C(M, N) + \frac{1}{\epsilon} P(M)$$

where $1/\epsilon$ is the proportionality constant in the log-linear choice rule.

When ϵ is small, the iid investment shocks are small. The invariant distribution will be concentrated near the global maximum of the potential function. The model is illustrated in the two graphs below.

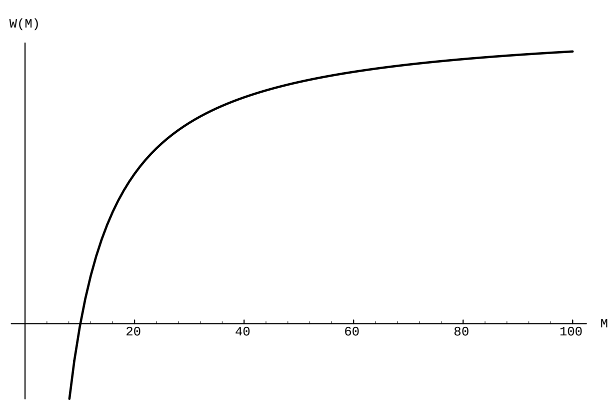

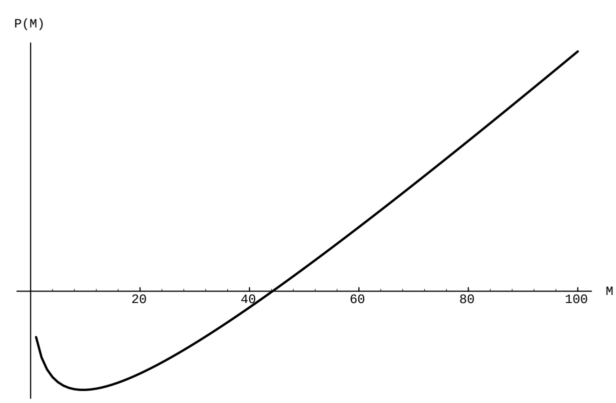

On the left is the expected return function $W(M)$, and the potential function $P(M)$ is on the right. The expected return $W(M)$ exceeds 0 for all $M > 10$. Thus viewing the model as a static investment model, if $N > 10$ there are two equilibrium states: A low-investment equilibrium at $M = 0$ and a high-investment equilibrium at $M = N$. If the iid expectations shocks are not large, the economy will spend a large fraction of time near the equilibrium of greatest potential value. The potential always has a local maximum at $M = 0$. If $N <= 10$ this is the only local maximum, and indeed there is only one equilibrium, with no active investors. When $N > 10$ there is a local maximum at $M = N$, corresponding to the high-investment equilibrium. But for $N < 45$ the global maximum is at $M = 0$, and so the economy never really gets off the ground. investment may occasionally spurt up, but ultimately it falls back to the ground state $M = 0$, and stays near there most of the time. For $N > 45$ the global maximum moves to $M = N$, and high investment levels predominate.

The model can be elaborated in several ways. Consider adding a fixed cost to investment. With a fixed cost, investments that generate small positive returns will not be undertaken, but active investors in such a state will not (on average) close their investments out. Consequently new equilibria are created, just where $W(M)$ goes positive; in this case at $M = 11$, $M = 12$, etc, depending on the magnitude of the fixed cost. These new equilibria have no significant dynamic consequences. They are not local maxima of the appropriate potential function, and so they can never be favored as ϵ becomes small.

This model exhibits two interesting properties. First, the high productivity equilibrium can only be sustained in the long run if the pool of potential investors is sufficiently large. If the economy has not achieved a critical mass, attempts to jump-start the economy by short-run programs to stimulate investment are doomed to fail. The second interesting property speaks directly to the original purpose of the model. The mere existence of multiple equilibria is not enough to conclude that significant cyclic activity will arise. The "cyclic" behavior of the model when expectation shocks are small is nothing like that observed in the actual economy. Dynamics here look as follows: The economy spend a negligible amount of time away from the equilibrium states, and it spends vastly more time in one equilibrium state than the other. Even without going through a calibration exercise it is clear that this description of dynamics does not correspond with any current picture of business cycle activity.

6. Conclusion

Two views of the invisible hand appear repeatedly in economic thought: Adam Smith's positive and benign invisible hand leads the merchant "to promote an end which is no part of his intention"; an end which represents the "interests of society". Austrian economists, most notably Menger and Hayek, see that institutions themselves evolve as a result of the flux of individual actions. "What Adam Smith and even those of his followers who have most successfully developed political economy can actually be charged with is . . . their defective understanding of the unintentionally created social institutions and their significance for economy."[3] There is no argument that the "spontaneous order" of the Austrians is necesarily beneficent. The invisible hand could equally well be Shakespeare's bloody and invisible hand of night as Adam Smith's hand of Pangloss.

Unfortunately it is hard to separate the notion of spontaneous order from that of a beneficient invisible hand. Robert Nozick (1974) coined the phrase "invisible hand explanation" to refer to exactly the spontaneous generation of order of which the Austrians wrote. "An invisible-hand explanation explains what looks to be the product of someone's intentional design, as not being brought about by anyone's intentions."[4] This explanation suggests some sort of optimization. But his leading example of invisible-hand explanations are evolutionary explanations of the traits of organisms and populations. Few biologists claim that evolution is a global optimization device.

The population games literature clarifies how the spontaneous evolution of order can arise as the unintended social consequence of many private actions. The resultant social order is frequently consistent with some kind of optimization hypothesis — the optimization of a potential. But the optimization of a potential does not imply optimization of social welfare. As coordination games demonstrate, the potential maximum may be the Pareto inferior equilibrium.

Invisible hand phenomena of the Austrian variety abound in social systems. Although population games are a useful tool for modelling invisible hand phenomena, more work is needed to get the most from them. A variety of methods exist for uncovering stochastically stable sets, but less is known about the asymptotic behavior of these models when ϵ is not small — when significant residual noise remains in the system. Little is known about the sample-path behavior of these processes. I argued above that the mean of the first-passage time distribution is not informative about short run behavior. Finally, application of these models to problems more complex than the choice of driving on the left or the right will require the development of computational techniques for calibration.

Further conceptual development is also needed. Population games are viewed as a branch of evolutionary game theory. This is true, to the extent that evolutionary game theory is that which is not eductive. But there is no real evolution going on in these models. The global environment established by the game and the interaction model is fixed, and players are equipped with a fixed set of rules. Some rules are better than others, and the better rules tend to be adopted. Consequently the local environment, determined by the rules in play, changes. The population game analysis tracks the interaction between individual play and the local environment. The appropriate biological analogy is an ecological model. In population games there is no innovation. The set of rules is fixed at the start. More important, there is no variation in the global environment. It would seem that an appropriate topic of a truly evolutionary game theory is evolution in the rules of the game. Such evolution is the ultimate source of the spontaneous order to which invisible hand explanations refer.

---

[3] Menger (1883), p 172.

[4] Nozick (1974), p 19.

REFERENCES

Bergin, J. and B. Lipman (1994) "Evolution with state-dependent mutations," unpublished, Queens University.

Binmore, K., L. Samuelson and R. Vaughn (1993) "Musical Chairs: noisy equilibrium selection," unpublished, University of Wisconsin.

Blume, L. (1993) "The statistical mechanics of strategic interaction," *Games and Economic Behavior,* 5, 387–426.

———— (1993b) "The statistical mechanics of best response strategy revision," Games and Economic Behavior, forthcoming, Economics Working Paper Archive ewp-game/9307001.

———— (1994) "How noise matters," unpublished, Cornell University, Economics Working Paper Archive ewp-game/9410001.

———— (1994b) "An evolutionary analysis of Keynsian coordination failure," unpublished, Cornell University.

———— (1995) "Evolutionary equilibrium with forward-looking players," unpublished, Cornell University.

Canning, D. (1992) "Average behavior in learning models," *Journal of Economic Theory,* 57, 442–472.

Cooper, R. and A. John (1988) "Coordinating coordination failures in Keynsian models," *Quarterly Journal of Economics,* 103, 441–63.

Diamond, P. (1982) "Aggregate demand management in search equilibrium," *Journal of Political Economy,* 90, 881–94.

Evans, R. (1993) "Observability, imitation and cooperation in the repeated prisoner's dilemma," unpublished, Cambridge Unibersity.

Ellison, G. (1993) "Learning, local interaction, and coordination," *Econometrica,* 61, 1047–1072.

———— (1995) "Basins of attraction and long run equilibria," unpublished, MIT.

Ethier, S. and T. Kurtz (1986) *Markov Processes,* New York: John Wiley and Sons.

Foster, D. and H. P. Young (1990) "Stochastic evolutionary game dynamics," *Theoretical Population Biology,* 38, 219–32.

Fudenberg, D. and C. Harris (1992) "Evolutionary dynamics with aggregate shocks," *Journal of Economic Theory,* 57, 420–441.

Kandori, M., G. Mailath and R. Rob (1993) "Learning, mutation, and long run equilibria in games," *Econometrica,* 61, 29–56.

Kandori, M. and R. Rob (1995) "Evolution of equilibria in the long run: A general theory and applications," *Journal of Economic Theory,* 65, 383–414.

Kiyotaki, N. and R. Wright (1989) "On money as a medium of exchange," *Journal of Political Economy,* 92, 927–54.

Matsuyama, K. (1992) "Custom versus fashion: Path-dependence and limit cycles in a random matching game," unpublished, Hoover Institution, Stanford University.

Menger C. (1883) *Untersuchungen über die Methode der Soczialwissenschaften und der Politischen Oekonomie inbesondere.* Leipzig: Duncker & Humboldt; English Translation, *investigations into the method of the social sciences with special reference to economics.* New York: New York University Press, 1985.

———— (1892) "On the origin of money," *Economic Journal,* 2, 239–55.

Monderer, D. and L. Shapley (1993) "Potential games," unpublished, The Technion.

Noldeke, G. and L. Samuelson (1994) "Learning to signal in markets," unpublished, University of Wisconsin, Economics Working Paper Archive ewp-game/9410001.

Nash, J. (1950) "*Non-cooperative games,*" unpublished Ph.D. thesis, Mathematics Department, Princeton University.

Nozick, R. (1974) *Anarchy, State, and Utopia,* New York: Basic Books.

Samuelson, L. (1994) "Stochastic stability in games with alternative best replies," *Journal of Economic Theory,* 64, 35–65.

Schelling, T. (1971) "Dynamic models of segregation," *Journal of Mathematical Sociology,* 1, 143–186.

Thurstone, L. (1927) "A law of comparative judgement," *Psychological Review,* 34, 273–286.

Weibull, J. (1994) "The mass-action interpretation of Nash equilibrium," unpublished, The Industrial Institute for Economic and Social Research, Stockholm.

Young, H. P. (1993) "The evolution of conventions," *Econometrica,* 61, 57–84.

———— (1993b) "An evolutionary model of bargaining," *Journal of Economic Theory,* 59, 145–168.