THE STATISTICAL MECHANICS OF BEST-RESPONSE STRATEGY REVISION

Lawrence E. Blume

March 1993
Revised January 1994

Version: 24 Jan 1994

The Statistical Mechanics of Strategic Interaction

Abstract

I continue the study, begun in Blume (1993), of stochastic strategy revision processes in large player populations where the range of interaction between players is small. Each player interacts directly with only a finite set of neighbors, but any two players indirectly interact through a finite chain of direct interactions. The purpose of this paper is to compare local strategic interaction with global strategic interaction when players update their choice according to the (myopic) best-response rule. I show that randomizing the order in which players update their strategic choice is sufficient to achieve coordination on the risk-dominant strategy in symmetric 2×2 coordination games. The "persistant randomness" which is necessary to achieve similar coordination when the range of interaction is global is replaced by spatial variation in choice in the initial condition when strategic interactions are local. An extension of the risk-dominance idea gives the same convergence result for $K \times K$ games with strategic complementarities. Similar results for $K \times K$ pure coordination games and potential games are also presented.

JEL Classification: C78

Correspondent:

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

Current address:

Professor Lawrence Blume
Eitan Berglas School of Economics
Naftali Building, 7th floor
Tel Aviv University
Ramat Aviv, Tel Aviv 69978
Israel
email: LB19@CORNELL.EDU

1. Introduction

Evolutionary game theory has tried to understand equilibrium in non-cooperative games as the outcome of a dynamic process of adaptation or evolution in a population of players. Work in this area has tried to identify some particular deterministic population processes, such as replicator dynamics, and study their behavior. Recently several scholars (Blume (1993), Canning (1990), Foster and Young (1990), Kandori, Mailath and Robb (1993), and Young (1993)) have observed that the qualitative behavior of population dynamics is strikingly different when stochastic perturbations are introduced to a deterministic adjustment process. In symmetric two-by-two coordination games, there is a sense in which behavior converges to that equilibrium which has the largest basin of attraction with respect to divergence by opponents from the equilibrium prescription. This is Harsanyi and Selten's *risk-dominant equilibrium.* The robustness of this result depends upon the way in which randomness is introduced, but the result is common over a rather large class of models. A typical statement is that the stochastic process of the population fraction selecting the risk-dominant strategy is ergodic, and as the noise term becomes small, the limit distribution converges to point mass at 1.

The introduction of persistant randomness is typically motivated either by appeal to fluctuations in the populations of players predisposed to play one way or another, or to bounded rationality in a fixed population of individual players, or to players' experimenting as they learn how to play. But whatever the source, the randomness must persist in order for the limit result to hold.[1] The interpretation of these results is not so straightforward. Even with very little noise, it will be the case that, infinitely often, the fraction of the population choosing the risk-dominated strategy will be arbitrarily large. Of course the same statement can be made for the risk-dominant strategy. The limit result merely reflects the fact that relatively more time will be spent near risk-dominance than away from it as the noise term shrinks. But the population's behavior never really settles down.

In this paper I will explore an alternative dynamic; that of *best-response strategy revision* (without stochastic perturbations) where each player interacts not with the entire player population but only with her neighbors, and players will always choose a best-reply to their immediate environment. I will compare local and global strategic interaction in best-response strategy revision in 2×2 and $K \times K$ coordination games, games with strategic complementarities and potential games. Although the processes studied here are not ergodic, we will see a strong tendency to settle at the risk-dominant equilibrium in 2×2 coordination games. Extensions of this criterion identify likely rest-points in the other game types.

This paper continues my study of local strategic interaction begun in Blume (1993). In

[1] In fact, the results in Blume (1993) still hold if the noise damps out at a sufficiently slow rate, but the analysis fails if the noise disappears quickly enough. This will also be true for global strategic interaction with a finite player population.

models with local interaction, each player interacts with only a few "neighboring players" rather than the entire population of players. However each player is indirectly connected to a every other player in the population through a chain of neighborhood relations. Thus it is possible for a particular strategic choice to propagate from some favorable locale through the entire population. The justification for studying economic models with local interaction is self-evident. Many economic institutions, such as markets for retail services and markets where transportation costs, are characterized by local interaction. In biological models local interaction can be viewed as a model of viscosity. The term "viscosity" was introduced by Hamilton (1964) to describe the phenomenon common to many organisms that most of their interactions with other organisms of the same species will be with members of their kin group. Since an organism's kin will be more likely to share the same characteristics than will the population at large, a consequence of viscosity is that organisms will tend to interact more with other organisms that are like themselves. In a game-theoretic approach to viscosity, Myerson, Pollack and Swinkels (1992) have modelled viscosity as an increased likelihood that a player in a symmetric game will be matched with another member of the player population playing the same strategy as he is. In the model studied here, members of a large player population interact only with a small fixed set of players. A different approach to viscosity in a dynamic model of strategic interaction is taken by Oechssler (1993). Bergstrom (1993) provides an evolutionary argument based on kin selection in favor of a criterion that looks much like risk-dominance in 2×2 games. Anderlini and Ianni (1993) compare some different dynamics for local strategic interaction. Deterministic dynamic models of strategy revision with local interaction have been studied by Berninghaus and Schwalbe (1993). Finally, Ellison (1993) has compared convergence rates of local and global interaction models with stochastic perturbations.

The class of strategy revision processes studied in this paper is introduced in the next section, and section 3 provides some general discussion of their invariant measures and asymptotic behavior. Two-by-two coordination games is the subject of section 4. The case of "nearest neighbor interaction" is worked out in some detail, and some general results are also shown. Global and local interaction are compared, and strategy revision processes with mixtures of local and global interaction are studied. Some simulation results are also presented to illustrate the results. Finally, rates of convergence are for local and global interaction are compared; it is shown that convergence is faster with global interaction than with local interaction. Section 5 takes up, in order, $K \times K$ coordination games, games with strategic complementarities and potential games. The more opaque proofs are left to section 6.

2. The Model

The model of strategic interaction presented here is that of Blume (1993), but specialized to best-response strategy revision. Each site on the 2-dimensional integer lattice Z^2 is the address of one player. There is a symmetric neighborhood relation $\sim$, where $s \sim t$ means that s and t are neighbors. The neighbors of site 0 are a set $V_0 = \{s \ : \ s \sim 0\}$. The neighborhood relationship is translation-invariant: $s \sim t$ if and only if $t - s \in V_0$. (Notice that translation invariance of the neighborhood relation and its symmetry imply that V_0 is symmetric to reflection through the origin.) The set of neighbors of s is denoted V_s. I will represent the neighborhood relation by a graph whose vertices are the sites in Z^2 and whose edges connect neighbors. At some points I will take the player set to be $B(N)^2 = [-N, N)^2$, the cube with side-length 2N centered at the origin in Z^2.

Players choose actions from the set $W = \{0, \ldots, K\}$. A configuration η is a map $\eta : Z^2 \to W$ which specifies what action each player is using. When the player set is $B(N)^2$, a configuration will be a map $\eta : B(N)^2 \to W$. The play of player s in configuration ϕ is $\phi(s)$, and the play of all players other than s is denoted $\phi(-s)$.

A player who has chosen an action w receives a payoff flow from each of his neighbors determined by w and by each neighbor's choice of action. She receives instantaneous payoff $G(w, v)$ from a given neighbor if she plays action w while that neighbor plays action v. Her instantaneous payoff from playing strategy w is the sum of the instantaneous payoffs received from playing w against each of her neighbors. The total payoff flow to player s from playing $w \in W$ when the play of the population is described by the configuration ϕ is

$$\sum_{t \in V_s} G\big(w, \phi(t)\big).$$

A *best-response strategy revision process* is a continuous time Markov process on the space of configurations which describes the evolution of players' choices through time. The process works as follows. All players have i.i.d. Poisson "alarm clocks". At randomly chosen moments (exponentially distributed with mean 1) a given player's alarm goes off. When it does, she responds to her neighbors' current configuration by choosing an action which maximizes her instantaneous payoff $\sum_{t \in V_s} G\big(w, \phi(t)\big)$, where ϕ is the current configuration of the population's play. When there is more than one best response, the player will draw from an equiprobable distribution on the set of best responses. However, this situation will not arise in this paper. I assume throughout the paper that *for all configurations* ϕ *there is a unique action* $\mathrm{br}(s, \phi)$ *which maximizes* $\sum_{t \in V_s} G\big(w, \phi(t)\big)$. On each $B(N)$ there is an open and dense set of matrices for which hypothesis will be true for all possible neighborhood specifications, and a residual set of matrices for which this is true for all finite neighborhoods V_0 in the plane Z^2. The kind of choice behavior described above is boundedly rational in several different ways. For discussion, see Blume (1993) and Kandori, Mailath and Robb (1993).

Configurations in which each player chooses a best response to the play of her neighbors are of special importance.

Definition 2.1: A configuration ϕ is a *Nash Configuration* if, for all s, $\phi(s) = \mathrm{br}(s, \phi)$.

Imagine a one-shot, simultaneous-move game wherein the payoff to player s of choosing

w when the play of all the other players is described by ϕ is $\sum_{t \in V_s} G\big(w, \phi(t)\big)$. A Nash configuration is a pure-strategy Nash equilibrium of this "lattice game".

The state space for all strategy revision processes is the Borel space X of configurations with the product discrete topology. The best-response strategy revision process is a Markov process on X. The formalities of constructing the process can be found in Blume (1993). On the finite box $B(N)^2$ the process is a Markov jump process which changes state at rate $4N^2$.

3. Invariant Distributions

The analysis of Markov jump processes on finite state spaces is not hard because they jump at discrete points in time. Let τ_i denote the time of the ith jump. If $X(t)$ is the (continuous time) Markov process, then $X(\tau_i)$ is a (discrete-time) Markov chain, sometimes referred to as the *embedded chain.* The invariant distributions and ergodic behavior of $X(\tau_i)$ are precisely that of $X(t)$.

Let P denote the transition matrix for the embedded chain:

$$P(\phi, \eta) = \frac{1}{4N^2} \#\{s : \phi(-s) = \eta(-s) \quad \text{and} \quad \phi(s) = \text{br}(s, \eta)\}$$

A distribution μ on configurations is invariant for this process if and only if for all η, $\mu(\eta) = \sum_{\phi} P(\eta, \phi)\mu(\phi)$. Consequently, μ is an invariant distribution for best-response strategy revision in the box $B(N)$ if and only if

$$\mu(\eta) = \frac{1}{4N^2} \sum_{\{s:\eta(s)=\text{br}(s,\eta)\}} \mu(\eta(-s))\,. \tag{3.1}$$

Notice that configurations in which no sites best-respond to the play of their neighbors have probability 0 under every invariant distribution. Thus configurations in which each player chooses a strongly dominated strategy are 0-probability configurations. Inductively it can be seen that any invariant distribution must concentrate all its mass on those configurations wherein each player chooses an iteratively undominated strategy.

As is usual with finite Markov Chains, the state space can be partitioned into the transient states and a finite number of ergodic sets. The ergodic distributions are the extreme points of the convex set of invariant measures, and there is a one-to-one correspondence between ergodic distributions and ergodic sets; the ith ergodic distribution puts all its mass on the ith ergodic set. The set of transient states is probability 0 under any invariant distribution. For example, consider the symmetric game with payoff matrix

$$\begin{bmatrix} 3 & 4 & 2 & 1 & 1 \\ 2 & 3 & 4 & 1 & 1 \\ 4 & 2 & 3 & 1 & 1 \\ 1 & 1 & 1 & 2 & 1 \\ 1 & 1 & 1 & 0 & 0 \end{bmatrix}$$

The top three-by-three submatrix is the payoff matrix for "Rock, Scissors Paper". Strategy 4 is a symmetric Nash equilibrium strategy, and strategy 5 is strongly dominated. The set $\{5\}$ is transient, and the ergodic sets are $\{1,2,3\}$ and $\{4\}$. The distribution μ_1 that assigns probability 1 to the Nash configuration $\phi(s) \equiv 4$ is invariant, and there is a unique invariant distribution μ_0 with support $\{1,2,3\}$. The invariant distributions for the best-response strategy revision process are the distributions $\mu_\alpha = \alpha\mu_1 + (1-\alpha)\mu_0$. In the games considered in this paper, the ergodic sets will be single Nash configurations, and interest centers on identifying the basins of attraction of each set.

The necessary and sufficient condition of equation (3.1) is not easy to manipulate. But similar computations give other interesting necessary conditions. One such is the following: If μ is invariant, then for each $w \in W$,

$$\mu\big(\phi(s) = w\big) = \mu\big(\mathrm{br}(s,\phi) = w\big)\,.$$

The probability of observing choice w by player s is precisely the probability that w is a best response to the play of s's neighbors. The two events whose probabilities are being measured on the right- and left-hand sides respectively, are identical for Nash configurations. More generally, for any set of players P and configuration $\psi(P)$ of the play of players in P,

$$\sum_{s\in P} \mu\big(\phi(P-s) = \psi(P-s)\ \&\ \phi(s) = \psi(s)\big) - \mu\big(\phi(P-s) = \psi(P-s)\ \&\ \psi(s) = \mathrm{br}(s,\phi)\big)\,.$$

Expressed in terms of conditional probabilities, this is

$$\sum_{s\in P} \Big(\mu\big(\phi(s) = \psi(s)\,|\,\phi(P-s) = \psi(P-s)\big) - \mu\big(\psi(s) = \mathrm{br}(s,\phi)\,|\,\phi(P-s) = \psi(P-s)\big)\Big)\cdot \mu\big(\phi(P-s) = \psi(P-s)\big) = 0\,.$$

Another technique for computing ergodic distributions is to apply the graph-theoretic techniques from the literature on stochastic perturbations of dynamical systems, introduced to game theorists by Foster and Young (1990), to the problem. Although best-response strategy revision has no stochastic perurbations, the restriction of the embedded chain to any ergodic class is a process to which the discrete-time developments of Kandori, Mailath and Robb (1993) and Young (1993) can be applied. It turns out that for the ergodic distribution on the i'th ergodic class X_i of configurations, the probability $\mu_i(\eta)$ of a given state $\eta \in X_i$ is proportional to the number of η-trees in which every transition has positive probability. (This is because in each η-tree, the product of the transitions is either 0 or $|X_i|^{1-|X_i|}$.) However I have found no way to actually count thes trees

In dynamic strategy revision models with persistent trembles, the focus of the analysis has been to find ergodic distributions. The ergodic theory of best-response strategy revision processes is typically hopeless — without trembles, there are many absorbing states. The goal of the analysis is to identify when particular absorbing states will be reached. It is generally true that processes with local interaction have more absorbing states then do processes with global interaction. When the box size N is large enough, the only Nash configurations of best-response strategy revision processes with global interaction are "symmetric configurations", where each player makes the same choice.[2] Any symmetric configuration which is Nash under global strategy revision is also Nash under local strategy revision, but when strategic interaction is local, asymmetric Nash configurations may exist as well.

4. 2×2 Coordination Games

The payoff matrix G describes a (symmetric) coordination game. That is, if

$$G = \begin{bmatrix} w & x \\ y & z \end{bmatrix} \qquad w > y \text{ and } z > x.$$

then $w > x$ and $z > y$. Throughout, the "top" strategy, noted strategy 0, will be *risk-dominant.* That is, $w - y > z - x$. In a two-person symmetric coordination game, strategy 0 will be risk-dominant if the probability a player assigns to the event that his opponent is playing strategy 0 which is necessary to leave that player indifferent between strategy 0 and 1 is less than 1/2. Let p denote that probability. Specifically

$$p = \frac{z - x}{w - y + z - x}.$$

Strategy 0 is risk-dominant if and only if $p < 1/2$, which is true if and only if $z - x < w - y$. I will refer to this threshold probability p as the *cut-off* for the payoff matrix G. *Throughout the remainder of this section I will assume that 0 is risk-dominant.*

4.1. Nearest Neighbor Interactions

First I take up the special case of l_1-nearest neighbor interactions on the plane and on the $4N^2$ lattice points in the square $B(N) = [-N, N)^2$. The neighbors of a given player are

[2] To see this, suppose that each player saw exactly the same distribution of play. Then, due to the unique best-response hypothesis, each player would choose exactly the same thing. Now a priori we do not know that each player sees the same thing, because no player counts himself in computing the distribution of his opponents' play. But if N is large, the beliefs of any two players are very close regardless of the configuration, and a consequence of the unique best response hypothesis again is that they will choose the same thing.

those players immediately above and below her, and immediately to her left and right. The cases $p < 1/4$ and $1/4 \leq p < 1/2$ require separate treatment. When the cut-off probability p is less than 1/4, a player will adopt strategy 1 if and only if all of her neighbors have chosen 1. When the cut-off probability p exceeds 1/4, a player will adopt strategy 1 if and only if at least three of her four neighbors play strategy 1.

When $p < 1/4$, the players *coordinate.* That is, either everyone ultimately plays 0 or everyone ultimately plays 1.

Theorem 4.1: For the best response strategy revision processes on $B(N)$ and Z^2 with nearest neighbor interaction on the plane and cut-off $p < 1/4$:

1. On the infinite lattice Z^2 or in any finite box $B(N)$, the best-response strategy revision process coordinates.
2. A sufficient condition for almost sure coordination on 0 is that the initial configuration has two neighboring players whose initial choice is strategy 0.
3. If the initial configuration on $B(N)$ has only one site s playing 0, then the probability of coordinating on 0 is 4/5.

Conclusion 2 states a strong growth property. If a pair of directly interacting players are playing strategy 0, they can never be induced to switch. Furthermore, their neighbors will switch to 0 when they have a chance, then their neighbors, and so on until the entire plane (or finite box) has switched to 0's. The conclusions of the theorem remain true for processes on d-dimensional lattices, but with critical probability $1/2^d$, and the probability of coordinating from a single site becomes $2^d/(2^d + 1)$

The interpretation of coordination in $B(N)$ differs somewhat from that of coordinating in Z^2. In both cases, "coordinating at 0" means that the limit distribution is point mass at the configuration $\phi(s) \equiv 0$. On $B(N)$, this configuration is actually reached in finite time with probability 1. But on Z^2 this will not be the case. If the initial configuration contains only a finite number of players playing 0, then even when the process coordinates at 0, the configuration at each moment of time contains only a finite number of players choosing 0.

Proof: Let ϕ^0 denote the configuration in which all players choose 0, and ϕ^1 the configuration in which all players choose 1. Let δ^i denote the point mass on the configuration ϕ^i.

Consider an initial configuration ϕ which has at least two neighboring players choosing 0. Neither of these players will ever switch to 1 because each has at least one neighboring 0, which is enough to make 0 a best response. Denote by A the collection of players such that if $s \in A$, then $\phi(s) = 0$ and there is a $t \in V_s$ such that $\phi(t) = 0$. No player in A will ever leave A. If $A = B(N)$ (or Z^2), the configuration ϕ is stable (pointmass at ϕ is invariant). At any revision opportunity for player s, all of her neighbors are choosing 0, and so she will continue to choose 0. If A does not include all the players, choose $s \notin A$ but with at least one neighbor in A. Then there is a time $\tau < \infty$ such that for all $t > \tau$, $s \in A$; s will switch to 0 at her next revision opportunity. Finally, if s is any player not in A, there is a finite chain of players $s_0, s_1, \ldots, s_K$ such that $s_0 \in A$, $s_K = s$ and $s_k \in V_{s_{k+1}}$. Let τ_1 denote the time of the first revision opportunity for s_1, and τ_k the time of the first revision opportunity for s_k after τ_{k-1}. Then the a.s.finite time $\tau_1 + \ldots + \tau_K$ is an upper bound on the time at which s will adopt strategy 0. Hence starting from any initial configuration with at least two adjacent 0's, the time t distribution of states converges (weakly in the case of Z^2 — the functions whose value is determined by the choice of play for a finite number of players are dense in the space of continuous functions on the space of configurations) to δ^0.

Now consider an initial configuration which has only one player s choosing 0 and all remaining players choosing 1. Let τ denote the time of the first revision opportunity of s or one of ss neighbors. Players outside $V_s \cup \{s\}$ are choosing 0 and all their neighbors are choosing 0, so $\phi_t = \phi_0$ for all $t < \tau$. With probability 4/5 the revision opportunity at τ will belong to some $t \in V_s$. The player t has a neighbor playing 0, and so she too will switch to 0. The resulting configuration ϕ_τ has two adjacent 0's, and so the process will converge to δ^0. With probability 1/5 s chooses first. She switches to 1. The distribution of states is now δ^1, which is invariant.

Now consider an initial configuration with $k < \infty$ isolated 0's. With probability $(1/5)^k$ all the 0's disappear and the process has reached invariant distribution δ^1, and with probability $1 - (1/5)^k$ a pair of 0's is created and the process converges to δ^0.

Finally, if the player set is Z^2 and the initial configuration contains an infinite number of isolated 0's, the creation of an adjacent pair of 0's somewhere in finite time is a probability 1 event, so the process converges weakly to δ^0. □

The updating rule for $1/2 > p \geq 1/4$ is different from that of the case $p < 1/4$. Our canonical player in the interior of the box will choose 0 if at least two of her neighbors choose 0. Players at the vertices of $B(N)$ will adopt strategy 1 if and only if both neighbors are playing 1, so for these sites, the appropriate rule is to choose r_{st} to be the minimum among the two neighbors. For players on the boundary but not at the corners, the right updating rule depends upon p. If $p > 1/3$ then the interior updating rule applies, while if $p < 1/3$ the vertex updating rule applies.

When $p \geq 1/4$ the population need no longer coordinate on a single strategy. Nonetheless, the best-response strategy revision process settles down into one of a number of absorbing states. Thus it is meaningful to talk about a limit configuration, and each player's limit choice. Furthermore, the geometry of limit configurations can be described.

Theorem 4.2: For the best-response strategy revision processes on $B(N)$ with nearest neighbor interaction on the plane and cut-off $1/4 \leq p < 1/2$,

1. A configuration is absorbing for the best response strategy revision process is 0 precisely on the union of a finite number of rectangles contained in $B(N)$ such that any two disjoint rectangles are at l_1-distance at least 3 from each other. The minimum dimension of any rectangle is 2, unless $p < 2/3$ in which case a 1-by-x or x-by-1 rectangle (with $x \geq 2$) can exist along the boundary.
2. An absorbing configuration is almost-surely reached from any initial configuration.

Proof: The initial problem is to describe all the absorbing states for the strategy revision process, and show that the process almost surely converges to one of them. The intuition is rather clear. Consider the process of 0's and 1's. Any square of 0's is stable. Any isolated 0 can be made to disappear. In a situation like that pictured below, the top 1 can be made to switch to a 0 (and subsequently, the 1 below it).

$$\begin{matrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{matrix}$$

Thus, in the interior of $B(N)$, stable sets of 0's are rectangles of side-length at least 2. Such sets are also stable on the boundary. In addition, any single row of 0's is stable if $p > 2/3$, but not if $p < 2/3$. If two disjoint rectangles are distance 2 apart, there is a player in between who is playing 1. But she has two neighbors, one from each rectangle, playing 0, and so she will switch to 0 at her next revision opportunity.

How do we know that an absorbing state must be reached? Let τ_i denote the time of the i'th alarm-clock ring for the threshold process. Then the process of configurations $\{\phi_{\tau_i}\}$ is a discrete-time Markov chain. Starting from any initial configuration of this process, we will see that there is a transition to an absorbing states in no more than $4N^2$ steps, with probability $1/(4N^2)!$, which is independent of the initial configuration. If this is true, it follows from the Borel-Cantelli lemma that the process must reach such a state infinitely often. (In considering rates of convergence it may be comforting to know that these estimates are really crude).

Here is how to reach an absorbing state. For a given initial configuration, ring the alarm clocks of all interior players choosing 1 who are adjacent to two or more 0's. Each player chosen will switch to 0, and perhaps add more players to the list of 1's adjacent to at least two 0's. Also ring the clocks of all boundary players who could switch to 0. Continue ringing until all such players are exhausted. Now ring the clocks of all players choosing 0 who are sufficiently isolated that they would switch to 1. When this is finished, the 0-1 configuration is stable. Notice that no site was visited twice, so at most $4N^2$ transitions have been made. $\square$

Best-response strategy revision with $p \geq 1/4$ does not have the stochastic growth property of best-response strategy revision with $p < 1/4$. When $p \geq 1/4$, a region of 0's in a sea of 1's will grow or shrink into a (possibly empty) rectangle, and the rectangle will be stable.

4.2. General Results on Local Interaction

When the radius of interaction is taken to be larger than 1, the character of the analysis for $p < 1/4$ remains the same. For neighborhoods which are l_1-spheres of radius k, squares of players with side-length at least $k+1$ cannot be made to disappear. Furthermore, any such square must grow. (Start from a player adjacent to a player at the center of a side. She must switch. Then either one of the two players next to her and adjacent to the side will switch ....) The following theorem is proved by a straightforward rewriting of the proof of Theorem 4.1.

Theorem 4.3: For the best-response strategy revision processes on $B(N)$ with interaction radius k and cut-off $p < 1/4$:

1. On the infinite lattice Z^2 or in any finite box $B(N)$, the best-response strategy revision process coordinates.
2. A sufficient condition for coordinating at 0 is that the initial configuration has a $(k+1) \times (k+1)$ box of players whose initial choice is strategy 0.

Surprisingly, this case of "stochastic growth" is the general case. This is the content of the following Theorem, which is the main result on coordination games. Here "general case" means that neighborhoods are any radially symmetric shape and sufficiently large. For a given finite set B of players, let ϕ^B denote the configuration which is identically 0 on B and 1 on B^c.

Definition 4.1: A finite set B of players is *stable* if, for all $s \in B$, $\Pr\{\phi_{t+u}(s) = 0 \,|\, \phi_t = \phi^B\} = 1$. A strategy revision process with parameter p *grows* from a finite set of players B if,for all $s \in \mathbf{Z^2}$, $\lim_t \Pr\{\phi_t(s) = 0 \,|\, \phi_0 = \phi^B\} = 1$.

Theorem 4.4: : Let C denote a radially symmetric (through 0) subset of $\mathbf{R^2}$ with positive Lebesgue measure and such that the elements of $C \cap \mathbf{Z^2}$ span $\mathbf{R^2}$. Let λ parametrize best-response strategy revision processes with neighborhood $V_0 = \lambda C \cap \mathbf{Z^2}$.

a) For all $\lambda \geq 1$, if $p \leq 1/2$ there exist finite stable sets.

b) For all $\lambda \geq 1$ there is a finite set B of players and a threshold $p(\lambda) < 1/2$ such that growth from B takes place for all $p < p(\lambda)$.

c) $\lim_{\lambda \to \infty} p(\lambda) = 1/2$.

d) For every fixed set B of players initially choosing 0, the critical threshold $p_B(\lambda)$ below which growth from this set takes place converges to 0 as λ increases.

The idea of the theorem is to fix a basic neighborhood shape and then make the neighborhoods large enough that the bumps due to the discreteness of the lattice are not too important. The Theorem then states that 0 will grow from some finite set of players. The theorem is proven for the player set $\mathbf{Z^2}$ to emphasize that the size of the neighborhoods can be arbitrarily small relative to the population size, and that the size of the initial region of 0's can also be arbitrarily small. Of course the conclusions hold for large enough boxes $B(N)$, and the critical probability $p(\lambda)$ is independent of N for N large enough.

In higher dimensions the arguments used to show Theorem 4.4 also show that for all $p < 1/2$ there is a $\lambda' > 0$ such that for all $\lambda > \lambda'$ there is a B such that if $V_0 = \lambda C$, then growth from B takes place. Conclusion *d*) holds as well. I do not have a proof for *a*) but I would be very surprised if it were false.

4.3. Global Interaction and Mixed Regimes

This section records the results for the more commonly studied case of global interaction, and also investigates the robustness of both local and global analysis with respect to mixing the regimes. The model presented in the last Section with players in the box $B(N)$ and an interaction radius so large that the set of neighbors for each player is all the remaining players in $B(N)$ is nearly the model studied by Kandori, Mailath and Robb (1993) but done in continuous time. This model is exactly equivalent to the discrete time model in which at each time t one player is chosen at random who then revises her choice according to the best-response strategy revision rule. The only information needed from the configuration for evaluating best-response strategy revision is the frequency of 0's and 1's. For purposes of comparison to local best-response strategy revision it is convenient to draw initial configurations randomly and ask after the probability of observing a sample path of the process which converges to the configuration $\phi \equiv 0$. I will assume that the initial configuration of the population is assigned by an independent draw for each player from a distribution with a probability θ of realizing strategy 0.

For any given cutoff p, a player will choose 0 when she has the opportunity if and only if at least $(4N^2 - 1)p$ other players have chosen 0. (Actually, we should be careful about the treatment of "ties" in the payoffs of the two strategies, but they do not really matter for what follows.) The process *coordinates.* That is, the process converges to a configuration where all players are choosing the same strategy. For given θ the probability that the population coordinates at 0 is the probability that the empirical frequency of 0's in a draw from the binomial distribution $b(4N^2, \theta)$ is at least p. As N gets large, this converges to 0 or 1 depending upon whether θ is less than or greater than p.

Theorem 4.5: For global best-response strategy revision on the plane Z^2 and given cut-off $p < 1/2$,

1. In the box $B(N)$ the process coordinates.
2. The process coordinates at 0 if the fraction of players initially choosing 0 exceeds p.

Comparing Theorem 4.4 with Theorem 4.5, we see that local best-response and global best-response strategy revision processes in coordination games are very different. With large enough neighborhoods in sufficiently large boxes, an initial small region of 0's will take over. For example, calculations show that if the radius of interaction is 5 and $p < .42$, an initial disk of 0's of radius 14 will grow and take over any given area. This disk contains 636 players. Making N, the size of the box, as large as we want, we see that there are initial configurations wherein the initial fraction of 0's is vanishingly small, $(26/N)^2$, which ultimately coordinate at 0 with local best-response. But with global best-response, the initial fraction of 0's must be at least p. Since the behavior of the two processes is qualitatively different, it is important to examine mixed regimes where both local and global interaction is present in order to test the robustness of the results.

There are many different ways of parametrizing mixed regimes. One simple way is to suppose that when the frequency of 0's in the entire population is $\hat{p}_G$ and the frequency of 0's in the neighborhood of player s is $\hat{p}_s$, player s best-responds to the distribution $\epsilon\hat{p}_G + (1-\epsilon)\hat{p}_s$. Global interaction corresponds to $\epsilon = 1$, and local interaction corresponds to $\epsilon = 0$. This corresponds to each player taking a weighted average of the play across all sites, where players in V_s are accorded weight $((1-\epsilon)/|V_s|) + \epsilon/(4N^2-1)$ and players not in V_s have weight $\epsilon/(4N^2-1)$.[3]

What one would most like to know about is the nature of the regime change as ϵ varies from 0 to 1. A first step is understanding the robustness of the local and global results. Do the conclusions of Theorems 4.5 and 4.4 hold up when ϵ is near 0 and 1, respectively.

Theorem 4.6: For mixed best-response strategy revision in the box $B(N)$, and given cut-off $p < 1/2$,

1. if $\epsilon < (2p-1)/2p$ the conclusion of Theorem 4.4 still holds: There is a k_0 and an r_0, both independent of N, such that if the interaction radius exceeds k_0 and if the initial configuration contains a (Euclidean) sphere of radius $r > r_0$ all playing 0's, then the process coordinates at 0.
2. For mixed best-response strategy revision in the box $B(N)$, and given cut-off $p < 1/2$, if the empirical frequency of 0's is $\hat{\theta} < p$ and if $\epsilon > (1-p)/(1-\hat{\theta})$, then the process almost surely coordinates at 1.

Proof: The first claim is a trivial consequence of Theorem 4.4. If $\epsilon < 1-2p$ then there is a p' between 1/2 and p such that if fraction p' of the players in the neighborhood V_s choose 1, then player s will switch to 1 regardless of the configuration $\phi(B(N)/V_s)$ outside her neighborhood.

[3] A generalization of this idea is to assume that the effects of interaction are distance-weighted. Here identical weights correspond to global interaction, and weights that are 1 for players inside V_s and 0 for players outside V_s corresponds to local interaction. The qualitative features of the effects of parameter variation here are more complicated, but not different in character from the simpler parametrization with ϵ.

The second claim follows from the fact that if ϵ, p and $\hat{\theta}$ are configured as indicated, then even if V_s contains only 0's, the average will have strategy 1 preferred to strategy 0. □

4.4. Simulations

The analysis of best-response strategy revision developed above derives conclusions based on the structure of the initial configuration. This analysis is nicely illustrated by simulations of the process on a large box. If the simulations required a separate run for each initial condition there would be no practical way to run them and no useful way to present the results. Fortunately it is possible to "couple" together into one simulation runs for many different initial configurations, and to display the results simultaneously in one picture.

The initial configurations for the simulations presented here have players within a given distance from the center of the box playing 0, and players farther away choosing 1. That is, $\phi_0(s) = 0$ if and only if $|s| < 1$. For a given specification of the neighborhoods V_s and threshold probability p, all of the different initial configurations can be simultaneously simulated. To see how, consider nearest neighbor strategy revision with $p < 0.25$. Suppose that all players within Euclidean distance k of the origin are initially playing 0, and the rest play 1. Consider the first switching opportunity, and suppose that it occurs for player (i, j) who is initially playing 1. He will continue to play 1 only if all of his neighbors are playing 1; that is, only if all of his neighbors are at least distance k from the origin. If, on the other hand, he is initially playing 0, he will continue to play 0 if and only if his neighbor nearest to the origin is playing 0. Now consider the next switch. The rules will be the same, unless the first player is in the second player's neighborhood. So here the rule of play is governed by the distances of the second player's neighbors from the origin and the distance from the origin of the neighbor of the first player nearest to the origin. This suggests the following scheme. Label each player with his distance from the origin. When a revision opportunity arises, relabel the chosen player with the lowest label from among his neighbors. So the first player chosen will be assigned the lowest label of his neighbors. The second player chosen will be assigned the lowest label of his neighbors. if the first player is not a neighbor of the second player, then the second player will be assigned the label of his neighbor nearest to the origin. If the first player is a neighbor of the second, the second player will be assigned the minimum distance to the origin from among his neighbors and the first player's neighbors, and so forth. Ultimately a labelling will be reached from which no further changes are possible. These are the data presented in the tables and figures. To interpret them, suppose we wanted to see what would have happened if the initial configuration had assigned 0's only to those players within Euclidean distance 5 from the origin. Map the labels into 0's and 1's by assigning a 0 to those sites with label less than or equal to five, and 1's to all remaining sites. To see what happened had the initial Euclidean ball of 0's had radius 10, assign 0's to all players with labels less than or equal to 10, and 1's to the rest. For larger interaction neighborhoods, the idea is the same: Order the neighborhood's sites by their labels and choose the pth fractile of the distribution.

The results of some simulations are reported in Table 1. Table 1 answers the question, "How big an initial neighborhood of 0's is required for the population of a 41×41 box to coordinate at 0?" for best-response strategy revision on a 41-by-41 box.

Table 1 reports the size of the initial Euclidean ball of 0's necessary to achieve coordination. The top number is the square of the radius of the Euclidean ball, and the bottom number is the percentage of players in the 41-by-41 box who are in this region. Asterisks indicate that no meaningful coordination was attained. Obviously in any finite box coordination is attained at some level, for instance, if the initial sphere is so large that it contains the entire box. The test of "meaningful coordination" employed here is that the the region of 0's is thicker than the neighborhood radius. Another useful test is to check if the simulation results are insensitive to box size. The table entries marked with double daggers changed only slightly when simulated in a 61-by-61 box. The table entries marked with a single dagger changed more dramatically with the increase in box size.

Table 1 illustrates the conclusions of Theorem 4.4. For any given threshold p, coordination at 0 is ultimately realized for some sufficiently large initial region of 0's. As the threshold probability p increases, the size of the required initial region of 0's grows. Reading down columns demonstrates part d of Theorem 4.4. The initial region of 0's necessary to sustain growth grows with the size of the neighborhood.

The asterisk entries in Table 1 are those combinations of radii and cutoff probabilities for which no meaningful coordination takes place. The picture of these simulations is similar to the situation described by Theorem 4.2. The limit configuration of radii form a nexted sequence of polygons, each only 1 or 2 units wider than the one preceding, which indicate convergence to distinct stable sets starting from the different initial conditions. The radius of the outermost stable set is very nearly half the side length of the box.

The data from Table 1 shows that as the cutoff probability increases, the percentage of players inthe initial 0 set grows at an ever increasing rate. However the growth curves become more nearly linear as as the radius of the interaction neighborhood increases. Notice that even when the neighborhood radius is large, the fraction of initial 0's is significantly less than the threshold probability. Table 2 also shows the effect of changing the neighborhood radius on the size of the set of initial 0's necessary to achieve coordination at 0. Again the qualitative picture is that the differential effects of various cutoff probabilities are most evident at larger radii.

Table 2 demonstrates the effects of different box sizes on processes with interaction radius 10 and three different cutoff probabilities: 0.325, 0.375, and 0.425. The table indicates both the size of the initial region of 0's (measured as the square of the Euclidean radius) necessary to achieve coordination, and the (smaller sizes) at which distinct stable sets emerge.

I have not reported any variations in these region sizes. This is for good reason.

Within limits, there is and should be very little. Suppose an initial region of 0's is given from which the set of 0's cannot shrink. Then it is easy to see that the limit configuration is independent of the order in which the different players have strategy revision opportunities. To see this, consider two distinct processes, process 1 and process 2, with identical initial conditions, and run them at the same time. Ultimately both will reach a limit configuration. The paths in configuration space for both processes have the property that the set of players playing 0 at time t is a subset of the set of players playing 0 at time $t' > t$. Consider the first player s_1 who switched from 1 to 0 in process 1. The neighbors playing 0 at the time of the switch are just members of the initial region of 0's, so they are 0's in the process 2 limit configuration. Hence if this player had a strategy revision opportunity in process 2, an almost certain event, she switched to 0. Suppose this switch took place at time t_1. Now consider s_2, the second player who switched to 0 in process 1. The set of players choosing 0 at the time she switched was the 0-set of the initial configuration and s_1. These players are all playing 0 in process 2 after time t_1. So if s_2 has a revision opportunity after time t_1, an almost certain event, then she will switch to 0. Proceeding inductively shows that the limit 0 set for process 2 contains the limit 0 set for process 1. Interchanging the roles of the two processes shows that the limit configurations must be identical.

4.5. Rates of Convergence

In comparing the local and global results, the obvious question is why is the basin of attraction for coordination at the risk-dominant strategy so much larger with local interactions than it is with global interaction. Intuitively, emergence of risk dominance with either local or global strategic interaction depends on the existence of an initial "critical density" of players choosing the risk-dominant strategy. But global interaction requires that this critical density be achieved globally; that is, the critical density must be achieved throughout the entire player space. Local interaction only requires that the critical density be achieved in some small locality in the player space. If this critical density is achieved anywhere, the neighborhood "fills in", so that localy the density of the risk-dominant strategy goes to 1. When the local density is sufficiently high, the critical density is achieved for neighborhoods of players which are just outside the initial locality. These players switch to the risk-dominant strategy, and it is in this manner that the risk-dominant strategy spreads.

The picture described by the proof of Theorem 4.4 has the boundary between 0's and 1's moving like a travelling wave. The ball of 0's spreads until it covers the entire box or plane. This contrasts with the evolution of states in global strategy revision. With global best-response strategy revision each player makes her final choice at the first strategy revision opportunity. At every subsequent revision opportunity she will decline to change her mind; her first choice is still a best response. This means that convergence is very fast. Let τ_g^N denote the first time that coordination is achieved in the box $B(N)$.

Theorem 4.7: Suppose that the initial fraction θ of the players playing strategy 0 exceeds p. Then $\tau_g^N/2\log N \to 1$ almost surely.

Theorem 4.7 shows that the waiting time until coordination at 0 in the box $B(N)$ grows like $\log N$. The story for local interaction is very different. Suppose that p is such that the set of players choosing 0 grows from a finite stable set of players B as described by Theorem 4.4. Let ϕ^B denote the configuration in which all players in the set B play 0 and all other players play 1. All revisions will be from 1 to 0. A player once having chosen 0 will never revert. For initial configurations in this class the best-response strategy revision is a growth process. The following Theorem can be proven in the same manner that Durrett (1988) proves the same theorem for a proportional stochastic growth model.

Theorem 4.8: Suppose that p is such that 0 grows in a given local best-response strategy revision process from an initial configuration ϕ^B (with B stable). There exists a closed, convex set C in $\mathbf{R^2}$ such that for all $\epsilon > 0$

$$\Pr\big((1-\epsilon)tC \subset \{s : \phi_t(s) = 0\} \subset (1+\epsilon)tC\big) \to 1 \quad \text{as} \quad t \to \infty.$$

The time it will take the process to coordinate at 0 in a large enough box $B(N)$ is bounded by the time it takes the sets $(1 \pm \epsilon)tC$ to cover the box, which grows like N. Letting τ_l^N denote the time at which coordination at 0 first occurs in the box $B(N)$ we see that

$$0 < \lim_{N\to\infty} \frac{1}{N}\tau_l^N < \infty \quad \text{a.s.}$$

Thus coordination at 0 takes longer to achieve with local interaction than it does with global interaction. This result may seem counterintuitive to those who have read Ellison (1993). Ellison shows that with stochastic perturbation, the convergence rate for local best-response strategy revision is much shorter than for global best-response strategy revision. Both Ellison's results and the rate results presented here have their origin in the same phenomenon, the fast rate of convergence of best-response strategy revision. When mutations are present, the convergence rates are governed by the time spent waiting for mutations. Suppose a stochastically perturbed best-response strategy revision process with infrequent deviations from best-response is at a state near coordination at 1. A certain number of mutations need to occur before 0 becomes a best-response. If the mutation rate is small, the waiting time for these mutations will be much greater than the time it takes the best response dynamic to drive the process near to coordination at 0. Ellison's observation is that, when interaction is local, it takes fewer mutations to make 0 a best response in some region of the player set from which it can grow. With global interaction, the waiting time for best-response to "take over" is much longer than it is with local interaction. Here I have shown that when best-response takes over, it works quicker with global interaction than with local interaction, but this is second-order to the differences in time spent waiting for enough mutations in Ellison's calculations.

5. K × K Games

The analysis presented so far considers only the very special case of two-by-two coordination games. That analysis has shown how viscosity drives the best-response dynamic to achieve a particular equilibrium configuration. On a finite box, the strategy revision process looks like a discrete-time Markov chain on a finite state space, and so it will have a limit distribution. In considering more general classes of games, two questions need to be addressed. When is this limit distribution connected to Nash equilibrium in some interesting way? Secnd, for those games in which the limit distribution is concentrated on Nash configurations, when is an equilibrium selection principle operative, and what is it?

Section 3 provides a general characterization of invariant distributions, but does not directly address either question. There are, however, three interesting classes of games for which the invariant distributions are concentrated on Nash configurations: K-by-K coordination games, potential games and games with strategic complementarities.

5.1. K-By-K Pure Coordination Games

The theory of 2-by-2 coordination games carries over in a straightforward way to K-by-K *pure coordination games.* A pure coordination game is a symmetric game in which all non-diagonal payoffs are 0, and diagonal payoffs are positive. For these games each strategy has its own cutoff p_i, the minimum probability with which i must appear in an opponent's play in order to induce a player to choose i regardless of the play of others. In this case, payoff dominance and risk dominance coincide — the smallest cutoff goes to the strategy with the highest payoff. Furthermore, this cutoff will be less than or equal to 1/2.

Suppose that the first strategy, strategy 0, has the lowest p_i. Then the conclusions of Theorem 4.4 still hold, an can be proved by exactly the same arguments. If the neighborhood size is large enough to smooth out discreteness effects, and if the initial disk of 0's is sufficiently large, then almost half the neighbors of any player near to this disk will be playing 0, and so that player too will switch to 0. More generally, consider any two-person game G such that $(0,0)$ is a Nash equilibrium. Now consider the lattice game and the payoff matrix G_α such that $G_\alpha)(i,j) = G(i,j)$ for $(i,j) \neq (0,0)$, and $G_\alpha(0,0) = G(0,0) + \alpha$. Then for α sufficiently large the conclusions of Theorem 4.4 are true.

5.2. Games with Strategic Complementarities

Games with strategic complementarities have proven to be important for a number of applications. A (symmetric) game G has strategic complementarities if for all strategy choices $0 \leq i < j \leq |W|$, $G(j,k) - G(i,k)$ is non-decreasing in k. I rely on the fact that for such games the best response correspondence is monotonic with respect to stochastic dominance. Let μ_1 and μ_2 be two probability distributions on W such that for all $w \in W$, $\mu_1\{v : v \geq w\} \leq \mu_2\{v : v \geq w\}$. Then μ_2 stochastically dominates (first order) μ_1, and the best responses to μ_2 are no smaller than the best responses to μ_1.

Theorem 5.1: The ergodic distributions of the best-response strategy revision process on $B(N)$ in a game of strategic complementarities are precisely the point masses on Nash configurations. For every initial distribution of states μ_0 there is a μ_∞ such that $\mu_t \to \mu_\infty$.

Proof: I will show that, starting from any configuration, a Nash configuration must be reached in finite time with probability 1. The convergence to a limit distribution follows from this fact, since the ergodic classes are singletons.

Order the players in the box in some given way, choose an initial configuration, and consider the following sequence of events. First, cycle through the players in order, selecting those players in order from among those players whose best response to their neighbors' play is less than their current play. Because of the monotonicity of best responses with respect to stochastic dominance , any player moving up may cause his neighbor to desire to move up, but never to move down. Now cycle through again in the same fashion, and continue to do so until no player wishes to move down any further. (The number of revisions is bounded by the number of players times the number of strategies.) Now cycle through the players in order, selecting those players who wish to move up. (Again monotonicity guarantees that any one player moving up will cause other players to move up, but not down. Continue cycling through until no one wishes to move up. (Again the number of necessary revisions is bounded by the number of players times the number of strategies.) The resulting configuration is a Nash configuration: No one wishes to move down, and no one wishes to move up. Since the number of player selections is uniformly bounded (independently of the starting configuration), this sequence of events must occur with a positive probability which can be positively bounded from below independently of the initial configuration. Conclude from the Borel-Cantelli lemma that it will happen sometime, and so an equilibrium will be reached. □

One consequence of this fact is that for stochastically perturbed best-response strategy revision, such as that in Blume (1993) and Kandori, Mailath and Robb (1993), if the perturbations are small the limit distribution of the process (which is unique because of the perturbations) puts most of its weight on the set of Nash configurations.

In games with strategic complementarities there is a Nash configuration in which all players choose the smallest iteratively undominated strategy, and a Nash configuration in which all players choose the largest iteratively undominated strategy. When the strategic complementarities are strict, all pure strategy Nash equilibria of the two-person game are symmetric, and these are the only Nash configurations that can arise with global strategic interaction. But when interaction is local, other asymmetric Nash configurations can arise which do not reflect any mixed Nash equilibrium of the two-person game. Consider the payoff matrix

$$\begin{bmatrix} 0 & 0 & 0 \\ -1 & 0 & 2 \\ -3 & 0 & 3 \end{bmatrix}$$

and a nearest-neighbor best-response strategy revision process on the interval $[0, 1, \ldots, N]$ in Z^1. Configurations such as $[0, \ldots, 0, 1, 2 \ldots, 2, 1, 0, \ldots]$ are Nash configurations, because 1 is the best response to the mixture of 0 with probability 1/2 and 2 with probability 1/2, 0 is the best response to the mixture of 0 with probability 1/2 and 1 with probability 1/2, and

2 is the best response to the mixture of 1 with probability 1/2 and 2 with probability 1/2. The only Nash equilibria for the two-person game are the pure strategy combinations $(0,0)$ and $(2,2)$. Thus if the neighborhood size is sufficiently large the only Nash configurations are $\phi(s) \equiv 0$ and $\phi(s) \equiv 2$.

This example illustrates the claim of section 3 that processes with local strategic interaction have more absorbing states than do processes in which the interaction is global. The following Theorem gives some information on where different configurations go. For configurations ϕ and η write $\phi \leq \eta$ if $\phi(s) \leq \eta(s)$ for all $s \in B(N)$.

Theorem 5.2: Let $\{\phi_t\}_{t\geq 0}$ and $\{\eta_t\}_{t\geq 0}$ denote two best-response strategy revision processes for a payoff matrix G with strategic complementarities. Let μ_t^ϕ and μ_t^η denote their respective time-t marginal distributions on the state space. If $\mu_0^\phi \leq \mu_0^\eta$, then μ_t^η first-order stochastically dominates μ_t^ϕ for all t.

This Theorem compares the paths of two processes starting from comparable initial conditions. If the η process initially has higher choices for each player than does the ϕ process, then at all subsequent dates, the probability of high choices by each player under the η process exceeds the probability of high choices under the ϕ process.[4]

Theorem 5.2 is a specific case of a more general Theorem whose statement and proof is in the last section. The conclusion is independent of the interaction structure, and is also robust to the addition of state-independent trembles.

The monotonicity of the best response correspondence with respect to stochastic dominance suggests that all players should move up or down together. The last theorem states that if the play at two sites is initially uncorrelated or positively correlated, the correlation cannot reverse over time. Formally, a distribution μ is said to have *positive correlations* if $\int fg\, d\mu \geq \int f\, d\mu \int g\, d\mu$ for all non-decreasing functions f and g. Suppose that a distribution of states in the best-response strategy revision process had positive correlations. Take $f(\eta)$ to be the indicator function for the event "s plays a strategy greater than or equal to v" and $g(\eta)$ to be the indicator function for the event "t plays a strategy greater than or equal to w". These functions are increasing. Applying the inequality, we get that $\mu(\phi(s) \geq v \mid \phi(t) \geq w) \geq \mu(\phi(s) \geq v)$, the statement that t choosing high does not decrease the probability that s will choose high.

Theorem 5.3: Suppose that $\{\phi_t\}_{t\geq 0}$ is a best-response strategy revision process for a payoff matrix with strategic complementarities. If the initial distribution of configura-

[4] Recall the definition of first-order stochastic dominance. Suppose the state-space is partially ordered. The distribution μ dominates ν if for every function f which is non-decreasing with respect to the order, $\int f\, d\nu \leq \int f\, d\mu$. In our setting, indicator functions which are 1 if player s chooses a strategy v at least as big as some given w and 0 otherwise are non-decreasing with respect to the partial order.

tions has positive correlations, then for all t the marginal distribution of ϕ_t has positive correlations.

Notice that positive correlations for μ_0 includes those distributions such that the random variables $\phi_0(s)$ for $s \in B(N)$ are mutually independent. This theorem too is a corollary of the general result on comparing two strategy revision processes, whose statement and proof can be found in the last section.

Just as in the analysis of pure coordination games, we would like to know when a local concentration of a particular choice is stable and when it can grow. The stochastic-dominance monotonicity of best-response correspondences allows for a sufficient variety of behaviors that it is not always the case that a single configuration will be singled out. The selection behavior observable in 2×2 games may not be present in larger games. Nonetheless, an extension of the risk-dominance idea does identify a sufficient condition for selecting a particular choice as an equilibrium. The idea is this: To test for convergence to w^*, first consider all pairs of strategies $v < w$ where $v < w^*$. Look at the mixed strategies that put weight only on v and w, and let $p(v, w)$ denote the minimum weight on w for the best response to be greater than v. Next consider all pairs $w < v$, where $v < w^*$, and let $p(v, w)$ denote the minimum weight on w which induces a best response less than v. Let p denote the maximum of the $p(v, w)$. The essence of the next theorem is that if p is less than $1/2$, then for sufficiently thick neighborhoods a process starting with a sufficiently large region of players choosing w^* will converge to everyone playing w^*. Just as in Theorem 4.4, the "sufficiently's" are all independent of the size of the player set providing it is large enough. The following theorem is for best-response strategy revision processes on $\mathbf{Z^2}$.

This condition generalizes the risk-dominance criterion of 2×2 games. In 2×2 games it *is* risk dominance, for in this case moving higher (or lower) means adopting the other strategy.

Formally, define $p(v, w)$ as follows: For all strategies $v < w$ such that $v < w^*$, let $p(v, w)$ denote the smallest q such that the best response to the $1 - q : q$ mixture of v and w exceeds v. For all strategies $v > w$ such that $v > w^*$, let $p(v, w)$ denote the smallest q such that the best response to the $1 - q : q$ mixture of v and w is exceeded by v. Following the definition of section 4, w^* *grows* from a finite set B of players in a strategy revision process if $\lim_t \Pr\{\phi_t(s) = w^* \mid \phi_0(B) = w^*\} = 1$ for all $s \in \mathbf{Z^2}$.

Theorem 5.4: Let C denote a radially symmetric subset of $\mathbf{R^2}$ with positive Lebesgue measure, and such that the elements of $C \cap \mathbf{Z^2}$ span $\mathbf{R^2}$. Let λ parametrize the best-response strategy revision processes with neighborhood $V_0 = \lambda C \cap \mathbf{Z^2}$. Then for all $\lambda > 1$ there is a finite set B of players and a threshold $p(\lambda) < 1/2$ such that w^* grows from B if $\max_{v<w} p(v, w) < p(\lambda)$, and $\lim_{\lambda \to \infty} p(\lambda) = 1/2$.

The proof of theorem 5.4 is a straightforward extension of the proof of theorem 4.4. The condition on best-responses is the proper measure of the basin of attraction, and the monotonicity of best-responses guaranteed by the strategic complementarity hypothesis ensures that inspecting a relatively small set of best-responses suffices. The conclusion of the Theorem implies the convergence of the distribution of ϕ_t in the weak convergence topology to point mass on the configuration in which all players choose w^*.

5.3. Potential Games

Let G be the payoff matrix for a two-player symmetric game. A potential function for the game G is a function $\Pi : W \times W \to \mathbf{R}$ such that

$$G(w_1, v) - G(w_2, v) = \Pi(w_1, v) - \Pi(w_2, v) = \Pi(v, w_1) - \Pi(v, w_2).$$

This strange-looking condition is exactly the specialization of Monderer and Shapley's (1993) definition of an *exact potential* to symmetric games. Simple calculations show that, without loss of generality, we can take

$$\Pi(v, w) = h(v) + h(w) + Q(v, w)$$

where Q is symmetric and $\sum_w Q(v, w) = \sum_w Q(w, v) = 0$ for all v, and $\sum_v h(v) = 0$.

Suppose now that a box $B(N)$ and neighborhoods V_s are given. Then

$$P(\phi) = 2 \sum_{s \in B(N)} |V_s| h\big(\phi(s)\big) + \sum_{s \in B(N)} \sum_{t \in V_s} Q\big(\phi(s), \phi(t)\big)$$

is a potential for the lattice game. It is easy to see that if ϕ and η are two configurations that differ only at player s, then $P(\phi) - P(\eta)$ measures the utility gain to player s from $\phi(s)$ over $\eta(s)$ in each of s's interactions.

Call a configuration ϕ a local maximum of the potential function P if changing the value of any one coordinate does not increase the potential. The local maxima of P are the Nash configurations. Since there are only a finite number of configurations, maxima, and hence pure strategy Nash equilibria, exist. Best-response strategy revision processes are stochastic hill-climbing processes; each step increases the value of the potential. They are not stochastic in the sense that sometimes the process slips down the hill, but stochastic in its choice of local improvement. In summary:

Theorem 5.5: Let G be a symmetric game, and suppose that G has a potential. Then every lattice game on $B(N)$ with neighborhoods $\{V_s\}_{s \in B(N)}$ is a potential game with potential P, and every such game has at least one Nash configuration. Furthermore the ergodic distributions are precisely the point masses on the Nash configurations.

Of course the global maxima of the potential function are appealing objects. They are not necessarily payoff dominant Nash configurations, however. Two-by-two coordination games are potential games, and the global maximum of the potential function has each player choosing the risk-dominant strategy. Blume (1993) examined 2×2 games on Z^d. There the potential function was used to introduce stochastic perturbations, and it was shown that as the perturbations become small, the limit probability on the set of "global maximizers" converges to 1.[5] On $B(N)$ it will also be the case that the introduction of small trembles pushes the strategy revision process towards the global maxima. In fact, if the perturbations are damped out over time at a sufficiently slow rate, the limit probability on the set of global maximizers is 1. This result is the topic of the simulated annealing literature, and a detailed description of the rate (including rate constants) is available. An interesting project (thankfully not attempted here) is to compare the critical damping rate with various tremble processes that allow for experimentation as part of a learning process.

The following example shows that local maxima which are not global maxima can be stable. The game with payoff matrix

$$\begin{bmatrix} 1 & 0 & -7 \\ 0 & 3 & 0 \\ -7 & 0 & 4 \end{bmatrix}$$

is a potential game. The potential has $h(0) = -4/3$, $h(1) = 5/3$ and $h(2) = -1/3$, and the matrix

$$Q = \frac{1}{3} \begin{bmatrix} 13 & 1 & -14 \\ 1 & 1 & -2 \\ -14 & -2 & 16 \end{bmatrix}$$

Now consider the lattice game on the one-dimensional interval $[0, \dots, N]$ with nearest neighbor interaction. Some checking establishes that the potential-maximizing configuration is $\phi(s) \equiv 2$. But any configuration of the form $\eta_n(s) = 0$ for $s < n$, $\eta(s) = 2$ for $s > n$, and $\eta(n) = 1$ is a local maximum for $1 < n < N - 1$. Although the configurations η_n and ϕ agree on the set of players $n + 1, \dots, N$, it takes two switches to move from η_n to a configuration with higher potential. An analysis of potential games along the lines suggested by the analysis of games with strategic complementarities is beyond the scope of this paper, but it is possible to derive conditions for potential games analogous to the risk-dominance conditions which guarantee convergence to a global maximum of the potential function.

I will conclude with the following conjecture about the general principle at work in processes with local interaction. It seems that, in large populations with local interaction, local variation in the configuration of play has the same role as stochastic perturbations do in processes with global interaction. The realization of a fortuitous sequence of mutations

[5] Global maximizers is in quotes here because, strictly speaking, the concept is not well-defined since the potential function is given by an infinite series that need not converge. See Blume (1993) for details.

is akin to the realization of a region of propitious play. This relationship between viscosity and random perturbations is suggested by the analysis in presented in this paper, but much hard analysis remains to be done to explore the generality of this conclusion.

6. Proofs

This section collects those proofs that were too tedious or unenlightening to include in the main text.

Proof of Theorem 4.4: *a*) It suffices to prove this statement for $p = 1/2$, since if a set C is stable for a given $p = p'$, it is clearly stable for $p < p'$. Fix λ. The corresponding neighborhoods V_s will all be radially symmetric. A consequence of radial symmetry is that any line which bisects V_s and goes through s will have at least half the points of V_s on each side (including points on the line).

Here is a recipe for constructing a convex polygon which is stable at $p = 1/2$. Consider the neighborhood V_0, and let B_0 denote the smallest square which strictly contains it (no points in V_0 are on the boundary of the square). This square has a side length s. Choose a point in the neighborhood V_0 and the corresponding ray r_1 through the point to the origin. Truncate r_1 so that its endpoint is a lattice point and so that the minimum of the absolute values of the coordinates of the endpoint is at least $2s$. Call this line l_1. Choose a point x_1 in the lattice, and translate l_1 so that its origin is x_1. Call the other endpoint x_2. Now imagine the neighborhood of x_2. Extend a ray from x_2 at an angle of π radians to the line l_1. The line and the ray bisect the neighborhood, and each half (counting the boundary, so there is overlap) has at least half the points in the neighborhood. If this ray is rotated clockwise the a sufficiently small amount, the number of neighborhood vertices contained in the convex cone formed by the two lines will be exactly half the number of points on the neighborhood. Rotate the ray clockwise the maximal amount so that the convex cone formed by the ray with the line 1_1 contains at least half the points in the neighborhood. In other words, rotate the ray until it intersects another vertex in the neighborhood. Call this ray r_2. Truncate r_2 according to the same rules used to truncate l_1, call this line l_2 and its new endpoint x_3. The radial symmetry of V_0 implies that, continuing in this manner to construct lines l_3, l_4, and so forth, the construction will eventually return to x_1. That is, there is a line l_N which has endpoints x_{N-1} and x_1. To see this, boserve that, starting from r_1, the ray is always rotated clockwise until it hits the next point. By radial symmetry, at some point in this procedure the aggregate angle of rotation must be π radians. Radial symmetry now implies that the continuation of the procedure will proceed as follows: The next rotation will be through the same angle as the first rotation, the one following will be through the same angle as the second rotation, and so forth. Eventually π additional radians will have been traversed, and the process will have returned to its starting point.

The area enclosed by these N lines is a stable polygon. To see this, observe that the neighborhood of every point in the polygon intersects the boundary along at most two edges (due to the size of the side lengths). The neighborhood of a point will intersect the polygon in the least number of points at a vertex, and the constructions guarantees that at every vertex x, the intersection of the polygon and the neighborhood V_x contains at least half the points of V_x.

b) Consider the stable polygon P constructed in *a*). For every vertex s within distance 1 of P let $p(s) = \#V_s \cap P/\#Vs$. Then for any $p \leq \min_{s\in P} p(s)$ the process grows. Again it is clear that if growth occurs for p, then it occurs for all $p' < p$.

c) We will rescale the set λC by looking at C and the lattice $(1/\lambda)\mathbf{Z}^2$. Choose p_1 and p_2 such that $1/2 > p_1 > p_2 > p$. First, due to radial symmetry, any hyperplane through the origin divides C into two pieces of equal area. There is a $\delta > 0$ such that any hyperplane coming within $\delta > 0$ of the origin contains at least fraction p_1 of the volume of C on each side.

Let $S(y, r)$ denote a disk of radius r centered at y. If r_0 is sufficiently large, than for any y and $r \geq r_0$, if $S(y, r)$ comes within distance δ of the origin, then the volume of $C \cap S(y, r)$ is at least p_2. Choose a y and an r.

Let $B(s, l)$ denote the square whose sides are parallel to the x- and y-axes with side length l and center s. Let $H(C) = \{s \in \lambda^{-1}\mathbf{Z}^2 : B(s, \lambda^{-1}) \subset C\}$. The volume of C is overestimated by $v_u(C, \lambda) = \lambda^{-2} \cdot \#C$. The volume of $C \cap S(y, r)$ is underestimated by $v_l\big(C \cap S(y, r), \lambda\big) = \lambda^{-2} \cdot \#H(C)$. Then

$$\frac{\#C \cap S(y,r) \cap \lambda^{-1}\mathbf{Z}^2}{\#C \cap \lambda^{-1}\mathbf{Z}^2} \geq \frac{v_l\big(C \cap S(y,r), \lambda\big)}{v_u(C, \lambda)} \uparrow_{\lambda\to\infty} \frac{\mathrm{vol}\big(C \cap S(y,r)\big)}{\mathrm{vol}(C)} \geq p_2 > p$$

Thus for large enough λ, the fraction of lattice points in C that are also in $\big(C \cap S(y, r)\big)$ exceeds p.

This argument works for the particular sphere $S(y, r)$. In order to complete the argument, we need to show the following claim: For λ sufficiently large, the fraction of lattice points in C that are also in any sphere S which intersects the closed ball of radius δ around 0 exceeds p. We need to show that for all $\epsilon > 0$ there is a λ' such that for all spheres S intersecting the ball of radius δ around 0, $v_l\big(C \cap S(\lambda), \lambda\big) > (1-\epsilon)\,\mathrm{vol}(C \cap S)$.

It suffices to consider those spheres of radius r whose boundaries intersect C. Suppose the claim is false. Then for all λ there is a sphere $S(\lambda$ of radius r whose boundary intersects C such that $v_l\big(C \cap S(\lambda), \lambda\big) \leq (1-\epsilon)\,\mathrm{vol}(C \cap S)$. We can extract a convergent subsequence of λ such that corresponding sequence of spheres converges to a limit sphere S of radius r and center x. Now consider the sphere S_ϵ of radius $r(\epsilon) < r$ and center x such that $\mathrm{vol}(S_\epsilon \cap C) = \sqrt{1-\epsilon/2}\,\mathrm{vol}(S \cap C)$. Thus $S_\epsilon \subset \mathrm{int}\, S$. For all λ sufficiently large, $S(\lambda) \supset S_\epsilon$, and so $v_l\big(S(\lambda) \cap C, \lambda\big) \geq v_l(S_\epsilon \cap C, \lambda)$. For λ sufficiently large, $v_l(S_\epsilon \cap C, \lambda) > \sqrt{1-\epsilon/2}\,\mathrm{vol}(S_\epsilon \cap C)$, so $v_l\big(S(\lambda) \cap C, \lambda\big) > (1-\epsilon/2)\,\mathrm{vol}(S \cap C)$. Hence for λ sufficiently large, $v_l\big(S(\lambda) \cap C, \lambda\big) > (1-\epsilon)\,\mathrm{vol}(S(\lambda) \cap C)$.

d) This claim is obvious, for if B is fixed and λ is sufficiently large, the neighborhood of any point "next to" B will contain all B, and so the frequency of 0's in the neighborhood will be $\#B/\#\lambda C$, which will go to 0 as λ grows. □

Proof of Theorem 4.7: All players initially choosing 0 will remain at 0 when a revision opportunity arises. Our concern is with the set A of $4(1-\theta)N^2$ players who are initially choosing 1. Let τ_i denote the smallest time at which i distinct players in A have had at least one revision opportunity. The waiting time τ_1 is exponentially distributed with mean $1/4N^2$. After $i-1$ players have switched, τ_i is the time of the first revision opportunity among the remaining $4N^2 - i$ players. Thus $\tau_i - \tau_{i-1}$ is distributed exponentially with mean $1/(4N^2+1-i)$. Furthermore the variables $\tau_i - \tau_{i-1}$ are mutually independent. By definition, $\tau_g^N = \tau_{4N^2}$. Summing over all the increments,

$$E(\tau_{4N^2}) = \sum_{k=1}^{4(1-\theta)N^2} \frac{1}{k}, \qquad \text{so that} \qquad \frac{1}{\log 4(1-\theta)N^2} E(\tau_g^N) \to 1$$

The variance of each increment is twice the square of the mean. This forms a convergent serices, so $\lim_{N\to\infty} \tau_g^N - E(\tau_g^N)$ converges to an (a.s. finite) limit random variable with variance $2\sum_{k=1}^{\infty} 1/k^2 < \infty$. Thus

$$\lim_{N\to\infty} \frac{1}{2\log N}\big(\tau_g^N - E(\tau_g^N)\big) \to 0$$

almost surely. □

The remaining results are for $K \times K$ games. Let S denote a given player set, and let $\{p_1(s,\eta)\}_{s\in S}$ and $\{p_2(s,\eta)\}_{s\in S}$ denote two families of probability distributions on W such that if $\phi \geq \eta$, then for all $s \in S$, $p_1(s,\phi) \geq p_2(s,\eta)$ in the sense of first order stochastic dominance. The embedded Markov chain that records jumps in the state has transition probabilities

$$P_i(\hat{\eta},\eta) = \frac{1}{|S|} p_i(s,\eta)\big(\hat{\eta}(s)\big)$$

For an initial distribution μ_0^i on configurations, define μ_t in the usual way;

$$\mu_t^i(\hat{\eta}) = \sum_{\eta} P_i(\hat{\eta},\eta)\mu_{t-1}(\eta).$$

The distribution μ_t^i describes the distribution of states just after the tth revision opportunity.

Theorem 6.1: Suppose that μ_0^1 first-order stochastically dominates μ_0^2. Then for all t, μ_t^1 first-order stochastically dominates μ_t^2.

Proof: The proof is a standard coupling argument and relies on the following Lemma which is well-known in the stochastic dominance literature (Strassen 1965). We shall use it thrice:

Lemma 6.1: Let X be a compact metric space with a partial ordering such that $K = \{(x,y) \in X \times X : x \geq y\}$ is closed in the product topology. The probability distribution μ_1 on X stochastically dominates probability distribution μ_2 if and only if there exists a probability distribution μ on $X \times X$ such that for all Borel sets A of X,

1. $\mu\{(x,y) : x \in A\} = \mu_1(A)$,
2. $\mu\{(x,y) : y \in A\} = \mu_2(A)$, and
3. $\mu(K) = 1$.

Whenever $\phi \geq \eta$, $p_1(s,\phi) \geq p_2(s,\eta)$. Let $\tilde{p}(s,\phi,\eta)(w,w')$ denote the coupling measure whose existence is the subject of Lemma 6.1. Now we build a Markov process on $X \times X$ with the following properties: (i) The marginal process on the first coordinate is the P_1 process; (ii) the marginal process on the second coordinate is the P_2 process, and if $(\phi_0,\eta_0) \in K$, then $(\phi_t,\eta_t) \in K$ for all t. Let

$$P(\hat{\phi},\hat{\eta},\phi,\eta) = \begin{cases} P_1(\hat{\phi},\phi)P_2(\hat{\eta},\eta) & \text{if } (\phi,\eta) \notin K, \\ \frac{1}{|S|}p(s,\phi,\eta)\big(\hat{\phi}(s),\hat{\eta}(s)\big) & \text{if } (\phi,\eta) \in K, \end{cases}$$

Intuitively, the coordinates evolve independently so long as the coordinate pair is not in K, and according to the coupling measure when they are in K. The coupling measure is such that states in K transit to K.

Nothing in the construction guarantee that the path from an arbitrarily given pair of coordinates will ever hit K. This comes from the remaining hypothesis. Suppose that μ_0^1, stochastically dominates μ_0^2. The lemma guarantees the existence of a coupling measure μ_0 such that $\mu_0(\phi \geq \eta) = 1$, and such that the marginal distribution on the first and second coordinates are μ_0^1 and μ_0^2, respectively. Now consider the evolution of the process on $X \times X$ with transition probability P and initial distribution μ. Since $\mu_0(\phi \geq \eta) = 1$, $\mu_t(\phi \geq \eta) = 1$ for all $t > 0$. The marginal distributions of μ_t on the first and second coordinates are μ_t^1 and μ_t^2, respectively, and we conclude from Lemma 6.1 that μ_t^1 stochastically dominates μ_t^2. □

The apparatus of infinitessimal operators allows this Theorem to be extended to processes on Z^d with infinitessimal generators like those described in Blume (1993). The application of this Theorem to Theorem 5.2 is clear. Suppose in addition that a distribution of trembles q is fixed: $q(w)$ is the probability that, in the event of a tremble, outcome w is drawn. Now consider dynamics in which every player at a revision opportunity chooses the best response with probability $1-\epsilon$ and draws from q with probability ϵ. Then the stochastic dominance hypothesis is still satisfied, and the conclusions of Theorem 5.2 continue to hold.

Theorem 6.1 also contains information about the distribution of play across the player set at any moment of time. A distribution μ is said to have *positive correlations* if $\int fg\,d\mu \geq \int f\,d\mu \int g\,d\mu$ for all increasing functions f and g. Suppose that a distribution of states in the best-response strategy revision process had positive correlations. Take $f(\eta)$ to be the indicator function for the event "s plays a strategy greater than or equal to v" and $g(\eta)$ to be the indicator function for the event "t plays a strategy greater than or equal to w". Applying the inequality, we get that $\mu(\phi(s) \geq v \mid \phi(t) \geq w) \geq \mu(\phi(s) \geq v)$, the statement that s is more likely to choose high when t does.

Theorem 6.2: If μ_0 has positive correlations, then the time t-marginal distribution for best-response strategy revision for a game G with strategic complementarities has positive correlations.

Proof: Define for each non-decreasing function g the probability distribution μ_0^g such that $\mu_0^g(\eta) = g(\eta)\mu(\eta)/\sum_\phi g(\phi)\mu(\phi)$. The distribution μ_0 has positive correlations if and only if μ_0^g stochastically dominates μ_0. If μ_0 has positive correlations, it follows from Theorem 5.2 that each μ_t^g stochastically dominates μ_t, so μ_t has positive correlations. □

The conclusion of this theorem applies whenever the initial distribution has assignments of play which are independent across players. This includes beginning at any point-mass.

Proof of Theorem 5.4: The proof of Theorem 5.4 follows the same path as the proof of Theorem with the necessary modifications for more than two strategies. The necessary modifications will involving coupling two processes together as in the proof of Theorem 6.1. The two processes will have the same transition rules but different initial conditions. The strategy of the proof is to use the coupling to show that for any s, $\lim_{t\to\infty} \Pr\{\eta_t(s) \geq w^*\} = 1$. A similar coupling shows that $\lim_{t\to\infty} \Pr\{\eta_t(s) \leq w^*\} = 1$, and this completes the proof.

Suppose that the strategies are ordered $0 < 1 < \ldots < w^* < \ldots < \bar{w}$. We start with a given λ, and let $r' = \max\{|x| : x \in \lambda C \cap \mathbf{Z^2}\}$. For a given neighborhood $V_0 = \lambda C$ there is a sphere B_{w^*} of radius r sufficiently large such that if s is within l_1 distance 1 from the sphere and if p is a sufficiently small number, $\#V_s \cap B_{w^*}/\#Vs \geq p$. Define $p(\lambda)$ to be the supremum of the numbers p for which this inequality is true. As before, as λ grows, $p(\lambda)$ converges to 1/2. We assume that Let $\sup_{v<w} p(v,w) < p(\lambda)$.

Successively, for $i \leq |w^*| - 1$, let B_{w^*-i} denote the disk around the origin of radius $r + ir'$, and let $R_i = B_{w^*-i}/B_{w^*-i+1}$. Then B_{w^*-i} contains at least fraction p of the neighbors of any player sufficiently near the boundary of B_{w^*-i}. Define the initial condition

$$\phi_0(s) = \begin{cases} w^* & \text{if } s \in B_{w^*}, \\ w^* - i & \text{if } s \in R_i \text{ for } i \leq w^* - 1, \\ 0 & \text{otherwise.} \end{cases}$$

This initial condition describes an interior disk of players choosing w^* surrounded by successive rings of players choosing $w^* - 1$, $w^* - 2$, and so forth, descernding to a plain of players choosing 0.

For the second process we shall take η_0 to be any configuration such that $\eta_0(s) = w^*$ for all $s \in B_1$. In other words, $\phi_0(s) > 0$ implies $\eta_0(s) = w^*$. Thus $\eta_0 \geq \phi_0$. Using the coupling of Theorem 6.1, it follows that if ϕ_t converges almost surely to the configuration in which all players choose w^*, then the same is true for the $\{\eta_t\}$ process.

To show that the $\{\phi_t\}$ process converges, notice first that it is monotone increasing; that is, $\phi_{t+1} \geq \phi_t$. To see this, observe that in the ϕ_0 configuration, no player wants to choose a lower strategy. The players who are most likely to want to move down are those inside, but at the very edge, of a B_j. Such a player has fraction p of her neighbors choosing $j-1$ while she chooses j — the width of the rings was chosen to be large enough that the neighborhood of no player can intersect more than two rings. The hypothesis of the theorem requires that such a player will choose $w > j-1$, and so w must be at least j. Moving more towards the inner edge of any ring (or into B_{w^*}) just increases the distribution of neighbors' play in the sense of first-order stochastic dominance, and so the best response never falls as we move in towards the center. Next, notice that if no player wants to revise downward and one player revises upward, monotonicity of best response with respect to scochastic dominance implies that still no player will wish to revise downward. Thus the process is monotone increasing.

Consider now a limit configuration ϕ of the $\{\phi_t\}$ process. It is a Nash configuration. Furthermore, on each set B_w ϕ is at least as big as w. Next we can see that $\phi(s) > 0$ for all s exactly by the arguments of the proof of theorem 4.4: Near the border of the largest sphere contained in the set of all players for which $\phi(s) > 0$, all players with revision opportunities confront mixtures of play in which no more than fraction p of their neighbors play 0. Hence they will want to switch. Thus in the limit this sphere must contain all of $B(N)$. Now we see that ϕ is at least 2 on B_2 and at least 1 everywhere else. Repeating the argument shows that ϕ must be at least 2 everywhere, and so on until we conclude that $\phi(s) \geq w^*$ for all s.

So far we have shown that $\lim_{t\to\infty} \Pr\{\eta_t(s) \geq w^*\} = 1$ for initial η_0 which has all players in B_{w^*} choosing w^*. Now run the same argument again, but with an initial condition which is w^* on some sphere B^{w^*}, w^*+1 on B^{w^*+1}/B^{w^*}, and so forth. The conclusion is that $\lim_{t\to\infty} \Pr\{\eta_t(s) \leq w^*\} = 1$ for any initial condition η_0 such that all players in B^{w^*} choose w^*. Thus it follows that for any initial condition η_0 such that all players in $B_{w^*} \cup B^{w^*}$ are choosing w^*, $\lim_{t\to\infty} \Pr\{\eta_t(s) = w^*\} = 1$. □

REFERENCES

Anderlini, L. and A. Ianni (1993). "Path dependence and learning from neighbors," unpublished.

Bergstrom, T. (1993). "On the evolution of altruistic ethical rules for siblings," unpublished.

Berninghaus, S. K. and U. Schwalbe (1993). "Evolution, interaction and Nash equilibrium," unpublished.

Blume, L. E. (1993). "The statistical mechanics of strategic interaction," *Games and Economic Behavior,* 4, 387–424.

Canning, D. (1990). "Social equilibrium," unpublished.

Durrett, R. (1988). *Lecture Notes on Particle Systems and Percolation.* Belmont CA: Wadsworth.

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

Foster, D. and P. Young (1990). "Stochastic evolutionary game dynamics," *Theoretical Population Biology,* 38, 219–232.

Hamilton, W. (1964). "The genetical evolution of social behavior," *Journal of Theoretical Biology,* 7, 1–52.

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

Myerson, R., G. Pollack and J. Swinkels (1991). "Viscous population equilibrium," *Games and Economic Behavior,* 3, 101–109.

Oechssler, J. (1993). "Competition among conventions," unpublished ms.

Strassen, V. (1965). "The existence of probability measures with given marginals," *Ann. Math. Statist.,* 36, 423–439.

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