GAMES AND ECONOMIC BEHAVIOR 5, 387–424 (1993)

The Statistical Mechanics of Strategic Interaction*

LAWRENCE E. BLUME

Department of Economics, Uris Hall, Cornell University, Ithaca, New York 14853

Received August 19, 1992

I study strategic interaction among players who live on a lattice. Each player interacts directly with only a finite set of neighbors, but any two players indirectly interact through a finite chain of direct interactions. I examine various *stochastic strategy revision processes*, including (myopic) best response and stochastic choice. I discuss both stationary distributions and the limit behavior of these Markov processes. Stationary distributions are partially characterized, and the asymptotic behavior of stochastic choice for those processes whose choice rule is nearly best-response is related to equilibrium selection in symmetric 2×2 and $n \times n$ coordination games. *Journal of Economic Literature* Classification Number: C78.

> Whoever is united to us by any connexion is always sure of a share of our love, proportion'd to the connexion, without enquiring into his other qualities.
>
> David Hume, *A Treatise of Human Nature*, Book II, Chap. 2, Pt. 4

1. INTRODUCTION

To most economists, the chief virtue of the market is its ability to *decentralize* the optimal resource allocation problem. As conceived in modern economic theory, decentralization has come to mean the "parallelization" of the planning problem. Decisions about consumption and

* I am grateful to Richard Boylan, David Easley, Drew Fudenberg, Andy McLennan, Roy Radner, Aldo Rustichini, Vince Crawford, two referees, and especially Morgan Kelly for many useful suggestions. The proximate cause of my writing this paper is the many conversations on population processes for strategic dynamics I had with David Canning at the EUI Workshop "Rational Behavior and Aggregation: Theory and Tests," in November 1990. Financial support from NSF Grant SES-8921415 is gratefully acknowledged.

0899-8256/93 $5.00

production are simultaneously made by small decisionmakers, whose decisions are coordinated by the price system. In the mechanism design literature, the market is praised for the parsimony of communication that is required to achieve this parallelization of the planning problem. The dimension of the message space is small.

Decentralization connotes more than the efficient parallelization of a complex computation. While the beauty of the Arrow–Debreu analytical framework is its freedom from particular descriptions of market institutions, it is not given that any set of market institutions will lead to competitive outcomes. Our informal description of the market mechanism involves a Walrasian auctioneer, or centralized price-setting institution. Although some markets, such as currency futures markets, are organized in this fashion, others, like the markets for many retail goods and services, are not. Competition in these markets is not centralized; each firm, providing services to a neighborhood clientele, may have only a few direct competitors. Nonetheless monopoly power is limited, in this case by extended chains of local competition. This description of economic interaction is a key feature of many economic phenomena, including community competition with local public goods and monopolistic competition. In the first instance, proximity is geographic. People in the same community interact in choosing the level of local public good consumption. In the second instance, proximity has to do with nearness of product. A firm interacts most directly with those firms that produce close substitutes. In each of these economic models, one studies how local behaviors propagate through the system to determine its global behavior.

The phenomenon of decentralization through local competition is even more prevalent when we turn to social choices that are not resolved by markets. The emerging literature on the evolution of conventions and norms is concerned with just this problem. A trivial example is the problem of determining which side of the road we drive on. Here an interaction occurs when two drivers approach each other on the road, travelling in different directions. Each driver will interact with only a small fraction of the set of all drivers and will interact most often with those who live or work nearest him. Nonetheless a convention is established.

Local interaction is a sufficiently widespread phenomenon that the small amount of attention such models have received is surprising. The earliest explicit treatment of which I am aware is due to Schelling (1971), who is concerned with processes leading to segregation in residential neighborhoods. An interesting part of Schelling's paper is his comparisons of models of "neighborhood tipping" where the interaction is very local with models where the interactions are more global.[1] The effects of locality have

[1] Schelling's well-known paper (1973) analyzes similar models with global interactions—where individuals care only about the behavior of population aggregates rather than the behavior of a few neighbors.

been studied most recently again by Ellison (1992). Whereas Schelling was more interested in how the size of neighborhoods affects the limit pattern of choices, Ellison is concerned with the speed of convergence.

In this paper I present some dynamic models of strategic interaction in a population of players where direct interaction is local but indirect interaction is global. The models I present are broadly concerned with the population dynamics of boundedly rational play. This paradigm is the point of departure for most recent work in evolutionary game theory. One view of this work is that it describes a population of players, each of whom "adapts" to the environment in which he plays. The players' environment is in turn determined by the choices made by the entire collection of players.[2] Recent work from this point of view includes Canning (1990), Ellison (1992), Fudenberg and Levine (1990), Kandori *et al.* (1993), Kandori and Rob (1992), Nachbar (1990), and Young (1993). These models of strategic interaction trade off rationality of the individual for the emergence of collective order in the population. (In this context, Peyton Young reminds us of Burke's dictum, "The individual is foolish but the species is wise.") The central role of sophisticated rationality hypotheses and common knowledge suppositions in the traditional analysis is played here by naive rationality notions and a detailed specification of the process by which players meet. One looks for the emergence of Nash-like play in the aggregate rather than at the level of the individual player.

The strategy revision dynamic has two features that represent a departure from continuous iteration of the (rational) best-response correspondence. First, players do not instantaneously react to their environment. Having made a decision, they are *locked in* for some (possibly short) period of time. Second, players are myopic in their decisionmaking. They respond to instantaneous reward rather than to some discounted value of the future reward stream.[3] A fully rational player would take into account not only the current play of his neighbors, but also forecasted future play over the period which he expects to be locked in, and discount appropriately. I specify two kinds of strategy revision processes. *Best-response dynamics* suppose that each player maximizes instantaneous payoff flow at each revision opportunity. *Stochastic-choice dynamics* suppose that at a revision opportunity, players choose their strategies from some probability distribution over choices whose character depends upon the payoffs each choice yields. This view is to be contrasted with that taken in much of the recent literature on population models, where the dynamics of strategy adoption are driven by biological processes of birth, death, and fertility.

[2] See Crawford (1991) for a discussion of this interpretation of evolutionary game theory.

[3] Best-response revision is dominant strategy play of the dynamic game when the discount factor is sufficiently low relative to the rate at which revision opportunities arrive, and hence fully rational.

In the model examined here, each player interacts directly with his neighbors, and although each player has few neighbors, all players interact indirectly through a claim of direct interactions. I study *stochastic strategy revision processes*, wherein each player revises his strategic choice in response to the play of his neighbors. These continuous-time Markov population processes are examples of infinite particle systems. For games with two strategies, best-response strategy revision is related to the "voter model." The log-linear strategy revision process is a "stochastic Ising model," a stochastic process related to the Ising model of statistical physics. [A good introduction to these and other infinite particle systems is Liggett (1985).] Although Section 3 contains some results on best-response strategy revision, the main results of this paper in Sections 5 and 6 address the equilibrium selection problem—the question of which Nash equilibria emerge in the asymptotic behavior of stochastic strategy revision processes, which are small perturbations of best-response processes. In two-strategy coordination games, the risk-dominant equilibrium is selected for. In games with more than two strategies, some sufficient conditions for equilibrium selection are given. Proofs and discussion of some ancillary technical issues are in the appendices.

2. The Model

Each site on the d-dimensional integer Z^d is the address of one player. Every site $s \in Z^d$ is directly connected to only a finite number of other sites. The set of sites directly connected to site s is the *neighborhood* of s. This nonempty set is called V_s. For any finite set of vertices $T \subset Z^d$ the *boundary* of T is the set $\partial T = \{\cup_{s \in T} V_s\} \backslash T$. This depiction of the neighborhood relation implicitly assumes that it is symmetric: If t is a neighbor of s, then s is a neighbor of t. (There are natural economic models where this is not the case.) I also assume that the neighborhoods are *translation-invariant* on the lattice: $V_s = \{s\} + V_0$. A convenient representation of the neighborhood arrangement is to imagine an (unoriented) graph whose vertices are the elements of Z^d and whose edges connect neighbors: There is an edge connecting s and t if and only if $t \in V_s$. Players are referred to by their addresses; player s is the player at site s.

Objects of choice for all players are actions in the set $W = \{0, \ldots, |W| - 1\}$. A *configuration* of the population is a function ϕ: $Z^d \mapsto W$. A configuration describes the strategy choices of the player population: $\phi(t)$ is the strategy employed by player t. For any $T \subset Z^d$ and configuration η, $\eta(T)$ denotes the restriction of η to T. For any finite set $T \subset Z^d$ of players, let $X(T)$ denote the set of configurations of sites in T. Let $\phi(-s)$ denote the configuration $\phi(Z^d \backslash \{s\})$ of all players other than player s, and similarly, $\phi(-T)$ will denote the configuration $\phi(Z^d \backslash T)$.

Finally, for a given configuration ϕ and action w, let ϕ_s^w denote the configuration identical to ϕ except that player s is using strategy w.

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

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

A *strategy-revision process* is a continuous-time Markov process on the space of configurations that describes the evolution of players' choices through time. A formal description of the process is given in terms of the generator of its semigroup in Appendix 1. Roughly speaking, 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, he responds to his neighbors' current configuration by choosing an action according to some choice rule. I consider here processes constructed from two classes of choice rules: First, simple *best response*, and second, *stochastic choice*.

With simple best response, the player chooses equiprobably from among those strategies that give the highest payoff flow given the current play of his neighbors. Formally, let $M(\eta, s)$ denote the set of best responses by s when the population configuration is η:

$$M(\eta, s) = \operatorname*{argmax}_{w \in W} \sum_{t \in V_s} G(w, \eta(t)).$$

Let $p_s(v \mid \phi)$ denote the probability that the player at site s will choose $v \in W$ given that the rest of the population is configured according to ϕ. Then

$$p_s(v \mid \phi) = \begin{cases} 1/|M(\phi, s)| & \text{if } v \in M(\phi, s), \\ 0 & \text{otherwise}. \end{cases} \tag{2.1}$$

Note that this definition depends only upon the coordinates $t \in V_s$ of ϕ, $\phi(V_s)$.

With stochastic choice, choice is random, as it can be with the best-response rule when the best response by player s is not unique, but the random draw is not confined to best responses. Two sources of random choice behavior are mistakes and unpredictable experimentation. These motivations are discussed, among other places, in Canning (1990), Kandori *et al.* (1993), and Young (1993). Let $\{p_s^{br}\}$ denote the selection probabilities given by Eq. (2.1). Let $\{q_s\}$ denote another system of selection probabilities: $q_s\{\cdot \mid \phi)$ is a *completely mixed* probability distribution on W that depends upon ϕ only through the coordinates in V_s. Then

$$p_s(w \mid \phi) = \frac{\beta}{1+\beta} p_s^{br}(w \mid \phi) + \frac{1}{1+\beta} q_s(w \mid \phi) \tag{2.2}$$

is a perturbation of best-response choice that places positive probability on each outcome. Obviously, as β grows large, this choice distribution converges to the best-response distribution. I make the following assumptions on the system $\{q_s\}$:

(i) Like the system $\{p_s\}$, they are shift invariant. If $\phi(V_s) = \phi(V_t)$, then for all $w \in W$, $q_s(w \mid \phi) = q_t(w \mid \phi)$.

(ii) For all s, w, and ϕ, $q_s(w \mid \phi) > 0$.

I refer to stochastic-choice rules described by Eq. (2.2) as *perturbed best-response rules.*

Another source of stochastic choice is as the outcome of a rational stochastic-choice theory. One such theory is the random utility model. Here the return to any action is random and varies across each player's choice opportunities. To the modeler with incomplete information about the return to actions at any instant, each player's choice given the modeler's information will appear to be random.

An important random utility model is the *log-linear model,* where the log-odds of the player at site s choosing strategy v to strategy w are proportional to some function of the payoffs that v and w achieve from the interaction of the player with his neighbors. Again, the choice of strategy by player s depends upon the strategic-choice processes of other players only in that payoffs are calculated at the current configuration of s's neighbors. It proves useful to parametrize choice functions by a multiplicative constant of proportionality β:

$$\log \frac{p_s(v \mid \phi)}{p_s(w \mid \phi)} = \beta \sum_{t \in V_s} G(v, \phi(t)) - G(w, \phi(t)). \tag{2.3}$$

The selection probabilities are

$$p_s(v \mid \phi) = \left\{ \sum_{w \in W} \exp \beta \sum_{t \in V_s} G(w, \phi(t)) - G(v, \phi(t)) \right\}^{-1}.$$

Again, this definition depends only upon the coordinates $t \in V_s$ of ϕ. The constant β parameterizes the sensitivity of choice. When $\beta = 0$, this distribution puts equal weight on all strategic choices. As $\beta \uparrow \infty$, this distribution converges to one that puts equal positive weight on all best responses and 0 weight on all other responses. Thus best-response revision arises as a limiting case of the stochastic-choice model as the proportionality constant becomes large. Stochastic-choice models like that described by Eq. (2.3) originated in Thurstone's study of comparative judgment. The characterization of these choice models as random utility models was first established by Block and Marschak (1960).

The class of stochastic choice models described by Eq. (2.3) is important in its own right, and it also provides comparison processes useful to the study of the perturbed best-response models described by Eq. (2.2). In what follows, I refer to the models of Eq. (2.3) as the log-linear rule. Unfortunately the class of stochastic choice rules described by Eq. (2.2) does not contain the log-linear rule. When the log-linear rule is rewritten in additive form, it is clear that as β grows, the additive perturbation $\{q_s\}$ changes. I have chosen the peculiar parametrization of Eq. (2.2) to simplify the statements of results. To further simplify statements of theorems, when referring to any stochastic choice rule I write $\beta = \infty$ to signify best-response strategy revision.

With both best-response rules and stochastic choice rules, the objects of choice are pure strategies, not mixed strategies. When a choice opportunity arrives, the player must choose a pure strategy and not a distribution. When he observes the play of his neighbors, he sees the pure strategy choices they have made. But even though players choose only pure strategies, players will confront distributions of actions. At a typical moment in time, the configuration of choices will be polymorphic—many choices will be represented in the population. Each player cares about the distribution of choice in his neighborhood, and this distribution may well be mixed.

The key features of both best response and stochastic choice are, first, lock-in, and second, bounded rationality. By lock-in I mean that once a player makes a choice, he is committed to it for some while. He cannot revise his choice until his alarm clock goes off again. It is easy to imagine how lock-in can arise by instituting decision costs or fixed costs for strategy revision. Technically speaking, the random lock-in assumption plays a critical role. A consequence of the Poisson assumption is that in any sufficiently small interval of time, it is unlikely that more than one player is making a decision. The dynamics of simultaneous strategy revision are very different.

Even if we assume that the lock-in is "technologically given," both decision rules still describe boundedly rational choice. For best-response dynamics, this bounded rationality arises because, at the moment of decision, players contemplate only the present rewards to each choice, and not the expected flow of future rewards. Full rationality entails each player forming beliefs about the future play of his neighbors given their play up to the present, and then maximizes an expected present discounted value of the utility stream computed with some given discount factor. The Markov property for a fully rational strategy revision process requires that information about the configuration before date t be irrelevant to choice at t. Typically this will not be the case: For a player s estimating the future play of his neighbor s' at time $t + t'$, it would be useful to know s''s neighbors' play at time t. These cannot be observed directly by the player except for the play of those neighbors whom he shares in common with s'. But information about s''s neighbors' play at t is carried by the play of s' at times before t. Markov play will arise when each player's discount factor is small relative to the mean waiting time between rings of the alarm clocks. In this case, future play is irrelevant, and full rationality entails maximizing the current instantaneous payoff flow, in other words, the simple best-response rule. Nonetheless I resist this fully rational interpretation of best-response dynamics in favor of one emphasizing players' myopia as a departure from rationality.

The state space for all stochastic strategy revision processes is the Borel space X of configurations with the product discrete topology. Let $\Pr_\mu(A)$ denote the probability of some measurable event A in the space of sample paths for a strategy revision process with initial distribution of configurations μ, and let $\Pr_\phi(\cdot)$ denote the case where μ is a point mass on configuration ϕ. A formal description in terms of its infinitesimal generator and some additional technical apparatus is given in Appendix 1.

DEFINITION 2.1. A probability measure μ on X is *stationary* for the process $\{\phi_t\}_{t\geq 0}$ if, for all $t \geq 0$ and every measurable set of states B, $\Pr_\mu(\phi_t \in B) = \mu(B)$.

The goal of the analysis in the remainder of this paper is to relate stationary distributions of strategy revision processes to equilibrium concepts for noncooperative games. But first I record the existence of stationary distributions and their continuity with respect to parameters. The first theorem is the standard stationary distribution existence theorem for Feller processes on compact domains:

THEOREM 2.1. *The set of stationary distributions of best-response and stochastic-choice strategy revision processes with choice probabilities* $\{p_s\}_{s\in Z^d}$ *is nonempty, convex and compact, and the closed convex hull of its extreme points. For both the perturbed best-response rule and the log-linear rule, the correspondence mapping* (β, G) *into the set of stationary measures of the corresponding process is pointwise upper hemicontinuous in the weak convergence topology for measures on* W^{Z^d}.

In particular, the limit of any weakly convergent sequence of stationary distributions as β grows large for either stochastic choice rule is a stationary distribution for the best-response strategy revision process.

3. Iterated Elimination of Dominated Strategies

In this section, stationary behavior and the dynamic behavior of strategy revision processes are seen to respect the iterated elimination of strongly dominated strategies.

Recall that a strategy w_i is strongly dominated (in the set of mixed strategies) for player i if and only if there is no probability distribution over the play of i's opponents for which w_i is a best response. In games on the lattice where each player is best responding to the play of his neighbors, it follows from the symmetry of the strategic situation and the additivity of payoff flows that only the numbers of occurrences, or frequencies, of choices by the neighbors of s can affect his choice at any of his revision opportunities. But because s has only $|V_s|$ neighbors, there are only a finite number of play distributions that he could possibly observe. For instance, if $|W| = 2$, the possible distributions are those where fraction $k/|V_s|$ of s's neighbors play strategy 1 and the remaining fraction, $(|V_s| - k)/|V_s|$ plays strategy 2, where k is an integer between 0 and V_s. This motivates the following definition:

Definition 3.1. Strategy w is *V_s-dominated* if there is no configuration $\phi(V_s)$ of play among the neighbors of s for which w is a best response. Let W_1 denote the set of V_s-undominated outcomes and define inductively U_k to be the set of outcomes $w \in W_{k-1}$, which are a best response to no configuration $\phi(V_s)$ taking values only in W_{k-1}, and $W_k = W_{k-1}/U_k$. The strategy w is *V_s-iteratively dominated* if $w \in U_k$ for some k.

In the two-player symmetric three-by-three game with payoff matrix

$$\begin{pmatrix} 1 & 0 & -10 \\ -2 & 0 & 1 \\ -1 & -1 & -1 \end{pmatrix}$$

the bottom strategy is not dominated. In particular it is a best response to the mixed strategy (0.7, 0.1, 0.2). However, it is V_s-dominated in the lattice game on the line with neighborhoods $V_s = \{s - 1, s + 1\}$ since there is no pure strategy combination for s's two neighbors, which has bottom as a best response. Clearly for a given game with rational payoffs, there is a neighborhood large enough that each strategy $w \in W$ is V_s-dominated if and only if it is dominated in the two-player game.

Only a finite number K of rounds of iterative elimination suffice to eliminate all V_s-dominated strategies, so the V_s-undominated stategies are precisely those in W_K. Note that if w is strictly dominated, then it is V_s-dominated. The following result connects stationary distributions with iterated elimination of V_s-dominated strategies. The theorem states that stationary distributions for best-response dynamics place all their mass on iteratively undominated strategies.

THEOREM 3.1. *Suppose that $w \in W$ is iteratively W_s-dominated in the game matrix G. If μ is an invariant distribution on X for best-response strategy revision, then $\mu\{\phi(s) = w\} = 0$.*

Proof. No player will ever switch to a V_s-dominated strategy, and every player will switch away from a V_s-dominated strategy. Suppose all iteratively V_s dominated strategies can be eliminated in at most K rounds of elimination. Fix s. Let $S_1 = V_s \cup \{s\}$, and proceeding inductively, let $S_k = S_{k-1} \cup \partial S_{k-1}$. The set S_K contains a finite number of players. The waiting time until all of the players in S_K have had at least one revision opportunity is almost surely finite. Call it t_K. At any time after t_K, all players in S_{K-1} respond only to configurations that take values in W_1. Let t_{K-1} denote the waiting time until all players in S_{K-1} have had a revision opportunity subsequent to time t_K. The waiting time t_{K-1} is finite. At any time after t_{K-1}, all players in S_{K-2} respond only to configurations that take values in W_2. Proceeding inductively, we establish a finite time t_1 after which s responds only to configurations in W_{K-1}. For all time beyond t_1 he will choose only strategies in W_K, the set of V_s-iteratively undominated strategies. ■

An immediate consequence of Theorem 3.1 and the upper hemicontinuity of the stationary distribution correspondence for both stochastic choice rules in the parameter β (Theorem 2.1) is:

COROLLARY 3.1. *Let $\{\mu_\beta\}$ denote a sequence of stationary distributions for either a log-linear or a perturbed best-response strategy revision process with fixed G as β becomes large. If w is iteratively V_s-dominated in the game matrix G, then $\lim_{\beta \uparrow \infty} \mu_\beta\{\phi(s) = w\} = 0$.*

Iteratively eliminated strategies tend not to arise in stationary strategy revision processes. Starting from an arbitrary initial (nonstationary) distribution of play, do iteratively eliminated strategies tend to disappear? Of course this must be true for ergodic processes, but not all strategy revision processes are ergodic. Nonetheless, the proof of Theorem 3.1 shows this to be true for $\beta = \infty$, and the proofs for the stochastic-choice rules can be found in Appendix 3.

THEOREM 3.2 *For any stochastic-choice rule and for all initial configurations* ϕ_0,

$$\liminf_{t\to\infty} \Pr_{\phi 0}\{\phi_t(s) \text{ is iteratively } V_s\text{-undominated}\}$$

equals 1 *at* $\beta = \infty$ *and converges to* 1 *as* β *grows large.*

4. NASH EQUILIBRIA OF LATTICE GAMES

This section begins the analysis of the relationship between stationary distributions of strategy revision processes and Nash equilibrium. First I identify a class of configurations called *Nash configurations*, which exhibit the best-response property of Nash equilibrium, and then I take up the questions of when they exist and when stationary distributions are concentrated on them.

Consider the one-shot game wherein each player $s \in Z^d$ chooses one strategy that is used simultaneously in play against all of his neighbors $t \in V_s$. A (pure strategy) Nash equilibrium for this infinite player game is a configuration wherein each player chooses an action that maximizes the sum of the returns against the play of each of his neighbors.

DEFINITION 4.1. The configuration η is a *Nash configuration* if for each player $s \in Z^d$, $\eta(s) \in M(\eta, s)$. The set of Nash configurations is denoted by $\mathcal{N}$.

The following example shows that Nash configurations may fail to exist.

EXAMPLE 4.1. *The game Rock, Scissors, and Paper, played on the one-dimensional integer lattice with nearest-neighbor interactions, shows that not all games have Nash configurations. Take* $d = 1$. *Each player's neighbors are those players nearest to him on either side:* $V_s = \{s - 1, s + 1\}$. *The game is given by the payoff matrix (for all* $s \in Z^d$)

$$G = \begin{bmatrix} 1 & 0 & 2 \\ 2 & 1 & 0 \\ 0 & 2 & 1 \end{bmatrix},$$

where each number g_{ij} *records the row player's payoff from the strategy combination* (i, j).

Player s *cannot play strategy* 1 *in any Nash configuration, for he will only play* 1 *if one of his neighbors plays* 3 *and the other does not play* 2.

But no player will play 3 if any of his neighbors plays 1. Similar arguments show that player s also cannot play 2 or 3 in any Nash configuration, so no Nash configurations exist.

This countable-player game has mixed strategy Nash equilibria; the standard existence proofs apply. It is pure strategy Nash equilibria that fail to exist.

In the one-shot game, a pure strategy Nash equilibrium may present a nontrivial distribution of strategies across the population. Consider the "anti-coordination" game

$$G = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

played on the integer line Z^1 with neighborhoods $V_s = \{s - 1, s + 1\}$. Two Nash configurations have sites alternating their play: $\phi(s) = s \bmod 2$, and $\eta(s) = s + 1 \bmod 2$. In these strategies, the fraction of the population playing "up" (0) equals $\frac{1}{2}$. Although in this game the population fractions correspond to a mixed strategy Nash equilibrium of the two-player game, this is a coincidental feature of this example. Changing any one of the 1's to $1 + \varepsilon$ reduces the weight that the mixed strategy Nash equilibrium places on the corresponding column, but ϕ and η will still be Nash configurations of the infinite-player one-shot game.

In addition to the regular-looking Nash configurations ϕ and η, this game has other Nash configurations that do not respect the periodic structure of the lattice. The configuration $\zeta(s)$, which equals $\phi(s)$ for $s < 0$ and $\eta(s)$ for $s \geq 0$, is also a Nash configuration. This is not delicate—ϕ is a Nash configuration for any game G where $G(0, 1) > G(1, 0)$ and the diagonal payoffs are sufficiently near 0.

The analysis of best-response strategy revision processes is made complicated by the fact that when a player has several best responses, he will choose among them with equal probability. In this case, for instance, point masses cannot be invariant. Fortunately, this situation typically fails to arise. If $|V_s|$ is finite, then player s could see only a finite number of distributions of play among his neighbors, and so for generic payoff matrices each player has a unique best response to every configuration of play.

DEFINITION 4.2. A configuration $\phi \in X$ has a *unique best response* for the game G on Z^d if for all s, there is a unique $w \in W$ such that $\Sigma_{t \in V_s} G(w, \phi(t)) \geq \Sigma_{t \in V_s} G(w', \phi(t))$ for all $w' \in W$. The game G on Z^d has the *unique best-response property* for the neighborhoods $\{V_s\}_{s \in Z^d}$ if all configurations $\phi \in X$ have unique best responses.

When the game G has the unique best-response property for the neighborhoods $\{V_s\}_{s \in Z^d}$, a complete characterization of stationary distributions for best-response strategy revision, as well as a large β approximation result for stochastic choice strategy revision, is possible.

THEOREM 4.1 *Suppose that G has the unique best-response property for the neighborhoods $\{V_s\}_{s\in Z^d}$. The distribution μ is a stationary distribution for the best-response strategy revision process if and only if for every finite set of sites $S = \{s_1, \ldots, s_J\}$ and configuration ϕ of S,*

$$\sum_j \{\mu(\eta(s_j) = \phi(s_j) \mid \eta(S_j) = \phi(S_j)) - \mu(M(\eta, s_j) = \phi(s_j) \mid \eta(S_j) \tag{4.1}$$

$$= \phi(S_j))\}\mu(\eta(S_j) = \phi(S_j)) = 0,$$

where $S_j = S/\{s_j\}$. Furthermore, for any stochastic-choice rule and for all $\varepsilon > 0$, there is a b such that, if $\beta > b$ then the expression in Eq. (4.1) *differs from* 0 *by at most ε for all finite sets of sites S and configurations ϕ.*

In the case where S is a singleton, Eq. (4.1) becomes $\mu\{w = M(\phi, s)\} = \eta\{\phi(s) = w\}$. The events $\{w = M(\phi, s)\}$ and $\{\phi(s) = w\}$ are identical for all s if and only if ϕ is a pure strategy Nash equilibrium for the game G on Z^d. But even if the lattice game has no pure strategy Nash equilibria, these pairs of events must occur with the same probability.

COROLLARY 4.1. *Suppose that G has the unique best-response property for the neighborhoods $\{V_s\}_{s\in Z^d}$. If ϕ is a Nash configuration, then δ_ϕ is a stationary distribution for the best-response strategy revision process.*

5. GAMES WITH TWO STRATEGIES

The analysis of stochastic-choice strategy revision for games with two strategies is particularly simple. All log-linear strategy revision processes fall into the class of stochastic processes known as *stochastic Ising models*. The proofs of most theorems on log-linear strategy revision follow directly from the mapping of log-linear strategy revision processes into stochastic Ising models, which I describe below. The analysis of perturbed best-response strategy revision works by comparing a given perturbed best-response strategy revision process to selected log-linear strategy revision processes. The comparison technique is described in Appendix 3.

5.1. *Stochastic Ising Models and Log-Linear Strategy Revision*

Stochastic Ising models are surveyed in Liggett (1985, Chap. 4), and the proofs of most Theorems in this section follow directly from the mapping of log-linear strategy revision processes into stochastic Ising models, which I now decribe.

While stochastic-choice models are described by selection probabilities, stochastic Ising models are described by "flip rates"; the rate at which a player changes from i to j, $i, j = 0, 1, i \neq j$. Let $W = \{0, 1\}$. (Note the departure from the convention of labeling the first strategy 1, to simplify notation.) The flip rate describes the rate at which a given player (site in the language of stochastic Ising models), flips from 0 to 1 or from 1 to 0. The rate at which player s flips is equal to $p_s(0 \mid \phi)$ if $\phi(s) = 1$, and $p_s(1 \mid \phi)$ if $\phi(s) = 0$. While the selection probabilities for player s in any strategy revision process are independent of the current play at s, this will not be the case for flip rates because they describe the probability of departing from current play.

It is clear from Eq. (2.3) that selection probabilities for a given game matrix G are unchanged if the matrix is perturbed in such a manner that the difference between any two elements in the same column are unchanged. Given

$$G = \begin{pmatrix} g_{11} & g_{12} \\ g_{21} & g_{22} \end{pmatrix}$$

define

$$a = \frac{g_{11} + g_{12} - g_{21} - g_{22}}{4}, \qquad b = \frac{g_{11} + g_{22} - g_{21} - g_{12}}{4}.$$

Then the matrix

$$\begin{pmatrix} b & -b \\ -b & b \end{pmatrix} + \begin{pmatrix} a & a \\ -a & -a \end{pmatrix}$$

has the same payoff differences against any strategy of the opponent. This canonical representation simplifies calculations and is used throughout this section.

For a given G, the selection probabilities for log-linear choice are

$$p_s(0 \mid \phi) = \left(1 + \exp - 2\beta \left\{ a|V_s| + b \sum_{t \in V_s} (2\,\phi(t) - 1) \right\}\right)^{-1}$$

$$p_s(1 \mid \phi = \left(1 + \exp 2\beta \left\{ a|V_s| + b \sum_{t \in V_s} (2\phi(t) - 1) \right\}\right)^{-1}.$$

Thus the flip rate at s is

$$c(s, \phi)$$
$$= \left(1 + \exp 2\beta \left\{a|V_s|(2\phi(s) - 1) + b \sum_{t \in V_s} (2\phi(t) - 1)(2\phi(s) - 1)\right\}\right)^{-1}.$$

These flip rates can be written in the form

$$c(s, \phi) = \left(1 + \exp 2 \sum_{R \ni s} J_R \prod_{t \in R} (2\phi(t) - 1)\right)^{-1}, \qquad (5.1)$$

where

$$J_R = \begin{cases} \beta_a|V_s| & \text{if } |R| = 1, \\ \beta b & \text{if } R = \{s, t\} \text{ for some } t \in V_s, \\ 0 & \text{otherwise,} \end{cases} \qquad (5.2)$$

and R indexes the finite subsets of Z^d. A collection $\{J_R\}$ of real numbers indexed by the finite subsets of Z^d is called a *potential* for the stochastic Ising model with flip rates given by Eq. (5.1).

Although there is no complete characterization of the stationary distributions for stochastic Ising models, much can be learned from a special class of stationary distributions called *Gibbs states.*

DEFINITION 5.1. A probability distribution μ on the space X of configurations is a *Gibbs state* for the potential $\{J_R\}$ of Eq. (5.2) if, for $w \in \{0, 1\}$,

$$\nu(\eta(s) = w|\eta_{-s}) = \left(1 + \exp - 2 \sum_{R \ni s} J_R \prod_{t \in R} (2\phi(t) - 1)\right)^{-1}. \qquad (5.3)$$

As the terminology indicates, Gibbs state is an equilibrium concept that has been extensively employed in physics. Föllmer (1974) has used Gibbs states to describe equilibrium of some stochastic general equilibrium models. In the present context, a Gibbs state is a stationary distribution for the stochastic strategy revision process that has a special property: The spatial distribution of equilibrium play is given by the system of conditional probability systems $\{p_s\}$.

Theorem 5.1. *A probability distribution μ on the space X of configurations is a Gibbs state for the system $\{p_s\}$ of log-linear choice probabilities* *iff*

$$\mu(\eta(s) = w \mid \eta(-s) = \phi(-s)) = p_s(w \mid \phi).$$

The utility of Gibbs states for the analysis of two-strategy games comes from the five facts which are summarized immediately below [see, for instance, Liggett (1985, Chap. 4)]. A translation on the lattice is an operator $\tau_s : X \to X$ such that $\tau_s(\eta)(t) = \eta(t - s)$. A distribution μ is translation-invariant on the lattice Z^d if for every lattice site s, measurable event $A \subset X$ and translation τ_s, $\mu(\tau_s A) = \mu(A)$.

Facts. (1) Gibbs states exist.

(2) Every Gibbs state is a stationary distribution.

(3) All translation-invariant stationary distributions are Gibbs states.

(4) If the lattice dimension d is 1 or 2, every stationary distribution is a Gibbs state.

(5) The set of Gibbs states varies upper hemicontinuously with β in the weak convergence topology.

If S is any subgroup of Z^d of finite index (e.g., $\{s : \Sigma_{i=1}^d s_i = 0 \bmod 2\}$), the set of measures invariant under all the τ_s are called *S-periodic*, and Fact (3) is still true (as is also Fact (10) of Section 4.2). The converse of Fact (3) is false for dimension $d \geq 3$. There exist Gibbs states that are not translation-invariant on Z^d when $d \geq 3$.

It is clear from Eq. (5.3) that, when β is large, the probability of any particular player's choice given the configuration of all other players is highest for those choices, which are best responses to the play of the neighbors. As β grows, this converges to play concentrated on best responses. This leads to a connection between Gibbs states and Nash configurations. Let $\mathcal{N}$ denote the set of Nash configurations.

THEOREM 5.2. *If distribution μ on X is the weak limit of a sequence of Gibbs states $\{\mu_\beta\}$ as $\beta \uparrow \infty$ for a log-linear strategy revision process with payoff matrix G, then μ is a stationary distribution for best-response strategy revision and $\mu(\mathcal{N}) = 1$.*

COROLLARY 5.1. *For every two-by-two lattice game such that $u_1(v, w) = u_2(v, w)$ for all v and w, pure strategy Nash equilibria exist; $\mathcal{N} \neq \emptyset$.*

5.2. *Ergodic Behavior with Log-Linear Strategy Revision*

One notion of ergodicity for Markov processes is the existence of a probability measure ν on X such that for any continuous real-valued function f on X,

$$\lim_{T\to\infty} T^{-1} \sum_{t=0}^{T-1} E(f(\phi_t)|\phi_0) = \int f(\eta)\, d\nu(\eta)$$

for all ϕ_0. A stronger requirement is that

$$\lim_t E(f(\phi_t)|\phi_0) = \int f(\eta)\, d\nu(\eta)$$

for all ϕ_0. For any Feller process on a compact state space, the first property is equivalent to the uniqueness of the stationary distribution. The second property is somewhat stronger. In this paper the word *ergodic* refers to the second property, which the literature alternately calls ergodicity or strong ergodicity.

Here are some facts about the ergodic behavior of the stochastic Ising model. Again see Liggett (1985, Chap. IV).

Facts. (6) If the log-linear strategy revision process is ergodic, then there is a unique Gibbs state.

(7) If $b \geq 0$ there is a β_c (possibly infinite) such that the log-linear strategy revision process is ergodic for all $\beta < \beta_c$ and nonergodic for all $\beta > \beta_c$.

(8) If $b \geq 0$, then the log-linear strategy revision process is ergodic for all β if $a \neq 0$ or if $d = 1$. As $\beta \uparrow \infty$, the unique stationary distribution converges weakly to point mass at $\eta(s) \equiv 0$ if $a > 0$, and to point mass at $\eta(s) \equiv 1$ if $a < 0$.

(9) Every log-linear strategy revision process is ergodic for β near 0.

(10) Let ν_0 denote a *translation-invariant* initial distribution for a log-linear strategy revision process. Then any subsequential limit of the set $\{\nu_t\}_{t\geq 0}$ is a translation-invariant Gibbs state (Künsch, 1984).

(11) If $d = 1$ and μ_∞ is the unique Gibbs state, then there is an $\varepsilon > 0$ such that for each continuous function f: $X \to$ R and initial configuration ϕ_0,

$$|E(f(\phi_t)|\phi_0) - \int f\, d\mu_\infty| \leq B(f) \exp\left\{\frac{-\varepsilon t}{\log t}\right\},$$

where $0 < B(f) < \infty$. If, in addition, $b > 0$, the convergence rate is exponential (Holley 1985, 1987).

In summary, ergodic log-linear strategy revision processes result from "most" games when $b \geq 0$. The ergodic behavior for stochastic Ising models with $b < 0$ is less well understood, but note that the asymptotic theory already discussed for iterative V_s-domination applies to games with $b < 0$. If $a > -b$, strategy 0 is a dominant strategy, and if $b > a$, then strategy 1 is dominant. According to these facts, Example 5.1 gives rise to an ergodic strategy revision process for $x \neq 1$, as does Example 5.2 for $a \geq -1$ and $a \neq 1$. When $|a| = 1$, the strategy revision process will be ergodic on Z^1, but not on Z^d for $d \geq 2$.

Fact (11) discusses rates of convergence in one dimension. The indicator function $1_A^\phi(\cdot)$, which takes the value 1 if $\eta(A) = \phi(A)$ and 0 otherwise, is continuous. Its expectation with respect to distribution μ_∞ is the limit probability of the players in A playing the configuration $\phi(A)$. In one dimension, the limit probability of a particular configuration in a given finite set of players exists, and convergence to it is at the rate given in Fact (11).

5.3. *Examples*

In a number of cases, Gibbs states can be computed. In the following examples, each player plays against his "nearest neighbors"; $V_s = \{t : \|t - s\| = 1\}$. The properties of Gibbs states and sometimes the asymptotics of the process are determined by the parameters a and b of the potential.

EXAMPLE 5.1. Consider the game matrix

$$\begin{bmatrix} 1 & 0 \\ 0 & x \end{bmatrix}.$$

Here $b = (1 + x)/4$ and $a = (1 - x)/4$. When $x < 0$, the dominated strategy is played with decreasing probability as β grows large. Suppose that $x \geq 0$ and fix β. When $d = 1$, the probability that any one given player plays "down" is analytic in x, near 1 for large x, equal to $\frac{1}{2}$ for $x = 1$, and near 0 for x sufficiently small. The same is true when $d = 2$ for sufficiently low β. But when β is sufficiently large, there is a discontinuity in this probability at $x = 1$. The limit as $x \downarrow 1$ of the probability of down is some number $\frac{1}{2} + \varepsilon$, and the limit as $x \uparrow 1$ is $\frac{1}{2} - \varepsilon$. As β grows big, ε goes to $\frac{1}{2}$.

EXAMPLE 5.2. Consider the game matrix

$$\begin{bmatrix} -1 - a & 1 - a \\ 1 + a & -1 + a \end{bmatrix}$$

and let $d = 1$. In the previous example, b was positive, indicating some tendency toward positive correlation of strategy choice among players. For the class of matrices described in this example, $b = -1$, indicating some tendency toward negative correlation of strategy choice among neighboring players. The asymptotic behavior of these stochastic strategy revision processes is very different from those of the preceding example. When b is positive, the value of a only determines which strategy the players will ultimately correlate on. When b is negative, the negative correlation effect can be overcome by an a term of sufficient magnitude. When there is a dominant strategy, $|a| > 1$ and there is a unique Gibbs state. The probability that player s will play down goes to 0 or 1 as β grows, depending on the sign of a. When $|a| < 1$, the behavior of the system is different. Let η_{ud} denote the configuration in which player s plays "up" if s is even, and down otherwise. Let η_{du} denote the configuration in which player s plays up if s is odd, and down otherwise. As β grows, the (unique) limit-invariant measure is that which puts mass $\frac{1}{2}$ on each of the configurations η_{ud} and η_{du}. The system alternately locks in near to all up or near to all down, each with equal probability, independent of the initial condition. The convergence rate, as described by fact (11), is not quite exponential. When $a = 1$, the unique invariant measure is the measure that, on the set of sites $\{-n, \ldots, n\}$, assigns equal conditional probability given any boundary behavior $(\phi(-n-1), \phi(n+1))$ to all configurations that have up at no two consecutive sites. A similar statement holds for the case $a = -1$. In higher dimensions with $|a| < 1$, there are multiple Gibbs states. For each of the two configurations $\eta(s) = (-1)^{s_1+\cdots+s_d}$ and $\phi(s) = (-1)^{1+s_1+\cdots+s_d}$ there is, for large β, a Gibbs state that is (weakly) near point mass on the respective configuration.

5.4 *Coordination Games*

One of the most interesting questions in the population dynamics literature is how (and if) coordination emerges in a population of players, each adapting to their environment. With stochastic-choice strategy revision, coordination emerges as limit behavior of the population over time if β is sufficiently large.

Consider the special case of Example 5.1 where $x = 2$:

$$\begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}.$$

In this game $b = \frac{3}{4}$ and $a = \frac{1}{4}$. In $d = 1$ and $d = 2$, the probability of playing down goes to 1 as $\beta \to +\infty$. Both the configurations "all play down" and "all play up" are Nash configurations, but the limit of boundedly rational play selects only the former, Pareto-superior equilibrium.

This example is particularly interesting because it illustrates how equilibrium selection arises from adaptive behavior of the population. One might hope that letting $\beta \to \infty$ always selects a Pareto-best equilibrium, but this is not always the case.

EXAMPLE 5.3. Consider the game matrix

$$\begin{bmatrix} 5 & 0 \\ 3 & 4 \end{bmatrix}.$$

here $b = \frac{3}{2}$ and $a = -\frac{1}{2}$. As $\beta \to +\infty$, the probability of playing down goes to 1.

In the preceding example the Harsanyi and Selten (1988) risk-dominant equilibrium was selected as β grew large. This is generally true in symmetric games with two strategies. Suppose a game is given with matrix

$$\begin{bmatrix} w & x \\ y & z \end{bmatrix}.$$

This game is a *coordination game* if the two "diagonal" outcomes are Nash equilibria of the two-player game. This entails $w > y$ and $z > x$. A calculation shows that $b > 0$ and that $a > 0$ or $a < 0$ as strategy 0 or 1, respectively, is risk dominant. Consequently [from Fact (8)]:

THEOREM 5.3. *Suppose that G describes a two-person coordination game, and that 0 is the unique risk-dominant equilibrium in the two-person game. Then the log-linear strategy revision process is ergodic and for all players s and initial configurations ϕ_0,*

$$\lim_{\beta \to \infty} \lim_{t \to \infty} \Pr_{\phi_0}(\phi_t(s) = 0) = 1.$$

This conclusion for log-linear strategy revision is the same as that found by Kandori *et al.* (1993) for stochastic perturbations of discrete-time best-response strategy revision in a finite population with global matching. The risk-dominant equilibrium is the equilibrium with the largest basin of attraction in the set of possible mixed strategies. This means that the number of mistakes necessary to leave the risk-dominant equilibrium is larger than the number of mistakes necessary to leave the other equilibrium. Kandori *et al.* (1993) and Young (1993) show that the number of mistakes necessary in transiting from one state of the population to another defines a *stochastic potential* that is minimized by the population dynamic. Theorem 5.3 works for similar reasons. The process acts to maximize a potential function defined in terms of the payoff structure of the game. Configurations with larger potential have larger basins of attraction, and so the process tends to the configurations with the largest basins of attraction, which is the risk-dominant equilibrium configuration.

In one dimension, convergence to equilibrium is at exponential rate, according to Fact (11). Let μ_β denote the unique invariant distribution of a log-linear strategy revision process with parameter β for a coordination game with a unique risk-dominant strategy. There is a $\lambda > 0$ such that for any finite set A of layers and configuration $\eta(A)$ of their play,

$$|\Pr_{\phi_0}(\phi_t(A) = \eta(A)) - \mu_\beta(\phi_t(A) = \eta(A))| \leq B(\eta(A)) \exp - \lambda t.$$

Note that only the multiplicative constant depends on the set of players; the convergence rate λ is independent of the size of the player set and the initial configuration. This is consistent with Ellison's (1992) simulation results that in large finite populations, convergence rates seem to be independent of population size.

5.5. *Perturbed Best-Response Strategy Revision*

Most "large β" asymptotic results that can be demonstrated for the log-linear strategy revision process also hold for the perturbed best-response process. Typical are the results on iterative V_s-domination showing that the iteratively V_s-undominated strategies emerge over time when β is large. A similar result is true for selecting risk-dominant Nash equilibria:

THEOREM 5.4. *Suppose that G describes a two-person two-strategy coordination game and that* 0 *is the unique risk-dominant equilibrium in the two-person game. For the perturbed best-response process, and for all players s and initial configurations* ϕ_0.

$$\lim_{\beta\to\infty} \liminf_{t\to\infty} Pr_{\phi_0}((\phi_t)(s) = 0) = 1.$$

This and other results here are proved by a method of comparing Markov processes. The idea is this: If process A is always more likely to switch toward state ϕ than is process B, and if process A is always less likely to leave state ϕ than is process B, then if process B converges to ϕ, then so does process A. The technique for making this intuition precise is called coupling and is used to prove many of the results in this paper. A coupling is the construction of two or more stochastic processes on the same probability space so as to connect their transition mechanisms. Statements about the behavior of one process can then be inferred from known statements about the behavior of the other process. The basic coupling used in this paper and proofs of the applications are in Appendix 3.

6. GAMES WITH MANY STRATEGIES

When it is applicable, the analysis of the previous section carries over in a straightforward way to those games for which $|W| > 2$. But it does not always apply. The source of the difficulty can be seen in Eq. (5.3). It is not true that any arbitrarily specified collection of conditional probability distributions is consistent in the sense that they derive from a common joint distribution. For symmetric two-by-two games, the log-linear choice rule generates consistent conditional probabilities, but perturbed best-response choice rules typically will not and neither will log-linear choice in games with more than two strategies.

6.1. *Gibbs States*

Gibbs states are defined in terms of a potential just as in the previous section; but the definition of a potential must be enlarged to accommodate more strategies for each player.

DEFINITION 6.1. A *potential* for a stochastic-choice strategy revision process is a collection $\{J_T\}$ of real-valued functions indexed by the finite subsets of T, each with domain $X(T)$, such that for all sites $S \in \mathcal{S}$ and actions $v, w \in W$,

$$\log \frac{p\{w \mid \phi(V_s)\}}{p\{v \mid \phi(V_s)\}} = \sum_{T \ni s} J_T(\phi_s^w(T)) - J_T(\phi_s^v(T)).$$

A Gibbs state for the potential $\{J_T\}$ is a probability distribution μ on the space X of configurations such that, for each finite set T of sites and given boundary condition $\phi(-T)$,

$$\mu\{\eta(T) \mid \phi(-T)\} = \frac{\exp \sum_{R:\, R\cap T \neq \varnothing} J_R(\phi^{\eta}_{R\cap T})}{\sum_{\psi(T)\in X(T)} \exp \sum_{R:\, R\cap T\neq\varnothing} J_R(\phi^{\psi}_{R\cap T})}, \tag{6.1}$$

where ϕ^{ζ}_S is the configuration that takes on the value $\zeta(s)$ at all $s \in S$, and $\phi(s)$ at all $s \notin S$.

This condition extends Eq. (5.3).

When $|W| > 2$, Gibbs states may not exist. Suppose that μ is a Gibbs state for the system of choice probabilities $\{p_s\}$. Then for any player t, configuration ϕ, and actions v, $w \in W$, a calculation shows that $\mu(\phi_t^v)/\mu(\phi_t^w) = p_t(v \mid \phi)/p_t(w \mid \phi)$. For any probability distribution, odds ratios must meet a consistency condition: $\Pr(A)/\Pr(C) = (\Pr(A)/\Pr(B))(\Pr(B)\Pr(C))$. For strategy revision processes, the odds ratio of any two configurations that differ only on a finite set of players can be computed by multiplying together the odds ratios for a sequence of one-player changes. In general, there will be many different sequences of one-player changes that connect the initial and final configuration, and the calculation should give the same answer no matter which route is chosen. For two players s and t, four (not necessarily distinct) actions u, v, w, $x \in W$ and a configuration ϕ, let

$$\phi_{st}^{uv}(r) = \begin{cases} u & \text{if } r = s, \\ v & \text{if } r = t, \\ \phi(r) & \text{otherwise.} \end{cases}$$

The consistency condition works out to be that

$$\frac{p_s(u \mid \phi_{st}^{wx})}{p_s(w \mid \phi_{st}^{wx})} \frac{p_t(v \mid \phi_{st}^{ux})}{p_t(x \mid \phi_{st}^{ux})} = \frac{p_t(v \mid \phi_{st}^{wx})}{p_t(x \mid \phi_{st}^{wx})} \frac{p_s(u \mid \phi_{st}^{wv})}{p_s(w \mid \phi_{st}^{wx})} \tag{6.2}$$

and that this ratio equals $\mu(\phi_{st}^{uv})/\mu(\phi_{st}^{wx})$. The existence of a Gibbs state implies the satisfaction of Eq. (6.2) for all sites s and t and actions u, v, w, $x \in W$. The left-hand side of Eq. (6.2) assumes a path along which first player s switches and then player t. On the right-hand side, the switching order is reversed; first t and then s. This consistency condition turns out to imply all of the other consistency conditions involving infinite numbers of players. In Appendix 2, it is shown that the satisfaction of Eq. (6.2) is sufficient for the existence of Gibbs states.

Example 4.1 describes a process where Eq. (6.2) will not be satisfied. Let ϕ denote the configuration that is identically 0, and consider the two paths $\phi \rightarrow \phi_0^1 \rightarrow \phi_{01}^{12}$ and $\phi \rightarrow \phi_1^2 \rightarrow \phi_{01}^{12}$. Computing along the first path, we would have for any Gibbs state that $\log \mu(\phi_{01}^{12}) - \log \mu(\phi) = 3\beta$. Computing along the second path, $\log \mu(\phi_{01}^{12}) - \log \mu(\phi) = 0$. So no Gibbs state can exist.

Whether or not a given payoff matrix has such distributions is largely a property of the best-response correspondence, even though the definition of the stochastic-choice strategy revision process involves the entire ordering of all the alternatives.

DEFINITION 6.2. The payoff matrices G and $G'$are *strongly best-response equivalent* if there exists numbers $\alpha > 0$ and β_j such that $G'(i, j) = \alpha G(i, j) + \beta_j$.

Contrast this notion with best-response equivalence. One can show that the two payoff matrices G and G' are *best-response equivalent*, that is, they have the same best-response correspondence for the two-person game, if all strong domination relations are preserved, and if the relationship of the definition holds for all strictly undominated rows of G.

DEFINITION 6.3. The matrix G *has a potential* if G is strongly best-response equivalent to a matrix $G' = B + A$, where B is symmetric with 0 column sums, A has column sums equal to 0, and all elements of any row of A are equal. The *potential for G* is:

$$J_T(\phi(T)) = \begin{cases} \beta |V_s| A(\phi(s), \phi(s)) & \text{if } T = \{s\}, \\ \beta\, B(\phi(s), \phi(t)) & \text{if } T = \{s, t\} \text{ and } t \in V_s, \\ 0 & \text{otherwise.} \end{cases} \tag{6.3}$$

Note that the matrices A and B are uniquely determined by G, and that, although the potential depends upon the neighborhoods $\{V_s\}_{s \in Z^d}$, the property of having a potential is independent of the neighborhoods.

The matrices A and B are best understood through the following algorithm used to find them: First, add a number β_j to each element in the jth column of G so that the column sums of the resulting matrix $\hat{G}$ are 0. This transformation always leaves both the best-response correspondence and selection probabilities for log-linear strategy revision unchanged. Next, construct B by subtracting from each element in row i the average payoff in that row of $\hat{G}$: $(1/K)\ \Sigma_j\ \hat{G}(i, j)$. This average is the common element in the ith row of A. The matrix B of deviations is symmetric if and only if G has a potential.

THEOREM 6.1. *The log-linear strategy revision process with parameter β has a Gibbs state μ if and only if the payoff matrix G has a potential.*

The potential in Eq. (6.3) is easy to work with, but it is better understood by an alternative expression in terms of the $\hat{G}$ matrix:

$$J'_T(\phi_T) = \begin{cases} \beta\, \hat{G}(\phi(s), \phi(t)) & \text{if } T = \{s, t\} \text{ and } t \in V_s, \\ 0 & \text{otherwise.} \end{cases} \tag{6.4}$$

Games that have a potential as described here are precisely the symmetric two-person games with a potential in the sense of Monderer and Shapley (1991). All symmetric 2×2 games have potentials, but most symmetric $n \times n$ games do not. The set of $n \times n$ payoff matrices having potentials is a linear subspace of the n^2-dimensional set of payoff matrices with dimension $(n^2 + 3n - 2)/2$.

6.2. *Results of Processes with Gibbs States*

For those payoff matrices for which Gibbs states exist, most of the results in Section 5 carry over. In particular, Facts (2) through (6), (9), and (10) are valid. (I do not know if the rate results in Fact (11) work for more than two strategies.) Again, the existence of Nash configurations is connected to the existence of Gibbs states:

THEOREM 6.2. *Suppose a payoff matrix G has a potential. If distribution μ on X is the weak limit of a sequence of Gibbs states $\{\mu_\beta\}$ as $\beta \uparrow \infty$ for a long-linear strategy revision process with payoff matrix G, then μ is a stationary distribution for best-response strategy revision and $\mu(\mathcal{N}) = 1$. Conversely, if the lattice game has no Nash configurations, then the system $\{p_s\}$ with log-linear choice rules has no Gibbs states for any β.*

In particular, if the lattice game has no pure strategy Nash equilibrium, then Gibbs states will not exist. The converse is obviously false. Take any game with three or more strategies having a potential and a strict Nash equilibrium for the lattice game, such as a pure coordination game. The set of nearby games preserving any particular (strict) Nash configuration has nonempty interior, but the set of games with a potential is lower-dimensional.

Although it is hard to get asymptotic results for stochastic strategy revision processes with $|W| > 2$, Fact (10) does have some strong consequences for games with a unique Gibbs state. And some games with unique Gibbs states are easy to identify. The following result can be viewed as a first approach to many-strategy coordination games, although it also applies to other games.

THEOREM 6.3. *Suppose that G has a potential, and also that $\hat{G}(0, 0) > \hat{G}(i, j)$ for all $(i, j) \neq (0, 0)$. Let ν_0 denote an initial distribution of configurations such that players' initial strategies are i.i.d. Then for log-linear strategy revision and for ν_0-almost all configurations ϕ_0,*

$$\lim_{\beta \to \infty} \liminf_{t \to \infty} \Pr_{\phi_0} (\phi_t(s) = 0) = 1.$$

Proof. The idea of the proof is to work on the space of initial configurations with probability distribution ν_0 and exploit the fact that the limit Gibbs distribution is an extreme point of the set of probability distributions on configurations.

The hypothesis of the Theorem guarantees that the potential has a unique "ground state"—a unique configuration that maximizes the potential among all configurations differing from it at only a finite number of sites. For large enough β, such potentials have a unique Gibbs state, which, in this case, converges to point-mass at the configuration $\phi(s) \equiv 0$ (Georgii, 1988, pp. 147–148). It then follows from Fact (10) that for any i.i.d. initial distribution ν_0, $\lim_{\beta \to \infty} \lim_t \nu_t$ is point mass at the configuration $\phi(s) \equiv 0$. Write

$$\begin{aligned} \nu_t(A) &= \sum_{\eta(B)} \Pr(\phi_t \in A \mid \phi_0(B) = \eta(B))\nu_0(\eta(B)) \\ &= \sum_{\eta(B)} \nu_0(\eta(B))S^*(t)\nu_0(A \mid \phi_0(B) = \eta(B))\nu_0(\eta(B)), \end{aligned}$$

where $S^*(t)$ is the adjoint operator on measures induced by the Markov operator $S(t)$ on the space of continuous real-valued functions of configurations. Since the point mass is an extreme point in the weakly compact and convex set of probability distributions on configurations, it follows that for any finite set of sites B and configuration ϕ_0,

$$\lim_{\beta\to\infty} \liminf_{t\to\infty} \Pr_{\phi_0} (\phi_t(s) = 0 \mid \phi_0(B)) = 1.$$

Let $D = \{\phi_0 : \lim_{\beta\to\infty} \liminf_{t\to\infty} \Pr_{\phi 0} (\phi_t(s) = 0) = 1\}$. Let $\mathcal{S}$ denote the product σ-field on X and let $\mathcal{S}_n$ denote the sub-σ-field generated by the events $X(-B_n) \times A$, where A is a subset of B_n, the cube with center 0, and edge length $2n + 1$. Then $E[1_D(\phi) \mid \mathcal{S}_n] = 1$, and from the martingale convergence theorem $\lim_{n\to\infty} E[1_D(\phi) \mid \mathcal{S}_n] = E[1_D(\phi) \mid \mathcal{S}]\nu_0$-almost surely. Thus $E[1_D(\phi) \mid \mathcal{S}]$ is ν_0-almost surely 1. ∎

This idea can be extended somewhat. Consider Example 5.2 with $|a| <$ 1. There for large β there is a unique translation-invariant Gibbs state that converges to $(\frac{1}{2})\delta_{\eta_{du}} - (\frac{1}{2})\delta_{\eta_{ud}}$. Using similar arguments, one can see that for any pair of players $\{s, s + 1\}$,

$$\lim_{\beta\to\infty} \liminf_{t\to\infty} \Pr (\phi_t(\{s, s + 1\}) \in \{(0, 1), (1, 0)\} \mid \phi_0(B)) = 1.$$

Note that the asymmetric Nash configurations discussed in Section 4 are not selected for as $\beta \to \infty$. Nonetheless, the application of this kind of argument is limited, and in general there is not much to say about the evolution of these processes from non translation-invariant initial conditions. There are two difficulties. First, when noninvariant Gibbs states exist, many Gibbs states exist. The set of stationary distributions must be large, even in one and two dimensions where the only stationary distributions are Gibbs states, and there is no hope of knowing where the process goes. However, the existence of noninvariant Gibbs states is a bit delicate. Certainly the hypothesis of Theorem 6.3 must fail. For 2×2 games, if $b \geq 0$, this requires $a = 0$. If $b < 0$, this requires that $|a|$ be sufficiently small. But even when there is a unique Gibbs state, a second problem arises. The arguments that prove ergodicity rely on translation invariance in a fundamental way (except in the case of a one-dimensional lattice). We are up against a constraint of technique. Even so, Fact (10) is informative. Theorem 6.3 and the next theorem demonstrate how order can arise from an initial disordered state. Fact (10) says that, when these conditions fail to hold, it is still true that disorder does not arise out of order. Moreover, in dimensions 1 and 2 where every stationary distribution is Gibbs, whenever there is a unique Gibbs state, time averages of the expectation of any continuous function of configurations, given the initial configuration, converge to the expectation of the function with respect to the Gibbs state.

6.3. *Coordination Games and Perturbed Best-Response Strategy Revision*

The analysis of coordination games with two strategies is exactly that of the attractive stochastic Ising models. The analysis of coordination games with an arbitrary finite number of strategies proceeds by coupling a many-state system to a two-state system. The results on coordination games are *not* restricted only to games with a potential. The theorem states a generalization of the Harsanyi–Selten risk-dominance criterion for two-by-two games, which is sufficient for asymptotic equilibrium selection in log-linear strategy revision processes.

DEFINITION 6.4. A symmetric two-person game with payoff matrix G is a *coordination game* if, for all pure strategies v and w, $g_{vv} > g_{wv}$.

In coordination games, each diagonal element is a Nash equilibrium outcome, and there are no pure strategy Nash equilibrium outcomes other than the diagonal elements.

THEOREM 6.4. *Suppose that G describes a two-person coordination game and that $\min_{k\geq 1}(G_{00} - G_{k0}) > \max_{k\geq 1}(G_{kk} - G_{0k})$. Then for both log-linear strategy revision and perturbed best-response strategy revision,*

$$\lim_{\beta\to\infty} \liminf_{t\to\infty} \Pr_{\phi_0}(\phi_t(s) = 0) = 1$$

for each initial configuration ϕ_0 and site s.

There will be games to which Theorem 6.3 applies that fail to satisfy the condition of Theorem 6.4. The importance of Theorem 6.4 lies in the fact that it addresses payoff matrices that do not have potentials.

The existence of a potential is not robust to the elimination of V_s-dominated strategies, but the results of Section 3 imply that their presence should not matter for determining large-β asymptotics. The following theorem is a straighforward application of the coupling technique discussed in Appendix 3. It states that mere best-response equivalence (rather than strong best-response equivalence) to a potential game suffices to characterize asymptotic behavior in terms of the potential.

THEOREM 6.5. *Suppose that payoff matrix G is best-response equivalent to matrix G', which has a potential. Let D denote an event that depends upon the configuration of only a finite number of players, and consider both log-linear and perturbed best-response strategy revision.*

For an initial distribution ν_0, if

$$\lim_{\beta\to\infty}\lim_{t\to\infty}\Pr_{\nu_0}(\phi_t \in D) = 1$$

for the process with matrix G', then for the process with matrix G,

$$\lim_{\beta\to\infty}\liminf_{t\to\infty}\Pr_{\nu_0}(\phi_t \in D) = 1.$$

7. Conclusion

With respect to the emergence of equilibrium in coordination games, my work is similar in spirit to Kandori *et al.* (1993) and Ellison (1992). Kandori *et al.* consider stochastic perturbations of a discrete-time deterministic population dynamic in two-strategy games that has the "Darwinian property," satisfied by best-response, that the frequency of a strategy in the population increases at date $t + 1$ if and only if it is the unique best-response to the current distribution of play. I motivate my model without appealing to a matching story, but if built on a finite lattice, the "global interaction" version of my perturbed best-response process is exactly a continuous-time version of the Kandori *et al.* model with best-response strategy revision with the one difference that, in my processes, players have randomly arriving revision opportunities.

In two-by-two coordination games I find the same selection principle at work that these authors do—selection for the risk-dominant equilibria. Section 6 contains some extensions to $n \times n$ games that are easily accessible with my analytical apparatus. My focus on local rather than global interactions is similar to Ellison's, and my analytical apparatus gives information on convergence rates. Ellison (1992) has shown that, in the Kandori *et al.* model with best-response strategy revision, rates of convergence for local interaction exceed those for global interaction. Furthermore, rates of convergence shrink as the radius of interaction grows. In the models presented here, the same phenomenon arises in a more dramatic fashion. I have computed some examples showing that in two-strategy coordination games with a unique risk-dominant Nash equilibrium and perturbed best-response strategy revision, where switch rates depend upon the average behavior of the population, the strategy revision process can fail to be ergodic. Both equilibria are possible limits, and limit behavior is completely determined by the initial distribution of configurations. This behavior difference is quite striking. It suggests that the global equilibrium selection results of Kandori *et al.* have an unsuspected delicacy. However, the general comparison of local versus global interaction in strategy revision processes is still an open question.

The research presented here raises other interesting questions. One is the study of best-response strategy revision. For the strategy revision process of Example 5.1 with $x = \frac{1}{2}$, log-linear strategy revision is ergodic for all β, and both perturbed best-response and log-linear strategy revision select the equilibrium in which all players choose up as β gets large. On the other hand, simulation of the simple best-response strategy revision process shows that the outcome depends strongly on the initial conditions of the process. If the initial frequency of down is sufficiently high, the process converges to all players choosing down. Although this behavior is very plausible, as yet there is no theoretical explanation for it.

Another question has to do with the specifics of the strategic interaction. I have assumed that the value of an alternative to any player is the sum of the values of the simultaneous interactions with each of his neighbors. The matching models so popular in much of evolutionary game theory do not work this way; there the value of an alternative depends upon the current match. It seems likely that the behavior of strategy revision processes is very sensitive to the interaction technology. Different interaction technologies lead to different kinds of infinite particle systems. This is good news, since the stochastic Ising models and their generalizations that have driven much of the analysis in this paper are particularly hard to work with. Other appealing matching technologies lead to percolation processes and processes akin to the voter model, which are easier to grasp.

Game theorists have already recognized that evolutionary forces govern population behavior. The next step is to move from a broad description of evolutionary outcomes to a more specific understanding of the mechanics of how the evolutionary forces work. In all likelihood, this will require a micro-level examination of the processes of strategic interaction similar, at least in spirit, to Ellison's model and the models presented here.

Appendix 1: Proofs

A stochastic strategy revision process is a continuous-time Markov process of strategic choices of the players as described in Section 2. Formally, the state space for the process is the space of configurations of the population, $X = \Pi_{s\in Z^d} W_s$, where W_s has the discrete topology, and X the product topology and its associated Borel σ-field.

A Markov strategy revision process describes a Markov semigroup $S(t)$: $C(X) \to C(X)$, where $C(X)$ denote the space of continuous functions on X (with the sup-norm topology). They are characterized by their semigroup, which is in turn characterized by its *infinitesimal generator*. Let $C_F(X)$ denote the subset of $C(X)$ containing those functions whose values depend on only finitely many coordinates: f is in $C_F(X)$ if there is a finite set T of sites such that, if $\phi(T) = \eta(T)$, then $f(\phi) = f(\eta)$. The set $C_F(X)$ is dense in $C^*(X)$. Let Ω: $C_F(X) \to C_F(X)$ be the linear map given by

$$\Omega f(\eta) = \sum_{s\in S} \sum_{w\in W} (f(\eta_s^w) - f(\eta))p_s(w \mid \eta(V_s)). \tag{A.1}$$

Straightforward calculations show that $\overline{\Omega}$, the closure of Ω in $C(X)$, is a generator of $S(t)$. Thus, for every initial distribution μ of configurations, there exists a Markov process corresponding to $S(t)$ with initial distribution μ and sample paths that are right-continuous with left limits and that has the strong Markov property with respect to the usual filtration (Ethier and Kurtz, 1986, Theorem 4.2.7).

The analysis of stochastic strategy revision processes depends on an additional concept, that of reversible distributions. A distribution μ on the state space of a Markov process is said to be *reversible* if the sample paths of the two processes $\{X(t), -\infty < t < \infty\}$ and $\{X(-t), -\infty < t < \infty\}$ have the same distribution. The semigroup characterization of reversibility is that for all $f, g \in C(X)$, $\int f S(t) g \, d\mu = \int g S(t) f \, d\mu$. From this characterization it is clear that if the distribution μ is reversible, then it is also stationary. In terms of the generator Ω of the semigroup, μ is reversible if and only if $\int f \Omega g \, d\mu = \int g \Omega f \, d\mu$. Reversible distributions are used to prove Theorems 6.1 and 6.2.

Proof of Theorem 4.1. Let $1_S^\phi(\eta)$ denote the function that takes on the value 1 when $\eta(S) = \phi(S)$ and 0 otherwise. These indicator functions are continuous and span the space of continuous functions with finite support, so $\int \Omega_x 1_S^\phi(\eta) \, d\mu(\eta) = 0$ for all such functions if and only if μ is stationary. Computing, this condition is

$$\int \sum_j \Pr_\mu (M(\eta, s_j) = \phi(s_j), \eta(s_j) \neq \phi(s_j), \eta_{S-j} = \phi_{S_j}) \\ - \Pr_\mu (M(\eta, s_j) \neq \phi(s_j), \eta(s_j) = \phi(s_j), \eta_{S-j} = \phi_{S_j}) \, d\mu(\eta) = 0.$$

This and algebraic manipulation of the statement

$$\Pr_\mu(\eta_S = \phi_S) = \Pr_\mu (M(\eta, s_j) = \phi(s_j), \eta_S = \phi_S) + \Pr_\mu (M(\eta, s_j) \neq \phi(s_j), \eta_S = \phi_S)$$

give

$$\sum_j \Pr_\mu(\eta_S = \phi_S) - \Pr_\mu(M(\eta, s_j) = \phi(s_j), \eta_S = \phi_S = 0,$$

which is Eq. (4.1), and conversely. ■

The analysis of this and the next appendix uses the following lemma:

LEMMA A.1. *The following statements are equivalent for a stochastic-choice strategy revision process:*

1. *μ on X is a reversible distribution;*
2. *for all $\phi \in X$,*

$$\mu\{\eta : \eta(s) = \phi(s) \mid \eta(t) = \phi(t) \text{ for } t \neq s\} = p_s\{\phi(s) \mid \phi(V_s)\}; \tag{A.2}$$

3. *for all $f \in D(X)$* and for each site $s \in Z^d$,

$$\int_\eta \sum_{w \in W} [f(\eta_s^w) - f(\eta)] p\{w \mid \eta(V_s)\} \, d\mu(\eta) = 0. \tag{A.3}$$

Condition 2 and Theorem 5.1 imply that the reversible distributions are the Gibbs states. Condition 3 is sometimes known as *detailed balance*. The proof of this Lemma follows along the lines of Liggett (1985, Proposition IV.2.7) and so is not presented here.

Proof of Theorems 5.2 *and* 6.2. Theorem 5.2 follows from Theorem 6.2 since all symmetric two-person games have a potential. To prove Theorem 6.2, consider first best-response strategy revision. If ϕ is a Nash configuration, a check of the definition shows that the point mass δ_ϕ is a reversible distribution. It follows from the detailed balance condition that if μ is reversible, then for all players $s \in S$,

$$\int \sum_{t \in V_s} \left[\sum_{w \in M(\phi, s)} G(w, \phi(t)) - G(\phi(s), \phi(t)) \right] p_s(w \mid \phi)\mu(d\phi) = 0.$$

(Recall that $M(\phi, s)$ is the set of best responses to configure ϕ by player s.) Each term inside the square brackets is nonnegative. For the left-hand side to equal 0, each term on the right must be 0 almost surely, and so we have that almost surely $\phi(s) \in M(\phi, s)$. Since Z^d is countable,

$$\mu\{\phi : \phi(s) \in M(\phi, s) \quad \text{for all } s \in Z^d\} = 1. \quad \blacksquare$$

Thus the reversible distributions for best-response strategy revision are precisely those that put all their mass on the set of Nash configurations.

It is not hard to show that for log-linear strategy revision, the correspondence that maps parameters (β, G) to reversible distributions is upper hemicontinuous. Lemma A.1 states that the reversible distributions are the Gibbs states, which proves the first statement of the theorem. Conversely, if $\mathcal{N} = \phi$, then for large β, Gibbs states do not exist. This means that the consistency conditions (6.2) fail. But these conditions are independent of β, so for no β do Gibbs states exist. $\blacksquare$

Appendix 2: Existence of Gibbs States

In this appendix, I discuss the existence of Gibbs states for log-linear strategy revision when $|W| > 2$. Not every game has Gibbs states. Section 5 notes that the Eq. (6.2) is necessary for the existence of Gibbs states. In this appendix, I show that this "path independence" condition is sufficient as well.

THEOREM A.1.: *Let $\{p_s\}$ denote a system of choice probabilities for a log-linear strategy revision process such that for all $w \in W$ and each player s and configuration ϕ, $p_s(w \mid \phi) > 0$. Then a Gibbs state exists for $\{p_s\}$ if and only if for all players s and t, any four actions $u, v, w, x \in W$ and any configuration $\phi \in X$, Equation* (6.2) *holds.*

This condition is basically Kolmogorov's condition for the reversibility of a Markov chain.

Choice probabilities are invariant to the addition of a column-specific term to each element of the matrix G: If $G'(v, w) = G(v, w) + h(w)$, then G and G' give rise to identical choice behavior in Eq. (2.3). Thus with no loss of generality we can assume in the following theorem that the column sums of G are 0: $\Sigma_w G(v, w) = 0$ for all $v \in W$.

THEOREM A.2. *The matrix G admits a reversible distribution with linear f if and only if the matrix $H(v, w)$ given by*

$$H(v, w) = G(v, w) - \frac{1}{K} \sum_{w' \in W} G(v, w')$$

is symmetric.

The proofs of these theorems (and that of Theorem 6.1) rely heavily on the symmetry of the neighborhood structure—that the neighborhood relation is symmetric. When externalities are directed in some way, upstream players can affect downstream players, but the converse is false. Effectively, each player has only upstream neighbors. It will be evident from the proof of Theorem A.1 that in such a case no reversible measures can exist.

Proof of Theorem A.1. Let $B(n) = [-n, n]^d$ denote the cube of side-length $2n + 1$ centered at the origin of Z^d, and fix a configuration $\eta(B(n)^c)$ off of B_n.

DEFINITION A.1. A collection $p\{\cdot \mid \cdot\}$: $W \times X[B(n) \cup \partial B(n)] \to [0, 1]$ of conditional probabilities is *consistent* on $B(n)$ if for each configuration $\eta(B(n)^c)$ of sites not in $B(n)$ there is a probability distribution $\mu_n\{\cdot\}$ on $X[B(n)]$ such that for all $\phi \in X(B(n))$, s in $B(n)$ and $v \in W$,

$$\mu_n\{\phi(s) = v \mid \phi(B(n)/\{s\})\} = p\{v \mid (\phi(t))_{t \in V_s \cap B(n)}, (\eta(t))_{t \in V_s \cap \partial B(n)}\}.$$

According to Lemma A.1, a probability distribution μ on configurations is reversible if and only if, at any site s, the conditional probability distribution on the actions at s given the configuration of s is equal to the selection probability distribution $p\{\phi(s) \mid \phi(V_s)\}$. This implies that the system of conditional probabilities described by p is consistent on each $B(n)$. Conversely, according to Preston (1973, Theorem 5.3), if the conditional probabilities are consistent on each $B(n)$, then there is probability distribution μ on X that satisfies Eq. (A.2) and so, from Lemma A.1, is reversible. Thus the existence problem for reversible probability distributions is equivalent to the consistency of the conditional probabilities p on each $B(n)$.

It is easy to see that the collection of conditional probabilities p given $\eta(\partial B(n))$ describes the conditional probabilities from some distribution μ_n on $B(n)$ if, and only if,

$$\begin{aligned} \log \mu_n\{\phi_s^v\} - \log \mu_n\{\phi_s^w\} &= \log p\{v \mid \phi(V_s \cap B(n)), \eta(V_s \cap \partial B(n))\} \\ &\quad - \log p\{w \mid \phi(V_s \cap B(n)), \eta(V_s \cap \partial B(n))\}. \end{aligned} \tag{A.4}$$

By chaining together configurations that differ at only one site, we define odds ratios from which we can construct a distribution.

LEMMA A.2. *Let X be a finite set, and suppose that for every pair of elements $x, y \in X$ a number $l(x, y)$ is given. Suppose that $l(x, y) = -l(y, x)$, and suppose that for all x, y and z, $l(x, y) + l(y, z) = l(x, z)$. Then and only then does there exist a probability distribution p on X such that for all x and y, $l(x, y)$ is the log-odds ratio* $\log p(x) - \log p(u)$.

Proof. This is a simple computation. ∎

DEFINITION A.2. Let ψ and ϕ denote two configurations of $B(n)$. A *path* from ψ tc ϕ is a finite sequence of configurations $\psi_0, \ldots, \psi_K$ such that

1. $\psi_0 = \psi$ and $\psi_K = \phi$;
2. ψ_k and ψ_{k-1} differ only at a single site $t_k \in B(n)$.

A *circuit* is a path such that $\psi = \psi_0 = \psi_K$.

LEMMA A.3. *Suppose that for every pair of configurations θ and ζ of $B(n)$, which differ at only one site, a number $l(\theta, \zeta)$ is given. There exists a probability distribution μ_n on $B(n)$ such that $l(\theta, \zeta) = \log \mu_n(\theta) - \log \mu_n(\zeta)$ if and only if, for any circuit $\psi_0, \ldots, \psi_K$, the sum $\Sigma_{k=1}^K l(\psi_k, \psi_{k-1}) = 0$.*

Proof. This condition is clearly necessary. To check sufficiency, suppose that the condition is true. Then if ψ and ϕ are any two configurations connected by the path $\psi_0, \ldots,$

ψ_K, define $l(\psi, \phi)$ to be the sum $\Sigma_{k=1}^{K} \log \mu(\psi_k) - \log \mu(\psi_{k-1}) = 0$. This sum is independent of the path chosen, for otherwise, by traversing from ψ to ϕ along one path and back along another for which the sum differs, we would have constructed a circuit for which the sum is not 0. Now it is easy to check that the odds ratios satisfy the conditions of Lemma A.2. ■

Next is a characterization of path independence.

LEMMA A.4. *Suppose that for every pair of configurations θ and ζ of $B(n)$, which differ at only one site, a number $l(\theta, \zeta)$ is given. The sum $\Sigma_{k=1}^{K} l(\psi_k, \psi_{k-1}) = 0$ along any circuit if and only if, for any configuration θ, sites s and t, and choices v_s, w_s, v_t, and w_t, such that $\theta(s) = v_s$ and $\theta(t) = v_t$,*

$$l(\theta_s^{w_s}, \theta) + l(\theta_{s,t}^{w_s, w_t}, \theta_s^{w_s}) = l(\theta_t^{w_t}, \theta) + l(\theta_{s,t}^{w_s, w_t}, \theta_t^{w_t}).$$

Theorem A.1 follows from Eq. (A.4) and Lemmas A.3 and A.4. Suppose that the collection

$$p\{\cdot \mid \cdot\}: W \times X[B(n) \cup \partial B(n)] \rightarrow [0, 1]$$

of conditional probabilities is consistent on $B(n)$. Then for configurations θ and ζ, which are identical except at one site, the log-odds $l(\theta, \zeta)$ must be given by Eq. (A.4), and the conditions of the theorem now follow from Lemmas A.3 and A.4. Conversely, suppose that the log differences

$$\log p\{v \mid \phi(V_s \cap B(n)), \eta(V_s \cap \partial B(n))\} - \log p\{w \mid \phi(V_s \cap B(n)), \eta(V_s \cap \partial B(n))\}$$

satisfy the hypotheses of the theorem. For θ and ζ, which are identical except perhaps at site s, define

$$l(\theta, \zeta) = \log p\{\theta(s) \mid \zeta(V_s \cap B(n)), \eta(V_s \cap \partial B(\eta))\} - \log p\{\zeta(s) \mid \zeta(V_s \cap B(n)), \eta(V_s \cap \partial B(n))\}.$$

Then the numbers $l(\theta, \zeta)$ satisfy the condition of Lemma A.4, so the sum along any circuit is 0, and the result follows from Lemma A.3.

The proof of Lemma A.4 is too long to include here, but the idea is very simple. Consider a circuit $\psi_0, \ldots, \psi_L$, and suppose that the condition of Lemma A.4 is satisfied. If a loop like that described by the lemma's condition is added or subtracted, the sum of the terms $l(\psi_k, \psi_{k-1})$ remains unchanged. By making such modifications, the circuit can be retracted to a constant circuit $\psi_0 = \psi_1 = \cdots$. The sum along the constant circuit is 0, so the sum along the original circuit must be 0 too.

Proof of Theorem A.2. Define the matrix $A = G - H$. Then $A_{v,w}$ is the average of the elements of the vth row of G. Since the columns of G sum to 0, so too do the columns of H and A. Since the elements of H are deviations from the row averages, the row sums of H are all 0. The conditions of Theorem A.1 gives, in this case,

$$(G(w, v') - G(v, v')) + (G(w', w) - G(v', w)) = (G(w', v) - G(v', v)) + (G(w, w') - G(v, w'))$$

for all v, w, v', and w' in W and $t - s \in V_0$. Thus

$$(H(w, v') - H(v, v')) + (H(w', w) - H(v', w)) = (H(w', v) - H(v', v)) + (H(w, w') - H(v, w')).$$

(Note that the A terms cancel out.) Summing over w and dividing by K gives

$$-H(v, v') = (H(w'v) - H(v', v)) - H(v, w').$$

Summing over w' and dividing by K gives

$$H(v, v') = H(v', v),$$

which proves the theorem. ■

Proof of Theorems 5.1 *and* 6.1. These Theorems follow immediately from the identification of Gibbs states with reversible distributions. Condition 2 of Lemma A.1 establishes that a distribution is reversible if and only if the conditional probabilities for behavior by any single player are those given by the potential from Eq. (6.1). Thus any Gibbs state is reversible. For any reversible distribution, these one-player conditional distributions uniquely determine the conditional distributions on all finite-player sets through the consistency conditions, and a calculation shows that these are exactly those given by Eq. (6.1). ■

Appendix 3: Coupling

This appendix details the coupling methods used to prove Theorems 3.2, 5.4, and 6.4. Coupling is a method of building two stochastic processes on the same probability space so as to compare their sample-path behavior.

A.3.1. *The Construction*

Given choice probabilities $p_s^1(\cdot \mid \cdot), p_s^2(\cdot|\cdot)$ and "alarm clock" rates $\lambda_1(\cdot)$ and $\lambda_2(\cdot)$, let

$$\bar{p}_s(w \mid \eta, \zeta) = \min\{\lambda_1(\eta)p_s^1(w \mid \eta), \lambda_2(\zeta)p_s^2(w \mid \zeta)\}.$$

For each s, let $D_s \subset W$. Define the sets

$$A = S/B \cup C,$$

where

$$B = \{s : \eta(s), \zeta(s) \in D_s\}$$
$$C = \{s : \eta(s), \zeta(s) \notin D_s\}.$$

Define

$$\begin{aligned}\bar{\Omega} f(\eta, \zeta) = &\sum_{s\in A} \sum_{w} \lambda_1(\eta)p_s^1(w \mid \eta)(f(\eta_s^w, \zeta) - f(\eta, \zeta)) \\ &+ \sum_{s\in A} \sum_{w} \lambda_2(\zeta)p_s^2(s \mid \zeta)(f(\eta, \zeta_s^w) - f(\eta, \zeta)) \\ &+ \sum_{s\in B\cup C} \sum_{w} p^-{}_s(w \mid \eta, \zeta)(f(\eta_s^w, \zeta_s^w) - f(\eta, \zeta))\end{aligned} \tag{A.5}$$

$$+ \sum_{s \in B \cup C} \sum_{w} (\lambda_1(\eta) p_s^1(w \mid \eta) - \bar{p}_s(w \mid \eta, \zeta))(f(\eta_s^w, \zeta) - f(\eta, \zeta))$$

$$+ \sum_{s \in B \cup C} \sum_{w} (\lambda_2(\zeta) p_s^2(w \mid \zeta) - \bar{p}_s(w \mid \eta, \zeta))(f(\eta, \zeta_s^w) - f(\eta, \zeta)).$$

Define $\eta \leq D\, \zeta$ if $\eta(s) \in D_s$ implies $\zeta(s) \in D_s$. Let $K_D = \{(\eta, \zeta) : \eta \leq D\, \zeta\}$.

THEOREM A.3. *Suppose that whenever $\eta \leq D\, \zeta$,*

$$\lambda_2(\zeta) p_s^2(w \mid \zeta) \leq \lambda_1(\eta) p_s^1(w \mid \eta) \qquad \text{for } w \notin D_s,\ \eta(s), \zeta(s) \in D_s$$

$$\lambda_2(\zeta) p_s^2(w \mid \zeta) \geq \lambda_1(\eta) p_s^1(s \mid \eta) \qquad \text{for } w \in D_s,\ \eta(s), \zeta(s) \in \Sigma / D_s.$$

Then for all $(\eta, \zeta) \in K_D$ and $t \geq 0$

$$\Pr((\eta_t, \zeta_t) \in K_D \mid (\eta_0, \zeta_0) \in K_D) = 1.$$

Proof. The proof of this theorem follows the proof of Ligget (1985, Theorem III.1.5). ■

Let $\mathcal{M}_D$ denote the set of all monotone continuous real-valued functions on X; that is, $f \in \mathcal{M}_D$ iff $\eta \leq D\, \zeta$ implies $f(\eta) \leq f(\zeta)$.

DEFINITION A.3. If μ_1 and μ_2 are probability measures on X, then $\mu_1 \leq D\, \mu_2$ iff $\int f\, d\mu_1 \leq \int f\, d\mu_2 \leq$ for all $f \in M$.

COROLLARY A.1. *Under the assumptions of Theorem A.3, if μ_1 and μ_2 are probability measures on X such that $\mu_1 \leq D\, \mu_2$, then $\mu_1 S_1(t) \leq D\, \mu_2 S_2(t)$.*

COROLLARY A.2. *If the process with semigroup S_1 is ergodic with stationary distribution μ, then for any $f \in M$ and distribution ν, $\int f\, d\mu \leq \liminf_{t \to \infty} \int S_2(t) f\, d\nu$.*

A.3.2. *Applications*

The coupling just discussed has several applications: Coordination games, iteratively V_s-undominated strategies, and stochastic perturbations of best-response strategy revision.

A.3.2.1. *Proof of Theorem* 6.4. First I prove the theorem for log-linear strategy revision. Rescale the payoff matrix. For all i let

$$H_{ij} = \begin{cases} G_{ij} & \text{if } j = 0, \\ G_{ij} + G_{11} - G_{ij} & \text{if } j > 0. \end{cases}$$

The matrix H has the same selection probabilities as does G, $H_{00} = G_{00}$, $H_{jj} = G_{11}$ for $j > 0$.

Take $D_s = \{0\}$, let $\{p_s^2\}$ denote the selection probabilities derived from the stochastic strategy revision process with parameter β and payoff matrix H, and let $\{p_s^1\}$ denote the selection probabilities derived from the stochastic strategy revision process with parameter β and payoff matrix $\hat{H}$, where

$$H_{ij} = \begin{cases} H_{00} & \text{if } i = j = 0 \\ \max_{k \geq 1} H_{k0} & \text{if } i \geq 1, j = 0, \\ \min_{k \geq 1} H_{0k} & \text{if } i = 0, j \geq 1, \\ H_{11} & \text{if } i, j \geq 1. \end{cases}$$

Computations show that this is a successful coupling in that Theorem A.3 applies. Consequently,

$$\liminf_{t\to\infty} \Pr^1(\phi_t(s) = 0) \leq \liminf_{t\to\infty} \Pr^2(\phi_t(s) = 0)$$

for all β. The point of this coupling is that the process with selection probabilities $\{p_s^1\}$ is essentially a stochastic Ising model, because the probability of selecting strategy 0 depends only on the fraction on neighbors that are 0 and neighbors that are not. Consider this process, and define the new process $\{\phi_t\}$ as follows. The state space is the space of all configurations ϕ: $Z^d \to \{0, 1\}$. Let $\phi_t(s) = 0$ iff $\eta_t(s) = 0$. The $\{\phi_t\}$ process is a stochastic Ising model. A computation shows that the potential is given by Eq. (5.2), where

$$\alpha = \tfrac{1}{4}\left(\min_{k\geq 1}\{g_{00} - g_{k0}\} - \max_{k\geq 1}\{g_{kk} - g_{0k}\} + \frac{2}{\beta|V_s|}\log(|W| - 1)\right)$$

$$b = \tfrac{1}{4}\left(\min_{k\geq 1}\{g_{00} - g_{k0}\} + \max_{k\geq 1}\{g_{kk} - g_{0k}\}\right)$$

For all coordination games, $b > 0$. By hypothesis, $\alpha > 0$ for large β. As β grows large, the set of Gibbs states converges to those (for corresponding β) of the stochastic Ising model with potential $b' = b$ and $a' = (\frac{1}{4})(\min_{k\geq 1}\{g_{00} - g_{k0}\} - \max_{k\geq 1}\{g_{kk} - g_{0k}\})$. By Fact (8), the $\{\phi_t\}$-process is ergodic, and its limit distribution converges to point mass on the configuration $\phi(s) \equiv 0$ as β grows large. The events $\{\phi(s)_t = 0\}$ and $\{\eta(s)_t = 0\}$ are identical for all t, so for each site s,

$$\lim_{\beta\to\infty}\lim_{t\to\infty} \Pr\{\eta(s)_t = 0\} = 1.$$

The result follows from Corollary A.2. The result for perturbed best-response strategy revision follows then from theorem A.3. below. ■

A.3.2.2. *Proof of Theorem* 3.2. Let the system $\{p_s\}$ be given, and let $V \subset W$ denote the set of iteratively V_s-undominated strategies. To prove Theorem 3.2, assume w.l.o.g. that $s = 0$. Choose a large cube $B \subset Z^d$ around 0 as in the proof of Theorem 3.1, a fixed strategy $v \in W$, and define the system $\{q_s\}$ as follows:

$$q_s(w|\eta) = \begin{cases} p_s(w|\eta) & \text{if } s \in B \text{ or } w = v, \\ 0 & \text{otherwise.} \end{cases}$$

Note that q_s has an alarm-clock rate that is different from 1 and is η-dependent. The sites in B behave under $\{q_s\}$ just as they do under $\{p_s\}$. Outside B, the sites freeze once they reach v. Let $\tau_0 = \inf\{t : \eta_t(\partial B) \equiv v\}$, and let τ_k denote the kth time after τ_0 that one of the Poisson alarm clocks inside B rings. The intervals $(\tau_k - \tau_{k-1})$ are independently and exponentially distributed with mean $1/|B|$. The process $\{\eta_{\tau_k}(B)\}_{k\geq 0}$ is a Markov chain with transition matrix Q_B^β. Because each $q_s(w \mid \eta) > 0$, it is irreducible and aperiodic, and therefore ergodic with stationary distribution μ_B. Thus the $\{q_s\}$-process is ergodic with a limit distribution whose projection onto $X(B)$ is μ_B.

The system $\{q_s\}$ can be coupled to the system $\{p_s\}$ as in Theorem A.3, where $D_s = W/\{v\}$ for $s \notin B$, $D_0 = V$, and $D_s = W$ for all remaining sites (taking $p^1 = q$ and $p^2 = p$.) The indicator function $f_0^V(\eta)$, which takes on the value 1 if $\eta(0) \in V$ and 0 otherwise, is in $\mathcal{M}$, so for any initial distribution ν on X, $\liminf_t \int S(t)f_0^V(\eta)\, d\nu \geq \int f_0^V(\eta)\, d\mu_B$.

It remains to see that given $\in > 0$, $\int f_0^V(\eta)\, d\mu_B > 1 - \varepsilon$ for all β sufficiently large. Consider the matrix Q_B^∞. Every irreducible recurrent class has $\eta(0) \in V$, so for any stationary distribution, $\Pr(\eta(0) \in V) = 1$. Since the stationary distribution correspondence is u.h.c, $\Pr(\eta(0) \in V) > 1 - \varepsilon$ for β sufficiently large. ■

A.3.2.3. *Coupling to other Perturbed Best-Response Models.* A general specification of stochastic choice involves stochastically perturbing best-response choice. Let $\{p_s\}$ denote a system of choice probabilities for best-response dynamics. Let $\{q_s\}$ denote arbitrary selection probabilities depending on the same neighborhoods V_s and such that $\min_{\eta,w} q_s(w \mid \eta) > 0$. Suppose that, like $\{p_s\}$, these selection probabilities are shift invariant. Let $\{r_s\}$ denote the system with

$$r_s(w \mid \eta) = (1 - \varepsilon)p_s(w \mid \eta) + \varepsilon\, q_s(w \mid \eta).$$

Let $\mathrm{Br}_s(\eta)$ denote the set of best responses to η at site s; the set of $w \in W$ for which $p_s(w \mid \eta) = 1$.

Suppose that ε is sufficiently small that $r_s(v \mid \eta) > r_s(w \mid \eta)$ for all $v \in \mathrm{Br}_s(\eta)$ and $w \notin \mathrm{Br}_s(\eta)$. Let $\{p_s^\beta\}$ denote the system with log-linear stochastic choice, with sensitivity parameter β. At $\beta = 0$ all choices are equally probable, so for all η,

$$r(w \mid \eta) \begin{Bmatrix} > \\ < \end{Bmatrix} p^0(w \mid \eta) \qquad \text{for } w \begin{Bmatrix} \in \\ \notin \end{Bmatrix} \mathrm{Br}_s(\eta).$$

Let $\bar\beta$ denote the largest β for which this is true for all η. Clearly, as ε goes to 0, $\bar\beta$ goes to ∞.

For $\beta = \infty$, log-linear dynamics correspond to best-response dynamics, so for all η,

$$r(w \mid \eta) \begin{Bmatrix} < \\ > \end{Bmatrix} p^\infty(w \mid \eta) \qquad \text{for } w \begin{Bmatrix} \in \\ \notin \end{Bmatrix} \mathrm{Br}_s(\eta).$$

Let $\hat\beta$ denote the smallest β for which this is true for all η. Again, $\hat\beta$ goes to ∞ as ε goes to 0.

Application of Theorem A.3 and Corollary A.1 to the systems $\{r_s\}$ and $\{p_s^\beta\}$, where β equals $\bar\beta$ and $\hat\beta$, respectively, gives the following result. Let $\Pr_\beta$ denote the probability distribution on sample paths for log-linear choice with parameter β, and $\Pr_\varepsilon$ the corresponding probability distribution for perturbed best-response choice with parameter ε.

THEOREM A.4. *Let $D_s \equiv D \subset W$. If $\lim_{\beta\to\infty} \Pr_\beta(\eta_t(s) \in D) = 1$, then $\lim_{\varepsilon\to 0} \Pr_\varepsilon(\eta_t(s) \in D) = 1$. If $\lim_{\beta\to\infty} \Pr_\beta)\eta_t(s) \in D) = 0$, then $\lim_{\varepsilon\to 0} \Pr_\varepsilon(\eta_t(s) \in D) = 0$. In the first instance, take $\{p_s^1\} = \{p_s^\beta$ and $\{p_s^2\} = \{r_s\}$. In the second, take $\{p_s^1\} = \{r_s\}$ and $\{p_s^2\} = \{p_s^\beta\}$.*

Theorems 5.4 and 6.4 are a consequence of Theorem A.4.

REFERENCES

BLOCK, H. D., AND MARSCHAK, J. (1960). "Random orderings and stochastic theories of responses," in *Contributions to Probability and Statistics* (I. Olkin, S. Ghuyre, W. Hoeffding, W. Madow, and H. Mann; Ed.) Stanford CA: Stanford Univ. Press.

CANNING, D. (1990). "Social Equilibrium," unpublished manuscript.

CRAWFORD, V. (1991). "An "Evolutionary" Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games Econ. Behav.* **3,** 25–59.

ELLISON, G. (1992). "Learning, Local Interaction and Coordination," unpublished manuscript.

ETHIER, S., AND KURTZ, T. (1986). *Markov Processes*. New York: Wiley.

FÖLLMER, H. (1974). "Random Economies with Many Interacting Agents," *J. Math. Econ.* **1,** 51–62.

FUDENBERG, D., AND LEVINE, D. (1990). "Steady State Learning and Self-Confirming Equilibrium," unpublished manuscript.

GEORGII, H.-O. (1988). *Gibbs Measures and Phase Transitions*. New York: de Gruyter.

HARSANYI, J. C., AND SELTEN, R. (1988). *A General Theory of Equilibrium Selection in Games*. Cambridge, MA. MIT Press.

HOLLEY, R. (1985). "Rapid Convergence to Equilibrium in One-Dimensional Stochastic Ising Models," *Ann. Probab.* **13,** 72–89.

HOLLEY, R. (1987). "One Dimensional Stochastic Ising Models," in by *Percolation Theory and Ergodic Theory of Infinite Particle Systems* (H. Kesten, Ed.). New York: Springer-Verlag.

KANDORI, M., AND ROB, R. (1992). "Evolution of Equilibria in the Long Run: A General Theory and Applications," unpublished, University of Pennsylvania.

KANDORI, M., MAILATH, G., AND ROB, R. (1993). "Learning, Mutation, and Long-Run Equilibria in Games," *Econometrica* **61,** 29–56.

KÜNSCH, H. (1984). "Non-Reversible Stationary Measures for Infinite Interacting Particle Systems," *Z. Wahrsch. Verw Gebiete* **66,** 407–424.

LIGGETT, T. M. (1985). *Interacting Particle Systems*. Berlin: Springer-Verlag.

MONDERER, D. AND SHAPLEY, L. (1991). "Potential Games," unpublished, UCLA.

NACHBAR, J. (1990). "'Evolutionary' Selection Dynamics in Games: Convergence and Limit Properties," *Int. J. Game Theory* **19,** 59–89.

PRESTON, C. (1973). "Random Fields," in *Lecture Notes in Mathematics*, No. 534. Berlin: Springer-Verlag.

SCHELLING, T. (1971). "Dynamic Models of Segregation," *J. Math. Soc.* **1,** 143–186.

SCHELLING, T. (1973). "Hockey Helmets, Concealed Weapons, and Daylight Saving. A Study of Binary Choices with Externalities," *J. Conflict Res.* **17,** 381–428.

YOUNG, H. P. (1993). "The Evolution of Choice Behavior," *Econometrica* **61,** 57–84.