

Evolution in Games with Randomly Disturbed Payoffs*

Josef Hofbauer
Department of Mathematics
University College London
Gower Street
London WC1E 6BT UK
j.hofbauer@ucl.ac.uk
<http://homepage.univie.ac.at/Josef.Hofbauer>

William H. Sandholm
Department of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706 USA
whs@ssc.wisc.edu
<http://www.ssc.wisc.edu/~whs>

July 28, 2005

* The authors gratefully acknowledge financial support from the Austrian Science Fund (P 15281), the Centre for Economic Learning and Social Evolution at University College London, and the National Science Foundation (SES-0092145).

Abstract

We consider a simple model of stochastic evolution in population games. In our model, each agent occasionally receives opportunities to update his choice of strategy. When such an opportunity arises, the agent selects a strategy that is currently optimal, but only after his payoffs have been randomly perturbed. We prove that the resulting evolutionary process converges to approximate Nash equilibrium in both the medium run and the long run in three general classes of population games: stable games, potential games, and supermodular games. We conclude by contrasting the evolutionary process studied here with stochastic fictitious play. *Journal of Economic Literature*
Classification Numbers: C72, C73.

Keywords: evolutionary game theory, convergence to approximate Nash equilibrium, equilibrium selection.

Running Title: Randomly Disturbed Payoffs

1. Introduction

Nash equilibrium is the cornerstone of noncooperative game theory. Nevertheless, the traditional theoretical justifications for predicting Nash equilibrium play, which are based on assumptions about the players' rationality and equilibrium knowledge, are not always convincing, since in many applications these assumptions seem unreasonably demanding. Because the strength of the equilibrium knowledge assumption increases as the number of players grows, the traditional justification of equilibrium behavior seems particularly questionable when the number of players is large.

Fortunately, the existence of large numbers of players enables us to consider alternative approaches to justifying the Nash prediction, in particular if the game is played repeatedly over time. In this situation, it is natural to introduce a model in which players only occasionally consider revising their behavior, utilizing simple myopic decision rules when revision opportunities arise.¹ With such a model in hand, one can attempt to justify the Nash prediction by showing that the resulting evolutionary process leads to equilibrium play.

In this paper, we study evolution in population games: games played by large numbers of agents whose payoff functions are continuously differentiable in the proportions of agents choosing each strategy. While this class of games includes the standard model of random matching in normal form games as a special case, it also allows one to capture nonlinearities in payoffs that arise in many applications.

Our model of evolution is quite simple. Each player occasionally receives opportunities to revise his choice of strategy. When such an opportunity arises, the player chooses a best response to the current population state. However, this choice is made only after the player's payoffs are randomly perturbed, with these perturbations occurring independently at each revision opportunity. These payoff perturbations are analogous to those introduced by Harsanyi [19] in his model of purification of mixed equilibrium: in both his model and in ours, players have a unique best response after almost every realization of payoffs.

Our main goal in this paper is to determine conditions under which this evolutionary process generates approximate Nash equilibrium play. We consider two notions of convergence: *convergence in the medium run*, which concerns the behavior of the population over long but finite time spans, and *convergence in the long run*, which

¹ While in some contexts myopia is untenable hypothesis, here inertia in opponents' behavior and the anonymity of individual agents render this assumption quite reasonable.

concerns its behavior over the infinite time horizon. We establish that evolution leads to equilibrium behavior under both notions of convergence for three general classes of games: stable games (Sandholm [37]), potential games, and supermodular games. Our convergence results do not depend on the distributions of the payoff perturbations, and, unlike many convergence results in the evolutionary literature, require no restrictions on the number of strategies in the underlying game.

To begin our analysis, we associate with our stochastic evolutionary process an ordinary differential equation that describes the process' expected motion. This equation, the *perturbed best response dynamic*, is a smoothly perturbed version of the best response dynamic of Gilboa and Matsui [18]; its rest points are approximate Nash equilibria of the underlying game. Building on the work of Hofbauer [22], Hofbauer and Hopkins [24], and Hofbauer and Sandholm [25] for random matching settings, we establish stability properties for the perturbed best response dynamic in the three classes of population games noted above. We then establish convergence results for the original stochastic process by relying on a variety of approximation theorems: our medium run convergence theorems use results on the convergence of sequences of Markov processes (Kurtz [30]), while our long run convergence theorems utilize techniques from stochastic approximation theory (Benaïm [2], Benaïm and Hirsch [6]).

A number of authors have obtained convergence results for unperturbed best response dynamics in normal form games. In stochastic, finite player frameworks, Monderer and Shapley [32] prove convergence to Nash equilibrium in potential games, while Kandori and Rob [28] establish convergence to equilibrium in supermodular games. In the deterministic, continuum of player framework of Gilboa and Matsui [18], Hofbauer [21, 22] proves convergence to equilibrium in zero sum games, games with an interior ESS, and potential games.

There are a variety of reasons to focus instead on perturbed best response dynamics. For one, the unperturbed dynamics require an extreme sensitivity of players' choices to the exact value of the population state. This sensitivity manifests itself in the fact that Gilboa and Matsui's [18] dynamic defines not a continuous differential equation, but rather a discontinuous differential inclusion. In contrast, perturbed best responses change smoothly in the population state, and so generate well-behaved deterministic dynamics. Moreover, unlike its counterpart for the unperturbed dynamic, the stochastic process underlying the perturbed best response dynamic is ergodic, with long run behavior described by a unique stationary distribution. Ergodicity simplifies our long run analysis, and also introduces the possibility of establishing strong equilibrium selection results, in the spirit of those proved by Foster and Young [16],

Young [43], Kandori, Mailath, and Rob [27], Kandori and Rob [28], Blume [9, 10], and especially Benaïm and Weibull [7].

In an earlier paper (Hofbauer and Sandholm [25]), we obtained convergence results for the learning process known as stochastic fictitious play (Fudenberg and Kreps [17]). In stochastic fictitious play, a group of n players repeatedly play an n player normal form game. During each discrete time period, each player plays a best response to the time average of his opponents' play, but only after his payoffs have been struck by random perturbations. Like those of the evolutionary process studied here, the limiting properties of stochastic fictitious play can be characterized in terms of the perturbed best response dynamic. But there are other respects in which the two processes are fundamentally different: the two processes are specified in terms of distinct types of state variable, and different limiting operations are employed in order to obtain convergence results. Furthermore, while our work on stochastic fictitious play concerned learning in normal form games, the present paper establishes convergence results in the more general context of population games. *Inter alia*, this broader framework enables us to establish global convergence to a unique equilibrium in all stable games, a class of games containing many examples of economic interest that fall outside the random matching framework. We discuss all of these issues in considerable detail in the final section of the paper.

Section 2 introduces our strategic framework and our model of stochastic evolution. Section 3 analyzes the perturbed best response dynamics in stable games, potential games, and supermodular games. Section 4 contains our results on convergence in the medium run and convergence in the long run. Section 5 concludes by contrasting stochastic evolution with stochastic fictitious play. All proofs are relegated to the Appendix.

2. The Model

2.1 Population Games

We begin by defining population games with continuous player sets. Let $\mathcal{P} = \{1, \dots, \bar{p}\}$ be a set of \bar{p} populations, where $\bar{p} \geq 1$. Population p is of mass m^p , and the total mass of all populations is $m = \sum_{p \in \mathcal{P}} m^p$; for convenience, we assume that each m^p is an integer.

Members of population p choose strategies from the set $S^p = \{1, \dots, n^p\}$, so the total number of pure strategies in all populations is $n = \sum_{p \in \mathcal{P}} n^p$. We let $\Delta^p = \{x^p \in \mathbf{R}_+^{n^p} :$

$\sum_{i \in S^p} x_i^p = 1$ denote the set of probability distributions over strategies in S^p . The set of strategy distributions for population p is denoted $X^p = m^p \Delta^p = \{x^p \in \mathbf{R}_+^{n^p} : \sum_{i \in S^p} x_i^p = m^p\}$, while $X = \{x = (x^1, \dots, x^{\bar{p}}) \in \mathbf{R}_+^n : x^p \in X^p\}$ is the set of overall strategy distributions. While the population's aggregate behavior is always described by a point in X , it is useful to define payoffs on the set $\bar{X} = \{x \in \mathbf{R}_+^n : m^p - \varepsilon \leq \sum_i x_i^p \leq m^p + \varepsilon \ \forall p \in \mathcal{P}\}$, where ε is a positive constant. This set contains the strategy distributions that arise if the populations' masses vary slightly. Defining payoffs on this set is useful because it enables us to speak directly about a player's marginal impact on his opponents' payoffs, but is not essential to our analysis.

The payoff function for strategy $i \in S^p$ is denoted $F_i^p : \bar{X} \rightarrow \mathbf{R}$, and is assumed to be continuously differentiable. Note that the payoffs to a strategy in population p can depend on the strategy distribution within population p itself. We let $F^p : \bar{X} \rightarrow \mathbf{R}^{n^p}$ refer to the vector of payoff functions for strategies belonging to population p , and we identify a *population game* with its payoff vector field $F : \bar{X} \rightarrow \mathbf{R}^n$.

We now introduce some examples of population games that we will revisit throughout the paper.

Random matching in normal form games. Suppose that a single unit mass population of players is randomly matched to play a symmetric normal form game with payoff matrix $A \in \mathbf{R}^{n \times n}$, where A_{ij} is the payoff a player obtains if he plays i and his opponent plays j . Then the payoffs for the corresponding population game are $F(x) = Ax$.

Alternatively, suppose that members of two unit mass populations are paired to play a normal form game with bimatrix $(A, B) \in \mathbf{R}^{n^1 \times n^2} \times \mathbf{R}^{n^1 \times n^2}$. If two matched players play strategies $i \in S^1$ and $j \in S^2$, they obtain payoffs of A_{ij} and B_{ij} , respectively. The corresponding population game has payoffs

$$F(x^1, x^2) = \begin{pmatrix} 0 & A \\ B' & 0 \end{pmatrix} \begin{pmatrix} x^1 \\ x^2 \end{pmatrix}.$$

Because of the linearity of the expectation operator, random matching yields population games with linear or multilinear payoffs, and in which a player's payoffs do not depend on the behavior of other members of his population (when $\bar{p} \geq 2$). Population games that are not based on random matching need not possess either of these properties. Our next class of examples provides a case in point.

Congestion games: Congestion games are a natural tool for modeling externalities, such as those arising in traffic networks (see Sandholm [35, 38]). In a congestion game,

each strategy $i \in S^p$ requires the use of some finite collection of facilities $\Phi_i^p \subseteq \Phi$. In traffic network models, each facility represents a link in the network, and each strategy corresponds to a path (i.e., a collection of links) connecting an origin/destination pair. In general, the set of facilities is simply an arbitrary finite set; in particular, there is no need to assume that a network structure on Φ exists.

Each facility ϕ has a cost function $c_\phi: \mathbf{R}_+ \rightarrow \mathbf{R}$ that describes the penalty (delay) from using the facility. The cost of facility ϕ is a function of its *utilization* u_ϕ , the total mass of the players who use the facility:

$$u_\phi(x) = \sum_{p \in \mathcal{P}} \sum_{i \in \rho^p(\phi)} x_i^p, \text{ where } \rho^p(\phi) = \{i \in S^p: \phi \in \Phi_i^p\}.$$

The congestion game is defined by the payoff functions

$$F_i^p(x) = - \sum_{\phi \in \Phi_i^p} c_\phi(u_\phi(x)).$$

In settings like traffic networks involving negative externalities, the cost functions c_ϕ are increasing; positive externalities lead to decreasing cost functions. Payoffs in congestion games depend on own-population behavior, and need only be linear if the underlying cost functions are linear themselves.

2.2 Evolution with Randomly Disturbed Payoffs

We now introduce our model of evolution with randomly disturbed payoffs. Models of this sort were first considered by Blume [9, 10] and Young [44] in a random matching setting under a specific parametric assumption on the disturbance distributions. Here we consider evolution in general population games, and place virtually no restrictions on the form that payoff disturbances take.²

Members of \bar{p} finite populations of sizes $(Nm^1, \dots, Nm^{\bar{p}})$ recurrently play the population game F . Players occasionally receive opportunities to switch strategies, with each player's opportunities arriving via independent, rate 1 Poisson processes. When a player from population p receives a revision opportunity, he evaluates the current expected payoff to each of his pure strategies, but his assessments are subject to random shocks that follow a given probability distribution ν^p on \mathbf{R}^{n^p} . The player selects the strategy that he evaluates as best.

² Blume and Young restrict attention to evolution under the logit choice rule, which we describe below. These authors also analyze models of local interaction, which we do not consider here.

Although payoff and choice shocks drawn at random in each period are now common features of evolutionary models, it is worthwhile to provide a direct justification for their use. Following Harsanyi [19], we can understand the payoff shocks as representing small, random influences on behavior; in this case, we consider distributions ν^p that place nearly all of their mass in a neighborhood of the origin. Large payoff shocks may be a more natural assumption, for example, in cases where preferences for variety are at least as strong as the preferences described by the payoffs of the underlying game. Of course, one can also take a middle course, under which payoff shocks are typically quite small but occasionally rather large, so that the more significant shocks only occur infrequently and irregularly.

Aggregate behavior in this model is described by a continuous time Markov chain $\{X_t^N\}_{t \geq 0}$, which takes values in the state space $\mathcal{X}^N = \{x \in X: Nx \in \mathbf{Z}^n\}$. The initial condition X_0^N is arbitrary. Let τ_k denote the random time at which the k th revision opportunity arises. For a switch from strategy $i \in S^p$ to strategy $j \in S^p$ to occur during this opportunity, the player granted the revision opportunity must be a member of population p who is playing strategy i , and the realization of his payoff disturbance must render strategy j his best response. Transitions of X_t^N are therefore described by

$$P\left(X_{\tau_{r+1}}^N = x + \frac{1}{N}(e_j^p - e_i^p) \mid X_{\tau_r}^N = x\right) = \frac{1}{m} x_i^p \nu^p(\varepsilon^p : \arg \max_{k \in S^p} F_k^p(x) + \varepsilon_k^p = j)$$

for $i \neq j$, where e_i^p and e_j^p are standard basis vectors. With the remaining probability of $\sum_{p \in \mathcal{P}} \sum_{i \in S^p} \frac{1}{m} x_i^p \nu^p(\varepsilon^p : \arg \max_{k \in S^p} F_k^p(x) + \varepsilon_k^p = i)$, no change in the state occurs.

To analyze this process, we introduce the notion of a perturbed best response function. To begin, define the *choice probability function* $C^p: \mathbf{R}^{n^p} \rightarrow \Delta^p$ by

$$(1) \quad C_i^p(\pi^p) = \nu^p\left(\varepsilon^p : i \in \arg \max_{j \in S^p} \pi_j^p + \varepsilon_j^p\right).$$

If a player currently faces a base payoff vector of π^p , then $C_i^p(\pi^p)$ represents the probability that the realized payoff perturbation leads him to choose strategy i . When ν^p places most of its mass near the origin, then $C^p(\pi^p)$ puts most of its mass on the maximizer of π^p , but places positive mass on all elements of S^p .

Our regularity condition on perturbation distributions is defined in terms of the function C^p . We call ν^p an *admissible distribution* if it admits a strictly positive density

on \mathbf{R}^{n^p} and is smooth enough that C^p is continuously differentiable. The profile $\nu = (\nu^1, \dots, \nu^{\bar{p}})$ is *admissible* if each of its components is admissible.³

Now let F be a population game, and let ν be a profile of admissible distributions. We define the *perturbed best response function* $\tilde{B}^p: X \rightarrow \Delta^p$ for the pair (F, ν) by the composition $\tilde{B}^p = C^p \circ F^p$.

With this definition in hand, we can express the transition rule above as follows:

$$\begin{aligned} P\left(X_{\tau_{r+1}}^N = x + \frac{1}{N}(e_j^p - e_i^p) \mid X_{\tau_r}^N = x\right) &= \frac{1}{m} x_i^p \nu^p(\varepsilon^p : \operatorname{argmax}_{k \in S^p} F_k^p(x) + \varepsilon_k^p = j) \\ &= \frac{1}{m} x_i^p C_j^p(F^p(x)) \\ &= \frac{1}{m} x_i^p \tilde{B}_j^p(x) \end{aligned}$$

The expected increment in X_t^N during a single revision opportunity is therefore described by

$$\begin{aligned} E\left(X_{\tau_{r+1}}^{N,p} - X_{\tau_r}^{N,p} \mid X_{\tau_r}^N = x\right) &= \sum_{i \in S^p} \sum_{j \in S^p} \frac{1}{N} (e_j^p - e_i^p) \frac{1}{m} x_i^p \tilde{B}_j^p(x) \\ &= \frac{1}{Nm} \left(\sum_{j \in S^p} e_j^p \tilde{B}_j^p(x) \sum_{i \in S^p} x_i^p - \sum_{i \in S^p} e_i^p x_i^p \sum_{j \in S^p} \tilde{B}_j^p(x) \right) \\ &= \frac{1}{Nm} \left(m^p \tilde{B}^p(x) - x^p \right) \end{aligned}$$

Since each of the Nm players' revision opportunities arrive according to independent rate 1 Poisson processes, the revision opportunities arriving in the society as a whole are described by the sum of these processes, which is a Poisson process with rate Nm . We therefore multiply the expression above by Nm to obtain the expected increment in X_t^N per unit of time. Writing the result as a differential equation, we obtain

$$(P) \quad \dot{x}^p = m^p \tilde{B}^p(x) - x^p \text{ for all } p \in \mathcal{P}.$$

We call this equation the *perturbed best response dynamic* for the pair (F, ν) .

We call $x \in X$ a *perturbed equilibrium* for (F, ν) if it is a fixed point of $(m^1 \tilde{B}^1(x), \dots, m^{\bar{p}} \tilde{B}^{\bar{p}}(x))$, or, equivalently, if it is a rest point of (P). We let $PE(F, \nu)$ denote the set of

³ The assumption that ν^p has full support on \mathbf{R}^n is stronger than necessary. Once we fix a game F , we can compute a finite bound M_F^p on the difference between the payoffs generated by any pair of strategies in S^p at any state in X . Using this bound, we can construct a smooth distribution $\hat{\nu}^p$ that generates the same choice probabilities as ν^p at all payoff vectors feasible under F but whose support is contained in a compact set (namely, a cube with sides of length $2n^p M_F^p$).

perturbed equilibria. One can show that if most of the mass in each distribution v^p is near the origin, then the perturbed equilibria of (F, v) approximate Nash equilibria of F .⁴

Our aim in this paper is to relate the behavior of the stochastic process X_t^N to solutions of the deterministic dynamic (P). Our analysis proceeds in three steps. In the following section, we investigate the behavior of the dynamic (P) in three classes of games. In Section 4.1, we combine the analysis of equation (P) with results on convergence of Markov processes to obtain finite horizon convergence results. In Section 4.2, the deterministic analysis and tools from stochastic approximation theory are employed to establish infinite horizon convergence results.

3. Analysis of the Perturbed Best Response Dynamic

We now introduce three classes of population games for which the behavior of perturbed best response dynamics can be well characterized: stable games, potential games, and supermodular games. These characterizations generalize results established by Hofbauer [22], Hofbauer and Hopkins [24], and Hofbauer and Sandholm [25] for random matching games to general population games. Our results for stable games substantially expand the set of games for which the dynamics are known to have a globally attracting state.

Our results for stable games and potential games rely on a discrete choice theorem from Hofbauer and Sandholm [25]. Recall that the choice probability function C^p from equation (1) is defined in terms of admissible stochastic payoff perturbations of the payoffs to each pure strategy. Theorem 2.1 of Hofbauer and Sandholm [25] shows that there is always an alternative representation of C^p that relies on a *deterministic* perturbation of the payoffs to each *mixed* strategy.

More specifically, we call the function $V^p: \text{int}(\Delta^p) \rightarrow \mathbf{R}$ an *admissible deterministic perturbation* if it is differentiable strictly convex and becomes infinitely steep near the boundary of Δ^p . Then if the function C^p is defined via equation (1) for some admissible distribution v^p , there is an admissible deterministic perturbation V^p such that

$$(2) \quad C^p(\pi^p) = \underset{y^p \in \text{int}(\Delta^p)}{\text{argmax}} (y^p \cdot \pi^p - V^p(y^p)).$$

On the other hand, the converse statement is false: there are choice functions defined by admissible deterministic perturbations that admit no stochastic representation.

⁴ See Proposition 3.1 of Hofbauer and Sandholm [25].

One can interpret the function V^p as a "control cost" that is larger for "purer" elements of Δ^p (van Damme [41, Ch. 4]). Thus, the representation theorem shows that the choice probability functions obtained from additive random utility models can always be represented using a framework in which mixed strategies are chosen directly, but in which this choice is subject to convex control costs. Further details on this result needed for our analysis are provided in the Appendix.

The best known example of a choice probability function is the *logit choice function*,

$$C_i(\pi) = \frac{\exp(\eta^{-1}\pi_i)}{\sum_j \exp(\eta^{-1}\pi_j)}.$$

By varying the *noise level* η from zero to infinity, one obtains behavior that varies from pure optimization to uniform randomization. It is well known that logit choice can be derived in terms of both stochastic and deterministic perturbations: equation (1) yields logit choice if the stochastic perturbations are i.i.d. with the extreme value distribution $\exp(-\exp(-\eta^{-1}x - \gamma))$ (where $\gamma \approx 0.5772$ is Euler's constant), while equation (2) yields logit choice if V^p is the entropy function $V^p(y^p) = \eta \sum_j y_j^p \ln y_j^p$. The theorem described above shows that such a dual representation is possible regardless of the joint distribution of the stochastic perturbations.

3.1 Stable Games

Let $TX = \{z \in \mathbf{R}^n : \sum_{i \in S^p} z_i^p = 0 \text{ for all } p \in \mathcal{P}\}$ be the set of directions tangent to the set of population states X , and for any function $f: X \rightarrow \mathbf{R}$ and direction $z \in TX$, let

$$\frac{\partial f}{\partial z}(x) = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon z) - f(x)}{\varepsilon}$$

denote the derivative of f at point x in direction z . Following Sandholm [37], we say that F is a *stable game* if it satisfies

$$(SE) \quad \sum_p \sum_i z_i^p \frac{\partial F_i^p}{\partial z}(x) \leq 0 \text{ for all } z \in TX \text{ and all } x \in X.$$

Equivalently, F is stable if it satisfies the negative semidefiniteness condition

$$z \cdot DF(x) z \leq 0 \text{ for all } x \in X \text{ and all } z \in TX.$$

Condition (SE) is called *self-defeating externalities*. It requires that if a small group of players switches strategies, then the improvements in payoffs of the strategies they switch to are exceeded by the improvement in payoffs of the strategies they abandon.

When the population game F is defined via random matching, condition (SE) is quite restrictive. For instance, in the two population random matching framework, it is easy to show that F is stable if and only if the underlying normal form game (A, B) is equivalent to a zero sum game. However, if payoffs can depend on own-population behavior, then condition (SE) is far less limiting. Indeed, congestion games with increasing facility costs (e.g., traffic network games) are all stable games, as are concave potential games, RL stable games (Cressman, Garay, and Hofbauer [14]), and negative diagonal dominant games. For a presentation of all of these examples and further discussion of condition (SE), see Sandholm [37].

If F is a stable game, then the set of all Nash equilibria of F is convex (see Hofbauer [23]); under a mild additional assumption, the Nash equilibrium of F is unique (see Sandholm [37]). In Theorem 3.1 below, we establish that all perturbed best response dynamics for stable games admit a single globally asymptotically stable rest point.

The construction we use to prove this result generalizes one introduced by Hofbauer [22] in a single population random matching setting. Consider the function $\Lambda: X \rightarrow \mathbf{R}_+$ defined by

$$\Lambda(x) = \sum_{p \in \mathcal{P}} m^p \left[\max_{y^p \in \text{int}(\Delta^p)} (y^p \cdot F^p(x) - V^p(y^p)) - \left(\frac{1}{m^p} x^p \cdot F^p(x) - V^p\left(\frac{1}{m^p} x^p\right) \right) \right],$$

where V^p is the deterministic perturbation associated with the distribution v^p .

Theorem 3.1: *Suppose F is a stable game and that v is admissible. Then*

(i) *The function Λ is a strict Lyapunov function for the dynamic (P): its value decreases strictly along every non-constant solution trajectory.*

(ii) *(F, v) admits a unique and globally asymptotically stable perturbed equilibrium, which is the lone state at which $\Lambda(x) = 0$.*

Theorem 3.1 shows that if F is a stable game and v a profile of admissible disturbance distributions, then the set of perturbed equilibria $PE(F, v)$ consists of a single state that is globally asymptotically stable under (P). To establish this, we first show that the positive function Λ is a strict Lyapunov function for (P), and that the zeros of Λ are the rest points of (P). We then use the stability of F and the strict convexity of V^p to prove that (P) admits exactly one rest point. Together, these assertions imply that there is a unique, globally asymptotically stable perturbed equilibrium of (F, v) .

To understand the Lyapunov function Λ , recall that the payoff vector for population p at population state x is $F^p(x)$. Fix this payoff vector, and suppose that the members of population p jointly choose a mixed strategy y^p in an attempt to maximize the difference between the aggregate payoff $y^p \cdot F^p(x)$ and the control cost $V^p(y^p)$. The bracketed expression in the definition of Λ is the gap between this maximized difference and the current difference, interpreting $\frac{1}{m^p} x^p \in \Delta^p$ as the population's current mixed strategy. Theorem 3.1 shows that the weighted sum of these gaps over all populations decreases under the dynamic (P). This sum is zero precisely when all populations maximize the difference between aggregate payoffs and control costs; the lone state where this occurs is the unique perturbed equilibrium of (F, v) .

3.2 Potential Games

We call the game F a *potential game* if it satisfies

$$(ES) \quad \frac{\partial F_i^p}{\partial x_j^q}(x) = \frac{\partial F_j^q}{\partial x_i^p}(x) \text{ for all } i \in S^p, j \in S^q, p, q \in \mathcal{P}, \text{ and } x \in X.$$

This requirement is stated more concisely as

$$DF(x) \text{ is symmetric for all } x \in X.$$

Condition (ES) is called *externality symmetry*. It requires that the effect on the payoffs to strategy $j \in S^q$ of introducing new players choosing strategy $i \in S^p$ is always equals the effect on the payoffs to strategy i of introducing new players choosing strategy j . Random matching games in which all players in a match receive the same payoff are potential games. More interesting examples arise in nonlinear settings: all congestion games are potential games, as are games generated by certain marginal externality pricing schemes. For further details on these examples, see Sandholm [35, 38].

Since the derivative of F is symmetric, every potential game F admits a *potential function* $f: \bar{X} \rightarrow \mathbf{R}$: that is, a function that satisfies $\nabla f(x) = F(x)$ for all $x \in X$. Hofbauer [22] and Sandholm [35] show that this potential function serves as a Lyapunov function for a wide range of unperturbed evolutionary dynamics, and so can be used to establish global convergence results. To obtain a Lyapunov function for the perturbed best response dynamics, one need only perturb the potential function by the deterministic perturbations V^p . Define

$$\Pi(x) = f(x) - \sum_{p \in \mathcal{P}} m^p V^p\left(\frac{1}{m^p} x^p\right).$$

Theorem 3.2: *If F is a potential game and v is admissible, then*

- (i) *Π is an (increasing) strict Lyapunov function for the dynamic (P).*
- (ii) *All solution trajectories of (P) converge to connected subsets of $PE(F, v)$, and $PE(F, v) = \{x \in X: x \text{ is a critical point of } \Pi\}$. If $PE(F, v)$ is a singleton it is globally asymptotically stable.*

3.3 Supermodular Games.

We say that F is a *supermodular game* if it satisfies

$$(SC) \quad \frac{\partial(F_{i+1}^p - F_i^p)}{\partial(e_{j+1}^q - e_j^q)}(x) \geq 0 \text{ for all } i < n^p, j < n^q, p, q \in P, \text{ and } x \in X.$$

When expanded, the leading inequality in this condition becomes

$$\frac{\partial F_{i+1}^p}{\partial x_{j+1}^q}(x) - \frac{\partial F_{i+1}^p}{\partial x_j^q}(x) \geq \frac{\partial F_i^p}{\partial x_{j+1}^q}(x) - \frac{\partial F_i^p}{\partial x_j^q}(x).$$

We call condition (SC) *strategic complementarity*. It states that if some players in population q switch from strategy j to strategy $j + 1$, the performance of strategy $i + 1 \in S^p$ improves relative to that of strategy i . This condition is an infinite player generalization of conditions for finite player games studied by Topkis [40], Vives [42], and Milgrom and Roberts [31]. These papers provide many microeconomic applications of supermodular games, while Cooper [13] offers a number of macroeconomic applications.

It is easiest to study perturbed best response dynamics for supermodular games after applying a change of coordinates. Define the linear operator $T^p: X^p \rightarrow \mathbf{R}^{n^p-1}$ by

$$(T^p x^p)_i = \sum_{j=i+1}^{n^p} x_j^p.$$

If we view x^p as a discrete density function on the set of pure strategies $S^p = \{1, \dots, n^p\}$ with total mass m^p , then $T^p x^p$ is the corresponding decumulative distribution function. Hence, z^p stochastically dominates x^p if and only if $T^p z^p \geq T^p x^p$. To compare complete population states, we let $Tx = (T^1 x^1, \dots, T^{\bar{p}} x^{\bar{p}})$.

Our goal is to show that when F is supermodular, the dynamic (P) is *strongly monotone with respect to the stochastic dominance order*: if $\{x_t\}_{t \geq 0}$ and $\{z_t\}_{t \geq 0}$ are two solutions to (P) with $Tz_0 \geq Tx_0$ and $z_0 \neq x_0$, then $Tz_t > Tx_t$ for all $t > 0$. Doing so is valuable because as we shall see, strongly monotone dynamics have appealing convergence properties.

To establish strong monotonicity, we require a mild additional assumption on the game F . Let $\hat{S} = \{(k, p): k \in S^p - \{n^p\}, p \in P\}$. We say that the supermodular game F is *irreducible* if for all states $x \in X$ and all nonempty proper subsets K of \hat{S} , there exist a pair $(k, p) \in K$, a strategy $i \in S^p - \{n^p\}$, and a pair $(j, q) \in \hat{S} - K$ such that condition (SC) holds strictly at x for the pairs (i, p) and (j, q) . Under this condition, a movement of

mass from strategy $j \in S^q$ to strategy $j + 1$ strictly improves the relative performance of some strategy belonging to the same population as strategy k .⁵

Theorem 3.3 shows that if F is supermodular and irreducible, then almost all solution trajectories of perturbed best response dynamics converge to perturbed equilibria.

Theorem 3.3: *If F is an irreducible supermodular game and ν is admissible, then the dynamic (P) is strongly monotone with respect to the stochastic dominance order. Hence, there is an open, dense, full measure set of initial conditions from which solutions to (P) converge to single points in $PE(F, \nu)$. If $PE(F, \nu)$ is a singleton it is globally asymptotically stable.*

The constructions used in Theorems 3.1 and 3.2 show that these results hold not only for perturbed best response dynamics based on stochastic payoff perturbations, but also for dynamics based on deterministic perturbations. In contrast, Theorem 3.3 *cannot* be extended to all perturbed best response dynamics based on deterministic perturbations, as the extra structure provided by the stochastic perturbations is needed to establish the monotonicity of the dynamics.

4. Convergence Theorems

In this section, we use the preceding analysis to prove two sets of convergence results for the Markov processes X_t^N . The first set, described in Theorem 4.1, shows that over finite time horizons, in the three classes of games studied above, the process X_t^N converges to the set of perturbed equilibria. The second set, stated in Theorem 4.2, demonstrates that over the infinite time horizon, X_t^N converges to the set of Lyapunov stable equilibria. While the medium analysis is simpler, the long run results evidently offer a more refined prediction of play. However, we shall see that the notions of convergence used in each case differ in subtle but important ways, lending each set of results its own unique appeal.

⁵ Irreducibility is a weaker assumption than strict supermodularity, the assumption utilized in Hofbauer and Sandholm [25] in the context of normal form games.

4.1 Convergence in the Medium Run

To state our finite horizon convergence result, we consider a sequence of Markov processes X_t^N whose initial conditions $X_0^N \in \mathcal{X}^N$ converge to some state $x_0 \in X$ as the population size N approaches infinity. We say that these processes *converge in the medium run* to the closed set $A \subseteq X$ from the initial condition x_0 if for each open set O containing A , there is a time $T_0 = T_0(x_0)$ such that for all $T \geq T_0$,

$$\lim_{N \rightarrow \infty} P(X_t^N \in O \text{ for all } t \in [T_0, T]) = 1.$$

In words: if a large group of players begins play near x_0 , then with probability close to 1, their behavior approaches the set A and remains nearby for a long, finite time span. We say that convergence is *uniform* if the time T_0 that the neighborhood of A is reached can be chosen independently of the initial condition $x_0 \in X$.

Theorem 4.1. *Consider stochastic evolution in the game F under the admissible disturbance distributions ν^p .*

(i) *If F is a stable game, then X_t^N converges in the medium run to the singleton $PE(F, \nu)$ from every initial condition $x_0 \in X$.*

(ii) *If F is a potential game, then X_t^N converges in the medium run to a connected subset of $PE(F, \nu)$ from every initial condition $x_0 \in X$.*

(iii) *If F is an irreducible supermodular game, then X_t^N converges in the medium run to an element of $PE(F, \nu)$ from an open, dense, full measure set of initial conditions $x_0 \in X$.*

In all cases, convergence is uniform whenever $PE(F, \nu)$ is a singleton.

The proof of the theorem is based on an analogue of the law of large numbers for sequences of Markov chains that has been studied in game theoretic contexts by Binmore and Samuelson [8], Sandholm [36], and Benaïm and Weibull [7]. It is presented in the Appendix.

4.2 Convergence in the Long Run

Theorem 4.1 cannot be extended to an infinite horizon result ($T = \infty$): since the process X_t^N is irreducible, all states in \mathcal{X}^N are visited infinitely often with probability one, and so large deviations from all rest points are certain to occur. But it is precisely this fact that enables us to obtain tighter predictions of behavior over this time span.

While all states are visited and abandoned infinitely often, one expects that only states near attractors of (P) will be visited with nonvanishing frequency. This observation is the basis for our infinite horizon convergence results.

We formally characterize infinite horizon behavior using the stationary distribution μ^N of the process X_t^N . Since X_t^N is irreducible and aperiodic, the stationary distribution is unique, and it describes the long run behavior of X_t^N in two distinct ways.⁶ Regardless of initial behavior, μ^N approximates the probability distribution of X_t^N after a long enough time has passed:

$$\lim_{t \rightarrow \infty} P(X_t^N \in A | X_0^N = x_0^N) = \mu^N(A) \text{ for all } x_0^N \in X^N.$$

More importantly, μ^N also describes the limiting time average of play:

$$P\left(\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T 1_{\{X_t^N \in A\}} dt = \mu^N(A) | X_0^N = x_0^N\right) = 1.$$

Our notion of infinite horizon convergence is defined in terms of the stationary distributions μ^N . We say that the processes X_t^N *converge in the long run* to the closed set $A \subseteq X$ if for each open set O containing A , we have that

$$\lim_{N \rightarrow \infty} \mu^N(O) = 1.$$

Our two notions of convergence differ not only in terms the time horizons under consideration, but also in terms of the fixedness of behavior at the predicted set A . Under medium horizon convergence, after the time T_0 at which a neighborhood of A is reached, the process X_t^N may not leave this neighborhood for a long, finite span. This form of convergence is appealing because of its stringency. However, the time scale on which this notion of convergence is useful is one that does not allow us to discard unstable rest points of (P).

By considering infinite horizon behavior, we are able to use the randomness of the process X_t^N to rule out unstable rest points. But the time scale that permits unstable rest points to be abandoned is also one on which convergence to stable rest points is temporary. This relative weakness is embodied in our convergence criterion. By defining our notion of long run convergence in terms of the stationary distributions μ^N , we concern ourselves with the time average of play. In doing so, we allow for

⁶ For more on these properties, see, e.g., Durrett [15].

departures from the predicted set A , so long as these departures are sufficiently uncommon.

To state Theorem 4.2, we recall that the rest point $x^* \in PE(F, \nu)$ is *Lyapunov stable* if for each neighborhood O of x^* , there is another neighborhood Q of x^* such that solutions to (P) that begin in Q remain in O for all positive times. Let $LS(F, \nu) \subseteq PE(F, \nu)$ denote the set of Lyapunov stable rest points of (P).

Theorem 4.2: *Consider stochastic evolution in the game F under the admissible disturbance distributions ν^p .*

(i) *Suppose that F is a stable game. Then X_t^N converges in the long run to the singleton $PE(F, \nu) = LS(F, \nu)$.*

(ii) *Suppose that F is a potential game and that $PE(F, \nu)$ is finite. Then X_t^N converges in the long run to $LS(F, \nu)$.*

(iii) *Suppose that F is an irreducible supermodular game. Then X_t^N converges in the long run to $LS(F, \nu)$.*

Part (i) of the theorem shows that if F is a stable game, then in the long run a large population is nearly always in a neighborhood of the unique perturbed equilibrium of (F, ν) . Part (ii) shows that if F is a potential game, then under a mild regularity condition, the population only stays near Lyapunov stable rest points of (P). Part (iii) shows that this conclusion also holds if F is supermodular and irreducible.⁷ The proof of the theorem, which combines our earlier analysis with stochastic approximation results due to Benaïm [2] and Benaïm and Hirsch [6], is provided in the Appendix.⁸

⁷ Benaïm and Hirsch [6] establish this last result for the case of normal form supermodular games with exactly two strategies per player.

⁸ It is worth noting that our convergence results, in particular our results for supermodular games, impose no restrictions on the number of strategies in the underlying game. In Hofbauer and Sandholm [25], our convergence theorem for stochastic fictitious play in normal form supermodular games requires the dimension of the state space to be no greater than 2. However, if a conjecture of Benaïm [4] is correct this dimensionality condition is actually not needed to establish convergence.

5. Contrasting Stochastic Evolution and Stochastic Fictitious Play

We conclude the paper by contrasting the stochastic evolutionary process studied here with stochastic fictitious play. In standard fictitious play (Brown [11]), a group of players repeatedly plays a normal form game. In every period, each player chooses a best response to his beliefs, which are given by the time average of past play. In stochastic fictitious play (Fudenberg and Kreps [17]), best responses are chosen after each player's payoffs are randomly perturbed. Like the process considered above, the expected motion of stochastic fictitious play is described by the perturbed best response dynamic (P). Using this observation, Fudenberg and Kreps [17], Kaniovski and Young [29] and Benaïm and Hirsch [5] prove that stochastic fictitious play converges in 2×2 games, while Hofbauer and Sandholm [25] establish convergence in games with an interior ESS, zero-sum games, potential games, and certain supermodular games.

While stochastic fictitious play is model of behavior in normal form games, stochastic evolution can be used to model behavior in any population game, allowing us to establish convergence results in a broader class of strategic settings. For example, in settings with two player roles, stochastic fictitious play converges to an interior equilibrium only in games that are essentially zero-sum (Hofbauer and Hopkins [24]). Theorems 4.1 and 4.2 show that in the evolutionary model, such convergence occurs in all stable games whose Nash equilibria are not on the boundary of the state space; these include, for example, games used to model highway congestion.

The most important differences between the two models lie in the definitions of their state variables and in the limits taken in establishing convergence results. The state variable of stochastic fictitious play is the time average of past play, so the increment in the state at time t is of size $\frac{1}{t}$. Because these increments become vanishingly small, one can obtain convergence results by simply studying the limit behavior of the state variable as t grows large. In contrast, the state variable under stochastic evolution describes the proportions of players choosing each strategy, so the increments of the state are of fixed size $\frac{1}{N}$. Since the state space of the process is a finite grid for each fixed value of N , proving convergence to perturbed equilibrium requires us to consider limits as the population size grows large. Because limits are taken in N rather than in t , it is possible to prove separate limit results for finite and infinite time horizons.

These distinguishing features also underlie a more subtle difference between the two processes. Suppose that the dynamic (P) has the phase diagram in Figure 1, flowing

clockwise around a circle except at a single rest point.⁹ Then the expected motions of both processes are described by this picture, although in each case actual motions are random.




Figure 1: A deterministic flow.

Under stochastic fictitious play, the state variable is the time average of past play. Here it proceeds clockwise on average, but moves quite slowly near the top of the circle. When the rest point is reached, the expected change in the state is zero, but since the actual increments are stochastic, the process eventually clears the rest point and begins another circuit. Consequently, while time averaged behavior under stochastic fictitious play can in principle converge to a single limit point, in this case the set of limit points is the entire state space.

Under stochastic evolution, the state represents the current proportions of players choosing each strategy, and this too perpetually rounds the circle. Because the evolutionary process is ergodic, convergence to a single limit point is impossible even in principle. Therefore, when studying long run behavior, we examine the stationary distribution of the process, which describes its limiting time average. Since the expected motion of the process becomes vanishingly slow only in a neighborhood of the rest point, in the long run the time average of play is concentrated entirely on this segment.

This difference in the strength of the convergence results is due to a reversal in the order of two operations: time averaging and deterministic approximation. Under stochastic fictitious play, the state variable is defined as the time average of play, and

⁹ Of course, the state space of (P) cannot be a circle, but pretending this is possible simplifies our discussion.

the dynamics of the time average are studied using a deterministic approximation. Under stochastic evolution, the state variable represents current behavior, the evolution of which is analyzed through a deterministic approximation, and only after this is a time average taken to describe long run play.

This distinction is reflected in the different notions of recurrence applied to the dynamic (P) when analyzing the two models. Benaïm and Hirsch [5] show that the limiting time average of stochastic fictitious play lies in the *chain recurrent set* of (P), a set containing those states that can occur repeatedly if the flow of (P) is subjected to small shocks at isolated points in time. In contrast, Benaïm and Weibull [7] show that the limiting stationary distribution under stochastic evolution is concentrated on the *minimal center of attraction* of (P).¹⁰ The chain recurrent set always contains the minimal center of attraction, and the example above shows that this inclusion can be strict. Thus, the basic prediction generated by the stochastic evolution model is finer than that derived from stochastic fictitious play.

Appendix

We begin by reviewing the discrete choice characterization theorem from Hofbauer and Sandholm [25]. Define the choice probability function $C^p: \mathbf{R}^{n^p} \rightarrow \Delta^p$ in terms of admissible distribution v^p , as in equation (1):

$$C_i^p(\boldsymbol{\pi}^p) = v^p\left(\boldsymbol{\varepsilon}^p : i \in \arg \max_{j \in S^p} \pi_j^p + \varepsilon_j^p\right).$$

We now summarize a number of properties of this function and provide an explicit formula for its deterministic representation (2).

- (P1) $C_i^p(\boldsymbol{\pi}^p + \lambda \mathbf{1}) = C_i^p(\boldsymbol{\pi}^p)$ for all $\boldsymbol{\pi}^p \in \mathbf{R}^{n^p}$ and $\lambda \in \mathbf{R}$.
- (P2) For all $\boldsymbol{\pi}^p \in \mathbf{R}^{n^p}$, $DC^p(\boldsymbol{\pi}^p)$ is symmetric, has positive diagonal elements and negative off-diagonal elements, and has rows and columns that sum to zero.
- (P3) C^p admits a potential function $W^p: \mathbf{R}^{n^p} \rightarrow \mathbf{R}$ (i.e., a function satisfying $C^p(\boldsymbol{\pi}^p) \equiv \nabla W^p(\boldsymbol{\pi}^p)$) that is convex, strictly so on $\mathbf{R}_0^{n^p}$.
- (P4) Let $V^p: \text{int}(\Delta^p) \rightarrow \mathbf{R}$ be the Legendre transform of $W^p: \mathbf{R}_0^{n^p} \rightarrow \mathbf{R}$:

¹⁰ The minimal center of attraction is closure of the union of the supports of all probability measures on X that are invariant under (P). This set is contained in (and often identical to) the more easily computed *Birkhoff center*, which is the closure of the set of recurrent points of (P). For more on notions of recurrence for deterministic flows, see Nemytskii and Stepanov [33], Conley [12], Akin [1], Robinson [34], and Benaïm [2, 3].

$$V^p(y^p) = \max_{\pi^p \in \mathbf{R}_0^{n^p}} (y^p \cdot \pi^p - W^p(\pi^p)).$$

Then V^p is an admissible deterministic perturbation that satisfies

$$\begin{aligned} W^p(\pi^p) &= \max_{y^p \in \text{int}(\Delta^p)} (y^p \cdot \pi^p - V^p(y^p)) \text{ and} \\ C^p(\pi^p) &= \arg \max_{y^p \in \text{int}(\Delta^p)} (y^p \cdot \pi^p - V^p(y^p)) \end{aligned}$$

for all $\pi^p \in \mathbf{R}_0^{n^p}$, and $\nabla V^p: \text{int}(\Delta^p) \rightarrow \mathbf{R}_0^{n^p}$ is the inverse of $C^p: \mathbf{R}_0^{n^p} \rightarrow \text{int}(\Delta^p)$.

To prove Theorems 3.1 and 3.2, it is convenient to define the *virtual payoffs* for the pair (F, ν) by

$$\hat{F}^p(x) = F^p(x) - \nabla V^p\left(\frac{1}{m^p} x^p\right),$$

where V^p is the deterministic perturbation corresponding to ν^p . The next two lemmas provide two justifications for this definition.

The proofs of these lemmas require two additional definitions. Let $\bar{F}^p(x) = \frac{1}{n^p} \sum_{i \in S^p} F_i^p(x)$ denote the average payoff obtained by population p strategies, and let $\tilde{F}^p(x) = F^p(x) - \bar{F}^p(x)\mathbf{1} \in \mathbf{R}_0^{n^p}$ be a normalized version of the payoff vector $F^p(x)$.

The first lemma shows that perturbed equilibria are those states that equalize virtual payoffs in each population.

Lemma A.1: $x \in PE(F, \nu)$ if and only if $\hat{F}^p(x) = c^p \mathbf{1}$ for some $c^p \in \mathbf{R}$ and all $p \in \mathcal{P}$.

Proof: Observe that by properties (P1) and (P4),

$$\begin{aligned} x \in PE(F, \nu) &\Leftrightarrow x^p = m^p \tilde{B}^p(x) \text{ for all } p \in \mathcal{P} \\ &\Leftrightarrow x^p = m^p C^p(F^p(x)) \text{ for all } p \in \mathcal{P} \\ &\Leftrightarrow x^p = m^p C^p(F^p(x) - \bar{F}^p(x)\mathbf{1}) \text{ for all } p \in \mathcal{P} \\ &\Leftrightarrow \nabla V^p\left(\frac{1}{m^p} x^p\right) = F^p(x) - \bar{F}^p(x)\mathbf{1} \text{ for all } p \in \mathcal{P} \\ &\Leftrightarrow \hat{F}^p(x) = \bar{F}^p(x)\mathbf{1} \text{ for all } p \in \mathcal{P}. \end{aligned}$$

This establishes the "only if" direction. To prove the "if" direction, note that since $\nabla V^p\left(\frac{1}{m^p} x^p\right) \in \mathbf{R}_0^{n^p}$ by property (P4), $\mathbf{1} \cdot \hat{F}^p(x) = \mathbf{1} \cdot F^p(x)$. Therefore, if $\hat{F}^p(x) = c^p \mathbf{1}$, then $\mathbf{1} \cdot F^p(x) = \mathbf{1} \cdot \hat{F}^p(x) = c^p n^p$, and so $c^p = \bar{F}^p(x)$. Thus, the "if" direction follows from the equivalence derived above. ■

In settings without perturbations, one appealing monotonicity property for evolutionary dynamics requires that each population's direction of motion always forms an acute angle with its payoff vector: in other words, that $\dot{x}^p \cdot F^p(x) \geq 0$ for all $x \in X$. Sandholm [38] calls this condition *positive correlation*. The next lemma, first proved by Hofbauer [22] for a single population setting, establishes a corresponding property for the perturbed best response dynamics expressed in terms of virtual payoffs. We use the properties listed above to provide a simple proof.

Lemma A.2: $(m^p \tilde{B}^p(x) - x^p) \cdot \hat{F}^p(x) \geq 0$ for all $p \in \mathcal{P}$ and $x \in X$, with equality only if $m^p \tilde{B}^p(x) = x^p$.

Proof: Since $m^p \tilde{B}^p(x) - x^p$ is a direction of motion through X^p , $(m^p \tilde{B}^p(x) - x^p) \cdot \mathbf{1} = 0$. Also, note that $y^p \equiv \tilde{B}^p(x) = C^p(F^p(x)) = C^p(\tilde{F}^p(x))$ by property (P1), so property (P4) implies that $\nabla V^p(y^p) = \tilde{F}^p(x)$. Using these observations in turn, we find that

$$\begin{aligned} (m^p \tilde{B}^p(x) - x^p) \cdot \hat{F}^p(x) &= (m^p \tilde{B}^p(x) - x^p) \cdot (F^p(x) - \nabla V^p(\frac{1}{m^p} x^p)) \\ &= (m^p \tilde{B}^p(x) - x^p) \cdot (\tilde{F}^p(x) - \nabla V^p(\frac{1}{m^p} x^p)) \\ &= m^p (y^p - \frac{1}{m^p} x^p) \cdot (\nabla V^p(y^p) - \nabla V^p(\frac{1}{m^p} x^p)), \end{aligned}$$

which is positive by the strict convexity of V^p , strictly so unless $m^p \tilde{B}^p(x) = x^p$. ■

The Proof of Theorem 3.1

We first prove part (i). Properties (P3) and (P2) and the definition of \tilde{B}^p imply that along any solution of (P),

$$\begin{aligned} \dot{\Lambda}(x) &= \frac{d}{dt} \sum_{p \in \mathcal{P}} m^p \left[\max_{y^p \in \text{int}(\Delta^p)} (y^p \cdot F^p(x) - V^p(y^p)) - \left(\frac{1}{m^p} x^p \cdot F^p(x) - V^p\left(\frac{1}{m^p} x^p\right) \right) \right] \\ &= \frac{d}{dt} \sum_{p \in \mathcal{P}} \left(m^p W^p(F^p(x)) - \left(x^p \cdot F^p(x) - m^p V^p\left(\frac{1}{m^p} x^p\right) \right) \right) \\ &= \sum_{p \in \mathcal{P}} \left(m^p \nabla C^p(F^p(x)) \cdot DF^p(x) \dot{x} - \left(x^p \cdot DF^p(x) \dot{x} + \dot{x}^p \cdot F^p(x) - \dot{x}^p \cdot \nabla V^p\left(\frac{1}{m^p} x^p\right) \right) \right) \\ &= \sum_{p \in \mathcal{P}} \left((m^p \tilde{B}^p(x) - x^p) \cdot DF^p(x) \dot{x} - \dot{x}^p \cdot (F^p(x) - \nabla V^p\left(\frac{1}{m^p} x^p\right)) \right) \\ &= \sum_{p \in \mathcal{P}} \left(\dot{x}^p \cdot DF^p(x) \dot{x} - \dot{x}^p \cdot \hat{F}^p(x) \right) \\ &= \dot{x} \cdot DF(x) \dot{x} - \sum_{p \in \mathcal{P}} (m^p \tilde{B}^p(x) - x^p) \cdot \hat{F}^p(x) \end{aligned}$$

The first term of the last expression is negative by condition (SE); the second term is negative by Lemma A.2, strictly so only if x is a rest point of (P). This establishes part (i) of the theorem.

We now prove part (ii). First, standard results (e.g., Theorem 7.6 of Hofbauer and Sigmund [26]) tell us that since (P) admits a strict Lyapunov function, all solution trajectories of (P) converge to connected sets of rest points of (P). By definition, these rest points are the perturbed equilibria of (F, ν) . Moreover, Lemma A.1 and property (P4) imply that

$$\begin{aligned} x \in PE(F, \nu) &\Leftrightarrow \nabla V^p(\frac{1}{m^p} x^p) = F^p(x) + c^p \mathbf{1} \text{ for all } p \in \mathcal{P} \\ &\Leftrightarrow \frac{1}{m^p} x^p = \underset{y^p \in \text{int}(\mathcal{A}^p)}{\text{argmax}} (y^p \cdot F^p(x) - V^p(y^p)) \text{ for all } p \in \mathcal{P} \\ &\Leftrightarrow \Lambda(x) = 0. \end{aligned}$$

It remains to show that $PE(F, \nu)$ is a singleton. For each $x \in X$ and $h \in TX$, define

$$\hat{f}_{x,h}(t) = h \cdot \hat{F}(x + th)$$

for all t such that $x + th \in X$. Since F is stable and each V^p is differentially strictly convex, we find that

$$\hat{f}'_{x,h}(t) = h \cdot D\hat{F}(x + th)h = h \cdot DF(x + th)h - \sum_{p \in \mathcal{P}} \frac{1}{m^p} h^p \cdot D^2 V^p(\frac{1}{m^p}(x^p + th^p))h^p < 0.$$

Thus, $\hat{f}_{x,h}(t)$ is decreasing in t .

If $x \in PE(F, \nu)$, then Lemma A.1 implies that $\hat{f}_{x,h}(0) = h \cdot \hat{F}(x) = 0$ for all $h \in TX$. Now let y be a state in X distinct from x , so that $y = x + t_y h_y$ for some $t_y > 0$ and nonzero $h_y \in TX$. Then $h_y \cdot \hat{F}(y) = h_y \cdot \hat{F}(x + t_y h_y) = \hat{f}_{x,h_y}(t_y) < 0$, and so y cannot be in $PE(F, \nu)$. We therefore conclude that $PE(F, \nu)$ is a singleton containing the unique state at which Λ equals zero, and that this state is globally asymptotically stable under (P). ■

The Proof of Theorem 3.2

Condition (ES) implies that along solutions of (P),

$$\begin{aligned} \dot{\Pi}(x) &= \nabla f(x) \cdot \dot{x} - \sum_{p \in \mathcal{P}} \nabla V^p(\frac{1}{m^p} x^p) \cdot \dot{x}^p \\ &= \sum_{p \in \mathcal{P}} (F^p(x) - \nabla V^p(\frac{1}{m^p} x^p)) \cdot \dot{x}^p \end{aligned}$$

$$= \sum_{p \in \mathcal{P}} \hat{F}^p(x) \cdot \dot{x}^p$$

By Lemma A.2, this expression is positive and equals zero only at rest points of (P). Hence, Π is a strict Lyapunov function for (P), implying global convergence of solution trajectories of (P) to connected subsets of $x \in PE(F, \varepsilon)$. Finally, Lemma A.1 tells us that

$$\begin{aligned} x \text{ is a critical point of } \Pi \text{ in } X &\Leftrightarrow F^p(x) - c^p \mathbf{1} = \nabla V^p\left(\frac{1}{m^p} x^p\right) \text{ for all } p \in \mathcal{P} \\ &\Leftrightarrow x \in PE(F, \varepsilon). \blacksquare \end{aligned}$$

The Proof of Theorem 3.3

It is useful to study the dynamic (P) after applying the change of variable T . To do so, we let $T^p[X^p] = \{v^p \in \mathbf{R}^{n^p-1} : m^p \geq v_1^p \geq \dots \geq v_{n^p-1}^p \geq 0\}$, so that $T[X] = \prod_p T^p[X^p]$ is the transformation of the state space X by T . Note that if $v \in T[X]$, the set of components of v is \hat{S} . If we then define $\hat{B}^p: T[X] \rightarrow T^p[\Delta^p]$ by $\hat{B}^p(v) = T^p \tilde{B}^p(T^{-1}v)$, the transformed dynamic is given by

$$(T) \quad \dot{v}^p = m^p \hat{B}^p(v) - v^p.$$

One can verify (P) and (T) are linearly conjugate: $\{x_t\}_{t \geq 0}$ solves (P) if and only if $\{Tx_t\}_{t \geq 0}$ solves (T).

Our goal is to show that the dynamic (T) is cooperative and irreducible. A differential equation $\dot{v} = g(v)$ on $T[X]$ is called *cooperative* if $\frac{\partial g_k^p}{\partial v_j^q}(v) \geq 0$ for all $v \in T[X]$ and all distinct pairs $(k, p), (j, q) \in \hat{S}$. The equation is *irreducible* if for each $v \in T[X]$ and each nonempty proper subset K of \hat{S} , there is an $(k, p) \in K$ and a $(j, q) \in \hat{S} - K$ such that $\frac{\partial g_k^p}{\partial v_j^q}(v) \neq 0$. Theorem 4.1.1 of Smith [39] shows that the flow of a cooperative irreducible dynamic is strongly monotone with respect to the standard vector order. Thus, if (T) is cooperative and irreducible, our first claim follows from this result and the conjugacy of (P) and (T), the second claim follows in turn from Theorem 2.4.7 of Smith [39] and Theorem 1.1 of Hirsch [20], and the third claim is proved as follows: Suppose that x^* is the unique perturbed equilibrium of (F, v) . Then if \underline{x} and \bar{x} are the minimal and maximal points in X , Theorem 1.2.1 of Smith [39] implies that the solutions to (P) from these points converge to rest points, and hence to x^* . Thus, for any $x \in X$, strong monotonicity implies that at all times t , the solutions to (P) starting from \underline{x} , x , and \bar{x} are ranked by T . Therefore, the solution to (P) from x must also converge to x^* , and x^* is Lyapunov stable.

We now show that (T) is cooperative and irreducible. Fix $v \in T[X]$, and let $x = T^{-1}v \in X$. Since $B^p(x) = C^p(F^p(x))$, the off-diagonal elements of the derivative matrix for (T) are given by $m^p \frac{\partial \hat{B}_k^p}{\partial v_j^q}(v)$, where

$$\begin{aligned} \frac{\partial \hat{B}_k^p}{\partial v_j^q}(v) &= \sum_{l=k+1}^{n^p} \frac{\partial B_l^p}{\partial (e_{j+1}^q - e_j^q)}(x) \\ &= \sum_{l=k+1}^{n^p} \sum_{i=1}^{n^p} \frac{\partial C_l^p}{\partial \pi_i^p}(F^p(x)) \frac{\partial F_i^p}{\partial (e_{j+1}^q - e_j^q)}(x) \\ &= \sum_{i=1}^{n^p} \frac{\partial F_i^p}{\partial (e_{j+1}^q - e_j^q)}(x) \sum_{l=k+1}^{n^p} \frac{\partial C_l^p}{\partial \pi_i^p}(F^p(x)) \\ &= \frac{\partial F_{n^p}^p}{\partial (e_{j+1}^q - e_j^q)}(x) \left[\sum_{l=k+1}^{n^p} \sum_{i=1}^{n^p} \frac{\partial C_l^p}{\partial \pi_i^p}(F^p(x)) \right] - \sum_{h=1}^{n^p} \frac{\partial (F_{h+1}^p - F_h^p)}{\partial (e_{j+1}^q - e_j^q)}(x) \left[\sum_{l=k+1}^{n^p} \sum_{i=1}^h \frac{\partial C_l^p}{\partial \pi_i^p}(F^p(x)) \right], \end{aligned}$$

where the last equality follows from the fact that

$$\sum_{i=1}^{n^p} f_i c_i = f_{n^p} \sum_{i=1}^{n^p} c_i - \sum_{h=1}^{n^p-1} (f_{h+1} - f_h) \sum_{i=1}^h c_i$$

for any pair of vectors $f, c \in \mathbf{R}^{n^p}$. Property (P2) implies that the first expression in brackets is zero and that the second expression in brackets is strictly negative for all $h < n^p$ and equals zero if $h = n^p$. Furthermore, condition (SC) implies that the directional derivative from the second term is always positive. Thus, $m^p \frac{\partial \hat{B}_k^p}{\partial v_j^q}(v) \geq 0$ for all distinct pairs $(\alpha, k), (\beta, j) \in \hat{S}$, and so (T) is cooperative.

To show that (T) is irreducible, fix a nonempty proper subset K of \hat{S} . Since F is irreducible by assumption, there exist a pair $(k, p) \in K$, a strategy $h \in S^p - \{n^p\}$, and a pair $(j, q) \in \hat{S} - K$ such that $\frac{\partial (F_{h+1}^p - F_h^p)}{\partial (e_{j+1}^q - e_j^q)}(x) > 0$. Hence, the reasoning above implies that $m^p \frac{\partial \hat{B}_k^p}{\partial v_j^q}(v) > 0$, so (T) is irreducible. This completes the proof of the theorem. ■

The Proof of Theorem 4.1

Theorem 4.1 of Sandholm [36], based on results of Kurtz [30], shows that over any finite horizon, the stochastic process X_i^N stays within $\frac{\varepsilon}{2}$ of the solution trajectory of (P) with the same initial condition with probability close to 1 when N is large. Theorems 3.1 and 3.2 show that in the games considered in parts (i) and (ii), all solution trajectories of (P) converge to $PE(F, v)$; Theorem 3.3 shows that in supermodular games,

this is true of trajectories starting from almost every initial condition. Combining these results proves parts (i), (ii), and (iii) of the theorem.

To prove the final claim, suppose that F has a unique equilibrium. Theorems 3.1, 3.2, and 3.3 imply that for our three classes of games, a unique equilibrium is globally asymptotically stable. The final claim then follows from this classical result from dynamical systems.

Lemma A.3: *Let x^* be globally asymptotically stable the flow ϕ on the compact set X . Fix $\gamma > 0$, and let $\tau(x) = \inf\{T: |\phi(t, x) - x^*| \leq \gamma \text{ for all } t \geq T\}$. Then $\sup_{x \in X} \tau(x) < \infty$.*

Proof: Since x^* is globally asymptotically stable, $\tau(x) < \infty$ for all $x \in X$. Now suppose that the lemma is false. Then there is a sequence of initial conditions $\{x^k\} \subset X$ such that $\lim_{k \rightarrow \infty} \tau(x^k) = \infty$. Since X is compact, this sequence has an accumulation point $\bar{x} \in X$. Because x^* is Lyapunov stable, there is an $\eta > 0$ such that whenever $|x - x^*| \leq \eta$, $|\phi(t, x) - x^*| \leq \gamma$ for all $t \geq 0$. Because x^* is a global attractor, there is a time $\bar{T} < \infty$ such that $|\phi(\bar{T}, \bar{x}) - x^*| \leq \frac{\eta}{2}$. Finally, since the flow is continuous in the initial condition x , we know that for all x sufficiently close to \bar{x} , $|\phi(\bar{T}, x) - \phi(\bar{T}, \bar{x})| \leq \frac{\eta}{2}$. Therefore, for all sufficiently large k , the triangle inequality implies that $|\phi(\bar{T}, x^k) - x^*| \leq \eta$, and hence that $|\phi(t, x^k) - x^*| \leq \gamma$ for all $t \geq \bar{T}$. But then $\tau(x^k) \leq \bar{T}$ for all sufficiently large k , contradicting the definition of the sequence $\{x^k\}$. \square

This completes the proof of Theorem 4.1. \blacksquare

The Proof of Theorem 4.2

The proof of parts (i) and (ii) rely on results from Benaïm [2] (hereafter B98). One can verify that Hypotheses 2.1, 2.3, and 3.4 of B98 are all satisfied (cf. B98 Example 1.1). Thus, part (i) of theorem follows directly from B98 Corollary 3.2 and our Theorem 3.1 (in particular, from the fact that the lone element of $PE(F, \nu)$ is the unique ω -limit point of (P)).

The proof of part (ii) utilizes B98 Theorem 4.3. Condition (i) of this theorem follows from Proposition 3.2 and the finiteness of $PE(F, \nu)$. Condition (ii) follows from B98 Remark 3.10 (iii) and the fact that all rest points of (P) are in $\text{int}(X)$. Condition (iii) follows from the fact that X_t^N is defined on X . Finally, since $PE(F, \nu)$ is finite, and since by our Theorem 3.2 (P) is gradient-like, the discussion on p.69 of B98 implies that the weakly stable equilibria are those that coincide with their own unstable manifolds; these are simply the local maximizers of Π , or equivalently the Lyapunov stable rest points $LS(F, \nu)$. This completes the proof of part (ii) of the theorem.

We now turn to the proof of part (iii). To begin, we establish a nondegeneracy condition on the motions of X_t^N . For each $x \in X$, let ξ^x be a random vector that is defined on an arbitrary probability space Ω and that describes the normalized increments of the process X_t^N from state x . The distribution of ξ^x is described by

$$\begin{aligned} P(\xi^x = e_j^p - e_i^p) &= \frac{1}{m} x_i^p \tilde{B}_j^p(x) \text{ whenever } i, j \in S^p, i \neq j \text{ and } p \in \mathcal{P}; \\ P(\xi^x = 0) &= \frac{1}{m} \sum_{p \in \mathcal{P}} \sum_{i \in S^p} x_i^p \tilde{B}_i^p(x). \end{aligned}$$

Let $\Sigma^x \in \mathbf{R}^{n \times n}$ denote the covariance matrix of ξ^x . Since Σ^x is symmetric, its eigenvalues are real. Let λ^x be the smallest eigenvalue of Σ^x corresponding to an eigenvector in TX . (One can show that the remaining eigenvectors are orthogonal to TX and have eigenvalues of zero.) We want to show that λ^x is uniformly bounded away from zero. Intuitively, this means that for any current state x and any direction of motion z in TX , the amount of randomness in the motion of the process X_t^N in the direction z is nonnegligible.

To establish the bound on λ^x , we let

$$\beta \equiv \min_{x \in \Delta} \min_{p \in \mathcal{P}} \min_{i \in S^p} \tilde{B}_i^p(x) > 0.$$

Lemma A.4: *For all $x \in X$, the minimum eigenvalue λ^x is at least $\frac{\beta}{m}$.*

Proof: Since Σ^x is symmetric, we know that if $\theta \in TX$ is a unit length eigenvector of Σ^x , the corresponding eigenvalue is $\theta \cdot \Sigma^x \theta = \text{Var}(\theta \cdot \xi^x)$. It is therefore sufficient to bound $\text{Var}(\theta \cdot \xi^x)$ for all unit length $\theta \in TX$.

Partition the probability space Ω into events I^p , where I^p is the event that the individual who receives the revision opportunity is from population p . Then $P(I^p) = \frac{m^p}{m}$, and all realizations of ξ^x involving nonzero increments for population p occur on I^p . Letting $y^p = \frac{1}{m^p} x^p$, we note the following conditional probabilities and expectations:

$$\begin{aligned}
P(\xi^x = e_j^p - e_i^p | I^p) &= y_i^p \tilde{B}_j^p(x) && \text{if } i \neq j; \\
P(\xi^x = e_j^q - e_i^q | I^p) &= 0 && \text{if } i \neq j \text{ and } q \neq p; \\
P(\xi^x = 0 | I^p) &= \sum_{i \in S^p} y_i^p \tilde{B}_i^p(x); \\
E(\xi_i^{x,p} | I^p) &= \tilde{B}_i^p(x) - y_i^p; \\
E(\xi_i^{x,p} \xi_j^{x,p} | I^p) &= -y_i^p \tilde{B}_i^p(x) - y_j^p \tilde{B}_j^p(x) && \text{if } i \neq j; \\
E((\xi_i^{x,p})^2 | I^p) &= y_i^p (1 - \tilde{B}_i^p(x)) + (1 - y_i^p) \tilde{B}_i^p(x); \\
\text{Cov}(\xi_i^{x,p}, \xi_j^{x,p} | I^p) &= -\tilde{B}_i^p(x) \tilde{B}_j^p(x) - y_i^p y_j^p && \text{if } i \neq j; \\
\text{Var}(\xi_i^{x,p} | I^p) &= -(\tilde{B}_i^p(x))^2 - (y_i^p)^2 + \tilde{B}_i^p(x) + y_i^p.
\end{aligned}$$

Fix a unit length $\theta \in TX$, and let \mathcal{F} be the σ -algebra generated by the events I^p . A standard decomposition of variance (see, e.g., Durrett [15]) yields

$$\begin{aligned}
\text{Var}(\theta \cdot \xi^x) &= E[\text{Var}[\theta \cdot \xi^x | \mathcal{F}]] + \text{Var}[E[\theta \cdot \xi^x | \mathcal{F}]] \\
&\geq E[\text{Var}[\theta \cdot \xi^x | \mathcal{F}]] \\
&= \sum_{p \in \mathcal{P}} \frac{m^p}{m} \text{Var}(\theta \cdot \xi^x | I^p) \\
&= \sum_{p \in \mathcal{P}} \frac{m^p}{m} \text{Var}(\theta^p \cdot \xi^{x,p} | I^p),
\end{aligned}$$

where the final equality follows from the fact that $P(\xi^x = e_j^q - e_i^q | I^p) = 0$ for $q \neq p$. Since y^p and $\tilde{B}^p(x)$ lie in the simplex Δ^p and since $\sum_{i \in S^p} \theta_i^p = 0$ for all p (because $\theta \in TX$), we can use the conditional probabilities and expectations above to compute that

$$\begin{aligned}
&\text{Var}(\theta^p \cdot \xi^{x,p} | I^p) \\
&= \theta^p \cdot \text{diag}(\tilde{B}^p(x)) \theta^p - \theta^p \cdot \tilde{B}^p(x) \tilde{B}^p(x) \cdot \theta^p + \theta^p \cdot \text{diag}(y^p) \theta^p - \theta^p \cdot y^p y^p \cdot \theta^p
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i \in S^p} (\theta_i^p)^2 \tilde{B}_i^p(x) - \left(\sum_{j \in S^p} \theta_j^p \tilde{B}_j^p(x) \right)^2 + \sum_{i \in S^p} (\theta_i^p)^2 y_i^p - \left(\sum_{j \in S^p} \theta_j^p y_j^p \right)^2 \\
&= \sum_{i \in S^p} \left(\theta_i^p - \sum_{j \in S^p} \theta_j^p \tilde{B}_j^p(x) \right)^2 \tilde{B}_i^p(x) + \sum_{i \in S^p} \left(\theta_i^p - \sum_{j \in S^p} \theta_j^p y_j^p \right)^2 y_i^p \\
&\geq \beta \sum_{i \in S^p} \left(\theta_i^p - \sum_{j \in S^p} \theta_j^p \tilde{B}_j^p(x) \right)^2 \\
&= \beta \left[\sum_{i \in S^p} (\theta_i^p)^2 - 2 \left(\sum_{i \in S^p} \theta_i^p \right) \left(\sum_{j \in S^p} \theta_j^p \tilde{B}_j^p(x) \right) + n^p \left(\sum_{j \in S^p} \theta_j^p \tilde{B}_j^p(x) \right)^2 \right] \\
&\geq \beta \sum_{i \in S^p} (\theta_i^p)^2.
\end{aligned}$$

Since each m^p is a positive integer and since $\sum_{p \in \mathcal{P}} \sum_{i \in S^p} (\theta_i^p)^2 = 1$, we conclude that

$$\text{Var}(\theta \cdot \xi^x) \geq \sum_{p \in \mathcal{P}} \frac{\beta m^p}{m} \sum_{i \in S^p} (\theta_i^p)^2 \geq \frac{\beta}{m}. \quad \square$$

Now, if we can show that the conditions supporting Theorem 1.5 of Benaïm and Hirsch [6] (henceforth BH) hold, part (iii) of our theorem immediately follows. The proof of Theorem 3.3 shows that after a linear transformation, the dynamics (P) form a cooperative, irreducible dynamical system on X , so BH Hypothesis 1.2 is satisfied. Since the increments are uniformly bounded above, and since λ^x is uniformly bounded below by Lemma A.4, BH Proposition 2.3 implies that BH Hypothesis 1.4 holds. Finally, since each X_t^N takes values in the compact set X , the tightness assumption in BH Theorem 1.5 is satisfied. Therefore, BH Theorem 1.5 implies that $\lim_{N \rightarrow \infty} \mu^N(O) = 1$ for any open set O containing the Lyapunov stable rest points of (P). ■

References

1. E. Akin, "The General Topology of Dynamical Systems," Providence, RI, American Mathematical Society, 1993.
2. M. Benaïm, Recursive algorithms, urn processes, and chaining number of chain recurrent sets, *Ergod. Th. & Dynam. Sys.* **18** (1998), 53-87.
3. M. Benaïm, Dynamics of stochastic approximation algorithms, In "Séminaire de Probabilités XXXIII" (J. Azéma *et. al.*, Eds.), Berlin, Springer, 1999.
4. M. Benaïm, Convergence with probability one of stochastic approximation algorithms whose average is cooperative, *Nonlinearity* **13** (2000), 601-616.
5. M. Benaïm and M. W. Hirsch, Mixed equilibria and dynamical systems arising from repeated games, *Games Econ. Behav.* **29** (1999), 36-72.
6. M. Benaïm and M. W. Hirsch, On stochastic approximation algorithms with constant step size whose average is cooperative, *Ann. Appl. Prob.* **30** (1999), 850-869.
7. M. Benaïm and J. Weibull, Deterministic approximation of stochastic evolution in games, *Econometrica* **71** (2003), 873-903.
8. K. J. Binmore and L. Samuelson, Evolutionary drift and equilibrium selection, *Rev. Econ. Stud.* **66** (1999), 363-393.
9. L. E. Blume, The statistical mechanics of strategic interaction, *Games Econ. Behav.* **5** (1993), 387-424.
10. L. E. Blume, Population games, In "The Economy as an Evolving Complex System II" (W. B. Arthur, S. N. Durlauf, and D. A. Lane, Eds.), Reading, MA, Addison-Wesley, 1997.
11. G. W. Brown, Iterative solutions of games by fictitious play, In "Activity Analysis of Production and Allocation" (T. C. Koopmans, Ed.), New York, Wiley, 1951.
12. C. C. Conley, "Isolated Invariant Sets and the Morse Index," Providence, RI, American Mathematical Society, 1978.
13. R. W. Cooper, "Coordination Games: Complementarities and Macroeconomics," Cambridge, Cambridge University Press, 1999.
14. R. Cressman, J. Garay, and J. Hofbauer, Evolutionary stability concepts for N -species frequency-dependent interactions, *J. Theor. Biol.* **211** (2001), 1-10.
15. R. Durrett, "Probability: Theory and Examples," 3rd ed., Belmont, CA, Thomson Brooks/Cole, 2005.
16. D. Foster and H. P. Young, Stochastic evolutionary game dynamics, *Theor. Pop. Biol.* **38** (1990), 219-232.

17. D. Fudenberg and D. M. Kreps, Learning mixed equilibria, *Games Econ. Behav.* **5** (1993), 320-367.
18. I. Gilboa and A. Matsui, Social stability and equilibrium, *Econometrica* **59** (1991), 859-867.
19. J. C. Harsanyi, Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points, *Int. J. Game Theory* **2** (1973), 1-23.
20. M. W. Hirsch, Systems of differential equations that are competitive or cooperative III: competing species, *Nonlinearity* **1** (1988), 51-71.
21. J. Hofbauer, Stability for the best response dynamics, unpublished manuscript, Universität Wien, 1995.
22. J. Hofbauer, From Nash and Brown to Maynard Smith: equilibria, dynamics, and ESS, *Selection* **1** (2000), 81-88.
23. J. Hofbauer, Remarks on stable games, unpublished manuscript, Universität Wien, 2003.
24. J. Hofbauer and E. Hopkins, Learning in perturbed asymmetric games, *Games Econ. Behav.* **52** (2005), 133-152.
25. J. Hofbauer and W. H. Sandholm, On the global convergence of stochastic fictitious play, *Econometrica* **70** (2002), 2265-2294.
26. J. Hofbauer and K. Sigmund, "The Theory of Evolution and Dynamical Systems," Cambridge, Cambridge University Press, 1988.
27. M. Kandori, G. J. Mailath, and R. Rob, Learning, mutation, and long run equilibria in games, *Econometrica* **61** (1993), 29-56.
28. M. Kandori and R. Rob, Evolution of equilibria in the long run: a general theory and applications, *J. Econ. Theory* **65** (1995), 383-414.
29. Y. Kaniovski and H. P. Young, Learning dynamics in games with stochastic perturbations, *Games Econ. Behav.* **11** (1995), 330-363.
30. T. G. Kurtz, Solutions of ordinary differential equations as limits of pure jump Markov processes, *J. Appl. Prob.* **7** (1970), 49-58.
31. P. Milgrom and J. Roberts, Rationalizability, learning, and equilibrium in games with strategic complementarities, *Econometrica* **58** (1990), 1255-1278.
32. D. Monderer and L. S. Shapley, Potential games, *Games Econ. Behav.* **14** (1996), 124-143.
33. V. V. Nemytskii and V. V. Stepanov, "Qualitative Theory of Differential Equations," Princeton, Princeton University Press, 1960.
34. C. Robinson, "Dynamical Systems: Stability, Symbolic Dynamics, and Chaos," Boca Raton, FL, CRC Press, 1995.

35. W. H. Sandholm, Potential games with continuous player sets, *J. Econ. Theory* **97** (2001), 81-108.
36. W. H. Sandholm, Evolution and equilibrium under inexact information, *Games Econ. Behav.* **44** (2003), 343-378.
37. W. H. Sandholm, Stable games, unpublished manuscript, University of Wisconsin, 2004.
38. W. H. Sandholm, Negative externalities and evolutionary implementation, forthcoming, *Review of Economic Studies*.
39. H. L. Smith, "Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems," Providence, RI, American Mathematical Society, 1995.
40. D. Topkis, Equilibrium points in nonzero-sum n -person submodular games, *SIAM J. Control Opt.* **17** (1979), 773-787.
41. E. van Damme, "Stability and Perfection of Nash Equilibria," 2nd ed., Berlin, Springer, 1991.
42. X. Vives, Nash equilibrium with strategic complementarities, *J. Math. Econ.* **19** (1990), 305-321.
43. H. P. Young, The evolution of conventions, *Econometrica* **61** (1993), 57-84.
44. H. P. Young, "Individual Strategy and Social Structure," Princeton, Princeton University Press, 1998.