Excess Payoff Dynamics, Potential Dynamics, and Stable Games

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

June 21, 2004

* This manuscript combines two working papers, one of the same title and another entitled "A Probabilistic Characterization of Integrability". I thank seminar audiences at UCLA, Chicago, Johns Hopkins, Penn, Princeton, Virginia Tech, Washington University, Wisconsin, the 2003 Econometric Society Summer Meeting, the 2003 Illinois Conference on Learning and Bounded Rationality, and the 2004 Kyoto Workshop on Game Dynamics for helpful comments. I am especially grateful to Martin Cripps, Drew Fudenberg, Josef Hofbauer, George Mailath, Aki Matsui, Larry Samuelson, Satoru Takahashi, Tymon Tatur, Jörgen Weibull, and Peyton Young for enlightening discussions of this work. The comments of two referees and an Associate Editor are also sincerely appreciated. Financial support from NSF Grant SES-0092145 is gratefully acknowledged.

Abstract

We consider a model of evolution in games in which agents occasionally receive opportunities to switch strategies, choosing between them using a probabilistic rule. Both the rate at which revision opportunities arrive and the probabilities with which each strategy is chosen are functions of current normalized payoffs. We call the aggregate dynamics induced by this model *excess payoff dynamics*. We prove that these dynamics satisfy existence, uniqueness, and continuity of solutions, respect a basic payoff monotonicity property, and have rest points at and only at the Nash equilibria of the underlying game. We show that the dynamics globally converge to Nash equilibrium in potential games. Finally, we introduce a new class of games called *stable games*, which include games with an interior ESS, zero sum games, and concave potential games as special cases. We show that while excess payoff dynamics can exhibit periodic behavior in stable games, a subset of these dynamics called *potential dynamics* always converge to equilibrium. We use this convergence result to prove that every stable game admits a unique component of Nash equilibria.

1. Introduction

Evolutionary game theory is the study of strategic interactions in large populations. Agents in these populations are assumed to base their decisions on simple myopic rules, and their aggregate behavior is described by a dynamic on the space of strategy distributions. A basic goal of the theory is to find plausible decision procedures that induce appealing aggregate dynamics. One can attempt to derive dynamics that satisfy certain general desiderata regardless of the strategic interaction in question, or that fulfill more demanding requirements when applied to certain specific classes of games.

In this paper, we introduce a simple model of behavior in population games, and show that the resulting dynamics satisfy desiderata of both of these sorts. In our model, agents occasionally receive opportunities to revise their strategies, and make a stochastic selection among the available strategies when such opportunities occur. Both the rate at which revision opportunities arrive and the probabilities with which the various strategies are chosen depend on the strategies' current relative payoffs—that is, on the the differences between the strategies' payoffs and the population's average payoff.

Our model is defined in terms of objects called *raw choice functions*. The inputs to these functions, called *excess payoff vectors*, capture the relative payoffs of each strategy. The outputs of raw choice functions are positive vectors, and define the revision protocol in a simple way: revision rates are determined by the *sum* of the components of the raw choice vector, while choice probabilities are *proportional* to the components of the raw choice vector.

Given the payoff functions that define the strategic interaction and the raw choice function that defines the revision protocol, the evolution of aggregate behavior can be described by the solutions to a certain differential equation. This equation, called the *mean dynamic*, is derived from the expected changes in aggregate behavior induced by the revision protocol and the underlying game.[1]

Our first goal in this paper is to find conditions on raw choice functions such that the resulting mean dynamics satisfy three broad desiderata, regardless of the nature of the underlying strategic interaction. The properties we consider are existence, uniqueness, and continuity of solution trajectories, positive correlation between strategies' growth rates and payoffs, and equivalence of stationary states and Nash equilibria.

[1] See Binmore and Samuelson, Sandholm (2003a), Benaïm and Weibull (2003), and Section 2.2 below.

A few words on the appeal of each of these properties is in order. Existence, uniqueness, and continuity (EUC) is important for a variety of reasons. For one, failures of uniqueness can be attributed to discontinuities in choice protocols, which in turn reflect an extreme sensitivity of the agents' behavior to the exact value of the population state. In most contexts, this level of sensitivity seems unrealistic, so it is natural to consider models that do not demand it. Taking the modeler's point of view, failure of (EUC) means that slight inaccuracies in information about initial conditions can generate large errors in predictions of future play, even over short time spans. In fact, models that fail (EUC) can exhibit very complicated behavior, and so can be difficult to analyze.[2] Thus, for purposes of predictive accuracy and tractability, condition (EUC) is desirable as well.

Positive correlation (PC) is a payoff monotonicity condition that is among the weakest considered in the evolutionary literature.[3] Some form of monotonicity seems necessary if the evolutionary process is to be interpreted as a model of informed myopic choice. Moreover, such a condition seems necessary for convergence to Nash equilibrium to occur. Nash stationarity (NS) requires a one-to-one link between the stationary states of the evolutionary dynamic and the strategy profiles from which no profitable unilateral deviations exist. Again, if our dynamics are to be understood as a describing informed myopic choice, then the identification of stationary states and Nash equilibria seems a natural property to require.

In Theorem 3.1, we show that properties (EUC), (PC), and (NS) are satisfied by a simple and broad class of evolutionary dynamics. In particular, the mean dynamic for the model above exhibits these three properties as long as the raw choice functions $\tilde{\sigma}$ satisfy two mild conditions: (*Lipschitz*) *continuity* and *acuteness*. The meaning of the former condition is clear. The latter condition requires that each excess payoff vector π with a positive component forms an acute angle with the corresponding raw choice vector $\tilde{\sigma}(\pi)$. This condition serves two distinct roles. Whenever payoff improvement opportunities exist (i.e., whenever some π_i is positive), acuteness ensures both that revision opportunities continue to arrive, and that during such opportunities, agents exhibit some tendency to select strategies that perform above average. We call the model of choice generated continuous and

2 See Hofbauer (1995).

3 The literature offers a variety of such conditions: see Nachbar (1990), Friedman (1991), Samuelson and Zhang (1992), Swinkels (1993), Ritzberger and Weibull (1995), and Hofbauer and Weibull (1996). Positive correlation is very similar to conditions proposed in Friedman (1991) and Swinkels (1993), and is less demanding than conditions introduced in the other papers.

acute raw choice functions one of *competent play*, and we call the induced class of evolutionary dynamics *excess payoff dynamics*.

Theorem 3.1 is worthy of note because few of the dynamics considered in the literature satisfy all three of the desiderata proposed above. The replicator dynamic, and indeed all imitative dynamics with monotone percentage growth rates, satisfy (EUC) and (PC) but fail (NS). While all Nash equilibria are stationary states of these dynamics, the dynamics also possess stationary boundary states that are not Nash equilibria.[4] The best response dynamic satisfies (PC) and (NS) but fails (EUC). This dynamic, like the correspondence on which it is based, is discontinuous, and for this reason multiple solutions can emanate from a single initial condition.[5] Finally, all perturbed best response dynamics satisfy (EUC) but fail (PC) and (NS). Because these dynamics are defined in terms of arbitrary perturbations of payoffs, the connection between growth rates and the underlying payoffs breaks down when there are strategies that are seldom used or payoff differences that are small; for the same reason, stationary states only approximate Nash equilibria.[6] Thus, while it is not difficult to find dynamics that satisfy some subset of the desiderata, it seems more challenging to construct classes of dynamics that satisfy all three of them at once.

There is one canonical dynamic that satisfies all three of our desiderata: namely, the Brown-von Neumann-Nash (BNN) dynamic.[7] Interestingly, the BNN dynamic is actually the simplest example of an excess payoff dynamic; it is generated when the raw choice function $\tilde{\sigma}$ takes a separable semilinear form. Our analysis therefore provides a microfoundation for the BNN dynamic, and it also shows that very little of the structure provided by this dynamic is needed for our desiderata to hold.

We noted earlier that the replicator dynamic and related imitative dynamics fail Nash stationarity. Happily, the constructions studied in this paper can be used to alleviate this difficulty in a minimally intrusive fahsion. In Theorem 4.1, we show that every imitative dynamic can be modified arbitrarily slightly in such a way that the modified dynamic satisfies all three of our desiderata. These modified dynamics are convex combinations of the imitative dynamic and an arbitrary excess payoff

4 See Taylor and Jonker (1978), Nachbar (1990), Friedman (1991), Samuelson and Zhang (1992), and Weibull (1995, Chapter 4), as well as Section 4 below.

5 See Gilboa and Matsui (1991) and Matsui (1992). For conditions (PC) and (NS) to be satisfied by the best response dynamic, they must be modified to account for the dynamic's multivalued form.

6 See Fudenberg and Levine (1998, Chapter 4), Hopkins (1999), Hofbauer (2000), Hofbauer and Hopkins (2003), and Hofbauer and Sandholm (2002, 2003).

7 See Brown and von Neumann (1950) and Section 2.3.1 below.

dynamic; they can be derived from choice protocols that usually rely on imitation but occasionally require competent play.

The second half of this paper addresses convergence to equilibrium in two specific classes of games. To begin, we prove in Theorem 5.1 that all excess payoff dynamics converge to Nash equilibrium whenever the underlying game is a potential game (Monderer and Shapley (1996), Sandholm (2001)). The fact that potential games have desirable convergence properties is well known, but as these games appear in a variety of applications (e.g., in models of oligopoly, congestion, externality pricing, and implementation), establishing explicit convergence results is worthwhile.

We then introduce a new class of games called *stable games*, a class characterized by a property called *self-defeating externalities*. This property requires that whenever a small group of agents changes strategies, the effect of this change on the strategies to which the agents switch is worse than the effect on the strategies that the agents abandon. Classes of games that satisfy this property include games with an interior ESS, zero-sum games, and concave potential games, and there are new classes of games of economic interest that satisfy this property as well.

Turning again to dynamics, we provide examples illustrating that excess payoff dynamics can exhibit periodic behavior in stable games. However, we show in Theorems 8.1 and 8.2 that if we impose a third condition on raw choice functions, we can prove that all Nash equilibria lie in a single connected component that is globally stable. The new condition on excess payoff functions we call *integrability*, and the resulting subclass of excess payoff dynamics we call *potential dynamics*.

Formally, integrability requires that the raw choice function $\tilde{\sigma}$ be expressible as the gradient $\nabla\psi$ of some *choice potential function* ψ. It is not immediately obvious whether this condition admits an economic interpretation. To address this question, we prove in Theorems A.1 and A.2 that integrability is equivalent to a lack of correlation between the raw choice weights on each strategy i and a statistic summarizing the performances of strategies other than i. Very roughly, integrability allows the raw choice weights on each strategy to depend on the excess payoffs of other strategies, but does not allow this dependence to take a systematic form.

Before proceeding to our model, we should point out connections between it and work of Hart and Mas-Colell (2001) on adaptive learning in repeated games. These authors construct a class of *consistent* repeated game strategies: strategies that ensure that in the long run, and for all possible sequences of opponents' plays, the payoff that a player obtains in the repeated game is as high as the best payoff he could have obtained had he known the empirical frequencies of his opponents' choices in advance.[8] The three conditions on raw choice functions considered in this paper—continuity, acuteness, and integrability—are nearly identical to those that Hart and Mas-Colell (2001) use to construct their consistent repeated game strategies. It is perhaps surprising that the decision rules introduced by Hart and Mas-Colell (2001) for an adaptive learning framework also prove fruitful in an evolutionary setting, despite substantial differences in the contexts, questions posed, and requisite analytical techniques.

2. The Model

2.1 A Random Matching Model

To introduce our evolutionary dynamics in the simplest possible setting, we describe a model in which a single population of agents is recurrently randomly matched to play a symmetric normal form game. We present a more general model of evolution in Section 2.4.

Let $S = \{1, \ldots, n\}$ be a set of strategies from which individual agents choose, and let $A \in \mathbf{R}^{n \times n}$ be a payoff matrix. Component A_{ij} represents the payoff obtained by an agent who chooses action i when his opponent chooses action j.

A large, finite population of agents is recurrently randomly matched to play the game with payoff matrix A. A *population state* is a vector x in the simplex $\Delta = \{x \in \mathbf{R}^n_+ \colon \sum_i x_i = 1\}$; component x_i represents the current proportion of agents choosing strategy i. More precisely, when the population size is N, the state is a point in the discrete grid $\{x \in \Delta \colon Nx \in \mathbf{Z}^n\}$.

If an agent chooses action i when the population state is x, his (expected) payoff is $F_i(x) = (Ax)_i = e_i \cdot Ax$; the average realized payoff at this population state is $\bar{F}(x) = x \cdot Ax$. We define the *excess payoff* of strategy i as the difference between the two:

$$\hat{F}_i(x) = F_i(x) - \bar{F}(x).$$

The *excess payoff vector* $\hat{F}(x) \in \mathbf{R}^n$ is given by

$$\hat{F}(x) = F(x) - \mathbf{1}\,\bar{F}(x),$$

[8] Hannan (1957) and Blackwell (1956) were the first to construct repeated game strategies that are consistent in this sense.

where $\mathbf{1} \in \mathbf{R}^n$ is a vector of ones.

2.2 Choice Rules and Revision Rates

We now introduce our model of competent play in games. In this model, agents receive revision opportunities via independent, variable rate Poisson processes. When an agent receives such an opportunity, he considers switching strategies. Both the rate at which agents receive revision opportunities and the probabilities with which they choose each strategy are functions of current excess payoffs.

Payoffs influence strategy choices in all evolutionary models. Allowing payoffs to influence revision rates is less common,[9] but seems reasonable in many contexts. For instance, the model below can be used in settings in which agents revise more frequently when the differences in strategies' payoffs are large than when these differences are small.

This revision process is defined in terms of a *raw choice function* $\tilde{\sigma}$, which is a map from excess payoff vectors $\pi \in \mathbf{R}^n_* = \mathbf{R}^n - \text{int}(\mathbf{R}^n_-)$ to nonnegative vectors $\tilde{\sigma}(\pi) \in \mathbf{R}^n_+$. We can leave $\tilde{\sigma}$ undefined on $\text{int}(\mathbf{R}^n_-)$ because an excess payoff vector cannot lie in this set: for this to occur, every strategy would need to earn a strictly below average payoff, which is clearly impossible. Note that $\text{int}(\mathbf{R}^n_*) = \mathbf{R}^n - \mathbf{R}^n_-$ is the set of excess payoff vectors under which at least one strategy has an above average payoff, while $\text{bd}(\mathbf{R}^n_*) = \text{bd}(\mathbf{R}^n_-)$ is the set of excess payoff vectors under which no strategy earns an above average payoff.

Given the raw choice function $\tilde{\sigma}$, revision rates and choice probabilities are determined as follows. When the excess payoff vector is π, each agent's revision opportunities arrive at a rate given by the *sum* of the components of $\tilde{\sigma}(\pi)$: that is, $\lambda(\pi) = \tilde{\sigma}_T(\pi) \equiv \sum_{j \in S} \tilde{\sigma}_j(\pi)$. After an agent receives a revision opportunity, he selects a strategy according to the *choice rule* σ: $\mathbf{R}^n_* \to \Delta$, the outputs of which are *proportional* to the raw choice vector:

$$\sigma(\pi) = \begin{cases} \dfrac{\tilde{\sigma}(\pi)}{\tilde{\sigma}_T(\pi)} & \text{if } \tilde{\sigma}_T(\pi) \neq 0; \\ \text{arbitrary} & \text{if } \tilde{\sigma}_T(\pi) = 0; \end{cases}$$

9 But see Björnerstedt and Weibull (1996) and Weibull (1995, Section 4.4).

Choice probabilities can be arbitrary when $\tilde{\sigma}_T(\pi) = 0$ since in this situation no revision opportunities arise.

We can obtain a direct interpretation of raw choice weights by describing the model in terms of the rates at which agents currently playing strategies other than i switch to strategy i. Using the (implicit) assumption that the arrivals of revision opportunities and the choices made thereafter are independent, we find that the rate of switching to strategy i from other strategies is $\lambda(\pi)\, \sigma_i(\pi) = \tilde{\sigma}_i(\pi)$. This formulation highlights a form of inertia built into our revision process: if for each $j \neq i$ the scalar $\tilde{\sigma}_j(\pi)$ is small, then agents playing strategy i rarely switch strategies.

To connect the agents' revision procedure with the underlying game, we impose two conditions on the raw choice function $\tilde{\sigma}$.

(C) $\tilde{\sigma}$ is Lipschitz continuous;

(A) $\tilde{\sigma}(\pi) \cdot \pi > 0$ whenever $\pi \in \text{int}(\mathbf{R}^n_*)$.

The first condition, *continuity*, asks that raw choice weights vary (Lipschitz) continuously with excess payoffs. Discontinuous raw choice functions exhibit an extreme sensitivity to the exact value of excess payoffs. In most applications, this level of sensitivity seems unrealistic, and so condition (C) precludes it.

The second condition, *acuteness*, requires that the excess payoff vector π and the raw choice weight $\tilde{\sigma}(\pi)$ form an acute angle whenever π lies in the interior of $\mathbf{R}^n_*$. Acuteness has distinct implications for revision rates and choice probabilities. For the former, condition (A) requires that whenever some strategy's excess payoff is strictly positive, the revision rate is strictly positive as well. Thus, acuteness implies a sort of persistence: as long as some agents would benefit from switching strategies, revision opportunities continue to arrive. Concerning choice probabilities, condition (A) requires that whenever some strategy achieves a strictly positive excess payoff, the expected value of a component of π chosen at random according to the probability distribution $\sigma(\pi)$ is strictly positive. Thus, *on average*, agents choose strategies with above average payoffs.

The simplest class of raw choice functions satisfying conditions (C) and (A) are the semilinear functions

(1) $$\tilde{\sigma}_i(\pi) = [\pi_i]_+.$$

Two increasingly general specifications are the truncated monomial forms

(2) $\tilde{\sigma}_i(\pi) = ([\pi_i]_+)^k, \; k \geq 1,$

and the separable forms

(3) $\tilde{\sigma}_i(\pi) = \phi_i(\pi_i),$ where ϕ_i: $\mathbf{R} \to \mathbf{R}_+$ is Lipschitz continuous, $\phi_i(\pi_i) = 0$ on $(-\infty, 0]$, and $\phi_i(\pi_i) > 0$ on $(0, \infty)$.

Separable raw choice functions only assign positive weights to strategies with positive excess payoffs. We now show that neither separability nor sign-preservation is implied by conditions (C) and (A). Consider the raw choice function

(4) $$\tilde{\sigma}_i(\pi) = \left((k+1)\sum_j \exp(c\pi_j)\right)([\pi_i]_+)^k + \left(c\sum_j ([\pi_j]_+)^{k+1}\right)\exp(c\pi_i).$$

Proposition 2.1: *Suppose that* $c > 0$, $k > 0$, *and* $(k + 1)\exp(k + 2) + 1 \geq n$. *Then the raw choice function* (4) *is nonseparable, generates strictly positive choice probabilities whenever* $\pi \in \text{int}(\mathbf{R}^n_*)$, *and satisfies conditions* (C) *and* (A).

Proof: In the Appendix.

The lower bound on the exponent k is quite weak: for example, we can let $k = 1$ as long as the number of pure strategies n does not exceed 41.

2.2 Evolutionary Dynamics

The evolutionary process defined above generates a Markov process on the simplex, with the realized sample path of this process depending on the realizations of each agent's revision opportunities and randomized choices. Using methods from the theory of convergence of Markov processes, Binmore and Samuelson (1999), Sandholm (2003a), and Benaïm and Weibull (2003) show that when the population size is large, the behavior of such processes is closely approximated by the solutions of a differential equation. This equation, the *mean dynamic* of the Markov process, is defined in terms of the *expected* changes in the population's behavior given the current population state.[10]

[10] More specifically, these papers show that during any finite time span, the actual behavior of the population stays within a narrow band surrounding the solution to the mean dynamic with high probability if the population size is sufficiently large.

To derive the mean dynamic for the present model, suppose that the current population state is x. Since there are N agents in the population, the expected number of agents receiving revision opportunities during the next dt time units is $N\ \lambda(\hat{F}(x))\ dt$. Since all agents are equally likely to receive revision opportunities, the expected number of opportunities received by agents currently choosing strategy i is $N\ \lambda(\hat{F}(x))\ x_i\ dt$. Finally, since choice probabilities are determined using the choice rule σ, the expected number of agents who receive opportunities and select strategy i is $N\ \lambda(\hat{F}(x))\ \sigma(\hat{F}(x))\ dt$. Therefore, the expected change in the number of agents choosing strategy i during the next dt time units is given by

$$N\ \lambda(\hat{F}(x))\ (\sigma_i(\hat{F}(x)) - x_i)\ dt.$$

The expected change in the *proportion* of agents choosing strategy i during the next dt time units is

$$\lambda(\hat{F}(x))\ (\sigma_i(\hat{F}(x)) - x_i)\ dt.$$

We therefore conclude that the mean dynamic for our Markov process is

$$\dot{x} = \lambda(\hat{F}(x))(\sigma(\hat{F}(x)) - x). \tag{5}$$

This dynamic has a simple interpretation: the population state always moves directly towards the "target state" defined by the current choice probability vector $\sigma(\hat{F}(x)) \in \Delta$, at a speed determined by the revision rate $\lambda(\hat{F}(x))$.

By substituting in the definitions of λ and σ, we can write this expression directly in terms of the raw choice function $\tilde{\sigma}$:

$$\dot{x} = \tilde{\sigma}(\hat{F}(x)) - \tilde{\sigma}_T(\hat{F}(x))x.$$

When $\tilde{\sigma}$ satisfies conditions (C) and (A), we call this differential equation an *excess payoff dynamic*.

2.3 Examples

2.3.1 The Brown-von Neumann-Nash Dynamic

If raw choice function take the truncated linear form (1), we obtain the excess payoff dynamic

$$\dot{x}_i = [\hat{F}_i(x)]_+ - \sum_{j \in S} [\hat{F}_j(x)]_+ \, x_i.$$

This equation is known as the *Brown-von Neumann-Nash* (*BNN*) *dynamic*. This dynamic was introduced in the context of symmetric zero-sum games by Brown and von Neumann (1950) (also see Nash (1951)), and then reintroduced by Skyrms (1990) and Swinkels (1992). For more recent treatments, see Weibull (1996), Hofbauer (2000), Berger and Hofbauer (2000), and Sandholm (2001).

We can use this dynamic to demonstrate the importance of allowing revision rates to vary. Had we fixed the revision rate fixed at 1, we would have obtained the mean dynamic

$$\dot{x}_i = \frac{[\hat{F}_i(x)]_+}{\sum_{j \in S} [\hat{F}_j(x)]_+} - x_i.$$

The initial term in this equation, representing current choice probabilities, is discontinuous: a small change in the state that causes a strategy's payoff to drop below average can force the probability with which the strategy is chosen to jump from 1 to 0.[11] It follows that the fixed rate dynamic is discontinuous as well. By allowing revision opportunities to arrive slowly when the benefits of switching strategies become small, we are able to ensure that our law of motion is Lipschitz continuous in the population state, thus ensuring the existence, uniqueness, and continuity of solution trajectories.

2.3.2 Connections with the Best Response Dynamic

The truncated monomial raw choice function (2) yields the choice rule

$$\sigma_i(\pi) = \frac{([\pi_i]_+)^k}{\sum_{j \in S} ([\pi_j]_+)^k}$$

whenever $\pi \in \text{int}(\mathbf{R}^n_*)$. If we let k approach infinity, then whenever the resulting limit exists it is described by the discontinuous choice rule

---

[11] For example, in a two strategy game, the choice probability for strategy 1 equals 1 if $F_1(x) > F_2(x)$, equals 0 if $F_1(x) < F_2(x)$, and is undefined otherwise. Hence, as long as neither strategy is dominant, a jump of the sort noted above must occur.

(7) $\sigma(\pi) = \operatorname*{argmax}_{y \in \Delta} y \cdot \pi$.

If we view equation (7) as a raw choice function, then the implied revision rate $\lambda(\pi) = \tilde{\sigma}_T(\pi)$ is fixed at one. Thus, since

$$\operatorname*{argmax}_{y \in \Delta} y \cdot \hat{F}(x) = \operatorname*{argmax}_{y \in \Delta} y \cdot Ax - x \cdot Ax = \operatorname*{argmax}_{y \in \Delta} y \cdot Ax \equiv B(x),$$

the resulting the mean dynamic is given by

$$\dot{x} \in B(x) - x.$$

This is the *best response dynamic* of Gilboa and Matsui (1991) and Matsui (1992).

Since the best response correspondence B is discontinuous, the best response dynamic possesses certain nonstandard properties. In particular, while solutions to this dynamic are certain to exist, they need not be unique; in certain cases, this multiplicity can be the source of quite complicated solution trajectories (Hofbauer (1995)). The discontinuity that is the source of these difficulties is a consequence of exact optimization. Under competent play, raw choice weights cannot depend too finely on payoff opportunities; this coarseness renders nonuniqueness of solution trajectories impossible.

2.4 Population Games

We conclude this section by introducing a more general class of games to which our analysis will apply. This new framework generalizes the symmetric random matching framework from Section 2.1 by allowing for multiple populations of agents (i.e., player roles) and by permitting payoffs to depend nonlinearly on the population state. While the games we define here are formally specified using continuous sets of players, one can interpret our results as providing approximate descriptions of the evolution of play in populations that are large but finite.

Let $\mathcal{P} = \{1, \dots, p\}$ denote the set of populations, where $p \geq 1$. Population masses are described by the vector $m = (m^1, \dots, m^p)$. The set of strategies for population p is denoted $S^p = \{1, \dots, n^p\}$, and $n = \sum_{p \in \mathcal{P}} n^p$ equals the total number of pure strategies.

The set of strategy distributions within population $p \in \mathcal{P}$ is denoted $X^p = \{x^p \in \mathbf{R}^{n^p}_+ : \sum_{i \in S^p} x^p_i = m^p\}$, while $X = \{x = (x^1, \dots, x^p) \in \mathbf{R}^n_+ : x^p \in X^p\}$ is the set of overall strategy distributions. Although behavior is always described by a point in X, it will be useful to define payoffs on the set $\overline{X} = \{x \in \mathbf{R}^n_+ \colon\ m^p - \varepsilon \le \sum_{i \in S^p} x_i^p \le m^p + \varepsilon \ \ \forall\ p \in \mathcal{P}\}$, where ε is a strictly positive constant. This set contains the strategy distributions that arise if there are slight changes in the populations' sizes. By defining payoffs on this set, we make it possible to speak directly about an agent's marginal impact on the payoffs of his opponents.

The payoff function for strategy $i \in S^p$ is denoted $F_i^p \colon \overline{X} \to \mathbf{R}$, and is assumed to be continuously differentiable. Observe that the payoffs to a strategy in population p can depend on the strategy distribution within population p itself. We let $F^p \colon \overline{X} \to \mathbf{R}^{n^p}$ refer to the vector of payoff functions for strategies belonging to population p and let $F \colon \overline{X} \to \mathbf{R}^n$ denote the vector of all payoff functions. Similar notational conventions are used throughout the paper. However, when we consider games with a single population, we assume that the population mass is one and omit the redundant superscript p.

The *average payoff* in population p is $\overline{F}^p(x) = \frac{1}{m^p}\, x^p \cdot F^p(x)$. Hence, the *excess payoff* to strategy $i \in S^p$ is $\hat{F}_i^p(x) = F_i^p(x) - \overline{F}^p(x)$, while $\hat{F}^p(x) = F^p(x) - \mathbf{1}\,\overline{F}^p(x)$ is the *excess payoff vector* for population p.

State $x \in X$ is a *Nash equilibrium* of F if all agents choose best responses to the current population state. Formally, x is a Nash equilibrium if

$$\text{For all } p \in \mathcal{P} \text{ and } i \in S^p,\ x_i^p > 0 \text{ implies that } i \in \operatorname{argmax}_{j \in S^p} F_j^p(x).$$

An evolutionary dynamic for a game F is a differential equation $\dot{x} = V(x)$ that describes the motion of the population through the set of population states X. The vector field V is a map from X to $TX = \{z \in \mathbf{R}^n \colon \sum_{i \in S^p} z_i^p = 0 \text{ for all } p \in \mathcal{P}\}$, the tangent space for the set X.

Suppose that agents in population p use a revision rate function λ^p and a choice rule σ^p derived from some raw choice function $\tilde{\sigma}^p$. The resulting mean dynamic is

$$\dot{x}^p = \lambda^p(\hat{F}^p(x))\Big(m^p \sigma^p(\hat{F}^p(x)) - x^p\Big) \ \text{ for all } p \in \mathcal{P},$$

Now let $\Delta^p = \{y^p \in \mathbf{R}^{n^p}_+ \colon\ \sum_{i \in S^p} x_i^p = 1\}$ denote the simplex in $\mathbf{R}^{n^p}$. Then under the dynamic above, the state variable for population p, $x^p \in X^p = m^p \Delta^p$, moves in the direction of the target state $m^p \sigma^p(r^p(x)) \in m^p \Delta^p$ at rate $\lambda^p(r^p(x))$. That is, the target state has the same relative weights as the probability vector $\sigma^p(r^p(x))$, but has a total mass of m^p.

We can once again rewrite our dynamic in terms of the raw choice functions $\tilde{\sigma}^p$:

(E) $$\dot{x}^p = m^p \tilde{\sigma}^p(\hat{F}^p(x)) - \tilde{\sigma}^p_T(\hat{F}^p(x))x^p \quad \text{for all } p \in \mathcal{P}.$$

Definition: *If the raw choice functions* $\tilde{\sigma}^p$ *satisfy conditions* (C) *and* (A), *we call equation* (E) *an* excess payoff dynamic.

3. Basic Properties of Excess Payoff Dynamics

We now define the three desiderata described informally in the introduction.

(EUC) $\dot{x} = V(x)$ admits a unique solution trajectory $\{x_t\}_{t\geq 0} = \{\phi_t(x)\}_{t\geq 0}$ from every initial condition $x \in X$, a trajectory that remains in X for all time. Moreover, for each $t \geq 0$, $\phi_t(x)$ is Lipschitz continuous in x.

(PC) For all $p \in \mathcal{P}$, $cov(V^p(x),\ F^p(x)) = \frac{1}{n^p}\big(V^p(x) \cdot F^p(x)\big) > 0$ whenever $V^p(x) \neq \mathbf{0}$.

(NS) $x \in X$ is a rest point if of V and only if it is a Nash equilibrium of F.

Condition (EUC) requires the *existence, uniqueness, and continuity* of solution trajectories. As we argued earlier, this condition ensures that predictions of behavior need are not overly sensitive to the exact value of the initial state, and it abrogates the analytical difficulties that discontinous dynamics present.

Condition (PC), *positive correlation*, requires that the growth rates and payoffs of strategies within each population be positively correlated, strictly so whenever the some growth rate is nonzero. To see that the equality stated in the condition is true under condition (EUC), note that the forward invariance of X implies that $\sum_{i\in S^p} V_i^p(x) = 0$, and hence that

$$\begin{aligned} cov(V^p(x),\ F^p(x)) &= \tfrac{1}{n^p}\sum_{i\in S^p}(V_i^p(x) - \tfrac{1}{n^p}\sum_{j\in S^p}V_j^p(x))(F_i^p(x) - \tfrac{1}{n^p}\sum_{j\in S^p}F_j^p(x)) \\ &= \tfrac{1}{n^p}\sum_{i\in S^p}(V_i^p(x) - 0)(F_i^p(x) - \tfrac{1}{n^p}\sum_{j\in S^p}F_j^p(x)) \\ &= \tfrac{1}{n^p}\Big(V^p(x)\cdot F^p(x) + \tfrac{1}{n^p}\sum_{i\in S^p}F_i^p(x)\sum_{i\in S^p}V_i^p(x)\Big) \\ &= \tfrac{1}{n^p}\Big(V^p(x)\cdot F^p(x)\Big). \end{aligned}$$

Conditions closely related to positive correlation have been proposed by Friedman (1991), Swinkels (1993), and Sandholm (2001). Requirements of this sort are the weakest used in the evolutionary literature, as they restrict each population's behavior using only a single scalar inequality.

Condition (NS), *Nash stationarity*, requires that the rest points of the dynamics and the Nash equilibria of the underlying game coincide. The condition captures the idea that there should be no impetus leading the population state to change in precisely those situations in which no agent can unilaterally improve his payoffs. It is worth noting that the replicator dynamic fails Nash stationarity; we discuss this fact in considerable detail in Section 4.

We now prove that all excess payoff dynamics satisfy these three desiderata.

Theorem 3.1: *Every excess payoff dynamic satisfies conditions* (EUC), (PC), *and* (NS).

Condition (EUC) is a direct consequence of the facts that excess payoff dynamics are Lipschitz continuous and point inward on the boundary of X (see Hirsch and Smale (1974, Chapter 8) and Ely and Sandholm (2004, Appendix I)). To establish the other two properties we prove three preliminary results.

Lemma 3.2: *Let $\dot{x} = V(x)$ be an excess payoff dynamic. Then for all $p \in \mathcal{P}$ and $x \in X$,*

(i) $x^p \cdot \hat{F}^p(x) = 0$;

(ii) *If $\hat{F}^p(x) \in \mathrm{int}(\mathbf{R}^{n^p}_*)$, then $V^p(x) \cdot F^p(x) > 0$.*

Part (i) of Lemma 3.2 observes that each population's state is always orthogonal to its excess payoff vector. Part (ii) shows that condition (PC) holds whenever some strategy earns an above average payoff.

Proof: (i) $x^p \cdot \hat{F}^p(x) = x^p \cdot (F^p(x) - \mathbf{1}\bar{F}^p(x)) = x^p \cdot F^p(x) - (x^p \cdot \mathbf{1})(\frac{1}{m^p} x^p \cdot F^p(x)) = 0.$

(ii) Suppose that $\hat{F}^p(x) \in \mathrm{int}(\mathbf{R}^{n^p}_*)$. Then the fact that $V(x) \in TX$, part (i) of the lemma, and acuteness imply that

$$\begin{aligned} V^p(x) \cdot F^p(x) &= V^p(x) \cdot (\hat{F}^p(x) + \mathbf{1}\bar{F}^p(x)) \\ &= (m^p \tilde{\sigma}^p(\hat{F}^p(x)) - x^p\, \tilde{\sigma}^p_T \hat{F}^p(x))) \cdot \hat{F}^p(x) \\ &= m^p \tilde{\sigma}^p(\hat{F}^p(x)) \cdot \hat{F}^p(x) - \tilde{\sigma}^p_T(\hat{F}^p(x)) x^p \cdot \hat{F}^p(x) \\ &= m^p \tilde{\sigma}^p(\hat{F}^p(x)) \cdot \hat{F}^p(x) > 0. \ \blacksquare \end{aligned}$$

The next lemma uses acuteness and continuity to restrict the action of raw choice functions on the boundary of $\mathbf{R}^{n^p}_*$: strategies whose payoffs are below average must receive zero weight, and a strategy whose payoff is exactly average can receive positive weight only if it is the only such action.

Lemma 3.3: *Let* $\tilde{\sigma}^p$ *satisfy properties* (C) *and* (A), *and let* $\pi^p \in \text{bd}(\mathbf{R}_*^{n^p})$, *so that the set of strategies earning average payoffs,* $Z^p(\pi^p) = \{i \in S^p\text{: } \pi_i^p = 0\}$, *is nonempty. Then*

(*i*) *If* $i \notin Z^p(\pi^p)$ (*i.e., if* $\pi_i^p < 0$), *then* $\tilde{\sigma}_i^p(\pi^p) = 0$;

(*ii*) *If* $Z^p(\pi^p) = \{j\}$, *then* $\tilde{\sigma}^p(\pi^p) = c\,e_j^p$ *for some* $c \geq 0$;

(*iii*) *If* $\#Z^p(\pi^p) \geq 2$, *then* $\tilde{\sigma}^p(\pi^p) = \mathbf{0}$.

Proof: For notational convenience, we only consider the case in which $p = 1$; the proof of the general case is an easy extension.

(*i*) Suppose that $\pi \in \text{bd}(\mathbf{R}_*^n)$, $i \notin Z(\pi)$, and $j \in Z(\pi)$. For $\varepsilon > 0$, let $\pi(\varepsilon) = \pi + \varepsilon e_j \in \text{int}(\mathbf{R}_*^n)$ (see Figure 1(*i*)). Then if $k \neq j$,

$$\tilde{\sigma}_k(\pi(\varepsilon))\,\pi_k(\varepsilon) = \tilde{\sigma}_k(\pi(\varepsilon))\,\pi_k \leq 0.$$

Moreover,

$$\lim_{\varepsilon \to 0} \tilde{\sigma}_j(\pi(\varepsilon))\,\pi_j(\varepsilon) = \lim_{\varepsilon \to 0} \tilde{\sigma}_j(\pi(\varepsilon))\,\varepsilon = 0.$$

Now were $\tilde{\sigma}_i(\pi)$ strictly greater than zero, it would follow from continuity that

$$\lim_{\varepsilon \to 0} \tilde{\sigma}_i(\pi(\varepsilon))\,\pi_i(\varepsilon) = \tilde{\sigma}_i(\pi)\pi_i < 0.$$

The last three expressions would then imply that $\tilde{\sigma}(\pi(\varepsilon)) \cdot \pi(\varepsilon) < 0$ for all sufficiently small ε, contradicting acuteness. Therefore, $\tilde{\sigma}_i(\pi) = 0$.

(*ii*) Follows immediately from part (*i*).

(*iii*) Suppose that $\pi \in \text{bd}(\mathbf{R}_*^n)$. If $i \notin Z(\pi)$, then $\tilde{\sigma}_i(\pi) = 0$ by part (*i*). So let i, $j \in Z(\pi)$, and suppose that $\tilde{\sigma}_i(\pi) > 0$.

Define $\pi(\varepsilon) = \pi - \varepsilon e_j + \varepsilon^2 e_j \in \text{int}(\mathbf{R}_*^n)$ (see Figure 1(*ii*)). If $k \notin \{i, j\}$, then

$$\tilde{\sigma}_k(\pi(\varepsilon))\,\pi_k(\varepsilon) = \tilde{\sigma}_k(\pi(\varepsilon))\,\pi_k \leq 0.$$

Thus,

$$\begin{aligned}\tilde{\sigma}(\pi(\varepsilon)) \cdot \pi(\varepsilon) &\leq \tilde{\sigma}_i(\pi(\varepsilon))\,\pi_i(\varepsilon) + \tilde{\sigma}_j(\pi(\varepsilon))\,\pi_j(\varepsilon) \\ &= -\varepsilon\,\tilde{\sigma}_i(\pi(\varepsilon)) + \varepsilon^2\,\tilde{\sigma}_j(\pi(\varepsilon)) \\ &= \varepsilon(-\tilde{\sigma}_i(\pi(\varepsilon)) + \varepsilon\,\tilde{\sigma}_j(\pi(\varepsilon))),\end{aligned}$$

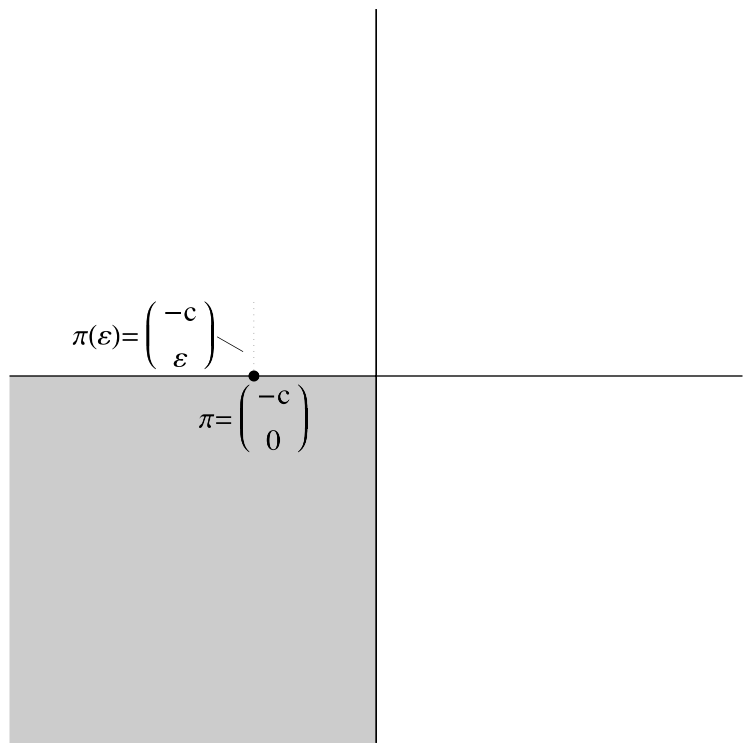

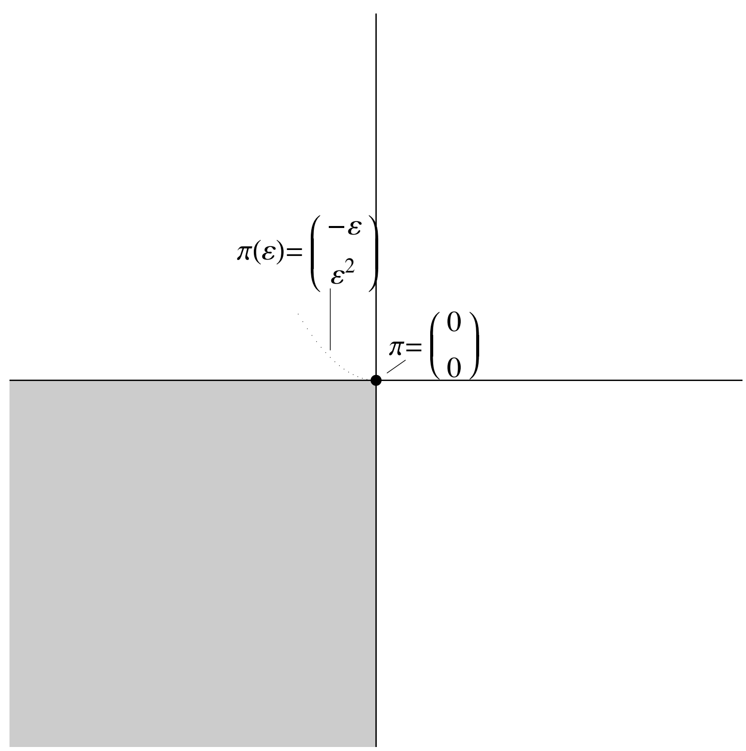

Figures 1(*i*) and 1(*ii*): Sequences of vectors that approach $\mathrm{bd}(\mathbf{R}^n_*)$

which by continuity must be strictly negative once ε small. This contradicts acuteness. We therefore conclude that $\tilde{\sigma}_i(\pi) = 0$. ■

The next proposition provides two alternate characterizations of states x at which the excess payoff vector $\hat{F}^p(x)$ lies on the boundary of $\mathbf{R}^{n^p}_*$. This result and the previous two immediately imply properties (PC) and (NS).

Proposition 3.4: *Let $\dot{x} = V(x)$ be an excess payoff dynamic, and fix $x \in X$ and $p \in \mathcal{P}$. Then the following are equivalent*:

(*i*) *For all $i \in S^p$, $x_i^p > 0$ implies that $i \in \arg\max_{j \in S^p} F_j^p(x)$;*

(*ii*) $\hat{F}^p(x) \in \text{bd}(\mathbf{R}^{n^p}_*)$;

(*iii*) $V^p(x) = \mathbf{0}$.

Proof: We first prove that (*i*) implies (*ii*). If condition (*i*) holds, then all strategies in the support of x^p yield the maximal payoff, which is therefore the population's average payoff: $\max_j F_j^p(x) = \bar{F}^p(x)$. It follows that $\hat{F}_i^p(x) = F_j^p(x) - \bar{F}^p(x) \le 0$ for all $i \in S^p$, with equality whenever $x_i^p > 0$. Hence, $\hat{F}^p(x) \in \text{bd}(\mathbf{R}^{n^p}_*)$.

Second, we show that (*ii*) implies (*i*). Suppose that $\hat{F}^p(x) \in \text{bd}(\mathbf{R}^{n^p}_*)$, and let i be a strategy in the support of x^p. If $\hat{F}_i^p(x) < 0$, then Lemma 3.2(*i*) implies that $\hat{F}_j^p(x) > 0$ for some action $j \in S^p$, contradicting the definition of $\hat{F}^p(x)$. Thus, $\hat{F}_i^p(x) = 0 = \max_{j \in S^p} \hat{F}_j^p(x)$. Since a strategy maximizes excess payoffs if and only if it also maximizes actual payoffs, we conclude that $i \in \arg\max_{j \in S^p} F_j^p(x)$.

Third, we prove that (*ii*) implies (*iii*). Let $\hat{F}^p(x) \in \text{bd}(\mathbf{R}^{n^p}_*)$, so that $Z^p(\hat{F}^p(x)) = \arg\max_{j \in S^p} \hat{F}_j^p(x) = \arg\max_{j \in S^p} F_j^p(x)$. We divide the analysis into two cases.

For the first case, suppose that $Z^p(\hat{F}^p(x)) = \{i\}$. Then since strategy i is the sole optimal strategy, statement (*i*) implies that $x_k^p = 0$ for all $k \neq i$, and so $x^p = m^p e_i^p$. Now Lemma 3.3(*ii*) tells us that $\tilde{\sigma}^p(\hat{F}^p(x)) = c\, e_i^p$ for some $c \ge 0$. Hence,

$$\begin{aligned} V^p(x) &= m^p \tilde{\sigma}_i^p(\hat{F}^p(x)) - x_i^p \tilde{\sigma}_T^p(\hat{F}^p(x)) \\ &= m^p(c\, e_i^p) - (m^p e_i^p)c = \mathbf{0}, \end{aligned}$$

which is statement (*iii*).

For the second case, suppose that $Z^p(\hat{F}^p(x)) \ge 2$. Then Lemma 3.3(*iii*) implies that $\tilde{\sigma}^p(\hat{F}^p(x)) = \mathbf{0}$, which immediately implies that $V^p(x) = \mathbf{0}$.

Fourth, we establish that (*iii*) implies (*ii*) by proving the contrapositive. Suppose that $\hat{F}^p(x) \in \text{int}(\mathbf{R}^{n^p}_*)$. Then Lemma 3.2(*ii*) implies that $V^p(x) \cdot F^p(x) > 0$, and hence that $V^p(x) \neq \mathbf{0}$. This completes the proof of the proposition. ■

With our preliminary results in hand we prove Theorem 3.1. Lemma 3.2(*ii*) shows that condition (PC) holds whenever $\hat{F}^p(x) \in \text{int}(\mathbf{R}^{n^p}_*)$, and Proposition 3.4 shows that condition (PC) holds when $\hat{F}^p(x) \in \text{bd}(\mathbf{R}^{n^p}_*)$, since it tells us that $V^p(x) = \mathbf{0}$ in this case. Furthermore, if the conditions in Proposition 3.4 are imposed on all populations at once, then statement (*i*) says that x is a Nash equilibrium, while statement (*iii*) says that x is a rest point of V. Since Proposition 4 tells us that these statements are equivalent, condition (NS) holds. This completes the proof of the theorem.

4. Well Behaved Approximations of Imitative Dynamics

The best known evolutionary dynamic is the *replicator dynamic*, defined by

$$\dot{x}^p_i = x^p_i \hat{F}^p_i(x).$$

This dynamic was introduced by Taylor and Jonker (1978) as a biological model of competition between species. More recently, Björnerstedt and Weibull (1996) and Schlag (1998) have shown that the replicator dynamic can be used to describe the behavior of agents who use decision procedures based on imitation, justifying the application of this dynamic in economic models.[12]

By allowing more general classes of imitative decision procedures, one obtains the class of *imitative dynamics*. These are smooth dynamics on X of the form

$$\dot{x}^p_i = I^p_i(x) = x^p_i g^p_i(x)$$

that exhibit *monotone percentage growth rates*:[13]

---

[12] Choice rules that generate the replicator dynamic must allow choice probabilities to depend not only on current payoffs, but also on the revising agent's current strategy; however, these more complicated choice rules can be paired with a constant revision rate.

[13] This property has appeared in the literature under a variety of names: *relative monotonicity* (Nachbar (1990)), *order compatibility of predynamics* (Friedman (1991)), *monotonicity* (Samuelson and Zhang (1992)), and *payoff monotonicity* (Weibull (1995)).

$$g_i^p(x) \geq g_j^p(x) \text{ if and only if } F_i^p(x) \geq F_j^p(x).$$

Since imitative dynamics are smooth, they admit unique solution trajectories from every initial condition. It is not difficult to show that these dynamics satisfy positive correlation as well.[14] But it is well known that imitative dynamics fail Nash stationarity: while every Nash equilibrium is a rest point of I, not all rest points of I are Nash equilibria. In fact, x is a rest point if and only if it is a *restricted equilibrium* of the underlying game: that is, if for each $p \in \mathcal{P}$, every strategy *in the support of* x^p achieves the same payoff. Thus, the extra rest points of imitative dynamics all lie on the boundary of the state space X. The reason for these extra rest points is clear: whenever all agents choose the same strategy, imitation accomplishes nothing. While such behavior is plausible in some economic contexts, in others it is more natural to expect that a successful strategy will eventually be played even if it is currently unused.

For this reason, it is common to introduce perturbed versions of the imitative dynamics under which the boundary of the state space is repelling. A typical formulation is the perturbed dynamic

$$\dot{x}^p = (1 - \alpha)\, I^p(x) + \alpha\, (m^p \bar{\sigma}^p - x^p),$$

where $\bar{\sigma}^p \in \text{int}(\Delta^{n^p})$ is some completely mixed strategy and α is a small positive constant. One interpretation of this dynamic is that each agent's revision opportunities are driven by two independent Poisson alarm clocks. Rings of the first clock lead to an application of an imitative choice rule of the kind mentioned above, while rings of the second clock, which arrive at a much slower rate, lead to a randomized choice according to mixed strategy $\bar{\sigma}^p$. This perturbation of the dynamic eliminates all rest points that are not Nash equilibria. Still, the assumption about behavior on which it is based seems rather ad hoc. It also has some negative consequences: under the perturbed dynamic, growth rates and payoffs are negatively correlated near the boundary of X and near the rest points that survive the perturbation; moreover, these surviving rest points need only approximate Nash equilibria.

The analysis in Section 3 leads us to consider a different modification of I. Let V be an excess payoff dynamic, and define a new dynamic C_α by

---

14 See Fudenberg and Levine (1998, Proposition 3.6) or Sandholm (2002, Lemma A3).

$$\dot{x} = C_\alpha(x) \equiv (1-\alpha)\, I(x) + \alpha\, V(x),$$

As before, one can interpret this dynamic in terms of pairs of Poisson alarm clocks; this time, the second alarm clock rings at a variable rate $\lambda(\cdot)$, and leads to the use of a choice rule $\sigma(\cdot)$ as defined above. Put differently, the dynamic C_α captures the behavior of agents whose decisions are usually based on imitation, but are occasionally based on efforts to choose a strategy that performs relatively well, regardless of its current level of use. Given the foregoing analysis, it is easy to show that this modification eliminates non-Nash rest points of imitative dynamics, but without disturbing these dynamics' other desirable properties.

Theorem 4.1: *The dynamic C_α satisfies* (EUC), (PC), *and* (NS) *for all* $\alpha \in (0, 1]$.

Proof: In the Appendix.

An intuition for this result is as follows. Out of our three desiderata for evolutionary dynamics, imitative dynamics only fail condition (NS), and then only on the boundary of the state space. It is therefore quite easy to introduce modifications of these dynamics that eliminate this failure, but typically at the cost of introducing other failures. Excess payoff dynamics are desirable modifications because they themselves satisfy (EUC), (PC) and (NS). For this reason, we are able to recover condition (NS) while preserving our other desiderata.

5. Potential Games

Our results in the following sections use the properties established above to prove global convergence to Nash equilibrium under competent play in two classes of games. In both cases, integrability plays a central role. In the this section, an integrability condition characterizes the class of games under study. In the subsequent sections, imposing an integrability condition directly on choice rules will be the key to our analysis.

Potential games are known to have appealing convergence properties. Games from this class were first used in studies of congestion (Beckmann, McGuire, and Winsten (1956), Rosenthal (1973)) and in population genetics models (Hofbauer and Sigmund (1988)). Monderer and Shapley (1996) provide a general definition of these games in a finite player context; the continuum of player version that we consider here is studied in Sandholm (2001). Recently, potential games have found application in evolutionary approaches to externality pricing and implementation —see Sandholm (2002, 2003b).

In the present context, potential games are defined by a condition on payoff derivatives. We say that the game F is a *potential game* if it satisfies

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

We call this condition *externality symmetry*. In words, this condition requires that the marginal impact of an agent who chooses strategy $j \in S^q$ on opponents choosing strategy $i \in S^p$ is always equal to the marginal impact of an agent who chooses strategy i on opponents who choose strategy j. If we let DF: $\mathbf{R}^n \to \mathbf{R}^{n\times n}$ denote the derivative matrix of the vector field F, then condition (ES) can be expressed succinctly as

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

Mathematically, condition (ES) is an integrability condition for the vector field F. It implies the existence of a function f: $\overline{X} \to \mathbf{R}$ satisfying

$$\frac{\partial f}{\partial x_i^p}(x) = F_i^p(x) \text{ for all } i \in S^p,\ p \in \mathcal{P}, \text{ and } x \in X.$$

The function f is called a *potential function* of the game F.

We now show that in potential games, all solution trajectories of excess payoff dynamics converge to Nash equilibria. To accomplish this, we show that the potential function f serves as a Lyapunov function for all such dynamics. Call the function L: $X \to \mathbf{R}$ a *strict Lyapunov function* for the dynamics $\dot{x} = V(x)$ if $\frac{d}{dt}L(x_t) \geq 0$ along every solution trajectory, with equality only at rest points of the dynamic.

Theorem 5.1: *Let V be an excess payoff dynamic for the potential game F. Then the potential function f is a strict Lyapunov function for V, and each solution to V converges to a connected set of Nash equilibria.*

Proof: Let $\dot{x} = V(x)$ be an excess payoff dynamic. Then Theorems 3.1 tells us that V satisfies conditions (PC) and (NS). The definition of potential and condition (PC) imply that

$$\tfrac{d}{dt} f(x_t) = \nabla f(x_t) \cdot \dot{x}_t = F(x_t) \cdot V(x_t) = \sum_p F^p(x_t) \cdot V^p(x_t) \geq 0.$$

Equality only holds if $V^p(x_t) = \mathbf{0}$ for all $p \in \mathcal{P}$, in which case x_t is a rest point of V. Thus, f is a strict Lyapunov function for V. Since X is compact, standard results (e.g., Theorem 7.6 of Hofbauer and Sigmund (1988)) imply that every solution trajectory of the dynamic must converge to a connected set of rest points. Condition (NS) tells us that such sets consist solely of Nash equilibria. ■

6. Stable Games

6.1 Definition

We now introduce a new class of games that are also defined by a condition on payoff derivatives.

Definition: *We call* F: $\overline{X} \to \mathbf{R}^n$ *a* stable game *if*

(SE) $$\sum_{p \in \mathcal{P}} \sum_{i \in S^p} z_i^p \, \frac{\partial F_i^p}{\partial z}(x) \leq 0 \quad \textit{for all } z \in TX \textit{ and all } x \in X.$$

We call condition (SE) *self-defeating externalities*. To interpret this condition, first observe that vectors $z \in TX$ represent directions of motion through the state space X. We can view such vectors as describing the aggregate effect on the population state of strategy revisions by some small group of agents. The derivative $\frac{\partial F_i^p}{\partial z}(x)$ represents the marginal effect that these revisions have on the payoffs of agents currently choosing strategy $i \in S^p$. Condition (SE) considers a weighted sum of these effects, with weights given by the changes in the use of each strategy. It requires that this weighted sum be negative.

Intuitively, a game exhibits self-defeating externalities if the improvements in the payoffs of strategies to which revising agents are switching are always exceeded by the improvements in the payoffs of strategies which revising agents are abandoning. For example, suppose the tangent vector z takes the form $z = e_j^p - e_i^p$. This vector represents switches by some members of population p from strategy i to strategy j. In this case, the requirement in condition (SE) reduces to $\frac{\partial F_j^p}{\partial z}(x) \leq \frac{\partial F_i^p}{\partial z}(x)$: in words, any performance gains that the switches create for the newly chosen strategy j are dominated by the performance gains created for the abandoned strategy i.[15]

We can also express condition (SE) in a more concise form. Since the derivative $\frac{\partial F_i^p}{\partial z}(x)$ equals $\nabla F_i^p(x) \cdot z$ by definition, we find that condition (SE) is equivalent to the requirement that

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

In other words, F is a stable game if for all population states x, the derivative matrix $DF(x)$ is negative semidefinite with respect to all tangent directions z.

6.2 Examples

Stable games subsume a number of interesting classes of games as special cases. The first two examples consider single population random matching games with payoff matrix A, so that $F(x) = Ax$ and $DF(x) = A$.

6.2.1 *Games with an interior ESS*. The state $x^* \in X$ is an *ESS* if $x^*{\cdot}Ax^* > x{\cdot}Ax^*$ for all $x \neq x^*$ in a neighborhood of x^*. It is well known (see, e.g., Hofbauer and Sigmund (1988, p. 122)) that if the game A admits an ESS in int(X), then $z{\cdot}Az < 0$ for all vectors $z \in \{v \in \mathbf{R}^n\colon\ \mathbf{1} \cdot v = 0\} = TX$. Hence, F is a stable game.

6.2.2 *Symmetric zero-sum games*. The symmetric game A is zero sum if A is skew-symmetric: that is, $A_{ij} = -A_{ji}$ for all $i, j \in S$. In this case, $z{\cdot}Az = 0$ for all vectors z, so F is a stable game.

6.2.3 *Asymmetric zero-sum games*. Consider a two player normal form game with bimatrix (A, B) played by two populations of unit mass. This defines the population game

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

---

[15] For a concrete example, consider a single population, two strategy, normal form game with payoff matrix $A \in \mathbf{R}^{2\times 2}$. Suppose that the off-diagonal elements of A are strictly positive, and while the diagonal elements of A are zero. In this two strategy case, all vectors in TX are proportional to either $e_1 - e_2$ or $e_2 - e_1$. If agents switch from strategy i to strategy j (i.e., if $z = e_j - e_i$), then the payoffs to strategy j fall at rate $\partial F_j / \partial z \equiv -A_{ji} < 0$, while the payoffs to strategy i rise at rate $\partial F_i / \partial z \equiv A_{ij} > 0$. Thus, the matrix A defines a stable population game.

The game (A, B) is zero-sum if $A = -B$. Then if $z = (z^1, z^2)$ is a vector in $\mathbf{R}^n = \mathbf{R}^{n^1+n^2}$, we find that $z \cdot DF(x)\, z = z^1 \cdot A\, z^2 + z^2 \cdot B'\, z^1 = -z^1 \cdot B\, z^2 + z^2 \cdot B'\, z^1 = 0$. Thus, F is a stable game.

6.2.4 *RL stable games*. Cressman, Garay, and Hofbauer (2001) study stability conditions for a model of random matches between members of p distinct species. Payoffs in this model are linear, and can be described by a p x p grid of matrices describing payoffs in all possible matches. If this grid forms a negative definite matrix, they call the corresponding game *RL* (*Replicator-Lyapunov*) *stable*. Cressman, Garay, and Hofbauer (2001) show that RL stable games admit a unique Nash equilibrium, and this equilibrium is globally stable with respect to interior initial conditions under the replicator dynamic. If we describe this model in our notation, then the derivative DF is identically equal to the p x p grid of matrices. Since this grid is negative definite, the game F is stable.[16]

6.2.5 *Concave potential games*. Potential games with concave potential functions arise in models of congestion in which congestion is a "bad" (e.g., models of highway congestion) and in applications of evolutionary techniques to implementation problems. Interestingly, all concave potential games are stable games. Suppose that F is a potential game whose potential function f is concave on X. Then since $F \equiv \nabla f$ by definition, we find that $z \cdot DF(x) z = z \cdot D^2 f(x) z \le 0$ for all $z \in TX$, and so F is a stable game.

6.2.6 *Negative dominant diagonal games*. We call F a *negative dominant diagonal game* if it satisfies the following three conditions for all $x \in X$, $i \in S^p$, and $p \in \mathcal{P}$:

(N1) $$\frac{\partial F_i^p}{\partial x_i^p}(x) \le 0;$$

(N2) $$\left|\frac{\partial F_i^p}{\partial x_i^p}(x)\right| \ge \sum_{(j,q)\neq(i,p)} \left|\frac{\partial F_j^q}{\partial x_i^p}(x)\right|;$$

(N3) $$\left|\frac{\partial F_i^p}{\partial x_i^p}(x)\right| \ge \sum_{(j,q)\neq(i,p)} \left|\frac{\partial F_i^p}{\partial x_j^q}(x)\right|.$$

16 Actually, RL stability only requires negative definiteness after a positive reweighting of the rows of the grid of payoff matrices. Our analysis can be extended to allow such reweightings as well.

Condition (N1) says that choosing strategy $i \in S^p$ imposes a negative externality on other users of this strategy; condition (N2) states that this externality exceeds the total externalities that strategy i imposes on other strategies, and condition (N3) states that this externality exceeds the aggregate externalities that other strategies impose on strategy i. If these three conditions hold, then the symmetric matrix $DF(x) + DF(x)'$ is diagonal dominant with weakly negative eigenvalues. This implies that $DF(x) + DF(x)'$ is negative semidefinite, and hence that $DF(x)$ is negative semidefinite. Therefore, F is a stable game.[17]

7. Cycling under Excess Payoff Dynamics in Stable Games

Many of the examples described above are known to have appealing evolutionary stability properties. One might therefore hope that excess payoff dynamics would globally converge to Nash equilibrium in all stable games. We now demonstrate that this is not the case.

Example 7.1: Consider the basic Rock-Scissors-Paper game, in which the winner of a match obtains a payoff of 1, the loser of a match obtains a payoff of –1, and in which draws yield 0 for both players. When a single population is randomly matched to play this game, the resulting payoff vector field is

$$F(x) = \begin{pmatrix} F_R(x) \\ F_S(x) \\ F_P(x) \end{pmatrix} = \begin{pmatrix} 0x_R + 1x_S - 1x_P \\ -1x_R + 0x_S + 1x_P \\ 1x_R - 1x_S + 0x_P \end{pmatrix}.$$

The unique Nash equilibrium of this game is $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.[18]

For each $\varepsilon > 0$, let $g^\varepsilon : \mathbf{R} \to \mathbf{R}$ be a continuous decreasing function that equals 1 on $(-\infty, 0]$, equals ε^2 on $[\varepsilon, \infty)$, and is linear on $[0, \varepsilon]$. Then consider the raw choice function

[17] Other classes of stable games can be obtained by considering other sufficient conditions for negative semidefiniteness of square matrices. For examples of such conditions, see Horn and Johnson (1985, Ch. 6).
[18] For analyses of the replicator dynamic and the best response dynamic for Rock-Scissors-Paper games, see Gaunersdorfer and Hofbauer (1995).

$$\tilde{\sigma}(\pi) = \begin{pmatrix} \tilde{\sigma}_R(\pi) \\ \tilde{\sigma}_S(\pi) \\ \tilde{\sigma}_P(\pi) \end{pmatrix} = \begin{pmatrix} [\pi_R]_+ g^\varepsilon(\pi_S) \\ [\pi_S]_+ g^\varepsilon(\pi_P) \\ [\pi_P]_+ g^\varepsilon(\pi_R) \end{pmatrix}.$$

Under $\tilde{\sigma}$, the weight placed on a strategy is proportional to positive part of the strategy's own excess payoff, but this weight is only of order ε^2 if the strategy it beats in Rock-Scissors-Paper has an excess payoff greater than ε. This raw choice function clearly satisfies the continuity condition (C). It also satisfies the acuteness condition (A), since

$$\tilde{\sigma}(\pi) \cdot \pi = ([\pi_R]_+)^2 g^\varepsilon(\pi_S) + ([\pi_S]_+)^2 g^\varepsilon(\pi_P) + ([\pi_P]_+)^2 g^\varepsilon(\pi_R),$$

which is strictly positive on $\text{int}(\mathbf{R}^n_*)$ and equals zero on $\text{bd}(\mathbf{R}^n_*)$.

Let V be the excess payoff dynamic defined by the game F and the raw choice function $\tilde{\sigma}$. Fix $\delta > 0$, and let $B^\delta(x^*)$ be a ball of radius δ around the equilibrium x^*. We then have

Proposition 7.2: (*i*) *When* $\varepsilon < .1094$, *there are initial conditions from which solutions to* V *converge to periodic orbits.*

(*ii*) *When* ε *is sufficiently small, solutions to* V *from all initial conditions outside of* $B^\delta(x^*)$ *converge to periodic orbits.*

The intuition behind this example can be explained as follows. In Figure 2, Scissors earns a positive payoff as soon as the trajectory from x^0 crosses segment ax^*, and becomes the sole strategy that does so once the segment e_Px^* is reached. However, the choice rule above puts very little probability on Scissors until Paper, the strategy it defeats, yields a payoff close to zero. As a result, the solution trajectory heads almost directly towards state e_P until Scissors becomes the sole strategy earning a payoff of ε. This extends the phase during which the solution approaches the vertex e_P before turning towards e_S, and thereby generates cycling.

Proof: Consider the trajectory that starts from some initial state $x^0 = (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2})$ that lies on segment e_Rx^* and satisfies $\alpha > \underline{\alpha} = \frac{1+\varepsilon}{3-3\varepsilon}$ (see Figure 2). This trajectory travels clockwise around the simplex. Our main task is to obtain an lower bound on the distance of this solution from state x^* when the solution crosses segment e_Px^*. Doing so enables us to bound the action of the Poincaré map of the dynamic on e_Rx^*, which in turn lets us use the Poincaré-Bendixson Theorem to demonstrate the existence of a periodic orbit.

To begin, note that F is derived from a symmetric zero-sum normal form game, which implies that the population's average payoff is always zero. It follows that a strategy's excess payoff is positive if and only if its actual payoff is positive.

When the current state lies in the triangle with vertices e_R, x^*, and $a = (0, \frac{1}{2}, \frac{1}{2})$, as it does at x^0, only strategy P has a positive payoff, so the target state under dynamic V is $\sigma(\hat{F}(x)) = e_P$. Therefore, the trajectory from x^0 leaves triangle $e_R x^* a$ at state $x^1 = (\frac{2\alpha}{1+3\alpha}, \frac{1-\alpha}{1+3\alpha}, \frac{2\alpha}{1+3\alpha})$. Since $\alpha > \underline{\alpha} = \frac{1+\varepsilon}{3-3\varepsilon}$, x^1 lies on the interior of segment az, where $z = (\frac{1+\varepsilon}{3}, \frac{1-2\varepsilon}{3}, \frac{1+\varepsilon}{3})$. For future reference, we observe that z is the intersection of segments ax^* and bc, where $b = (\frac{1+\varepsilon}{2}, \frac{1-\varepsilon}{2}, 0)$ and $c = (\varepsilon, 0, 1 - \varepsilon)$.

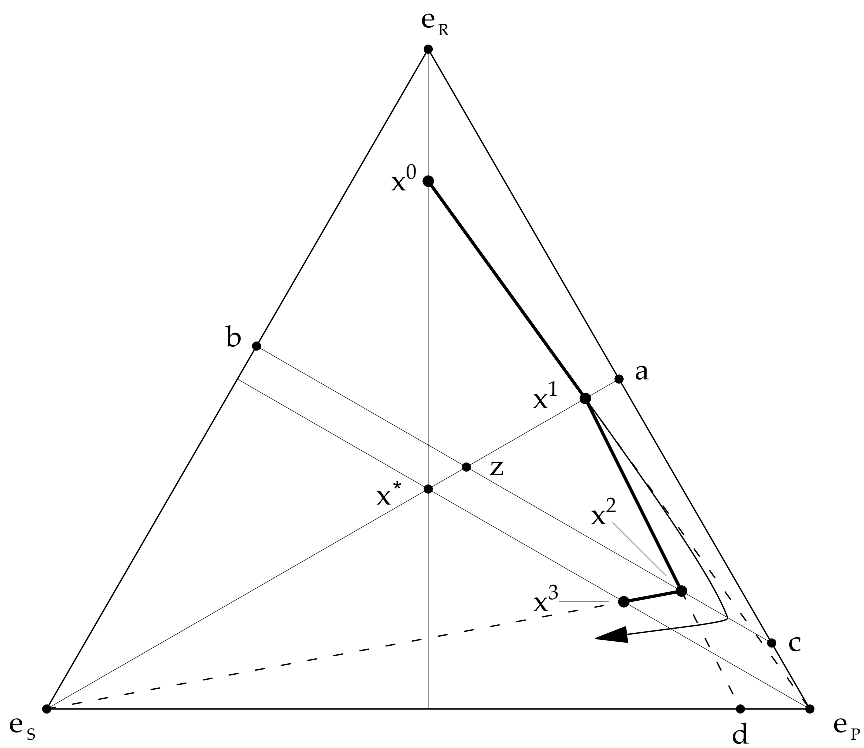

Figure 2

In triangle $e_P x^* a$, only strategies P and S earn positive payoffs. By construction, $\tilde{\sigma}_S(\hat{F}(x)) = \varepsilon^2[\hat{F}_S(x)]_+$ as long as the payoff to P is at least ε, which is the case in triangle $e_R bc$. The intersection of these two triangles is the triangle azc. When the current state x is in this region, the target state is always a point $(0, \sigma_P(\hat{F}(x)), \sigma_R(\hat{F}(x)))$ at which

$$\begin{aligned}\sigma_S(\hat{F}(x)) &= \frac{\tilde{\sigma}_S(\hat{F}(x))}{\tilde{\sigma}_S(\hat{F}(x)) + \tilde{\sigma}_P(\hat{F}(x))} \\ &= \frac{[\hat{F}_S(x)]_+ g^\varepsilon(\hat{F}_P(x))}{[\hat{F}_S(x)]_+ g^\varepsilon(\hat{F}_P(x)) + [\hat{F}_P(x)]_+ g^\varepsilon(\hat{F}_R(x))} \\ &\leq \frac{1\times\varepsilon^2}{(1\times\varepsilon^2)+(\varepsilon\times 1)} \\ &= \frac{\varepsilon}{\varepsilon+1}.\end{aligned}$$

Now the ray from point x^1 through point $d = (0, \frac{\varepsilon}{1+\varepsilon}, \frac{1}{1+\varepsilon})$ intersects segment bc at $x^2 = (\frac{2\alpha\varepsilon(2+\varepsilon)}{3\alpha(1+2\varepsilon)-1}, \frac{\varepsilon(1+\alpha-4\alpha\varepsilon)}{3\alpha(1+2\varepsilon)-1}, \frac{\alpha(3+\varepsilon+2\varepsilon^2)-\varepsilon-1}{3\alpha(1+2\varepsilon)-1})$. Hence, the inequality above implies that the solution trajectory from x^1 (and hence the one from x^0) hits segment zc at a point between x^2 and c.

Finally, consider the behavior of solution trajectories passing through the polygon ce_Px^*z. In this region, the target point is always on segment e_Se_P. In fact, once the solution hits segment e_Px^*, strategy S becomes the sole strategy earning a positive payoff, so the target point must be e_S. Thus, the solution starting from x^2 must hit e_Px^* no closer to x^* than $x^3 = (\frac{2\alpha\varepsilon(2+\varepsilon)}{(1+\varepsilon)3\alpha(1+2\varepsilon)-1}, \frac{2\alpha\varepsilon(2+\varepsilon)}{(1+\varepsilon)3\alpha(1+2\varepsilon)-1}, \frac{\alpha(3+\varepsilon+2\varepsilon^2)-\varepsilon-1}{(1+\varepsilon)3\alpha(1+2\varepsilon)-1})$, the point where a ray from x^2 through e_S crosses segment e_Px^*. Since the solution starting from x^0 hits segment zc to the right of x^2, it too must hit e_Px^* to the right of x^3. We have thus established a lower bound of $\beta(\alpha) = \frac{\alpha(3+\varepsilon+2\varepsilon^2)-\varepsilon-1}{(1+\varepsilon)3\alpha(1+2\varepsilon)-1}$ on the value of x_P at the point where the solution starting from $x^0 = (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2})$ intersects segment e_Px.

The function β is an increasing hyperbola whose asymptotes lie at $\alpha = \frac{1}{3+9\varepsilon+6\varepsilon^2}$ and $\beta = \frac{3+\varepsilon+2\varepsilon^2}{3+9\varepsilon+6\varepsilon^2}$. It intersects the 45° line at

$$\alpha_\pm = \frac{2+\varepsilon+\varepsilon^2 \pm \sqrt{1-8\varepsilon-10\varepsilon^2-4\varepsilon^3+\varepsilon^4}}{3+9\varepsilon+6\varepsilon^2}.$$

whenever the expression under the square root is positive. This is true whenever $\varepsilon < .1094$. In this case, $(\alpha_-, \alpha_+) \subset (\frac{1}{3}, 1)$, and β is above the 45° line on the former interval. Hence, any solution that begins at a point $x^0 = (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2})$ with $\alpha > \max\{\underline{\alpha}, \alpha_-\}$ will hit segment e_Px^* at some point y with $y_P > \beta(\alpha) \in (\alpha, \alpha_+)$. It then follows from the symmetry of the game and of the choice rule that that the region bounded on the inside by the solution from x^0 to y, its 120° and 240° rotations about x^*, and the pieces of e_Px^*, e_Sx^*, and e_Rx^* that connect the three solutions, and on

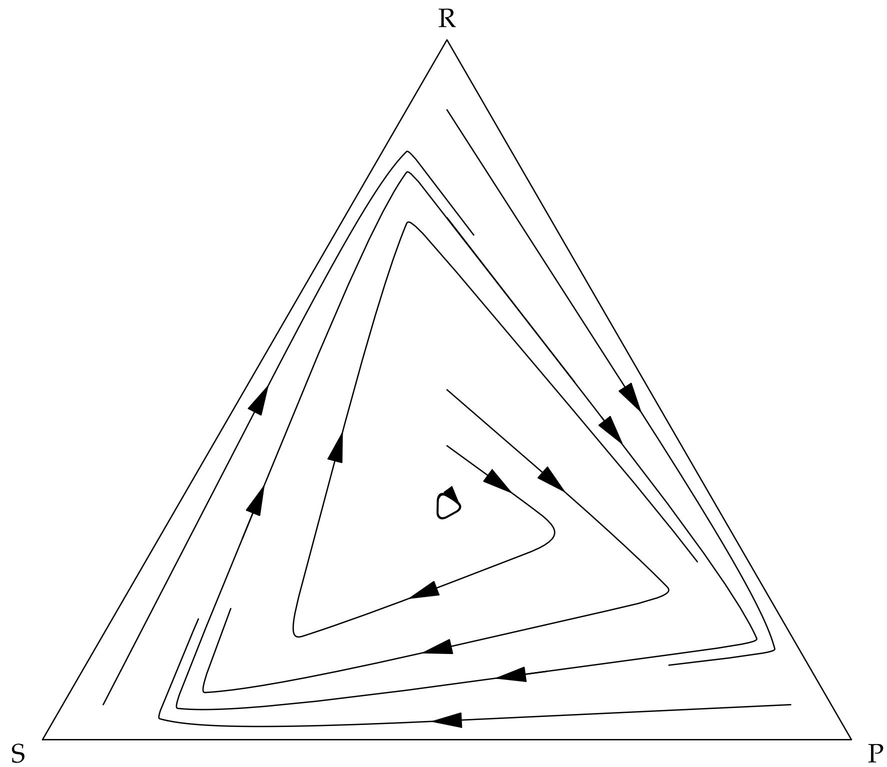

Figure 3: Cycling in (1, –1, 0)-Rock-Scissors-Paper

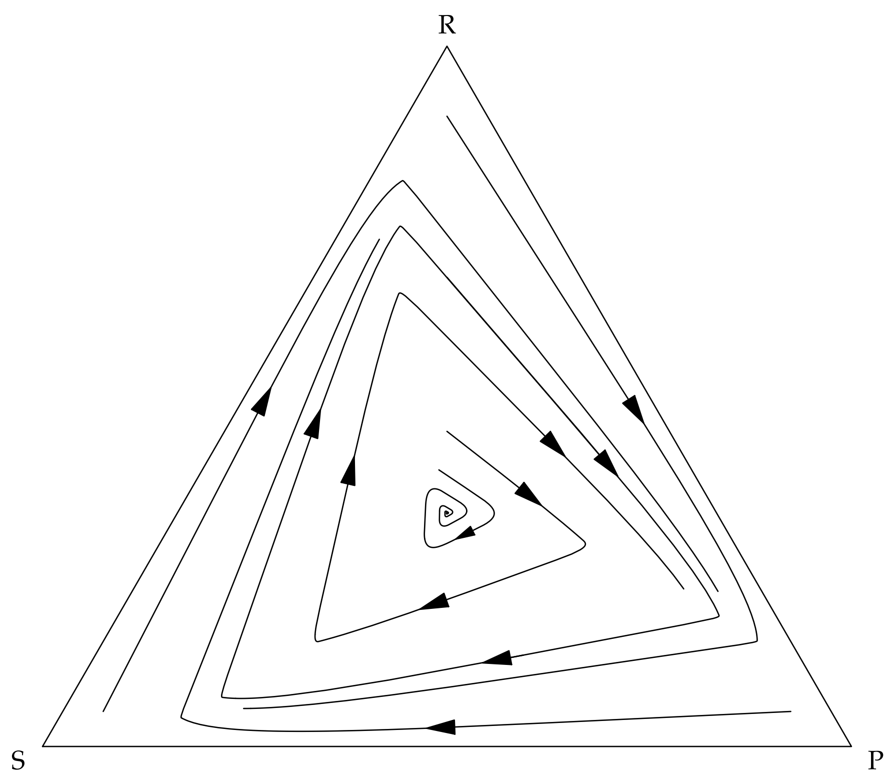

Figure 4: Cycling in (2, –1, 0)-Rock-Scissors-Paper

the outside by the boundary of X is a trapping region for the dynamic V. Theorem 3.3 tells us that the only rest point of the dynamic is the Nash equilibrium x^*, which lies outside of this region. Therefore, the Poincaré-Bendixson Theorem (Hirsch and Smale (1974, Theorem 11.4)) implies that every solution with an initial condition in the region converges to a periodic orbit. If we take ε to zero, $\underline{\alpha}$ and α_{-} approach $\frac{1}{3}$, which implies that the radius of the ball around x^* from which convergence to a periodic orbit is not guaranteed vanishes. ■

Figure 3 presents some numerical solutions to V under the assumption that $\varepsilon = \frac{1}{10}$. Convergence to periodic orbits occurs for initial states $x^0 = (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2})$ with $\alpha >$ .36. That so few initial conditions lead to equilibrium play suggests that the requirement that $\varepsilon < .1094$ is stronger than necessary to obtain cycling. Numerical analysis indicates that cycling occurs for values of ε up to .17.

We considered a zero sum game in this example in order to keep the algebra manageable, but it is clear that a similar analysis would establish the existence of limit cycles in any game with reasonably similar payoffs. In particular, limit cycles under excess payoff dynamics can occur in stable games in which the negative semidefiniteness condition holds strictly. In Figure 4, we present numerical solutions for the dynamic defined by the raw choice and revision rate functions from the previous example and the strictly stable game

$$F(x) = \begin{pmatrix} F_R(x) \\ F_S(x) \\ F_P(x) \end{pmatrix} = \begin{pmatrix} 0x_R + 2x_S - 1x_P \\ -1x_R + 0x_S + 2x_P \\ 2x_R - 1x_S + 0x_P \end{pmatrix},$$

which also has its unique equilibrium at $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Convergence to a periodic orbit occurs from all initial states $x^0 = (\alpha, \frac{1-\alpha}{2}, \frac{1-\alpha}{2})$ with $\alpha > .40$.

8. Potential Dynamics and Stable Games

8.1 Potential Dynamics

Example 7.1 shows that continuity and acuteness of the underlying raw choice functions are not enough to ensure convergence of excess payoff dynamics in stable games. The periodic behavior in this example can be viewed as a consequence of correlations between the choice weights placed on each strategy and the excess payoffs earned by *other* strategies. For acuteness to hold, larger choice weights have some tendency to be assigned to strategies with larger excess payoffs. The choice rule from the example builds in a different sort of dependence: for instance, when the excess payoff to Paper, π_P, is above ε, the choice weight on Scissors, $\tilde{\sigma}_S(\pi)$, must be low. Cycling occurs because of the way in which the choice weight placed on each strategy is made to depend on the performance of the previous strategy in the cycle.

This discussion suggests that global convergence results might be obtained by introducing a condition that rules out systematic dependence between the choice weights on each strategy and the excess payoffs of others. With this motivation, we offer the following condition on raw choice functions $\tilde{\sigma}$.

(I) There exists a continuously differentiable function ψ: $\mathbf{R}_*^{n^p} \to \mathbf{R}$ such that $\tilde{\sigma} \equiv \nabla\psi$.

This condition, *integrability*, demands that the raw choice function be expressible as the gradient of some *choice potential function* ψ.

To begin our discussion of this condition, we first note that all separable raw choice functions are Integrable: every raw choice function of form (3) admits the choice potential

$$\psi(\pi) = \sum_i \int_0^{\pi_i} \phi_i(s_i)\, ds_i .$$

But while separability is sufficient for integrability, it is far from necessary. For example, although the raw choice functions of form (4) are clearly not separable, they are integrable with choice potential

$$\psi(\pi) = \left(\sum_i ([\pi_i]_+)^{k+1}\right)\left(\sum_j \exp(c\pi_j)\right).$$

Separable raw choice functions are those for which the choice weight assigned to each strategy is independent of the excess payoffs of other strategies. It is not immediately obvious whether some natural generalization of this requirement characterizes integrability. Nevertheless, Theorems A.1 and A.2 in the Appendix provide just such a characterization.

Our characterization theorems are stated in terms of the action of $\tilde{\sigma}$ on closed curves through $\mathbf{R}^n$. While the exact statement of these results requires a substantial investment in notation, we can describe them informally as follows. Each closed curve $C \subset \mathbf{R}^n$ can be given a "parameterization" $\rho^{[i]}$: $[0, 1] \to C$ whose ith component, $\rho_i^{[i]}$, changes at a constant speed: in other words, $\left|\dot{\rho}_i^{[i]}\right| = v^{[i]}$, where $v^{[i]}$ is the total variation in component i over the curve C. By construction, the sign of $\dot{\rho}_i^{[i]}$ provides no information about the value of $\rho_i^{[i]}$ itself, and so can be viewed as a statistic that summarizes information about the values of the other components of the vector $\rho^{[i]}$.

Theorem A.1 shows that the vector field $\tilde{\sigma}$ is separable if and only if for each strategy $i \in S$, $\tilde{\sigma}_i(\rho^{[i]})$ and $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ define *independent* random variables. Furthermore, weakening this condition in two distinct ways yields a characterization of integrability. Theorem A.2 shows that the vector field $\tilde{\sigma}$ is integrable if and only if $\tilde{\sigma}_\iota(\rho^{[\iota]})$ and $\mathrm{sgn}(\dot{\rho}_\iota^{[\iota]})$ are *uncorrelated* random variables, where the index ι is *random* with a distribution proportional to the vector of total variations $(v^{[1]}, \ldots, v^{[n]})$.

Put differently, separability requires that learning the choice weight on one strategy conveys no information about the excess payoffs of other strategies. Integrability allows some information to be conveyed, but our characterization theorem shows that this information cannot be systematic: over each curve C, the *expected* impact of such information is zero. Formal statements and additional discussion of these results can be found in the Appendix.

With our integrability condition in hand, we can state our final definition.

Definition: *If the raw choice functions* $\tilde{\sigma}^p$ *satisfy conditions* (C), (A), *and* (I), *we call equation* (E) *a* potential dynamic.

8.2 Global Stability and Uniqueness of Nash Equilibria

Both the classes of examples that exhibit self-defeating externalities and the nature of the property itself suggest that it is a dynamic stability condition. However, Proposition 7.1 shows that excess payoff dynamics need not converge to equilibrium in stable games. In light of the previous discussion, one might hope to establish global convergence by imposing the integrability condition (C3)—in other words, by narrowing the scope of our analysis from the class of excess payoff dynamics to the class of potential dynamics.

Theorem 5.1 established global convergence of excess payoff dynamics in potential games. In its proof, we showed that the potential function of the underlying game serves as a Lyapunov function for all excess payoff dynamics. In the current context, the game itself does not provide us with a candidate Lyapunov function. Fortunately, candidates are suggested by the dynamics themselves. Theorem 8.1 shows that given any potential dynamic V and any stable game F, one can construct a strict Lyapunov function out of the choice potentials that define the V and the excess payoff functions induced by F, enabling us to establish a global convergence result.

Theorem 8.1: *Let F be a stable game, let V be the potential dynamic for this game defined by the choice potentials $\psi^1, \dots, \psi^p$, and define the function*

$$\Lambda(x) = \sum_{p \in P} m^p \psi^p(\hat{F}^p(x)).$$

Then Λ is a decreasing strict Lyapunov function for V. Hence, every solution trajectory of (D) *converges to a connected set of Nash equilibria of F.*

Proof: Recall that the excess payoff function $\hat{F}_i^p$ is defined by $\hat{F}_i^p(x) = F_i^p(x) - \frac{1}{m^p} x^p \cdot F^p(x)$. Since $\tilde{\sigma}^p \equiv \nabla \psi^p$ by condition (I), we can compute the time derivative of the choice potential ψ^p as

$$\begin{aligned}
\dot{\psi}^p(\hat{F}^p(x)) &= \sum_{i \in S^p} \frac{\partial \psi^p}{\partial \pi_i^p}(\hat{F}^p(x))\, \dot{\hat{F}}_i^p(x) \\
&= \sum_{i \in S^p} \tilde{\sigma}_i^p(\hat{F}^p(x))\Big((e_i^p - \tfrac{1}{m^p} x^p) \cdot DF^p(x)\dot{x} - \tfrac{1}{m^p} \dot{x}^p \cdot F^p(x)\Big) \\
&= \Big(\tilde{\sigma}^p(\hat{F}^p(x)) - \tfrac{1}{m^p} \tilde{\sigma}_T^p(\hat{F}^p(x)) x^p\Big) \cdot DF^p(x)\dot{x} - \tfrac{1}{m^p} \tilde{\sigma}_T^p(\hat{F}^p(x))\Big(\dot{x}^p \cdot F^p(x)\Big)
\end{aligned}$$

If V is a potential dynamic, the first parenthesized expression equals $\frac{1}{m^p} \dot{x}^p$, so

$$\begin{aligned}
\dot{\Lambda}(x) &= \sum_{p \in P} m^p \dot{\psi}^p(\hat{F}^p(x)) \\
&= \dot{x} \cdot DF(x)\dot{x} - \sum_{p \in P} \tilde{\sigma}_T^p(\hat{F}^p(x))\Big(\dot{x}^p \cdot F^p(x)\Big).
\end{aligned}$$

The first term in this expression is negative by condition (SE); it equals zero at rest points of V. Since Theorem 3.1 tells us that V satisfies condition (PC), the second term in the expression is positive, equaling zero only at rest points of V. Combining these observations, we see that Λ is a decreasing strict Lyapunov function for V. The proof is completed in the same fashion as that of Theorem 5.1. ■

An interpretation of the Lyapunov function Λ can be found in Section A.2 of the Appendix.

Theorem 8.1 builds on a result of Hofbauer (2000, Theorem 6.1). Hofbauer (2000) studies evolution in single population normal form games in which the payoff matrix A satisfies the negative semidefiniteness condition $z \cdot Az \leq 0$ for all tangent directions z. He proves a global stability theorem for members of a certain class of dynamics; in our terminology, they are the potential dynamics based on raw choice functions of the separable, sign preserving form (3). Theorem 8.1 shows that the separability and the sign preserving property of the choice rule are unnecessary to establish a global convergence result, and that the restrictions to single population games and to linear payoff functions are inessential as well.

In both potential games and stable games, solutions to broad classes of dynamics converge to connected sets of equilibria. Potential games generally admit many components of equilibria, and solutions from different initial conditions converge to different components of equilibria.[19] In contrast, Theorem 8.2 shows that every stable game possesses a unique connected component of Nash equilibria, and that under a mild additional condition, the Nash equilibrium is unique. Theorem 8.1 then implies that this component is globally stable under all potential dynamics. In fact, the proof of the uniqueness provided here is based on the characterization of dynamics provided by the previous theorem.[20]

Theorem 8.2: (*i*) *If F is a stable game, then all Nash equilibria of F lie in a single connected component, which is therefore globally stable.*

(*ii*) *If x is a Nash equilibrium of the stable game F such that $DF(x)$ is negative definite with respect to TX, then x is the unique Nash equilibrium of F.*

Proof: (*i*) Consider the BNN dynamic, which is the potential dynamic defined by the choice potential $\psi(\pi) = \frac{1}{2}\sum_i ([\pi_i]_+)^2$. Theorem 3.1 tells us that the Nash equilibria of F are precisely the rest points of the this dynamic. Moreover, Proposition 3.4 shows that these rest points are precisely the states x satisfying $\hat{F}^p(x) \in bd(\mathbf{R}^{n^p}_*)$ for all $p \in \mathcal{P}$; at all other states, $\hat{F}^p_i(x) > 0$ for some $i \in S^p$ and $p \in \mathcal{P}$. It

[19] The Nash equilibria of a potential game are those states that satisfy the Kuhn-Tucker first order conditions for maximizing potential on the state space X, while the locally stable states are those that locally maximize potential (Sandholm (2001)).

[20] In fact, it has recently been established that the set of Nash equilibria of a stable game is not only connected, but also *convex*: see Hofbauer and Sandholm (2004).

therefore follows from the definition of ψ that the rest points of the BNN dynamic are precisely those states where the Lyapunov function $\Lambda(x) = \sum_{p \in \mathcal{P}} m^p \psi^p(\hat{F}^p(x))$ takes the value zero, and that Λ is strictly positive at all other states.

Let $E = \{x \in X\colon\ \Lambda(x) = 0\}$ be the set of Nash equilibria of F, and suppose that E is not connected. Then by definition, there exists a partition of E into two sets, E_1 and E_2, neither of which intersects the closure of the other. Now if $\{x_k\}$ is a convergent sequence in E_i, the continuity of Λ implies that its limit is in E, and hence in E_i. That is, E_1 and E_2 are actually disjoint *closed* sets.

Since the closed set E_i admits a Lyapunov function, it is asymptotically stable (Weibull (1995, Theorem 6.3)), and so possesses a basin of attraction $B(E_i)$ that is open relative to X (Hirsch and Smale (1974, p. 190)). By definition, $B(E_1)$ and $B(E_2)$ are disjoint. Let y and z be points in $B(E_1)$ and $B(E_2)$ respectively, and define $\hat{x}$ to be the point on the segment yz that is closest to y among points on the segment that are not in $B(E_1)$. This point exists because $B(E_1)$ and $B(E_2)$ are disjoint and open, and in fact $\hat{x}$ is on the boundary of $B(E_1)$. Indeed, $\hat{x} \notin B(E_1) \cap B(E_2)$. Since $\hat{x}$ is in neither basin of attraction, the solution starting from $\hat{x}$ does not converge to either E_1 or E_2. But since all Nash equilibria are in either E_1 or E_2, this contradicts Theorem 8.1. Hence, E must be connected, establishing part (*i*) of the theorem.

(*ii*) We begin by stating a simple characterization of Nash equilibrium in general population games. To do so, we define the set of "good" (i.e., inward pointing) directions at $x \in X$ as

$$
\begin{aligned}
G(x) &= \{h \in TX\colon\ x + th \in X \text{ for some } t > 0\} \\
&= \{h \in TX\colon\ x_i^p = 0 \text{ implies that } h_i^p \geq 0\}.
\end{aligned}
$$

Lemma 8.3: *Let F be a population game. Then $x \in X$ is a Nash equilibrium of F if and only if for each $h \in G(x)$, the inequality $h^p \cdot F^p(x) \leq 0$ holds for all $p \in \mathcal{P}$.*

Proof: x is a Nash equilibrium of F

$$
\begin{aligned}
&\Leftrightarrow\ x^p \cdot F^p(x) \geq y^p \cdot F^p(x)\ \ \forall y \in X, p \in \mathcal{P} \\
&\Leftrightarrow\ (y^p - x^p) \cdot F^p(x) \leq 0\ \ \forall y \in X, p \in \mathcal{P} \\
&\Leftrightarrow\ h^p \cdot F^p(x) \leq 0\ \ \forall h \in G(x), p \in \mathcal{P}. \ \blacksquare
\end{aligned}
$$

We now continue with the proof of part (*ii*). Let $y \in X$ be a state distinct from x; we will show that y is not a Nash equilibrium of x. Let $h = y - x$. Then $h \in G(x)$ and $-h \in G(y)$. Next, define $f_{x,h}(t) = h \cdot F(x + th)$. Then the stability of F implies that $f'_{x,h}(t) = h \cdot DF(x + th)h \leq 0$ for $t \in [0, 1]$, while the strict stability of F at x shows that $f'_{x,h}(0) <$

0. Moreover, since x is a Nash equilibrium, the lemma tells us that $h^p \cdot F^p(x) \le 0$ for all $p \in \mathcal{P}$, and hence that $f_{x,h}(0) = h{\cdot}F(x) \le 0$. Thus, integrating yields

$$h{\cdot}F(y) = f_{x,h}(1) = f_{x,h}(0) + \int_0^1 f'_{x,h}(t)\,dt < f_{x,h}(0) \le 0.$$

We can rewrite this inequality as $(-h){\cdot}F(y) > 0$, which implies that $(-h^p) \cdot F^p(y) > 0$ for some $p \in \mathcal{P}$. Since $-h \in G(y)$, we conclude from the lemma that y is not a Nash equilibrium of F. ■

Appendix A: A Probabilistic Characterization of Integrability

A.1 Two Characterization Theorems

Let $\tilde{\sigma}$ be a continuous vector field on $\mathbf{R}^n$. Suppose that $\tilde{\sigma}$ is *separable*: in other words, that $\tilde{\sigma}_i(\pi)$ only depends on π_i. This property can be characterized in terms of the independence of certain appropriately constructed random variables. For example, consider any product set Π in $\mathbf{R}^n$ whose components are nonempty compact intervals, and endow this set with the uniform probability measure. Then separability clearly implies that the random variable $\tilde{\sigma}_i(\pi)$ is independent of the random vector $\pi_{-i} = (\pi_1, \ldots, \pi_{i-1}, \pi_{i+1}, \ldots, \pi_n)$. In fact, one can show that this independence condition on product sets characterizes the separable vector fields.

We now construct an alternative independence condition that is also equivalent to separability. This condition is stated in terms of random variables that are defined using piecewise smooth closed curves through $\mathbf{R}^n$. Call the set of such curves $\mathbf{C}$, and fix a curve $C \in \mathbf{C}$. If we let γ: [0, 1] → C be some parameterization of C, then the total variation in component γ_i along C is $v^{[i]} = \int_\Omega |\dot{\gamma}_i(t)|\,dt$. The value of $v^{[i]}$ does not depend on the choice of parameterization. For each index i for which $v^{[i]} > 0$, introduce a function $\rho^{[i]}$: [0, 1] → C, a "parameterization" of C whose ith component changes at a constant speed. In particular, $\rho^{[i]}$ satisfies these six properties: (*i*) $\rho^{[i]}$ is right continuous with left limits; (*ii*) $\rho^{[i]}$ is smooth on any interval upon which it is continuous; (*iii*) $\rho^{[i]}$ has the same orientation as γ; (*iv*) $\rho^{[i]}$ is one-to-one (though not necessarily onto); (*v*) the ith component of C is constant between $\rho^{[i]}(t^-)$ and $\rho^{[i]}(t)$ whenever these two points differ; and, (*vi*) $|\dot{\rho}_i^{[i]}| = v^{[i]}$ at all but finitely many $t \in [0, 1]$. It is easy to verify that such a function always exists.

We can view each $\rho^{[i]}$ as a random variable by supposing that the parameter t is determined via a uniform random draw from the unit interval. Our alternative characterization of separability is stated in terms of relationships between the ith component of $\tilde{\sigma}(\rho^{[i]})$ and the direction of motion of the ith component of $\rho^{[i]}$.

Theorem A.1: *The vector field $\tilde{\sigma}$ is separable if and only if for all $C \in \mathbf{C}$ and all $i \in \{1, \dots, n\}$, $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ and $\tilde{\sigma}_i(\rho^{[i]})$ are independent random variables.*

For intuition, fix a vector field $\tilde{\sigma}$ and an index i, and consider the following question: what would knowledge of components of the vector π other than component π_i tell us about the value of $\tilde{\sigma}_i(\pi)$? If the vector field $\tilde{\sigma}$ is separable, no information about $\tilde{\sigma}_i(\pi)$ is provided at all. Earlier, we expressed this idea by considering uniform draws of the vector π from product sets in $\mathbf{R}^n$. Here, we instead consider choices of π from closed curves through $\mathbf{R}^n$.

Fix a "parameterization" $\rho^{[i]}$ of the curve C, and suppose that a parameter t will be drawn at random from the unit interval. Then since C is closed and since $\left|\dot{\rho}_i^{[i]}\right|$ is constant, each rightward motion of $\rho_i^{[i]}$ through a portion of some interval I can be paired with a corresponding leftward motion. Hence, the event $\{\rho_i^{[i]} \in I\} \subseteq [0, 1]$ has its mass evenly divided between the events $\{\mathrm{sgn}(\dot{\rho}_i^{[i]}) = 1\}$ and $\{\mathrm{sgn}(\dot{\rho}_i^{[i]}) = -1\}$. It follows that while learning the sign of $\dot{\rho}_i^{[i]}$ provides information about the values of components of $\rho^{[i]}$ besides $\rho_i^{[i]}$, it provides no information about $\rho_i^{[i]}$ itself. If the vector field $\tilde{\sigma}$ is separable, this implies that $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ is uninformative about $\tilde{\sigma}_i(\rho^{[i]})$, since the value of the latter only depends on the value of $\rho_i^{[i]}$. Conversely, if $\tilde{\sigma}$ is not separable, our proof shows how one can construct a curve C such that $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ is informative about $\tilde{\sigma}_i(\rho^{[i]})$.

Since independence of $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ and $\tilde{\sigma}_i(\rho^{[i]})$ for all i and C is equivalent to separability of $\tilde{\sigma}$, it is natural to ask whether some weakening of independence corresponds to integrability of $\tilde{\sigma}$: that is, to the requirement that $\tilde{\sigma} \equiv \nabla\psi$ for some potential function ψ: $\mathbf{R}^n \to \mathbf{R}$. One plausible possibility to consider is to weaken the requirement of independence to that of zero correlation. To obtain our characterization theorem, we make this modification along with one additional change: we require not only that the parameter t fed into the "parameterization" $\rho^{[i]}$ be chosen at random, but also that the index i of the "parameterization" utilized be chosen at random at well.

Consider the following two step procedure for selecting a point on the closed curve C. First, randomly draw an index i from the set S. It is natural to define the probability of drawing index i to be proportional to the variation $v^{[i]}$, which represents the distance traversed along coordinate i during one circuit of the curve. One the index i is determined, use the "parameterization" $\rho^{[i]}$ to randomly select a point on C as before. To represent this procedure formally, we let

$$V = \sum_{i=1}^{n} v^{[i]};$$

$$\iota(\omega) = i \text{ for } \omega \in \left(V^{-1}\sum_{j=1}^{i-1} v^{[i]}, V^{-1}\sum_{j=1}^{i} v^{[i]} \right];$$

$$\tau(\omega) = (v^{\iota(\omega)})^{-1}\left(\omega - \sum_{j=1}^{\iota(\omega)-1} v^{[i]} \right); \text{ and}$$

$$\rho^{[C]}(\omega) = \rho^{[\iota(\omega)]}(\tau(\omega)).$$

If $\omega \in [0, 1]$ is the realized state, then $\iota(\omega)$ is the index of the "parameterization" we consider, and $\tau(\omega)$ is the argument inserted in $\rho^{[\iota(\omega)]}(\cdot)$ to determine the point in C we choose. If the state ω is obtained via a uniform draw from the unit interval, then one can verify that these definitions capture the process described in words above.

Zero correlation between $\text{sgn}(\dot{\rho}_\iota^{[C]})$ and $\tilde{\sigma}_\iota(\rho^{[C]})$ can be interpreted as follows. As before, learning $\text{sgn}(\dot{\rho}_i^{[i]})$ is informative about $\rho_{-i}^{[i]}$, but is uninformative about $\rho_i^{[i]}$. But as $\text{sgn}(\dot{\rho}_i^{[i]})$ and $\tilde{\sigma}_i(\rho^{[i]})$ may not be independent, knowledge about the realization of $\rho_{-i}^{[i]}$ obtained by learning $\text{sgn}(\dot{\rho}_i^{[i]})$ may provide information about the value of $\tilde{\sigma}_i(\rho^{[i]})$. That $\text{sgn}(\dot{\rho}_\iota^{[C]})$ and $\tilde{\sigma}_\iota(\rho^{[C]})$ are uncorrelated imposes a restriction on the nature of this information. In particular, it must be the case that after averaging over the possible realizations of the index ι, the information about components $\rho_{-\iota}^{[\iota]}$ provided by the sign of $\dot{\rho}_\iota^{[C]}$ does not change one's assessment of the *expected value* of $\tilde{\sigma}_\iota(\rho^{[C]})$. More precisely, we show in the course of the proof below that there is zero correlation between $\text{sgn}(\dot{\rho}_\iota^{[C]})$ and $\tilde{\sigma}_\iota(\rho^{[C]})$ if and only if $E[\tilde{\sigma}_\iota(\rho^{[C]}) \mid \text{sgn}(\dot{\rho}_\iota^{[C]}) = 1] = E[\tilde{\sigma}_\iota(\rho^{[C]}) \mid \text{sgn}(\dot{\rho}_\iota^{[C]}) = -1] = E\,\tilde{\sigma}_\iota(\rho^{[C]})$.

Our main result establishes that the absence of correlation between $\text{sgn}(\dot{\rho}_\iota^{[C]})$ and $\tilde{\sigma}_\iota(\rho^{[C]})$ for all closed curves C fully characterizes the integrable vector fields.

Theorem A.2: *The vector field $\tilde{\sigma}$ is integrable if and only if for all $C \in \mathbf{C}$, $\text{sgn}(\dot{\rho}_\iota^{[C]})$ and $\tilde{\sigma}_\iota(\rho^{[C]})$ are uncorrelated random variables.*

The proofs of Theorems A.1 and A.2 can be found in Appendix B.

A.2 Discussion

To provide intuition about the role of integrability in establishing convergence to equilibrium in stable games, we reconsider evolution in Rock-Scissors-Paper in light of Theorems A.1 and A.2.

Consider the boundary of the simplex, bd(Δ), which is a piecewise linear closed curve. Since the payoffs $F\colon \Delta \to \mathbf{R}^3$ of Rock-Paper-Scissors define a full rank linear transformation on Δ, $C = F(\text{bd}(\Delta))$ is piecewise linear closed curve through $\mathbf{R}^3$. In fact, since Rock-Paper-Scissors is a zero-sum game, its payoffs and excess payoffs are identical, and so C can also be viewed as a an excess payoff trajectory. It is easily verified that the pivot points of C are given by the columns of the payoff matrix: $\pi^R = (\pi^R_R, \pi^R_S, \pi^R_P) = (0, -1, 1)$, $\pi^S = (1, 0, -1)$, and $\pi^P = (-1, 1, 0)$. The closed curve C is pictured in Figure 5 below.

Because of the symmetry of the curve C and of the choice weight functions $\tilde{\sigma}$ we will consider, the covariances cov(sgn($\dot{\rho}^{[i]}_i$), $\tilde{\sigma}_i(\rho^{[i]})$) will be the same for each strategy $i \in \{R, S, P\}$; cov(sgn($\dot{\rho}^{[C]}_\iota$), $\tilde{\sigma}_\iota(\rho^{[C]})$) therefore equals their common value. This fact simplifies our discussion by enabling us to focus on the expression cov(sgn($\dot{\rho}^{[R]}_R$), $\tilde{\sigma}_R(\rho^{[R]})$), the covariance corresponding to the strategy Rock.

To compute this covariance, we first parameterize the curve C so that the payoff to Rock changes at a constant rate. If C is traversed in the order $\pi^P \to \pi^R \to \pi^S \to \pi^P$, the payoff to R increases as one travels from $\pi^P = (-1, 1, 0)$ to $\pi^R = (0, -1, 1)$ to $\pi^S = (1, 0, -1)$ and then falls as one continues from π^S back to π^P. Thus, since each segment of C is linear, the constant speed parameterization $\rho^{[R]}$ has pivot points $\rho^{[R]}(0) = \pi^P$, $\rho^{[R]}(\frac{1}{4}) = \pi^R$, $\rho^{[R]}(\frac{1}{2}) = \pi^S$, and $\rho^{[R]}(1) = \pi^P$, and is linear between these points. Figure 5 illustrates this parameterization of C, while Figure 6 graphs the individual components $\rho^{[R]}_R$, $\rho^{[R]}_S$, and $\rho^{[R]}_P$ of the parameterization.

On the interval $I = [0, \frac{1}{2})$, the payoff to Rock, $\rho^{[R]}_R$, is increasing, and hence sgn($\dot{\rho}^{[R]}_R$) = 1; similarly, sgn($\dot{\rho}^{[R]}_R$) = –1 on the interval $D = (\frac{1}{2}, 1]$. Figure 6 shows that $\rho^{[R]}_S$, the payoff to Scissors, tends to take above average values when the event D occurs ($E[\rho^{[R]}_S \mid D] = \frac{1}{2} > \frac{1}{8} = E\rho^{[R]}_S$), while $\rho^{[R]}_P$, the payoff to strategy Paper, tends to take below average values on this event ($E[\rho^{[R]}_P \mid D] = -\frac{1}{2} < -\frac{1}{8} = E\rho^{[R]}_P$).[21] Thus, the value of cov(sgn($\dot{\rho}^{[R]}_R$), $\tilde{\sigma}_R(\rho^{[R]})$) indicates the degree to which high values of the payoff to Scissors and low values of the payoff to Paper are associated with high values of the choice weight on Rock along the curve C.

---

[21] The definition of the constant speed parameterization implies that $E[\rho^{[R]}_R \mid D] = E[\rho^{[R]}_R \mid I] = E\rho^{[R]}_R$. In the present case, this common expectation equals zero.

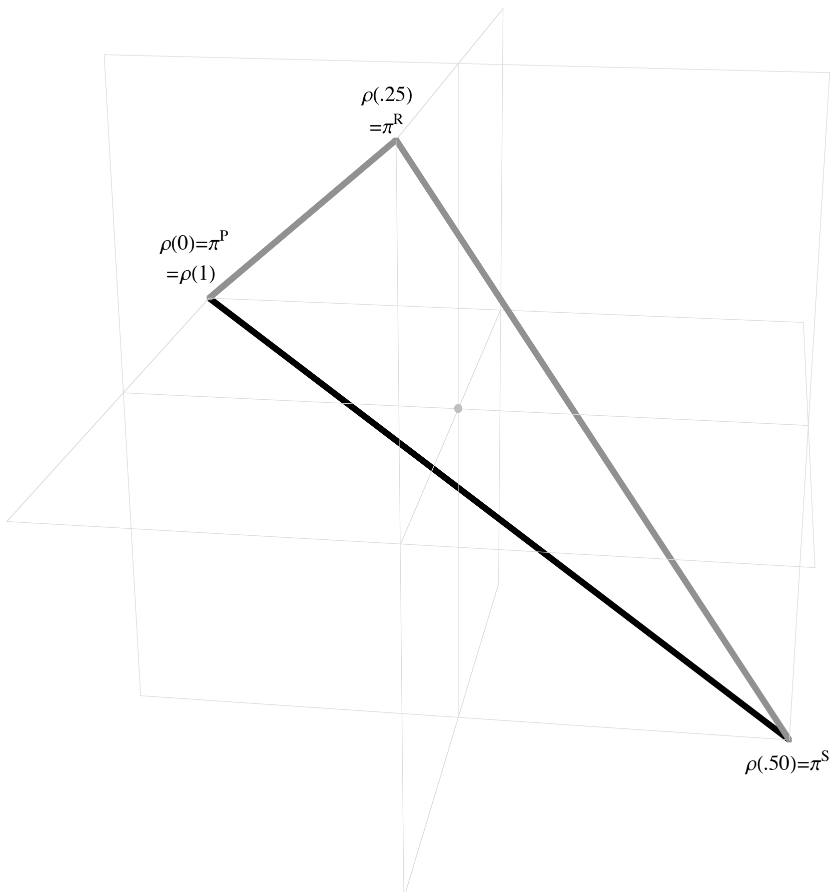

Figure 5 : A closed curve through the space of (excess) vectors.

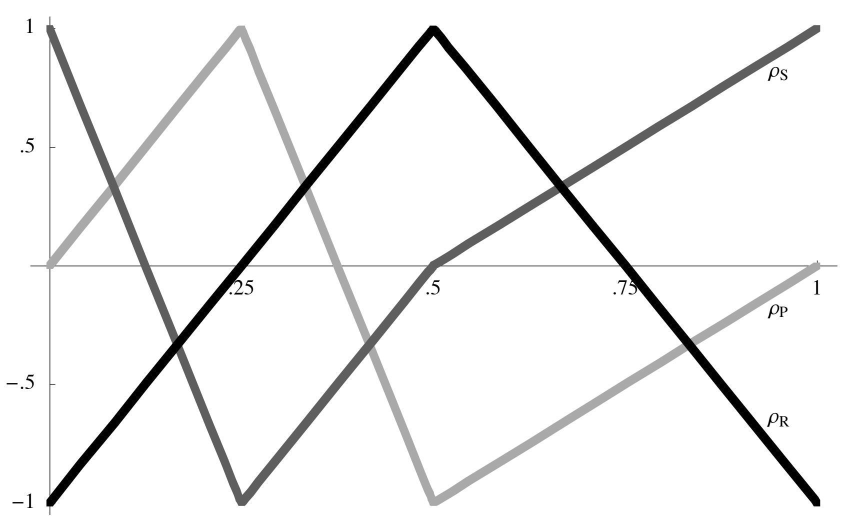

Figure 6 : The same curve, component by component.

We consider excess payoff dynamics generated by three specifications of the choice weight function $\tilde{\sigma}$. These specifications are drawn from our examples from Sections 2 and 7.

(*i*) $\tilde{\sigma}_i(\pi) = [\pi_i]_+$;

(*ii*) $\tilde{\sigma}_i(\pi) = \left(2\sum_j \exp(\pi_j)\right)[\pi_i]_+ + \left(\sum_j ([\pi_j]_+)^2\right)\exp(\pi_i)$;

(*iii*) $\tilde{\sigma}_i(\pi) = [\pi_i]_+ \, g^{\varepsilon}(\pi_{(i+1) \bmod 3})$.

In Figure 7, we plot the choice weights $\tilde{\sigma}_R(\rho^{[R]})$ along parameterization $\rho^{[R]}$ for each specification of $\tilde{\sigma}$; in specification (*iii*), we let $\varepsilon = \frac{1}{5}$.

In the truncated linear specification (*i*), the choice weight placed on each strategy is given by the positive part of its payoff. As this specification of $\tilde{\sigma}$ is separable, Theorem A.1 implies that $\text{sgn}(\dot{\rho}_R^{[R]})$ and $\tilde{\sigma}_R(\rho^{[R]})$ are independent random variables. This property is clearly visible in Figure 3(*i*): the distribution of $\tilde{\sigma}_R(\rho^{[R]})$ conditional on event $I = [0, \frac{1}{2})$ is the same as its distribution conditional on $D = (\frac{1}{2}, 1]$.

In specification (*ii*), the choice weight placed on a strategy is a weighted sum of the positive part of its payoff and the exponential of its payoff. This choice of $\tilde{\sigma}$ is clearly not separable, and Figure 3(*ii*) shows that the distribution of $\tilde{\sigma}_R(\rho^{[R]})$ conditional on I differs from its distribution conditional on D. But since $\tilde{\sigma}$ is integrable, Theorem A.2 and the symmetry of our example imply that $\text{sgn}(\dot{\rho}_R^{[R]})$ and $\tilde{\sigma}_R(\rho^{[R]})$ are uncorrelated. Zero correlation between $\text{sgn}(\dot{\rho}_R^{[R]})$ and $\tilde{\sigma}_R(\rho^{[R]})$ is equivalent to the equality of the conditional means $E[\tilde{\sigma}_R(\rho^{[R]}) \mid D]$ and $E[\tilde{\sigma}_R(\rho^{[R]}) \mid I]$.[22] This property appears consistent with Figure 3(*ii*), and in fact these conditional means have a common value of approximately 2.3926.[23] We can summarize this discussion somewhat loosely as follows: when a vector field is integrable but not separable, the values of alternative payoffs π_{-i} are informative about choice weights $\tilde{\sigma}_i(\pi)$, but the *average* influence of π_{-i} on $\tilde{\sigma}_i(\pi)$ over any closed curve C must equal zero.

To explain the connection between zero correlation and convergence, let us first recall a result on evolution in zero sum games due to Hofbauer (1995) (also see

---

22 See the claim immediately preceding Theorem A.2.

23 Still, since the two conditional distributions differ, the occurrence of D or I provides other information about the choice weight placed on Rock. For example, learning that D has occurred increases the variance in one's beliefs about the choice weight $\tilde{\sigma}_R(\rho^{[R]})$: one can compute that the initial variance of $\tilde{\sigma}_R(\rho^{[R]})$ is 8.3405, but that its conditional variances are given by $\text{var}[\tilde{\sigma}_R(\rho^{[R]}) \mid D] \approx 8.6371$ and $\text{var}[\tilde{\sigma}_R(\rho^{[R]}) \mid I] \approx 8.0440$.

Brown (1951)): in these games, the maximum payoff function $M(x) = \max_{i \in S} F_i(x)$ is a Lyapunov function for the best response dynamic. Intuitively, this result shows that in a strictly competitive game played by a single population of agents, the act of switching to an optimal strategy reduces that strategy's payoff. It follows that in the long run, the payoffs to all strategies in use are equalized, and so Nash equilibrium is attained.

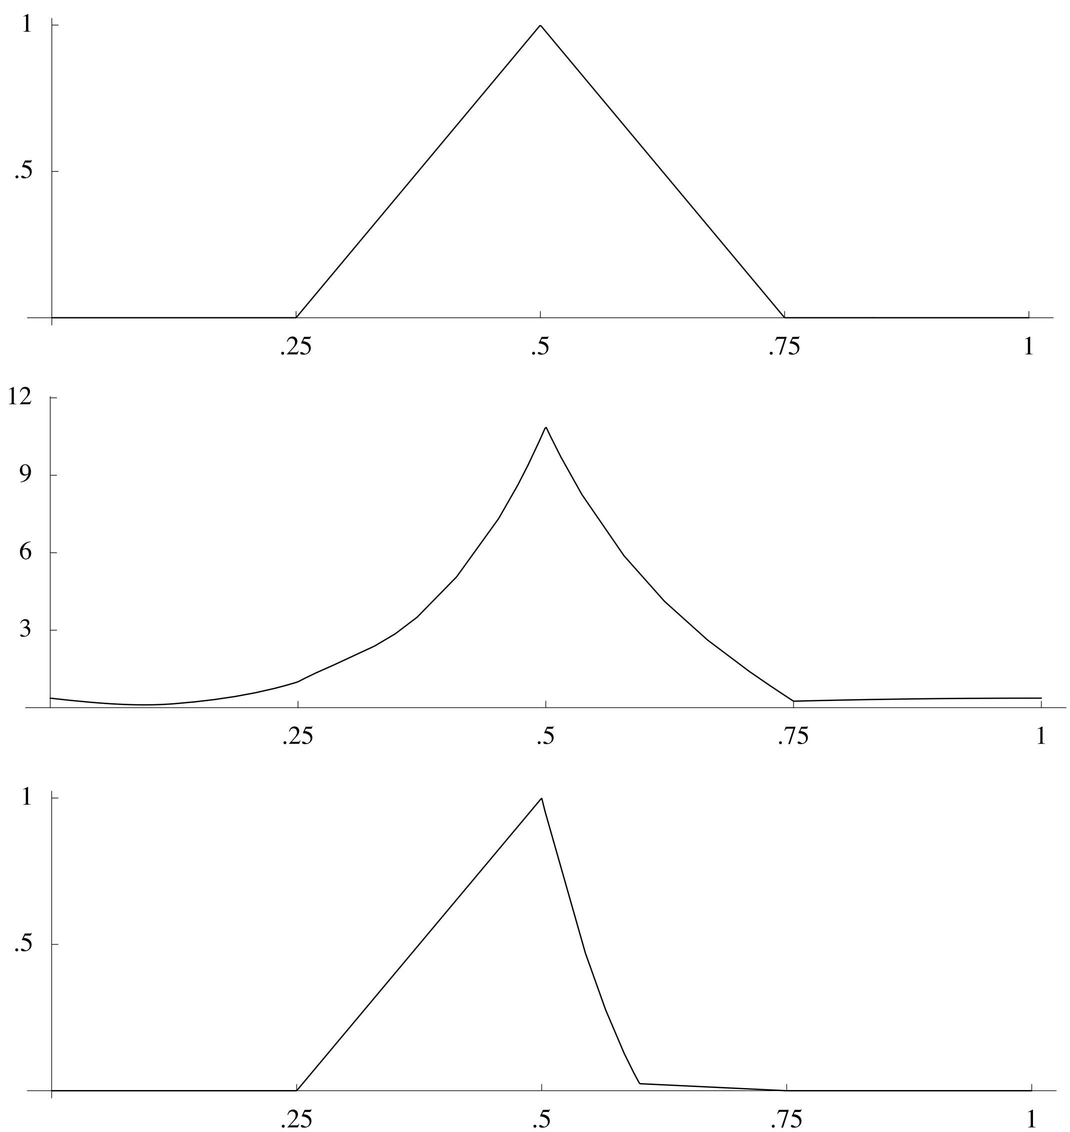

Figures 7(i), 7(ii), and 7(iii): Plots of $\tilde{\sigma}_R(\rho^{[R]})$ for three specifications of $\tilde{\sigma}_R$

The best response dynamic can be derived from the raw choice function $\tilde{\sigma}_i(\pi) = 1_{\{i = \arg\max_j \pi_j\}}$, under which the optimal strategy is always chosen. But the preceding argument can be extended to settings where choices are not always optimal, but where choice weights $\tilde{\sigma}$ are still acute and separable: that is, $\tilde{\sigma}_i(\pi) = \phi_i(\pi_i)$ as in equation (3). In this case, Theorem 8.1 implies that the function $\Phi(x) = \sum_{i \in S} \int_0^{\hat{F}_i(x)} \phi_i(s_i)\, ds_i$ serves as a Lyapunov function for the corresponding potential dynamic. Thus, as long as choice weights depend positively and separably on performance, there is still negative feedback from choices to a well chosen measure of payoff opportunities. Because all strategies whose payoffs are above average are chosen, the payoff opportunity measure that replaces M depends upon the payoffs to all such strategies. Still, the spirit of the previous analysis is preserved here: in both cases the fact that agents switch to "good" strategies proves detrimental to a measure of the performance of "good" strategies.

If $\tilde{\sigma}$ is acute but not separable, then the choice weight $\tilde{\sigma}_i(\pi)$ can depend on the payoffs π_{-i} earned by other strategies. If $\tilde{\sigma}$ is not integrable, this dependence can be systematic, in the sense that it does not vanish after averaging. In particular, the effects of this dependence can accumulate as one traverses closed curves, disrupting the negative feedback from choices to payoff opportunities that underlies Theorem 8.1.

For example, under the raw choice function defined in equation (*iii*), the choice weight placed on Rock is $[\pi_R]_+$ when the payoff to Scissors is negative, but is only $\varepsilon^2[\pi_R]_+$ when the payoff to Scissors exceeds ε. Now along the parameterization $\rho^{[R]}$, Scissors generally has a payoff above ε when event D occurs but not when event I occurs (see Figure 6). Consequently, in Figure 7(*iii*) we see that the choice weight $\tilde{\sigma}_R(\rho^{[R]})$ is typically lower when D occurs than when I occurs: indeed, $E[\tilde{\sigma}_R(\rho^{[R]}) \mid D] \approx .0932 < \frac{1}{4} = E[\tilde{\sigma}_R(\rho^{[R]}) \mid I]$. Theorem A.2 implies that a choice function that exhibits this sort of systematic dependence on the payoffs of alternative strategies cannot be integrable. And as Proposition 7.1 illustrates, there may be no measure of payoff opportunities whose value falls over time in cases where such dependence exists

On the other hand, if $\tilde{\sigma}$ is acute and integrable as in equations (*ii*) and (*iii*), Theorem A.2 shows that such systematic relationships between choice weights $\tilde{\sigma}_i(\pi)$ and alternative payoffs π_{-i} cannot exist. In particular, since the influence of alternative payoffs π_{-i} on choice weights $\tilde{\sigma}_i(\pi)$ averages to zero on each closed curve C, the negative feedback from choices to payoff opportunities required to rule out cycling through C is preserved. Indeed, Theorem 8.1 shows that when $\tilde{\sigma}$ is acute and integrable, its potential function ψ can be used to construct the Lyapunov function $\Lambda(x) = \psi(\hat{F}(x))$, a measure of payoff opportunities whose value falls over time. The existence of this Lyapunov function implies the global asymptotic stability of Nash equilibrium.

Appendix B: Additional Proofs

The Proof of Proposition 2.1

Lipschitz continuity, nonseparability, and strict positivity clearly hold. To check acuteness, we compute that

$$\begin{aligned}\tilde{\sigma}(\pi)\cdot\pi &= \left((k+1)\sum_j \exp(c\pi_j)\right)\left(\sum_i \pi_i([\pi_i]_+)^k\right) + \left(c\sum_j([\pi_j]_+)^{k+1}\right)\left(\sum_i \pi_i \exp(c\pi_i)\right) \\ &= \left(\sum_i \exp(c\pi_i)(c\pi_i + k + 1)\right)\left(\sum_j([\pi_j]_+)^{k+1}\right).\end{aligned}$$

The second summation is strictly positive on $\text{int}(\mathbf{R}^n_*)$. To sign the first summation, note that the derivative of its ith term, $c\exp(c\pi_i)(c\pi_i + k + 2)$, has the same sign as $\pi_i + \frac{k+2}{c}$. Thus, the ith term itself is minimized when $\pi_i = -\frac{k+2}{c}$, where it takes the value $-\exp(-(k+2))$. Now any vector in $\text{int}(\mathbf{R}^n_*)$ has at least one strictly positive component π_j. The corresponding component of the first summation must strictly exceed $k + 1$. Since each of the remaining $n - 1$ components the summation is bounded below by $-\exp(-(k+2))$, the summation will be strictly positive whenever $-(n-1)\exp(-(k+2)) + (k+1) \geq 0$, and hence whenever $(k+1)\exp(k+2) + 1 \geq n$. ■

The Proof of Proposition 4.1

It is easy to see that the properties we appealed to in the proof of Theorem 3.1 in proving existence and uniqueness of solutions are satisfied not only by V, but also by I, and that these properties are closed under convex combination. Thus, C_α satisfies condition (EU). It is also simple to verify that condition (PC) is closed under convex combination, so Lemma A3 of Sandholm (2002) and Theorem 3.2 above imply that C_α satisfies this condition. To establish condition (NS), recall that the rest points of V are precisely the Nash equilibria of the underlying game (by Theorem 3.3), and that the rest points of I include the Nash equilibria of F. It follows immediately that all Nash equilibria are rest points of C_α, and that non-Nash rest points of I are not rest points of C_α. To complete the proof, suppose that x is neither a rest point of V nor a rest point of I. Then since both of these dynamics satisfy condition (PC), we know that $V^p(x) \cdot F^p(x) > 0$ and $I^p(x) \cdot F^p(x) > 0$ for all $p \in \mathcal{P}$. Hence, $C^p_\alpha(x) \cdot F^p(x) > 0$, and so x is not a rest point of C_α. We therefore conclude that C_α satisfies (NS). ■

The Proof of Theorem A.1

We begin with a preliminary result.

Lemma B.1: $\int_{(\rho_i^{[i]})^{-1}(I)} \dot{\rho}_i^{[i]}(\tau)\,d\tau = 0$ *for all half-infinite intervals* $I = (-\infty, c] \subseteq \mathbf{R}$.

Proof: The result is obvious if the closed curve C is contained in $I \times \mathbf{R}^{n-1}$, so suppose this is not the case. Since C is piecewise smooth, the points where its ith component is contained in I can be divided into (at most) countably many connected components. Each connected component of $(\rho_i^{[i]})^{-1}(I) \subseteq [0, 1]$ corresponds to exactly one of the connected components of C, although discontinuities in $\rho^{[i]}$ may render this correspondence many-to-one. Since the connected components of C enter and leave $I \times \mathbf{R}^{n-1}$ through the same boundary $\{c\} \times \mathbf{R}^{n-1}$, the corresponding components of $(\rho_i^{[i]})^{-1}(I)$ do not contribute to the integral above. ❏

Now suppose that $\tilde{\sigma}$ is separable; we would like to show that $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ and $\tilde{\sigma}_i(\rho^{[i]})$ are independent random variables. Since $\tilde{\sigma}_i(\rho^{[i]})$ only depends on $\rho^{[i]}$ through $\rho_i^{[i]}$, it is enough to show that $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ and $\rho_i^{[i]}$ are independent. Moreover, as $\mathrm{sgn}(\dot{\rho}_i^{[i]})$ takes values in $\{-1, 1\}$, it is sufficient to check that

$$P(\mathrm{sgn}(\dot{\rho}_i^{[i]}) = 1,\ \rho_i^{[i]} \in I) = P(\mathrm{sgn}(\dot{\rho}_i^{[i]}) = 1)\ P(\rho_i^{[i]} \in I)$$

for all half-infinite intervals I.

Because C is a closed curve and $\left|\dot{\rho}_i^{[i]}\right|$ is constant, $P(\{\mathrm{sgn}(\dot{\rho}_i^{[i]}) = 1\}) = P(\{\mathrm{sgn}(\dot{\rho}_i^{[i]}) = -1\}) = \frac{1}{2}$, so the condition above reduces to

$$P(\mathrm{sgn}(\dot{\rho}_i^{[i]}) = 1,\ \rho_i^{[i]} \in I) = \tfrac{1}{2} P(\rho_i^{[i]} \in I).$$

But Lemma B.1 implies that

$$\begin{aligned} 0 &= \int_{(\rho^{[i]})^{-1}(I)} \dot{\rho}_i^{[i]}(\tau)\,d\tau \\ &= \int_{\{\tau:\mathrm{sgn}(\dot{\rho}_i^{[i]}(\tau))=1,\rho_i^{[i]}(\tau)\in I\}} v^{[i]}\,d\tau + \int_{\{\tau:\mathrm{sgn}(\dot{\rho}_i^{[i]}(\tau))=-1,\rho_i^{[i]}(\tau)\in I\}} (-v^{[i]})\,d\tau \\ &= v^{[i]}\Big(P\big(\mathrm{sgn}(\dot{\rho}_i^{[i]}) = 1, \rho_i^{[i]} \in I\big) - P\big(\mathrm{sgn}(\dot{\rho}_i^{[i]}) = -1, \rho_i^{[i]} \in I\big)\Big). \end{aligned}$$

Therefore, $P(\mathrm{sgn}(\dot{\rho}_i^{[i]}) = 1,\ \rho_i^{[i]} \in I) = P(\mathrm{sgn}(\dot{\rho}_i^{[i]}) = -1,\ \rho_i^{[i]} \in I) = \frac{1}{2} P(\rho_i^{[i]} \in I)$, completing the proof of independence.

To prove that independence implies separability, we establish the contrapositive. Suppose that for some index i, there exist π_i, π_{-i}, and $\hat{\pi}_{-i}$ such that $\tilde{\sigma}_i(\pi_i, \pi_{-i}) > s > \tilde{\sigma}_i(\pi_i, \hat{\pi}_{-i})$. Then since $\tilde{\sigma}$ is continuous, these inequalities remain true if we replace π_i with any $\hat{\pi}_i \in [\pi_i, \pi_i + \varepsilon]$, where ε is some small positive number. Now let C be the rectangle with vertices (π_i, π_{-i}), $(\pi_i + \varepsilon, \pi_{-i})$, $(\pi_i + \varepsilon, \hat{\pi}_{-i})$, and $(\pi_i, \hat{\pi}_{-i})$, and suppose that $\rho^{[i]}$ traverses these points in this same order. Then it is easily verified that $\tilde{\sigma}_i(\rho^{[i]}) > s$ if and only if $\text{sgn}(\dot{\rho}_i^{[i]}) = 1$, which implies that $\tilde{\sigma}_i(\rho^{[i]})$ and $\text{sgn}(\dot{\rho}_i^{[i]})$ are not independent. ■

The Proof of Theorem A.2

Since C is a closed curve and since $\left|\dot{\rho}_i^{[i]}\right|$ is constant, $P(\{\text{sgn}(\dot{\rho}_i^{[i]}) = 1\}) = P(\{\text{sgn}(\dot{\rho}_i^{[i]}) = -1\}) = \frac{1}{2}$. Hence, by construction, $P(\{\text{sgn}(\dot{\rho}_\iota^{[C]}) = 1\}) = P(\{\text{sgn}(\dot{\rho}_\iota^{[C]}) = -1\}) = \frac{1}{2}$, which implies that $E\text{sgn}(\dot{\rho}_\iota^{[C]}) = 0$, and so that

$$\begin{aligned}\text{cov}(\text{sgn}(\dot{\rho}_\iota^{[C]}),\ \tilde{\sigma}_\iota(\rho^{[C]})) &= E\text{sgn}(\dot{\rho}_\iota^{[C]})\,\tilde{\sigma}_\iota(\rho^{[C]}) - E\text{sgn}(\dot{\rho}_\iota^{[C]})\, E\,\tilde{\sigma}_\iota(\rho^{[C]}) \\ &= E\text{sgn}(\dot{\rho}_\iota^{[C]})\,\tilde{\sigma}_\iota(\rho^{[C]}) \\ &= \tfrac{1}{2} E[\tilde{\sigma}_\iota(\rho^{[C]}) \mid \text{sgn}(\dot{\rho}_\iota^{[C]}) = 1] - \tfrac{1}{2} E[\tilde{\sigma}_\iota(\rho^{[C]}) \mid \text{sgn}(\dot{\rho}_\iota^{[C]}) = -1].\end{aligned}$$

Observe that if $\text{sgn}(\dot{\rho}_\iota^{[C]})$ and $\tilde{\sigma}_\iota(\rho^{[C]})$ are uncorrelated, then $E[\tilde{\sigma}_\iota(\rho^{[C]}) \mid \text{sgn}(\dot{\rho}_\iota^{[C]}) = 1]$ and $E[\tilde{\sigma}_\iota(\rho^{[C]}) \mid \text{sgn}(\dot{\rho}_\iota^{[C]}) = -1]$ are equal, and so both equal $E\,\tilde{\sigma}_\iota(\rho^{[C]})$ as noted above.

Now since $\rho^{[C]} \equiv \rho^{[\iota]}(\tau)$ and since $\text{sgn}(\dot{\rho}_\iota^{[C]}) \equiv \text{sgn}(\dot{\rho}_\iota^{[\iota]}(\tau))$, we see that

$$\begin{aligned}\text{cov}(\text{sgn}(\dot{\rho}_\iota^{[C]}),\ \tilde{\sigma}_\iota(\rho^{[C]})) &= E\text{sgn}(\dot{\rho}_\iota^{[C]})\,\tilde{\sigma}_\iota(\rho^{[C]}) \\ &= \int_0^1 \text{sgn}(\dot{\rho}_{\iota(\omega)}^{[C]}(\omega))\,\tilde{\sigma}_{\iota(\omega)}(\rho^{[C]}(\omega))\,d\omega \\ &= \int_0^1 \text{sgn}(\dot{\rho}_{\iota(\omega)}^{[\iota(\omega)]}(\tau(\omega)))\,\tilde{\sigma}_{\iota(\omega)}(\rho^{[\iota(\omega)]}(\tau(\omega)))\,d\omega \\ &= V^{-1} \sum_{i:\,v^{[i]}>0} v^{[i]} \int_0^1 \text{sgn}(\dot{\rho}_i^{[i]}(\tau))\,\tilde{\sigma}_i(\rho^{[i]}(\tau))\,d\tau.\end{aligned}$$

Let γ: $[0, 1] \to C$ be some bijective, piecewise smooth parameterization of C. If we perform the change of variable implicitly defined by $\gamma(t) = \rho^{[i]}(\tau)$, then taking derivatives of the ith component of each side of the change of variable equation yields $\dot{\gamma}_i(t)\,dt = \dot{\rho}_i^{[i]}(\tau)\,d\tau = \text{sgn}(\dot{\rho}_i^{[i]}(\tau))\,v^{[i]}d\tau$. Therefore, since $\dot{\gamma}_j \equiv 0$ whenever $v^{[j]} = 0$, we find that

$$\text{cov}(\text{sgn}(\dot{\rho}_\iota^{[C]}),\ \tilde{\sigma}_\iota(\rho^{[C]})) = V^{-1} \sum_{i:\,v^{[i]}>0} \int_0^1 \tilde{\sigma}_i(\gamma(t))\,\dot{\gamma}_i(t)\,dt$$

$$= V^{-1}\oint \tilde{\sigma}(\gamma)\cdot d\gamma \, .$$

Thus, $\mathrm{sgn}(\dot{\rho}_t^{[C]})$ and $\tilde{\sigma}_t(\rho^{[C]})$ are uncorrelated for all $C \in \mathbf{C}$ if and only if the line integral of $\tilde{\sigma}$ over each $C \in \mathbf{C}$ evaluates to zero. The latter condition is necessary and sufficient for the integrability of $\tilde{\sigma}$. ■

References

Beckmann, M., C. B. McGuire, and C. B. Winsten (1956). *Studies in the Economics of Transportation*. New Haven: Yale University Press.

Benaïm, M., and J. W. Weibull (2003). "Deterministic Approximation of Stochastic Evolution in Games," *Econometrica* 71, 873-903.

Berger, U., and J. Hofbauer (2000). "The Nash Dynamics," mimeo, Universität Wien.

Binmore, K. J., and L. Samuelson (1999). "Evolutionary Drift and Equilibrium Selection," *Rev. Econ. Stud.* 66, 363-393.

Björnerstedt, J., and J. W. Weibull (1996). "Nash Equilibrium and Evolution by Imitation," in K. J. Arrow *et. al.* (eds.), *The Rational Foundations of Economic Behavior*. New York: St. Martin's Press.

Blackwell, D. (1956). "Controlled Random Walks." In *Proceedings of the International Conference of Mathematicians 1954, Vol. 3*, J. C. H. Gerretsen and J. De Groot, Eds. Amsterdam: North Holland.

Brown, G. W. (1951). "Iterative Solutions of Games by Fictitious Play," in *Activity Analysis of Production and Allocation*, T. C. Koopmans, Ed., New York: Wiley.

Brown, G. W. and J. von Neumann (1950). "Solutions of Games by Differential Equations," In H. W. Kuhn and A. W. Tucker, Eds., *Contributions to the Theory of Games I*. Annals of Mathematics Studies 24. Princeton: Princeton University Press.

Cressman, R., J. Garay, and J. Hofbauer (2001). "Evolutionary Stability Concepts for *N*-species Frequency-dependent Interactions," *J. Theor. Biol.* 211, 1-10.

Ely, J. C., and W. H. Sandholm (2004). "Evolution in Bayesian Games I: Theory," unpublished manuscript, Boston University and University of Wisconsin.

Friedman, D. (1991). "Evolutionary Games in Economics," *Econometrica* 59, 637-666.

Fudenberg, D., and D. M. Kreps (1993). "Learning Mixed Equilibria," *Games Econ. Behav.* 5 (1993), 320-367.

Fudenberg, D., and D. K. Levine (1998). *Theory of Learning in Games*. Cambridge: MIT Press.

Gaunersdorfer, A., and J. Hofbauer (1995). “Fictitious Play, Shapley Polygons, and the Replicator Equations,” *Games Econ. Behav.* 11, 279-303.

Gilboa, I., and A. Matsui (1991). "Social Stability and Equilibrium," *Econometrica* 59, 859-867.

Hannan, J. (1957). "Approximation to Bayes Risk in Repeated Play," In M. Dresher *et. al.*, Eds., *Contributions to the Theory of Games III*. Annals of Mathematics Studies 39. Princeton: Princeton University Press.

Hart, S., and A. Mas-Colell (2001). "A General Class of Adaptive Strategies," *J. Econ. Theory* 98, 26-54.

Hirsch, M. W., and S. Smale (1974). *Differential Equations, Dynamical Systems, and Linear Algebra*. San Diego: Academic Press.

Hofbauer, J. (1995). "Stability for the Best Response Dynamics," mimeo, Universität Wien.

Hofbauer, J. (2000). "From Nash and Brown to Maynard Smith: Equilibria, Dynamics, and ESS," *Selection* 1, 81-88.

Hofbauer, J., and E. Hopkins (2003). “Learning in Perturbed Asymmetric Games,” unpublished manuscript, University College London and University of Edinburgh.

Hofbauer, J., and W. H. Sandholm (2002). "On the Global Convergence of Stochastic Fictitious Play," *Econometrica* 70, 2265-2294 .

Hofbauer, J., and W. H. Sandholm (2003). "Evolution in Games with Randomly Disturbed Payoffs," unpublished manuscript, University College London and University of Wisconsin.

Hofbauer, J., and W. H. Sandholm (2004). “Equilibrium and Evolution in Stable Games,” mimeo, University College London and University of Wisconsin.

Hofbauer, J., and K. Sigmund (1988). *Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge University Press.

Hofbauer, J., and J. Weibull (1996). “Evolutionary Selection against Dominated Strategies,” *J. Econ. Theory* 71, 558-573.

Hopkins, E. (1999). “A Note on Best Response Dynamics,” *Games Econ. Behav.* 29, 138-150.

Horn, R. A., and C. R. Johnson (1985). *Matrix Analysis*. Cambridge: Cambridge University Press.

Matsui, A. (1992). "Best Response Dynamics and Socially Stable Strategies," *J. Econ. Theory* 57, 343-362.

Monderer, D., and L. Shapley (1996). "Potential Games," *Games Econ. Behav.* 14, 124-143.

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

Nash, J. F. (1951). "Non-cooperative Games," *Ann. Math.* 54, 287-295.

Ritzberger, K. and J. W. Weibull (1995). "Evolutionary Selection in Normal Form Games," *Econometrica* 63, 1371-1399.

Rosenthal, R. W. (1973). "A Class of Games Possessing Pure Strategy Nash Equilibria," *Int. J. Game Theory*, **2**, 65-67.

Samuelson, L., and J. Zhang (1992). "Evolutionary Stability in Asymmetric Games," *J. Econ. Theory* 57, 363-391.

Sandholm, W. H. (2001). "Potential Games with Continuous Player Sets," *J. Econ. Theory* 97, 81-108.

Sandholm, W. H. (2002). "Evolutionary Implementation and Congestion Pricing," *Rev. Econ. Stud.* 69, 667-689.

Sandholm, W. H. (2003a). "Evolution and Equilibrium under Inexact Information," *Games Econ. Behav.* 44, 343-378.

Sandholm, W. H. (2003b). "Negative Externalities and Evolutionary Implementation," mimeo, University of Wisconsin.

Schlag, K. H. (1998). "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits," *J. Econ. Theory* 78, 130-156.

Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Cambridge: Harvard University Press.

Swinkels, J. M. (1993). "Adjustment Dynamics and Rational Play in Games," *Games Econ. Behav.* 5, 455-484.

Taylor, P. D., and L. Jonker (1978). "Evolutionarily Stable Strategies and Game Dynamics," *Math. Biosci.* 40, 145-156.

Vickrey, W. S. (1969). "Congestion Theory and Transport Investment," *Amer. Econ. Rev. Papers and Proceedings*, 59, 251-260.

Walters, A. A. (1961). "The Theory and Measurement of Private and Social Cost of Highway Congestion," *Econometrica* 29, 676-699.

Weibull, J. W. (1995). *Evolutionary Game Theory*. Cambridge: MIT Press.

Weibull, J. W. (1996). In "The Work of John Nash in Game Theory," *J. Econ. Theory* 69, 153-185.