The Projection Dynamic and the Geometry of Population Games*

Ratul Lahkar† and William H. Sandholm‡

February 1, 2008

Abstract

The projection dynamic is an evolutionary dynamic for population games. It is derived from a model of individual choice in which agents abandon their current strategies at rates inversely proportional to the strategies' current levels of use. The dynamic admits a simple geometric definition, its rest points coincide with the Nash equilibria of the underlying game, and it converges globally to Nash equilibrium in potential games and in stable games. *JEL classification*: C72, C73.

1. Introduction

Population games describe strategic interactions in settings with large numbers of agents. In such environments, the knowledge assumptions that underlie the direct assumption of equilibrium play are quite strong, making it desirable to provide explicitly dynamic models of behavior. To accomplish this, one can posit a primitive specification of how individual agents respond to their current incentives, and from this specification derive a stochastic process that describes the evolution of the population's aggregate behavior. By computing the expected increments of this process, one can derive a deterministic differential equation—the *mean dynamic*—that closely captures the evolution of aggregate behavior over finite time spans.[1]

*We thank Emin Dokumacı, two anonymous referees, an anonymous Associate Editor, and many seminar audiences for helpful comments. Financial support from NSF Grants SES-0092145 and SES-0617753 is gratefully acknowledged.

†Department of Mathematics and ELSE, University College London, Gower Street, London WC1E6BT, UK. e-mail: r.lahkar@ucl.ac.uk, website: http://rlahkar.googlepages.com.

‡Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA. e-mail: whs@ssc.wisc.edu, website: http://www.ssc.wisc.edu/~whs.

[1]For formal statements of this idea, see Benaïm and Weibull (2003) and Sandholm (2003).

When a deterministic evolutionary dynamic is derived in this fashion, the form of the dynamic depends on the nature of the revision protocol individual agents follow when they consider switching strategies. When agents imitate successful opponents, aggregate behavior is described by the replicator dynamic, or by other related dynamics.[2] If instead agents directly assess the payoffs of each alternative strategy, one obtains *direct* (or *innovative*) *dynamics*; examples of these include best response and perturbed best response dynamics, the Brown-von Neumann-Nash (BNN) dynamic, and the Smith dynamic.[3] After specifying the evolutionary dynamic of interest, one can turn to questions of prediction, both by establishing connections between the dynamic's rest points and the Nash equilibria of game at hand, and by proving convergence results for games whose payoffs satisfy appropriate structural assumptions.[4]

In this paper, we study the *projection dynamic*, an evolutionary game dynamic introduced in the transportation science literature by Nagurney and Zhang (1997).[5] We provide microfoundations for the projection dynamic by constructing a revision protocol that generates this dynamic as its mean dynamic. The protocol relies on the direct assessment of payoffs of alternative strategies. What is most novel about the protocol is that the rate at which an agent actively reconsiders his choice of strategy depends directly on the popularity of his current strategy: more precisely, each agent considers abandoning his current strategy at a rate inversely proportional to the strategy's current level of use. One can therefore designate the projection dynamic as capturing "revision driven by insecurity", as it describes the behavior of agents who are especially uncomfortable choosing strategies not used by many others.

While the projection dynamic is discontinuous at the boundary of the set of population states, it nevertheless exhibits desirable game-theoretic properties. In particular, its rest points are precisely the Nash equilibria of the underlying game, and it converges to equilibrium from all initial conditions in potential games and in stable games.

Given our description of the revision protocol that generates the projection dynamic, it seems natural to expect that the behavior of the projection dynamic should be closest to that of other direct dynamics—especially the BNN and Smith dynamics, since these also require active reevaluation of strategies to occur at a variable rate.[6] Considered from

[2]See Weibull (1995), Björnerstedt and Weibull (1996), Hofbauer (1995), and Schlag (1998).

[3]See Gilboa and Matsui (1991), Fudenberg and Levine (1998), Hofbauer and Sandholm (2007), Brown and von Neumann (1950), Smith (1984), and Sandholm (2005, 2006).

[4]See Hofbauer (2000), Sandholm (2001), and Hofbauer and Sandholm (2008).

[5]Dupuis and Nagurney (1993) and Nagurney and Zhang (1996) use differential equations defined in terms of projections to study a variety of economic applications outside the context of population games.

[6]In contrast, best response and perturbed best response dynamics have agents reconsider their strategies at a fixed rate—see Section 4.2.

this point of view, it is surprising that the dynamic with the closest links to the projection dynamic is actually the replicator dynamic (Taylor and Jonker (1978)), which is founded not on direct evaluation of payoffs, but on imitation of opponents. The many connections between these two dynamics are explored in a companion paper, Sandholm et al. (2008).

The projection dynamic not only has appealing game-theoretic foundations and convergence properties, but also admits a strikingly simple geometric definition. To best avail ourselves of this fact, we begin the paper by presenting population games from a geometric point of view. In a population game, there are one or more continuous populations of agents, agents in each population choose strategies from a finite set, and each agent's payoffs depend on his own choice and on the other agents' aggregate behavior.[7] Payoffs are described using a continuous map $F : X \to \mathbf{R}^n$, where n is the total number of strategies in all populations, and $X \subset \mathbf{R}^n$ is the set of social states;[8] component $F_i^p(x)$ of the vector $F(x) \in \mathbf{R}^n$ represents the payoff to population p's strategy i at social state x. With this background in place, we define the projection dynamic for population game F by the differential equation $\dot{x} = V^F(x)$, where $V^F(x) \in \mathbf{R}^n$ is the best approximation of the payoff vector $F(x)$ by a feasible direction of motion from the current state x through the state space X.[9] This geometric definition enables us to establish the projection dynamic's game-theoretic properties in a straightforward way.

Still, since the projection dynamic is discontinuous at the boundary of X, more fundamental properties of the dynamic—namely, the existence and uniqueness of its solutions—do not follow from standard results. Fortunately, theorems due to Henry (1973) and Aubin and Cellina (1984) (see also Dupuis and Ishii (1991) and Dupuis and Nagurney (1993)) imply that solutions to the projection dynamic exist, are unique, and are Lipschitz continuous in their initial conditions. But by virtue of its discontinuities, solutions to the projection dynamic have some properties quite different from those of standard dynamics: its solutions can merge in finite time, and can enter and exit the boundary of X repeatedly as time passes.

Section 2 introduces population games, and Section 3 investigates their geometric properties. Section 4 defines the projection dynamic and describes its microfoundations, comparing them to those of other fundamental dynamics from evolutionary game theory. Section 5 establishes the existence and uniqueness of solution trajectories, and uses examples to illustrate some of their novel properties. Section 6 provides the game-theoretic

[7] Notice that this framework differs from that considered by Rosen (1965) and Gabay and Moulin (1980), in which a finite number of players select from continuous sets of strategies.

[8] X can also be described as the set of empirical strategy distributions. In a single-population game, X is the unit simplex in $\mathbf{R}^n$; if there are multiple populations, X is a product of simplices.

[9] Formally, $V^F(x)$ is the closest point projection of $F(x)$ onto $TX(x)$, the tangent cone of the set X at point x; see Section 4.1.

properties of the dynamic, including Nash stationarity and global convergence in potential games and in stable games. Proofs omitted from the text are provided in an appendix.

2. Population Games

2.1 Definitions

Let $\mathcal{P} = \{1, \ldots, p\}$ be a *society* consisting of $p \geq 1$ *populations* of agents. Agents in population p form a continuum of *mass* $m^p > 0$. Agents in population p choose strategies from the set $S^p = \{1, \ldots, n^p\}$. The total number of strategies in all populations is $n = \sum_{p \in \mathcal{P}} n^p$.

During play, each agent in population p selects a (pure) strategy from S^p. The set of *population states* (or *strategy distributions*) for population p is $X^p = \{x^p \in \mathbf{R}^{n^p}_+ : \sum_{i \in S^p} x^p_i = m^p\}$. The scalar $x^p_i \in \mathbf{R}_+$ represents the mass of players in population p choosing strategy $i \in S^p$. Elements of $X = \prod_{p \in \mathcal{P}} X^p = \{x = (x^1, \ldots, x^p) : x^p \in X^p\}$, the set of *social states*, describe behavior in all p populations at once. When there is only one population, we omit the superscript p from our notation and assume that the population's mass is 1.

We generally take the sets of populations and strategies as fixed and identify a game with its payoff function. A *payoff function* $F : X \to \mathbf{R}^n$ is a Lipschitz continuous map that assigns each social state a vector of payoffs, one for each strategy in each population. $F^p_i : X \to \mathbf{R}$ denotes the payoff function for strategy $i \in S^p$, while $F^p : X \to \mathbf{R}^{n^p}$ denotes the payoff functions for all strategies in S^p. Similar notational conventions are used throughout the paper.

Social state $x \in X$ is a *Nash equilibrium* of F if all agents in all populations play best responses. Formally, $x \in NE(F)$ if

$$x^p_i > 0 \Rightarrow i \in \underset{j \in S^p}{\operatorname{argmax}}\, F^p_j(x) \text{ for all } i \in S^p \text{ and } p \in \mathcal{P}.$$

2.2 Examples

Example 2.1. Random matching in symmetric normal form games. An n-strategy symmetric normal form game is defined by a payoff matrix $A \in \mathbf{R}^{n \times n}$. A_{ij} is the payoff a player obtains when he chooses strategy i and his opponent chooses strategy j; this payoff does not depend on whether the player in question is called player 1 or player 2. When a population of agents are randomly matched to play this game, the (expected) payoff to strategy i at population state is x is $F_i(x) = \sum_{j \in S} A_{ij} x_j$; hence, the population game associated with A is the linear game $F(x) = Ax$. §

Example 2.2. Congestion games. Congestion games provide a basic model of multilateral externalities. For concreteness, we describe these games using the context of highway network congestion.[10] Consider a collection of towns is connected by a network of *links* L. For each ordered pair $p \in \mathcal{P}$ of towns, there is a population of agents, each of whom needs to commute from the first town in the pair (where he lives) to the second (where he works). To accomplish this, the agent must choose a *path*: each path $i \in S^p$ consists of a set of links $L_i^p \subseteq L$ connecting the towns in pair p. An agent's payoff from choosing path i is the negation of the delay on this path. The delay on a path is the sum of the delays on its links, and the delay on link is a function of the number of agents using that link. Formally,

$$F_i^p(x) = -\sum_{l \in L_i^p} c_l(u_l(x)), \text{ where } u_l(x) = \sum_{p \in \mathcal{P}} \sum_{i \in S^p :\, l \in L_i^p} x_i^p.$$

When congestion games are used to model highway networks or other environments in which externalities are negative, the cost functions c_l are increasing in the utilization levels u_l. But one can also use congestion games to capture positive externalities by assuming that the cost functions c_l are decreasing. §

3. The Geometry of Population Games

In this section, we present population games from a geometric point of view, noting the key roles played by projections onto tangent spaces and tangent cones.

3.1 Tangent Spaces and Orthogonal Projections for Population Games

3.1.1 Definitions

The *tangent space* of X^p, denoted TX^p, is the smallest subspace of $\mathbf{R}^{n^p}$ that contains all vectors describing motions between points in X^p. In other words, if $x^p, y^p \in X^p$, then $y^p - x^p \in TX^p$, and TX^p is the span of all vectors of this form. Evidently, TX^p contains exactly those vectors in $\mathbf{R}^{n^p}$ whose components sum to zero: $TX^p = \{z^p \in \mathbf{R}^{n^p} : \sum_{i \in S^p} z_i^p = 0\}$. All directions of motion between points in the set of social states X are contained in its tangent space, the product set $TX = \prod_{p \in \mathcal{P}} TX^p$.

Projections—in particular, projections of payoff vectors onto sets of feasible directions of motion—play a central role throughout this paper. When the target space of a projection is a linear subspace, like the tangent spaces TX^p and TX, the appropriate notion

[10] See Beckmann et al. (1956), Rosenthal (1973), Monderer and Shapley (1996), and Sandholm (2001).

of projection is *orthogonal projection*. The geometric definition of orthogonal projection is well-known; algebraically, orthogonal projections are the linear operations represented by symmetric idempotent matrices.

We represent the orthogonal projection onto the subspace $TX^p \subset \mathbf{R}^{n^p}$ by the matrix $\Phi \in \mathbf{R}^{n^p \times n^p}$, defined by $\Phi = I - \frac{1}{n^p}\mathbf{1}\mathbf{1}'$. Here $\mathbf{1} = (1, ..., 1)'$ is the vector of ones, so $\frac{1}{n^p}\mathbf{1}\mathbf{1}'$ is the matrix whose entries are all $\frac{1}{n^p}$.

The projection Φ has a simple interpretation. If v^p is a payoff vector in $\mathbf{R}^{n^p}$, the projection of v^p onto TX^p is

$$\Phi v^p = v^p - \frac{1}{n^p}\mathbf{1}\mathbf{1}' v^p = v^p - \mathbf{1}\left(\frac{1}{n^p}\sum_{k \in S^p} v^p_k\right).$$

Thus, the ith component of the vector Φv^p is the difference between the actual payoff to strategy i and the unweighted average payoff of all strategies in S^p. Put differently Φv^p discards information about the absolute level of payoffs under v^p while retaining information about relative payoffs. This is interesting from a game-theoretic point of view, since incentives, and hence Nash equilibria, only depend on payoff differences. Therefore, when incentives (as opposed to, e.g., efficiency) are our main concern, the projected payoff vectors Φv^p are sufficient statistics for the actual payoff vectors v^p.

Since $TX = \prod_{p \in \mathcal{P}} TX^p$ is a product set, the orthogonal projection onto TX is represented by a block diagonal matrix, $\mathbf{\Phi} = \text{diag}(\Phi, \dots, \Phi) \in \mathbf{R}^{n \times n}$. If we apply $\mathbf{\Phi}$ to the payoff vector $v = (v^1, \dots, v^p)$, the resulting vector $\mathbf{\Phi} v = (\Phi v^1, \dots, \Phi v^p)$ lists the relative payoffs in each population.

3.1.2 Drawing Population Games

Low-dimensional population games can be presented in pictures. Consider these two single-population, two-strategy games, a coordination game and a Hawk-Dove game:

$$F(x) = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \end{pmatrix} \quad \text{and} \quad F(x) = \begin{pmatrix} -1 & 2 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} x_H \\ x_D \end{pmatrix} = \begin{pmatrix} 2x_D - x_H \\ x_D \end{pmatrix}.$$

Figures 1(i) and 1(ii) present these two games, along with their projected versions,

$$\Phi F(x) = \Phi\begin{pmatrix} x_1 \\ 2x_2 \end{pmatrix} = \begin{pmatrix} \frac{1}{2}x_1 - x_2 \\ -\frac{1}{2}x_1 + x_2 \end{pmatrix} \quad \text{and} \quad \Phi F(x) = \Phi\begin{pmatrix} 2x_D - x_H \\ x_D \end{pmatrix} = \begin{pmatrix} \frac{1}{2}(x_D - x_H) \\ \frac{1}{2}(x_H - x_D) \end{pmatrix}.$$

The figures are synchronized with the payoff matrix by using the vertical coordinate to represent the mass on the first strategy and the horizontal coordinate to represent the

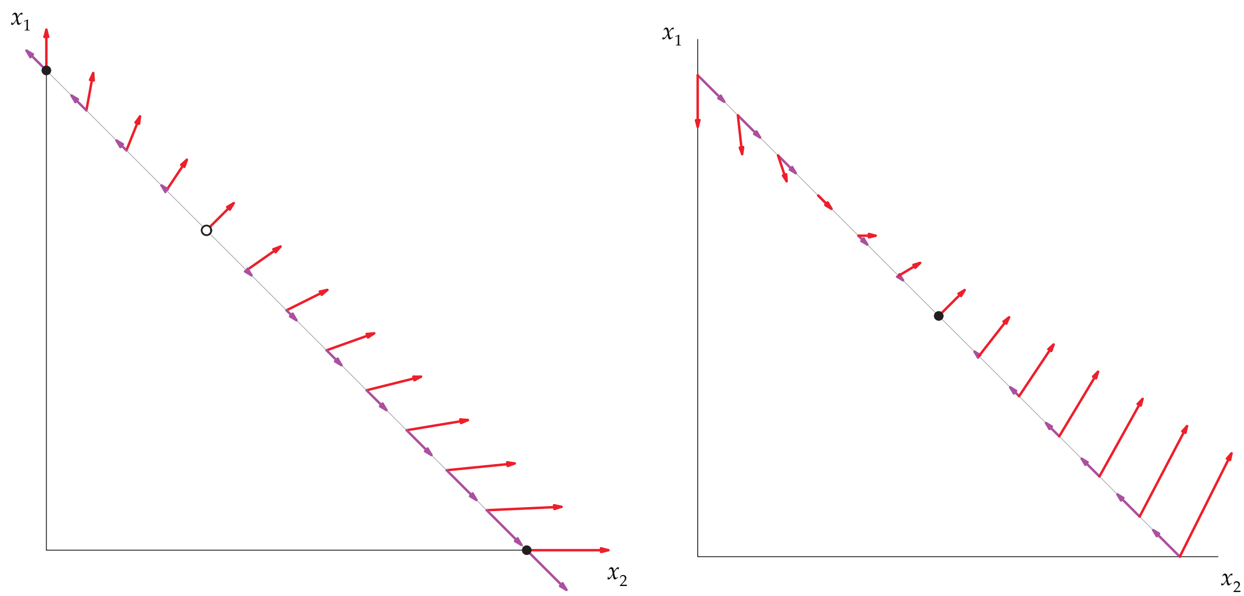

(i) Coordination (ii) Hawk-Dove

Figure 1: The payoff vector field $F(\cdot)$ and its projection $\Phi F(\cdot)$ in two two-strategy games.

mass on the second strategy. At various states x, we draw (scaled down) versions of the corresponding payoff vectors $F(x)$ and projected payoff vectors $\Phi F(x)$.

Let us focus first on Figure 1(i), representing the coordination game. At the pure state $e_1 = (1, 0)$, at which all agents play strategy 1, the payoffs to the two strategies are $F_1(e_1) = 1$ and $F_2(e_1) = 0$, so the payoff vector $F(e_1)$ points directly *upward*. At the interior Nash equilibrium $x^* = (x_1^*, x_2^*) = (\frac{2}{3}, \frac{1}{3})$, each strategy earns a payoff of $\frac{2}{3}$; the arrow representing payoff vector $F(x^*) = (\frac{2}{3}, \frac{2}{3})$ is drawn at a right angle to the simplex, implying that the projected payoff vector $\Phi F(x^*) = (0, 0)$ is null. Similar logic explains how the payoff vectors are drawn at other states, and how Figure 1(ii) is constructed as well.

These diagrams help us visualize the incentives faced by agents playing these games. In the coordination game, the payoff vectors "push outward" toward the two axes, reflecting an incentive structure that drives the population toward the two pure Nash equilibria. In contrast, payoff vectors in the Hawk-Dove game "push inward", away from the axes, reflecting forces leading the population toward the interior Nash equilibrium $x^* = (\frac{1}{2}, \frac{1}{2})$.

Representing three-strategy games in two-dimensional pictures requires more care. Figure 2 presents a "three-dimensional" picture of the simplex X situated in its ambient plane $\text{aff}(X) = \{x \in \mathbf{R}^3 : \sum_{i \in S} x_i = 1\}$ in $\mathbf{R}^3$, known as the *affine hull* of X. When we draw the simplex on a sheet of paper as an equilateral triangle, the paper represents a portion of this plane. Each payoff vector $F(x)$ in a three-strategy game is an element of $\mathbf{R}^3$, but the

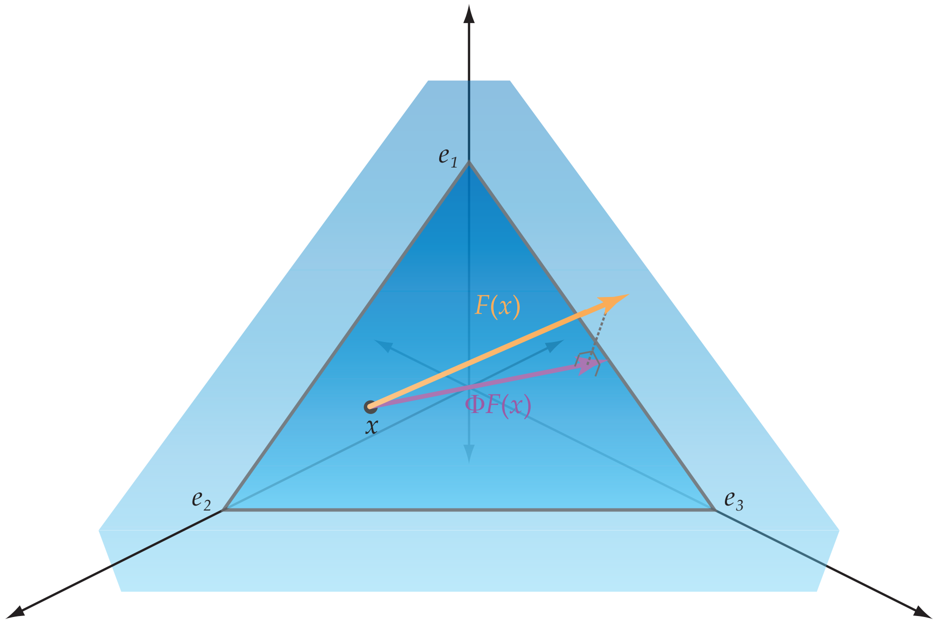

Figure 2: The simplex $X \subset \mathbf{R}^3$ and its affine hull.

projected payoff vector $\Phi F(x)$ lies in the two-dimensional tangent space TX. If the vector $\Phi F(x)$ is drawn as an arrow rooted at state x, then by construction this arrow will lie in the plane aff(X) (see Figure 2).

In Figures 3(i) and 3(ii), we use this approach to draw pictures of two single-population games with three strategies: a coordination game and a standard Rock-Paper-Scissors game:

$$F(x) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_3 \end{pmatrix} \quad \text{and} \quad F(x) = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \end{pmatrix} = \begin{pmatrix} x_S - x_P \\ x_R - x_S \\ x_P - x_R \end{pmatrix}.$$

Of course, these drawings are actually of the projected payoffs

$$\Phi F(x) = \Phi \begin{pmatrix} x_1 \\ 2x_2 \\ 3x_3 \end{pmatrix} = \begin{pmatrix} \frac{1}{3}(2x_1 - 2x_2 - 3x_3) \\ \frac{1}{3}(-x_1 + 4x_2 - 3x_3) \\ \frac{1}{3}(-x_1 - 2x_2 + 6x_3) \end{pmatrix} \quad \text{and} \quad \Phi F(x) = \Phi \begin{pmatrix} x_S - x_P \\ x_R - x_S \\ x_P - x_R \end{pmatrix} = \begin{pmatrix} x_S - x_P \\ x_R - x_S \\ x_P - x_R \end{pmatrix}.$$

But since standard RPS is symmetric zero-sum (i.e., since its payoff matrix is skew-symmetric), the original and projected payoffs for this game are identical.

Much like Figure 1(i), Figure 3(i) shows that in the three-strategy coordination game, the projected payoff vectors push outward toward the extreme points of the simplex.

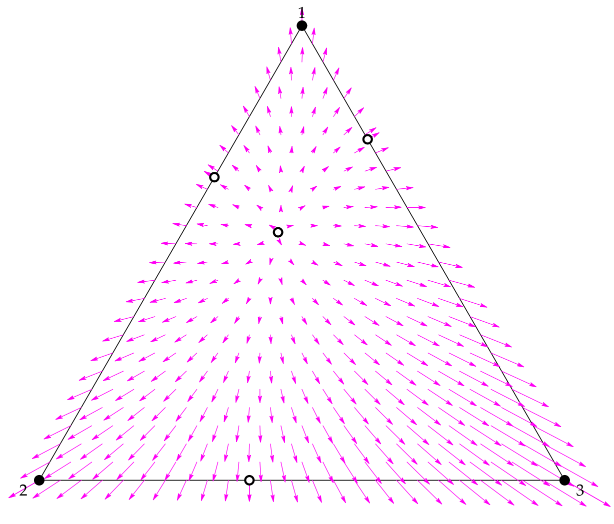

(i) Coordination (ii) Standard Rock-Paper-Scissors

Figure 3: The projected payoff vector field $\Phi F(\cdot)$ in two three-strategy games.

Figure 3(ii) exhibits a property that is only possible when there are three or more strategies: instead of heading toward Nash equilibria, the vectors in this figure describe cycle around state $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, the unique Nash equilibrium of standard RPS.

3.2 Tangent Cones and Normal Cones for Convex Sets

All vectors $z \in TX$ represent feasible motions from social states in the (relative) interior of X. To represent a feasible motion from a boundary state, a vector cannot cause an unused strategy to lose mass. To describe sets of such motions, we introduce the notion of a tangent cone to a convex set. Defining closest point projections onto these sets will enable us to define the projection dynamic. For further background on these topics, see, e.g., Hiriart-Urruty and Lemaréchal (2001).

3.2.1 Definitions

The set $K \subseteq \mathbf{R}^n$ is a *cone* if whenever it contains the vector z, it contains each vector αz with $\alpha > 0$. If K is a closed convex cone, its *polar cone* K° is a new closed convex cone:

$$K^\circ = \{y \in \mathbf{R}^n : y'z \leq 0 \text{ for all } z \in K\}.$$

In words, K° contains all vectors that form a weakly obtuse angle with each vector in K.

If the closed convex cone K is symmetric, in the sense that $K = -K$, then K is actually a linear subspace of $\mathbf{R}^n$; in this case, $K^\circ = K^\perp$, the orthogonal complement of K. More generally, polarity defines an *involution* on the set of closed convex cones: that is, $(K^\circ)^\circ = K$ for any closed convex cone K.

If $C \subset \mathbf{R}^n$ is a closed convex set, then the *tangent cone* of C at state $x \in C$, denoted $TC(x)$, is the closed convex cone

$$TC(x) = \mathrm{cl}\Big(\{z \in \mathbf{R}^n : z = \alpha(y - x) \text{ for some } y \in C \text{ and some } \alpha \geq 0\}\Big).$$

If $C \subset \mathbf{R}^n$ is a *polytope* (i.e., the convex hull of a finite number of points), then the closure operation is redundant. In this case, $TC(x)$ is the set of directions of motion from x that initially remain in C; more generally, $TC(x)$ also contains the limits of such directions. If x is in the (relative) interior of C, then $TC(x)$ is just TC, the tangent space of C; otherwise, $TC(x)$ is a strict subset of TC.

The *normal cone* of C at x is the polar of the tangent cone of C at x: that is, $NC(x) = TC(x)^\circ$. By definition, $NC(x)$ is a closed convex cone, and it contains every vector that forms a weakly obtuse angle with every feasible displacement vector at x.

3.2.2 Normal Cones and Nash Equilibria

When X is the set of social states of a population game, each tangent cone $TX(x)$ contains the feasible directions of motion from social state $x \in X$. In multipopulation cases $TX(x)$ can be decomposed population by population.[11]

In Figures 4(i) and 4(ii), we sketch examples of tangent cones and normal cones when X is the state space of a two-strategy game and of a three-strategy game. Since Figure 4(ii) is two-dimensional, with the sheet of paper representing the plane aff(X), the figure actually displays the projected normal cones $\Phi NX(x)$.

At first glance, normal cones might appear to be less relevant to game theory than tangent cones. Theorem 3.1 shows that this impression is false, and provides us with a simple geometric description of Nash equilibria of population games. Versions of this result can be found in the literature on variational inequalities—see Harker and Pang (1990) and Nagurney and Zhang (1996).

Theorem 3.1. *Let F be a population game. Then $x \in NE(F)$ if and only if $F(x) \in NX(x)$.*

Proof. $x \in NE(F) \Leftrightarrow [x_i^p > 0 \Rightarrow F_i^p(x) \geq F_j^p(x)]$ for all $i, j \in S^p, p \in \mathcal{P}$

[11] That is: since $X = \prod_{p \in \mathcal{P}} X^p$ is a product set, $TX(x) = \prod_{p \in \mathcal{P}} TX^p(x^p)$ is a product set as well. Similarly, we have that $NX(x) = \prod_{p \in \mathcal{P}} NX^p(x^p)$; this fact is used in the proof of Theorem 3.1 below.

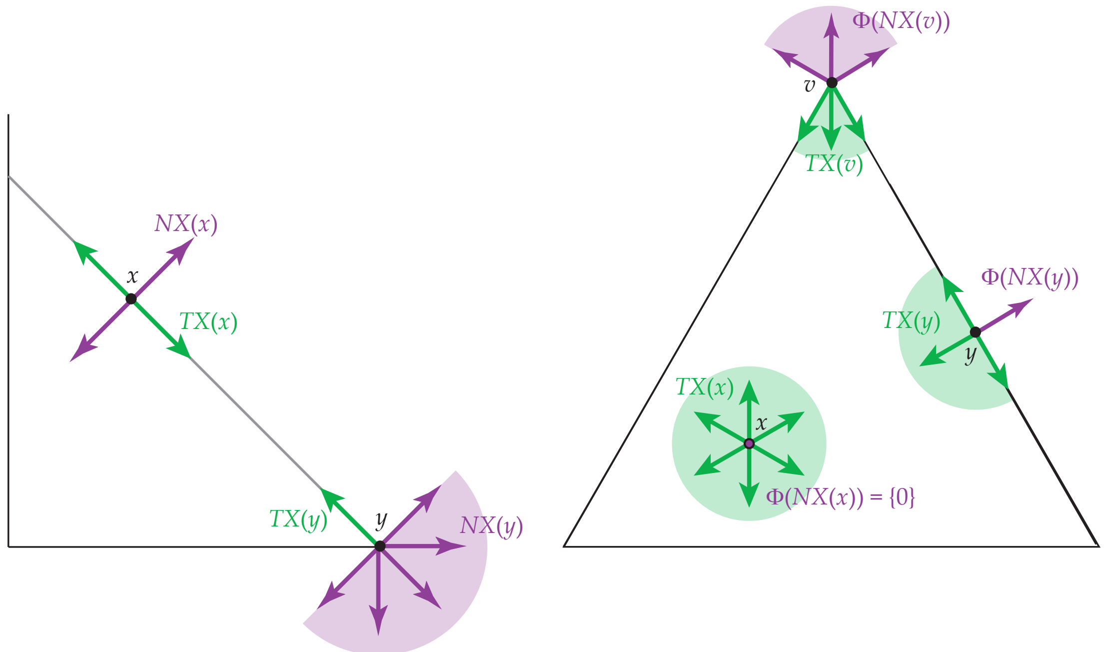

(i) $n = 2$ (ii) $n = 3$

Figure 4: Tangent cones and (projected) normal cones in two- and three-strategy games.

$$
\begin{aligned}
&\Leftrightarrow (x^p)'F^p(x) \geq (y^p)'F^p(x) \text{ for all } y^p \in X^p, p \in \mathcal{P} \\
&\Leftrightarrow (y^p - x^p)'F^p(x) \leq 0 \text{ for all } y^p \in X^p, p \in \mathcal{P} \\
&\Leftrightarrow (z^p)'F^p(x) \leq 0 \text{ for all } z^p \in TX^p(x), p \in \mathcal{P} \\
&\Leftrightarrow F^p(x) \in NX^p(x^p) \text{ for all } p \in \mathcal{P} \\
&\Leftrightarrow F(x) \in NX(x). \blacksquare
\end{aligned}
$$

In the Figures 1 and 3 above, Nash equilibria are marked with dots. In the two-strategy games, the Nash equilibria are those states x at which the payoff vector $F(x)$ lies in the normal cone $NX(x)$, as Theorem 3.1 shows. In the three-strategy games, the Nash equilibria are those states x at which the projected payoff vector $\Phi F(x)$ lies in the projected normal cone $\Phi NX(x)$; this is an easy corollary of Theorem 3.1.

3.3 Closest Point Projections onto Convex Cones

3.3.1 Definition and Characterization

To define the projection dynamic on the boundary of X, we need to introduce projections onto convex cones. If $K \subset \mathbf{R}^n$ is a closed convex cone, then the *closest point projection*

$\Pi_K : \mathbf{R}^n \to K$ is defined by

$$\Pi_K(v) = \operatorname*{argmin}_{z \in K} \left| z - v \right|.$$

If K is a subspace of $\mathbf{R}^n$ (i.e., if $K = -K$), then the closest point projection onto K is simply the orthogonal projection onto K.

The fundamental result about projections onto closed convex cones is the *Moreau Decomposition Theorem*, which generalizes the notion of orthogonal decomposition to the "one-sided" world of convex cones. In words, this theorem tells us that the projections of the vector v onto K and K° are the unique vectors in K and K° that are orthogonal to one another and that sum to v. For a proof, see Hiriart-Urruty and Lemaréchal (2001)

Theorem 3.2 (The Moreau Decomposition Theorem)**.** *Let $K \subseteq \mathbf{R}^n$ and $K^\circ \subseteq \mathbf{R}^n$ be a closed convex cone and its polar cone, and let $v \in \mathbf{R}^n$. Then the following are equivalent:*

(i) $v_K = \Pi_K(v)$ *and* $v_{K^\circ} = \Pi_{K^\circ}(v)$.

(ii) $v_K \in K, v_{K^\circ} \in K^\circ, (v_K)' v_{K^\circ} = 0,$ *and* $v = v_K + v_{K^\circ}$.

3.3.2 Projecting Payoff Vectors onto Tangent Cones of X

In Figures 5(i) and 5(ii), we draw the projected payoff functions $V(\cdot) = \Pi_{TX(\cdot)}F(\cdot)$ for our two three-strategy games. In these figures, each payoff vector $F(x)$ is represented by $\Pi_{TX(x)}(F(x))$, the best approximation by a feasible direction of motion from x. Evidently, the states x at which the projected payoff vector is null are precisely the Nash equilibria of the underlying game. That this is true in general is an immediate consequence of Theorems 3.1 and 3.2:

Corollary 3.3. *Let F be a population game. Then $x \in NE(F)$ if and only if* $\Pi_{TX(x)}(F(x)) = \mathbf{0}$.

3.3.3 An Explicit Formula for $\Pi_{TX(x)}(v)$

In general, explicitly computing a closest point projection onto a convex cone requires solving a quadratic program. The next result, Theorem 3.4, shows that the projection onto $TX^p(x^p)$ admits a simple explicit description. Since the projection onto $TX(x)$ can be decomposed population by population, this formula is sufficient to determine $\Pi_{TX(x)}(v)$.[12]

[12]More explicitly: since $TX(x) = TX^1(x^1) \times \cdots \times TX^p(x^p)$ is a product set, we have that $\Pi_{TX(x)}(v) = \Pi_{TX^1(x^1)}(v^1) \times \cdots \times \Pi_{TX^p(x^p)}(v^p)$.

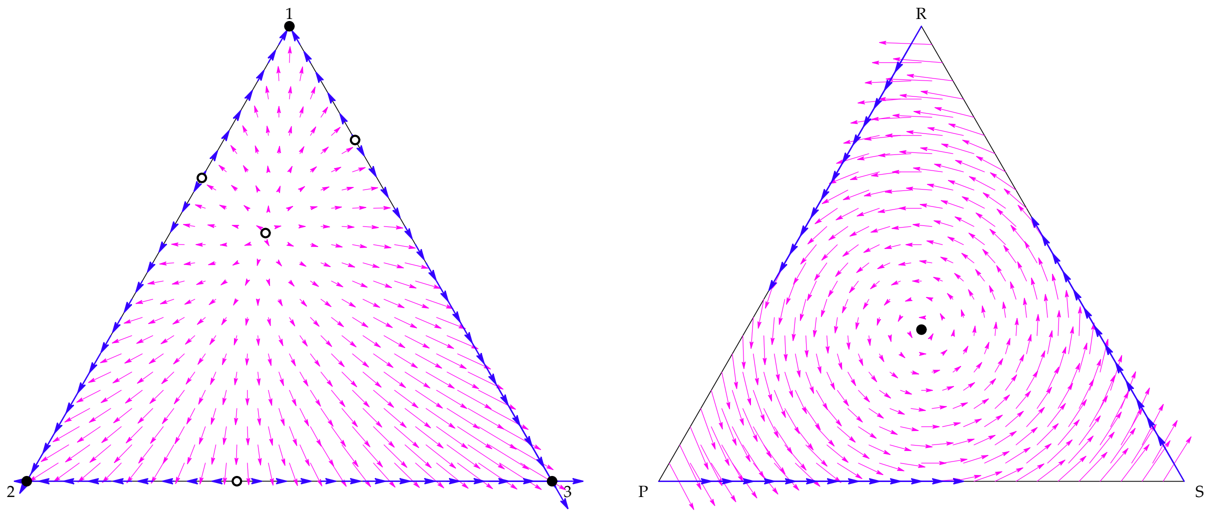

(i) Coordination (ii) Standard Rock-Paper-Scissors

Figure 5: Vector fields $V(\cdot) = \Pi_{TX(\cdot)}F(\cdot)$ obtained by projecting payoffs onto tangent cones.

Theorem 3.4. *The projection $\Pi_{TX^p(x^p)}(v^p)$ can be expressed as follows:*

$$(\Pi_{TX^p(x^p)}(v^p))_i = \begin{cases} v_i^p - \frac{1}{\#\mathcal{S}^p(v^p,x^p)} \sum_{j\in\mathcal{S}^p(v^p,x^p)} v_j^p & \textit{if } i \in \mathcal{S}^p(v^p,x^p), \\ 0 & \textit{otherwise.} \end{cases}$$

Here, the set $\mathcal{S}^p(v^p, x^p) \subseteq S^p$ contains all strategies in $\text{support}(x^p)$, *along with any subset of* $S^p - \text{support}(x^p)$ *that maximizes the average* $\frac{1}{\#\mathcal{S}^p(v^p,x^p)} \sum_{j\in\mathcal{S}^p(v^p,x^p)} v_j^p$.

The proof of Theorem 3.4 can be found in the appendix.

To explain Theorem 3.4, let us avoid superscripts by focusing on the single population case. Imagine that $v \in \mathbf{R}^n$ is the vector of payoffs earned by strategies in $S = \{1, \ldots, n\}$ at state $x \in X$. When x is in the interior of X, the tangent cone $TX(x)$ is the just the subspace TX; therefore, the closest point projection onto $TX(x)$ is the orthogonal projection $\Phi = I - \frac{1}{n}\mathbf{11}'$ from Section 3.1, which subtracts the average payoff under v from each component of v:

$$(1) \qquad (\Pi_{TX(x)}(v))_i = (\Phi v)_i = v_i - \frac{1}{n}\sum_{j\in S} v_j.$$

If instead there is exactly one unused strategy at state x, say strategy n, then the tangent cone $TX(x)$ consists of vectors in TX whose nth component is nonnegative. In this case, if strategy n earns an above average payoff, then $(\Phi v)_n \geq 0$, and Theorem 3.4 tells us that formula (1) still applies. But if strategy n earns a below average payoff, then $(\Phi v)_n < 0$, so $\Pi_{TX(x)}(v)$ cannot equal Φv. Instead, according to Theorem 3.4, the nth component of $\Pi_{TX(x)}(v)$ is set to 0, while the remaining components of $\Pi_{TX(x)}(v)$ are obtained from those of v by subtracting $\frac{1}{n-1}(v_1 + \ldots + v_{n-1})$ from each. More generally, the components of $\Pi_{TX(x)}(v)$ corresponding to "bad" unused strategies are set to 0, while the remaining components are obtained from v by normalizing away the average of these components only.

4. The Projection Dynamic: Definition and Microfoundations

4.1 Definition

An *evolutionary dynamic* is a map that assigns each population game F a dynamical system on the set of social states X. Typically the dynamical system is described by a differential equation $\dot{x} = V^F(x)$. To define the projection dynamic, we suppose that the vector $V^F(x)$ is the best approximation of the payoff vector $F(x)$ by a feasible direction of motion from x. As we noted in the introduction, this dynamic first appears in the transportation science literature in the work of Nagurney and Zhang (1997); see also Nagurney and Zhang (1996, Ch. 8).

Definition. *The* projection dynamic *assigns each population game F the differential equation*

$$\text{(P)} \qquad \dot{x} = \Pi_{TX(x)}(F(x)).$$

When $x \in \text{int}(X)$, the tangent cone $TX(x)$ is just the subspace TX, so the explicit formula for (P) is simply $\dot{x} = \mathbf{\Phi}F(x)$. When $x \in \text{bd}(X)$, the explicit formula for (P) is given by Theorem 3.4.

A relative of the projection dynamic from the economics literature is the *linear dynamic* of Friedman (1991, p. 643 and 661). Like the projection dynamic, the linear dynamic is defined by $\dot{x} = \mathbf{\Phi}F(x)$ on the interior of X. But at boundary states, the linear dynamic posits that all unused strategies have growth rates of zero, making the boundary forward invariant, while the growth rate of each strategy in use is the difference between its payoff and the average payoff of the strategies in use. By Theorem 3.4, the linear dynamic is identical to the projection dynamic if and only if at each state $x \in \text{bd}(X)$ and in all populations $p \in \mathcal{P}$, every unused strategy earns a payoff no greater than the average payoff of the strategies in use.

4.2 Microfoundations

We can provide microfoundations for evolutionary dynamics by deriving them from explicit models of individual choice. In this section, we provide such a foundation for the projection dynamic, and use this foundation to relate the dynamic to other dynamics from the literature. To conserve on notation, we focus on the single population setting; the analysis of multipopulation settings is a straightforward extension.

A *revision protocol* $\rho : \mathbf{R}^n \times X \to \mathbf{R}^{n \times n}_+$ describes the process through which individual agents in an evolutionary model make decisions. As time passes, agents are chosen at random from the population and granted opportunities to switch strategies. When an i player receives such an opportunity, he switches to strategy j with probability proportional to the conditional switch rate $\rho_{ij}(F(x), x)$. For future reference, we observe that the sum $\sum_{j \in S} \rho_{ij}(F(x), x)$ describes the rate at which strategy i players actively reevaluate their choice of strategy.

When a large population of agents using revision protocol ρ plays game F, its aggregate behavior is described by the *mean dynamic*[13]

$$\text{(M)} \qquad \dot{x}_i = V_i^F(x) = \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x).$$

The first term in (M) captures the inflow of agents to strategy i from other strategies, while the second term captures the outflow of agents from strategy i to other strategies.

In Table I, we present the revision protocols corresponding to five fundamental evolutionary dynamics: the replicator dynamic (Taylor and Jonker (1978)), the best response dynamic (Gilboa and Matsui (1991)), the logit dynamic (Fudenberg and Levine (1998)), the BNN dynamic (Brown and von Neumann (1950)), and the Smith dynamic (Smith (1984)). In the first and fourth of these examples, the function

$$\text{(2)} \qquad \hat{F}_i(x) = F_i(x) - \sum_{k \in S} x_k F_k(x)$$

describes the *excess payoff* to strategy i: that is, the difference between strategy i's payoff and the average payoff obtained by members of the population.

The replicator dynamic is the prime example of an *imitative dynamic*. Its revision protocol has agents choose new strategies by imitating opponents, as reflected by the x_j term in the definition of ρ_{ij}.[14] By contrast, the other dynamics in the table are *direct* (or

[13] See Benaïm and Weibull (2003) and Sandholm (2003, 2006).

[14] The revision protocol in the table, called *pairwise proportional imitation*, is due to Schlag (1998). Other revision protocols that generate the replicator dynamic can be found in Björnerstedt and Weibull (1996),

Revision protocol	Evolutionary dynamic	Name of dynamic
$\rho_{ij} = x_j[F_j - F_i]_+$	$\dot{x}_i = x_i\hat{F}_i(x)$	replicator
$\rho_{i\cdot} \in \underset{y\in X}{\operatorname{argmax}}\, y'F(x)$	$\dot{x} \in \underset{y\in X}{\operatorname{argmax}}\, y'F(x) - x$	best response
$\rho_{ij} = \dfrac{\exp(\eta^{-1}F_j)}{\sum_{k\in S}\exp(\eta^{-1}F_k)}$	$\dot{x}_i = \dfrac{\exp(\eta^{-1}F_i(x))}{\sum_{k\in S}\exp(\eta^{-1}F_k(x))} - x_i$	logit
$\rho_{ij} = [\hat{F}_j]_+$	$\dot{x}_i = [\hat{F}_i(x)]_+ - x_i \sum_{j\in S}[\hat{F}_j(x)]_+$	BNN
$\rho_{ij} = [F_j - F_i]_+$	$\dot{x}_i = \sum_{j\in S} x_j[F_i(x) - F_j(x)]_+ \; - x_i \sum_{j\in S}[F_j(x) - F_i(x)]_+$	Smith

Table I: Five evolutionary dynamics and their revision protocols.

innovative) *dynamics*, as the choices of new strategies under the protocols that generate them are not mediated through imitation of opponents.

We can also classify the dynamics according to whether agents actively reconsider their choices of strategy at a rate that is both fixed and independent of the strategy currently chosen. We call dynamics based on protocols of this sort *simple dynamics*. The best response and logit dynamics are simple in this sense, while the other dynamics listed in the table are not.[15]

Our goal in this section is to introduce a revision protocol whose mean dynamic is the projection dynamic, and to relate it to the protocols noted above. As a step toward this goal, consider the following imitative protocol:

$$\rho_{ij} = \begin{cases} [\hat{F}_i(x)]_- \cdot \dfrac{x_j\,[\hat{F}_j(x)]_+}{\sum_{k\in S} x_k\,[\hat{F}_k(x)]_+} & \text{if } \sum_{k\in S} x_k\,[\hat{F}_k(x)]_+ > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Under protocol (3), an agent actively reevaluates his choice of strategy only if his current

Weibull (1995), and Hofbauer (1995); also see Proposition 4.1 below.

[15]More explicitly, under the protocols for the best response and logit dynamics, $\rho_{ij}(F(x), x)$ does not depend on the current strategy i, and so can be written as $\sigma_j(F(x), x)$; furthermore, the vector $\sigma(F(x), x)$ is not only nonnegative, but is also a probability vector: $\sigma(F(x), x) \in X$ for all $x \in X$. It thus follows from equation (M) that these dynamics can be written in the "target" form $\dot{x} = \sigma(F(x), x) - x$, meaning that the vector of motion can be drawn with its tail at the current state x and its head at the state $\sigma(F(x), x)$ defined by the current choice probabilities. The replicator, BNN, and Smith dynamics cannot be expressed in this form.

payoff is below average; he does so at a rate proportional to how far below average. In this event, the agent chooses an opponent at random, switching to the opponent's strategy only if the opponent's payoff is above average, with probability proportional to how far above average; with the complementary probability, the agent chooses a new opponent at random and repeats the procedure.

Proposition 4.1, whose proof is provided in the Appendix, shows that the mean dynamic of protocol (3) is the replicator dynamic.

Proposition 4.1. *Revision protocol* (3) *generates the replicator dynamic* (R) *as its mean dynamic.*

We return at last to the projection dynamic by modifying protocol (3). To start, we define

$$\tilde{F}_i^{\mathcal{S}}(x) = F_i(x) - \frac{1}{\#\mathcal{S}(F(x), x)} \sum_{k \in \mathcal{S}(F(x),x)} F_k(x). \tag{4}$$

Like the excess payoff $\hat{F}_i$ defined in equation (2), $\tilde{F}_i^{\mathcal{S}}$ is given by the difference between strategy i's payoff and an average payoff. But here the latter is an unweighted average, and it only includes strategies in the set $\mathcal{S}(F(x), x)$ defined in Theorem 3.4.

Next we define the revision protocol

$$\rho_{ij} = \begin{cases} \dfrac{[\tilde{F}_i^{\mathcal{S}}(x)]_-}{x_i} \cdot \dfrac{[\tilde{F}_j^{\mathcal{S}}(x)]_+}{\sum_{k \in \mathcal{S}(F(x),x)} [\tilde{F}_k^{\mathcal{S}}(x)]_+} & \text{if } \sum\limits_{k \in \mathcal{S}(F(x),x)} x_i \, [\tilde{F}_k^{\mathcal{S}}(x)]_+ > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

Protocol (5) differs from protocol (3) in three ways. First, where in (3) conditional switch rates depend on weighted average payoffs, in (5) they depend on unweighted average payoffs. Second, while in (3) the decision to actively reevaluate only depends on payoffs, in (5) this decision also depends on the current strategy's popularity. In particular, the rate at which agents abandon strategy i is inversely proportional to x_i, and so can be understood as a model of "revision driven by insecurity".[16] Finally, once an agent using protocol (5) decides to actively consider switching strategies, his decision about which strategy to choose next is not based on imitation as in (3), but rather on the direct choice of possible alternatives. In this respect, the protocol resembles those of the BNN and Smith dynamics than that of the replicator dynamic.

[16] For one foundation for this functional form, suppose that at every moment in time, each agent chooses members of his population at random until he draws someone playing his own strategy, and then considers switching strategies at a rate proportional to the number of draws. As this number of draws has a *geometric*(x_i) distribution, the expected number of draws is $\frac{1}{x_i}$.

Proposition 4.2, whose proof is in the Appendix, confirms that the mean dynamic of revision protocol (5) is the projection dynamic.

Proposition 4.2. *Revision protocol* (5) *generates the projection dynamic* (P) *as its mean dynamic.*

Despite the fact that one of them is imitative and the other direct, the similarity between protocols (3) and (5) suggests the possibility of other connections between the replicator dynamic and the projection dynamic. There are in fact many such connections, and they are explored in the companion paper, Sandholm et al. (2008).

To conclude this section, we note that foundations recently have been suggested for another game dynamic defined in terms of closest point projections. The dynamic

$$\text{(TP)} \qquad \dot{x} = P_X(x + F(x)) - x.$$

first appears in the transportation literature in the work of Friesz et al. (1994), and is dubbed the *target projection dynamic* in Sandholm (2005), where a common mathematical ancestry is provided for this dynamic and the projection dynamic (P). Tsakas and Voorneveld (2007) show that the dynamic (TP) can be derived from a revision protocol based on best responses, if payoffs to mixed strategies are subjected to state-dependent perturbations:

$$\text{(6)} \qquad \rho_{i\cdot} = \operatorname*{argmax}_{y \in X} \left(y'F(x) - \tfrac{1}{2} \left| y - x \right|^2 \right).$$

Under protocol (6), each agent cares directly about the distance between his chosen mixed strategy and the current population state. Such preferences seem difficult to justify. At the same time, protocol (6) suggests a role for the target projection dynamic as a model of myopic learning in normal form games, as studied, e.g., by Börgers and Sarin (1997) and Hopkins (2002).

5. Basic Properties of Solution Trajectories

5.1 Existence, Uniqueness, and Continuity of Forward Solutions

Since the projection dynamic (P) is discontinuous at the boundary of X, standard results on the existence of solutions to differential equations do not apply to it. Indeed, the appropriate notion of solution for this equation must allow for kinks at boundary states: we call the trajectory $\{x_t\}_{t \geq 0} \subset X$ a (*Carathéodory*) *solution* to (P) if it is absolutely continuous and satisfies equation (P) at almost every $t \geq 0$. Theorem 5.1 shows that despite

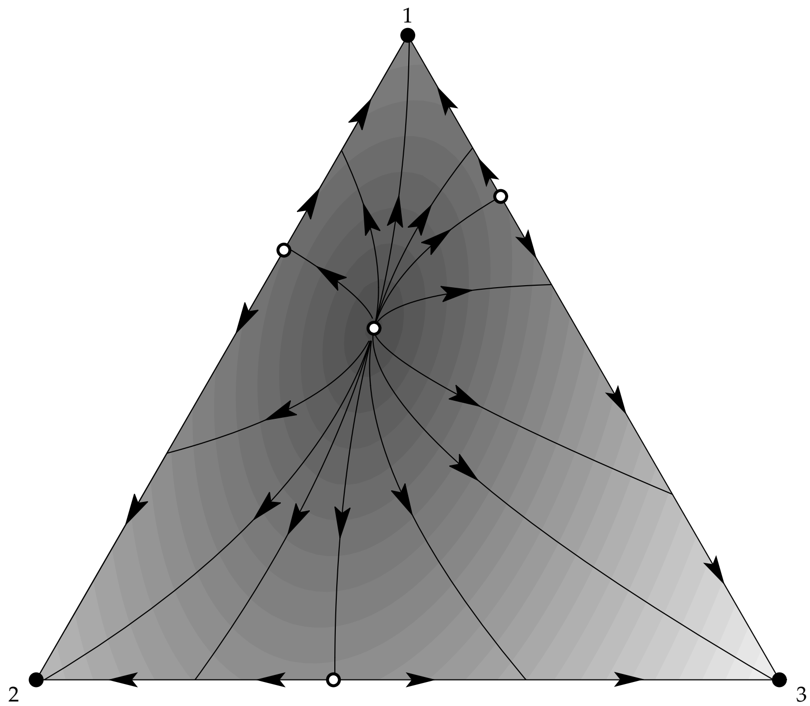

Figure 6: Phase diagram of (P) for a coordination game. Grayscale represents speeds: lighter shades indicate faster motion.

these inconveniences, forward-time solutions to (P) exist, are unique, and are Lipschitz continuous in their initial conditions.

Theorem 5.1. *Let a Lipschitz continuous population game F and an initial condition $\xi \in X$ be given. Then there exists a unique solution $\{x_t\}_{t\geq 0}$ to the projection dynamic* (P) *with $x_0 = \xi$. Solutions to* (P) *are Lipschitz continuous in their initial conditions: if $\{x_t\}_{t\geq 0}$ and $\{y_t\}_{t\geq 0}$ are solutions to* (P)*, then $\left|y_t - x_t\right| \leq \left|y_0 - x_0\right| e^{Kt}$ for all $t \geq 0$, where K is the Lipschitz coefficient for F.*

The existence result in Theorem 5.1 follows from more general existence results proved by Henry (1973) and Aubin and Cellina (1984, Sec. 5.6), and later rediscovered by Dupuis and Ishii (1991) and Dupuis and Nagurney (1993). In the appendix, we briefly present a proof of Theorem 5.1 based on the Viability Theorem for differential inclusions.

5.2 Examples

While the projection dynamic admits a unique solution from every initial condition, these solutions differ from solutions to standard Lipschitz differential equations in a number of important respects, as the following examples illustrate.

Example 5.2. A three-strategy coordination game. In Figure 6, we present the phase diagram for the projection dynamic in the three-strategy coordination game

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix}.$$

This phase diagram provides the solution trajectories generated by the vector field pictured in Figure 5(i). In both Figure 6 and Figure 7, the background color represents the speed of motion: regions where motion is fastest are red, while regions where motion is slowest are blue. (Since the dynamic changes discontinuously at the boundary of X, the colors are only guaranteed to represent speeds at interior states.)

As one travels away from the completely mixed equilibrium, the speed of motion increases until the boundary of the state space is reached; thus, $\text{bd}(X)$ is reached in finite time. At this point, the solution changes direction, merging with a solution that travels along the boundary of the simplex, implying that backward-time solutions from boundary states are not unique. All solutions reach one of the seven symmetric Nash equilibria of A in finite time, with solutions from almost all initial conditions leading to one of the three strict equilibria. §

Example 5.3. Rock-Paper-Scissors games. Consider the payoff matrix

$$A = \begin{pmatrix} 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{pmatrix}.$$

with $w, l > 0$. A is a *good RPS game* if $w > l$ (that is, if the winner's profit is higher than the loser's loss). A is a *standard RPS game* if $w = l$, and A is a *bad RPS game* if $w < l$. In all cases, the unique symmetric Nash equilibrium of A is $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

Figure 7 presents phase diagrams for the projection dynamic in good RPS ($w = 3, l = 2$), standard RPS ($w = l = 1$), and bad RPS ($w = 2, l = 3$). In all three games, solutions spiral around the Nash equilibrium in a counterclockwise direction. In good RPS (Figure 7(i)), all solutions converge to the Nash equilibrium. Solutions that start at an interior state close to a vertex first hit and then travel along the boundary of X; they then reenter $\text{int}(X)$ and spiral inward toward x^*. In standard RPS (Figure 7(ii)), all solutions enter closed orbits at a fixed distance from x^*. Solutions starting at distance $\frac{1}{\sqrt{6}}$ or greater from x^* (i.e., all solutions at least as far from x^* as the state $(0, \frac{1}{2}, \frac{1}{2})$) merge with the closed orbit at distance $\frac{1}{\sqrt{6}}$ from x^*; other solutions maintain their initial distance from x^* forever. In bad

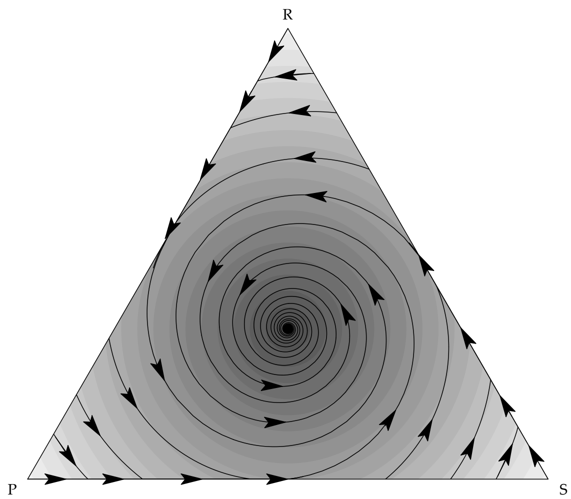

(i) Good RPS ($w = 3, l = 2$).

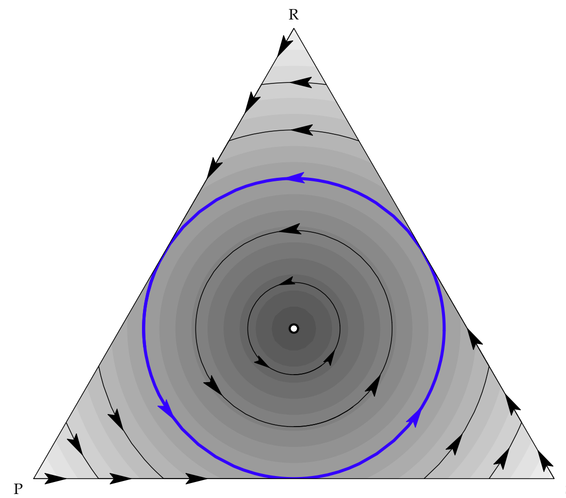

(ii) Standard RPS ($w = l = 1$).

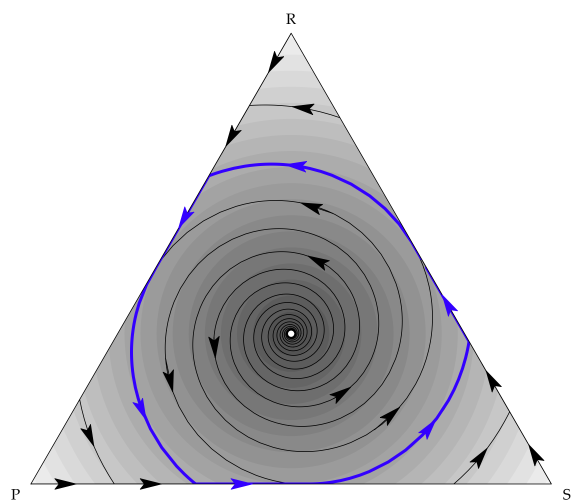

(iii) Bad RPS ($w = 2, l = 3$).

Figure 7: Phase diagrams of (P) for three Rock-Paper-Scissors games. Grayscale represents speeds: lighter shades indicate faster motion.

RPS (Figure 7(iii)), all solutions other than the one starting at x^* enter the same closed orbit. This orbit alternates between segments through the interior of X and segments that traverse the boundaries of X.

In all three versions of RPS, there are solution trajectories starting in $\text{int}(X)$ that reach $\text{bd}(X)$ in finite time. By Theorem 3.4, solutions leave $\text{bd}(X)$ at the point where the unused strategy's payoff exceeds the average payoff of the two strategies in use. §

6. Equilibrium and Convergence Properties

6.1 Nash Stationarity and Positive Correlation

Next, we first establish two basic game-theoretic properties of the projection dynamic. To state these properties, suppose that F is a population game and $\dot{x} = V(x)$ an evolutionary dynamic for this game. Define $RP(V) = \{x \in X : V(x) = \mathbf{0}\}$ to be the set of rest points of V.

(NS)	*Nash stationarity*	$RP(V) = NE(F)$.
(PC)	*Positive correlation*	$[V^p(x) \neq \mathbf{0} \Rightarrow V^p(x)'F^p(x) > 0]$ for all $p \in \mathcal{P}$.

Nash stationarity (NS) requires that the Nash equilibria of the game F and the rest points of the dynamic V coincide. Dynamics satisfying this condition provide the strongest support for the fundamental solution concept of noncooperative game theory. Positive correlation (PC) imposes restrictions on disequilibrium dynamics. It requires that whenever population $p \in \mathcal{P}$ is not at rest, there is a positive correlation between the growth rates and payoffs of strategies in S^p. In geometric terms, (PC) demands that the direction of motion $V^p(x)$ and the payoff vector $F^p(x)$ form acute angles with one another whenever $V^p(x)$ is not null. This property, versions of which have been studied by Friedman (1991), Swinkels (1993), and Sandholm (2001), is an important ingredient in establishing global convergence results—see Section 6.2 below.

Both (NS) and (PC) are simple consequences of the developments in Section 3.

Proposition 6.1. *The projection dynamic satisfies Nash stationarity* (NS) *and positive correlation* (PC).

Proof. Property (NS) is a restatement of Corollary 3.3. To prove property (PC), we take the Moreau decomposition of the payoff vector $F^p(x)$:

$$\begin{aligned} V^p(x)'F^p(x) &= \Pi_{TX^p(x^p)}(F^p(x))' \left(\Pi_{TX^p(x^p)}(F^p(x)) + \Pi_{NX^p(x^p)}(F^p(x))\right) \\ &= \left|\Pi_{TX^p(x^p)}(F^p(x))\right|^2 \end{aligned}$$

$$\geq 0.$$

The inequality is strict if and only if $\Pi_{TX^p(x^p)}(F^p(x)) = V^p(x) \neq \mathbf{0}$. ■

6.2 Global Convergence in Potential Games

In the remainder of this section, we show that the projection dynamic converges to Nash equilibrium from all initial conditions in two important classes of games: potential games (Monderer and Shapley (1996), Sandholm (2001, 2008)) and stable games (Hofbauer and Sandholm (2008)).

In a potential game, information about all strategies' payoffs is encoded in a single scalar-valued function on the state space X. Following Sandholm (2008), we call the population game $F : X \to \mathbf{R}^n$ a *potential game* if it admits a *potential function* $f : X \to \mathbf{R}$ satisfying

$$(7) \qquad \nabla f(x) = \mathbf{\Phi} F(x) \text{ for all } x \in X.$$

Since the domain of f is X, the gradient vector $\nabla f(x)$ is the unique vector in TX that represents the derivative of f at X, in the sense that $f(y) = f(x) + \nabla f(x)'(y - x) + o(|y - x|)$ for all $y \in X$. Definition (7) requires that this gradient vector always equal the projected payoff vector $\mathbf{\Phi} F(x)$. Common interest games, congestion games, and games defined by variable externality pricing schemes are all potential games.

In potential games, all limit points of the projection dynamic are Nash equilibria.

Theorem 6.2. *Let F be a potential game with potential function f. Then f is a strict Lyapunov function for the projection dynamic* (P) *on X. Therefore, each solution to* (P) *converges to a connected set of Nash equilibria of F.*

Proof. Property (PC) and the fact that $V(x) \in TX$ imply that

$$\dot{f}(x) = \nabla f(x)'\dot{x} = (\mathbf{\Phi} F(x))' V(x) = F(x)' V(x) = \sum_{p \in \mathcal{P}} F^p(x)' V^p(x) \geq 0,$$

and that $\dot{f}(x) = 0$ if and only if $V(x) = 0$. Therefore, standard results (e.g., Theorem 7.6 of Hofbauer and Sigmund (1988)) imply that every solution of (P) converges to a connected set of rest points of V. By Nash stationarity, these rest points are all Nash equilibria. ■

6.3 Global Convergence and Cycling in Stable Games

The population game F is a *stable game* (Hofbauer and Sandholm (2008)) if

$$(y - x)'(F(y) - F(x)) \leq 0 \text{ for all } x, y \in X. \tag{8}$$

If inequality (8) is strict whenever $y \neq x$, then F is a *strictly stable game*; if (8) is always satisfied with equality, then F is a *null stable game*. The set of Nash equilibria of any stable game is convex; if F is strictly stable, then $NE(F)$ is a singleton. Games with an interior ESS, wars of attrition, and congestion games in which congestion is a bad are all stable games, while zero-sum games are null stable games.[17]

Let

$$E_{x^*}(x) = \left|x - x^*\right|^2,$$

the squared Euclidean distance from the Nash equilibrium x^*. Nagurney and Zhang (1997) show that E_{x^*} is a Lyapunov function in any stable game. The analysis to follow includes a streamlined proof of their result.

Theorem 6.3. *Let x^* be a Nash equilibrium of F.*

(*i*) *If F is a stable game, then E_{x^*} is a Lyapunov function for* (P)*, so x^* is Lyapunov stable under* (P)*.*

(*ii*) *If F is a strictly stable game, then E_{x^*} is a strict Lyapunov function for* (P)*, so $NE(F) = x^*$ is globally asymptotically stable under* (P)*.*

(*iii*) *If F is a null stable game and $x^* \in \text{int}(X)$, then E_{x^*} defines a constant of motion for* (P) *on* $\text{int}(X)$*.*

Proof. The proof of Theorem 3.1 shows that x^* is a Nash equilibrium if and only if

$$(x - x^*)'F(x^*) \leq 0 \text{ for all } x \in X. \tag{9}$$

By adding this inequality to inequality (8), Hofbauer and Sandholm (2008) show that if F is a stable game, then $x^* \in NE(F)$ if and only if x^* is a *globally neutrally stable state* of F:

$$(x - x^*)'F(x) \leq 0 \text{ for all } x \in X; \tag{10}$$

while if F is a strictly stable game, then its unique Nash equilibrium x^* is also its unique

[17] In the variational inequality literature, condition (8) is known as *monotonicity*—see Minty (1967), Kinderlehrer and Stampacchia (1980), Harker and Pang (1990), and Nagurney (1999). Further discussion of condition (8) can be found in Hofbauer and Sandholm (2008).

globally evolutionarily stable state:

(11) $\quad (x - x^*)'F(x) < 0$ for all $x \in X - \{x^*\}$.

Now suppose that F is stable. Using the Moreau decomposition, equation (10), and the fact that $x^* - x \in TX(x)$, we compute the time derivative of E_{x^*} over a solution to (P):

$$
\begin{aligned}
(12) \qquad \dot{E}_{x^*}(x) &= \nabla E_{x^*}(x)'\dot{x} \\
&= 2(x - x^*)'\Pi_{TX(x)}(F(x)) \\
&= 2(x - x^*)'F(x) + 2(x^* - x)'\Pi_{NX(x)}(F(x)) \\
&\leq 2(x^* - x)'\Pi_{NX(x)}(F(x)) \\
&\leq 0.
\end{aligned}
$$

Thus, E_{x^*} is a Lyapunov function for (P), which implies that x^* is Lyapunov stable. If F is strictly stable, then equation (11) implies that the first inequality in (12) is strict; thus, E_{x^*} is a strict Lyapunov function for (P), and so x^* is globally asymptotically stable.

Finally, suppose that F is null stable and that x^* is an interior Nash equilibrium. The first of these assumptions tells us that equation (8) always holds with equality, while the second implies that all pure strategies are optimal, and hence that equation (9) always holds with equality. Adding these two equalities shows that equation (10), and hence the first inequality in (12), always holds with equality. If $x \in \text{int}(X)$, then $NX(x)$ and $TX(x)$ are orthogonal, so the the second inequality in (12) holds with equality as well. Therefore, $\dot{E}_{x^*}(x) = 0$ on $\text{int}(X)$, and so E_{x^*} defines a constant of motion for (P) on this set. ■

To conclude this section, we show that at interior states, the squared speed of motion under (P) also serves as a Lyapunov function for (P). Unlike that of the distance function E_{x^*}, the definition of this function does not directly incorporate the Nash equilibrium x^*.

Theorem 6.4. *Let F be continuously differentiable. If F is a stable game, then $L(x) = |\mathbf{\Phi}F(x)|^2$ is a Lyapunov function for* (P) *on* $\text{int}(X)$*. If F is a null stable, then L defines a constant of motion for* (P) *on* $\text{int}(X)$*.*

Proof. Since F is C^1, the Fundamental Theorem of Calculus implies that F is stable if and only if

(13) $\quad z'DF(x)z \leq 0$ for all $z \in TX$ and $x \in X$,

and that F is null stable game if and only if inequality (13) always binds.

When $x \in \text{int}(x)$, the projection dynamic is given by $\dot{x} = \mathbf{\Phi}F(x)$. Therefore, since $D\mathbf{\Phi}F(x) = \mathbf{\Phi}DF(x)$ and $\mathbf{\Phi}^2 = \mathbf{\Phi}$, we find that

$$\dot{L}(x) = \nabla L(x)'\dot{x} = 2(\mathbf{\Phi}F(x))'D\mathbf{\Phi}F(x)(\mathbf{\Phi}F(x)) = 2(\mathbf{\Phi}F(x))'DF(x)(\mathbf{\Phi}F(x)).$$

Of course, $\mathbf{\Phi}F(x) \in TX$. Therefore, if F is stable, then $\dot{L}(x) \leq 0$ on $\text{int}(X)$, and if F is null stable, then $\dot{L}(x) = 0$ on $\text{int}(X)$. ■

The results in this section are illustrated by our phase diagrams for the projection dynamic in Rock-Paper-Scissors games (Example 5.3). Good RPS defines a strictly stable population game. In its phase diagram (Figure 7(i)), we see that distance from the equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ falls over time, and that the speed of motion falls over time on $\text{int}(X)$. Standard RPS is a zero-sum game, and so defines a null stable population game. In this game's phase diagram (Figure 7(ii)), both distance from equilibrium and the speed of motion remain fixed over time on $\text{int}(X)$.

As with its behavior in potential games, the behavior of the projection dynamic in stable games provides a link between it and the replicator dynamic. This link and others are explored in the companion paper, Sandholm et al. (2008).

A. Appendix

The Proof of Theorem 3.4

In this section, we derive the formula for the projection map $\Pi_{TX^p(x^p)}(v^p)$ stated in Theorem 3.4. To eliminate superscripts, we suppose that $p = 1$. Then Theorem 3.4 is an immediate consequence of the two propositions to follow.

Proposition A.1. *The vector $z^* = \Pi_{TX(x)}(v)$ is defined by*

$$z_i^* = \begin{cases} v_i - \mu & \text{if } x_i > 0, \text{ or if } x_i = 0 \text{ and } v_i > \mu, \\ 0 & \text{if } x_i = 0 \text{ and } v_i \leq \mu \end{cases} \tag{14}$$

for some $\mu = \mu(v, x)$.

Proof. Let $Y = \{i \in S : x_i > 0\}$ and $N = \{i \in S : x_i = 0\}$ denote the sets of strategies that are used and unused at x. Then z^* is the solution to the following quadratic program:

$$\min_{z \in \mathbf{R}^n} \tfrac{1}{2} \sum_{i \in S} (z_i - v_i)^2 \qquad \text{s.t.} \qquad \sum_{i \in S} z_i = 0$$
$$\text{and} \qquad z_i \geq 0 \text{ for all } i \in N.$$

The Lagrangian for this program is

$$L(z,\mu,\gamma) = -\tfrac{1}{2}\sum_{i\in S}(z_i - v_i)^2 - \mu\sum_{i\in S} z_i + \sum_{j\in N}\lambda_j z_j,$$

so the Kuhn-Tucker first order conditions are

$$\begin{array}{lll}
\text{(15a)} & z_i^* = v_i - \mu & \text{for all } i \in Y, \\
\text{(15b)} & z_j^* = v_j - \mu + \lambda_j & \text{for all } j \in N, \\
\text{(15c)} & \lambda_j \geq 0 & \text{for all } j \in N, \text{ and} \\
\text{(15d)} & \lambda_j z_j^* = 0 & \text{for all } j \in N.
\end{array}$$

If we let $N^+ = \{i \in N : z_i^* > 0\}$ and $N^0 = \{j \in N : z_j^* = 0\}$ we can rewrite the Kuhn-Tucker conditions (15a-15d) as

$$\begin{array}{lll}
\text{(16a)} & z_i^* = v_i - \mu > 0 & \text{for all } i \in Y \cup N^+, \\
\text{(16b)} & v_j^* = z_j - \mu + \lambda_j = 0 & \text{for all } j \in N^0, \text{ and} \\
\text{(16c)} & \lambda_j \geq 0 & \text{for all } j \in N^0.
\end{array}$$

It follows immediately from equations (16a-16c) that

$$\text{(17)} \qquad N^+ = \{i \in N : v_i > \mu\} \text{ and } N^0 = \{i \in N : v_i \leq \mu\}.$$

Equations (16a-16c) and (17) together yield equation (14). ■

To complete our description of $z^* = \Pi_{TX(x)}(v)$, and hence of the projection dynamic (P), we need to determine the value of $\mu = \mu(x,v)$. To accomplish this, we permute the names of the strategies in S so that strategies 1 through $n_Y = \#Y$ are in Y and strategies $n_Y + 1$ through n are in N, with the latter strategies ordered so that $v_{n_Y+1} \geq v_{n_Y+2} \geq \cdots \geq v_n$. For $k \in \{1,\ldots,n\}$, let

$$\mu_k = \frac{1}{k}\sum_{i=1}^{k} v_i$$

be the average of the first k components of v. Throughout the analysis, we make use of the recursion

$$\mu_{k+1} = \frac{1}{k+1}v_{k+1} + \frac{k}{k+1}\mu_k,$$

which tells us that μ_{k+1} and v_{k+1} deviate from μ_k in the same direction. Finally, to simplify the statement of the result to follow, we set $\mu_{n+1} = -\infty$.

Proposition A.2. *The Lagrange multiplier μ is equal to μ_{k^*}, where*

$$k^* = \min\{k \in \{n_Y, \ldots, n\} : \mu_{k+1} \le \mu_k\}.$$

Therefore, the set of unused strategies with positive components in z^* is $N^+ = \{n_Y + 1, \ldots, k^*\}$.

Proof. To prove the proposition, it is enough to show that μ_{k^*} satisfies

$$v_i > \mu_{k^*} \text{ for all } i \in N^+, \text{ and}$$
$$\mu_{k^*} \ge v_j \text{ for all } j \in N^0.$$

Therefore, by our ordering assumption, we need only show that

(18) if $k^* > n_Y$, then $v_{k^*} > \mu_{k^*}$, and

(19) if $k^* < n$, then $\mu_{k^*} \ge v_{k^*+1}$.

We divide the analysis into four cases. If $n_Y = n$, then $k^* = n$ as well, and conditions (18) and (19) are vacuous. (In this case, we obtain $\mu = \mu_n = \frac{1}{n}\sum_{i=1}^n v_i$ and $z^* = \Phi v$, as expected.) Therefore, for the remainder of the proof we can suppose that $n_Y < n$.

If $k^* = n_Y$, then condition (18) is vacuous. Since $\mu_{k^*+1} \le \mu_{k^*}$ by the definition of k^*, it follows from the recursion that $v_{k^*+1} \le \mu_{k^*}$, which is condition (19).

If $k^* \in \{n_Y, \ldots, n-1\}$, then the definition of k^* tells us that $\mu_{k^*} > \mu_{k^*-1}$. It then follows from the recursion that $v_{k^*} > \mu_{k^*-1}$, and hence that

$$v_{k^*} > \frac{1}{k^*} v_{k^*} + \frac{k^*-1}{k^*}\mu_{k^*-1} = \mu_{k^*},$$

which is condition (18). The definition of k^* also tells us that $\mu_{k^*+1} \le \mu_{k^*}$; the recursion then reveals that $v_{k^*+1} \le \mu_{k^*}$, which is condition (19).

Finally, if $n_Y < n$ and $k^* = n$, then condition (19) is vacuous, and repeating the corresponding proof from the previous case gives us condition (18). ■

The Proof of Proposition 4.1

Observe that the weighted average of the components of the excess payoff vector $\hat{F}(x)$

is 0:

$$\sum_{k\in S} x_k\hat{F}_k(x) = \sum_{k\in S} x_k\left(F_k(x) - \sum_{l\in S} x_l F_l(x)\right) = \sum_{k\in S} x_k F_k(x) - \sum_{l\in S} x_l F_l(x) = 0.$$

It follows directly that

$$\sum_{k\in S} x_k[\hat{F}_k(x)]_+ = \sum_{k\in S} x_k[\hat{F}_k(x)]_-. \tag{20}$$

Now fix a state $x \in X$. If both sides of equation (20) equal 0 at x, then the mean dynamic generated by protocol (3) has a rest point at x; but since in this case $x_k\hat{F}_k(x) = 0$ for all $k \in S$, the dynamic (R) has a rest point at x as well. On the other hand, if both sides of equality (20) are positive at state x, we can use this equality to compute as follows:

$$\begin{aligned}
\dot{x}_i &= \sum_{j\in S} x_j\rho_{ji} - x_i\sum_{j\in S}\rho_{ij} \\
&= \sum_{j\in S} x_j\left([\hat{F}_j(x)]_-\frac{x_i[\hat{F}_i(x)]_+}{\sum_{k\in S} x_k[\hat{F}_k(x)]_+}\right) - x_i\sum_{j\in S}\left([\hat{F}_i(x)]_-\frac{x_j[\hat{F}_j(x)]_+}{\sum_{k\in S} x_k[\hat{F}_k(x)]_+}\right) \\
&= x_i[\hat{F}_i(x)]_+ - x_i[\hat{F}_i(x)]_- \\
&= x_i\hat{F}(x),
\end{aligned}$$

This too agrees with equation (R), completing the proof. ∎

The Proof of Proposition 4.2

Write $\mathcal{S}$ for $\mathcal{S}(F(x), x)$. Since

$$\sum_{k\in\mathcal{S}} \tilde{F}^{\mathcal{S}}_k(x) = \sum_{k\in\mathcal{S}}\left(F^{\mathcal{S}}_k(x) - \frac{1}{\#\mathcal{S}}\sum_{l\in\mathcal{S}} F^{\mathcal{S}}_l(x)\right) = 0,$$

we see that

$$\sum_{k\in\mathcal{S}}[\tilde{F}^{\mathcal{S}}_k(x)]_+ = \sum_{k\in\mathcal{S}}[\tilde{F}^{\mathcal{S}}_k(x)]_-. \tag{21}$$

Also, we note these two implications of Theorem 3.4:

$$\text{if } j \in \mathcal{S} \text{ and } x_j = 0, \text{ then } [\tilde{F}^{\mathcal{S}}_j(x)]_- = 0; \tag{22}$$

$$\text{if } j \notin \mathcal{S}, \text{ then } [\tilde{F}^{\mathcal{S}}_j(x)]_+ = 0. \tag{23}$$

Fix a state $x \in X$. If both sides of equation (21) equal 0 at x, then the mean dynamic generated by protocol (5) has a rest point at x; but since in this case $\tilde{F}^{\mathcal{S}}_k(x) = 0$ for all $k \in \mathcal{S}$, the dynamic (P) has a rest point at x as well.

Suppose instead that both sides of equation (21) are positive at state x. If $x_i > 0$, we can use equations (22), (23), and (21) to compute as follows:

$$\begin{aligned}
\dot{x}_i &= \sum_{j \in S} x_j \rho_{ji} - x_i \sum_{j \in S} \rho_{ij} \\
&= \sum_{j:\, x_j > 0} x_j \left(\frac{[\tilde{F}^{\mathcal{S}}_j(x)]_-}{x_j} \cdot \frac{[\tilde{F}^{\mathcal{S}}_i(x)]_+}{\sum_{k \in \mathcal{S}} [\tilde{F}^{\mathcal{S}}_k(x)]_+} \right) - x_i \sum_{j \in S} \left(\frac{[\tilde{F}^{\mathcal{S}}_i(x)]_-}{x_i} \cdot \frac{[\tilde{F}^{\mathcal{S}}_j(x)]_+}{\sum_{k \in S} [\tilde{F}^{\mathcal{S}}_k(x)]_+} \right) \\
&= \left(\frac{\sum\limits_{j:\, x_j > 0} [\tilde{F}^{\mathcal{S}}_j(x)]_-}{\sum\limits_{k \in \mathcal{S}} [\tilde{F}^{\mathcal{S}}_k(x)]_+} \right) [\tilde{F}^{\mathcal{S}}_i(x)]_+ - \left(\frac{\sum\limits_{j \in S} [\tilde{F}^{\mathcal{S}}_j(x)]_+}{\sum\limits_{k \in \mathcal{S}} [\tilde{F}^{\mathcal{S}}_k(x)]_+} \right) [\tilde{F}^{\mathcal{S}}_i(x)]_- \\
&= \left(\frac{\sum\limits_{j \in \mathcal{S}} [\tilde{F}^{\mathcal{S}}_j(x)]_-}{\sum\limits_{k \in \mathcal{S}} [\tilde{F}^{\mathcal{S}}_k(x)]_+} \right) [\tilde{F}^{\mathcal{S}}_i(x)]_+ - \left(\frac{\sum\limits_{j \in \mathcal{S}} [\tilde{F}^{\mathcal{S}}_j(x)]_+}{\sum\limits_{k \in \mathcal{S}} [\tilde{F}^{\mathcal{S}}_k(x)]_+} \right) [\tilde{F}^{\mathcal{S}}_i(x)]_- \\
&= [\tilde{F}^{\mathcal{S}}_i(x)]_+ - [\tilde{F}^{\mathcal{S}}_i(x)]_- \\
&= \tilde{F}^{\mathcal{S}}(x).
\end{aligned}$$

This agrees with equation (P). If $x_i = 0$, then the second term in the previous calculation drops out immediately, and the calculation of the first term shows that $\dot{x}_i = [\tilde{F}^{\mathcal{S}}_i(x)]_+$, again in agreement with equation (P). This completes the proof of the theorem. ■

The Proof of Theorem 5.1

We first sketch a proof of existence of solutions to (P) due to Aubin and Cellina (1984). Define the multivalued map $V : X \to \mathbf{R}^n$ by

$$V(x) = \bigcap_{\varepsilon > 0} \mathrm{cl}\left(\mathrm{conv}\left(\bigcup_{y \in X: |y - x| \le \varepsilon} \Pi_{TX(y)}(F(y)) \right) \right).$$

In words, $V(x)$ is the closed convex hull of all values of $\Pi_{TX(y)}(F(y))$ that obtain at points y arbitrarily close to x. It is easy to check that V is upper hemicontinuous with closed convex values. Moreover, $V(x) \cap TX(x)$, the set of feasible directions of motion from x contained in $V(x)$, is always equal to $\{\Pi_{TX(x)}(F(x))\}$, and so in particular is nonempty. Because $V(x) \cap TX(x) \neq \emptyset$, the Viability Theorem for differential inclusions (Aubin and Cellina (1984, Theorem 4.2.1)) implies that for each $\xi \in X$, a Carathéodory solution $\{x_t\}_{t \ge 0}$

to $\dot{x} \in V(x)$ exists. But since $V(x) \cap TX(x) = \{\Pi_{TX(x)}(F(x))\}$, this solution must also solve the original equation (P).

Our proof of uniqueness and Lipschitz continuity of solutions to (P) follows Cojocaru and Jonker (2004). Let $\{x_t\}$ and $\{y_t\}$ be solutions to (P). Using the chain rule, Theorem 3.2, and the Lipschitz continuity of F, we see that

$$\begin{aligned}
\tfrac{d}{dt}\left|y_t - x_t\right|^2 &= 2(y_t - x_t)'(\Pi_{TX(y_t)}(F(y_t)) - \Pi_{TX(x_t)}(F(x_t))) \\
&= 2(y_t - x_t)'(F(y_t) - F(x_t)) - 2(y_t - x_t)'(\Pi_{NX(y_t)}(F(y_t)) - \Pi_{NX(x_t)}(F(x_t))) \\
&= 2(y_t - x_t)'(F(y_t) - F(x_t)) + 2(x_t - y_t)'\Pi_{NX(y_t)}(F(y_t)) \\
&\quad + 2(y_t - x_t)'\Pi_{NX(x_t)}(F(x_t)) \\
&\leq 2(y_t - x_t)'(F(y_t) - F(x_t)) \\
&\leq 2K\left|y_t - x_t\right|^2,
\end{aligned}$$

and hence that

$$\left|y_t - x_t\right|^2 \leq \left|y_0 - x_0\right|^2 + \int_0^t 2K\left|y_s - x_s\right| ds.$$

Gronwall's inequality then implies that

$$\left|y_t - x_t\right|^2 \leq \left|y_0 - x_0\right|^2 \mathrm{e}^{2Kt}.$$

Taking square roots yields the inequality stated in the theorem. ■

References

Aubin, J.-P. and Cellina, A. (1984). *Differential Inclusions*. Springer, Berlin.

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

Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.

Björnerstedt, J. and Weibull, J. W. (1996). Nash equilibrium and evolution by imitation. In Arrow, K. J. et al., editors, *The Rational Foundations of Economic Behavior*, pages 155–181. St. Martin's Press, New York.

Börgers, T. and Sarin, R. (1997). Learning through reinforcement and the replicator dynamics. *Journal of Economic Theory*, 77:1–14.

Brown, G. W. and von Neumann, J. (1950). Solutions of games by differential equations. In Kuhn, H. W. and Tucker, A. W., editors, *Contributions to the Theory of Games I*, volume 24 of *Annals of Mathematics Studies*, pages 73–79. Princeton University Press, Princeton.

Cojocaru, M.-G. and Jonker, L. B. (2004). Existence of solutions to projected differential equations in Hilbert spaces. *Proceedings of the American Mathematical Society*, 132:183–193.

Dupuis, P. and Ishii, H. (1991). On Lipschitz continuity of the solution mapping to the Skorokhod problem, with applications. *Stochastics and Stochastics Reports*, 35:31–62.

Dupuis, P. and Nagurney, A. (1993). Dynamical systems and variational inequalities. *Annals of Operations Research*, 44:9–42.

Friedman, D. (1991). Evolutionary games in economics. *Econometrica*, 59:637–666.

Friesz, T. L., Bernstein, D., Mehta, N. J., Tobin, R. L., and Ganjalizadeh, S. (1994). Day-to-day dynamic network disequilibria and idealized traveler information systems. *Operations Research*, 42:1120–1136.

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

Gabay, D. and Moulin, H. (1980). On the uniqueness and stability of Nash equilibria in noncooperative games. In Bensoussan, A., Kleindorfer, P., and Tapiero, C. S., editors, *Applied Stochastic Control in Econometrics and Management Science*, pages 271–293. North Holland, Amsterdam.

Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.

Harker, P. T. and Pang, J.-S. (1990). Finite-dimensional variational inequality and nonlinear complementarity problems: A survey of theory, algorithms, and applications. *Mathematical Programming*, 48:161–220.

Henry, C. (1973). An existence theorem for a class of differential equations with multivalued right-hand side. *Journal of Mathematical Analysis and Applications*, 41:179–186.

Hiriart-Urruty, J.-B. and Lemaréchal, C. (2001). *Fundamentals of Convex Analysis*. Springer, Berlin.

Hofbauer, J. (1995). Imitation dynamics for games. Unpublished manuscript, University of Vienna.

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

Hofbauer, J. and Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132:47–69.

Hofbauer, J. and Sandholm, W. H. (2008). Stable population games and integrability for evolutionary dynamics. Unpublished manuscript, University of Vienna and University of Wisconsin.

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

Hopkins, E. (2002). Two competing models of how people learn in games. *Econometrica*, 70:2141–2166.

Kinderlehrer, D. and Stampacchia, G. (1980). *An Introduction to Variational Inequalities and their Applications*. Academic Press, New York.

Minty, G. J. (1967). On the generalization of a direct method of the calculus of variations. *Bulletin of the American Mathematical Society*, 73:315–321.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.

Nagurney, A. (1999). *Network Economics: A Variational Inequality Approach*. Kluwer, Dordrecht, second edition.

Nagurney, A. and Zhang, D. (1996). *Projected Dynamical Systems and Variational Inequalities with Applications*. Kluwer, Dordrecht.

Nagurney, A. and Zhang, D. (1997). Projected dynamical systems in the formulation, stability analysis, and computation of fixed demand traffic network equilibria. *Transportation Science*, 31:147–158.

Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33:520–534.

Rosenthal, R. W. (1973). A class of games possessing pure strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67.

Sandholm, W. H. (2001). Potential games with continuous player sets. *Journal of Economic Theory*, 97:81–108.

Sandholm, W. H. (2003). Evolution and equilibrium under inexact information. *Games and Economic Behavior*, 44:343–378.

Sandholm, W. H. (2005). Excess payoff dynamics and other well-behaved evolutionary dynamics. *Journal of Economic Theory*, 124:149–170.

Sandholm, W. H. (2006). Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. Unpublished manuscript, University of Wisconsin.

Sandholm, W. H. (2008). Potential functions for normal form games and for population games. Unpublished manuscript, University of Wisconsin.

Sandholm, W. H., Dokumacı, E., and Lahkar, R. (2008). The projection dynamic and the replicator dynamic. *Games and Economic Behavior*, forthcoming.

Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78:130–156.

Smith, M. J. (1984). The stability of a dynamic model of traffic assignment—an application of a method of Lyapunov. *Transportation Science*, 18:245–252.

Swinkels, J. M. (1993). Adjustment dynamics and rational play in games. *Games and Economic Behavior*, 5:455–484.

Taylor, P. D. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.

Tsakas, E. and Voorneveld, M. (2007). The target projection dynamic. Unpublished manuscript, Göteborg University and Stockholm School of Economics.

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