Excess Payoff Dynamics and Other Well-Behaved Evolutionary Dynamics

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

February 11, 2005

* I thank many seminar audiences, as well as Martin Cripps, Josef Hofbauer, Ratul Lahkar, Aki Matsui, Larry Samuelson, Satoru Takahashi, Jörgen Weibull, and Peyton Young, for helpful discussions. The comments of two referees and an Associate Editor are also sincerely appreciated. Financial support from NSF Grant SES-0092145 is gratefully acknowledged. Versions of Sections 2 through 4 of this article first appeared in a working paper entitled "Excess Payoff Dynamics, Potential Dynamics, and Stable Games".

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 show that every excess payoff dynamic is *well-behaved*: regardless of the underlying game, each excess payoff dynamic admits unique solution trajectories that vary continuously with the initial state, identifies rest points with Nash equilibria, and respects a basic payoff monotonicity property. We show how excess payoff dynamics can be used to construct well-behaved modifications of imitative dynamics, and relate them to two other well-behaved dynamics based on projections.

1. Introduction

Evolutionary game theory is the study of strategic interactions in large populations whose members base decisions on simple myopic rules. This approach to game theory stands in contrast to traditional approaches based on the assumption of equilibrium play. Of fundamental interest is the connection between the two approaches: To what extent do evolutionary models support traditional predictions of play? What sorts of myopic decision rules sustain this link?

In this paper, we seek evolutionary dynamics that exhibit three attractive properties regardless of the strategic environment at hand. *Existence, uniqueness, and continuity of solutions* (EUC) requires a dynamic to admit exactly one solution from each initial state, and requires solutions to change continuously as one varies the initial state. Failures of (EUC) mean that slightly inaccurate information about initial behavior can spawn large errors in predictions of future behavior, even over short spans of time. *Nash stationarity* (NS) requires a one-to-one link between the stationary states of an evolutionary dynamic and the Nash equilibria of the underlying game. This condition provides the basic link between the evolutionary dynamic and the predictions of the standard theory. Finally, *positive correlation* (PC) requires that strategies' growth rates and payoffs be positively correlated. In so doing, the condition ensures that out-of-equilibrium dynamics reflect strategic incentives in a reasonable way.[1] We call any dynamic that respects properties (EUC), (NS), and (PC) *well-behaved*.

Interestingly, neither of the two best known evolutionary dynamics are well-behaved in the sense defined above. The replicator dynamic satisfies (EUC) and (PC), but fails (NS): while all Nash equilibria are rest points of this dynamic, the dynamic also admits boundary rest points that are not Nash equilibria.[2] The best response dynamic satisfies modified versions of (PC) and (NS), but fails (EUC). Since this dynamic's law of motion is discontinuous, even its behavior over short time spans is quite sensitive to its initial state. Thus, while solutions to the best response dynamic exist and are upper hemicontinuous in their initial conditions, multiple solution trajectories can emanate from a single initial condition.[3]

[1] To understand what is at issue here, compare the replicator dynamic to the dynamic defined by its negation. The solution trajectories of this new dynamic are identical to those of the replicator dynamic after a reversal of time. Both of these dynamics satisfy property (EUC), and both generate identical rest points. But the dynamics have very different out-of-equilibrium properties. For instance, while strict Nash equilibria are asymptotically stable under the replicator dynamic, they are unstable under the time-reversed version.

[2] See, for example, Section 3.3 of Weibull (1995).

[3] Upper hemicontinuity follows from standard results on differential inclusions—see, for example, Theorem 4.11 of Smirnov (2002). For examples of nonuniqueness of solutions, see Gilboa and Matsui

Both the replicator dynamic and the best response dynamic can be derived from models of individual choice. The former dynamic describes aggregate behavior in certain models of imitation of successful agents,[4] while the latter dynamic is derived from a model of optimal choice.

In this paper, we introduce a new paradigm for individual choice in evolutionary models: in place of imitation or optimization, we consider *moderation*. Under this paradigm, agents exert moderate levels of effort to find strategies that perform well. The measures taken to select good strategies are sufficiently modest that suboptimal and even subaverage choices are made with nonnegligible probability. Moderation obviously differs from optimization, and it is distinct from imitation in that it is based on direct rather than indirect evaluation of strategies' payoffs. Still, we believe that moderation aptly describes choice behavior in many situations naturally modeled using an evolutionary approach.

Our formal analysis starts with a simple model of individual choice. Agents receive opportunities to choose new strategies according to independent Poisson processes, choosing stochastically from the available strategies when such opportunities arise. Both revision rates and choice probabilities are functions of the strategies' *excess payoffs*—that is, on the differences between the strategies' payoffs and the population's average payoff. More precisely, the agents' revision protocols are defined in terms of objects called *raw choice functions* $\tilde{\sigma}$, which map each excess payoff vector π to a nonnegative vector $\tilde{\sigma}(\pi)$ called a *raw choice vector*. Revision rates are determined by the *sum* of the components of the raw choice vector, while choice probabilities are *proportional* to the components of this vector.

Rather than make specific assumptions about functional forms, we only require that raw choice functions satisfy two mild conditions that embody the notion of moderation described above. To rule out extreme sensitivity of decisions to the exact values of payoffs, we require raw choice functions to be *Lipschitz continuous*. In most contexts where evolutionary models are appropriate, excessive sensitivity of choice rules to payoffs seems unrealistic, making it natural to consider models that do not demand it. To link choices to payoffs, we impose a condition called *acuteness*: each excess payoff vector π and corresponding raw choice vector $\tilde{\sigma}(\pi)$ have a positive inner product. This condition ensures that whenever payoff improvement opportunities exist, revision

(1991), Matsui (1992), and Hofbauer (1995). Perturbed versions of the best response dynamic are discussed below.

4 See Björnerstedt and Weibull (1996) and Schlag (1998).

opportunities continue to arrive, and that revising agents show some tendency to select strategies with above average payoffs.

A game, a revision protocol, and an initial state define a Markov process through the space of strategy distributions. If the population size is large, this process is well-approximated by the solutions to a certain differential equation. This equation, called the *mean dynamic*, is defined in terms of the expected change in the population's behavior at the current population state.[5] We call the class of mean dynamics derived from our model *excess payoff dynamics*.

The main result of this paper shows that every excess payoff dynamic is well-behaved, in the sense of satisfying properties (EUC), (NS), and (PC). Therefore, unlike dynamics based on imitation and optimization, dynamics based on moderation satisfy all three of our evolutionary desiderata.

There is one canonical dynamic that satisfies all three of our desiderata: namely, the Brown-von Neumann-Nash (BNN) dynamic.[6] Interestingly, the BNN dynamic is 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 thus provides a microfoundation for the BNN dynamic, and it also shows that very little of the structure of this dynamic is needed to support the properties we seek.

One way to contend with the discontinuities inherent in the best response dynamic is to introduce perturbations of payoffs. The resulting smooth dynamics are known as perturbed best response dynamics; within this class, the logit dynamic is the best known special case.[7] Unlike the original best response dynamic, perturbed best response dynamics satisfy condition (EUC). However, they fail conditions (NS) and (PC): because of the payoff disturbances, the rest points of perturbed best response dynamics differ from the Nash equilibria of the underlying game, and the growth rates of these dynamics fail to be positively correlated with payoffs in a variety of regions of the state space.

In very rough terms, these violations of properties (NS) and (PC) are "small" when the payoff perturbations are "small", so that the perturbed best response dynamic is "close" to the exact best response dynamic. This observation may seem to suggest that how near a dynamic comes to satisfying these two desiderata depends on how close choices are to being optimal, or, alternatively, on how sensitive choices are to the exact

5 See Binmore and Samuelson (1999), Sandholm (2003), Benaïm and Weibull (2003), and Section 2.2 below.
6 See Brown and von Neumann (1950) and Section 2.3 below.
7 See Fudenberg and Levine (1998, Chapter 4), Hopkins (1999), Hofbauer (2000), Hofbauer and Hopkins (2003), and Hofbauer and Sandholm (2002, 2003).

value of the state.[8] The results in this paper show that these conclusions are false. The choice rules that underlie excess payoff dynamics can be quite remote from exact optimization, yet these dynamics always satisfy all three of our desiderata exactly.

The replicator dynamic and other imitative dynamics fail to be well-behaved because they violate Nash stationarity. Happily, we can use excess payoff dynamics to remedy this difficulty in a minimally intrusive fashion. In Section 4, we show that by modifying any imitative dynamic to an arbitrarily small extent, we can create new dynamics that satisfy all three of our desiderata. These new dynamics are convex combinations of the imitative dynamic and an arbitrary excess payoff dynamic; they can be derived from choice protocols that usually are based on imitation, but that occasionally rely on moderation.

Section 5 concludes the paper by describing two additional well-behaved dynamics, which we call the *target projection dynamic* and the *projection dynamic*.[9] We argue that both these two dynamics and all excess payoff dynamics can be derived from a common ancestor. We hope that this pedigree provides some preliminary steps toward a full characterization of well-behaved dynamics.

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 $S = \{1, \dots, n\}$ 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 $\Delta = \{x \in \mathbf{R}^n_+ : \sum_i x_i = 1\}$; 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 $\{x \in \Delta\colon\ Nx \in \mathbf{Z}^n\}$.

8 The latter sort of sensitivity might be measured, for example, by the dynamic's Lipschitz constant.

9 For the former dynamic, see Friesz *et. al.* (1994); for the latter, see Dupuis and Nagurney (1993), Nagurney and Zhang (1996), and Lahkar and Sandholm (2004).

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

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

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

2.2 Choice Rules and Revision Rates

We now describe our revision protocol. 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, but seems reasonable in many contexts.[10] 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, $\lambda(\pi)$

---

[10] Björnerstedt and Weibull (1996) and Weibull (1995, Section 4.4) derive the replicator dynamic and other imitative dynamics using choice protocols with variable revision rates.

$= \tilde{\sigma}_T(\pi) \equiv \sum_{j\in S} \tilde{\sigma}_j(\pi)$. 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:

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

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

To interpret the raw choice function directly, consider the rate at which agents currently playing strategies other than i switch to strategy i. Under the (implicit) assumption that the arrivals of revision opportunities and the choices made thereafter are independent, this rate is given by $\lambda(\pi)\ \sigma_i(\pi) = \tilde{\sigma}_i(\pi)$. Describing the model in this manner 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 to other strategies.

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

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

The first condition, *Lipschhitz continuity*, asks that raw choice weights be Lipschitz continuous functions of 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 (LC) precludes it.

The second condition, *acuteness*, requires that the excess payoff vector π and the raw choice vector $\tilde{\sigma}(\pi)$ have a positive inner product whenever π lies in the interior of $\mathbf{R}^n_*$. This condition 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. In other words, 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 raw choice function satisfying conditions (LC) and (A) takes a separable semilinear form:

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

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 (LC) 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 (2003), 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.[11]

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 $N\ \lambda(\hat{F}(x))\ \sigma_i(\hat{F}(x))\ dt$. 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 toward the "target state" defined by the current choice probability vector $\sigma(\hat{F}(x)) \in \Delta$, at a speed determined by the revision rate $\lambda(\hat{F}(x)) \in \mathbf{R}_+$.

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 (LC) and (A), we call this differential equation an *excess payoff dynamic*.

[11] 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.

2.3 Examples

2.3.1 The Brown-von Neumann-Nash Dynamic

If raw choice function takes 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), more recently reintroduced by Skyrms (1990), Swinkels (1992), and Weibull (1996), and further investigated by Hofbauer (2000), Berger and Hofbauer (2001), and Sandholm (2001).[12]

We can use this dynamic to demonstrate the importance of allowing revision rates to vary. Had we fixed the revision rate fixed at one, 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.[13] 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

---

[12] For the connection with Nash (1951), see Section 5.

[13] 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. As long as neither strategy is dominant, a jump of the sort noted above must occur.

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

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

(7) $$\sigma(\pi) = \underset{y\in\Delta}{\arg\max}\ y \cdot \pi.$$

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

$$\underset{y\in\Delta}{\arg\max}\ y \cdot \hat{F}(x) = \underset{y\in\Delta}{\arg\max}\ y \cdot Ax - x \cdot Ax = \underset{y\in\Delta}{\arg\max}\ y \cdot Ax \equiv B(x),$$

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 discontinuities that spawn these difficulties are consequences of exact optimization. Under moderation, raw choice weights cannot depend too finely on payoff opportunities; this coarseness ensures that solution trajectories are not only unique, but also continuous in the initial state.

2.4 Population Games

We conclude this section by introducing a more general class of strategic environments 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 $X^p = \{x^p \in \mathbf{R}_+^{n^p}: \sum_{i\in S^p} x_i^p = m^p\}$, while $X = \{x = (x^1, \dots, x^p) \in \mathbf{R}_+^n: x^p \in X^p\}$ is the set of overall strategy distributions.

The payoff function for strategy $i \in S^p$ is denoted $F_i^p: X \to \mathbf{R}$, 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 $F^p: X \to \mathbf{R}^{n^p}$ refer to the vector of payoff functions for strategies belonging to population p and let $F: X \to \mathbf{R}^n$ 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 $\bar{F}^p(x) = \frac{1}{m^p} x^p \cdot F^p(x)$. Hence, the *excess payoff* to strategy $i \in S^p$ is $\hat{F}_i^p(x) = F_i^p(x) - \bar{F}^p(x)$, while $\hat{F}^p(x) = F^p(x) - \mathbf{1}\bar{F}^p(x)$ is the *excess payoff vector* for population p.

State $x \in X$ is a *Nash equilibrium* of F if each strategy used at x is a best response to x. 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 $TX = \{z \in \mathbf{R}^n: \sum_{i\in S^p} z_i^p = 0 \text{ for all } p \in \mathcal{P}\}$, 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

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

Now let $\Delta^p = \{y^p \in \mathbf{R}_+^{n^p}: \sum_{i\in S^p} x_i^p = 1\}$ 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) $\qquad \dot{x}^p = m^p\,\tilde{\sigma}^p(\hat{F}^p(x)) - \tilde{\sigma}^p_T(\hat{F}^p(x))\,x^p$ for all $p \in \mathcal{P}$.

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

3. Properties of Excess Payoff Dynamics

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

(EUC) $\quad \dot{x} = V(x)$ 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.

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

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

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

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 if and only if no agent can unilaterally improve his payoffs.

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, notice that

$$\begin{aligned} cov(V^p(x), F^p(x)) &= \tfrac{1}{n^p}\textstyle\sum_{i\in S^p}(V^p_i(x) - \tfrac{1}{n^p}\sum_{j\in S^p}V^p_j(x))(F^p_i(x) - \tfrac{1}{n^p}\sum_{j\in S^p}F^p_j(x)) \\ &= \tfrac{1}{n^p}\textstyle\sum_{i\in S^p}(V^p_i(x) - 0)(F^p_i(x) - \tfrac{1}{n^p}\sum_{j\in S^p}F^p_j(x)) \end{aligned}$$

$$= \tfrac{1}{n^p}\Big(V^p(x)\cdot F^p(x)+\tfrac{1}{n^p}\sum_{i\in S^p}F_i^p(x)\sum_{i\in S^p}V_i^p(x)\Big)$$
$$= \tfrac{1}{n^p}\Big(V^p(x)\cdot F^p(x)\Big),$$

where the second and fourth equalities follow from the fact that $V(x) \in TX$. 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.[14]

We call a dynamic *well-behaved* if it satisfies properties (EUC), (NS), and (PC) regardless of the population game F being played. With this definition in hand, we can state our main result.

Theorem 3.1: *Every excess payoff dynamic is well-behaved.*

Property (EUC) is a direct consequence of the facts that excess payoff dynamics are Lipschitz continuous and that they are inward pointing on the boundary of X.[15] 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*) $x^p \cdot \hat{F}^p(x) = x^p \cdot (F^p(x) - \mathbf{1}\bar{F}^p(x)) = x^p \cdot F^p(x) - (x^p \cdot \mathbf{1})(\tfrac{1}{m^p}\; x^p \cdot F^p(x)) = 0.$

(*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

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

[14] For stronger monotonicity conditions, see Nachbar (1990), Samuelson and Zhang (1992), Ritzberger and Weibull (1995), and Hofbauer and Weibull (1996).

[15] To be inward pointing means that $V_i^p(x) \geq 0$ whenever $x_i^p = 0$. For a proof that Lipschitz continuity and the inward pointing property imply (EUC), see Appendix A.1 of Ely and Sandholm (2004).

$$= m^p \tilde{\sigma}^p(\hat{F}^p(x)) \cdot \hat{F}^p(x) - \tilde{\sigma}_T^p(\hat{F}^p(x)) x^p \cdot \hat{F}^p(x)$$
$$= m^p \tilde{\sigma}^p(\hat{F}^p(x)) \cdot \hat{F}^p(x) > 0. \blacksquare$$

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* (LC) *and* (A), *and let $\pi^p \in$ bd($\mathbf{R}_*^{n^p}$), so that the set of strategies earning average payoffs, $Z^p(\pi^p) = \{i \in S^p : \pi_i^p = 0\}$, 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 $\pi \in$ bd($\mathbf{R}_*^n$), $i \notin Z(\pi)$, and $j \in Z(\pi)$. For $\varepsilon > 0$, let $\pi(\varepsilon) = \pi + \varepsilon e_j \in$ int($\mathbf{R}_*^n$) (see Figure 1). 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 $\pi \in$ 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 $\pi(\varepsilon) = \pi - \varepsilon e_j + \varepsilon^2 e_j \in$ int($\mathbf{R}_*^n$) (see Figure 2). If $k \notin \{i, j\}$, then

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

Thus,

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

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

[[Insert Figures 1 and 2 about here]]

Figures 1 and 2: Sequences of excess payoff vectors that approach $\mathrm{bd}(\mathbf{R}^n_*)$.

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 imply properties (NS) and (PC).

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

(*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 $\hat{F}_i^p(x) = F_j^p(x) - \bar{F}^p(x) \leq 0$ for all $i \in S^p$, with equality whenever $x_i^p > 0$. Hence, $\hat{F}^p(x) \in \mathrm{bd}(\mathbf{R}^{n^p}_*)$.

Second, we show that (*ii*) implies (*i*). Suppose that $\hat{F}^p(x) \in \mathrm{bd}(\mathbf{R}^{n^p}_*)$, 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 $\hat{F}^p(x) \in \mathrm{bd}(\mathbf{R}^{n^p}_*)$, so that $Z^p(\hat{F}^p(x)) = \operatorname{argmax}_{j\in S^p} \hat{F}^p_j(x) = \operatorname{argmax}_{j\in S^p} F^p_j(x)$. 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 $x^p_k = 0$ for all $k \neq i$, and so $x^p = m^p\, e^p_i$. Now Lemma 3.3(*ii*) tells us that $\tilde{\sigma}^p(\hat{F}^p(x)) = c\,e^p_i$ for some $c \geq 0$. Hence,

$$\begin{aligned} V^p(x) &= m^p\tilde{\sigma}^p_i(\hat{F}^p(x)) - x^p_i\tilde{\sigma}^p_T(\hat{F}^p(x)) \\ &= m^p(c\,e^p_i) - (m^p e^p_i)c = \mathbf{0}, \end{aligned}$$

which is statement (*iii*).

For the second case, suppose that $Z^p(\hat{F}^p(x)) \geq 2$. 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 $\hat{F}^p(x) \in \mathrm{int}(\mathbf{R}^{n^p}_*)$. 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 $\hat{F}^p(x) \in \mathrm{int}(\mathbf{R}^{n^p}_*)$, and Proposition 3.4 shows that condition (PC) holds when $\hat{F}^p(x) \in \mathrm{bd}(\mathbf{R}^{n^p}_*)$, 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 describes the behavior of agents who use decision procedures based on imitation, justifying the application of this dynamic in economic models.[16]

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

$$\dot{x}_i^p = I_i^p(x) = x_i^p g_i^p(x)$$

that exhibit *monotone percentage growth rates*:[17]

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

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.[18] 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 imitative dynamics. A typical formulation of a perturbed dynamic is

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

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

16 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.

17 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)).

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

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

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

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 moderate efforts to choose a strategy that performs well, whether or not it is currently in use. As Theorem 4.1 shows, this modification eliminates non-Nash rest points of the imitative dynamic, and does so without disturbing the dynamic's other desirable properties.

Theorem 4.1: *If $\alpha \in (0, 1]$, the dynamic C_α is well-behaved.*

Proof: In the Appendix.

The intuition behind 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). This fact allows us to recover condition (NS) while preserving our other desiderata.

In addition to being well-behaved, many combined dynamics C_α have another appealing property: the local stability of their rest points only depends on the rest points' stability under the imitative dynamic I. As an example, consider a single population game F, and the excess payoff dynamic $V(x) = \tilde{\sigma}(\hat{F}(x)) - \tilde{\sigma}_T(\hat{F}(x))x$ driven by the raw choice function $\tilde{\sigma}_i(\pi) = ([\pi_i]_+)^k$, where the exponent k is strictly greater than one. Then $\tilde{\sigma}$ is well defined on all of $\mathbf{R}^n$, and one can check that its derivative matrix $D\tilde{\sigma}(\pi)$ is the zero matrix whenever the excess payoff vector π lies on the boundary of $\mathbf{R}^n_*$. Proposition 3.4 shows that the excess payoff vector $\hat{F}(x)$ lies on this boundary whenever x is a Nash equilibrium. It follows that the derivative matrix $DV(x)$ of the dynamic V is the zero matrix at any equilibrium, and consequently that

$$DC_\alpha(x) = (1-\alpha)\, DI(x) + \alpha\, DV(x) = (1-\alpha)\, DI(x).$$

Therefore, if the Nash equilibrium x is a hyperbolic rest point of the imitative dynamic I,[19] then the eigenvalues of $DI(x)$ determine the stability of x not only under I, but also under the combined dynamic C_α.

We express this idea in somewhat greater generality in the following proposition.

Proposition 4.2: *Let x be a Nash equilibrium of F. Suppose that the raw choice functions $\tilde{\sigma}^p$ used to define the excess payoff dynamic V have null derivative matrices at $\hat{F}(x) \in \mathrm{bd}(\mathbf{R}^n_*)$, and that x is a hyperbolic rest point of the imitative dynamic I. Then x is an asymptotically stable rest point of C_α if and only if it is an asymptotically stable rest point of I, and it is an unstable rest point of C_α if and only if it is an unstable rest point of I.*

5. Additional Well-Behaved Dynamics

The properties that define excess payoff dynamics are sufficient conditions for an evolutionary dynamic to be well-behaved. Are there simple necessary conditions for a dynamic to be well-behaved? How close are these necessary conditions to the sufficient conditions studied here?

To conclude this paper, we present two additional well-behaved dynamics, neither of which are excess payoff dynamics. While doing so, we argue that excess payoff dynamics and these two new dynamics can all be viewed as descendents of a common ancestor. This genealogy may provide a first step toward answering the questions raised above.

For convenience, we consider games played by a single population, so that population states are elements of the simplex X. The common ancestor of all the dynamics we consider is the discrete-time "ur-dynamic"

[19] The rest point x is *hyperbolic* if the eigenvalues of the derivative matrix $DI(x)$ corresponding to eigenvectors in the tangent space TX have nonzero real parts.

(U) $$x_{t+1} \approx x_t + \left(F(x_t) - \mathbf{1}\pi_t\right).$$

In this expression, $\mathbf{1}$ is the vector of ones, and π_t represents a reference payoff at time t.

The right hand side of (U) need not define an element of the simplex, and we use the symbol "≈" to emphasize that (U) does not specify a well-defined dynamic. Yet there is a sense in which this expression captures a property intimately connected with well-behavedness: strategies whose payoffs are high should tend to be come more common, while strategies whose payoffs are low should tend to become less so.

If $\pi_t = \bar{F}(x_t) = x_t \cdot F(x_t)$ equals the average payoff in the population at time t, then (U) becomes

$$x_{t+1} \approx x_t + \hat{F}(x_t).$$

To turn this expression into a legitimate dynamic, we replace the excess payoff vector $\hat{F}(x_t)$ with a nonnegative proxy: the raw choice vector $\tilde{\sigma}(\hat{F}(x_t))$.

(8) $$x_{t+1} \approx x_t + \tilde{\sigma}(\hat{F}(x_t)).$$

The right hand side of expression (8) is a positive vector, so normalizing its components to sum to one yields a legitimate discrete-time dynamic:

(9) $$x_{t+1} = \frac{x_t + \tilde{\sigma}(\hat{F}(x_t))}{1 + \tilde{\sigma}_T(\hat{F}(x_t))}.$$

It is not difficult to show that if $\tilde{\sigma}$ is continuous and acute, then the rest points of (9) are the Nash equilibria of F. In fact, if $\tilde{\sigma}$ returns the positive parts of excess payoffs (i.e., if $\tilde{\sigma}_i(\pi) = [\pi_i]_+$), then equation (9) is the mapping used in Nash's (1951) proof of existence of equilibrium.

To obtain a continuous-time dynamic, we shorten the time increment and the state increment in expression (8) in this natural way:

$$x_{t+\varepsilon} \approx x_t + \varepsilon\tilde{\sigma}(\hat{F}(x_t)).$$

Normalizing again yields a well-defined discrete-time dynamic:

$$x_{t+\varepsilon} = \frac{x_t + \varepsilon\,\tilde{\sigma}(\hat{F}(x_t))}{1 + \varepsilon\,\tilde{\sigma}_T(\hat{F}(x_t))}.$$

To obtain a continuous-time dynamic, we rearrange this expression to obtain

$$\frac{x_{t+\varepsilon} - x_t}{\varepsilon} = \frac{\tilde{\sigma}(\hat{F}(x_t)) - \tilde{\sigma}_T(\hat{F}(x_t))x_t}{1 + \varepsilon\,\tilde{\sigma}_T(\hat{F}(x_t))}.$$

Then taking limits of both sides of this equation as ε goes to zero yields

(E) $$\dot{x}_t = \tilde{\sigma}(\hat{F}(x_t)) - \tilde{\sigma}_T(\hat{F}(x_t))x_t,$$

the excess payoff dynamic derived from excess payoff function $\tilde{\sigma}$. If $\tilde{\sigma}_i(\pi) = [\pi_i]_+$, then equation (E) is the BNN dynamic.

We obtain our two new well-behaved dynamics through a different method of salvaging the ur-dynamic (U). This time, we replace the vector $x_t + \left(F(x_t) - \mathbf{1}\,\pi_t\right)$ in (U) with the closest point to this vector in the simplex X. That is,

$$x_{t+1} = \Pi_X\left(x_t + \left(F(x_t) - \mathbf{1}\,\pi_t\right)\right),$$

where Π_X denotes the closest-point projection onto the closed convex set X. Actually, since the vector $\mathbf{1}$ is orthogonal to the simplex, the previous equation can be rewritten as

(10) $$x_{t+1} = \Pi_X\left(x_t + F(x_t)\right),$$

regardless of the specification of the reference payoff π_t.[20]

One way to derive a continuous-time dynamic from equation (10) is to consider stepping only ε of the way from x_t to $\Pi_X(x_t + F(x_t))$ during the first time interval of length ε:

$$x_{t+\varepsilon} = (1 - \varepsilon)x_t + \varepsilon\Pi_X(x_t + F(x_t)).$$

Rearranging this equation yields

$$\frac{x_{t+\varepsilon} - x_t}{\varepsilon} = \Pi_X(x_t + F(x_t)) - x_t,$$

and so taking ε to zero yields

(TP) $$\dot{x}_t = \Pi_X(x_t + F(x_t)) - x_t.$$

---

[20] Equation (10) is not new to game theory. Gul, Pearce, Stacchetti (1993) observe that the fixed points of this equation are the Nash equilibria of F. These authors credit Hartman and Stampacchia (1966) for introducing this map in a more general mathematical context.

We call the dynamic (TP) the *target projection dynamic*. To the best of our knowledge, this dynamic first appeared in the transportation science literature in the work of Friesz *et. al*. (1994). It is not difficult to verify that the target projection dynamic is well-behaved. Unfortunately, we do not know of an appealing way of deriving this dynamic from a model of individual choice.

Alternatively, we can derive a continuous time dynamic from equation (10) by reducing the size of the increment in the state *before* employing the projection Π_X:[21]

$$x_{t+\varepsilon} = \Pi_X(x_t + \varepsilon F(x_t)).$$

Subtracting x_t and dividing by ε on each side of this equation yields

$$\frac{x_{t+\varepsilon} - x_t}{\varepsilon} = \frac{\Pi_X(x_t + \varepsilon F(x_t)) - x_t}{\varepsilon}.$$

By taking the limit as ε approaches zero and appealing to a geometrically obvious fact from convex analysis,[22] we obtain the following differential equation:

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

The right hand side of equation (P) is the projection of the payoff vector $F(x_t)$ onto $TX(x_t)$, the cone of feasible directions of motion from state x_t. On the interior of the simplex, equation (P) immediately reduces to

$$\dot{x}_t = \hat{F}^u(x_t),$$

where

$$\hat{F}_i^u(x_t) = F_i(x_t) - \tfrac{1}{n}\mathbf{1} \cdot F(x)$$

is the excess payoff to strategy i over the *unweighted* average payoff. On the boundary of the simplex, describing the dynamic without using projections requires some

---

[21] This order of operations is analogous to the one we used when deriving the excess payoff dynamic earlier in this section. In that case, reversing the order of the two initial operations (i.e., normalizing before reducing the step size) results in a dynamic that is equivalent to the excess payoff dynamic up to a reparameterization of time.

[22] See, e.g., Proposition 3.5.3.5 of Hiriart-Urruty and Lemaréchal (1993).

additional effort. Lahkar and Sandholm (2004) call the dynamic (P) the *projection dynamic*.[23]

Since the set of feasible directions $TX(x_t)$ changes discontinuously as the boundary of the simplex is reached, the dynamic (P) is discontinuous. Nevertheless, results of Dupuis and Nagurney (1993), who study projection dynamics in a broader mathematical context, imply that the dynamic (P) nevertheless satisfies our existence, uniqueness, and continuity property (EUC). With this property in hand, it is not difficult to show that the projection dynamic respects properties (NS) and (PC), and so that it is well-behaved.

Unlike the target projection dynamic (TP), the projection dynamic (P) can be derived from a natural model of individual choice. It also has the attractive property of eliminating all iteratively strictly dominated strategies of the underlying game. For detailed analyses of the payoff projection dynamic, we refer the reader to Nagurney and Zhang (1996) and Lahkar and Sandholm (2004).

We have now described two well-behaved dynamics that are not excess payoff dynamics, but that like excess payoff dynamics are descended from the ur-dynamic (U). Necessary and sufficient conditions for dynamics to be well-behaved, as well as the relationship between these conditions and equation (U), are important topics for future research.

Appendix: Additional Proofs

The Proof of Proposition 2.1

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

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

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 $\pi_i + \frac{k+2}{c}$.

[23] On the interior of the simplex, the dynamic (P) is equivalent to Friedman's (1991) *linear dynamic* (see his Appendix A.1). However, there are important differences between the definitions of the two dynamics on the boundary of the simplex. For example, while Friedman's linear dynamic admits non-Nash rest points on bd(X), the projection dynamic satisfies Nash stationarity, and is even well-behaved. See Lahkar and Sandholm (2004) for further discussion.

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). ■

References

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 Brown-von Neumann-Nash Dynamics," unpublished manuscript, Vienna University of Economics and University of Vienna.

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.

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.

Dupuis, P., and A. Nagurney (1993). "Dynamical Systems and Variational Inequalities," *Ann. Oper. Res.* 44, 9-42.

Ely, J. C., and W. H. Sandholm (2004). "Evolution in Bayesian Games I: Theory," forthcoming, *Games Econ. Behav*.

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

Friesz, T. L., D. Bernstein, N. J. Mehta, R. L. Tobin, and S. Ganjalizadeh (1994). "Day-to-day Dynamic Network Disequilibria and Idealized Traveler Information Systems," *Oper. Res*. 42, 1120-1136.

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

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

Gul, F., D. Pearce, and E. Stacchetti (1993). "A Bound on the Proportion of Pure Strategy Equilibria in Generic Games," *Math. Oper. Res.* 18, 548-552.

Hartman, P., and G. Stampacchia (1966). "On Some Nonlinear Elliptic Differential Functional Equations," *Acta Math*. 115, 271-310.

Hirriart-Urruty, J.-B., and C. Lemaréchal (1993). *Convex Analysis and Minimization Algorithms I*. Berlin: Springer.

Hofbauer, J. (1995). "Stability for the Best Response Dynamics," 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 E. Hopkins (2003). "Learning in Perturbed Asymmetric Games," Forthcoming, *Games Econ. Behav*.

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 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.

Lahkar, R., and W. H. Sandholm (2004). "The Projection Dynamic," unpublished manuscript, University of Wisconsin.

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

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

Nagurney, A., and D. Zhang (1996). *Projected Dynamical Systems and Variational Inequalities with Applications*. Boston: Kluwer Academic Publishers.

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.

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

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.

Smirnov, G. V. (2002). *Introduction to the Theory of Differential Inclusions*. Providence, RI: American Mathematical Society.

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.

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.

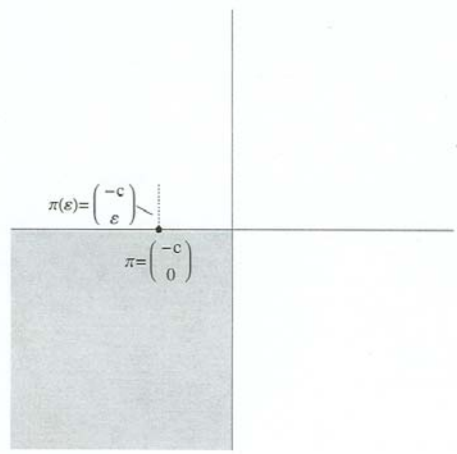

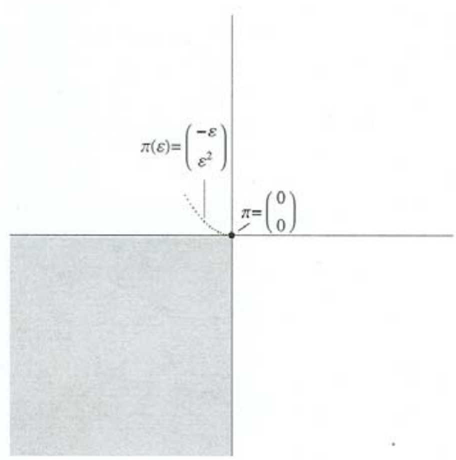

Figures 1 and 2