

Survival of Dominated Strategies under Evolutionary Dynamics*

Josef Hofbauer[†] and William H. Sandholm[‡]

September 7, 2009

Abstract

We prove that any deterministic evolutionary dynamic satisfying four mild requirements does not eliminate strictly dominated strategies in some games. We also show that existing elimination results for evolutionary dynamics are not robust to small changes in the specifications of the dynamics. Numerical analysis reveals that dominated strategies can persist at nontrivial frequencies even when the level of domination is not small.

1. Introduction

One fundamental issue in evolutionary game theory concerns the relationship between its predictions and those provided by traditional, rationality-based solution concepts. Indeed, much of the early interest in the theory among economists is due to its ability to justify traditional equilibrium predictions as consequences of myopic decisions made by simple agents.

Some of the best known results in this vein link the rest points of a deterministic evolutionary dynamic with the Nash equilibria of the game being played. Under most dynamics considered in the literature, the set of rest points includes all Nash equilibria of

*We thank Drew Fudenberg, Larry Samuelson, two anonymous referees, a Managing Editor, and many seminar audiences for helpful comments, and Emin Dokumaci for outstanding research assistance. Financial support from NSF Grants SES-0092145, SES-0617753, and SES-0851580, the Centre for Economic Learning and Social Evolution (ELSE) at University College London, and the University of Vienna is gratefully acknowledged.

[†]Department of Mathematics, University of Vienna, Nordbergstrasse 15, A-1090 Vienna, Austria. e-mail: josef.hofbauer@univie.ac.at; website: <http://homepage.univie.ac.at/josef.hofbauer>.

[‡]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>.

the underlying game, and under many of these dynamics the sets of rest points and Nash equilibria are identical.¹

In order to improve upon these results, one might look for dynamics that converge to Nash equilibrium from most initial conditions regardless of the game at hand. Such a finding would provide a strong defense of the Nash prediction, as agents who began play at some disequilibrium state could be expected to find their way to Nash equilibrium. Unfortunately, results of this kind cannot be proved. Hofbauer and Swinkels (1996) and Hart and Mas-Colell (2003) show that no reasonable evolutionary dynamic converges to Nash equilibrium in all games: there are some games in which cycling or more complicated limit behavior far from any Nash equilibrium is the only plausible long run prediction.

These negative results lead us to consider a more modest question. Rather than seek evolutionary support for equilibrium play, we instead turn our attention to a more basic rationality requirement: namely, the avoidance of strategies that are strictly dominated by a pure strategy.

Research on this question to date has led to a number of positive results. Two of the canonical evolutionary dynamics are known to eliminate strictly dominated strategies, at least from most initial conditions. Akin (1980) shows that starting from any interior population state, the replicator dynamic (Taylor and Jonker (1978)) eliminates strategies that are strictly dominated by a pure strategy. Samuelson and Zhang (1992), building on work of Nachbar (1990), extend this result to a broad class of evolutionary dynamics driven by imitation: namely, dynamics under which strategies' percentage growth rates are ordered by their payoffs.² Elimination results are also available for dynamics based on traditional choice criteria: the best response dynamic (Gilboa and Matsui (1991)) eliminates strictly dominated strategies by construction, as under this dynamic revising agents always switch to optimal strategies. Since the elimination of strategies strictly dominated by a pure strategy is the mildest requirement employed in standard game-theoretic analyses, it may seem unsurprising that two basic evolutionary dynamics obey this dictum.

In this paper, we argue that evolutionary support for the elimination of dominated strategies is more tenuous than the results noted above suggest. In particular, we prove that all evolutionary dynamics satisfying four mild conditions—continuity, positive correlation, Nash stationarity, and innovation—must fail to eliminate strictly dominated strategies in some games. Dynamics satisfying these conditions include not only well-

¹For results of the latter sort, see Brown and Neumann (1950), Smith (1984), and Sandholm (2005, 2008b).

²Samuelson and Zhang (1992) and Hofbauer and Weibull (1996) also introduce classes of imitative dynamics under which strategies strictly dominated by a mixed strategy are eliminated.

known dynamics from the evolutionary literature, but also slight modifications of the dynamics under which elimination is known to occur. In effect, this paper shows that the dynamics known to eliminate strictly dominated strategies in all games are the only ones one should expect to do so, and that even these elimination results are knife-edge cases.

An important predecessor of this study is the work of Berger and Hofbauer (2006), who present a game in which a strictly dominated strategy survives under the Brown-von Neumann-Nash (BNN) dynamic (Brown and von Neumann (1950)). We begin the present study by showing how Berger and Hofbauer's (2006) analysis can be extended to a variety of other dynamics, including the Smith dynamic (Smith (1984)) as well as generalizations of both the BNN and Smith dynamics (Hofbauer (2000), Sandholm (2005, 2008b)). While this analysis is relatively simple, it is not general, as it depends on the functional forms of the dynamics at issue. Since in practice it is difficult to know exactly how agents will update their choices over time, a more compelling elimination result would require only minimal structure.

Our main theorem provides such a result. Rather than specifying functional forms for the evolutionary dynamics under consideration, the theorem allows for any dynamic satisfying four mild conditions. The first, *continuity*, asks that the dynamic change continuously as a function of the payoff vector and the population state. The second, *positive correlation*, is a weak monotonicity condition: it demands that away from equilibrium, the correlation between strategies' payoffs and growth rates always be positive. The third condition, *Nash stationarity*, asks that states that are not Nash equilibria—that is, states where payoff improvement opportunities are available—are not rest points of the dynamic. The final condition, *innovation*, is a requirement that has force only at non-Nash boundary states: if at such a state some unused strategy is a best response, the growth rate of this strategy must be positive. The last three conditions rule out the replicator dynamic and the other purely imitative dynamics noted above; at the same time, they allow arbitrarily close approximations of these dynamics, under which agents usually imitate successful opponents, but occasionally select new strategies directly.

To prove the main theorem, we construct a four-strategy game in which one strategy is strictly dominated by another pure strategy. We show that under any dynamic satisfying our four conditions, the strictly dominated strategy survives along solution trajectories starting from most initial conditions.

Because evolutionary dynamics are defined by nonlinear differential equations, our formal results rely on topological properties, and so provide limited quantitative information about the conditions under which dominated strategies survive. We therefore supplement our formal approach with numerical analysis. This analysis reveals that

dominated strategies with payoffs substantially lower than those of their dominating strategies can be played at nontrivial frequencies in perpetuity.

Since elimination of dominated strategies is a basic requirement of traditional game theory, the fact that such strategies can persist under evolutionary dynamics may seem counterintuitive. A partial resolution of this puzzle lies in the fact that survival of dominated strategies is intrinsically a disequilibrium phenomenon.

To understand this point, remember that evolutionary dynamics capture the aggregate behavior of agents who follow simple myopic rules. These rules lead agents to switch to strategies whose current payoffs are good, though not necessarily optimal.

When a solution trajectory of an evolutionary dynamic converges, the payoffs to each strategy converge as well. Because payoffs become fixed, even simple rules are enough to ensure that only optimal strategies are chosen. In formal terms, the limits of convergent solution trajectories must be Nash equilibria; it follows *a fortiori* that when these limits are reached, strictly dominated strategies are not chosen.

Of course, it is well understood that solutions of evolutionary dynamics need not converge, but instead may enter limit cycles or more complicated limit sets.³ When solutions do not converge, payoffs remain in flux. In this situation, it is not obvious whether choice rules favoring strategies whose current payoffs are relatively high will necessarily eliminate strategies that perform well at many states, but that are never optimal. To the contrary, the analysis in this paper demonstrates that if play remains in disequilibrium, even strategies that are strictly dominated by other pure strategies can persist indefinitely.

One possible reaction to our results is to view them as an argument against the relevance of evolutionary dynamics for modeling economic behavior. If an agent notices that a strategy is strictly dominated, then he would do well to avoid playing it, whatever his rule of thumb might suggest. We agree with the latter sentiment: we do not expect agents, even simple ones, to play strategies they know to be dominated. At the same time, we feel that the ability to recognize dominated strategies should not be taken for granted. In complicated games with large numbers of participants, it may not always be reasonable to expect agents to know the payoffs to all strategies at every population state, or to be able to make all the comparisons needed to identify a dominated strategy. It is precisely in such large, complex games that agents might be expected to make decisions by applying rules of thumb. Our analysis suggests that if agents cannot directly exclude dominated strategies from their repertoire of choices, then these strategies need not fade from use through a lack of positive reinforcement.

³For a survey of nonconvergence results, see Sandholm (2008a). Some specific references are offered in the next two footnotes.

To prove our main result, we must show that for each member of a large class of deterministic evolutionary dynamics, there is a game in which dominated strategies survive. To accomplish this most directly, we use the same construction for all dynamics in the class. We begin by introducing a three-strategy game with nonlinear payoffs—the *hypnodisk game*—under which solution trajectories of all dynamics in the class enter cycles from almost all initial conditions. We then modify this game by adding a dominated fourth strategy, and show that the proportion of the population playing this strategy stays bounded away from zero along solutions starting from most initial conditions.

Since the game we construct to ensure cyclical behavior is rather unusual, one might wonder whether our survival results are of practical relevance, rather than being a mere artifact of a pathological construction. In fact, while introducing a special game is quite convenient for proving the main result, we feel that our basic message—that in the absence of convergence, myopic heuristics need not root out dominated strategies in large games—is of broader relevance. In Section 5.1, we explain why the proof of the main theorem does not depend on the introduction of a complicated game in an essential way. Analyses there and elsewhere in the paper suggest that in any game for which some dynamic covered by the theorem fails to converge, there are augmented games with dominated strategies that the dynamic allows to survive. There are many simple examples of nonconvergent evolutionary dynamics in the theoretical⁴ and applied theory⁵ literatures, suggesting that there is adequate ground for our analysis to have some practical import.

Section 2 introduces population games and evolutionary dynamics. Section 3 establishes the survival results for excess payoff dynamics and pairwise comparison dynamics, families that contain the BNN and Smith dynamics, respectively. Section 4 states and proves the main result. Section 5 presents our numerical analyses, and illustrates the sensitivity of existing elimination results to slight modifications of the dynamics in question. Section 6 concludes. Auxiliary results and proofs omitted from the text are provided in the Appendix.

⁴See Shapley (1964), Jordan (1993), Gaunersdorfer and Hofbauer (1995), Hofbauer and Swinkels (1996), Hart and Mas-Colell (2003), and Sparrow et al. (2008).

⁵See especially Hopkins and Seymour (2002) and Lahkar (2007), who argue that multiplicity of prices in models of price dispersion (Burdett and Judd (1983), Varian (1980)) is the result of cyclical rather than equilibrium behavior. For supporting empirical evidence, see Eckert (2003) and Noel (2007).

2. The Model

2.1 Population Games

We consider games played by a single unit mass population of agents.⁶ All agents choose from the finite set of strategies $S = \{1, \dots, n\}$. The set of population states is therefore the simplex $X = \{x \in \mathbf{R}_+^n : \sum_{i \in S} x_i = 1\}$, where x_i is the proportion of agents choosing strategy $i \in S$. The standard basis vector $e_i \in \mathbf{R}^n$ represents the state at which all agents choose strategy i .

If we take the set of strategies as fixed, we can identify a game with a Lipschitz continuous *payoff function* $F : X \rightarrow \mathbf{R}^n$, which assigns each population state $x \in X$ a vector of payoffs $F(x) \in \mathbf{R}^n$. The component $F_i : X \rightarrow \mathbf{R}$ represents the payoffs to strategy i alone. We also let $\bar{F}(x) = \sum_{i \in S} x_i F_i(x)$ denote the population's average payoff, and let $B^F(x) = \operatorname{argmax}_{y \in X} y'F(x)$ denote the set of (mixed) best responses at population state x .

The simplest examples of population games are generated by 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 agents are randomly matched to play this game, the (expected) payoff to strategy i at population state x is $F_i(x) = \sum_{j \in S} A_{ij}x_j$; hence, the population game associated with A is the linear game $F(x) = Ax$.

While random matching generates population games with linear payoffs, many population games that arise in applications have payoffs that are nonlinear in the population state—see Section 5.1. Games with nonlinear payoff functions will play a leading role in the analysis to come.

2.2 Evolutionary Dynamics

An evolutionary dynamic assigns each population game F an ordinary differential equation $\dot{x} = V^F(x)$ on the simplex X . One simple and general way of defining an evolutionary dynamic is via a growth rate function $g : \mathbf{R}^n \times X \rightarrow \mathbf{R}^n$; here $g_i(\pi, x)$ represents the (absolute) growth rate of strategy i as a function of the current payoff vector $\pi \in \mathbf{R}^n$ and the current population state $x \in X$. Our notation suppresses the dependence of g on the number of strategies n .

To ensure that the simplex is forward invariant under the induced differential equa-

⁶Versions of our results can also be proved in multipopulation models.

tions, the function g must satisfy

$$g_i(\pi, x) \geq 0 \text{ whenever } x_i = 0, \text{ and } \sum_{i \in S} g_i(\pi, x) = 0.$$

In words: strategies that are currently unused cannot become less common, and the sum of all strategies' growth rates must equal zero. A growth rate function g satisfying these conditions defines an evolutionary dynamic as follows:

$$\dot{x}_i = V_i^F(x) = g_i(F(x), x).$$

One can also build evolutionary dynamics from a more structured model that not only provides explicit microfoundations for the dynamics, but also is inclusive enough to encompass all dynamics considered in the literature.⁷ In this model, the growth rate function g is replaced by a *revision protocol* $\rho : \mathbf{R}^n \times X \rightarrow \mathbf{R}_+^{n \times n}$, which describes the process through which individual agents 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 $\rho_{ij}(\pi, x)$. Aggregate behavior in the game F is then described by the differential equation

$$(1) \quad \dot{x}_i = V_i^F(x) = \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x),$$

which is known as the *mean dynamic* generated by ρ and F . The first term in (1) captures the inflow of agents into strategy i from other strategies, while the second term captures the outflow of agents from strategy i to other strategies.

Table I presents four basic examples of evolutionary dynamics, along with revision protocols that generate them. Further discussion of these dynamics can be found in Sections 3.1, 5.3, and 5.4 below.

3. Survival under the BNN, Smith, and Related Dynamics

Using a somewhat informal analysis, Berger and Hofbauer (2006) argue that strictly dominated strategies can survive under the BNN dynamic (Brown and von Neumann (1950)). To prepare for our main result, we formalize and extend Berger and Hofbauer's (2006) arguments to prove a survival result for two families of evolutionary dynamics;

⁷For explicit accounts of microfoundations, see Benaïm and Weibull (2003) and Sandholm (2003).

Revision protocol	Evolutionary dynamic	Name	Origin
$\rho_{ij} = x_j[F_j - F_i]_+$	$\dot{x}_i = x_i(F_i(x) - \bar{F}(x))$	replicator	Taylor and Jonker (1978)
$\rho_{ij} = B_j^F(x)$	$\dot{x} \in B^F(x) - x$	best response	Gilboa and Matsui (1991)
$\rho_{ij} = [F_j - \bar{F}]_+$	$\dot{x}_i = [F_i(x) - \bar{F}(x)]_+$ $-x_i \sum_{j \in S} [F_j(x) - \bar{F}(x)]_+$	BNN	Brown and von Neumann (1950)
$\rho_{ij} = [F_j - F_i]_+$	$\dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+$ $-x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+$	Smith	Smith (1984)

Table I: Four evolutionary dynamics and their revision protocols.

these families include the BNN dynamic and the Smith dynamic (Smith (1984)) as their simplest members.

3.1 Excess Payoff Dynamics and Pairwise Comparison Dynamics

The two families of dynamics we consider are based on revision protocols of the forms

$$(2) \quad \rho_{ij} = \phi(F_j - \bar{F}) \text{ and}$$

$$(3) \quad \rho_{ij} = \phi(F_j - F_i),$$

where in each case, $\phi : \mathbf{R} \rightarrow \mathbf{R}_+$ is a Lipschitz continuous function satisfying

$$(4) \quad \text{sgn}(\phi(u)) = \text{sgn}([u]_+) \text{ and } \left. \frac{d}{du^+} \phi(u) \right|_{u=0} > 0.$$

The families of evolutionary dynamics obtained by substituting expressions (2) and (3) into the mean dynamic (1) are called *excess payoff dynamics* (Weibull (1996); Hofbauer (2000); Sandholm (2005)) and *pairwise comparison dynamics* (Sandholm (2008b)), respectively. The BNN and Smith dynamics are the prototypical members of these two families: examining Table I, we see that these two dynamics are those obtained from protocols (2) and (3) when ϕ is the semilinear function $\phi(u) = [u]_+$.

Protocols of forms (2) and (3) describe distinct revision processes. Under (2), an agent who receives a revision opportunity has a positive probability of switching to any strategy

whose payoff exceeds the population’s average payoff; the agent’s current payoff has no bearing on his switching rates. Under (3), an agent who receives a revision opportunity has a positive probability of switching to any strategy whose payoff exceeds that of his current strategy. While the latter protocols lead to mean dynamics with more complicated functional forms (compare the BNN and Smith dynamics in Table I), they also seem more realistic than those of form (2): protocols satisfying (3) make an agent’s decisions depend on his current payoffs, and do not require him to know the average payoff obtained in the population as a whole.

3.2 Theorem and Proof

Theorem 3.1 shows that excess payoff dynamics and pairwise comparison dynamics allow dominated strategies to survive in some games.⁸

Theorem 3.1. *Suppose that V is an evolutionary dynamic based on a revision protocol ρ of form (2) or (3), where the function ϕ satisfies condition (4). Then there is a game F_d such that under V^{F_d} , along solutions from most initial conditions, there is a strictly dominated strategy played by a fraction of the population that is bounded away from 0 and that exceeds $\frac{1}{6}$ infinitely often as time approaches infinity.*

While the computations needed to prove Theorem 3.1 differ according to the dynamic under consideration, the three main steps are always the same. First, we show that for each of the relevant dynamics, play converges to a limit cycle in the bad Rock-Paper-Scissors game (Figure 1). Second, we introduce a new strategy, Twin, which duplicates the strategy Scissors, and show that in the resulting four-strategy game, solutions to the dynamic from almost all initial conditions converge to a cycling attractor; this attractor sits on the plane where Scissors and Twin are played by equal numbers of agents, and has regions where both Scissors and Twin are played by more than one sixth of the population (Figure 2). Third, we uniformly reduce the payoff of the new strategy by d , creating a “feeble twin”, and use a continuity argument to show that the attractor persists (Figure 3). Since the feeble twin is a strictly dominated strategy, this last step completes the proof of the theorem.

We now present the proof in more detail, relegating some details of the argument to the Appendix.

Proof. Fix a dynamic V (i.e., a map from population games F to differential equations

⁸In the statements of Theorems 3.1 and 4.1, “most initial conditions” means all initial conditions outside an open set of measure ε , where $\varepsilon > 0$ is specified before the choice of the game F_d .




Figure 1: The Smith dynamic in bad RPS. Lighter shading represents higher speeds of motion.

$\dot{x} = V^F(x)$ generated by a revision protocol ρ that satisfies the conditions of the theorem. We will construct a game F_d in which a dominated strategy survives under V^{F_d} .

To begin, we introduce the bad Rock-Paper-Scissors game:

$$G(x) = Ax = \begin{pmatrix} 0 & -b & a \\ a & 0 & -b \\ -b & a & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}, \quad \text{where } b > a > 0.$$

(Since $b > a$, the cost of losing a match exceeds the benefit of winning a match.) For any choices of $b > a > 0$, the unique Nash equilibrium of G is $y^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Although our proof does not require this fact, it can be shown as a corollary of Lemma 3.2 below that y^* is unstable under the dynamic V^G .

Next, following Berger and Hofbauer (2006), we introduce a four-strategy game F , which we obtain from bad RPS by introducing an “identical twin” of Scissors.

$$(5) \quad F(x) = \tilde{A}x = \begin{pmatrix} 0 & -b & a & a \\ a & 0 & -b & -b \\ -b & a & 0 & 0 \\ -b & a & 0 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$

The set of Nash equilibria of F is the line segment $NE = \{x^* \in X : x^* = (\frac{1}{3}, \frac{1}{3}, \alpha, \frac{1}{3} - \alpha)\}$.




Figure 2: The Smith dynamic in “bad RPS with a twin”.




Figure 3: The Smith dynamic in “bad RPS with a feeble twin”.

We now present two lemmas that describe the behavior of the dynamic V^F for game F . The first lemma concerns the local stability of the set of Nash equilibria NE .

Lemma 3.2. *The set NE is a repeller under the dynamic V^F : there is a neighborhood U of NE such that all trajectories starting in $U - NE$ leave U and never return.*

The proof of this lemma, which is based on construction of appropriate Lyapunov functions, is presented in Appendix B.

Since V is an excess payoff dynamic or a pairwise comparison dynamic, the rest points of V^F are precisely the Nash equilibria of F (see Sandholm (2005, 2008b)). Therefore, Lemma 3.2 implies that solutions of V^F from initial conditions outside NE do not converge to rest points. Our next lemma constrains the limit behavior of these solutions.

Since the revision protocol ρ treats strategies symmetrically, and since Scissors and Twin always earn the same payoffs ($F_3(x) \equiv F_4(x)$), it follows that

$$\rho_{j3}(F(x), x) = \rho_{j4}(F(x), x) \quad \text{and} \quad \rho_{3j}(F(x), x) = \rho_{4j}(F(x), x) \quad \text{for all } x \in X.$$

These equalities yield a simple expression for the rate of change of the difference in utilizations of strategies 3 and 4:

$$(6) \quad \begin{aligned} \dot{x}_3 - \dot{x}_4 &= \left(\sum_{j \in S} x_j \rho_{j3} - x_3 \sum_{j \in S} \rho_{3j} \right) - \left(\sum_{j \in S} x_j \rho_{j4} - x_4 \sum_{j \in S} \rho_{4j} \right) \\ &= -(x_3 - x_4) \sum_{j \in S} \rho_{3j}(F(x), x). \end{aligned}$$

Since conditional switch rates ρ_{ij} are nonnegative by definition, equation (6) implies that the plane $P = \{x \in X : x_3 = x_4\}$ on which the identical twins receive equal weight is invariant under V^F , and that distance from P is nonincreasing under V^F . In fact, we can show further that

Lemma 3.3. *Solutions of the dynamic V^F starting outside the set NE converge to the plane P .*

Proving Lemma 3.3 is straightforward when ρ is of the excess payoff form (2), since in this case it can be shown that $\dot{x}_3 < \dot{x}_4$ whenever $x_3 > x_4$ and $x \notin NE$, and that $\dot{x}_3 - \dot{x}_4 > 0$ whenever $x_3 < x_4$ and $x \notin NE$. But when ρ is of the pairwise comparison form (3), one needs to establish that solutions to V^F cannot become stuck in regions where $\dot{x}_3 = \dot{x}_4$. The proof of Lemma 3.3 is provided in Appendix B.

Lemmas 3.2 and 3.3 imply that all solutions of V^F other than those starting in NE converge to an attractor \mathcal{A} , a set that is compact (see Appendix A), is disjoint from the set

NE , is contained in the invariant plane P , and encircles the Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{6}, \frac{1}{6})$ (see Figure 2). It follows that there are portions of \mathcal{A} where more than one sixth of the population plays Twin.

Finally, we modify the game F by making Twin “feeble”: in other words, by uniformly reducing its payoff by d :

$$F_d(x) = \tilde{A}_d x = \begin{pmatrix} 0 & -b & a & a \\ a & 0 & -b & -b \\ -b & a & 0 & 0 \\ -b-d & a-d & -d & -d \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$

If $d > 0$, strategy 4 is strictly dominated by strategy 3.

Increasing d from 0 continuously changes the game from F to F_d , and so continuously changes the dynamic from V^F to V^{F_d} (where continuity is with respect to the supremum norm topology). It thus follows from results on continuation of attractors (Theorem A.1 in Appendix A) that for small domination levels d , the attractor \mathcal{A} of V^F continues to an attractor \mathcal{A}_d that is contained in a neighborhood of \mathcal{A} , and that the basin of attraction of \mathcal{A}_d contains all points outside of a thin tube around the set NE .

On the attractor \mathcal{A} , the speed of rotation under V^F around the segment NE is bounded away from 0. Therefore, by continuity, the attractor \mathcal{A}_d of V^{F_d} must encircle NE , and so must contain states at which x_4 , the weight on the strictly dominated strategy Twin, is more than $\frac{1}{6}$. By the same logic, solutions of V^{F_d} that converge to \mathcal{A}_d have ω -limit sets with these same properties. In conclusion, we have shown that most solutions of V^{F_d} converge to the attractor \mathcal{A}_d , a set on which x_4 is bounded away from 0, and that these solutions satisfy $x_4 > \frac{1}{6}$ infinitely often in the long run. This completes the proof of Theorem 3.1. ■

It is worth noting that the number $\frac{1}{6}$, the bound that the weight on the dominated strategy continually exceeds, is not as large as possible. By replacing A , a cyclically symmetric version of bad Rock-Paper-Scissors, with an asymmetric version of this game, we can move the unstable Nash equilibrium from $y^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ to a state where the fraction of the population choosing Scissors is as close to 1 as desired (see Gaunersdorfer and Hofbauer (1995)). Then repeating the rest of the proof above, we find that the bound of $\frac{1}{6}$ in the statement of Theorem 3.1 can be replaced by any number less than $\frac{1}{2}$.

The analysis above makes explicit use of the functional forms of excess payoff and pairwise comparison dynamics. This occurs first in the proof of Lemma 3.2, which states that the set of Nash equilibria of “bad RPS with a twin” is a repeller. The Lyapunov functions used to prove this lemma depend on the dynamics’ functional forms; indeed,

there are evolutionary dynamics for which the equilibrium of bad RPS is attracting instead of repelling. Functional forms are also important in proving Lemma 3.3, which states that almost all solutions to dynamics from the two classes lead to the plane on which the identical twins receive equal weights. For arbitrary dynamics, particularly ones that do not respect the symmetry of the game, convergence to this plane is not guaranteed. To establish our main result, in which nothing is presumed about functional forms, both of these steps from the proof above will need to be replaced by more general arguments.

4. The Main Theorem

4.1 Statement of the Theorem

While the proof of Theorem 3.1 takes advantage of the functional forms of excess payoff and pairwise comparison dynamics, the survival of dominated strategies is a more general phenomenon. We now introduce a set of mild conditions that are enough to yield this result.

(C)	<i>Continuity</i>	g is Lipschitz continuous.
(PC)	<i>Positive correlation</i>	If $V^F(x) \neq \mathbf{0}$, then $V^F(x)'F(x) > 0$.
(NS)	<i>Nash stationarity</i>	If $V^F(x) = \mathbf{0}$, then $x \in NE(F)$.
(IN)	<i>Innovation</i>	If $x \notin NE(F)$, $x_i = 0$, and $e_i \in B^F(x)$, then $V_i^F(x) > 0$.

Continuity (C) requires that small changes in aggregate behavior or payoffs not lead to large changes in the law of motion $V^F(x) = g(F(x), x)$. Since discontinuous revision protocols can only be executed by agents with extremely accurate information, this condition seems natural in most contexts where evolutionary models are appropriate. Of course, this condition excludes the best response dynamic from our analysis, but it does not exclude continuous approximations thereof—see Section 5.4.

Positive correlation (PC) is a mild payoff monotonicity condition. It requires that whenever the population is not at rest, there is a positive correlation between strategies' growth rates and payoffs.⁹ From a geometric point of view, condition (PC) requires that the directions of motion $V^F(x)$ and the payoff vectors $F(x)$ always form acute angles with one another. This interpretation will be helpful for understanding the constructions to come.

⁹Requiring growth rates to respect payoffs appears to work against the survival of dominated strategies. At the same time, some structure must be imposed on the dynamics in order to make headway with our analysis, and we would hesitate to consider a dynamic that did not satisfy a condition in the spirit of (PC) as a general model of evolution in games.

Nash stationarity (NS) requires that the dynamic V^F only be at rest at Nash equilibria of F . This condition captures the idea that agents eventually recognize payoff improvement opportunities, preventing the population from settling down at a state where such opportunities are present.¹⁰

In a similar spirit, *innovation* (IN) requires that when a non-Nash population state includes an unused optimal strategy, this strategy's growth rate must be strictly positive. In other words, if an unplayed strategy is sufficiently rewarding, some members of the population will discover it and select it.

A few further comments about conditions (PC), (NS), and (IN) may be helpful in interpreting our results. First, condition (PC) is among the weakest monotonicity conditions proposed in the evolutionary literature.¹¹ Thus, our arguments that appeal to this condition are robust, in that they apply to any dynamic that respects the payoffs from the underlying game to some weak extent.

Second, since condition (PC) requires a positive correlation between growth rates and payoffs at all population states, it rules out evolutionary dynamics under which the boundary of the state space is repelling due to "mutations" or other forms of noise. Consequently, condition (PC) excludes the possibility that a dominated strategy survives for trivial reasons of this sort.

Third, conditions (PC), (NS), and (IN) all rule out dynamics based exclusively on imitation. At the same time, all of these conditions are satisfied by dynamics under which agents usually imitate, but occasionally evaluate strategies in a more direct fashion. We will present this idea in some detail in Section 5.3.

The main result of this paper is Theorem 4.1.

Theorem 4.1. *Suppose the evolutionary dynamic V satisfies (C), (PC), (NS), and (IN). Then there is a game F_a such that under V^{F_a} , along solutions from most initial conditions, there is a strictly dominated strategy played by a fraction of the population bounded away from 0.*

Before proceeding, we should point out that the conclusion of Theorem 4.1 is weaker than that of Theorem 3.1 in one notable respect: while Theorem 3.1 ensured that at least $\frac{1}{6}$ of the population would play the dominated strategy infinitely often, Theorem 4.1 only ensures that the strategy is always used by a proportion of the population bounded away from 0. The reason for this weaker conclusion is the absence of any assumption that the dynamic V^F treats different strategies symmetrically. Adding such a symmetry

¹⁰The converse of this condition, that all Nash equilibria are rest points, follows easily from condition (PC); see Sandholm (2001).

¹¹Conditions similar to (PC) have been proposed, for example, in Friedman (1991), Swinkels (1993), and Sandholm (2001).

assumption would allow us to recover the stronger conclusion. See Section 4.2.4 for further discussion.¹²

4.2 Proof of the Theorem

As we noted earlier, the proof of Theorem 3.1 took advantage of the functional forms of the dynamics at issue. Since Theorem 4.1 provides no such structure, its proof will require some new ideas.

Our first task is to construct a replacement for the bad RPS game. More precisely, we seek a three-strategy game in which dynamics satisfying condition (PC) will fail to converge to Nash equilibrium from almost all initial conditions. Our construction relies on the theory of potential games, developed in the normal form context by Monderer and Shapley (1996) and Hofbauer and Sigmund (1998) and in the population game context by Sandholm (2001, 2009).

4.2.1 Potential Games

A population game F is a *potential game* if there exists a continuously differentiable function $f : \mathbf{R}_+^n \rightarrow \mathbf{R}$ satisfying

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

Put differently, each strategy's payoff function must equal the appropriate partial derivative of the *potential function*:

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

Games satisfying this condition include common interest games and congestion games, among many others. A basic fact about potential games is that reasonable evolutionary dynamics increase potential: if the dynamic V^F satisfies condition (PC), then along each solution trajectory $\{x_t\}$ we have that

$$\frac{d}{dt}f(x_t) = \nabla f(x_t)' \dot{x}_t = F(x_t)' V^F(x_t) \geq 0,$$

¹²The proof of Theorem 4.1 establishes that the dynamic V^{F_d} for the game F_d admits an attractor on which the proportion of agents using a dominated strategy is bounded away from zero, and whose basin contains all initial conditions in X outside a set of small but positive measure. It therefore follows from Theorem A.1 that the dominated strategy would continue to survive if the dynamic were subject to small perturbations representing "evolutionary drift", as studied by Binmore and Samuelson (1999).

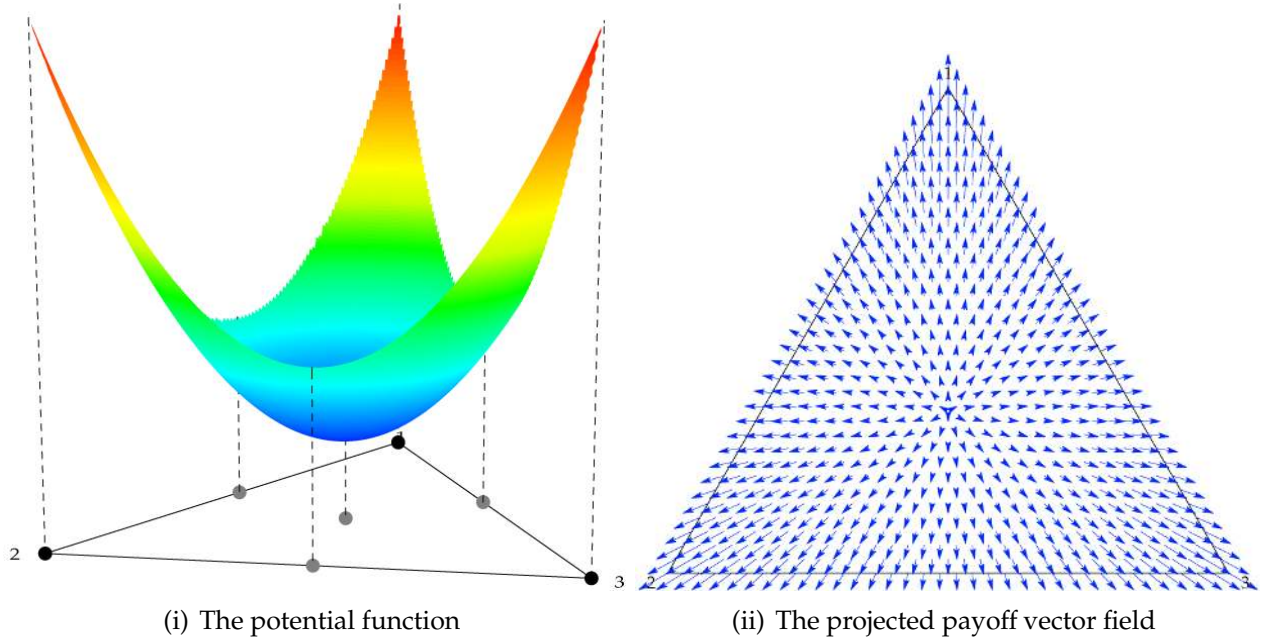


Figure 4: A coordination game.

with equality only at Nash equilibria. This observation along with standard results from dynamical systems imply that each solution trajectory of V^F converges to a connected set of Nash equilibria—see Sandholm (2001).

As an example, suppose that agents are randomly matched to play the pure coordination game

$$C = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The resulting population game, $F^C(x) = Cx = x$, is a potential game; its potential function, $f^C(x) = \frac{1}{2}x'Cx = \frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3)^2)$, is the convex function pictured in Figure 4(i). Solutions to any evolutionary dynamic satisfying condition (PC) ascend this function. Indeed, solutions from almost all initial conditions converge to a vertex of X —that is, to a strict equilibrium of F^C .

The ability to draw the game F^C itself will prove useful in the analysis to come. Notice that F^C is a map from the simplex $X \subset \mathbf{R}^3$ to \mathbf{R}^3 , and so can be viewed as a vector field. Rather than draw F^C as a vector field in \mathbf{R}^3 , we draw a projected version of F^C on the hyperplane in \mathbf{R}^3 that contains the simplex.¹³ The vectors drawn in Figure 4(ii) represent

¹³More precisely, we draw the vector field ΦF^C , where $\Phi = I - \frac{1}{3}\mathbf{1}\mathbf{1}' \in \mathbf{R}^{3 \times 3}$ is the orthogonal projection of \mathbf{R}^3

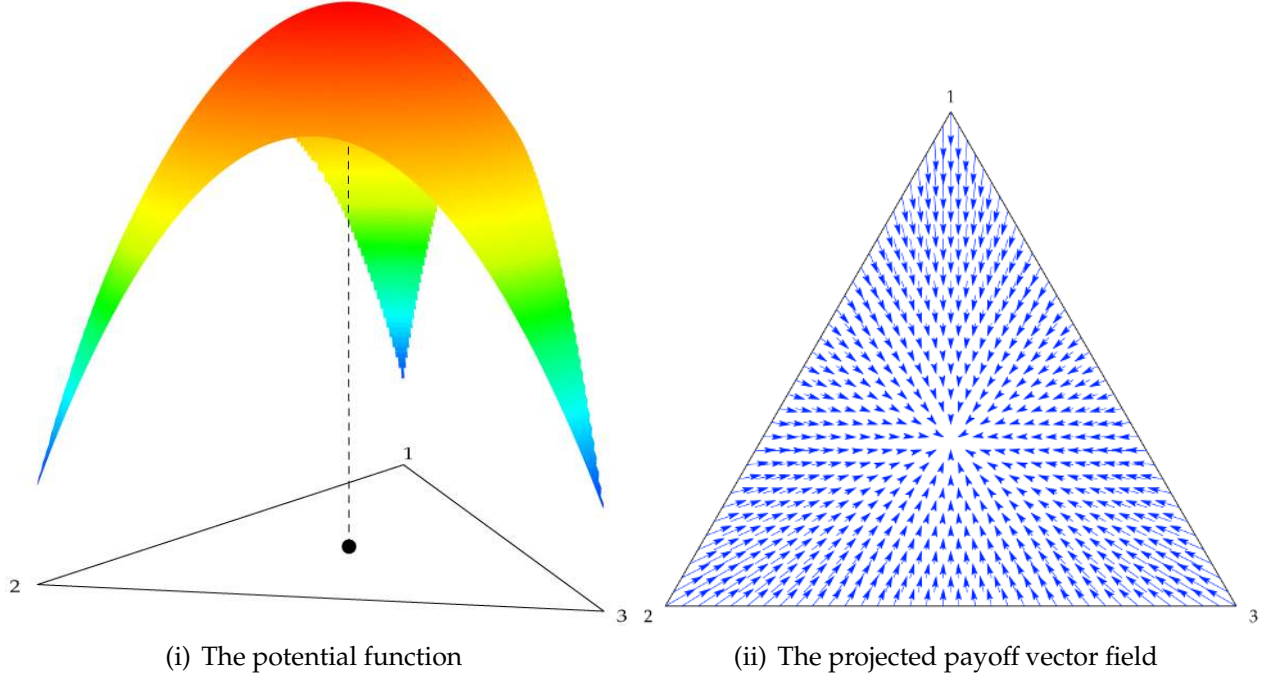


Figure 5: An anticoordination game.

the directions of maximal increase of the function f^C , and so point outward from the center of the simplex. Dynamics satisfying condition (PC) always travel at acute angles to the vectors in Figure 4(ii), and so tend toward the vertices of X , and solutions from almost all initial conditions converge to a vertex of X


As a second example, suppose that agents are randomly matched to play the anticoordination game $-C$. In Figures 5(i) and 5(ii), we draw the resulting population game $F^{-C}(x) = -Cx = -x$ and its concave potential function $f^{-C}(x) = -\frac{1}{2}x'Cx = -\frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3)^2)$. Both pictures reveal that under any evolutionary dynamic satisfying condition (PC), all solution trajectories converge to the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

4.2.2 The Hypnodisk Game


We now use the coordination game F^C and the anticoordination game F^{-C} to construct our replacement for bad RPS. While F^C and F^{-C} are potential games with linear payoffs, our new game will have neither of these properties.

The construction is easiest to describe in geometric terms. Begin with the coordination game $F^C(x) = Cx$ pictured in Figure 4(ii). Then draw two circles centered at state $x^* =$

onto $TX = \{z \in \mathbf{R}^3 : \sum_{i \in S} z_i = 0\}$, the tangent space of the simplex X . The projection Φ forces the components of $\Phi F^C(x)$ to sum to zero while preserving their differences, so that $\Phi F^C(x)$ preserves all information about incentives contained in payoff vector $F^C(x)$.



(i) Projected payoff vector field for the coordination game



(ii) Projected payoff vector field for the hypnodisk game

Figure 6: Construction of the hypnodisk game.

$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ with radii $0 < r < R < \frac{1}{\sqrt{6}}$, as shown in Figure 6(i); the second inequality ensures that both circles are contained in the simplex. Finally, twist the portion of the vector field lying outside of the inner circle in a clockwise direction, excluding larger and larger circles as the twisting proceeds, so that the outer circle is reached when the total twist is 180° . The resulting vector field is pictured in Figure 6(ii). It is described analytically by

$$H(x) = \cos(\theta(x)) \begin{pmatrix} x_1 - \frac{1}{3} \\ x_2 - \frac{1}{3} \\ x_3 - \frac{1}{3} \end{pmatrix} + \frac{\sqrt{3}}{3} \sin(\theta(x)) \begin{pmatrix} x_2 - x_3 \\ x_3 - x_1 \\ x_1 - x_2 \end{pmatrix} + \frac{1}{3} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix},$$

where $\theta(x)$ equals 0 when $|x - x^*| \leq r$, equals π when $|x - x^*| \geq R$, and varies linearly in between. We call the game H the *hypnodisk game*.

What does this construction accomplish? Inside the inner circle, H is identical to the coordination game F^C . Thus, solutions to dynamics satisfying (PC) starting at states in the inner circle besides x^* must leave the inner circle. At states outside the outer circle, the drawing of H is identical to the drawing of the antcoordination game F^{-C} .¹⁴ Therefore, solutions to dynamics satisfying (PC) that begin outside the outer circle must enter the outer circle. Finally, at each state x in the annulus bounded by the two circles, $H(x)$ is not a componentwise constant vector. Therefore, states in the annulus are not Nash equilibria, and so are not rest points of dynamics satisfying (PC). We assemble these observations in the following lemma.

Lemma 4.2. *Suppose that V is an evolutionary dynamic that satisfies conditions (C) and (PC), and let H be the hypnodisk game. Then every solution to V^H other than the stationary solution at x^* enters the annulus with radii r and R and never leaves.*

In fact, since there are no rest points in the annulus, the Poincaré-Bendixson Theorem implies that every nonstationary solution to V^H converges to a limit cycle.

4.2.3 The Twin

Now, let F be the four-strategy game obtained from H by adding a twin: $F_i(x_1, x_2, x_3, x_4) = H_i(x_1, x_2, x_3 + x_4)$ for $i \in \{1, 2, 3\}$, and $F_4(x) = F_3(x)$. The set of Nash equilibria of F is the line segment

$$NE = \left\{ x^* \in X : x_1^* = x_2^* = x_3^* + x_4^* = \frac{1}{3} \right\}.$$

¹⁴At states x outside the outer circle, $H(x) = -x + \frac{2}{3}\mathbf{1} \neq -x = F^{-C}(x)$. But since $\Phi H(x) = -x + \frac{1}{3}\mathbf{1} = \Phi F^{-C}(x)$ at these states, the pictures of H and F^{-C} , and hence the incentives in the two games, are the same.




Figure 7: Regions O , I , and $D = O - I$.

Let

$$I = \left\{ x \in X : (x_1 - \frac{1}{3})^2 + (x_2 - \frac{1}{3})^2 + (x_3 + x_4 - \frac{1}{3})^2 \leq r^2 \right\} \text{ and}$$

$$O = \left\{ x \in X : (x_1 - \frac{1}{3})^2 + (x_2 - \frac{1}{3})^2 + (x_3 + x_4 - \frac{1}{3})^2 \leq R^2 \right\}$$

be concentric cylindrical regions in X surrounding NE , as pictured in Figure 7. By construction, we have that

$$F(x) = \tilde{C}x = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$

at all $x \in I$. Therefore, solutions to dynamics satisfying (PC) starting in $I - NE$ ascend the potential function $f^{\tilde{C}}(x) = \frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3 + x_4)^2)$ until leaving the set I . At states outside the set O , we have that $F(x) = -\tilde{C}x$, so solutions starting in $X - O$ ascend $f^{-\tilde{C}}(x) = -f^{\tilde{C}}(x)$ until entering O . In summary:




Figure 8: The best response correspondence of the hypnodisk game.

Lemma 4.3. *Suppose that V is an evolutionary dynamic that satisfies conditions (C) and (PC), and let F be the “hypnodisk with a twin” game. Then every solution to V^F other than the stationary solutions at states in NE enter region $D = O - I$ and never leave.*

4.2.4 The Feeble Twin

To prove Theorem 3.1, we argued in Lemma 3.3 that under any of the dynamics addressed by the theorem, nonstationary solution trajectories equalize the utilization levels of identical twin strategies. If we presently focus on dynamics that not only satisfy conditions (C), (PC), and (IN), but also treat different strategies symmetrically, we can argue that in the “hypnodisk with a twin” game F , all nonstationary solutions of V^F converge not only to region D , but also to the plane $P = \{x \in X : x_3 = x_4\}$. Continuing with the argument from Section 3 then allows us to conclude that in F_d , the game obtained from F by turning strategy 4 into a feeble twin (that is, by reducing the payoff to strategy 4 uniformly by $d > 0$), the fraction x_4 playing the feeble twin exceeds $\frac{1}{6}$ infinitely often.

Since we would like a result that imposes as little structure as possible on permissible evolutionary dynamics, Theorem 4.1 avoids the assumption that different strategies are treated symmetrically. Since this means that agents may well be biased against choosing the dominated strategy, we can no longer prove that the fraction playing it will repeatedly exceed $\frac{1}{6}$. But we can still prove that the dominated strategy survives. To accomplish

this, it is enough to show that in game F , most solutions of the dynamic V^F converge to a set on which x_4 is bounded away from 0. If we can do this, then repeating the continuity argument that concluded the proof of Theorem 3.1 shows that in the game F_d , the dominated strategy 4 survives.

A complete proof that most solutions of V^F converge to a set on which x_4 bounded away from 0 is presented in Appendix C. We summarize the argument here. To begin, it can be shown that all solutions to V^F starting outside a small neighborhood of the segment of Nash equilibria NE converge to an attractor \mathcal{A} , a compact set that is contained in region D and that is an invariant set of the dynamic V^F .

Now suppose by way of contradiction that the attractor \mathcal{A} intersects $Z = \{x \in X : x_4 = 0\}$, the face of X on which Twin is unused. The Lipschitz continuity of the dynamic V^F implies that Z is backward invariant under V^F . Since \mathcal{A} is itself invariant, the fact that \mathcal{A} intersects Z implies the existence of a closed orbit $\gamma \subset \mathcal{A} \cap Z$ that circumnavigates the disk $I \cap Z$. Examining the best response correspondence of the hypnodisk game (Figure 8), we find that such an orbit γ must pass through a region in which strategy 3 is a best response. But since the twin strategy 4 is also a best response in this region, innovation (IN) tells us that solutions passing through this region must reenter the interior of X , contradicting that the attractor \mathcal{A} intersects the face Z .

5. Discussion

5.1 Constructing Games in which Dominated Strategies Survive

If an evolutionary dynamic satisfies monotonicity condition (PC), all of its rest points are Nash equilibria. It follows that dominated strategies can only survive on solution trajectories that do not converge to rest points. To construct games in which dominated strategies can survive, one should first look for games in which convergence rarely occurs.

The hypnodisk game, the starting point for the proof of the main theorem, is a population game with nonlinear payoff functions. Such games were uncommon in the early literature on evolution in games, which focused on random matching settings. But population games with nonlinear payoffs are more common now, in part because of their appearance in applications. For example, the standard model of driver behavior in a highway network is a congestion game with nonlinear payoff functions, as delays on each network link are increasing, convex functions of the number of drivers using the link.¹⁵

¹⁵Congestion games with a continuum of agents are studied by Beckmann et al. (1956) and Sandholm (2001). For finite player congestion games, see Rosenthal (1973) and Monderer and Shapley (1996).

For this reason, we do not view the use of a game with nonlinear payoffs as a shortcoming of our analysis. But despite this, it seems worth asking whether our results could be proved within the linear, random matching framework.

In Section 3, where we considered dynamics with prespecified functional forms, we were able to prove survival results within the linear setting. More generally, if we fix an evolutionary dynamic before seeking a population game, finding a linear game that exhibits cycling seems a feasible task. Still, a virtue of our analysis in Section 4 is that it avoids this case-by-case analysis: the hypnodisk game generates cycling under all of the relevant dynamics simultaneously, enabling us to prove survival of dominated strategies under all of these dynamics at once.

Could we have done the same using linear payoffs? Consider the following game due to Hofbauer and Swinkels (1996) (see also (Hofbauer and Sigmund, 1998, Sec 8.6)):

$$F^\varepsilon(x) = A^\varepsilon x = \begin{pmatrix} 0 & 0 & -1 & \varepsilon \\ \varepsilon & 0 & 0 & -1 \\ -1 & \varepsilon & 0 & 0 \\ 0 & -1 & \varepsilon & 0 \end{pmatrix} x.$$

When $\varepsilon = 0$, the game F^0 is a potential game with potential function $f(x) = -(x_1x_3 + x_2x_4)$. It has two components of Nash equilibria: one is a singleton containing the completely mixed equilibrium $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$; the other is the closed curve γ containing edges $\overline{e_1e_2}$, $\overline{e_2e_3}$, $\overline{e_3e_4}$, and $\overline{e_4e_1}$. The former component is a saddle point of f , and so is unstable under dynamics that satisfy (PC); the latter component is the maximizer set of f , and so attracts most solutions of these dynamics.

If ε is positive but sufficiently small, Theorem A.1 implies that most solutions of dynamics satisfying (PC) lead to an attractor near γ . But once ε is positive, the unique Nash equilibrium of F^ε is the mixed equilibrium x^* . Therefore, the attractor near γ is far from any Nash equilibrium.

If we now introduce a feeble twin, we expect that this dominated strategy would survive in the resulting five-strategy game. But in this case, evolutionary dynamics run on a four-dimensional state space. Proving survival results when the dimension of the state space exceeds three is very difficult, even if we fix the dynamic under consideration in advance. This points to another advantage of the hypnodisk game: it allows us to work with dynamics on a three-dimensional state space, where the analysis is still tractable.

5.2 How Dominated Can Surviving Strategies Be?

Since the dynamics we consider are nonlinear, our proofs of survival of dominated strategies are topological in nature, and so do not quantify the level of domination that is consistent with a dominated strategy maintaining a significant presence in the population. We can provide a sense of this magnitude by way of numerical analysis.

Our analysis considers the behavior of the BNN and Smith dynamics in the following version of “bad RPS with a feeble twin”:

$$(7) \quad F_d(x) = \tilde{A}_d x = \begin{pmatrix} 0 & -2 & 1 & 1 \\ 1 & 0 & -2 & -2 \\ -2 & 1 & 0 & 0 \\ -2-d & 1-d & -d & -d \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$


Figure 9(i) presents the maximum, time-average, and minimum weight on the dominated strategy in the limit cycle of the BNN dynamic, with these weights being presented as functions of the domination level d . The figure shows that until the dominated strategy Twin is eliminated, its presence declines at roughly linear rate in d . Twin is played recurrently by at least 10% of the population when $d \leq .14$, by at least 5% of the population when $d \leq .19$, and by at least 1% of the population when $d \leq .22$.

Figure 9(ii) shows that under the Smith dynamic, the decay in the use of the dominated strategy is much more gradual. In this case, Twin is recurrently played by at least 10% of the population when $d \leq .31$, by at least 5% of the population when $d \leq .47$, and by at least 1% of the population when $d \leq .66$. These values of d are surprisingly large relative to the base payoff values of 0, -2 , and 1: even strategies that are dominated by a significant margin can be played in perpetuity under common evolutionary dynamics.


The reason for the difference between the two dynamics is easy to explain. As we saw in Section 3, the BNN dynamic describes the behavior of agents who compare a candidate strategy’s payoff with the average payoff in the population. For its part, the Smith dynamic is based on comparisons between the candidate strategy’s payoff and an agent’s current payoff. The latter specification makes it relatively easy for agents obtaining a low payoff from Paper or Rock to switch to the dominated strategy Twin.

5.3 Exact and Hybrid Imitative Dynamics

An important class of dynamics excluded by our results are imitative dynamics, a class that includes the replicator dynamic as its best-known example. In general, imitative



(i) BNN



(ii) Smith

Figure 9: The maximum, time-average, and minimum weight on the dominated strategy in the limit cycles of the BNN and Smith dynamics. These weights are presented as functions of the domination level d in game (7).

dynamics are derived from revision protocols of the form

$$\rho_{ij}(\pi, x) = x_j r_{ij}(\pi, x).$$

The x_j term reflects the fact that when an agent receives a revision opportunity, he selects an opponent at random, and then decides whether or not to imitate this opponent's strategy. Substituting ρ into equation (1), we see that imitative dynamics take the simple form

$$(8) \quad \begin{aligned} \dot{x}_i &= x_i \sum_{j \in S} x_j (r_{ji}(F(x), x) - r_{ij}(F(x), x)) \\ &\equiv x_i p_i(F(x), x). \end{aligned}$$

In other words, each strategy's absolute growth rate \dot{x}_i is proportional to its level of utilization x_i .

To see the consequences of this for dominated strategies, use equation (8) and the quotient rule to obtain

$$(9) \quad \frac{d}{dt} \left(\frac{x_i}{x_j} \right) = \frac{x_i}{x_j} (p_i(F(x), x) - p_j(F(x), x)).$$

Now suppose that *percentage growth rates are monotone*, in the sense that

$$p_i(\pi, x) \geq p_j(\pi, x) \text{ if and only if } \pi_i \geq \pi_j,$$

Then if strategy i strictly dominates strategy j , the right hand side of equation (9) is positive at all $x \in \text{int}(X)$. We can therefore conclude that the dominated strategy j vanishes along every interior solution trajectory of (8). This is Samuelson and Zhang's (1992) elimination result.¹⁶

Equation (9) can be used to explain why elimination results for imitative dynamics are fragile. Suppose now that strategies i and j always earn the same payoffs. In this case, the right hand side of equation (9) is identically zero on $\text{int}(X)$, implying that the ratio $\frac{x_i}{x_j}$ is constant along every interior solution trajectory. For instance, in Figure 10(i), a phase diagram of the replicator dynamic in "standard RPS with a twin", we see that the planes on which the ratio $\frac{x_S}{x_T}$ is constant are invariant sets. If we make the twin feeble by lowering

¹⁶Sandholm et al. (2008) establish close links between the replicator dynamic and the projection dynamic of Nagurney and Zhang (1997). They show that on the interior of the simplex, these two dynamics share a property called "inflow-outflow symmetry", which ensures that dominated strategies lose ground to the strategies that dominate them. But the projection dynamic is discontinuous at the boundary of the simplex, and its behavior on the boundary can allow dominated strategies to survive.




Figure 10: The replicator dynamic in two games.

its payoff uniformly by d , we obtain the dynamics pictured in Figure 10(ii): now the ratio $\frac{x_S}{x_T}$ increases monotonically, and the dominated strategy is eliminated.

The existence of a continuum of invariant hyperplanes in games with identical twins is crucial to this argument. At the same time, dynamics with a continuum of invariant hyperplanes are structurally unstable. If we fix the game but slightly alter the agents' revision protocol, these invariant sets can collapse, overturning the elimination result.

As an example, suppose that instead of always following an imitative protocol, agents occasionally use a protocol that allows switches to unused strategies. This situation is illustrated in Figure 11(i), which contains the phase diagram for a "bad RPS with a twin" game under a convex combination of the replicator and Smith dynamics.¹⁷ While Figure 10(i) displayed a continuum of invariant hyperplanes, Figure 11(i) shows almost all solution trajectories converging to a limit cycle on the plane where $x_S = x_T$. If we then make the twin feeble, the limit cycle moves slightly to the left, and the dominated strategy survives (Figure 11(ii)).

¹⁷In particular, we consider the bad RPS game with payoffs $0, -\frac{11}{10}$, and 1 , and the combined dynamic that puts weight $\frac{9}{10}$ on the replicator dynamic and weight $\frac{1}{10}$ on the Smith dynamic. This dynamic is generated by the corresponding convex combination of the underlying revision protocols: $\rho_{ij} = \frac{9}{10}x_j[F_j - F_i]_+ + \frac{1}{10}[F_j - F_i]_+$.

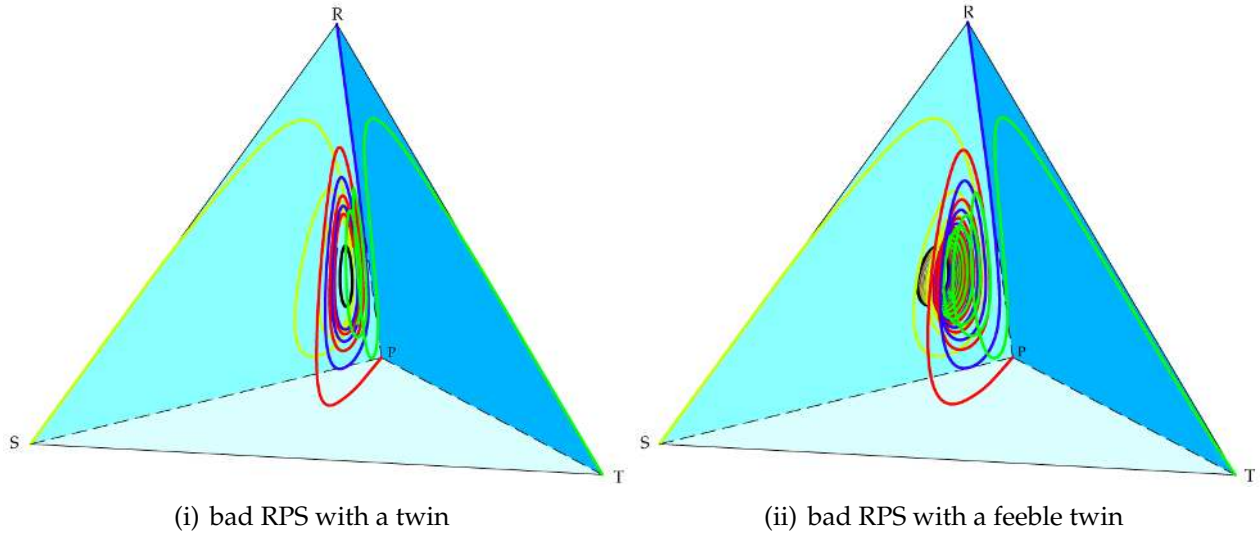


Figure 11: The $\frac{9}{10}$ replicator + $\frac{1}{10}$ Smith dynamic in two games.

5.4 Exact and Perturbed Best Response Dynamics

Of the basic evolutionary dynamics presented in Table I, it remains to consider the *best response dynamic* of Gilboa and Matsui (1991). The best response dynamic is defined by

$$\dot{x} \in B^F(x) - x, \quad \text{where } B^F(x) = \operatorname{argmax}_{y \in X} y'F(x)$$

is the set of mixed best responses to population state x . This dynamic describes the behavior of agents who occasionally receive opportunities to switch strategies, and switch to a best response whenever such an opportunity arises. It is obvious that the best response dynamic eliminates any strictly dominated strategy: since such strategies are never best responses, the weight on them vanishes at an exponential rate.

The best response dynamic is defined by a discontinuous differential inclusion. If we approximate the best response dynamic by a continuous differential equation—for instance, by a dynamic defined in terms of perturbed best responses—the resulting limit behavior can change dramatically, allowing dominated strategies to endure in significant proportions.¹⁸

¹⁸As an aside, we note that the limit behavior of the best response dynamic itself can change discontinuously as we change the payoffs of the underlying game. For instance, in game (7), a positive dominance level d leads the best response dynamic to have a triangular limit cycle on the face of the simplex where the dominated strategy Twin is unused (cf Gaunersdorfer and Hofbauer (1995)), while a negative value of d transfers this limit cycle to the face where the now-dominated strategy Scissors is unused. But when d is zero, so that Scissors and Twin are identical, the union of the three planes that connect the corresponding sides of the triangles is an attractor, and any point on this surface can be reached from any other. For more




Figure 12: The maximum weight on the dominated strategy in limit cycles of the $\text{logit}(\eta)$ dynamic, $\eta = .01, .05, .10$, and $.20$, in game (7). Weights are presented as functions of the domination level d .

To illustrate this, we consider the *logit dynamic* of Fudenberg and Levine (1998):

$$\dot{x} \in L^F(x) - x, \quad \text{where } L_i^F(x) = \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))}.$$

The logit dynamic describes the aggregate behavior of agents who update strategies by applying the logit rule L^F . When the *noise level* $\eta > 0$ is small, the logit rule places nearly all probability on optimal strategies, but this rule always places positive probability on all available strategies.¹⁹

Figure 12 presents the results of a numerical analysis of the logit dynamic in the “bad RPS with a feeble twin” game from equation (7). The four curves represent the maximum weight on the dominated strategy Twin in the stable limit cycle of the $\text{logit}(\eta)$ dynamic for noise levels $\eta = .01, .05, .10$, and $.20$. In each case, the weight on Twin is presented as a function of the domination level d in game (7).

When both the noise level η and the domination level d are very close to 0, the weight

on attractors of differential inclusions, see Benaïm et al. (2005).

¹⁹To interpret the analysis below, note that a noise level of η corresponds to the introduction of i.i.d. extreme-value distributed payoff disturbances with standard deviation $\pi\eta/\sqrt{6} \approx 1.28\eta$; see Anderson et al. (1992) or Hofbauer and Sandholm (2002).

on the dominated strategy Twin recurrently approaches values of nearly $\frac{2}{7}$.²⁰ For small fixed values of η , the maximum weight on the dominated strategy falls rapidly as the domination d level increases.

Higher values of η introduce more randomness into agents' choices, creating a force pushing the population state toward the center of the simplex. This inward force reduces the maximal weight on the dominated strategy at low values of d , but allows the dominated strategy to maintain a significant presence at considerably higher values of d .

6. Conclusion

Traditional game-theoretic analyses rule out strictly dominated strategies, as playing such strategies is inconsistent with decision-theoretic rationality. This paper argues that in settings where evolutionary game models are appropriate, the justification for eliminating dominated strategies is far less secure. When evolutionary dynamics converge, their limits are equilibria of the underlying game, and so exclude strictly dominated strategies. But guarantees of convergence are available only for a few classes of games. When dynamics fail to converge, the payoffs of the available strategies remain in flux. If agents are not exact optimizers, but instead choose among strategies whose current payoffs are reasonably high, dominated strategies may be played by significant numbers of agents in perpetuity.

Appendix

A. Continuation of Attractors

Let X be a compact metric space and let ϕ be a semi-flow on X ; thus, $\phi : [0, \infty) \times X \rightarrow X$ is a continuous map satisfying $\phi^0(x) = x$ and $\phi^t(\phi^s(x)) = \phi^{t+s}(x)$ for all $s, t \geq 0$ and $x \in X$. A set $A \subset X$ is an *attractor* of ϕ if there is a neighborhood U of A such that $\omega(U) = A$ (see Conley (1978)). Here the ω -*limit set* of U is defined as $\omega(U) = \bigcap_{t>0} \text{cl}(\phi^{[t, \infty)}(U))$, where for $T \subset \mathbf{R}$ we let $\phi^T(U) = \bigcup_{t \in T} \phi^t(U)$. An attractor is compact and *invariant* ($\phi^t(A) = A$ for all t). Observe that an attractor can strictly contain another attractor.

The *basin* of the attractor is defined as $B(A) = \{x : \omega(x) \subseteq A\}$. For each open set U with $A \subset U \subset \bar{U} \subset B(A)$ we have $\omega(\bar{U}) = A$, see Section II.5.1.A of Conley (1978). Furthermore,

²⁰To see why, note that under the best response dynamic for bad RPS, the maximum weight on Scissors in the limit cycle is $\frac{4}{7}$ (see Gaunersdorfer and Hofbauer (1995)). If we move from the best response dynamic to a low-noise logit dynamic, and introduce a slightly dominated strategy Twin, a total weight of approximately $\frac{4}{7}$ is split nearly evenly between Scissors and Twin.

if $\phi^t(\bar{U}) \subset U$ holds for some $t > 0$ and for some open set U (which is then called a *trapping region*), then $\omega(U)$ is an attractor, see Section II.5.1.C of Conley (1978).

For a flow $(\phi^t, t \in \mathbf{R})$, the complement of the basin $B(A)$ of the attractor A is called the *dual repeller* of A . For all $x \in B(A) - A$, $\phi^t(x)$ approaches this dual repeller as $t \rightarrow -\infty$.

Consider now a one-parameter family of differential equations $\dot{x} = V_\varepsilon(x)$ in \mathbf{R}^n (with unique solutions $x(t) = \Phi_\varepsilon^t(x(0))$) such that $(\varepsilon, x) \mapsto V_\varepsilon(x)$ is continuous. Then $(\varepsilon, t, x) \mapsto \Phi_\varepsilon^t(x)$ is continuous as well. Suppose that $X \subset \mathbf{R}^n$ is compact and forward invariant under the semi-flows Φ_ε . For $\varepsilon = 0$ we omit the subscript in Φ^t .

The following continuation theorem for attractors is part of the folklore of dynamical systems; compare, e.g., Proposition 8.1 of Smale (1967).

Theorem A.1. *Let A be an attractor for Φ with basin $B(A)$. Then for each small enough $\varepsilon > 0$ there exists an attractor A_ε of Φ_ε with basin $B(A_\varepsilon)$, such that the map $\varepsilon \mapsto A_\varepsilon$ is upper hemicontinuous and the map $\varepsilon \mapsto B(A_\varepsilon)$ is lower hemicontinuous.*

Upper hemicontinuity cannot be replaced by continuity in this result. Consider the family of differential equations $\dot{x} = (\varepsilon + x^2)(1 - x)$ on the real line. The semi-flow Φ corresponding to $\varepsilon = 0$ admits $A = [0, 1]$ as an attractor, but when $\varepsilon > 0$ the unique attractor of Φ_ε is $A_\varepsilon = \{1\}$. This example shows that perturbations can cause attractors to implode; the theorem shows that perturbations cannot cause attractors to explode.

Theorem A.1 is a direct consequence of the following lemma, which is sufficient to prove the results in Sections 3 and 4.

Lemma A.2. *Let A be an attractor for Φ with basin $B(A)$, and let U_1 and U_2 be open sets satisfying $A \subset U_1 \subseteq U_2 \subseteq \bar{U}_2 \subseteq B(A)$. Then for each small enough $\varepsilon > 0$ there exists an attractor A_ε of Φ_ε with basin $B(A_\varepsilon)$, such that $A_\varepsilon \subset U_1$ and $U_2 \subset B(A_\varepsilon)$.*

In this lemma, one can always set $U_1 = \{x : \text{dist}(x, A) < \delta\}$ and $U_2 = \{x \in B(A) : \text{dist}(x, X - B(A)) > \delta\}$ for some small enough $\delta > 0$.

Proof of Lemma A.2. Since A is an attractor and $\omega(\bar{U}_2) = A$, there is a $T > 0$ such that $\Phi^t(\bar{U}_2) \subset U_1$ for $t \geq T$. By the continuous dependence of the flow on the parameter ε and the compactness of $\Phi^T(\bar{U}_2)$, we have that $\Phi_\varepsilon^T(\bar{U}_2) \subset U_1 \subseteq U_2$ for all small enough ε . Thus, U_2 is a trapping region for the semi-flow Φ_ε , and $A_\varepsilon \equiv \omega(U_2)$ is an attractor for Φ_ε . Moreover, $A_\varepsilon \subset U_1$ (since $A_\varepsilon = \Phi_\varepsilon^T(A_\varepsilon) \subseteq \Phi_\varepsilon^T(\bar{U}_2) \subset U_1$) and $U_2 \subset B(A_\varepsilon)$. ■

B. Proofs Omitted from Section 3

B.1 The Proof of Lemma 3.2

Lemma 3.2 states that the set of Nash equilibria $NE = \{x^* \in X : x^* = (\frac{1}{3}, \frac{1}{3}, \alpha, \frac{1}{3} - \alpha)\}$ in the “bad RPST” game $F(x) = \tilde{A}x$ is a repellor under the dynamics defined in Theorem 3.1. These dynamics are generated by revision protocols of the forms

$$(2) \quad \rho_{ij} = \phi(F_j - \bar{F}) \text{ and}$$

$$(3) \quad \rho_{ij} = \phi(F_j - F_i),$$

where $\phi : \mathbf{R} \rightarrow \mathbf{R}_+$ is a Lipschitz continuous function satisfying

$$(4) \quad \text{sgn}(\phi(u)) = \text{sgn}([u]_+) \text{ and } \left. \frac{d}{du^+} \phi(u) \right|_{u=0} > 0.$$

Equation (4) implies that for any $\bar{d} > 0$, there are positive constants k_1 and k_2 such that

$$(10) \quad k_1 d \leq \phi(d) \leq k_2 d \text{ for all } d \in [0, \bar{d}].$$

Choosing $\bar{d} = 2 \max_{i,x} |F_i(x)|$ ensures that bound (10) holds for all values of d relevant to our analysis. If we define

$$(11) \quad \psi(d) = \int_0^d \phi(u) du,$$

then integrating (10) shows that

$$(12) \quad \psi(d) \leq \frac{k_2}{2} d^2 \text{ for all } d \in [0, \bar{d}].$$

We now split the analysis into two parts, according to whether the dynamic is based on a protocol of form (2) or (3).

Lemma B.1. *Let $F(x) = \tilde{A}x$ be the “bad RPS with a twin” game from equation (5), and let V be the excess payoff dynamic generated by protocol (2), with ϕ satisfying condition (4). Then the set NE is a repellor under V^F .*

Proof. Define the excess payoff to strategy i by $\hat{F}_i(x) = F_i(x) - \bar{F}(x)$. Then using equation (1), we can express the dynamic V^F as

$$(13) \quad \dot{x}_i = \phi(\hat{F}_i(x)) - x_i \sum_{j \in S} \phi(\hat{F}_j(x)).$$

The rest points of (13) are the Nash equilibria of F ; moreover, if we let

$$\Phi(x) = \sum_{j \in S} \phi(\hat{F}_j(x)),$$

then $\Phi(x) \geq 0$, with equality if and only if x is a Nash equilibrium of F (see Sandholm (2005)).

Consider the Lyapunov function

$$U(x) = \sum_{i \in S} \psi(\hat{F}_i(x)),$$

where ψ is defined in equation (11). Hofbauer (2000) and Hofbauer and Sandholm (2009) show that $U(x) \geq 0$, with equality holding if and only if x is a Nash equilibrium of F . The proof of this theorem shows that the time derivative of U under the dynamic (13) can be expressed as

$$(14) \quad \dot{U}(x) = \dot{x}' \tilde{A} \dot{x} - \Phi(x) F(x)' \dot{x}.$$

To prove our lemma, we need to show that $\dot{U}(x) > 0$ whenever $x \notin NE$ and $\text{dist}(x, NE)$ is sufficiently small.

Let $TX = \{z \in \mathbf{R}^4 : z' \mathbf{1} = 0\}$, the tangent space of the simplex X , so that $\dot{x} \in TX$, and suppose that $z \in TX$. Then letting $(\zeta_1, \zeta_2, \zeta_3) = (z_1, z_2, z_3 + z_4)$, we have that

$$(15) \quad \begin{aligned} z' \tilde{A} z &= (a - b) (z_1 z_2 + z_2 (z_3 + z_4) + (z_3 + z_4) z_1) \\ &= (a - b) (\zeta_1 \zeta_2 + \zeta_2 \zeta_3 + \zeta_3 \zeta_1) \\ &= \frac{b - a}{2} \left(\left(\sum_{i=1}^3 \zeta_i \right)^2 - 2 \sum_{1 \leq i < j \leq 3} \zeta_i \zeta_j \right) \\ &= \frac{b - a}{2} \sum_{i=1}^3 \zeta_i^2. \\ &= \frac{b - a}{2} \left((z_1)^2 + (z_2)^2 + (z_3 + z_4)^2 \right). \end{aligned}$$

Now if $x \notin NE$, we can write (13) as

$$(16) \quad \dot{x} = \Phi(x) (\sigma(x) - x)$$

where $\sigma(x) \in X$ is given by $\sigma_i(x) = \phi(\hat{F}_i(x))/\Phi(x)$. Since $x \notin NE$, some strategy i has a below

average payoff ($F_i(x) < \bar{F}(x)$), implying that $\sigma_i(x) = 0$, and hence that $\sigma(x) \in \text{bd}(X)$. In fact, since strategies 3 and 4 always earn the same payoff, we have that $\sigma_3(x) = 0$ if and only if $\sigma_4(x) = 0$.

If we now write $y = (x_1, x_2, x_3 + x_4)$ and $\tau(x) = (\sigma_1(x), \sigma_2(x), \sigma_3(x) + \sigma_4(x))$, equation (16) becomes

$$\dot{y} = \Phi(x)(\tau(x) - y).$$

The arguments in the previous paragraph show that $\tau(x)$ is on the boundary of the simplex in \mathbf{R}^3 . Therefore, if we fix a small $\varepsilon > 0$ and assume that $\text{dist}(x, NE) < \varepsilon$, then $|y - (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})| < \varepsilon$, giving us a uniform bound on the distance between $\tau(x)$ and y , and hence a uniform lower bound on $|\dot{y}|$:

$$|\dot{y}| \geq c \Phi(x)$$

for some $c > 0$. By squaring and rewriting in terms of \dot{x} , we obtain

$$(17) \quad \dot{x}_1^2 + \dot{x}_2^2 + (\dot{x}_3 + \dot{x}_4)^2 \geq c^2 \Phi(x)^2.$$

Thus, combining equations (15) and (17) shows that if $\text{dist}(x, NE) < \varepsilon$, then

$$(18) \quad \dot{x}' \tilde{A} \dot{x} \geq \frac{b-a}{2} c^2 \Phi(x)^2.$$

To bound the second term of equation (14), use equation (10) to show that

$$\begin{aligned} (19) \quad \Phi(x) F(x)' \dot{x} &= \Phi(x) (\hat{F}(x) - \bar{F}(x) \mathbf{1})' \dot{x} \\ &= \Phi(x) \hat{F}(x)' \dot{x} && \text{since } \mathbf{1}' \dot{x} = 0 \\ &= \Phi(x) \sum_{i \in S} \hat{F}_i(x) (\phi(\hat{F}_i(x)) - x_i \Phi(x)) \\ &= \Phi(x) \sum_{i \in S} \hat{F}_i(x) \phi(\hat{F}_i(x)) && \text{since } \hat{F}(x)' x = 0 \\ &\geq \Phi(x) k_1 \sum_{i \in S} \hat{F}_i(x)^2 \\ &\geq \Phi(x) \frac{k_1}{n} \left(\sum_{i \in S} \hat{F}_i(x) \right)^2 \\ &\geq \Phi(x) \frac{k_1}{nk_2^2} \left(\sum_{i \in S} \phi(\hat{F}_i(x)) \right)^2 \end{aligned}$$

$$= \frac{k_1}{nk_2^2} \Phi(x)^3.$$

Combining inequalities (18) and (19) with equation (14), we find that for x close enough to NE ,

$$\dot{U}(x) \geq \frac{b-a}{2} c^2 \Phi(x)^2 - \frac{k_1}{nk_2^2} \Phi(x)^3.$$

Since $\Phi(x) \geq 0$, with equality only when $x \in NE$, we conclude that $\dot{U}(x) > 0$ whenever $x \notin NE$ is close enough to NE , and therefore that NE is a repellor under (13). ■

Lemma B.2. *Let $F(x) = \tilde{A}x$ be the “bad RPS with a twin” game from equation (5), and let V be the pairwise comparison dynamic generated by protocol (3), with ϕ satisfying condition (4). Then the set NE is a repellor under V^F .*

Proof. Using equation (1), we express the dynamic V^F as

$$(20) \quad \dot{x}_i = \sum_{j \in S} x_j \phi(F_i(x) - F_j(x)) - x_i \sum_{j \in S} \phi(F_j(x) - F_i(x)).$$

Sandholm (2008b) shows that the rest points of (20) are the Nash equilibria of F .

Our analysis relies on the following Lyapunov function:

$$\Psi(x) = \sum_{i \in S} \sum_{j \in S} x_i \psi(F_j(x) - F_i(x)),$$

where ψ is defined in equation (11). Hofbauer and Sandholm (2009) (also see Smith (1984)) show that $\Psi(x) \geq 0$, with equality holding if and only if x is a Nash equilibrium of F . The proof of that theorem shows that the time derivative of Ψ under the dynamic (20) can be expressed as

$$(21) \quad \begin{aligned} \dot{\Psi}(x) &= \dot{x}' \tilde{A} \dot{x} + \sum_{i \in S} \sum_{j \in S} \left(x_j \phi(F_i(x) - F_j(x)) \sum_{k \in S} \left(\psi(F_k(x) - F_i(x)) - \psi(F_k(x) - F_j(x)) \right) \right) \\ &\equiv T_1(x) + T_2(x). \end{aligned}$$

Equation (15) tells us that $T_1(x) \geq 0$, with equality when $x \in NE$ (i.e., when $\dot{x} = 0$). Hofbauer and Sandholm (2009) show that $T_2(x) \leq 0$, with equality only when $x \in NE$. To prove the lemma, we must show that $T_1(x) + T_2(x) > 0$ whenever $x \notin NE$ and $\text{dist}(x, NE)$ is sufficiently small.

To begin, observe that since F is linear, we have that

$$(22) \quad [F_j(x) - F_i(x)]_+ \leq c_1 \operatorname{dist}(x, NE).$$

for some $c_1 > 0$. Equations (10), (12), and (22) immediately yield a cubic bound on T_2 :

$$(23) \quad |T_2(x)| \leq c_2 \operatorname{dist}(x, NE)^3$$

for some $c_2 > 0$.

To obtain a lower bound on T_1 , first note that the linearity of F implies that

$$(24) \quad \max_{i \in S} F_i(x) - \min_{j \in S} F_j(x) \geq c_3 \operatorname{dist}(x, NE)$$

for some $c_3 > 0$. If $F_1(x) \geq F_2(x) \geq F_3(x) = F_4(x)$, then equations (10) and (24) imply that

$$\dot{x}_1 = \sum_{j=2}^4 x_j \phi(F_1(x) - F_j(x)) \geq (x_3 + x_4) \phi(F_1(x) - F_3(x)) \geq (x_3 + x_4) c_3 k_1 \operatorname{dist}(x, NE).$$

Similarly, if $F_1(x) \leq F_2(x) \leq F_3(x) = F_4(x)$, then

$$|\dot{x}_1| = x_1 \sum_{j=2}^4 \phi(F_j(x) - F_1(x)) \geq x_1 \phi(F_3(x) - F_1(x)) \geq x_1 c_3 k_1 \operatorname{dist}(x, NE).$$

Obtaining bounds on $|\dot{x}_1|$ and on $|\dot{x}_2|$ for the remaining four cases in like fashion, we find that for some $c_4 > 0$ and some $\varepsilon > 0$, for any x with $\operatorname{dist}(x, NE) \leq \varepsilon$ (and hence $|x_1 - \frac{1}{3}| \leq \varepsilon$, $|x_2 - \frac{1}{3}| \leq \varepsilon$, and $|(x_3 + x_4) - \frac{1}{3}| \leq \varepsilon$), we have that

$$(25) \quad \dot{x}_1^2 + \dot{x}_2^2 + (\dot{x}_3 + \dot{x}_4)^2 \geq \max(\dot{x}_1^2, \dot{x}_2^2) \geq c_4 \operatorname{dist}(x, NE)^2.$$

Thus, equations (15) and (25) together imply that for such x ,

$$(26) \quad T_1 = \dot{x}' \tilde{A} \dot{x} \geq \frac{b-a}{2} c_4 \operatorname{dist}(x, NE)^2.$$

Combining equations (21), (23), and (26), we find that

$$\dot{\Psi}(x) \geq \frac{b-a}{2} c_4 \operatorname{dist}(x, NE)^2 - c_2 \operatorname{dist}(x, NE)^3$$

whenever $\operatorname{dist}(x, NE) \leq \varepsilon$. We therefore conclude that $\dot{\Psi}(x) > 0$ whenever x is sufficiently close to but not in NE , and so that NE is a repeller under (20). ■

B.2 The Proof of Lemma 3.3

Let $\{x_t\}_{t \geq 0}$ be the solution to V^F from initial condition $\xi \in X - NE$, and suppose without loss of generality that $\xi_3 \geq \xi_4$. Then equation (6) implies that $(x_3 - x_4)_t$ is positive and is nonincreasing in t , while Lemma 3.2 reveals that there is a neighborhood U of NE and a time $T \geq 0$ such that $\{x_t\} \in X - U$ for all $t \geq T$.

We split the remaining analysis into two cases. Suppose first that the revision protocol ρ is of form (2), so that V is an excess payoff dynamic. In this case, equation (6) becomes

$$\dot{x}_3 - \dot{x}_4 = -(x_3 - x_4) \sum_{j \in S} \phi(F_j(x) - \bar{F}(x)) \equiv -(x_3 - x_4) \Phi(x).$$

Sandholm (2005) shows that $\Phi(x) \geq 0$, with equality if and only if x is a Nash equilibrium of F . Since U is an open neighborhood of NE , it follows that $m = \min_{x \in X - U} \Phi(x) > 0$. Thus $\frac{d}{dt}(x_3 - x_4)_t \leq -m(x_3 - x_4)_t$ for all $t \geq T$, from which we conclude that $\lim_{t \rightarrow \infty} (x_3 - x_4)_t = 0$.

Now suppose that ρ is of form (3), so that V is a pairwise comparison dynamic. In this case, equation (6) becomes

$$(27) \quad \dot{x}_3 - \dot{x}_4 = -(x_3 - x_4) \sum_{j \in S} \phi(F_j(x) - F_3(x)).$$

Since $(x_3 - x_4)_t$ is positive and is nonincreasing in t , it converges to some nonnegative limit l . Thus, $\omega(\xi)$, the ω -limit set of trajectory $\{x_t\}$, is contained in the plane $P_l = \{x \in X : x_3 - x_4 = l\}$. Also, $\omega(\xi)$ is compact and invariant by definition (see, e.g., Robinson (1995)), and $\omega(\xi)$ is disjoint from NE by Lemma 3.2.

Now assume that l is positive. Then equation (27), combined with the fact that $\omega(\xi)$ is invariant and a subset of P_l , implies that $\sum_{j \in S} \phi(F_j(x) - F_3(x)) = 0$ for all $x \in \omega(\xi)$. It then follows from condition (4) (namely, from the fact that $\text{sgn}(\phi(u)) = \text{sgn}([u]_+)$) that strategy 3 is optimal throughout $\omega(\xi)$; the identical twin strategy 4 must be so as well. But it is easy to check that under any pairwise comparison dynamic V^F for the “bad RPS with a twin” game F , optimal strategies have positive growth rates at any non-Nash state. Since $\omega(\xi)$ is disjoint from NE , we conclude that $V_3^F(x) + V_4^F(x) > 0$ whenever $x \in \omega(\xi)$.

At the same time, since $\omega(\xi)$ is compact, $\chi = \arg\max_{x \in \omega(\xi)} x_3 + x_4$ exists, and since $\omega(\xi)$ is also invariant, it must be that $V_3^F(\chi) + V_4^F(\chi) \leq 0$. This is a contradiction. We therefore conclude that $l = 0$, and so that $\{x_t\}$ converges to the plane $P_0 = P$. ■

C. Proofs Details Omitted from Section 4.2.4

Our analysis relies on the notion of attractor-repellor pairs introduced by Conley (1978); see Robinson (1995) for a textbook treatment. Define the flow from the set $U \subseteq X$ under the dynamic V^F by

$$\phi_t(U) = \left\{ \xi \in X : \text{there is a solution } \{x_s\} \text{ to } \dot{x} = V^F(x) \text{ with } x_0 \in U \text{ and } x_t = \xi. \right\}$$

In words, $\phi_t(U)$ contains the time t positions of solutions to V^F whose initial conditions are in U .

Recall that solutions to V^F starting in $I - NE$ ascend the function $f^{\tilde{C}}$ until leaving the set I . It follows that the set NE is a *repellor* under V^F : all backward-time solutions to V^F that begin near NE converge to NE . More precisely, there is a neighborhood U of NE such that

$$\bigcap_{t < 0} \phi_t(\text{cl}(U)) = NE.$$

The *dual attractor* \mathcal{A} of the repellor NE is the forward-time limit of the flow of V^F starting from the complement of $\text{cl}(U)$:

$$\mathcal{A} = \bigcap_{t > 0} \phi_t(X - \text{cl}(U)).$$

Lemma 4.3 tells us that $\mathcal{A} \subset D$. Three other key properties of the attractor \mathcal{A} are noted next.

Lemma C.1. *The attractor \mathcal{A} is nonempty, compact, and forward and backward invariant under V^F .*

These properties of attractors are well known—for proofs, see Conley (1978) or Robinson (1995).

We now show that \mathcal{A} is contained in the interior of the simplex. To do so, we let $Z = \{x \in X : x_4 = 0\}$ be the face of X on which the twin strategy is unused. We prove

Lemma C.2. *The attractor \mathcal{A} and the face Z are disjoint.*

Proof. Recall that $V^F(x) = g(F(x), x)$, where the growth rate function g is Lipschitz continuous (by condition (C)) and satisfies $g_i(\pi, x) \geq 0$ whenever $x_i = 0$. It follows that solutions to V^F that start in $X - Z$ cannot approach Z more than exponentially quickly, and

in particular cannot reach Z in finite time. Equivalently, backward solutions to V^F starting from states in Z cannot enter $\text{int}(X)$.

Now suppose by way of contradiction that there exists a state ξ in $\mathcal{A} \cap Z$. Then by Lemma C.1 and the previous paragraph, the entire backward orbit from ξ is also contained in $\mathcal{A} \cap Z$, and hence in $D \cap Z$. Since the latter set contains no rest points (by condition (PC)), the Poincaré-Bendixson Theorem implies that the backward orbit from ξ converges to a closed orbit γ in $D \cap Z$ that circumnavigates $I \cap Z$.

By construction, the annulus $D \cap Z$ can be split into three regions: one in which strategy 1 is the best response, one in which strategy 2 is the best response, and one in which strategy 3 (and hence strategy 4) is a best response. Each of these regions is bounded by a simple closed curve that intersects the inner and outer boundaries of the annulus. Therefore, the closed orbit γ , on which strategy 4 is unused, passes through the region in which strategy 4 is optimal. This contradicts innovation (IN). ■

References

- Akin, E. (1980). Domination or equilibrium. *Mathematical Biosciences*, 50:239–250.
- Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge.
- Beckmann, M., McGuire, C. B., and Winsten, C. B. (1956). *Studies in the Economics of Transportation*. Yale University Press, New Haven.
- Benaïm, M., Hofbauer, J., and Sorin, S. (2005). Stochastic approximation and differential inclusions. *SIAM Journal on Control and Optimization*, 44:328–348.
- Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.
- Berger, U. and Hofbauer, J. (2006). Irrational behavior in the Brown-von Neumann-Nash dynamics. *Games and Economic Behavior*, 56:1–6.
- Binmore, K. and Samuelson, L. (1999). Evolutionary drift and equilibrium selection. *Review of Economic Studies*, 66:363–393.
- 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.
- Burdett, K. and Judd, K. L. (1983). Equilibrium price dispersion. *Econometrica*, 51:955–969.

- Conley, C. (1978). *Isolated Invariant Sets and the Morse Index*. American Mathematical Society, Providence, RI.
- Eckert, A. (2003). Retail price cycles and presence of small firms. *International Journal of Industrial Organization*, 21:151–170.
- Friedman, D. (1991). Evolutionary games in economics. *Econometrica*, 59:637–666.
- Fudenberg, D. and Levine, D. K. (1998). *Theory of Learning in Games*. MIT Press, Cambridge.
- Gaunersdorfer, A. and Hofbauer, J. (1995). Fictitious play, Shapley polygons, and the replicator equation. *Games and Economic Behavior*, 11:279–303.
- Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica*, 59:859–867.
- Hart, S. and Mas-Colell, A. (2003). Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93:1830–1836.
- Hofbauer, J. (2000). From Nash and Brown to Maynard Smith: Equilibria, dynamics, and ESS. *Selection*, 1:81–88.
- Hofbauer, J. and Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, 70:2265–2294.
- Hofbauer, J. and Sandholm, W. H. (2009). Stable games and their dynamics. *Journal of Economic Theory*, 144:1665–1693.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge.
- Hofbauer, J. and Swinkels, J. M. (1996). A universal Shapley example. Unpublished manuscript, University of Vienna and Northwestern University.
- Hofbauer, J. and Weibull, J. W. (1996). Evolutionary selection against dominated strategies. *Journal of Economic Theory*, 71:558–573.
- Hopkins, E. and Seymour, R. M. (2002). The stability of price dispersion under seller and consumer learning. *International Economic Review*, 43:1157–1190.
- Jordan, J. S. (1993). Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior*, 5:368–386.
- Lahkar, R. (2007). The dynamic instability of dispersed price equilibria. Unpublished manuscript, University College London.
- Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.
- Nachbar, J. H. (1990). 'Evolutionary' selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19:59–89.

- 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.
- Noel, M. D. (2007). Edgeworth price cycles: Evidence from the Toronto retail gasoline market. *Journal of Industrial Economics*, 55:69–92.
- Robinson, C. (1995). *Dynamical Systems: Stability, Symbolic Dynamics, and Chaos*. CRC Press, Boca Raton, FL.
- Rosenthal, R. W. (1973). A class of games possessing pure strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67.
- Samuelson, L. and Zhang, J. (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391.
- 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. (2008a). Evolutionary game theory. In Meyers, R. A., editor, *Encyclopedia of Complexity and System Science*. Forthcoming, Springer, Heidelberg.
- Sandholm, W. H. (2008b). Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. Unpublished manuscript, University of Wisconsin.
- Sandholm, W. H. (2009). Large population potential games. *Journal of Economic Theory*, 144:1710–1725.
- Sandholm, W. H., Dokumacı, E., and Lahkar, R. (2008). The projection dynamic and the replicator dynamic. *Games and Economic Behavior*, 64:666–683.
- Shapley, L. S. (1964). Some topics in two person games. In Dresher, M., Shapley, L. S., and Tucker, A. W., editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–28. Princeton University Press, Princeton.
- Smale, S. (1967). Differentiable dynamical systems. *Bulletin of the American Mathematical Society*, 73:747–817.
- 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.

- Sparrow, C., van Strien, S., and Harris, C. (2008). Fictitious play in 3×3 games: The transition between periodic and chaotic behaviour. *Games and Economic Behavior*, 63:259–291.
- 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.
- Varian, H. (1980). A model of sales. *American Economic Review*, 70:651–659.
- Weibull, J. W. (1996). The mass action interpretation. Excerpt from 'The work of John Nash in game theory: Nobel Seminar, December 8, 1994'. *Journal of Economic Theory*, 69:165–171.