

Population Games and Deterministic Evolutionary Dynamics*

William H. Sandholm[†]

January 9, 2014

Abstract

Population games describe strategic interactions among large numbers of small, anonymous agents. Behavior in these games is typically modeled dynamically, with agents occasionally receiving opportunities to switch strategies, basing their choices on simple myopic rules called *revision protocols*. Over finite time spans the evolution of aggregate behavior is well approximated by the solution of a differential equation. From a different point of view, every revision protocol defines a map—a *deterministic evolutionary dynamic*—that assigns each population game a differential equation describing the evolution of aggregate behavior in that game.

In this chapter, we provide an overview of the theory of population games and deterministic evolutionary dynamics. We introduce population games through a series of examples and illustrate their basic geometric properties. We formally derive deterministic evolutionary dynamics from revision protocols, introduce the main families of dynamics—imitative/biological, best response, comparison to average payoffs, and pairwise comparison—and discuss their basic properties. Combining these streams, we consider classes of population games in which members of these families of dynamics converge to equilibrium; these classes include potential games, contractive games, games solvable by iterative solution concepts, and supermodular games. We relate these classes to the classical notion of an evolutionarily stable state (ESS) and to recent work on deterministic equilibrium selection. We present a variety of examples of cycling and chaos under evolutionary dynamics, as well as a general result on survival of strictly dominated strategies. Finally, we provide connections to other approaches to game dynamics, and indicate applications of evolutionary game dynamics to economics and social science.

Keywords: Evolutionary game theory, learning in games, population games, revision protocols, deterministic evolutionary dynamics, global convergence, local stability, nonconvergence, stochastic processes, dynamical systems. *JEL* classification codes: C72, C73.

*I thank Man Wah Cheung for meticulous proofreading and an anonymous reviewer for helpful comments. Financial support from NSF Grant SES-1155135 is gratefully acknowledged.

[†]Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706. e-mail: whs@ssc.wisc.edu; website: www.ssc.wisc.edu/~whs.

Contents

1	Introduction	4
2	Population Games	6
2.1	Definitions	6
2.2	Examples	7
2.3	The Geometry of Population Games	9
3	Revision Protocols and Mean Dynamics	12
3.1	Revision Protocols	13
3.2	Information Requirements for Revision Protocols	14
3.3	The Stochastic Evolutionary Process and Mean Dynamics	15
3.4	Finite Horizon Deterministic Approximation	17
4	Deterministic Evolutionary Dynamics	18
4.1	Definition	18
4.2	Incentives and Aggregate Behavior	18
5	Families of Evolutionary Dynamics	20
5.1	Imitative Dynamics	23
5.2	The Best Response Dynamic and Related Dynamics	28
5.3	Excess Payoff and Pairwise Comparison Dynamics	33
6	Potential Games	36
6.1	Population Games and Full Population Games	37
6.2	Definition, Characterization, and Interpretation	37
6.3	Examples	38
6.4	Characterization of Equilibrium	40
6.5	Global Convergence and Local Stability	42
6.6	Local Stability of Strict Equilibria	43
7	Evolutionarily Stable States and Contractive Games	44
7.1	Evolutionarily Stable States	44
7.2	Contractive Games	46
7.3	Examples	48
7.4	Equilibrium in Contractive Games	48
7.5	Global Convergence and Local Stability	50
8	Iterative Solution Concepts, Supermodular Games, and Equilibrium Selection	54
8.1	Iterated Strict Dominance and Never-a-Best-Response	54
8.2	Supermodular Games and Perturbed Best Response Dynamics	56
8.3	Iterated p -Dominance and Equilibrium Selection	58

9	Nonconvergence of Evolutionary Dynamics	61
9.1	Examples	61
9.2	Survival of Strictly Dominated Strategies	67
10	Connections and Further Developments	70
10.1	Connections with Stochastic Stability Theory	70
10.2	Connections with Models of Heuristic Learning	71
10.3	Games with Continuous Strategy Sets	72
10.4	Extensive Form Games and Set-Valued Solution Concepts	74
10.5	Applications	75
	References	77

1. Introduction

Consider a population of commuters, each of whom must select a path through a highway network from his home to his workplace. Each commuter's payoff depends not only on the route he chooses, but also on the distribution of the route choices of other drivers, as these will determine the delays the commuter will face.

Population games are a powerful tool for modeling strategic environments like traffic networks, in which the number of agents is large, each agent is small, and agents are anonymous, with each agent's payoffs depending on his own strategy and the distribution of others' strategies. One typically imposes further restrictions on the agents' diversity: there are a finite number of populations, and agents in each population are identical, in that they choose from the same set of strategies and have identical payoff functions. Despite their simplicity, population games offer a powerful tool for applications in economics, computer science, biology, sociology, and other fields that study interactions among large numbers of participants.¹

The traditional approach to prediction in noncooperative games, equilibrium analysis, is based on strong assumptions about what players know. Such assumptions—that players fully understand the game they are playing, and that they are able to correctly anticipate how others will act—are overly demanding in many applications, particularly those in which the number of participants is large.

An alternative approach, one especially appropriate for recurring interactions with many agents, proceeds through a dynamic, disequilibrium analysis. One assumes that agents occasionally receive opportunities to switch strategies. A modeling device called a *revision protocol* specifies when and how they do so. The definition of the protocol reflects what information is available to agents when they make decisions, and how this information is used. Protocols can capture imitation, optimization, or any other criterion that agents can employ to respond to current strategic conditions.

Together, a population game, a population size, and a revision protocol generate a Markov process on the set of population states—that is, of distributions over pure strategies. While this Markov process can be studied directly, particularly powerful conclusions can be reached by evaluating this process over a fixed time horizon in the large population limit.² A suitable law of large numbers implies that as the population

¹Research on population games is distinguished from the literature on large noncooperative games (Schmeidler (1973), Khan and Sun (2002), Balder (2002), Carmona and Podczeck (2009)) by both the limited diversity assumption and the central role played by disequilibrium dynamics.

²The other main approach to studying this process focuses on its infinite horizon behavior, as either a noise parameter or the population size approaches its limit. This approach, known as *stochastic stability analysis*, is the subject of the chapter by Wallace and Young in this volume, and is discussed in Section 10.1

size grows large, the sample paths of the Markov process can be approximated arbitrarily well by a deterministic trajectory. This trajectory is obtained as a solution of an ordinary differential equation: the *mean dynamic* induced by the game and the protocol.

We can view this development in a different light by fixing the revision protocol, thereby obtaining a map from population games to differential equations. This map—and by the usual synecdoche, the output of this map—is known as a *deterministic evolutionary dynamic*. Deterministic evolutionary dynamics reflect the character of the protocols that generate them; for example, dynamics based on imitation are readily distinguished from those based on optimization. Nevertheless, in some classes of games, dynamics derived from a variety of choice principles exhibit qualitatively similar behavior.

Evolutionary dynamics are nonlinear differential equations, so explicit formulas for solution trajectories are almost never available. Instead, methods from the theory of dynamical systems are used to establish local stability, global convergence, and nonconvergence results. Software for computing and visualizing numerical solutions offers a powerful tool for understanding the behavior of evolutionary dynamics, particularly in cases of cyclical and chaotic behavior.³

Section 2 introduces population games, provides some basic examples, and offers geometric intuitions. Section 3 introduces revision protocols and mean dynamics, and describes the deterministic approximation theorem that allows one to move from one to the other. Section 4 formalizes the notion of deterministic evolutionary dynamics, and introduces criteria that relate their behavior to incentives in the underlying games. Section 5 introduces the main families of evolutionary dynamics that have been studied in the literature: imitative dynamics (and biological dynamics), the best response dynamic and its variants, and dynamics based on comparisons to average payoffs and pairwise comparisons.

With these basic ingredients in place, the next four sections present convergence and nonconvergence results for various combinations of games and dynamics. Section 6 considers the class of potential games, for which the most general convergence results are available. Section 7 introduces the notion of an evolutionarily stable state (ESS), which provides a general sufficient condition for local stability. It also studies the class of contractive games, which are motivated by a version of the ESS condition, and which admit global convergence results. Section 8 addresses iterative solution concepts, supermodular games, and deterministic equilibrium selection. Section 9 provides examples of cycling and chaos under evolutionary dynamics, and presents a general result on survival of

below. For complete treatments, see Young (1998) and Sandholm (2010c, Chapters 11 and 12).

³The phase diagrams in this chapter were created using the *Dynamo* software suite (Sandholm et al. (2012)). See Franchetti and Sandholm (2013) for an introduction.

strictly dominated strategies.

While the material above is the core of the theory of population games and deterministic evolutionary dynamics, many other topics in the theory have been explored to various degrees. Section 10 provides snapshots of work on these topics, connections to other approaches to game dynamics, and a brief summary of applications in economics and social science.

Other treatments of evolutionary game theory with a focus on deterministic dynamics include survey papers by Hofbauer and Sigmund (2003), Sandholm (2009a), Hofbauer (2011) and Cressman (2011), and books by Hofbauer and Sigmund (1988, 1998), Weibull (1995), Cressman (2003), and Sandholm (2010c). The last reference includes detailed treatments of the mathematics used in the analysis of evolutionary game dynamics, and provides full accounts of many of the topics introduced below.

2. Population Games

This section introduces games played by a single population, with all agents sharing the same strategy set and payoff function. Extending the model to allow multiple populations is simple, but complicates the notation; see Sandholm (2010c) for details.

2.1 Definitions

We consider a unit mass of agents, each of whom chooses a pure strategy from the set $S = \{1, \dots, n\}$. The aggregate behavior of these agents is described by a *population state*. This is an element of the simplex $X = \{x \in \mathbb{R}_+^n : \sum_{j \in S} x_j = 1\}$, with x_j representing the proportion of agents choosing pure strategy j . The standard basis vector $e_i \in \mathbb{R}^n$ represents the *pure population state* at which all agents choose strategy i .

We identify a *population game* with a continuous vector-valued payoff function $F: X \rightarrow \mathbb{R}^n$. The scalar $F_i(x)$ represents the payoff to strategy i when the population state is x .

A population state is a *Nash equilibrium* of F , denoted $x \in NE(F)$, if no agent can improve his payoff by unilaterally switching strategies. More explicitly, x^* is a Nash equilibrium if

$$(1) \quad x_i^* > 0 \text{ implies that } F_i(x^*) \geq F_j(x^*) \text{ for all } j \in S.$$

By representing Nash equilibria as fixed points of the game's best response correspondence (see Section 5.2), one can use the Kakutani fixed point theorem to prove that equilibrium exists.

Theorem 2.1. *Every population game admits at least one Nash equilibrium.*

As we argued in the introduction, the direct assumption of equilibrium play rests on strong knowledge assumptions, assumptions that are particularly suspect in games with large numbers of participants. But in large population settings, equilibrium predictions can sometimes be justified by consideration of disequilibrium dynamics. Sections 6, 7, and 8 will show that for certain interesting classes of population games, various dynamic adjustment processes lead to equilibrium play. It is not coincidental that in these classes of games, existence of equilibrium can be established by elementary methods—in particular, without recourse to fixed point theorems. In games for which the full power of fixed point theory is needed to establish existence of equilibrium, the direct assumption of equilibrium play is least convincing.

2.2 Examples

We now introduce some basic examples of population games that we revisit throughout the chapter.

Example 2.2. Matching in two-player symmetric normal form games. In a symmetric two-player normal form game, each of the two players chooses a (pure) strategy from the finite set $S = \{1, \dots, n\}$. The game's payoffs are described by the matrix $A \in \mathbb{R}^{n \times n}$. Entry 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.

Suppose that the unit mass of agents are randomly matched to play the symmetric normal form game A (or, alternatively, that each agent is matched once with each possible opponent). The expected payoff to strategy i at population state x is described by the linear function $F_i(x) = \sum_{j \in S} A_{ij}x_j$. The payoffs to all strategies can be expressed concisely as $F(x) = Ax$. ♦

Population games based on matching in normal form games are by far the most commonly studied in the literature. There are many reasons for this attention. Linearity of payoffs in the population state makes these games mathematically simple. Furthermore, some research on population games and evolutionary dynamics is motivated not by a direct interest in large games per se, but rather by the possibility of obtaining equilibrium selection results for (reduced) normal form games by embedding them in a larger context. But if one is genuinely interested in modeling behavior in large populations of strategically interacting agents, focusing on the linear case is unnecessarily restrictive. Indeed, there are some applications in which nonlinearity is essential.

Example 2.3. Congestion games. Consider the following model of highway congestion, due to Beckmann et al. (1956). A pair of towns, Home and Work, are connected by a network of *links*. To commute from Home to Work, an agent must choose a *path* $i \in S$ connecting the two towns. The payoff the agent obtains is the negation of the delay on the path he takes. The delay on the path is the sum of the delays on its links, and the delay on a link is a function of the number of agents who use that link.

To formalize this environment as a *congestion game*, let \mathcal{L} be the collection of links in the highway network. Each strategy $i \in S$ is a path from Home to Work, and so is identified with a set of links $\mathcal{L}_i \subseteq \mathcal{L}$. Each link ℓ is assigned a *cost function* $c_\ell: \mathbb{R}_+ \rightarrow \mathbb{R}$, whose argument is link ℓ 's *utilization level* u_ℓ :

$$u_\ell(x) = \sum_{i \in S(\ell)} x_i, \text{ where } S(\ell) = \{i \in S: \ell \in \mathcal{L}_i\}.$$

The payoff of choosing path i is the negation of the total delays on its links:

$$F_i(x) = - \sum_{\ell \in \mathcal{L}_i} c_\ell(u_\ell(x)).$$

Since highway congestion involves negative externalities, cost functions in models of highway congestion are increasing. They are typically convex as well: delays are essentially fixed until a link becomes congested, at which point they increase quickly. Congestion games can also be used to model positive externalities, like the choice between different technological standards; in this case, cost functions are decreasing in the utilization levels. ♦

Population games also provide a useful framework for macroeconomic applications.

Example 2.4. Search with positive externalities. Agents in a unit mass population choose levels of search effort from the set $S = \{1, \dots, n\}$. Stronger efforts increase the likelihood of finding trading partners, so that payoffs are increasing both in own search effort and in aggregate search effort. This search model is represented by a population game with payoffs $F_i(x) = m_i b(a(x)) - c_i$, where $a(x) = \sum_{k=1}^n kx_k$ represents aggregate search effort, the increasing function $b: \mathbb{R}_+ \rightarrow \mathbb{R}$ represents the benefits of search as a function of aggregate effort, the increasing function $m: S \rightarrow \mathbb{R}$ is the benefit multiplier, and the arbitrary function $c: S \rightarrow \mathbb{R}$ captures search costs. ♦

2.3 The Geometry of Population Games

An important advantage of the population game framework is the possibility of representing games' incentive structure geometrically, at least in low-dimensional cases. Intuition obtained in these cases often carries over to higher dimensions.

To present population games in pictures, we introduce the matrix $\Phi = I - \frac{1}{n}\mathbf{1}\mathbf{1}' \in \mathbb{R}^{n \times n}$, where $\mathbf{1} \in \mathbb{R}^n$ is the vector of ones. This matrix represents the orthogonal projection of \mathbb{R}^n onto the subspace $TX = \{z \in \mathbb{R}^n : \sum_{i \in S} z_i = 0\}$, which is the tangent space of the simplex X . For a payoff vector $\pi \in \mathbb{R}^n$, the projected payoff vector

$$\Phi\pi = \pi - \mathbf{1} \left(\frac{1}{n} \sum_{i \in S} \pi_i \right) \equiv \pi - \mathbf{1}\bar{\pi}$$

is obtained by subtracting the population average payoff $\bar{\pi}$ from each component. Thus applying the projection Φ to a payoff vector eliminates information about average payoffs, while preserving information about payoff differences, and hence about the incentives that revising agents face.


Example 2.5. Drawing two-strategy games. Figure 1 presents the payoff vectors and projected payoff vectors for the two-strategy coordination game F^{C2} and the Hawk-Dove game F^{HD} :

$$F^{C2}(x) = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ 2x_2 \end{pmatrix}; \quad F^{HD}(x) = \begin{pmatrix} -1 & 2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_H \\ x_D \end{pmatrix} = \begin{pmatrix} 2x_D - x_H \\ x_D \end{pmatrix}.$$


In each case, the first strategy is represented on the vertical axis, so as to agree with the payoff matrix, and the projected payoff vectors are those running parallel to the simplex. In the coordination game, the payoff vectors push outward, away from the mixed equilibrium $x^* = (\frac{2}{3}, \frac{1}{3})$ and toward the pure equilibria e_1 and e_2 . In Hawk-Dove, the payoff vectors push inward, toward the unique Nash equilibrium $x^* = (\frac{1}{2}, \frac{1}{2})$.

We will see that the basic payoff monotonicity condition for disequilibrium dynamics (see Section 4.2) requires the vector describing the motion of the population state to agree with the payoff vector, in the weak sense that the angle between the two is acute. In a two-strategy game, this condition and feasibility completely determine the direction in which evolution should proceed: the state should move in the direction indicated by the projected payoff vector. Thus, Figure 1 shows that evolution sends the population toward a pure equilibrium in the coordination game, and toward the mixed equilibrium in the Hawk-Dove game. ♦

In games with three strategies, it is no longer possible to draw the payoff vectors




(i) 12 Coordination




(ii) Hawk-Dove

Figure 1: Payoffs and projected payoffs in two-strategy games.



(i) 123 Coordination



(ii) standard Rock-Paper-Scissors

Figure 2: Projected payoffs in three-strategy games.

directly; only the projected payoff vectors may be drawn. Moreover, while these vectors continue to describe incentives, they are no longer enough to determine the direction in which the population state will move, which will depend on the details of the revision procedure.

Example 2.6. Drawing three-strategy games. Figure 2 illustrates projected payoffs in the three-strategy coordination game F^{C^3} and the standard Rock-Paper-Scissors game F^{RPS} .

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

In the coordination game, the payoff vectors again push outward toward the pure equilibria. We will see in Section 6.5 that for a large class of evolutionary dynamics, play will converge to one of these equilibria from most initial states.

In the standard Rock-Paper-Scissors game, the projected payoffs cycle around the simplex. This suggests that evolutionary dynamics in this game need not converge to Nash equilibrium. As we will see below, whether or not convergence occurs depends on the revision protocol agents employ. ♦

In Figures 1 and 2, the Nash equilibria of the games in question are drawn as dots. While these Nash equilibria can be found by checking definition (1), they can also be discovered geometrically.

For each population state x , we define the *tangent cone* of X at x to be the set of directions of motion from x that do not cause the state to leave the simplex X :

$$TX(x) = \{z \in \mathbb{R}^n : z = \alpha(y - x) \text{ for some } y \in X \text{ and some } \alpha \geq 0\}.$$

The *normal cone* of X at x is the polar of the tangent cone of X at x . In other words, it is the set of directions that form an obtuse or right angle with every vector in $TX(x)$:

$$NX(x) = (TX(x))^\circ = \{y \in \mathbb{R}^n : y'z \leq 0 \text{ for all } z \in TX(x)\}.$$

The possible forms of the tangent and normal cones in two- and three-strategy games are presented in Figure 3. In the latter case this is done after applying the projection Φ .

Observe that

$$\begin{aligned} x \in NE(F) &\Leftrightarrow [x_i > 0 \Rightarrow F_i(x) \geq F_j(x)] \text{ for all } i, j \in S \\ &\Leftrightarrow x'F(x) \geq y'F(x) \text{ for all } y \in X \end{aligned}$$

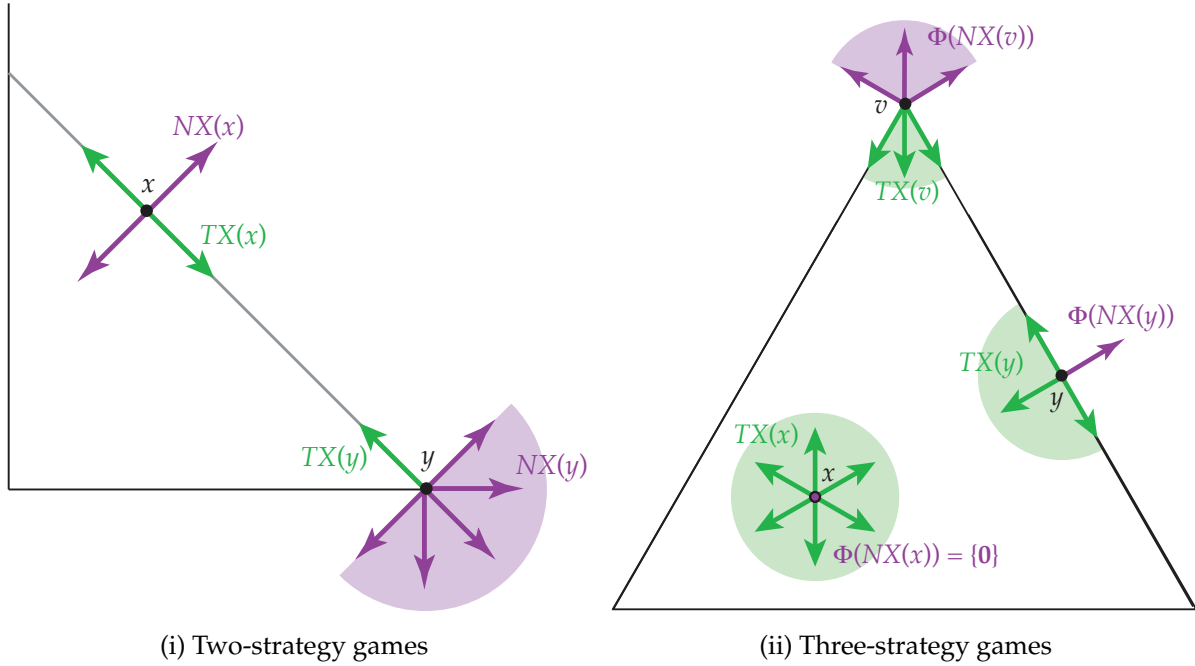


Figure 3: Tangent cones and normal cones.

$$(2) \quad \Leftrightarrow (y - x)'F(x) \leq 0 \text{ for all } y \in X.$$

We thus have the following geometric characterization of Nash equilibrium.

Theorem 2.7. *Let F be a population game. Then $x \in NE(F)$ if and only if $F(x) \in NX(x)$.*

According to the theorem, the Nash equilibria drawn in Figures 1 and 2 are precisely those states whose payoff vectors lie in the relevant normal cones, as drawn in Figure 3.

3. Revision Protocols and Mean Dynamics

Evolutionary dynamics for population games are designed to capture two basic assumptions. The first, *inertia*, means that agents only occasionally consider switching strategies. This assumption is natural when the environment being modeled is one of many in which the agents participate, so that the agents only pay limited attention to each. The second assumption, *myopia*, says that agents do not attempt to forecast future behavior, but instead base their decisions on the information they have about the current strategic environment. These two assumptions are mutually reinforcing: myopic behavior is most sensible when opponents' behavior adjusts slowly, so that strategies that perform

well now are likely to continue to do so.⁴

To specify the adjustment process most transparently, we consider how individual agents make decisions. Formally, this is accomplished by means of objects called revision protocols. Revision protocols describe when and how agents decide to switch strategies, and they implicitly specify what information agents use to make these decisions.

Together, a population game and a revision protocol generate a differential equation called the mean dynamic. This dynamic describes the evolution of aggregate behavior when the revision protocol is employed during recurrent play of the game.

This section introduces revision protocols and mean dynamics, and justifies the use of the latter to describe the evolution of the population state. This approach to defining deterministic evolutionary dynamics via microfoundations was first developed for imitative dynamics by Björnerstedt and Weibull (1996) and Weibull (1995), and later expanded by Benaïm and Weibull (2003) and Sandholm (2003, 2010b).

3.1 Revision Protocols

In the most general case, a *revision protocol* ρ is a map that assigns each population game F a function $\rho^F: X \rightarrow \mathbb{R}_+^{n \times n}$, which maps population states $x \in X$ to collections of *conditional switch rates* $\rho_{ij}^F(x)$. A population game F , a revision protocol ρ , and a finite population size N together define a stochastic evolutionary process—a Markov process—that runs on the discrete grid $\mathcal{X}^N = X \cap \frac{1}{N}\mathbb{Z}^n = \{x \in X : Nx \in \mathbb{Z}^n\}$.⁵

This process, which we define formally below, can be described as follows. Each agent in the society is equipped with a “stochastic alarm clock”. The times between rings of an agent’s clock are independent, each with a rate R exponential distribution, where $R = R(\rho^F) \geq \max_{x,i} \sum_{j \neq i} \rho_{ij}^F(x)$,⁶ and different agents’ clocks are independent of one another. The ringing of a clock signals the arrival of a revision opportunity for the clock’s owner. If an agent playing strategy $i \in S$ receives a revision opportunity, he switches to strategy $j \neq i$ with probability $\frac{1}{R}\rho_{ij}^F(x)$, and he continues to play strategy i with probability $1 - \frac{1}{R} \sum_{j \neq i} \rho_{ij}^F(x)$. This decision is made independently of the timing of the clocks’ rings. If a switch occurs, the population state changes accordingly, from the old state x to a new

⁴Other approaches to dynamics for population games build on different assumptions. The most notable example, the *perfect foresight dynamics* of Matsui and Matsuyama (1995), are obtained by retaining the assumption of inertia, but assuming that agents are forward looking. While evolutionary dynamics describe disequilibrium adjustment, perfect foresight dynamics represent adjustments that occur within a dynamic equilibrium.

⁵This process can also be formulated in discrete time, and can be adjusted to account for finite-population effects; see Benaïm and Weibull (2003, 2009) and Sandholm (2010c, Section 10.3). One can also incorporate finite-population effects into the definition of population games; see Sandholm (2010c, Section 11.4).

⁶As we will see, the mean dynamic is not affected by the value of the clock rate R .

state y that accounts for the agent's choice.

This interpretation of the stochastic evolutionary process has the advantage of working for any revision protocol. But once a protocol is fixed, one can use the structure provided by that protocol to provide a simpler interpretation of the process—see Example 3.2 below.

According to the description above, the diagonal elements $\rho_{ii}^F(x)$ of a revision protocol are irrelevant. An important exception occurs when the protocol has unit row sums, $\sum_{j \in S} \rho_{ij}^F(x) = 1$ for all $x \in X$ and $i \in S$, so that $\rho_i^F(x) = (\rho_{i1}^F(x), \dots, \rho_{in}^F(x))$ is a probability vector. In this case, we set the clock rate R at 1, and refer to $\rho_{ij}^F(x)$ as a *conditional switch probability*.

3.2 Information Requirements for Revision Protocols

The model above describes the behavior of agents who only occasionally consider switching strategies. To assess how well a particular revision protocol agrees with this general approach, we introduce some restrictions concerning the information that revision protocols may require.

Most of the protocols studied in the literature are of the form $\rho^F(x) = \rho(F(x), x)$, so that the current conditional switch rates only depend on the game by way of the current payoff. We call such protocols *reactive*. The remaining protocols, which we call *prospective*, are more demanding, in that they require agents to know enough about the payoff functions to engage in counterfactual reasoning. We will only consider prospective protocols in Section 8.3, where we will see that dynamics based on them can differ markedly from those based on reactive protocols.

Protocols can also be distinguished by the *amount and types of data* about the current strategic environment that they require. The least demanding protocols require knowledge of only one payoff, that of the current or the candidate strategy; slightly more demanding ones require both. Still more demanding ones require knowledge of the current payoffs to all strategies. Protocols that require further information—say, information about the average payoffs in the population—may be regarded as too demanding for typical applications.

A third distinction separates *continuous* and *discontinuous* protocols. Under continuous protocols, agents' choices do not change abruptly after small changes in the strategic environment. This property accords well with the evolutionary paradigm. Discontinuous protocols, which require exact information about the current strategic environment, are perhaps less natural under this paradigm, but are used to define important ideal cases.

3.3 The Stochastic Evolutionary Process and Mean Dynamics

Formally, the game F , the protocol ρ , and a finite population size N define a Markov process $\{X_t^N\}_{t \geq 0}$ on the finite state space \mathcal{X}^N . Since each of the N agents receives revision opportunities at rate R , revision opportunities arrive in the population as a whole at rate NR , which is thus the expected number of revision opportunities arriving during each unit of clock time. Letting τ_k denote the arrival time of the k th revision opportunity, we can describe the transition law of the process $\{X_t^N\}$ by

$$\mathbb{P}\left(X_{\tau_{k+1}}^N = y \mid X_{\tau_k}^N = x\right) = \begin{cases} \frac{x_i \rho_{ij}^F(x)}{R} & \text{if } y = x + \frac{1}{N}(e_j - e_i), j \neq i, \\ 1 - \sum_{i \in S} \sum_{j \neq i} \frac{x_i \rho_{ij}^F(x)}{R} & \text{if } y = x, \\ 0 & \text{otherwise.} \end{cases}$$

While one can analyze this Markov process directly, we instead consider its limiting behavior as the population size N becomes large. Notice that all of the randomness in the process $\{X_t^N\}$ is idiosyncratic: both the assignments of revision opportunities and the randomizations performed by the agents are independent of past events conditional on the current state. Therefore, taking the limit as N grows large enables us to approximate $\{X_t^N\}$ by a deterministic trajectory—namely, a solution of the so-called *mean dynamic* generated by ρ and F . We introduce the mean dynamic next, saving the formal approximation result for the next section.

To derive the mean dynamic, we consider the behavior of the process $\{X_t^N\}$ over the next dt time units, starting from state x . Since each of the N agents receives revision opportunities at rate R , the expected number of opportunities arriving during these dt time units is $NR dt$. Each opportunity is equally likely to go to each agent, so the expected number of these opportunities that are received by current strategy i players is $Nx_i R dt$; the expected number of these that lead to switches to strategy j is $Nx_i \rho_{ij}^F(x) dt$. Hence, the expected change in the proportion of agents using strategy i is

$$\left(\sum_{j \in S} x_j \rho_{ji}^F(x) - x_i \sum_{j \in S} \rho_{ij}^F(x) \right) dt.$$

The *mean dynamic* induced by population game F and revision protocol ρ is thus

$$(3) \quad \dot{x} = V^F(x), \quad \text{where } V_i^F(x) = \sum_{j \in S} x_j \rho_{ji}^F(x) - x_i \sum_{j \in S} \rho_{ij}^F(x).$$

Here $\dot{x} = \dot{x}_t = \frac{d}{dt}x_t$ denotes the time derivative of the solution trajectory $\{x_t\}_{t \geq 0}$.

The mean dynamic (3) is an ordinary differential equation defined on the simplex X . Evidently, $V_i^F(x) \geq 0$ whenever $x_i = 0$, so $V(x) \in TX(x)$ for all $x \in X$; that is, V^F never points outward from the boundary of X . So long as V^F is Lipschitz continuous,⁷ existence and uniqueness of solutions to (3) follow from standard results.

Theorem 3.1. *If V^F is Lipschitz continuous, then the mean dynamic (3) admits a unique forward solution from each $\xi \in X$: that is, a trajectory $\{x_t\}_{t \geq 0}$ satisfying $x_0 = \xi$ and $\frac{d}{dt}x_t = V^F(x_t)$ for all $t \geq 0$.*

Typically, Lipschitz continuity of V^F is ensured by assuming that the game F is Lipschitz continuous and that the protocol $\rho^F(x) = \rho(F(x), x)$ is reactive and Lipschitz continuous.

Example 3.2. Pairwise proportional imitation and the replicator dynamic. Helbing (1992) and Schlag (1998) introduce the following reactive protocol, called *pairwise proportional imitation*:

$$(4) \quad \rho_{ij}^F(x) = \rho_{ij}(F(x), x) = x_j[F_j(x) - F_i(x)]_+.$$

Under protocol (4), an agent who receives a revision opportunity chooses an opponent at random. This opponent is a strategy j player with probability x_j . The agent imitates the opponent only if the opponent's payoff is higher than his own, doing so with probability proportional to the payoff difference.

Substituting protocol (4) into formula (3), we find that the mean dynamic associated with this protocol is

$$\begin{aligned} \dot{x}_i &= \sum_{j \in S} x_j x_i [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} x_j [F_j(x) - F_i(x)]_+ \\ &= x_i \sum_{j \in S} x_j (F_i(x) - F_j(x)) \\ &= x_i \left(F_i(x) - \sum_{j \in S} x_j F_j(x) \right). \end{aligned}$$

This is the *replicator dynamic* of Taylor and Jonker (1978), the best known dynamic in evolutionary game theory. Under this dynamic, the percentage growth rate \dot{x}_i/x_i of each strategy currently in use is equal to the difference between that strategy's payoff and

⁷The function $V^F: X \rightarrow \mathbb{R}^n$ is *Lipschitz continuous* if there exists a constant $K > 0$ such that $|V^F(y) - V^F(x)| \leq K|y - x|$ for all $x, y \in X$.

the average payoff obtained in the population; unused strategies always remain so. We discuss this dynamic further in Example 5.1 below. ♦

3.4 Finite Horizon Deterministic Approximation

The basic link between the Markov processes $\{X_t^N\}$ and the mean dynamic (3) is provided by the following theorem.

Theorem 3.3. *Suppose that the mean dynamic V^F is Lipschitz continuous. Let the initial conditions $X_0^N = x_0^N$ converge to state $x_0 \in X$, and let $\{x_t\}_{t \geq 0}$ be the solution to the mean dynamic (3) starting from x_0 . Then for all $T < \infty$ and $\varepsilon > 0$,*

$$(5) \quad \lim_{N \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [0, T]} |X_t^N - x_t| < \varepsilon \right) = 1.$$

Theorem 3.3, due to Kurtz (1970), says that when the population size N is large, nearly all sample paths of the Markov process $\{X_t^N\}$ stay within ε of a solution of the mean dynamic (3) through time T . In particular, by choosing N large enough, we can ensure that with probability close to one, X_t^N and x_t differ by no more than ε at all times t between 0 and T (Figure 4).⁸

The intuition for this result comes from the law of large numbers. At each revision opportunity, the increment in the process $\{X_t^N\}$ is stochastic. But if we fix dt , then when N is large enough, the expected number of revision opportunities arriving during time interval $I = [t, t + dt]$, namely $NR dt$, is large as well. Since each opportunity leads to an increment of the state of size $\frac{1}{N}$, the overall change in the state during interval I is only of order $R dt$. Thus, during this interval, there are many of revision opportunities, each involving nearly the same transition probabilities. The law of large numbers therefore suggests that the change in $\{X_t^N\}$ during this interval should be almost completely determined by the expected motion of $\{X_t^N\}$, as described by the mean dynamic (3).

It is important to note that Theorem 3.3 cannot be extended to an infinite horizon result. If, for instance, the conditional switch rates $\rho_{ij}^F(x)$ are always positive, then the process $\{X_t^N\}$ is irreducible, and thus must visit every state in $\{X_t^N\}$ infinitely often with probability one. Even so, one can use the mean dynamic (3) to obtain restrictions on the infinite horizon behavior of the process $\{X_t^N\}$, a point we explain in Section 10.1.

⁸Benaïm and Weibull (2003) prove that the rate of convergence in (5) is exponential in N . Roth and Sandholm (2013) extend the theorem to allow the mean dynamic to be an upper hemicontinuous differential inclusion, as in the case of the best response dynamic (see Section 5.2). Hwang et al. (2013) consider stochastic evolutionary processes for large games on spatial domains with long-range interactions, and establish convergence of these processes to solutions of deterministic integro-differential equations.




Figure 4: Deterministic approximation of the Markov process $\{X_t^N\}$.

4. Deterministic Evolutionary Dynamics

4.1 Definition

We are now prepared to state a formal definition of deterministic evolutionary dynamics. Let \mathcal{F} be a set of population games $F: X \rightarrow \mathbb{R}^n$ (with some fixed number of strategies n). Let \mathcal{D} be the class of Lipschitz continuous ordinary differential equations $\dot{x} = V(x)$ on the simplex X , where the vector field $V: X \rightarrow \mathbb{R}^n$ satisfies $V(x) \in TX(x)$ for all $x \in X$.⁹ A map that assigns each game $F \in \mathcal{F}$ a differential equation in \mathcal{D} is called a *deterministic evolutionary dynamic*.

Every well-behaved revision protocol implicitly defines a deterministic evolutionary dynamic. Specifically, suppose that the revision protocol ρ is such that for each $F \in \mathcal{F}$, the function $\rho^F: X \rightarrow \mathbb{R}_+^{n \times n}$ is Lipschitz continuous. Then ρ defines an evolutionary dynamic $\dot{x} = V^F(x)$ by way of equation (3).

4.2 Incentives and Aggregate Behavior

In order to draw links between deterministic evolutionary dynamics and traditional game-theoretic analyses, we must introduce conditions that relate the evolution of aggregate behavior under the dynamics to the incentives in the underlying game. The two most important conditions are these:

⁹In what follows, we often identify the differential equation $\dot{x} = V(x)$ with the vector field V .

- (PC) *Positive correlation* $V^F(x) \neq \mathbf{0} \Rightarrow V^F(x)'F(x) > 0.$
(NS) *Nash stationarity* $V^F(x) = \mathbf{0} \Leftrightarrow x \in NE(F).$

Positive correlation (PC) is the basic restriction on disequilibrium dynamics. In game-theoretic terms, it requires that there be a positive correlation between growth rates and payoffs under the uniform probability distribution on strategies: under this distribution,

$$\mathbb{E}(V^F(x)) = \sum_{k \in S} \frac{1}{n} V_k^F(x) = 0, \text{ and so}$$

$$\text{Cov}(V^F(x), F(x)) = \mathbb{E}(V^F(x)F(x)) - \mathbb{E}(V^F(x))\mathbb{E}(F(x)) = \frac{1}{n}V^F(x)'F(x).$$

Geometrically, (PC) requires that whenever the growth rate vector $V^F(x)$ is nonzero, the angle it forms with the payoff vector $F(x)$ is acute. Thus, away from Nash equilibrium (see Proposition 4.2 below), (PC) restricts the direction of motion to a half space. In this sense, it is as weak a condition as one could hope to be useful. We will see in Sections 6 and 7 that it is useful indeed.¹⁰

Example 4.1. Drawing positive correlation. Figure 5 presents projected payoff vectors and vectors of motion under the replicator dynamic in 123 Coordination (Example 2.6). Evidently, the direction of motion makes an acute angle with the payoff vector whenever the population is not at rest. This illustrates the general fact that the replicator dynamic satisfies positive correlation (PC) (see Section 5.1.3). ♦

Nash stationarity (NS) requires that the Nash equilibria of the game F and the rest points of the dynamic V^F coincide. It can be split into two distinct restrictions. First, (NS) asks that every Nash equilibrium of F be a rest point of V^F . If state x is a Nash equilibrium, then no agent benefits from switching strategies; (NS) demands that in this situation, aggregate behavior is at rest under V^F . This does not imply that individual agents' behavior is also at rest—remember that V^F only describes the *expected* motion of the underlying stochastic process.¹¹

In any case, this direction of Nash stationarity is implied by positive correlation:

Proposition 4.2. *If V^F satisfies (PC), then $x \in NE(F)$ implies that $V^F(x) = \mathbf{0}$.*

¹⁰Versions of this condition are considered by Friedman (1991), Swinkels (1993), Sandholm (2001b), and Demichelis and Ritzberger (2003). For other monotonicity conditions, some of which are particular to imitative dynamics, see Nachbar (1990), Friedman (1991), Samuelson and Zhang (1992), Swinkels (1993), Ritzberger and Weibull (1995), Hofbauer and Weibull (1996), and Viossat (2011).

¹¹This distinction is important in local stability analyses of the underlying stochastic process—see Sandholm (2003).




Figure 5: Condition (PC): Projected payoff vectors (blue) and vectors of motion (violet) under the replicator dynamic in 123 Coordination.

The proof of this result is simple if we use the geometric ideas from Section 2.3. If $x \in NE(F)$, then $F(x) \in NX(x)$ (by Theorem 2.7). But $V^F(x) \in TX(x)$ (since it is a feasible direction of motion from x). Thus $V^F(x)'F(x) \leq 0$ (by the definition of normal cones), so (PC) implies that $V^F(x) = \mathbf{0}$.

Second, Nash stationarity asks that every rest point of V^F be a Nash equilibrium of F . If the current population state is not a Nash equilibrium, then there are agents who would benefit from switching strategies. (NS) requires that enough agents avail themselves of this opportunity that aggregate behavior is not at rest.

5. Families of Evolutionary Dynamics

Evolutionary dynamics are defined in families, with members of a family being derived from qualitatively similar revision protocols. This approach addresses the fact that in practice, one does not expect to know the protocols agents employ with much precision. If we can show that qualitatively similar protocols lead to qualitatively similar aggregate dynamics, then knowing which family a protocol comes from may be enough to draw conclusions about aggregate play.

We present five basic examples of revision protocols and their mean dynamics in Table 1. As the protocols in this section are all reactive, we write them as $\rho_{ij}(\pi, x)$, so that

Revision protocol	Mean dynamic	Name
$\rho_{ij} = x_j[\pi_j - \pi_i]_+$	$\dot{x}_i = x_i \hat{F}_i(x)$	replicator
$\rho_i = M(\pi)$	$\dot{x} \in M(F(x)) - x$	best response
$\rho_{ij} = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}$	$\dot{x}_i = \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} - x_i$	logit
$\rho_{ij} = [\pi_j - \sum_{k \in S} x_k \pi_k]_+$	$\dot{x}_i = [\hat{F}_i(x)]_+ - x_i \sum_{j \in S} [\hat{F}_j(x)]_+$	BNN
$\rho_{ij} = [\pi_j - \pi_i]_+$	$\dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+$	Smith

Table 1: Five basic deterministic dynamics.

$$\rho_{ij}^F(x) = \rho_{ij}(F(x), x).$$

We also introduce some additional notation. We define $\hat{F}_i(x) = F_i(x) - \bar{F}(x)$ to be the *excess payoff* to strategy i ; this is the difference between strategy i 's payoff and the population average payoff $\bar{F}(x) = \sum_{i \in S} x_i F_i(x)$. In addition, we let $M: \mathbb{R}^n \Rightarrow X$ denote the (*mixed*) *maximizer correspondence*, $M(\pi) = \operatorname{argmax}_{y \in X} y' \pi$.

Figure 6 presents phase diagrams for the five basic dynamics when agents are matched to play standard Rock-Paper-Scissors (Example 2.6). In the phase diagrams, colors represent speed of motion: within each diagram, motion is fastest in the red regions and slowest in the blue ones.

The phase diagram of the replicator dynamic displays closed orbits around the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Since this dynamic is based on imitation (or on reproduction), each face and each vertex of the simplex X is an invariant set: a strategy initially absent from the population will never subsequently appear.

The other four dynamics presented in the figure are based on protocols that allow agents to select unused strategies. Under these dynamics, the Nash equilibrium is the sole rest point, and attracts solutions from all initial conditions.¹² Under best response dynamic, solution trajectories quickly change direction and then accelerate when the best response changes; under the Smith dynamic, solutions approach the Nash equilibrium in a less angular, more gradual fashion.

¹²In the case of the logit dynamic, the rest point happens to coincide with the Nash equilibrium only because of the symmetry of the game—see Section 5.2.3.

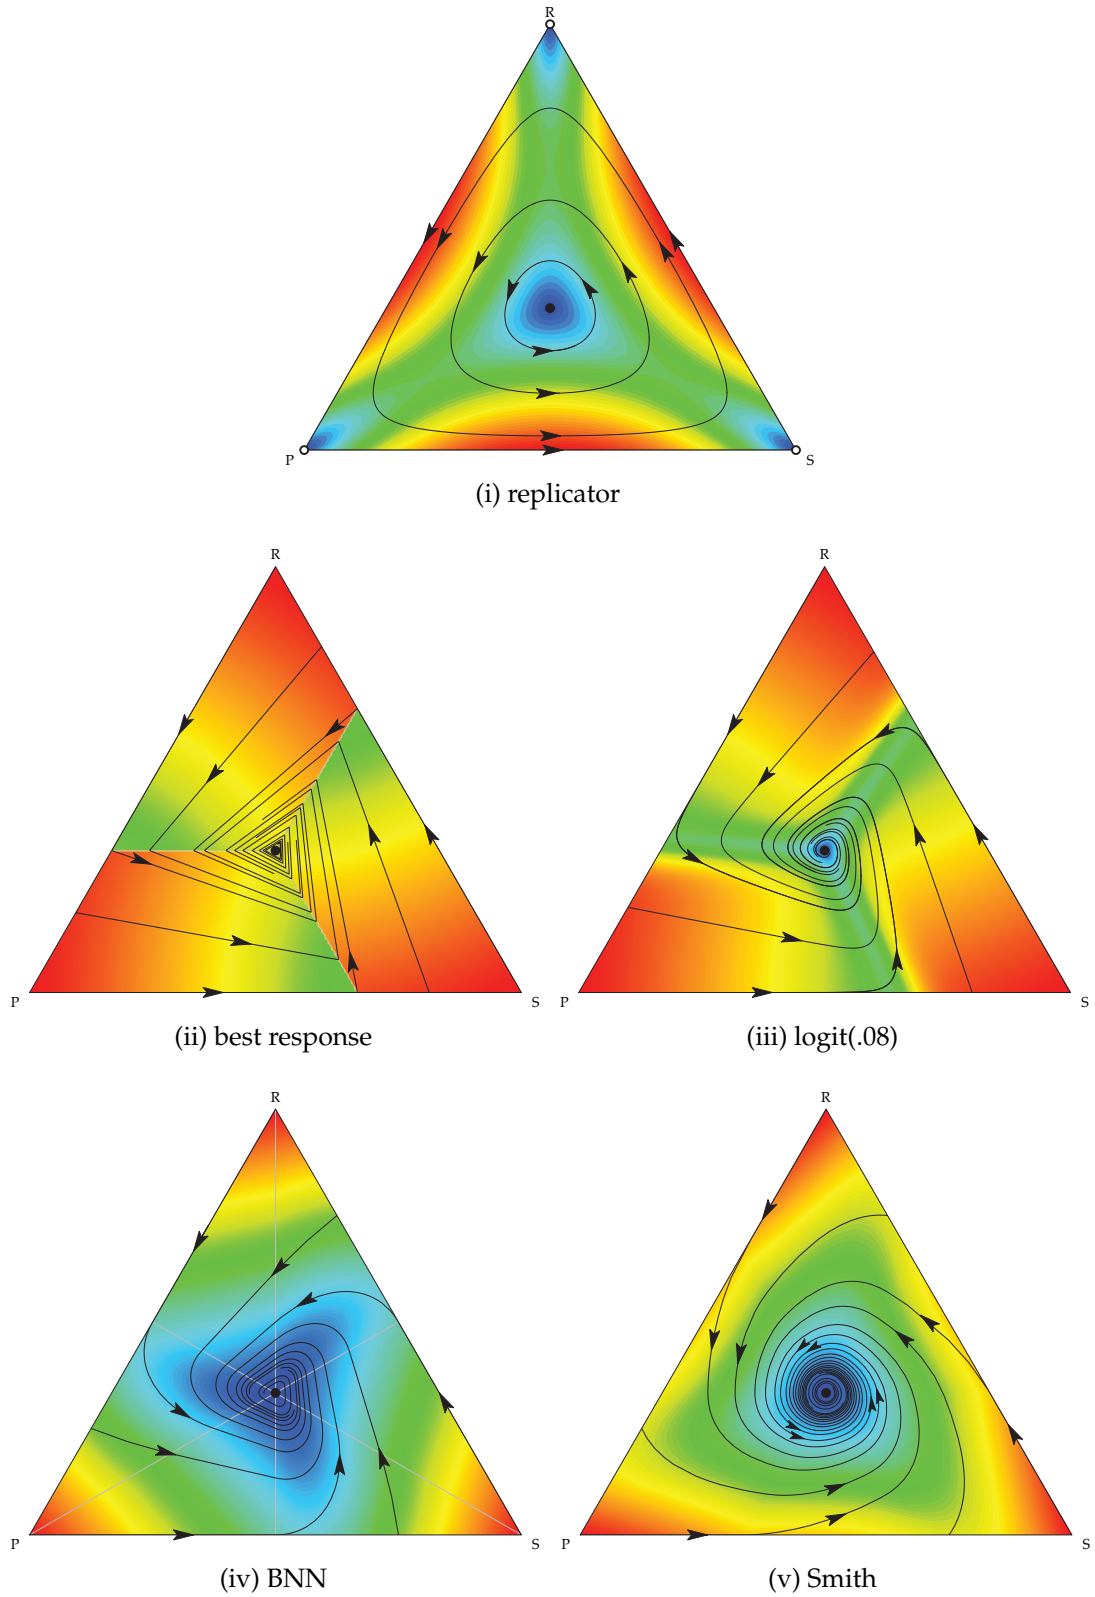


Figure 6: Five basic deterministic dynamics in standard Rock-Paper-Scissors. Colors represent speeds: red is fastest, blue is slowest.

5.1 Imitative Dynamics

Imitative dynamics are the most thoroughly studied dynamics in evolutionary game theory. They are the descendants, or more accurately, a reinterpretation of the game dynamics studied in biology,¹³ and they predominated in the early economic literature on evolutionary game dynamics.¹⁴

Imitative dynamics are derived from protocols under which agents consider switching to the strategy of a randomly sampled opponent, with the ultimate decision depending on payoff considerations. This specification leads imitative dynamics to have a particularly convenient functional form, in which strategies' absolute growth rates are proportional to their current levels of use. This functional form is just as apposite from the biological point of view: there the dynamics reflect births and deaths of different genetic types, and so are naturally expressed relative to current population shares.

5.1.1 Definition

The Lipschitz continuous revision protocol ρ is an *imitative protocol* if

$$(6a) \quad \rho_{ij}(\pi, x) = x_j r_{ij}(\pi, x), \text{ where}$$

$$(6b) \quad \pi_j \geq \pi_i \iff [r_{kj}(\pi, x) - r_{jk}(\pi, x) \geq r_{ki}(\pi, x) - r_{ik}(\pi, x) \text{ for all } i, j, k \in S].$$

The values of r_{ij} are called *conditional imitation rates*. Condition (6b), called *net monotonicity of conditional imitation rates*, says that when strategy j has a higher payoff than strategy i , the net rate of imitation from any strategy k to j exceeds the net rate of imitation from k to i . This condition is enough to ensure that the resulting mean dynamic has the monotonicity properties we desire.

Substituting the functional form (6a) into the general equation (3) for mean dynamics, we obtain

$$(7) \quad \dot{x}_i = x_i \sum_{j \in S} x_j (r_{ji}(F(x), x) - r_{ij}(F(x), x)).$$

When the conditional imitation rates satisfy net monotonicity (6b), we call (7) an *imitative dynamic*. Under these dynamics, strategies' (absolute) growth rates are proportional to their current levels of use, while unused strategies remain so.

¹³See, for instance, Taylor and Jonker (1978), Maynard Smith (1982), and Hofbauer and Sigmund (1988, 1998).

¹⁴See, for example, Nachbar (1990), Samuelson and Zhang (1992), Björnerstedt and Weibull (1996), Weibull (1995), and Hofbauer (1995a).

Formula	Restriction	Interpretation
$\rho_{ij}(\pi, x) = x_j \phi(\pi_j - \pi_i)$	$\text{sgn}(\phi(d)) = \text{sgn}([d]_+)$	imitation via pairwise comparisons
$\rho_{ij}(\pi, x) = a(\pi_i) x_j$	a decreasing	imitation driven by dissatisfaction
$\rho_{ij}(\pi, x) = x_j c(\pi_j)$	c increasing	imitation of success
$\rho_{ij}(\pi, x) = \frac{x_j w(\pi_j)}{\sum_{k \in S} x_k w(\pi_k)}$	w increasing	imitation of success with repeated sampling

Table 2: Some specifications of imitative revision protocols.

Table 2 lists a variety of specifications of imitative protocols consistent with condition (6b). Under the first three, the agent observes the strategy of a single opponent and decides whether to imitate him; under the last, the agent repeatedly draws opponents until he deems one worthy of imitating. The protocols also require different pieces of data. Under imitation via pairwise comparisons, the agent compares his own payoff to that of the opponent he observes; under imitation driven by dissatisfaction, he only observes his own payoff, and under both forms of imitation of success, he only considers the payoffs of those he observes.

5.1.2 Examples

We now consider some important instances of these protocols and the dynamics they induce.

Example 5.1. The replicator dynamic. Example 3.2 introduced the replicator dynamic, which we can express concisely using the notation introduced above as

$$(8) \quad \dot{x}_i = x_i \hat{F}_i(x).$$

Example 3.2 derived this dynamic from imitation via pairwise comparisons with the semilinear functional form $\rho_{ij}(\pi, x) = x_j[\pi_j - \pi_i]_+$. It can also be derived from the linear versions of imitation driven by dissatisfaction ($\rho_{ij}(\pi, x) = (K - \pi_i)x_j$ with K sufficiently large) and of imitation of success ($\rho_{ij}(\pi, x) = x_j(\pi_j - K)$ with K sufficiently small).

While in social science contexts the replicator dynamic is best understood as describing the aggregate consequences of imitation, the dynamic first appeared as a biological model: it was introduced by Taylor and Jonker (1978) to provide a dynamic foundation for

Maynard Smith and Price's (1973) notion of an evolutionarily stable strategy.¹⁵

The replicator dynamic has deep connections with other dynamic models from biology. Hofbauer (1981) shows that the replicator dynamic is equivalent after a nonlinear change of variable to the Lotka-Volterra equation, a fundamental model of the dynamics of ecological systems. Schuster and Sigmund (1983) observe that basic models of population genetics (Crow and Kimura (1970)) and of biochemical evolution (Eigen and Schuster (1979)) can be viewed as special cases of the replicator dynamic; they are also the first to refer to the dynamic by this name.¹⁶ ♦

Example 5.2. The Maynard Smith replicator dynamic. Suppose that agents employ imitation of success with repeated sampling, that payoffs are positive, and that $w(\pi_j) = \pi_j$.¹⁷ In this case, the mean dynamic (7) becomes

$$(9) \quad \dot{x}_i = \frac{x_i \hat{F}_i(x)}{\bar{F}(x)}.$$

This is known as the *Maynard Smith replicator dynamic*, after Maynard Smith (1982). In the current single-population context, dynamics (8) and (9) differ only by a change of speed. With multiple populations this is no longer true, since the changes of speed differ across populations; consequently, the two dynamics have different stability properties in some multipopulation games. ♦

Example 5.3. The imitative logit dynamic. If agents again employ imitation of success with repeated sampling, using the exponential transformation $w(\pi_j) = \exp(\eta^{-1}\pi_j)$ with *noise level* $\eta > 0$, then (7) becomes

$$\dot{x}_i = \frac{x_i \exp(\eta^{-1}F_i(x))}{\sum_{k \in S} x_k \exp(\eta^{-1}F_k(x))} - x_i.$$

This is the *imitative logit dynamic* of Björnerstedt and Weibull (1996) and Weibull (1995). When the noise level η is small, behavior under this dynamic resembles that under the best response dynamic (Section 5.2), at least away from the boundary of the simplex. ♦

¹⁵However, the connection between the two constructions is looser than one might expect: see Sections 7.1 and 10.3.

¹⁶For more on the links among these models, see Hofbauer and Sigmund (1988).

¹⁷In the biology literature, the stochastic evolutionary process generated by this protocol is called the *frequency-dependent Moran process*, after Moran (1962). See Nowak (2006) for references and discussion.

5.1.3 Basic properties

We now consider some general properties of imitative dynamics. Later sections will describe their behavior in various classes of games and interesting examples.

Theorem 3.1 showed that in general, Lipschitz continuous mean dynamics admit unique forward solutions $\{x_t\}_{t \geq 0}$ from every initial condition in the simplex. For imitative dynamics (7), under which \dot{x}_i is proportional to x_i , more is true: solutions $\{x_t\}_{t \in (-\infty, \infty)}$ exist in both forward and backward time, and along any solution the support of x_t does not change.

Turning to payoff monotonicity properties, let us rewrite the dynamic (7) as

$$(10) \quad \dot{x}_i = V_i(x) = x_i G_i(x), \text{ where } G_i(x) = \sum_{k \in S} x_k (r_{ki}(F(x), x) - r_{ik}(F(x), x)).$$

If strategy $i \in S$ is in use, then $G_i(x) = V_i(x)/x_i$ represents the *percentage growth rate* of the number of agents using this strategy. Condition (6b) implies that strategies' percentage growth rates are ordered by their payoffs,

$$(11) \quad G_i(x) \geq G_j(x) \text{ if and only if } F_i(x) \geq F_j(x),$$

a property called *monotonicity of percentage growth rates*. This property, a strong restriction on strategies' percentage growth rates, can be shown to imply positive correlation (PC), a weak restriction on strategies' absolute growth rates.¹⁸

It is easy to show that the rest points of any imitative dynamic V_F include the Nash equilibria of F ; in fact, this follows from the previous claim and Proposition 4.2. But imitative dynamics may also have non-Nash rest points. For instance, any pure state e_i is a rest point under (7): since everyone is playing the same strategy, imitation leads to stasis. In fact, x is a rest point of (7) if and only if it is a *restricted equilibrium* of F , meaning that it is a Nash equilibrium of a restricted version of F in which only strategies in the support of x can be played. Non-Nash rest points of (7) are not natural predictions of play: they cannot be locally stable, nor can they be approached by any interior solution trajectory (Bomze (1986), Nachbar (1990)). Even so, continuous dynamics move slowly near rest points, so escape from the vicinity of non-Nash rest points is necessarily slow.

¹⁸See Sandholm (2010c, Theorem 5.4.9).

5.1.4 Inflow-outflow symmetry

To introduce a final property, let us compare the general equation for mean dynamics with that for imitative dynamics:

$$(3) \quad \dot{x}_i = \sum_{j \in S} x_j \rho_{ji}^F(x) - x_i \sum_{j \in S} \rho_{ij}^F(x),$$

$$(7) \quad \dot{x}_i = x_i \sum_{j \in S} x_j \left(r_{ji}(F(x), x) - r_{ij}(F(x), x) \right).$$

In general, mean dynamics exhibit an asymmetry between their inflow and outflow terms: the rate of switches from strategy j to strategy i is proportional to x_j , while the rate of switches from strategy i to strategy j is proportional to x_i . But under imitative dynamics, both of these rates are proportional to both x_i and x_j . This property, called *inflow-outflow symmetry*, underlies a number of properties of imitative dynamics that fail for other continuous dynamics. Figure 6 illustrates this point in the standard Rock-Paper-Scissors game: while the other dynamics converge, the replicator dynamic exhibits a continuum of closed orbits. A more surprising distinction concerns the behaviors of the dynamics in games with dominated strategies: see Sections 8.1 and 9.2.

Another dynamic that satisfies inflow-outflow symmetry is the *projection dynamic* (Nagurney and Zhang (1997), Lahkar and Sandholm (2008)). On the interior of the simplex, this dynamic is defined by

$$(12) \quad \dot{x} = \Phi F(x),$$

so that the direction of motion is given by the projected payoff vector.¹⁹ It can be derived from revision protocols reflecting “revision driven by insecurity”, under which the conditional switch rate ρ_{ij} is inversely proportional to x_i . The rate of switches from strategy i to strategy j under (12) is proportional to neither x_i nor x_j . By virtue of this inflow-outflow symmetry, the projection dynamic exhibits close connections to the replicator dynamic, at least in the interior of the simplex (Sandholm et al. (2008)). While the projection dynamic is mathematically appealing, the discontinuities in the dynamic and its protocols at the boundary of the simplex raise doubts about its appropriateness for applications.

¹⁹The dynamic is defined globally by $\dot{x} = \text{Proj}_{TX(x)}(F(x))$, where the right hand side represents the closest point projection of $F(x)$ onto $TX(x)$. This definition ensures that unique forward solution trajectories exist from every initial condition in X .

5.2 The Best Response Dynamic and Related Dynamics

Traditionally, choice in game theory is based on optimization. The original evolutionary dynamic embodying this paradigm is the best response dynamic of Gilboa and Matsui (1991) (see also Matsui (1992) and Hofbauer (1995b)). This dynamic can also be derived as a continuous-time version of the well-known *fictitious play* process of Brown (1949, 1951) (see Section 10.2). We now introduce the best response dynamic and some interesting variations, most notably those defined by means of payoff perturbations.

5.2.1 Target protocols and target dynamics

Under the revision protocol for the best response dynamic, an agent's conditional switch rates do not depend on his current strategy. Protocols with this feature have identical rows: $\rho_{i\hat{i}}^F(x) = \rho_{i\hat{i}}^F(x)$ for all $x \in X$ and $i, \hat{i}, j \in S$. In this case, we use $\tau^F \equiv \rho_{i\hat{i}}^F$ to refer to the common row of ρ^F , and call τ^F a *target protocol*.

Substitution into equation (3) shows that target protocols generate mean dynamics of the form

$$(13) \quad \dot{x}_i = \tau_i^F(x) - x_i \sum_{j \in S} \tau_j^F(x),$$

which we call *target dynamics*. When $\lambda^F(x) = \sum_{j \in S} \tau_j^F(x)$ is not zero, we can define $\sigma^F(x) \in X$ by $\sigma_j^F(x) = \tau_j^F(x) / \lambda^F(x)$ and express the target dynamic (13) as

$$(14) \quad \dot{x} = \lambda^F(x)(\sigma^F(x) - x).$$

Geometrically, equation (14) says that the population state moves from its current position x toward the target state $\sigma^F(x)$ at rate $\lambda^F(x)$.

A further simple property of the best response dynamic's protocol is that its entries sum to one, $\sum_{j \in S} \tau_j^F(x) = 1$ for all $x \in X$, so that the conditional switch rates τ_j^F are actually conditional switch probabilities (see Section 3.1). To highlight this property, we denote the protocol by $\sigma^F(x)$, and express the target dynamic as

$$(15) \quad \dot{x} = \sigma^F(x) - x.$$

5.2.2 The best response dynamic

Under the best response protocol, agents receive revision opportunities at a unit rate, and use these opportunities to switch to a current best response. This protocol is reactive

($\rho^F(x) = \rho(F(x), x)$), and it is a target protocol with unit row sum ($\rho_i(F(x), x) = \sigma(F(x), x) \in X$); moreover, it does not condition directly on the population state ($\sigma(F(x), x) = \sigma(F(x))$). It is defined formally by

$$(16a) \quad \sigma(\pi) = M(\pi),$$

where the map $M: \mathbb{R}^n \Rightarrow X$, defined by

$$(16b) \quad M(\pi) = \operatorname{argmax}_{y \in X} y' \pi,$$

is the (*mixed*) *maximizer correspondence*. Substituting (16a) and (16b) into the mean dynamic equation (15), we obtain the *best response dynamic*,

$$(17) \quad \dot{x} \in M(F(x)) - x.$$

This dynamic can also be expressed as

$$\dot{x} \in B^F(x) - x,$$

where $B^F = M \circ F$ is the (*mixed*) *best response correspondence* for F .

Since the maximizer correspondence M is set-valued and discontinuous, the best response dynamic is a *differential inclusion*, and lies outside the framework developed in Section 3. Thus, the basic results on existence and uniqueness of solutions (Theorem 3.1) and on deterministic approximation (Theorem 3.3) do not apply here. Fortunately, versions of both of these results are available for the current setting. Since M is a convex-valued and upper hemicontinuous correspondence, results from the theory of differential inclusions imply the existence of a *Carathéodory solution* from every initial condition: a Lipschitz continuous trajectory $\{x_t\}_{t \geq 0}$ that satisfies $\dot{x}_t \in V(x_t)$ at all but a measure zero set of times.²⁰ But solutions are generally not unique: for instance, starting from a mixed equilibrium of a coordination game, there is not only a stationary solution, but also solutions that head immediately toward a pure equilibrium, as well as solutions that do so after an initial delay. Regarding deterministic approximation, one can define versions of the stochastic evolutionary process $\{X_t^N\}$ that account for the multivaluedness of M , and whose sample paths are approximated by solutions of (17).²¹

Despite these technical complications, most solutions of the best response dynamic

²⁰The formulation of the best response dynamic as a differential inclusion is due to Hofbauer (1995b). For an introduction to differential inclusions, see Smirnov (2002).

²¹See Roth and Sandholm (2013).

take a very simple form. To draw a phase diagram of (17) for a given game F , one divides the state space X into the best response regions for each strategy. When the state is in the best response region for strategy i , solutions of (17) move directly toward pure state e_i . At states admitting multiple best responses, more than one direction of motion is possible in principle, though not always in practice.

Example 5.4. The best response dynamic for standard RPS. The best response dynamic for standard Rock-Paper-Scissors (Example 2.6) is illustrated in Figure 6(ii). Within each best response region, solutions proceed toward the relevant vertex. At the boundaries of these regions, the only continuation path consistent with (17) turns instantly toward the next vertex in the best response cycle. In the limit, the process converges to the mixed equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ after an infinite number of turns.²² ♦

Despite its absence in this example, multiplicity of solutions under the best response dynamic is quite common. We noted earlier that there are always multiple solutions emanating from mixed equilibria of coordination games. In other cases, multiplicity of solutions can lead to complicated dynamics, including solutions that cycle in and out of equilibrium in perpetuity (Hofbauer (1995b)).

It is easy to verify that the best response dynamic satisfies versions of positive correlation (PC) and Nash stationarity (NS) suitable for differential inclusions. For the former, note that if $y \in M(F(x)) = B^F(x)$ is any best response to x , we have

$$(y - x)'F(x) = \max_{j \in S} F_j(x) - \bar{F}(x) = \max_{j \in S} \hat{F}_j(x) \geq 0,$$

with equality only when x is a Nash equilibrium (see Proposition 5.6 below). And it is clear from (17) that stationary solutions coincide with Nash equilibria.

While the best response dynamic is the simplest dynamic based on the idea of exact optimization, it is not the only such dynamic. Balkenborg et al. (2013) consider the *refined best response dynamic*, under which agents only switch to best responses that are robust, in the sense that they are unique best responses at some nearby state. Under the *tempered best response dynamics* of Zusai (2011), all revising agents switch to best responses, but the rate at which an agent revises declines with the payoff of his current strategy. A third alternative, the *sampling best response dynamics* of Oyama et al. (2012), is introduced in Section 8.3.

²²For a detailed analysis, see Gaunersdorfer and Hofbauer (1995).

5.2.3 Perturbed best response dynamics

For both technical convenience and realism, it is natural to consider variants of the best response dynamic that define smooth dynamical systems. The most important alternative, introduced by Fudenberg and Levine (1998) (see also Fudenberg and Kreps (1993)), supposes that agents optimize after their payoffs have been subject to perturbations. The perturbations ensure that agents' choice probabilities vary smoothly with payoffs, and so lead to differentiable dynamics. This allows us to avail ourselves of the full toolkit for analyzing smooth dynamics, including both linearization around rest points and the theory of cooperative differential equations (Section 8.2). The resulting perturbed best response dynamics have important connections with work in stochastic stability theory (Section 10.1) and models of heuristic learning in games (Section 10.2); they also describe the dynamics of aggregate behavior in heterogeneous populations of exact optimizers (Ellison and Fudenberg (2000), Ely and Sandholm (2005)).

Perturbed best response protocols are target protocols of the form

$$(18) \quad \sigma(\pi) = \tilde{M}(\pi),$$

where the *perturbed best response function* $\tilde{M}: \mathbb{R}^n \rightarrow \text{int}(X)$ is a smooth approximation of the maximizer correspondence M . This function is defined most conveniently in terms of a smooth *deterministic perturbation* $v: \text{int}(X) \rightarrow \mathbb{R}$ of the payoff to each *mixed strategy*:

$$(19a) \quad \tilde{M}^v(\pi) = \underset{y \in \text{int}(X)}{\text{argmax}} (y' \pi - v(y)), \text{ where}$$

$$(19b) \quad z' \nabla^2 v(y) z > 0 \text{ for all } z \in TX \text{ and } y \in \text{int}(X), \text{ and}$$

$$(19c) \quad \lim_{k \rightarrow \infty} y_k \in \text{bd}(X) \Rightarrow \lim_{k \rightarrow \infty} |\nabla v(y_k)| = \infty.$$

The convexity condition (19b) and the steepness condition (19c) ensure that the optimal solution to (19a) is unique and lies in the interior of X . Indeed, taking the first order condition for (19a) shows that

$$(20) \quad \tilde{M}^v(\pi) = (\nabla v)^{-1}(\Phi \pi).$$

From a game-theoretic point of view, it is more natural to define perturbed maximization using *stochastic perturbations* of the payoff of each *pure strategy*:

$$(21) \quad \tilde{M}_i^\varepsilon(\pi) = \mathbb{P} \left(i = \underset{j \in S}{\text{argmax}} \pi_j + \varepsilon_j \right),$$

where ε is a random vector that admits a density function that is positive throughout \mathbb{R}^n . While the explicit expression for \tilde{M}^ε is quite cumbersome, one can avoid working with it directly: Hofbauer and Sandholm (2002) show that under appropriate smoothness conditions, any maximizer function of form (21) can be represented using a deterministic perturbation as in (19a)–(19c).

Example 5.5. Logit choice. The best known perturbed maximizer function is the *logit choice function* with noise level $\eta > 0$:

$$(22) \quad \tilde{M}_i(\pi) = \frac{\exp(\eta^{-1}\pi_i)}{\sum_{j \in S} \exp(\eta^{-1}\pi_j)}.$$

When η is large, $\tilde{M}(\pi)$ is close to a uniform probability vector. When η is close to 0, $\tilde{M}(\pi)$ is a close approximation to the maximizer $M(\pi)$, but places positive probability on every strategy. The logit choice function can be expressed in form (21) using stochastic perturbations that are i.i.d. with a double exponential distribution. It can also be expressed in form (5.2.3) by letting the deterministic perturbation v be the negated entropy function $v(y) = \eta \sum_{j \in S} y_j \log y_j$. ♦

If agents revise using a perturbed best response protocol (5.2.3), then aggregate behavior evolves according to the *perturbed best response dynamic*

$$(23) \quad \dot{x} = \tilde{M}^v(F(x)) - x.$$

When the perturbations that generate \tilde{M}^v are small in a suitable sense, the dynamic (23) is a smooth approximation of the best response dynamic (17).

Because its definition uses payoff perturbations, perturbed best response dynamics cannot satisfy positive correlation (PC) and Nash stationarity (NS) exactly. They do, however, satisfy perturbed versions of these conditions. Considering the latter first, observe that the rest points of (23) are the fixed points of the *perturbed best response function* $\tilde{B}^v = \tilde{M}^v \circ F$. These *perturbed equilibria* approximate Nash equilibria when the perturbations are small.²³

To obtain a useful alternative characterization, define the *virtual payoffs* $\tilde{F}^v: X \rightarrow \mathbb{R}^n$ associated with the pair (F, v) by $\tilde{F}^v(x) = F(x) - \nabla v(x)$. By way of interpretation, note that the convexity and steepness of v ensure that strategies played by few agents have high virtual payoffs. Equation (20) implies that the perturbed equilibria for the pair (F, v) are precisely those states for which $\tilde{F}^v(x)$ is a constant vector.

²³In the experimental literature, perturbed equilibria are known as *quantal response equilibria*; see McKelvey and Palfrey (1995) and Goeree et al. (2008).

Virtual payoffs are also used to define the appropriate analogue of positive correlation. Hofbauer and Sandholm (2002, 2007) show that the perturbed best response dynamics (23) satisfy *virtual positive correlation*

$$(24) \quad V^{E,v}(x) \neq \mathbf{0} \text{ implies that } V^{E,v}(x)' \tilde{F}^v(x) > 0.$$

This condition is just what is needed to extend stability and convergence results for the best response dynamic to perturbed best response dynamics—see Sections 6.5 and 7.5.

5.3 Excess Payoff and Pairwise Comparison Dynamics

A basic question addressed by evolutionary game dynamics is whether Nash equilibrium can be interpreted as stationary behavior among agents who employ simple myopic rules. Indeed, a version of this interpretation was offered by Nash himself in his doctoral dissertation.²⁴ Imitative dynamics fail to satisfy Nash stationarity, and so do not provide an ideal basis for this interpretation. Best response dynamics are also not ideal for this purpose: they are based on discontinuous revision protocols, which require more precise information than simple agents should be expected to possess.

In this section, we introduce two classes of continuous dynamics that satisfy Nash stationarity, and so provide the interpretation of Nash equilibrium we seek.

5.3.1 Excess payoff dynamics

The excess payoff function $\hat{F}: X \rightarrow \mathbb{R}^n$ for game F , defined by

$$\hat{F}_i(x) = F_i(x) - \bar{F}(x) = F_i(x) - \sum_{j \in S} x_j F_j(x)$$

describes the performance of each strategy relative to the population average. Clearly, the excess payoff vector $\hat{F}(x)$ cannot lie in the interior of the negative orthant \mathbb{R}_-^n , as this would mean that all strategies receive a worse-than-average payoff. In fact, it is not difficult to establish that in this context, the boundary of the negative orthant plays a special role:

Proposition 5.6. $x \in NE(F)$ if and only if $\hat{F}(x) \in \text{bd}(\mathbb{R}_-^n)$.

To explain this proposition, we let $\mathbb{R}_*^n = \mathbb{R}^n \setminus \text{int}(\mathbb{R}_-^n)$ denote the set of vectors in \mathbb{R}^n with at least one nonnegative component. If the excess payoff vector lies on the boundary of this set, $\text{bd}(\mathbb{R}_*^n) = \text{bd}(\mathbb{R}_-^n)$, then the maximal component of the excess payoff vector is

²⁴See Nash (1950) and Weibull (1996).

zero. Viewed in this light, Proposition 5.6 says that the Nash equilibria are just those states at which no strategy receives an above-average payoff.

Excess payoff protocols are target protocols (Section 5.2.1) under which conditional switch rates are expressed as functions of the excess payoff vector $\hat{F}(x)$.²⁵ Specifically, we call a target protocol τ an *excess payoff protocol* if it is Lipschitz continuous and satisfies

$$(25a) \quad \tau_j(\pi, x) = \tau_j(\hat{\pi}), \text{ where } \hat{\pi}_i = \pi_i - x'_i \pi, \text{ and}$$

$$(25b) \quad \hat{\pi} \in \text{int}(\mathbb{R}_*^n) \Rightarrow \tau(\hat{\pi})' \hat{\pi} > 0.$$

Condition (25b), called *acuteness*, requires that away from Nash equilibrium, strategies with higher growth rates tend to be those with higher excess payoffs.

Substituting equation (25a) into the mean dynamic (13) yields the corresponding class of evolutionary dynamics, the *excess payoff dynamics*:

$$(26) \quad \dot{x}_i = \tau_i(\hat{F}(x)) - x_i \sum_{j \in S} \tau_j(\hat{F}(x)).$$

Example 5.7. The BNN dynamic. If the protocol τ takes the semilinear form

$$(27) \quad \tau_i(\hat{\pi}) = [\hat{\pi}_i]_+,$$

we obtain the *Brown-von Neumann-Nash (BNN) dynamic*:

$$(28) \quad \dot{x}_i = [\hat{F}_i(x)]_+ - x_i \sum_{j \in S} [\hat{F}_j(x)]_+.$$

This dynamic was introduced in the early days of game theory by Brown and von Neumann (1950) in the context of symmetric zero-sum games. Nash (1951) used a discrete-time analogue of this dynamic as the basis for a simple proof of existence of equilibrium via Brouwer's theorem. The dynamic was then forgotten for 40 years before being reintroduced by Skyrms (1990), Swinkels (1993), Weibull (1996), and Hofbauer (2000). ♦

It is not difficult to verify that the BNN dynamic satisfies both positive correlation (PC) and Nash stationarity (NS). Sandholm (2005a) shows that these properties are satisfied by all excess payoff dynamics.

²⁵Best response protocols and perturbed best response protocols can also be expressed in this way—a point we return to in Section 7.5.

5.3.2 Pairwise comparison dynamics

Excess payoffs dynamics are not completely satisfactory as a model of behavior in population games. To use an excess payoff protocol, an agent needs to know the vector of excess payoffs, and so, implicitly, the average payoff obtained in the population. Unless this information were provided by a planner, it is not information that agents could easily obtain.

As a more credible alternative, we consider dynamics based on pairwise comparisons. When an agent receives a revision opportunity, he selects an alternative strategy at random. He compares its payoff to that of his current strategy, and considers switching only if the former exceeds the latter.

Formally, *pairwise comparison protocols* are Lipschitz continuous protocols $\rho: \mathbb{R}^n \times X \rightarrow \mathbb{R}_+^{n \times n}$ that satisfy *sign preservation*:

$$(29) \quad \text{sgn}(\rho_{ij}(\pi, x)) = \text{sgn}([\pi_j - \pi_i]_+) \text{ for all } i, j \in S.$$

The resulting evolutionary dynamics, described by equation (3), are called *pairwise comparison dynamics*.

Example 5.8. The Smith dynamic. Suppose the revision protocol ρ takes the semilinear form

$$(30) \quad \rho_{ij}(\pi, x) = [\pi_j - \pi_i]_+,$$

Inserting this formula into the mean dynamic (3), we obtain the *Smith dynamic*:

$$(31) \quad \dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+.$$

This dynamic was introduced by Smith (1984), who used it to model disequilibrium adjustment by drivers in highway networks. ♦

Protocol (30) is closely related to the pairwise proportional imitation protocol from Example 3.2, $\rho_{ij}(\pi, x) = x_j [\pi_j - \pi_i]_+$. Indeed, the protocols only differ in how candidate strategies are chosen. Under the imitative protocol, a revising agent obtains a candidate strategy by observing the choice of an opponent. Under protocol (30), he obtains a candidate strategy by choosing uniformly from a list of all strategies.

This difference between imitative and pairwise comparison protocols has clear consequences for the dynamics of aggregate behavior. Under imitative dynamics, rare strategies are unlikely to be chosen, and unused strategies, even optimal ones, are never chosen. Pairwise comparison dynamics have neither of these properties, and as a consequence

Family	Example	Continuity	Data Req.	(PC)	(NS)
imitation	replicator	yes	weak	yes	no
optimization	best response	no	moderate	yes	yes
perturbed optimization	logit	yes	moderate	approx.	approx.
excess payoff	BNN	yes	strong	yes	yes
pairwise comparison	Smith	yes	weak	yes	yes

Table 3: Families of revision protocols and evolutionary dynamics, and their properties.

satisfy not only positive correlation (PC), but also Nash stationarity (NS): see Smith (1984) and Sandholm (2010b).

In fact, pairwise comparison protocols can be used in combination with imitative protocols to improve the performance of the latter. Sandholm (2005a, 2010b) considers *hybrid protocols* that combine imitation with consideration of unused strategies via protocols of form (25) or (29). The resulting *hybrid dynamics* satisfy both positive correlation (PC) and Nash stationarity (NS). It is thus not imitation per se, but rather the exclusive use of imitation, that allows non-Nash rest points to exist.

Our main conclusions about families of revision protocols and evolutionary dynamics are summarized in Table 3. For each family, the table describes the continuity and data requirements of the protocols, and the incentive properties of the corresponding dynamics.

6. Potential Games

In this section and the following two, we consider classes of population games that have attractive theoretical properties and are useful in applications. Games in these classes admit simple characterizations of Nash equilibrium, and ensure global convergence under various evolutionary dynamics.

The most general convergence results are available for potential games, in which all information about incentives can be captured by a scalar-valued function defined on the set of population states. Dynamics satisfying positive correlation (PC) and Nash stationarity (NS) ascend this function and converge to Nash equilibrium.

The first appearance of potential functions in game theory is in the work of Beckmann et al. (1956), who use potential function arguments to analyze congestion games. Rosenthal (1973) introduced congestion games with finite numbers of players, motivating Monderer and Shapley’s (1996) definition of finite-player potential games. Potential

function arguments have long been used in models from population genetics that are equivalent to the replicator dynamic in a normal form potential game; see Kimura (1958), Shahshahani (1979), Akin (1979), and Hofbauer and Sigmund (1988, 1998). Our presentation of potential games played by continuous populations of agents follows Sandholm (2001b, 2009b).

6.1 Population Games and Full Population Games

Some of the classes of games introduced in the coming sections can be characterized in terms of the externalities that players of different strategies impose on one another. In discussing these externalities, it is natural to consider the effect of adding new agents playing strategy j on the payoffs of agents currently choosing strategy i . In principle, this effect should be captured by the partial derivative $\frac{\partial F_i}{\partial x_j}$. However, since payoffs in population games are only defined on the simplex, this partial derivative does not exist.

This difficulty can be circumvented by considering *full population games*, in which payoffs are defined on the positive orthant \mathbb{R}_+^n . We can interpret the extended payoff functions as describing the payoffs that would arise were the population size to change. Although there are some subtleties involved in the use of these extensions, in the end they are harmless, and we employ them here with little further comment.

6.2 Definition, Characterization, and Interpretation

Let $F : \mathbb{R}_+^n \rightarrow \mathbb{R}^n$ be a (full) population game. We call F a *potential game* if there exists a continuously differentiable function $f : \mathbb{R}_+^n \rightarrow \mathbb{R}$, called a *potential function*, satisfying

$$\begin{aligned} \nabla f(x) &= F(x) \text{ for all } x \in \mathbb{R}_+^n, \text{ or equivalently} \\ \frac{\partial f}{\partial x_i}(x) &= F_i(x) \text{ for all } i \in S \text{ and } x \in \mathbb{R}_+^n. \end{aligned}$$

Thus the partial derivatives of the potential function are the payoff functions of the game.

If the payoff function F is continuously differentiable, then well-known results from calculus tell us that F is a potential game if and only if

$$\begin{aligned} DF(x) &\text{ is symmetric for all } x \in \mathbb{R}_+^n, \text{ or equivalently} \\ \frac{\partial F_i}{\partial x_j}(x) &= \frac{\partial F_j}{\partial x_i}(x) \text{ for all } i, j \in S \text{ and } x \in \mathbb{R}_+^n. \end{aligned}$$

This condition, which we call *full externality symmetry*, has a simple game-theoretic in-

terpretation: the effect on the payoffs to strategy i of introducing new agents choosing strategy j always equals the effect on the payoffs to strategy j of introducing new agents choosing strategy i .²⁶

For intuition concerning the potential function itself, suppose that some members of the population switch from strategy i to strategy j , so that the population state moves in direction $z = e_j - e_i$. If these switches improve the payoffs of those who switch, then

$$\frac{\partial f}{\partial z}(x) = \nabla f(x)'z = F(x)'z = F_j(x) - F_i(x) > 0.$$

Thus profitable strategy revisions increase potential.

For more general sorts of adjustment, we have the following simple lemma:

Lemma 6.1. *Let F be a potential game with potential function f , and suppose the dynamic V^F satisfies positive correlation (PC). Then along any solution trajectory $\{x_t\}$, we have $\frac{d}{dt}f(x_t) > 0$ whenever $\dot{x}_t \neq \mathbf{0}$.*

The proof follows directly from the chain rule:

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

where equality holds only if $V^F(x_t) = \mathbf{0}$.

6.3 Examples

Potential games admit a number of important applications, three of which are described next.

Example 6.2. Matching in games with common interests. Example 2.2 considered matching in the two-player symmetric normal form game $A \in \mathbb{R}^{n \times n}$, which generates the population game $F(x) = Ax$. The game A is said to exhibit *common interests* if A is a symmetric matrix ($A_{ij} = A_{ji}$), so that two matched players always receive the same payoffs. Matching in two-player games with common interests defines a fundamental model from population genetics; the common interest assumption reflects the shared fate of two genes that inhabit the same organism.²⁷

²⁶If we only defined payoffs on the simplex, the corresponding condition would be *externality symmetry*, which requires that $DF(x)$ be symmetric with respect to $TX \times TX$ (in other words, that $z'DF(x)\hat{z} = \hat{z}'DF(x)z$ for all $z, \hat{z} \in TX$ and $x \in X$.)

²⁷See Hofbauer and Sigmund (1988, 1998).

Since $DF(x) = A$, the population game derived from a game with common interests is a potential game. Its potential function, $f(x) = \frac{1}{2}x'Ax = \frac{1}{2}\bar{F}(x)$, is one-half of the average payoff function. Thus Lemma 6.1 implies that in common interest games, evolutionary dynamics satisfying positive correlation (PC) improve social outcomes. ♦

Example 6.3. Congestion games. Example 2.3 introduced congestion games, whose payoff functions are of the form

$$F_i(x) = - \sum_{\ell \in \mathcal{L}_i} c_\ell(u_\ell(x)).$$

In the context of highway congestion, each ℓ represents a link in the highway network, and \mathcal{L}_i is the set of links that make up path i .

In a congestion game, the marginal effect of adding an agent to path i on the payoffs to drivers on path j is due to the marginal increases in congestion on the links the two paths have in common. The marginal effect of adding an agent to path i on the payoffs to drivers on path j is the same:

$$\frac{\partial F_i}{\partial x_j}(x) = - \sum_{\ell \in \mathcal{L}_i \cap \mathcal{L}_j} c'_\ell(u_\ell(x)) = \frac{\partial F_j}{\partial x_i}(x).$$

Thus congestion games satisfy externality symmetry, and so are potential games. Their potential functions are of the form

$$(32) \quad f(x) = - \sum_{\ell \in \mathcal{L}} \int_0^{u_\ell(x)} c_\ell(z) dz. \quad \blacklozenge$$

The potential function (32) of a congestion game is generally unrelated to its average payoff function,

$$\bar{F}(x) = - \sum_{\ell \in \mathcal{L}} u_\ell(x) c_\ell(u_\ell(x)).$$

Dafermos and Sparrow (1969) observed that if each cost function is a monomial of the same degree $\eta \geq 0$, so that $c_\ell(u) = a_\ell u^\eta$, then the potential function is proportional to average payoffs: $f(x) = \frac{1}{\eta+1} \bar{F}(x)$. For general full potential games, it follows from Euler's theorem that the potential function f is proportional to aggregate payoffs if and only if F is a homogeneous function of degree $k > -1$ (Hofbauer and Sigmund (1988), Sandholm (2001b)).

Congestion games have received considerable attention in the computer science liter-

ature. Much of this work focuses on so-called “price of anarchy” results, which bound the ratio of the total delay in the network in Nash equilibrium to the minimal feasible delay. The most basic and best-known result in this literature, due to Roughgarden and Tardos (2002, 2004) (see also Correa et al. (2004, 2008)), shows that in congestion games with nonnegative, increasing, affine cost functions, total delays in Nash equilibrium must be within a factor of $\frac{4}{3}$ of the minimal feasible delay. In fact, tight bounds can be obtained for quite general classes of cost functions.

Example 6.4. Games generated by variable pricing schemes. Population games can be viewed as models of externalities for environments with many agents. One way to force agents to internalize the externalities they impose upon others is to introduce pricing schemes. Given an arbitrary population game F with average payoff function \bar{F} , we define a new game \tilde{F} by

$$\tilde{F}_i(x) = F_i(x) + \sum_{j \in S} x_j \frac{\partial F_j}{\partial x_i}(x).$$

We interpret the second term as a price (either subsidy or tax) imposed by a planner. It represents the marginal effect that a strategy i player has on the payoffs of his opponents.

Observe that

$$\frac{\partial \bar{F}}{\partial x_i}(x) = \frac{\partial}{\partial x_i} \sum_{j \in S} x_j F_j(x) = F_i(x) + \sum_{j \in S} x_j \frac{\partial F_j}{\partial x_i}(x) = \tilde{F}_i(x).$$

In words, the augmented game \tilde{F} is a full potential game, and its full potential function is the average payoff function of the original game F . Thus, when individual agents switch strategies in response to the combination of original payoffs and prices, average payoffs in the population increase. Sandholm (2002, 2005b, 2007b) uses this construction as the basis for an evolutionary approach to implementation theory. ♦

6.4 Characterization of Equilibrium

The fact that evolutionary processes increase potential suggests a connection between local maximization of potential and Nash equilibrium. We therefore consider the problem of maximizing potential over the set of population states:

$$\max f(x) \quad \text{subject to} \quad \sum_{j \in S} x_j = 1 \quad \text{and} \quad x_i \geq 0 \quad \text{for all } i \in S.$$

The Lagrangian for this maximization problem is

$$L(x, \mu, \lambda) = f(x) + \mu \left(1 - \sum_{i \in S} x_i \right) + \sum_{i \in S} \lambda_i x_i,$$

so the Kuhn-Tucker first-order necessary conditions for maximization are

$$(33a) \quad \frac{\partial f}{\partial x_i}(x) = \mu - \lambda_i \quad \text{for all } i \in S,$$

$$(33b) \quad \lambda_i x_i = 0, \quad \text{for all } i \in S, \text{ and}$$

$$(33c) \quad \lambda_i \geq 0 \quad \text{for all } i \in S.$$

These conditions characterize the Nash equilibria of a potential game.

Theorem 6.5. *Let F be a potential game with potential function f . Then x is a Nash equilibrium of F if and only if (x, μ, λ) satisfies (33a)–(33c) for some $\lambda \in \mathbb{R}^n$ and $\mu \in \mathbb{R}$.*

The proof is simple: the multiplier μ is the maximal payoff, and the multiplier λ_i is the payoff deficit of strategy i relative to the optimal strategy.

Theorem 6.5 provides a simple proof of existence of equilibrium in potential games. Since f is continuous and X is compact, the former attains its maximum on the latter, and the theorem implies that any maximizer is a Nash equilibrium.

Example 6.6. Nash equilibria in a potential game. The coordination game F^{C3} from Example 2.6 is generated by a common interest game, and so is a potential game. Its potential function is the convex function $f^{\text{C3}}(x) = \frac{1}{2}((x_1)^2 + 2(x_2)^2 + 3(x_3)^2)$. Figure 7 presents a graph and a contour plot of this function. The three pure states locally maximize potential, and so are Nash equilibria. There are four additional states that do not maximize potential, but that satisfy the Kuhn-Tucker first-order conditions. Geometrically, this means that if we linearize the potential function at these states, the linearized functions admit no feasible direction of increase from these states. By Theorem 6.5, these points too are Nash equilibria, although Theorem 6.8 will show they are not locally stable under typical evolutionary dynamics. ♦

The potential function in the previous example was convex. If instead a game's potential function is concave—as is true, for instance, in congestion games with increasing cost functions—then only the global maximizers of potential are Nash equilibria.




Figure 7: Graph and contour plot of the potential function $f^{\text{C}^3}(x) = \frac{1}{2}((x_1)^2 + 2(x_2)^2 + 3(x_3)^2)$, with Nash equilibria of game F^{C^3} marked.

6.5 Global Convergence and Local Stability

Lemma 6.1 tells us that evolutionary dynamics satisfying positive correlation (PC) will ascend the potential function whenever they are not at rest. In dynamical systems terminology, the potential function is a *strict (global) Lyapunov function* for all such dynamics.

Stating the consequences of this fact for the global behavior of the dynamics requires a formal definition. Let $\dot{x} = V(x)$ be a Lipschitz continuous differential equation that is forward invariant on X . For each initial condition $\xi \in X$, we define the ω -limit set $\omega(\xi)$ to be the set of all points that the solution trajectory $\{x_t\}_{t \geq 0}$ starting from $x_0 = \xi$ approaches arbitrarily closely infinitely often:

$$\omega(\xi) = \left\{ y \in X : \text{there exists } \{t_k\}_{k=1}^{\infty} \text{ with } \lim_{k \rightarrow \infty} t_k = \infty \text{ such that } \lim_{k \rightarrow \infty} x_{t_k} = y \right\}.$$

In general, ω -limit points include rest points, periodic orbits, and chaotic attractors. However, standard results from dynamical systems theory show that the existence of a strict global Lyapunov function rules out the latter possibilities. We therefore have

Theorem 6.7. *Let F be a potential game, and let V^F be an Lipschitz continuous evolutionary dynamic that satisfies positive correlation (PC). Then all ω -limit points of V^F are rest points. If in*

addition V^F satisfies Nash stationarity (NS), then all ω -limit points of V^F are Nash equilibria.

Lemma 6.1 also suggests that only local maximizers of f should be locally stable. To present such a result, we introduce some formal definitions of local stability for the differential equation $\dot{x} = V(x)$ above. These definitions are provided for a single state $y \in X$; replacing y with a closed set $Y \subset X$ yields the definitions for closed sets of states.

State y is *Lyapunov stable* if for every neighborhood O of y , there exists a neighborhood O' of y such that every solution $\{x_t\}_{t \geq 0}$ that starts in O' is contained in O : that is, $x_0 \in O'$ implies that $x_t \in O$ for all $t \geq 0$. State y is *attracting* if there is a neighborhood O of y such that every solution that starts in O converges to y ; if we can choose $O = X$, y is *globally attracting*. Finally, state y is *asymptotically stable* if it is Lyapunov stable and attracting, and it is *globally asymptotically stable* if it is Lyapunov stable and globally attracting.

With these definitions, we have

Theorem 6.8. *Let F be a potential game with potential function f , and let V^F be an evolutionary dynamic that satisfies positive correlation (PC) and Nash stationarity (NS). Then state $x \in X$ is asymptotically stable under V^F if and only if it is a local maximizer of f and is isolated in $NE(F)$.*

Versions of Theorem 6.8 remain true if Nash stationarity is not assumed, or if we consider local maximizer sets instead of points.²⁸

The results above are stated for Lipschitz continuous dynamics that satisfy positive correlation (PC). Analogous results hold for the best response dynamic, and also for perturbed best response dynamics. For the analysis of the latter case, one introduces the perturbed potential function $\tilde{f}(x) = f(x) - v(x)$. Since the gradient of this function is the virtual payoff function $\tilde{F}(x) = F(x) - \nabla v(x)$, condition (24) implies that perturbed best response dynamics ascend this function and converge to perturbed equilibria.

6.6 Local Stability of Strict Equilibria

The existence of a potential function allows one to establish convergence and stability results for all dynamics satisfying positive correlation (PC). Similar reasoning can be used to establish local stability of any *strict equilibrium*: that is, any pure state e_k such that $F_k(e_k) > F_j(e_k)$ for all $j \neq k$.

This result requires a slightly stronger restriction on the dynamics. We say that a dynamic V^F for game F satisfies *strong positive correlation* in $Y \subseteq X$ if

There exists a $c > 0$ such that for all $x \in Y$,

²⁸See Sandholm (2001b).

$$V^F(x) \neq \mathbf{0} \text{ implies that } \text{Corr}(V^F(x), F(x)) = \frac{V^F(x)' \Phi F(x)}{|V^F(x)| |\Phi F(x)|} \geq c.$$

That is, the correlation between strategies' growth rates and payoffs, or equivalently, the cosine of the angle between the growth rate and projected payoff vectors, must be bounded away from zero on Y . This condition is satisfied by imitative dynamics, excess payoff dynamics, and pairwise comparison dynamics in a neighborhood of any strict equilibrium.

Sandholm (2014) shows that the function

$$L(x) = (e_k - x)' F(e_k),$$

representing the payoff deficit of mixed strategy x at the strict equilibrium e_k , is a *strict local Lyapunov function* for state e_k : its value decreases along solutions of V^F in a neighborhood of e_k , and it is nonnegative in this neighborhood, equalling zero only when $x = e_k$. It thus follows from standard results that if e_k is an isolated rest point of V^F , then it is asymptotically stable under V^F .

The next section introduces a more general criterion for local stability. This criterion is defined for states throughout the simplex, but it requires more structure on the dynamics for stability to be assured.

7. Evolutionarily Stable States and Contractive Games

7.1 Evolutionarily Stable States

The birth of evolutionary game theory can be dated to the definition of an *evolutionarily stable strategy* (ESS) by Maynard Smith and Price (1973). Their model is one of monomorphic populations, whose members all choose the same mixed strategy in a symmetric normal form game. The notion of an evolutionarily stable strategy is meant to capture the capacity of a monomorphic population to resist invasion by a monomorphic mutant group whose members play some alternative mixed strategy.

This framework is quite different from the polymorphic, pure-strategist model we consider here. Nevertheless, if we reinterpret Maynard Smith and Price's (1973) conditions as constraints on population states—and call a point that satisfies these conditions an *evolutionarily stable state*—we obtain a sufficient condition for local stability under a variety

of evolutionary dynamics.^{29,30}

There are three equivalent ways of defining evolutionarily stable states. The simplest one, and the most useful for our purposes, was introduced by Hofbauer et al. (1979).³¹ It defines an ESS to be an (*infinitesimal*) *local invader*:

$$(34) \quad \text{There is a neighborhood } O \text{ of } x \text{ such that } (y - x)'F(y) < 0 \text{ for all } y \in O \setminus \{x\}.$$

To interpret this condition, fix the candidate state x , and consider some nearby state y . Condition (34) says that if the current population state is y , and an infinitesimal group of agents with strategy distribution x joins the population, then the average payoff of the agents in this group, $x'F(y) = \sum_{j \in S} x_j F_j(y)$, exceeds the average payoff in the population as a whole, $y'F(y)$. Of course, forming predictions based on group average payoffs runs counter to the individualistic approach that defines noncooperative game theory. Thus in the present context, the definition of ESS is not of interest directly, but only instrumentally, as a sufficient condition for stability under evolutionary dynamics.

The second definition, introduced by Taylor and Jonker (1978) and Bomze (1991), defines an ESS as a state that possesses *uniform invasion barrier*:

$$(35) \quad \text{There is an } \bar{\varepsilon} > 0 \text{ such that } (y - x)'F(\varepsilon y + (1 - \varepsilon)x) < 0 \text{ for all } y \in X \setminus \{x\} \text{ and } \varepsilon \in (0, \bar{\varepsilon}).$$

In contrast to condition (34), definition (35) looks at invasions by groups with positive mass, and compares the average payoffs of the incumbent and invading groups in the post-entry population.

The third definition, the original one of Maynard Smith and Price (1973), shows what restrictions ESS adds to Nash equilibrium.

$$(36a) \quad x \text{ is a Nash equilibrium: } (y - x)'F(x) \leq 0 \text{ for all } y \in X.$$

$$(36b) \quad \text{There is a neighborhood } O \text{ of } x \text{ such that for all } y \in O \setminus \{x\}, \\ (y - x)'F(x) = 0 \text{ implies that } (y - x)'F(y) < 0.$$

The stability condition (36b) says that if a state y near x is an alternative best response to x ,

²⁹The distinction between evolutionarily stable strategies and evolutionarily stable states is emphasized by Thomas (1984). General references on ESS theory include the survey of Hines (1987) and the monographs of Bomze and Pötscher (1989) and Cressman (1992).

³⁰How one ought to extend the ESS definition to multipopulation settings depends on which of the interpretations above one has in mind. For the mixed-strategist environment, the appropriate extension is *Cressman ESS* (Cressman (1992, 2006), Cressman et al. (2001)), while for population games, it is *Taylor ESS* (Taylor (1979)). See Sandholm (2010c, Chapter 8) for further discussion.

³¹See also Pohley and Thomas (1983) and Thomas (1985).

then an incumbent population with strategy distribution y obtains a lower average payoff against itself than an infinitesimal group of invaders with strategy distribution x obtains against the incumbents.

The following theorem confirms the equivalence of these definitions.

Theorem 7.1. *The following are equivalent:*

- (i) x satisfies condition (34).
- (ii) x satisfies condition (35).
- (iii) x satisfies conditions (36a) and (36b).

A state that satisfies these conditions is called an evolutionarily stable state (ESS).

For certain local stability results, we need a slightly stronger condition than ESS. We call state x a *regular ESS* (Taylor and Jonker (1978)) if

$$(37a) \quad x \text{ is a } \textit{quasistrict equilibrium}: F_i(x) = \bar{F}(x) > F_j(x) \text{ whenever } x_i > 0 \text{ and } x_j = 0.$$

$$(37b) \quad z' DF(x) z < 0 \text{ for all } z \in TX \setminus \{0\} \text{ such that } z_i = 0 \text{ whenever } x_i = 0.$$

The quasistrictness condition (37a) strengthens the Nash equilibrium condition (36a) by requiring unused strategies to be suboptimal. The first-order stability condition (37b) strengthens the stability condition (36b) by requiring the strict inequality to hold even at the level of a linear approximation. If $x = e_k$ is a pure state, then condition (37a) requires x to be a strict equilibrium, and condition (37b) is vacuous. At the other extreme, if x is in the interior of the simplex, then condition (37a) is equivalent to Nash equilibrium (36a), and condition (37b) requires that the derivative matrix $DF(x)$ be negative definite with respect to the tangent space TX .

This last requirement is the motivation for our next class of games.

7.2 Contractive Games

The population game $F: X \rightarrow \mathbb{R}^n$ is a (*weakly*) *contractive game*³² if

$$(38) \quad (y - x)'(F(y) - F(x)) \leq 0 \text{ for all } x, y \in X.$$

If the inequality in condition (38) holds strictly whenever $x \neq y$, F is *strictly contractive*, while if this inequality always binds, F is *conservative* (or *null contractive*).

³²Also known as *stable games* (Hofbauer and Sandholm (2009)) or *negative semidefinite games* (Hopkins (1999b)). In the convex analysis literature, an F satisfying condition (38) (sometimes with the inequality reversed) is called a *monotone operator*—see Hiriart-Urruty and Lemaréchal (2001).

For a first intuition, notice that if F is a potential game, so that $F \equiv \nabla f$, then (38) says that the potential function f is concave, ensuring the convexity and global stability of the set of Nash equilibria.

For a more general intuition, consider again the projection dynamic (12), defined on $\text{int}(X)$ by $\dot{x} = \Phi F(x)$. Solutions to this dynamic “follow the payoff vectors” to the greatest extent possible. If we run this dynamic from two initial states x_0 and y_0 , the squared distance between the states changes according to

$$(39) \quad \frac{d}{dt} |y_t - x_t|^2 = 2(y_t - x_t)'(\dot{y}_t - \dot{x}_t) = 2(y_t - x_t)'(F(y_t) - F(x_t)).$$

Thus, in a contractive game, following the payoff vectors brings states (weakly) closer together.³³

The connection between these games and the notion of ESS follows from the characterization for the differentiable case. By the fundamental theorem of calculus, if F is continuously differentiable, it is contractive if and only if it satisfies *self-defeating externalities*:

$$(40) \quad DF(x) \text{ is negative semidefinite with respect to } TX \text{ for all } x \in X.$$

This condition, which can be rewritten as

$$\sum_{i \in S} z_i \frac{\partial F_i}{\partial z}(x) \leq 0 \text{ for all } z \in TX \text{ and } x \in X,$$

provides an economic interpretation of contractive games. The vector $z \in TX$ represents the aggregate effect of revisions by a group of agents on the population state. Condition (40) requires that improvements in the payoffs of strategies to which revising agents are switching are always exceeded by the improvements in the payoffs of strategies which revising agents are abandoning. For instance, when $z = e_j - e_i$, this condition becomes

$$\frac{\partial F_j}{\partial(e_j - e_i)}(x) \leq \frac{\partial F_i}{\partial(e_j - e_i)}(x).$$

According to this inequality, any gains that the switches create for the newly chosen strategy j are dominated by gains for the abandoned strategy i .

³³This remains true on the boundary of the simplex if the dynamic is defined there via closest-point projection—see Lahkar and Sandholm (2008).

7.3 Examples

Example 7.2. Matching in symmetric zero-sum games. A symmetric two-player normal form game A is *symmetric zero-sum* if A is skew-symmetric: $A_{ji} = -A_{ij}$ for all $i, j \in S$. If agents are matched to play this game, then the resulting population game $F(x) = Ax$ satisfies $z'DF(x)z = z'Az = 0$ for all vectors $z \in \mathbb{R}^n$, and so is conservative. ♦

Example 7.3. Matching in Rock-Paper-Scissors. Let $F(x) = Ax$ with

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

Here $w > 0$ and $l > 0$ represent the benefit from a win and the cost of a loss, respectively. When $w = l$, we refer to A as (*standard*) *RPS*; when $w > l$, we refer to A as *good RPS*, and when $w < l$, we refer to A as *bad RPS*. In all cases, the unique Nash equilibrium of $F(x) = Ax$ is $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

To evaluate condition (40) on the derivative matrix $DF(x) = A$, notice that the symmetric matrix $A + A'$ has eigenvalue $l - w$ with geometric multiplicity 2, corresponding to the eigenspace TX . It follows that F is conservative in standard RPS, and strictly contractive in good RPS. In the latter case, Nash equilibrium x^* is a regular ESS. ♦

Example 7.4. (Perturbed) concave potential games. A potential game F with a strictly concave potential function satisfies condition (38) with a strict inequality, and so is a strictly contractive game. If we slightly perturb F , then the new game is quite unlikely to be a potential game. But (38) will continue to hold, so the new game is still a strictly contractive game. ♦

7.4 Equilibrium in Contractive Games

The contraction condition (39) suggests that obedience of incentives will push the population toward some “central” equilibrium state. To work toward a confirmation of this idea, we say that x is a *globally neutrally stable state* of F , denoted $x \in \text{GNSS}(F)$, if

$$(41) \quad (y - x)'F(y) \leq 0 \text{ for all } y \in X.$$

GNSS is the global analogue of Maynard Smith’s (1982) notion of a *neutrally stable strategy*, which is obtained by replacing the strict inequality in definition (34) of ESS with a weak one. Rewriting inequality (41) as $(x - y)'F(y) \geq 0$, we obtain a simple geometric interpre-




Figure 8: The GNSS of standard RPS.

tation: proceeding from any state y an infinitesimal distance in the direction specified by payoff vector $F(y)$ moves the state (weakly) closer to the GNSS x .

Example 7.5. GNSS in standard RPS. Starting from a selection of states $y \in X$, Figure 8 presents projected payoff vectors $\Phi F(y)$ from standard RPS, along with vectors $x^* - y$ leading to the Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. All such pairs of vectors are orthogonal, so x^* is a GNSS. ♦

The set $GNSS(F)$ is an intersection of half spaces, and so is convex. A simple geometric argument shows that it only contains Nash equilibria.

Proposition 7.6. $GNSS(F) \subseteq NE(F)$.

To prove this inclusion, let $x \in GNSS(F)$ and let $y \neq x$. Define $x_\varepsilon = \varepsilon y + (1 - \varepsilon)x$. Since x is a GNSS, $(x - x_\varepsilon)'F(x_\varepsilon) \geq 0$ for all $\varepsilon \in (0, 1]$: that is, motion from x_ε in direction $F(x_\varepsilon)$ is weakly toward x (see Figure 9). Taking ε to zero yields $(y - x)'F(x) \leq 0$ for all $y \in X$, which is definition (2) of Nash equilibrium.

While every GNSS is a Nash equilibrium, not every game has a GNSS. But the next result shows that in contractive games, a GNSS must exist, because every Nash equilibrium is a GNSS.

Theorem 7.7. *If F is a contractive game, then $NE(F) \subseteq GNSS(F)$. Thus $NE(F) = GNSS(F)$, and so is a convex set.*




Figure 9: Why every GNSS is a Nash equilibrium.

To prove the theorem, add the definition (38) of contractive games to definition (2) of Nash equilibrium; the sum is the definition (41) of GNSS.

If inequality (41) is strict whenever $y \neq x$, we call x a *globally evolutionarily stable state*. This condition is the global analogue of ESS condition (34). A variation on the arguments above shows that in strictly contractive games, a Nash equilibrium is a GESS, and that the Nash equilibrium is unique.

In light of Theorem 7.7 and the foregoing analysis, it is natural to ask whether one can prove existence of Nash equilibrium in contractive games using tools from convex analysis, avoiding an appeal to fixed point theory. Indeed, adapting an argument of Minty (1967), one can provide a direct proof of existence of equilibrium using the minmax theorem and a compactness argument.³⁴

7.5 Global Convergence and Local Stability

The discussions so far suggest that contractive games should admit global convergence results under evolutionary dynamics, and that ESS should serve as a sufficient condition for local stability. This section verifies these intuitions.

7.5.1 Imitative dynamics

The original global convergence result for strictly contractive games, due to Hofbauer et al. (1979) and Zeeman (1980), is for the replicator dynamic (8). The analysis proceeds as

³⁴See Sandholm (2010c, Sec. 3.3.5).

follows: If F is a strictly contractive game, it admits a GESS x^* . We can therefore consider the function

$$H_{x^*}(x) = \sum_{i \in \text{supp}(x^*)} x_i^* \log \frac{x_i^*}{x_i},$$

which is defined on the set of states $X_{x^*} = \{x \in X: \text{supp}(x^*) \subseteq \text{supp}(x)\}$ whose supports contain the support of x^* . Jensen's inequality implies that H_{x^*} is nonnegative, and equals 0 precisely when $x = x^*$. Moreover, the time derivative of this function under the replicator dynamic (8) is

$$\dot{H}_{x^*}(x) = \nabla H_{x^*}(x)' \dot{x} = - \sum_{i \in \text{supp}(x^*)} \frac{x_i^*}{x_i} \cdot x_i \hat{F}_i(x) = - \sum_{i \in S} x_i^* \hat{F}_i(x) = -(x^* - x)' F(x).$$

Thus, since x^* is a GESS, the value of H_{x^*} falls over time. Standard results on Lyapunov functions then imply that x^* is asymptotically stable, attracting solutions from every initial condition in X_{x^*} . Whether it is possible to prove global convergence results for contractive games under other imitative dynamics is an open question.

The only property of the game F used in this calculation was the fact that x^* is a GESS. Therefore, virtually the same argument implies that any ESS is locally asymptotically stable under the replicator dynamic.

One can obtain more general conclusions about local stability using linearization. Taylor and Jonker (1978) show that when the derivative matrix $DV(x^*)$ of the replicator dynamic is evaluated at a regular ESS x^* , its eigenvalues corresponding to directions in TX have negative real part, implying that solutions starting near the ESS converge to it at an exponential rate. Remarkably, Cressman (1997b) shows that in hyperbolic cases,³⁵ the linearizations of all imitative dynamics at a given restricted equilibrium are positive multiples of one another, and consequently have the same stability properties. Among other things, this implies the local stability of any regular ESS under typical imitative dynamics.

7.5.2 Target and pairwise comparison dynamics: Global convergence in contractive games

We saw earlier that in potential games, the potential function serves as a Lyapunov function for any evolutionary dynamic satisfying positive correlation (PC). Contractive

³⁵A rest point x^* is *hyperbolic* if the eigenvalues of the derivative matrix $DV(x^*)$ have nonzero real part. The Hartman-Grobman theorem (see Robinson (1995)) shows that near a hyperbolic rest point, the flow of a nonlinear dynamic is topologically conjugate to the flow of its linearization.

games, by contrast, do not come equipped with candidate Lyapunov functions. In order to construct Lyapunov functions for dynamics in these games, we must look to the dynamics themselves to obtain the necessary structure, following Smith (1984), Hofbauer (1995b, 2000), and Hofbauer and Sandholm (2007, 2009).

The best response dynamic, perturbed best response dynamics, and excess payoff dynamics all can be derived from target protocols $\tau: \mathbb{R}^n \rightarrow \mathbb{R}_+^n$ that condition on the excess payoff vector $\hat{F}(x)$. Since target protocols are maps from \mathbb{R}^n into itself, it makes sense to ask whether a target protocol is integrable: that is, whether there exists a *revision potential* $\gamma: \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying $\tau(\hat{\pi}) = \nabla \gamma(\hat{\pi})$ for all $\hat{\pi} \in \mathbb{R}^n$. In this case, the revision potential provides a building block for constructing a Lyapunov function to analyze the dynamic at hand.³⁶

For example, it is easy to verify that the revision protocol (27) for the BNN dynamic is integrable:

$$\tau_i(\hat{\pi}) = [\hat{\pi}_i]_+ \implies \gamma(\hat{\pi}) = \frac{1}{2} \sum_{i \in S} [\hat{\pi}_i]_+^2.$$

To make use of the revision potential γ , we introduce the Lyapunov function $\Gamma: X \rightarrow \mathbb{R}$, defined by

$$(42) \quad \Gamma(x) = \gamma(\hat{F}(x)).$$

This function is nonnegative, and by Proposition 5.6, it equals zero precisely at the Nash equilibria of F . A computation shows that along solutions to the BNN dynamic, the value of Γ obeys

$$\dot{\Gamma}(x) = \dot{x}' DF(x) \dot{x} - (\tau(\hat{F}(x)))' \mathbf{1} (F(x)' \dot{x}).$$

If x is not a rest point, then the contribution of the second term here is negative by positive correlation (PC), and if F is contractive, the contribution of the first term is nonpositive by condition (40). Thus Γ is a (decreasing) strict global Lyapunov function, and so standard results imply that the set $NE(F)$ is globally asymptotically stable under the BNN dynamic. The specific functional form of τ is not too important here: as long as the protocol is acute (25b) and integrable, all solutions to the corresponding excess payoff dynamic converge to the set of Nash equilibria.

Similar results hold for the best response dynamic and all perturbed best response

³⁶For a game-theoretic interpretation of this integrability condition, see Sandholm (2013).

Dynamic	Lyapunov function for contractive games
replicator	$H_{x^*}(x) = \sum_{i \in \text{supp}(x^*)} x_i^* \log \frac{x_i^*}{x_i}$
best response	$G(x) = \max_{i \in S} \hat{F}_i(x)$
logit	$\tilde{G}(x) = \max_{y \in \text{int}(X)} \left(y' \hat{F}(x) - \eta \sum_{i \in S} y_i \log y_i \right) + \eta \sum_{i \in S} x_i \log x_i$
BNN	$\Gamma(x) = \frac{1}{2} \sum_{i \in S} [\hat{F}_i(x)]_+^2$
Smith	$\Psi(x) = \frac{1}{2} \sum_{i \in S} \sum_{j \in S} x_i [F_j(x) - F_i(x)]_+^2$

Table 4: Lyapunov functions for five basic deterministic dynamics in contractive games.

dynamics, as they too can be derived from integrable target protocols:³⁷

$$M(\hat{\pi}) = \underset{y \in X}{\operatorname{argmax}} y' \hat{\pi} \quad \implies \quad \mu(\hat{\pi}) = \max_{y \in X} y' \hat{\pi},$$

$$\tilde{M}(\hat{\pi}) = \underset{y \in \text{int}(X)}{\operatorname{argmax}} y' \hat{\pi} - v(y) \quad \implies \quad \tilde{\mu}(\hat{\pi}) = \max_{y \in \text{int}(X)} y' \hat{\pi} - v(y).$$

Therefore, Lyapunov functions analogous to (42) can be constructed in these cases as well.

Pairwise comparison dynamics cannot be derived from target protocols. Even so, Smith (1984) constructs a more complicated Lyapunov function for his dynamics for contractive games, and Hofbauer and Sandholm (2009) show that a similar construction can be used for a broader class of pairwise comparison dynamics.³⁸

The Lyapunov functions for the five basic dynamics in contractive games are summarized in Table 4.

7.5.3 Target and pairwise comparison dynamics: Local stability of regular ESS

These Lyapunov functions for target and pairwise comparison dynamics can also be used to establish local stability of regular ESS. In the case of an interior ESS, the contractive game condition (40) holds in a neighborhood of the ESS, so local stability follows immediately from the previous analysis.

³⁷We can replace actual payoffs with excess payoffs as the arguments of M and \tilde{M} , since maximizers of actual payoffs are also maximizers of excess payoffs. Also, since M is a correspondence, the function μ is not a potential function in the usual calculus sense, but in the sense of convex analysis: we have $\partial\mu \equiv M$, where the correspondence $\partial\mu: \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is the subdifferential of μ .

³⁸The condition that is needed, *impartiality*, requires that the function of the payoff difference $\pi_j - \pi_i$ that describes the conditional switch rate from i to j does not depend on an agent's current strategy i .

The analysis is trickier for a regular ESS on the boundary of the simplex. The relevant contraction condition, (37b), only applies in directions tangent to the face on which the ESS lies, reflecting switches among strategies currently in use. For its part, quasistrictness (37a) requires that the strategies in use are strictly better than those that are not. Intuitively, one expects that quasistrictness should drive the state toward the face of the simplex, and that once the state is close enough to the face, the contraction condition should be enough to ensure stability. To formalize this argument, Sandholm (2010a) augments the Lyapunov functions introduced above by a term that penalizes mass on strategies outside the support of the ESS. The values of these augmented functions decrease over time, establishing local stability of regular ESS in each case.

8. Iterative Solution Concepts, Supermodular Games, and Equilibrium Selection

This section considers connections between evolutionary game dynamics and traditional iterative solution concepts. In some cases iteration is used directly to study the dynamics, but in the case of supermodular games a rather different tool, the theory of cooperative differential equations, is employed. This section also introduces a new class of dynamics, sampling best response dynamics, which select unique equilibria in certain games with multiple strict equilibria.

8.1 Iterated Strict Dominance and Never-a-Best-Response

To begin, we show that traditional iterative solution concepts are respected by imitative dynamics and the best response dynamic. Although these dynamics are the most commonly studied, the properties described here do not extend to other dynamics, as Section 9.2 will show.

For imitative dynamics, we consider strict dominance, following Akin (1980) and Nachbar (1990). Strategy i is *strictly dominated* by strategy j if $F_j(x) > F_i(x)$ for all $x \in X$.³⁹

Theorem 8.1. *Let $\{x_t\}$ be an interior solution trajectory of an imitative dynamic in game F . If strategy $i \in S$ is strictly dominated in F , then $\lim_{t \rightarrow \infty} (x_t)_i = 0$.*

To prove this result, suppose that strategy i is strictly dominated by strategy j , and write $r = x_i/x_j$. If we express the imitative dynamic in the percentage growth rate form

³⁹For results on strict dominance by mixed strategies, see Samuelson and Zhang (1992), Hofbauer and Weibull (1996), and Viossat (2011).

$\dot{x}_i = x_i G_i(x)$ as in (10), then monotonicity condition (11) and a compactness argument imply that $G_j(x) - G_i(x) > c > 0$ for all $x \in X$. Applying the quotient rule yields

$$\frac{d}{dt} r = \frac{d}{dt} \frac{x_i}{x_j} = \frac{\dot{x}_i x_j - \dot{x}_j x_i}{(x_j)^2} = \frac{x_i G_i(x) x_j - x_j G_j(x) x_i}{(x_j)^2} = r(G_i(x) - G_j(x)) \leq -cr,$$

and so x_i vanishes.

Of course, once we remove a dominated strategy from a game, other strategies may become dominated in turn. A simple continuity argument shows that the conclusion of Theorem 8.1 extends to strategies eliminated via iterated dominance.

It is easy to see that under the best response dynamic, not only does any strictly dominated strategy vanish; so too does any strategy i that is never a best response, in the sense that for every $x \in X$ there is a $j \in S$ such that $F_j(x) > F_i(x)$.⁴⁰ In fact, such a strategy must vanish at an exponential rate: $(x_t)_i = (x_0)_i e^{-t}$ for any $x_0 \in X$, so that elimination occurs not just from interior initial conditions, but from all initial conditions. And by an argument based on continuity of payoffs, any strategy eliminated by iterative removal of never-a-best-response strategies eventually vanishes at an exponential rate.

It follows that at each stage of this elimination, the face of the simplex corresponding to the surviving strategies is asymptotically stable under the best response dynamic. The next example shows that the corresponding claim for imitative dynamics is false.

Example 8.2. Dominated strategies and local instability under imitative dynamics. Consider the game $F(x) = Ax$ with

$$A = \begin{pmatrix} 1 & 1 & 1 \\ 2 & 2 & 2 \\ 0 & 4 & 0 \end{pmatrix}.$$

Since strategy 1 is strictly dominated in this game, Theorem 8.1 implies that under any imitative dynamic, x_1 vanishes along any interior solution trajectory. However, if strategy 2 is omitted, strategy 1 strictly dominates strategy 3, so on the interior of the $e_1 e_3$ face, solutions of imitative dynamics converge to e_1 . We therefore conclude that the $e_2 e_3$ face is neither Lyapunov stable nor attracting. ♦

However, if a game is dominance solvable, with only strategy k surviving, then state e_k is a strict equilibrium, and hence asymptotically stable—see Section 6.6.

⁴⁰Any strategy that is strictly dominated by a pure strategy is never a best response, but because we only consider pure strategies here, the converse statement is false, even when $F(x) = Ax$ is linear.

8.2 Supermodular Games and Perturbed Best Response Dynamics

Supermodular games, introduced in finite player contexts by Topkis (1979), Vives (1990), and Milgrom and Roberts (1990), are defined by the property that higher choices by one's opponents make one's own higher strategies look relatively more desirable. This implies the monotonicity of the best response correspondence, which in turn implies the existence of minimal and maximal Nash equilibria. We consider the large population version of these games. Following Hofbauer and Sandholm (2002, 2007), we use techniques from the theory of cooperative dynamical systems to establish almost global convergence results for perturbed best response dynamics.

To define supermodular games, we introduce the stochastic dominance matrix $\Sigma \in \mathbb{R}^{(n-1) \times n}$, defined by $\Sigma_{ij} = 1_{j>i}$. Then $(\Sigma x)_i = \sum_{j=i+1}^n x_j$ equals the total mass on actions greater than i at population state x , and $\Sigma y \geq \Sigma x$ if and only if y stochastically dominates x .

We call F a *supermodular game* if

$$(43) \quad \Sigma y \geq \Sigma x \text{ implies that } F_{i+1}(y) - F_i(y) \geq F_{i+1}(x) - F_i(x) \text{ for all } i < n.$$

That is, if y stochastically dominates x , then for any action $i < n$, the payoff advantage of $i + 1$ over i is greater at y than at x . If F is continuously differentiable, then it is supermodular if and only if

$$(44) \quad \frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) \geq 0 \text{ for all } i < n, j < n, \text{ and } x \in X.$$

In words, if some agents switch from strategy j to strategy $j + 1$, the performance of strategy $i + 1$ improves relative to that of strategy i . Conditions (43) and (44) are both called *strategic complementarity*.

Example 8.3. Search with positive externalities. Example 2.4 introduced a population game model of search: $F_i(x) = m_i b(a(x)) - c_i$, where $a(x) = \sum_{k=1}^n kx_k$ represents aggregate search effort, the increasing function b represents the benefits of search as a function of aggregate effort, the increasing function m is the benefit multiplier, and the arbitrary function c captures search costs. Since

$$\frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) = (m_{i+1} - m_i) b'(a(x)) \geq 0,$$

F is a supermodular game. ♦

It is intuitively clear that supermodular games must have increasing best response correspondences: when opponents choose higher strategies, an agent's own higher strategies look relatively better, so his best strategies must be higher as well. The next result makes this observation precise, and presents its implications for the structure of the Nash equilibrium set.

Let $\underline{B}: X \rightarrow X$ and $\bar{B}: X \rightarrow X$, defined by $\underline{B}(x) = \min B(x)$ and $\bar{B}(x) = \max B(x)$, be the minimal and maximal best response functions, where the minimum and maximum are defined with respect to the stochastic dominance order. Evidently, these functions always evaluate to pure states. For states $\underline{x}, \bar{x} \in X$ satisfying $\Sigma \underline{x} \leq \Sigma \bar{x}$, we define the interval $[\underline{x}, \bar{x}]$ as the set of states lying between \underline{x} and \bar{x} in the stochastic dominance order: $[\underline{x}, \bar{x}] = \{x \in X : \Sigma \underline{x} \leq \Sigma x \leq \Sigma \bar{x}\}$.

Theorem 8.4. *Suppose F is a supermodular game. Then*

- (i) \underline{B} and \bar{B} are increasing in the stochastic dominance order: if $\Sigma x \leq \Sigma y$, then $\Sigma \underline{B}(x) \leq \Sigma \underline{B}(y)$ and $\Sigma \bar{B}(x) \leq \Sigma \bar{B}(y)$.
- (ii) The sequences of iterates $\{\underline{B}^k(e_1)\}_{k \geq 0}$ and $\{\bar{B}^k(e_n)\}_{k \geq 0}$ are monotone sequences of pure states, and so converge within n steps to their limits, \underline{x}^* and \bar{x}^* .
- (iii) $\underline{x}^* = \underline{B}(\underline{x}^*)$ and $\bar{x}^* = \bar{B}(\bar{x}^*)$, so \underline{x}^* and \bar{x}^* are pure Nash equilibria of F .
- (iv) $NE(F) \subseteq [\underline{x}^*, \bar{x}^*]$. Thus if $\underline{x}^* = \bar{x}^*$, then this state is the unique Nash equilibrium of F .

Given the monotonicity of the best response correspondence in supermodular games, it is natural to look for convergence results under the best response dynamic. It follows from the results in the previous section that this dynamic must converge to the interval $[\underline{x}^*, \bar{x}^*]$ whose endpoints are the minimal and maximal Nash equilibria. But to prove convergence to Nash equilibrium, one must impose additional structure on the game.⁴¹

General convergence results can be proved using the theory of (*strongly*) cooperative differential equations. These are continuously differentiable differential equations $\dot{\chi} = \mathcal{V}(\chi)$ whose derivative matrices have positive off-diagonal elements, so that increases in any component of χ increase the growth rates of other components.

Since cooperative differential equations must be smooth, this theory cannot be applied to the best response dynamic. However, Hofbauer and Sandholm (2002, 2007) show that the needed monotonicity is preserved by perturbed best response dynamics, although only those derived from stochastic perturbations.

Specifically, they consider the perturbed best response dynamic

⁴¹For instance, Berger (2007) proves that in two-population games generated by two-player normal form games that are supermodular and satisfy a diminishing returns condition, most solution trajectories of the best response dynamic converge to pure Nash equilibria.

$$(45) \quad \dot{x} = \tilde{M}^\varepsilon(F(x)) - x,$$

where the perturbed best response function \tilde{M}^ε is defined in terms of the stochastic perturbation ε , as in equation (21). Equation (45) is transformed into a cooperative differential equation using the change of variable defined by the stochastic dominance operator Σ . Let $\mathcal{X} = \Sigma X \subset \mathbb{R}^{n-1}$ denote the image of X under Σ , and let $\tilde{\Sigma}: \mathcal{X} \rightarrow X$ denote the (affine) inverse of the map Σ . Then the change of variable Σ converts (45) into the dynamic

$$(46) \quad \dot{\chi} = \Sigma \tilde{M}^\varepsilon(F(\tilde{\Sigma}\chi)) - \chi,$$

on \mathcal{X} .

Combining strategic complementarity (44) with the properties of the derivative matrix $D\tilde{M}^\varepsilon(\pi)$, Hofbauer and Sandholm (2002, 2007) show that equation (46) is strongly cooperative. Results from the theory of cooperative differential equations (Hirsch (1988), Smith (1995)) then yield the following result.

Theorem 8.5. *Let F be a C^1 strictly supermodular game, and let $\dot{x} = V^{F,\varepsilon}(x)$ be a stochastically perturbed best response dynamic for F . Then*

- (i) *States $\underline{x}^* \equiv \omega(\underline{x})$ and $\bar{x}^* \equiv \omega(\bar{x})$ exist and are the minimal and maximal elements of the set of perturbed equilibria. Moreover, $[\underline{x}^*, \bar{x}^*]$ contains all ω -limit points of $V^{F,\varepsilon}$ and is globally asymptotically stable.*
- (ii) *Solutions to $\dot{x} = V^{F,\varepsilon}(x)$ from an open, dense, full measure set of initial conditions in X converge to perturbed equilibria.*

8.3 Iterated p -Dominance and Equilibrium Selection

The evolutionary dynamics we have studied so far have been based on reactive protocols $\rho^F(x) = \rho(F(x), x)$, with switch rates depending on the game only through current payoffs. This formulation leads naturally to dynamics satisfying positive correlation (PC) and Nash stationarity (NS). While this approach has many advantages, it also imposes restrictions on what the dynamics can achieve. For instance, in coordination games, all pure equilibria are locally stable, implying that predictions of play must depend on initial conditions.

Following Oyama et al. (2012) (see also Sandholm (2001a)), we now argue that dynamics based on prospective revision protocols may have quite different properties, and in particular may lead to unique predictions in certain games with multiple strict equilibria.⁴²

⁴²Dynamics based on reactive protocols with a moderate amount of noise can have similar properties—

We consider an analog of the best response dynamic in which agents do not know the population state, and so base their decisions on information from samples. A revising agent obtains information by drawing a sample of size k from the population. The agent then plays a best response to the empirical distribution of strategies in his sample. This is a prospective revision protocol: agents are not reacting to current payoffs, but to the payoffs that would obtain if their sample were representative of behavior in the population at large.

Let $\mathbb{Z}_+^{n,k} = \{z \in \mathbb{Z}_+^n : \sum_{i \in S} z_i = k\}$ be the set of possible outcomes of samples of size k . The k -sampling best response protocol is the target protocol

$$(47a) \quad \sigma^F(x) = B^{F,k}(x), \text{ where}$$

$$(47b) \quad B^{F,k}(x) = \sum_{z \in \mathbb{Z}_+^{n,k}} \binom{k}{z_1 \dots z_n} (x^{z_1} \dots x^{z_n}) B^F\left(\frac{1}{k}z\right).$$

is the k -sampling best response correspondence.⁴³ The mean dynamic of this protocol is the k -sampling best response dynamic

$$(48) \quad \dot{x} \in B^{F,k}(x) - x.$$

Sampling best response dynamics agree with standard dynamics in one basic respect: strict equilibria of F are rest points of (48). However, while strict equilibria are locally stable under standard dynamics (Section 6.6), this need not be true under sampling best response dynamics.

For $p \in [0, 1]$, we call strategy $i \in S$ a p -dominant equilibrium of F (Morris et al. (1995)) if it is the unique optimal strategy at any state x satisfying $x_i \geq p$. Thus 1-dominant equilibria are strict equilibria, while 0-dominant equilibria correspond to strictly dominant strategies. The basic selection result for sampling best response dynamics is

Theorem 8.6. *Suppose that strategy i is $\frac{1}{k}$ -dominant in game F , and let V^F be the k -sampling best response dynamic. Then state e_i is asymptotically stable, and attracts solutions from all initial conditions with $x_i > 0$.*

To explain this result, it is sufficient to consider a coordination game with strategy set $S = \{0, 1\}$ and in which strategy 1 is $\frac{1}{k}$ -dominant but not $\frac{1}{k+1}$ -dominant. An example is the game $F(x) = Ax$ with

see Oyama et al. (2012) and Kreindler and Young (2013).

⁴³While in general $B^{F,k}$ is multivalued, it is single valued when every possible sample of size k generates a unique best response in F . For linear games $F(x) = Ax$, this is true for generic choices of the matrix A .

$$A = \begin{pmatrix} 1 & 0 \\ 0 & k - 1 + \varepsilon \end{pmatrix} \text{ and } \varepsilon \in (0, 1).$$

Since strategy 1 is $\frac{1}{k}$ -dominant and the sample size is k , an agent will only choose strategy 0 if all of the agents he samples choose strategy 0. Thus if $x_1 \in (0, 1)$,

$$\dot{x}_1 = B_1^{F,k}(x) - x_1 = \left(1 - (1 - x_1)^k\right) - x_1 = (1 - x_1) - (1 - x_1)^k.$$

Since this expression is positive whenever $x_1 \in (0, 1)$, state $x_1 = 1$ attracts all interior initial conditions.

The use of a fixed sample size is not essential: the same result holds so long as enough probability is placed on sample sizes no greater than k . Moreover, the selection result extends to iterated versions of p -dominance (Tercieux (2006), Oyama and Tercieux (2009)), as we illustrate through an example.

Example 8.7. Selection of iterated $\frac{1}{2}$ -dominant equilibrium. Consider the following 3×3 coordination game of Young (1993):

$$F(x) = Ax, \text{ where } A = \begin{pmatrix} 6 & 0 & 0 \\ 5 & 7 & 5 \\ 0 & 5 & 8 \end{pmatrix}.$$

The phase diagram of the best response dynamic for this game, presented in Figure 10(i), shows that each of the three strict equilibria of F has a nonnegligible basin of attraction.

Examining the best response regions of F as depicted in Figure 10(i), it is clear that no strategy in this game is $\frac{1}{2}$ -dominant. However, the set $\{2, 3\}$ is a $\frac{1}{2}$ -best response set: if at least half the population plays strategies from this set, then all pure best responses are elements of the set. Moreover, once strategy 1 is removed, strategy 3 is $\frac{1}{2}$ -dominant in the game that remains. As Figure 10(ii) shows, state e_3 is asymptotically stable under the 2-sampling best response dynamic in the original game, and attracts solutions from all states x with $x_3 > 0$. ♦

To establish asymptotic stability in Example 8.7, one uses a transitivity theorem for asymptotic stability due to Conley (1978). In the present context, this theorem says that if edge e_2e_3 is an asymptotically stable invariant set in X , and state e_3 is asymptotically stable with respect to the dynamic restricted to edge e_2e_3 , then state e_3 is asymptotically stable in X . This theorem may prove useful elsewhere for transferring the logic of iteration from the analysis of a game to the analysis of a corresponding evolutionary dynamic.

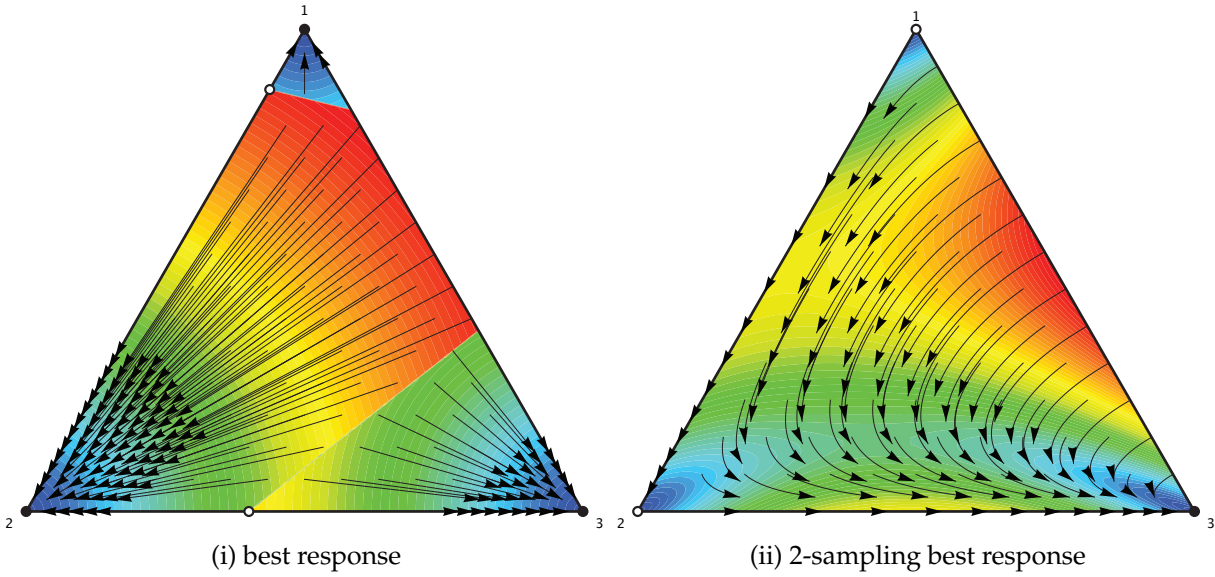


Figure 10: Phase diagrams for dynamics in Young's game.

9. Nonconvergence of Evolutionary Dynamics

The analyses in the previous sections have highlighted combinations of classes of games and dynamics for which convergence to equilibrium can be assured. Beyond these combinations, there are few general guarantees of convergence. Indeed, a key reason for studying evolutionary game dynamics is to understand when equilibrium analysis should be augmented or supplanted by an analysis of persistent disequilibrium behavior. In this section, we offer a selection of examples in which convergence to equilibrium fails, and present an implication of this possibility for the survival of strictly dominated strategies.

9.1 Examples

Example 9.1. Bad RPS. Example 7.3 introduced Rock-Paper-Scissors:

$$F(x) = Ax, \text{ where } A = \begin{pmatrix} 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{pmatrix}.$$

In *bad Rock-Paper-Scissors*, $l > w > 0$, so that the cost of a loss is higher than the benefit of a win. Figure 11 presents phase diagrams of the five basic evolutionary dynamics for bad RPS with $l = 2$ and $w = 1$. Under the replicator dynamic, interior solutions approach a heteroclinic cycle along the three edges of the simplex. Under the other four dynamics,

solutions approach interior limit cycles, with the precise position of the cycle depending on the dynamic in question.


Gaunersdorfer and Hofbauer (1995) provide a full analysis of the replicator and best response dynamics in Rock-Paper-Scissors games. They dub the triangular closed orbit of the best response dynamic in bad RPS a *Shapley polygon*, after Shapley (1964), who constructed the first example of cycling under the fictitious play process (Section 10.2). Benaïm et al. (2009) establish the existence of attracting Shapley polygons in higher dimensional examples with a similar cyclic structure.

Since the state space for bad RPS is two-dimensional, the existence of limit cycles for the remaining dynamics can be proved using the *Poincaré-Bendixson theorem*, which states that in planar systems, forward invariant regions containing no rest points must contain closed orbits; see Berger and Hofbauer (2006), Sandholm (2007a), and Hofbauer and Sandholm (2011). ♦


Example 9.2. Mismatching Pennies. *Mismatching Pennies* (Jordan (1993)) is a three-player normal form game in which each player has two strategies, Heads and Tails. Player p receives a payoff of 1 for choosing a different strategy than player $p + 1$ and a payoff of 0 otherwise, where players are indexed modulo 3. The unique Nash equilibrium of this game has each player play each of his strategies with equal probability.

Figure 12 presents phase diagrams of the replicator dynamic and the best response dynamic for the three-population game obtained by matching triples of agents to play *Mismatching Pennies*. Almost all solutions of the replicator dynamic (Figure 12(i)) converge to a six-sided heteroclinic cycle on the boundary of the state space; this cycle follows the best response cycle of the normal form game. Almost all solutions of the best response dynamic (Figure 12(ii)) converge to a Shapley polygon, here a six-sided closed orbit in the interior of the state space. The behaviors of the two dynamics are intimately related: Gaunersdorfer and Hofbauer (1995) show that the time averages of interior solutions of the replicator dynamic approach the Shapley polygon of the best response dynamic. This link is not accidental: quite general connections between the time-averaged replicator dynamic and the best response dynamic are established by Hofbauer et al. (2009b).


Hart and Mas-Colell (2003) show that the failures of convergence exhibited in Figure 12 are not exceptional. They consider three-population games based on normal form games whose unique Nash equilibrium is completely mixed. Under the weak assumption that each population's revision protocol does not condition on other populations' payoffs, they prove that any hyperbolic rest point of the resulting evolutionary dynamic must be unstable. It follows that if the corresponding dynamics satisfy Nash stationarity (NS), solutions from almost all initial conditions do not converge. ♦




(i) replicator




(ii) best response



(iii) logit(.1)




(iv) BNN




(v) Smith

Figure 11: Five basic deterministic dynamics in bad Rock-Paper-Scissors.



(i) replicator



(ii) best response

Figure 12: Evolutionary dynamics in Mismatching Pennies.




Figure 13: Projected payoffs in the hypnodisk game

Example 9.3. The hypnodisk game. *Hypnodisk* (Hofbauer and Sandholm (2011)) is a three-strategy population game with nonlinear payoffs. In a small circle centered at $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, the payoffs are those of the coordination game

$$F^C(x) = Cx = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix},$$

a potential game with convex potential function $f^C(x) = \frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3)^2)$. Outside a larger circle centered at x^* , the payoffs are those of the antcoordination game $F^{-C}(x) = -Cx$, a potential game with concave potential function $f^{-C}(x) = -\frac{1}{2}((x_1)^2 + (x_2)^2 + (x_3)^2)$. In between, payoffs are defined so that F is continuous, and so that x^* is the game's unique Nash equilibrium. Geometrically, this is accomplished by starting with the vector field for F^C , and then twisting the portion of the vector field outside 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 projected payoffs of the resulting game are illustrated in Figure 13.

Now consider evolution under an evolutionary dynamic satisfying positive correlation




Figure 14: A chaotic attractor under the replicator dynamic.

(PC) and Nash stationarity (NS). By the analysis in Section 6, solutions starting at states besides x^* in the inner disk must leave the inner disk. Similarly, solutions from states outside the outer disk must enter the outer disk. Since there are no Nash equilibria, and hence no rest points, in the annulus bounded by the circles, the Poincaré-Bendixson theorem implies that every solution other than the one at x^* must converge to a limit cycle in the annulus. ♦

Example 9.4. Chaotic dynamics. In population games with four or more strategies, and hence state spaces with three or more dimensions, solution trajectories of game dynamics can converge to complicated sets called *chaotic attractors*. Central to most definitions of chaos is *sensitive dependence on initial conditions*: solution trajectories starting from close together points on the attractor move apart at an exponential rate. Chaotic attractors can also be recognized in phase diagrams by their intricate appearance.

Following Arneodo et al. (1980) and Skyrms (1992), we consider evolution in the game

$$F(x) = Ax = \begin{pmatrix} 0 & -12 & 0 & 22 \\ 20 & 0 & 0 & -10 \\ -21 & -4 & 0 & 35 \\ 10 & -2 & 2 & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix},$$

whose lone interior Nash equilibrium is $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Figure 14 presents a solution to the replicator dynamic for this game from initial condition $x_0 = (.24, .26, .25, .25)$. This solution spirals clockwise about x^* . Near the rightmost point of each circuit, where the value of x_3 gets close to zero, solutions sometimes proceed along an “outside” path on which the value of x_3 surpasses .6. But they sometimes follow an “inside” path on which x_3 remains below .4, and at other times do something in between. Which of these alternatives occurs is difficult to predict from approximate information about the previous behavior of the system. While this game has a complicated payoff structure, in multipopulation contexts one can find chaotic evolutionary dynamics in simple games with three strategies per player.⁴⁴ ♦

9.2 Survival of Strictly Dominated Strategies

In Section 8.1, we argued that the best response dynamic and all imitative dynamics eliminate dominated strategies, at least along solutions starting from most initial conditions. These conclusions may seem unsurprising given the fundamental role played by dominance arguments in traditional game-theoretic analyses. In fact, these conclusions are quite special, as the following example illustrates.


Example 9.5. Survival of strictly dominated strategies. Following Berger and Hofbauer (2006) and Hofbauer and Sandholm (2011), we first consider the Smith dynamic for “bad RPS with a twin” (Figure 15(i)):

$$F(x) = Ax = \begin{pmatrix} 0 & -2 & 1 & 1 \\ 1 & 0 & -2 & -2 \\ -2 & 1 & 0 & 0 \\ -2 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \\ x_T \end{pmatrix}.$$


The Nash equilibria of F are the states on line segment $\{x^* \in X : x^* = (\frac{1}{3}, \frac{1}{3}, c, \frac{1}{3} - c)\}$, which is a repeller under the Smith dynamic. Away from Nash equilibrium, strategies gain players at rates that depend on their payoffs, but lose players at rates proportional to their current usage levels. The proportions of players choosing the twin strategies are therefore equalized, with the state approaching the plane $\mathcal{P} = \{x \in X : x_S = x_T\}$. Since F is based on bad RPS, solutions on plane \mathcal{P} approach a closed orbit away from any Nash equilibrium.⁴⁵

⁴⁴See Sato et al. (2002) for examples of chaos under the replicator dynamic, and Cowan (1992), Sparrow et al. (2008), and van Strien and Sparrow (2011) for analyses of chaos under the best response dynamic.

⁴⁵Under the replicator dynamic in this game, the state space X foliates into a continuum of invariant planes on which the ratio x_S/x_T is constant. This foliation, which is in part a consequence of inflow-outflow



(i) bad RPS with a twin



(ii) bad RPS with a feeble twin

Figure 15: Survival of strictly dominated strategies: the Smith dynamic in two games.

Figure 15(ii) presents the Smith dynamic in “bad RPS with a feeble twin”,

$$F^d(x) = A^d x = \begin{pmatrix} 0 & -2 & 1 & 1 \\ 1 & 0 & -2 & -2 \\ -2 & 1 & 0 & 0 \\ -2-d & 1-d & -d & -d \end{pmatrix} \begin{pmatrix} x_R \\ x_P \\ x_S \\ x_T \end{pmatrix}$$

with $d = \frac{1}{10}$. Evidently, the attractor from Figure 15(i) moves slightly to the left, reflecting the fact that the payoff to Twin has gone down. But since the new attractor is in the interior of X , the strictly dominated strategy Twin is always played by fractions of players bounded away from zero. These fractions need not be small: 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 $d \leq .66$. ♦

Using the Poincaré-Bendixson theorem and results on continuation of attractors, Hofbauer and Sandholm (2011) show that the phenomenon illustrated in Example 9.5 is quite general. They consider continuous evolutionary game dynamics satisfying positive correlation (PC), Nash stationarity (NS), and a third condition, *innovation*, which requires that away from Nash equilibrium, unused optimal strategies have positive growth rates. They show that for any such dynamic, there are games in which strictly dominated strategies survive in perpetuity.

We should emphasize two features of the evolutionary process that underlie this result. First, it is important that agents base their decisions on the strategies’ present performances. If agents were able to compare different strategies’ payoff functions in their entirety, they could recognize strictly dominated strategies and avoid them. While the knowledge needed to make such global comparisons is assumed implicitly in traditional analyses, this assumption is quite demanding in a large game. Without it, dominated strategies need not disappear from use.

Second, the possibility that dynamics do not converge is crucial to the survival result. Under dynamics satisfying Nash stationarity (NS), any solution trajectory that converges must converge to a Nash equilibrium, implying that dominated strategies eventually go unplayed. Example 9.5 illustrates that when solutions do not converge, so that payoffs remain in flux, evolutionary dynamics need not eliminate strategies that often perform well, but are never optimal.

symmetry (Section 5.1.4), is structurally unstable, helping illustrate why elimination results for imitative dynamics (Theorem 8.1) should be viewed as knife-edge cases.

10. Connections and Further Developments

10.1 Connections with Stochastic Stability Theory

Theorem 3.3 showed that deterministic dynamics can be used to approximate the stochastic evolutionary process $\{X_t^N\}$ in a particular limit, one that fixes a finite time horizon T and takes the population size N to infinity. To focus on very long run behavior, one can consider the opposite order of limits, fixing the population size N and taking the time horizon T to infinity. If the revision protocol is such that all strategies are always chosen with positive probability, then over an infinite time horizon, the process $\{X_t^N\}$ will visit each state in the finite grid \mathcal{X}^N infinitely often. The very long run behavior of the process is summarized by its unique stationary distribution, μ^N , which describes the limiting distribution of the process, as well as its limiting empirical distribution on almost every sample path.

Pioneering work by Foster and Young (1990), Kandori et al. (1993), and Young (1993) showed that over very long time spans, stochastic evolutionary processes often admit unique predictions of play, even in games with multiple strict equilibria. While the stationary distribution μ^N places positive mass on every state in \mathcal{X}^N , it typically concentrates its mass on a particular subset of states, often just a single state. To obtain clean results, certain limits are taken to make this concentration as stark as possible; the population states selected in this fashion are said to be *stochastically stable*.

The most common approach to stochastic stability considers revision protocols parameterized by a noise level, whether a fixed mistake probability, or a parameter reflecting the size of payoff perturbations, as in the logit choice rule (Example 5.5). States that retain mass in the stationary distribution as the noise level approaches zero are said to be *stochastically stable in the small noise limit*. This approach to stochastic stability theory is described in the chapter by Wallace and Young in this volume.

Another approach to stochastic stability, introduced by Binmore et al. (1995) and Binmore and Samuelson (1997), considers the limiting behavior of the stationary distributions μ^N as the population size N approaches infinity. This analysis of *stochastic stability in the large population limit* can be viewed as the infinite horizon analogue of the finite horizon analysis of $\{X_t^N\}$ provided by deterministic dynamics.

In view of this connection, it is natural to expect the behavior of deterministic evolutionary dynamics to have consequences for stochastic stability in the large population limit. Connections of this sort can be established using tools from (*constant step size*) *stochastic approximation theory*. Benaïm and Weibull (2003) show that states retaining mass in the stationary distributions μ^N in the large population limit must satisfy a basic no-

tion of recurrence for the mean dynamic (3): they must be elements of its *Birkhoff center*, which contains rest points, closed orbits, and more complicated limit sets.⁴⁶ Some states in the Birkhoff center—for instance, unstable rest points—seem unlikely candidates for stochastic stability. Benaïm (1998) provides conditions of a global nature ensuring that mass in μ^N does not accumulate near such states, implying that these states cannot be stochastically stable. To go further than this—to actually identify the component of the Birkhoff center containing the stochastically stable states—requires a detailed analysis of large deviations properties of the stochastic process $\{X_t^N\}$; for work in this direction, see Benaïm et al. (2013).

10.2 Connections with Models of Heuristic Learning

This chapter has considered disequilibrium dynamics in games played by large populations of agents. There are close connections between these evolutionary game dynamics and models of *heuristic learning in games*, which consider disequilibrium dynamics during repeated play of games with small numbers of players—for instance, when there is one individual in each player role of a recurring normal form game.⁴⁷

The original model of disequilibrium dynamics in games is *fictitious play*, introduced by Brown (1949, 1951) as a method of computing equilibria. During each period of the fictitious play process, each agent plays a best response to the time average of his opponents' past choices. The state variable of this process, the time average of play, has increments that become smaller over time, since choices late in the process have small effects on the overall time average. After accounting for this reparameterization of time, the fictitious play process is essentially a discrete-time version of the best response dynamic (17). Indeed, convergence results developed in one context often have analogues in the other.⁴⁸

Since best responses are generically pure, convergence of time averages in the fictitious play process does not imply convergence of players' period-by-period intended play. To generate such a link, Fudenberg and Kreps (1993) introduced *stochastic fictitious play*, in which players play best responses only after their payoffs have been subjected to perturbations. The result is a stochastic process with decreasing step size whose expected motion is described by perturbed best response dynamics (23). By combining analyses of these dynamics with tools from (*decreasing step size*) *stochastic approximation theory* (Benaïm

⁴⁶See also Benaïm (1998), Sandholm (2007c), and Roth and Sandholm (2013).

⁴⁷For surveys of this literature, see Young (2004) and Hart (2005).

⁴⁸See Hofbauer (1995b), Harris (1998), Berger (2005, 2007, 2008, 2012), and Benaïm et al. (2005).

(1999)), one can establish global convergence results for stochastic fictitious play.⁴⁹

The decision rule used in stochastic fictitious play generates an ε -consistent strategy for repeated play of a normal form game: a player using this rule could not have improved his average payoff by more than ε by having played a single pure strategy in all previous periods. There are other subtle links between consistency in repeated games and evolutionary game dynamics. For instance, the convergence results for contractive games in Section 7.5.2 are directly inspired by a class of consistent repeated game strategies introduced by Hart and Mas-Colell (2001).

There are also connections between evolutionary game dynamics and models of *reinforcement learning*, in which each player's mixed strategy weights are determined by his strategies' aggregated past performances. Börgers and Sarin (1997) show that when players' choices are based on the reinforcement learning rule of Cross (1973), then in the continuous-time limit, the players' mixed strategies evolve according to the replicator dynamic (8).⁵⁰ Still other connections exist with hybrid models, in which large populations of agents employ heuristic learning rules that condition on statistics of the history of play.⁵¹

10.3 Games with Continuous Strategy Sets

Many applications of game theory are naturally modeled using a continuous strategy space. Population dynamics for this setting were first used by Hines (1980) and Zeeman (1981) to study evolution in populations of agents playing mixed strategies, a generalization of the environment in which Maynard Smith and Price (1973) introduced the ESS concept (see Section 7.1). Bomze (1990, 1991) continued this line of research by defining the replicator dynamic for general games with continuous strategy sets.

Analyzing population dynamics in this setting introduces novel technical and conceptual issues. If the strategy set is a convex subset of \mathbb{R}^n , then a population state is a probability measure on this space. Since the space of these probability measures is a function space, defining evolutionary dynamics requires one not only to write down an ordinary differential equation, but also to specify a norm with respect to which to this

⁴⁹See Fudenberg and Kreps (1993), Kaniovski and Young (1995), Benaïm and Hirsch (1999), Hofbauer and Sandholm (2002), Benaïm and Faure (2012), and Perkins and Leslie (2013).

⁵⁰Hopkins (2002) (see also Hofbauer et al. (2009b)) uses surprising relations between the replicator and logit dynamics to establish strong links between behavior under reinforcement learning and period-by-period choices under stochastic fictitious play. Kosfeld et al. (2002), Tsakas and Voorneveld (2009), and Laraki and Mertikopoulos (2013) study mixed strategy dynamics for normal form games derived from other heuristic principles.

⁵¹See Hopkins (1999a), Ellison and Fudenberg (2000), Ramsza and Seymour (2010), Fudenberg and Takahashi (2011), and Lahkar and Seymour (2013).

equation's time derivative is defined.

Moreover, while the meaning of local stability is unambiguous when strategy sets are finite, this is no longer the case with a continuous strategy sets, a point emphasized by Oechssler and Riedel (2002). With finite strategy sets, population states are close to one another if they assign similar probabilities to each pure strategy. With continuous strategy sets, this notion of closeness is captured by using a variational norm on probability measures. But in the latter context, there is a second sense in which population states might be close. For instance, one might want to regard two pure population states as close together if the strategies played in these states are close together in the Euclidean metric on \mathbb{R}^n .

Therefore, before studying local stability under evolutionary dynamics, one must first choose the definition of a neighborhood in the space of population states. If one hews closely to the finite case by defining neighborhoods in terms of the variational norm, one can obtain sufficient conditions for local stability under the replicator dynamic using suitable generalizations of the ESS concept.⁵² Alternatively, if the definition of a neighborhood captures only the second sort of closeness, allowing large numbers of agents to make small changes in strategy, then quite distinct conditions for local stability based on comparisons to the performances of nearby strategies become relevant.⁵³ Finally, one might want to define neighborhoods in a way that captures both sorts of closeness. To accomplish this, Oechssler and Riedel (2002) introduce a condition called *evolutionary robustness*, which uses neighborhoods derived from the weak topology on probability measures, and show that it is sufficient for local stability under a class of imitative dynamics.⁵⁴

Once one allows that certain pairs of strategies are closer together than others, it becomes natural to ask whether not only notions of stability, but also the very definition of the dynamic, should take distances between strategies into account. Friedman and Ostrov (2008, 2010, 2013) consider a framework in which each agent adjusts his strategy continuously in response to the incentives he faces; the use of continuous adjustment is justified as an optimal response to adjustment costs. Rather than obeying an ordinary differential equation, the evolution of aggregate behavior in this framework follows a partial differential equation. This formulation introduces a wide range of new dynamic phenomena, and requires analytical techniques different from those considered in this

⁵²See Bomze (1990), Oechssler and Riedel (2001), and Norman (2008).

⁵³Eshel and Motro (1981), Eshel (1983), and Apaloo (1997) introduce such conditions, and Eshel and Sansone (2003), Cressman (2005), Cressman et al. (2006), and van Veelen and Spreij (2009) show them to be sufficient for local stability under the replicator dynamic.

⁵⁴See Cressman and Hofbauer (2005) for further stability results for the replicator dynamic. For definitions and stability analyses of the BNN, logit, and Smith dynamics with continuous strategy sets, see Hofbauer et al. (2009a), Lahkar and Riedel (2013), and Cheung (2013), respectively.

chapter, leaving many open questions for future research.

10.4 Extensive Form Games and Set-Valued Solution Concepts

To describe behavior in populations whose members are engaged in sequential interactions, one can define population games via matching in extensive form games, and evaluate the properties of evolutionary dynamics in these population games.

Even in the traditional theory, both the notion of a strategy for an extensive form game and the modeling of beliefs about events off of the path of play raise conceptual difficulties. Variations on these difficulties persist in large population contexts. But new difficulties arise as well. For instance, it is well known that in extensive form games with perfect recall, there is a many-to-one map between mixed strategies and outcome equivalent behavior strategies, with equivalent mixed strategies exhibiting different correlations among choices at temporally unordered information sets. Chamberland and Cressman (2000) show that under evolutionary dynamics, statistical correlation among choices at different information sets can interfere with the forces of selection, leading to unexpected forms of dynamic instability.

Traditionally, a fundamental idea in the analysis of extensive form games is *sequential rationality*: the requirement that players' equilibrium strategies specify actions that are rational from the perspective of the moments the actions would be taken, even if these moments should not occur during equilibrium play. It is a basic question whether the predictions of evolutionary dynamics agree with solution concepts embodying sequential rationality. The strongest agreements occur in games of (stagewise) perfect information, for equilibria in which all information sets are reached: Cressman (1997a) and Cressman and Schlag (1998) show that in such cases, asymptotic stability under the replicator and best response dynamics agree with subgame perfection in the extensive form. Although positive results are available for certain further instances,⁵⁵ connections between sequential rationality and asymptotic stability have not been established in great generality.⁵⁶

In extensive form games, the possibility of unreached information sets gives rise to connected sets of outcome-equivalent equilibria, which differ only in the choices they specify off the path of play. For this reason, it is natural to consider the stability of sets of equilibria—for instance, by defining set-valued extensions of the notion of ESS. In single-population contexts, Thomas's (1985) notion of an *evolutionarily stable set* (or *ES set*) provides a sufficient condition for asymptotic stability under the replicator dynamic.⁵⁷ In

⁵⁵See, e.g., Cressman (1996).

⁵⁶See Cressman (2003) for a comprehensive treatment of this literature.

⁵⁷See Balkenborg and Schlag (2001) and van Veelen (2012) for alternative characterizations and related

multipopulation contexts, Balkenborg and Schlag (2007) show that Balkenborg's (1994) notion of a *strict equilibrium set* (or *SE set*) plays an analogous role.

Using ideas from differential topology, Demichelis and Ritzberger (2003) develop a powerful approach to evaluating local stability of sets of Nash equilibria.⁵⁸ They assign each component of Nash equilibria an integer, the *index* of the component, that is determined by the behavior of the payoff vector field in a neighborhood of the component. For any evolutionary dynamic satisfying positive correlation (PC) and Nash stationarity (NS), a necessary condition for the component to be asymptotically stable is that its index agree with the component's Euler characteristic. In typical (specifically, contractible) cases, this means that a component can only be asymptotically stable if its index is 1. Through this result, one can establish the instability of a component of equilibria under a large class of evolutionary dynamics—the same class in which convergence to Nash equilibrium is ensured in potential games—using only information about the payoffs of the game.

10.5 Applications

Population games provide a formal framework for applications in which large numbers of participants make interdependent decisions. As we argued at the start, such applications are precisely the ones in which traditional equilibrium assumptions seem questionable. By taking evolutionary dynamics to applications, one can both justify equilibrium predictions, and, more interestingly, call attention to settings in which disequilibrium predictions are warranted.

This latter possibility is well illustrated by analyses of price dispersion. While in idealized markets, homogeneous goods sell at a single price, in actual markets one often observes a variety of prices for the same good. Research in information economics has shown that with heterogeneously informed consumers, price dispersion can occur as an equilibrium phenomenon: different sellers choose different prices, and each seller, knowing the choices of the others, is content with his own choice.⁵⁹

Using tools from evolutionary game theory, Hopkins and Seymour (2002), Lahkar (2011), Hahn (2012), and Lahkar and Riedel (2013) argue that equilibrium is not the best explanation for the price dispersion we see. They show that equilibria with price dispersion are unstable, but that price dispersion can persist through disequilibrium cycles. Indeed, empirical and experimental evidence supports the conclusion that price

notions.

⁵⁸See also Ritzberger (1994), Demichelis and Germano (2000, 2002), and McLennan (2012).

⁵⁹See Varian (1980) and Burdett and Judd (1983).

dispersion may be a cyclical phenomenon rather than an equilibrium phenomenon.⁶⁰

Deterministic evolutionary dynamics have been used to study a variety of other topics in economics and social science. A partial list includes auctions (Louge and Riedel (2012)); fiat money (Sethi (1999)); conspicuous consumption (Friedman and Ostrov (2008)); common resource use (Sethi and Somanathan (1996)); cultural evolution (Bisin and Verdier (2001), Sandholm (2001c), Kuran and Sandholm (2008), Montgomery (2010)); the evolution of language (Pawlowitsch (2008)); implementation problems (Cabrales and Ponti (2000), Sandholm (2002, 2005b), Fujishima (2012)); international trade (Friedman and Fung (1996)); residential segregation (Dokumacı and Sandholm (2007)); preference evolution (Sandholm (2001c), Heifetz et al. (2007), Norman (2012)); and theories of mind (Mohlin (2012)). Further applications of deterministic evolutionary game dynamics to these topics and others offer fascinating directions for future research.

⁶⁰See Lach (2002), Eckert (2003), Noel (2007), and Cason et al. (2005).

References

- Akin, E. (1979). *The Geometry of Population Genetics*. Springer, Berlin.
- Akin, E. (1980). Domination or equilibrium. *Mathematical Biosciences*, 50:239–250.
- Apaloo, J. (1997). Revisiting strategic models of evolution: The concept of neighborhood invader strategies. *Theoretical Population Biology*, 52:71–77.
- Arneodo, A., Couillet, P., and Tresser, C. (1980). Occurrence of strange attractors in three-dimensional Volterra equations. *Physics Letters*, 79A:259–263.
- Balder, E. J. (2002). A unifying pair of Cournot–Nash equilibrium existence results. *Journal of Economic Theory*, 102:437–470.
- Balkenborg, D. (1994). Strictness and evolutionary stability. Discussion paper 52, Center for the Study of Rationality, Hebrew University.
- Balkenborg, D., Hofbauer, J., and Kuzmics, C. (2013). Refined best-response correspondence and dynamics. *Theoretical Economics*, 8:165–192.
- Balkenborg, D. and Schlag, K. H. (2001). Evolutionarily stable sets. *International Journal of Game Theory*, 29:571–595.
- Balkenborg, D. and Schlag, K. H. (2007). On the evolutionary selection of sets of Nash equilibria. *Journal of Economic Theory*, 133:295–315.
- Beckmann, M., McGuire, C. B., and Winsten, C. B. (1956). *Studies in the Economics of Transportation*. Yale University Press, New Haven.
- Benaïm, M. (1998). Recursive algorithms, urn processes, and the chaining number of chain recurrent sets. *Ergodic Theory and Dynamical Systems*, 18:53–87.
- Benaïm, M. (1999). Dynamics of stochastic approximation algorithms. In Azéma, J., Émery, M., Ledoux, M., and Yor, M., editors, *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pages 1–68. Springer, Berlin.
- Benaïm, M. and Faure, M. (2012). Stochastic approximations, cooperative dynamics, and supermodular games. *Annals of Applied Probability*, 22:2133–2164.
- Benaïm, M. and Hirsch, M. W. (1999). Mixed equilibria and dynamical systems arising from fictitious play in perturbed games. *Games and Economic Behavior*, 29:36–72.
- Benaïm, M., Hofbauer, J., and Hopkins, E. (2009). Learning in games with unstable equilibria. *Journal of Economic Theory*, 144:1694–1709.
- Benaïm, M., Hofbauer, J., and Sorin, S. (2005). Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44:328–348.

- Benaïm, M., Sandholm, W. H., and Staudigl, M. (2013). Large deviations and stochastic stability in the large population limit. Unpublished manuscript, Université de Neuchâtel, University of Wisconsin, and Bielefeld University.
- Benaïm, M. and Weibull, J. W. (2003). Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903.
- Benaïm, M. and Weibull, J. W. (2009). Mean-field approximation of stochastic population processes in games. Unpublished manuscript, Université de Neuchâtel and Stockholm School of Economics.
- Berger, U. (2005). Fictitious play in $2 \times n$ games. *Journal of Economic Theory*, 120:139–154.
- Berger, U. (2007). Two more classes of games with the continuous-time fictitious play property. *Games and Economic Behavior*, 60:247–261.
- Berger, U. (2008). Learning in games with strategic complementarities revisited. *Journal of Economic Theory*, 143:292–301.
- Berger, U. (2012). Non-algebraic convergence proofs for continuous-time fictitious play. *Dynamic Games and Applications*, 2:4–17.
- 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. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74:235–265.
- Binmore, K., Samuelson, L., and Vaughan, R. (1995). Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35. Erratum, 21 (1997), 325.
- Bisin, A. and Verdier, T. (2001). The economics of cultural transmission and the dynamics of preferences. *Journal of Economic Theory*, 97:298–319.
- Björnerstedt, J. and Weibull, J. W. (1996). Nash equilibrium and evolution by imitation. In Arrow, K. J. et al., editors, *The Rational Foundations of Economic Behavior*, pages 155–181. St. Martin's Press, New York.
- Bomze, I. M. (1986). Non-cooperative two-person games in biology: A classification. *International Journal of Game Theory*, 15:31–57.
- Bomze, I. M. (1990). Dynamical aspects of evolutionary stability. *Monatshefte für Mathematik*, 110:189–206.
- Bomze, I. M. (1991). Cross entropy minimization in uninvadable states of complex populations. *Journal of Mathematical Biology*, 30:73–87.
- Bomze, I. M. and Pötscher, B. M. (1989). *Game Theoretical Foundations of Evolutionary Stability*. Springer, Berlin.

- Börgers, T. and Sarin, R. (1997). Learning through reinforcement and the replicator dynamics. *Journal of Economic Theory*, 77:1–14.
- Brown, G. W. (1949). Some notes on computation of games solutions. Report P-78, The Rand Corporation.
- Brown, G. W. (1951). Iterative solutions of games by fictitious play. In Koopmans, T. C. et al., editors, *Activity Analysis of Production and Allocation*, pages 374–376. Wiley, New York.
- 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.
- Cabrales, A. and Ponti, G. (2000). Implementation, elimination of weakly dominated strategies, and evolutionary dynamics. *Review of Economic Dynamics*, 3:247–282.
- Carmona, G. and Podczeck, K. (2009). On the existence of pure-strategy equilibria in large games. *Journal of Economic Theory*, 144:1300–1319.
- Cason, T. N., Friedman, D., and Wagener, F. (2005). The dynamics of price dispersion, or Edgeworth variations. *Journal of Economic Dynamics and Control*, 29:801–822.
- Chamberland, M. and Cressman, R. (2000). An example of dynamic (in)consistency in symmetric extensive form evolutionary games. *Games and Economic Behavior*, 30:319–326.
- Cheung, M. W. (2013). Pairwise comparison dynamics for games with continuous strategy space. Unpublished manuscript, University of Wisconsin.
- Conley, C. (1978). *Isolated Invariant Sets and the Morse Index*. American Mathematical Society, Providence, RI.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. (2004). Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29:961–976.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. (2008). A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64:457–469.
- Cowan, S. G. (1992). *Dynamical Systems Arising from Game Theory*. PhD thesis, University of California, Berkeley.
- Cressman, R. (1992). *The Stability Concept of Evolutionary Game Theory: A Dynamic Approach*. Springer, Berlin.
- Cressman, R. (1996). Evolutionary stability in the finitely repeated prisoner’s dilemma game. *Journal of Economic Theory*, 68:234–248.

- Cressman, R. (1997a). Dynamic stability in symmetric extensive form games. *International Journal of Game Theory*, 26:525–547.
- Cressman, R. (1997b). Local stability of smooth selection dynamics for normal form games. *Mathematical Social Sciences*, 34:1–19.
- Cressman, R. (2003). *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge.
- Cressman, R. (2005). Stability of the replicator equation with continuous strategy space. *Mathematical Social Sciences*, 50:127–147.
- Cressman, R. (2006). Uninvadability in N -species frequency models for resident-mutant systems with discrete or continuous time. *Theoretical Population Biology*, 69:253–262.
- Cressman, R. (2011). Beyond the symmetric normal form: Extensive form games, asymmetric games, and games with continuous strategy spaces. In Sigmund, K., editor, *Evolutionary Game Dynamics*, pages 27–59. American Mathematical Society, Providence, RI.
- Cressman, R., Garay, J., and Hofbauer, J. (2001). Evolutionary stability concepts for N -species frequency-dependent interactions. *Journal of Theoretical Biology*, 211:1–10.
- Cressman, R. and Hofbauer, J. (2005). Measure dynamics on a one-dimensional continuous trait space: Theoretical foundations for adaptive dynamics. *Theoretical Population Biology*, 67:47–59.
- Cressman, R., Hofbauer, J., and Riedel, F. (2006). Stability of the replicator equation for a single-species with a multi-dimensional continuous trait space. *Journal of Theoretical Biology*, 239:273–288.
- Cressman, R. and Schlag, K. H. (1998). On the dynamic (in)stability of backwards induction. *Journal of Economic Theory*, 83:260–285.
- Cross, J. G. (1973). A stochastic learning model of economic behavior. *Quarterly Journal of Economics*, 87:239–266.
- Crow, J. F. and Kimura, M. (1970). *An Introduction to Population Genetics Theory*. Harper and Row, New York.
- Dafermos, S. and Sparrow, F. T. (1969). The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards B*, 73:91–118.
- Demichelis, S. and Germano, F. (2000). On the indices of zeros of Nash fields. *Journal of Economic Theory*, 94:192–217.
- Demichelis, S. and Germano, F. (2002). On (un)knots and dynamics in games. *Games and Economic Behavior*, 41:46–60.

- Demichelis, S. and Ritzberger, K. (2003). From evolutionary to strategic stability. *Journal of Economic Theory*, 113:51–75.
- Dokumacı E. and Sandholm, W. H. (2007). Schelling redux: An evolutionary model of residential segregation. Unpublished manuscript, University of Wisconsin.
- Eckert, A. (2003). Retail price cycles and presence of small firms. *International Journal of Industrial Organization*, 21:151–170.
- Eigen, M. and Schuster, P. (1979). *The Hypercycle: A Principle of Natural Self-Organization*. Springer, Berlin.
- Ellison, G. and Fudenberg, D. (2000). Learning purified mixed equilibria. *Journal of Economic Theory*, 90:84–115.
- Ely, J. C. and Sandholm, W. H. (2005). Evolution in Bayesian games I: Theory. *Games and Economic Behavior*, 53:83–109.
- Eshel, I. (1983). Evolutionary and continuous stability. *Journal of Theoretical Biology*, 103:99–111.
- Eshel, I. and Motro, U. (1981). Kin selection and strong evolutionary stability of mutual help. *Theoretical Population Biology*, 19:420–433.
- Eshel, I. and Sansone, E. (2003). Evolutionary and dynamic stability in continuous population games. *Journal of Mathematical Biology*, 46:445–459.
- Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232. Corrigendum, 51 (1997), 77-78.
- Franchetti, F. and Sandholm, W. H. (2013). An introduction to *Dynamo*: Diagrams for evolutionary game dynamics. *Biological Theory*, forthcoming.
- Friedman, D. (1991). Evolutionary games in economics. *Econometrica*, 59:637–666.
- Friedman, D. and Fung, K. C. (1996). International trade and the internal organization of firms: An evolutionary approach. *Journal of International Economics*, 41:113–136.
- Friedman, D. and Ostrov, D. N. (2008). Conspicuous consumption dynamics. *Games and Economic Behavior*, 64:121–145.
- Friedman, D. and Ostrov, D. N. (2010). Gradient dynamics in population games: Some basic results. *Journal of Mathematical Economics*, 46:691–707.
- Friedman, D. and Ostrov, D. N. (2013). Evolutionary dynamics over continuous action spaces for population games that arise from symmetric two-player games. *Journal of Economic Theory*, 148:743–777.

- Fudenberg, D. and Kreps, D. M. (1993). Learning mixed equilibria. *Games and Economic Behavior*, 5:320–367.
- Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press, Cambridge.
- Fudenberg, D. and Takahashi, S. (2011). Heterogeneous beliefs and local information in stochastic fictitious play. *Games and Economic Behavior*, 71:100–120.
- Fujishima, S. (2012). Evolutionary implementation with estimation of unknown externalities. Unpublished manuscript, Washington University.
- 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.
- Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2008). Quantal response equilibrium. In Blume, L. E. and Durlauf, S. N., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, Basingstoke, U.K., second edition.
- Hahn, M. (2012). An evolutionary analysis of Varian’s model of sales. *Dynamic Games and Applications*, 2:71–96.
- Harris, C. (1998). On the rate of convergence of continuous-time fictitious play. *Games and Economic Behavior*, 22:238–259.
- Hart, S. (2005). Adaptive heuristics. *Econometrica*, 73:1401–1430.
- Hart, S. and Mas-Colell, A. (2001). A general class of adaptive strategies. *Journal of Economic Theory*, 98:26–54.
- Hart, S. and Mas-Colell, A. (2003). Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93:1830–1836.
- Heifetz, A., Shannon, C., and Spiegel, Y. (2007). What to maximize if you must. *Journal of Economic Theory*, 133:31–57.
- Helbing, D. (1992). A mathematical model for behavioral changes by pair interactions. In Haag, G., Mueller, U., and Troitzsch, K. G., editors, *Economic Evolution and Demographic Change: Formal Models in Social Sciences*, pages 330–348. Springer, Berlin.
- Hines, W. G. S. (1980). Strategic stability in complex populations. *Journal of Applied Probability*, 17:600–610.
- Hines, W. G. S. (1987). Evolutionary stable strategies: A review of basic theory. *Theoretical Population Biology*, 31:195–272.

- Hiriart-Urruty, J.-B. and Lemaréchal, C. (2001). *Fundamentals of Convex Analysis*. Springer, Berlin.
- Hirsch, M. W. (1988). Systems of differential equations that are competitive or cooperative III: Competing species. *Nonlinearity*, 1:51–71.
- Hofbauer, J. (1981). On the occurrence of limit cycles in the Volterra-Lotka equation. *Nonlinear Analysis*, 5:1003–1007.
- Hofbauer, J. (1995a). Imitation dynamics for games. Unpublished manuscript, University of Vienna.
- Hofbauer, J. (1995b). 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. (2011). Deterministic evolutionary game dynamics. In Sigmund, K., editor, *Evolutionary Game Dynamics*, pages 61–79. American Mathematical Society, Providence, RI.
- Hofbauer, J., Oechssler, J., and Riedel, F. (2009a). Brown-von Neumann-Nash dynamics: The continuous strategy case. *Games and Economic Behavior*, 65:406–429.
- 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. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132:47–69.
- Hofbauer, J. and Sandholm, W. H. (2009). Stable games and their dynamics. *Journal of Economic Theory*, 144:1665–1693.
- Hofbauer, J. and Sandholm, W. H. (2011). Survival of dominated strategies under evolutionary dynamics. *Theoretical Economics*, 6:341–377.
- Hofbauer, J., Schuster, P., and Sigmund, K. (1979). A note on evolutionarily stable strategies and game dynamics. *Journal of Theoretical Biology*, 81:609–612.
- Hofbauer, J. and Sigmund, K. (1988). *Theory of Evolution and Dynamical Systems*. Cambridge University Press, Cambridge.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge.
- Hofbauer, J. and Sigmund, K. (2003). Evolutionary game dynamics. *Bulletin of the American Mathematical Society (New Series)*, 40:479–519.

- Hofbauer, J., Sorin, S., and Viossat, Y. (2009b). Time average replicator and best reply dynamics. *Mathematics of Operations Research*, 34:263–269.
- Hofbauer, J. and Weibull, J. W. (1996). Evolutionary selection against dominated strategies. *Journal of Economic Theory*, 71:558–573.
- Hopkins, E. (1999a). Learning, matching, and aggregation. *Games and Economic Behavior*, 26:79–110.
- Hopkins, E. (1999b). A note on best response dynamics. *Games and Economic Behavior*, 29:138–150.
- Hopkins, E. (2002). Two competing models of how people learn in games. *Econometrica*, 70:2141–2166.
- Hopkins, E. and Seymour, R. M. (2002). The stability of price dispersion under seller and consumer learning. *International Economic Review*, 43:1157–1190.
- Hwang, S.-H., Katsoulakis, M., and Rey-Bellet, L. (2013). Deterministic equations for stochastic spatial evolutionary games. *Theoretical Economics*, 8:829–874.
- Jordan, J. S. (1993). Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior*, 5:368–386.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56.
- Kaniovski, Y. M. and Young, H. P. (1995). Learning dynamics in games with stochastic perturbations. *Games and Economic Behavior*, 11:330–363.
- Khan, M. A. and Sun, Y. (2002). Non-cooperative games with many players. In Aumann, R. J. and Hart, S., editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 46, pages 1761–1808. North Holland.
- Kimura, M. (1958). On the change of population fitness by natural selection. *Heredity*, 12:145–167.
- Kosfeld, M., Droste, E., and Voorneveld, M. (2002). A myopic adjustment process leading to best reply matching. *Journal of Economic Theory*, 40:270–298.
- Kreindler, G. E. and Young, H. P. (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior*, 80:39–67.
- Kuran, T. and Sandholm, W. H. (2008). Cultural integration and its discontents. *Review of Economic Studies*, 75:201–228.
- Kurtz, T. G. (1970). Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7:49–58.

- Lach, S. (2002). Existence and persistence of price dispersion: An empirical analysis. *Review of Economics and Statistics*, 84:433–444.
- Lahkar, R. (2011). The dynamic instability of dispersed price equilibria. *Journal of Economic Theory*, 146:1796–1827.
- Lahkar, R. and Riedel, F. (2013). The continuous logit dynamic and price dispersion. Unpublished manuscript, Institute for Financial Management and Research and Bielefeld University.
- Lahkar, R. and Sandholm, W. H. (2008). The projection dynamic and the geometry of population games. *Games and Economic Behavior*, 64:565–590.
- Lahkar, R. and Seymour, R. M. (2013). Reinforcement learning in population games. *Games and Economic Behavior*, 8:10–38.
- Laraki, R. and Mertikopoulos, P. (2013). Higher order game dynamics. *Journal of Economic Theory*, 148:2666–2695.
- Louge, F. and Riedel, F. (2012). Evolutionary stability in first price auctions. *Dynamic Games and Applications*, 2:110–128.
- Matsui, A. (1992). Best response dynamics and socially stable strategies. *Journal of Economic Theory*, 57:343–362.
- Matsui, A. and Matsuyama, K. (1995). An approach to equilibrium selection. *Journal of Economic Theory*, 65:415–434.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38.
- McLennan, A. (2012). Advanced fixed point theory for economics. Unpublished manuscript, University of Queensland.
- Milgrom, P. and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58:1255–1278.
- Minty, G. J. (1967). On the generalization of a direct method of the calculus of variations. *Bulletin of the American Mathematical Society*, 73:315–321.
- Mohlin, E. (2012). Evolution of theories of mind. *Games and Economic Behavior*, 75:299–318.
- Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.

- Montgomery, J. (2010). Intergenerational cultural transmission as an evolutionary game. *American Economic Journal: Microeconomics*, 2:115–136.
- Moran, P. A. P. (1962). *The Statistical Processes of Evolutionary Theory*. Clarendon Press, Oxford.
- Morris, S., Rob, R., and Shin, H. S. (1995). p -dominance and belief potential. *Econometrica*, 63:145–157.
- 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.
- Nash, J. F. (1950). *Non-Cooperative Games*. PhD thesis, Princeton. Reprinted in *The Essential John Nash*, H. W. Kuhn and S. Nasar, Eds., p. 51-83, Princeton University Press, Princeton (2002).
- Nash, J. F. (1951). Non-cooperative games. *Annals of Mathematics*, 54:287–295.
- Noel, M. D. (2007). Edgeworth price cycles: Evidence from the Toronto retail gasoline market. *Journal of Industrial Economics*, 55:69–92.
- Norman, T. W. L. (2008). Dynamically stable sets in infinite strategy spaces. *Games and Economic Behavior*, 62:610–627.
- Norman, T. W. L. (2012). Equilibrium selection and the dynamic evolution of preferences. *Games and Economic Behavior*, 74:311–320.
- Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap/Harvard, Cambridge.
- Oechssler, J. and Riedel, F. (2001). Evolutionary dynamics on infinite strategy spaces. *Economic Theory*, 17:141–162.
- Oechssler, J. and Riedel, F. (2002). On the dynamic foundation of evolutionary stability in continuous models. *Journal of Economic Theory*, 107:223–252.
- Oyama, D., Sandholm, W. H., and Tercieux, O. (2012). Sampling best response dynamics and deterministic equilibrium selection. Unpublished manuscript, University of Tokyo, University of Wisconsin, and Paris School of Economics.
- Oyama, D. and Tercieux, O. (2009). Iterated potential and robustness of equilibria. *Journal of Economic Theory*, 144:1726–1769.
- Pawlowitsch, C. (2008). Why evolution does not always lead to an optimal signaling system. *Games and Economic Behavior*, 63:203–226.

- Perkins, S. and Leslie, D. S. (2013). Stochastic fictitious play with continuous action sets. Unpublished manuscript, University of Bristol.
- Pohley, H.-J. and Thomas, B. (1983). Non-linear ESS models and frequency dependent selection. *BioSystems*, 16:87–100.
- Ramsza, M. and Seymour, R. M. (2010). Fictitious play in an evolutionary environment. *Games and Economic Behavior*, 68:303–324.
- Ritzberger, K. (1994). The theory of normal form games from the differentiable viewpoint. *International Journal of Game Theory*, 23:207–236.
- Ritzberger, K. and Weibull, J. W. (1995). Evolutionary selection in normal form games. *Econometrica*, 63:1371–1399.
- 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.
- Roth, G. and Sandholm, W. H. (2013). Stochastic approximations with constant step size and differential inclusions. *SIAM Journal on Control and Optimization*, 51:525–555.
- Roughgarden, T. and Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM*, 49:236–259.
- Roughgarden, T. and Tardos, É. (2004). Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 49:389–403.
- Samuelson, L. and Zhang, J. (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391.
- Sandholm, W. H. (2001a). Almost global convergence to p -dominant equilibrium. *International Journal of Game Theory*, 30:107–116.
- Sandholm, W. H. (2001b). Potential games with continuous player sets. *Journal of Economic Theory*, 97:81–108.
- Sandholm, W. H. (2001c). Preference evolution, two-speed dynamics, and rapid social change. *Review of Economic Dynamics*, 4:637–639.
- Sandholm, W. H. (2002). Evolutionary implementation and congestion pricing. *Review of Economic Studies*, 69:667–689.
- Sandholm, W. H. (2003). Evolution and equilibrium under inexact information. *Games and Economic Behavior*, 44:343–378.

- Sandholm, W. H. (2005a). Excess payoff dynamics and other well-behaved evolutionary dynamics. *Journal of Economic Theory*, 124:149–170.
- Sandholm, W. H. (2005b). Negative externalities and evolutionary implementation. *Review of Economic Studies*, 72:885–915.
- Sandholm, W. H. (2007a). Evolution in Bayesian games II: Stability of purified equilibria. *Journal of Economic Theory*, 136:641–667.
- Sandholm, W. H. (2007b). Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132:367–382.
- Sandholm, W. H. (2007c). Simple formulas for stationary distributions and stochastically stable states. *Games and Economic Behavior*, 59:154–162.
- Sandholm, W. H. (2009a). Evolutionary game theory. In Meyers, R. A., editor, *Encyclopedia of Complexity and Systems Science*, pages 3176–3205. Springer, Heidelberg.
- Sandholm, W. H. (2009b). Large population potential games. *Journal of Economic Theory*, 144:1710–1725.
- Sandholm, W. H. (2010a). Local stability under evolutionary game dynamics. *Theoretical Economics*, 5:27–50.
- Sandholm, W. H. (2010b). Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. *Games*, 1:3–17.
- Sandholm, W. H. (2010c). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge.
- Sandholm, W. H. (2013). Probabilistic interpretations of integrability for game dynamics. *Dynamic Games and Applications*, forthcoming.
- Sandholm, W. H. (2014). Local stability of strict equilibrium under evolutionary game dynamics. *Journal of Dynamics and Games*, forthcoming.
- Sandholm, W. H., Dokumacı, E., and Franchetti, F. (2012). Dynamo: Diagrams for evolutionary game dynamics. Software. <http://www.ssc.wisc.edu/~whs/dynamo>.
- Sandholm, W. H., Dokumacı, E., and Lahkar, R. (2008). The projection dynamic and the replicator dynamic. *Games and Economic Behavior*, 64:666–683.
- Sato, Y., Akiyama, E., and Farmer, J. D. (2002). Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99:4748–4751.
- Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78:130–156.

- Schmeidler, D. (1973). Equilibrium points of non-atomic games. *Journal of Statistical Physics*, 7:295–300.
- Schuster, P. and Sigmund, K. (1983). Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538.
- Sethi, R. (1999). Evolutionary stability and media of exchange. *Journal of Economic Behavior and Organization*, 40:233–254.
- Sethi, R. and Somanathan, E. (1996). The evolution of social norms in common property resource use. *American Economic Review*, 86:766–788.
- Shahshahani, S. (1979). A new mathematical framework for the study of linkage and selection. *Memoirs of the American Mathematical Society*, 211.
- 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.
- Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge.
- Skyrms, B. (1992). Chaos in game dynamics. *Journal of Logic, Language, and Information*, 1:111–130.
- Smirnov, G. V. (2002). *Introduction to the Theory of Differential Inclusions*. American Mathematical Society, Providence, RI.
- Smith, H. L. (1995). *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, Providence, RI.
- 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. (1979). Evolutionarily stable strategies with two types of players. *Journal of Applied Probability*, 16:76–83.
- Taylor, P. D. and Jonker, L. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.
- Tercieux, O. (2006). p -best response set. *Journal of Economic Theory*, 131:45–70.

- Thomas, B. (1984). Evolutionary stability: States and strategies. *Theoretical Population Biology*, 26:49–67.
- Thomas, B. (1985). On evolutionarily stable sets. *Journal of Mathematical Biology*, 22:105–115.
- Topkis, D. (1979). Equilibrium points in nonzero-sum n -person submodular games. *SIAM Journal on Control and Optimization*, 17:773–787.
- Tsakas, E. and Voorneveld, M. (2009). The target projection dynamic. *Games and Economic Behavior*, 67:708–719.
- van Strien, S. and Sparrow, C. (2011). Fictitious play in 3×3 games: Chaos and dithering behavior. *Games and Economic Behavior*, 73:262–286.
- van Veelen, M. (2012). Robustness against indirect invasions. *Games and Economic Behavior*, 74:382–393.
- van Veelen, M. and Spreij, P. (2009). Evolution in games with a continuous action space. *Economic Theory*, 39:355–376.
- Varian, H. (1980). A model of sales. *American Economic Review*, 70:651–659.
- Viossat, Y. (2011). Monotonic dynamics and dominated strategies. Unpublished manuscript, Université Paris-Dauphine.
- Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19:305–321.
- Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge.
- 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.
- Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.
- Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.
- Young, H. P. (2004). *Strategic Learning and Its Limits*. Oxford University Press, Oxford.
- Zeeman, E. C. (1980). Population dynamics from game theory. In Nitecki, Z. and Robinson, C., editors, *Global Theory of Dynamical Systems (Evanston, 1979)*, number 819 in Lecture Notes in Mathematics, pages 472–497, Berlin. Springer.
- Zeeman, E. C. (1981). Dynamics of the evolution of animal conflicts. *Journal of Theoretical Biology*, 89:249–270.
- Zusai, D. (2011). The tempered best response dynamic. Unpublished manuscript, Temple University.