

Evolutionary Game Theory*

William H. Sandholm
University of Wisconsin[†]

November 12, 2007

Article Outline

Glossary	2
I Definition and Importance	3
II Introduction	4
III Normal Form Games	5
IV Static Notions of Evolutionary Stability	8
V Population Games	10
VI Revision Protocols	12
VI.1 Definition	12
VI.2 Examples	13
VII Deterministic Dynamics	14
VII.1 Mean Dynamics	15
VII.2 Examples	16
VII.3 Evolutionary Justification of Nash Equilibrium	20
VII.4 Local Stability	21

*The figures in Sections VII and IX were created using Dynamo [184] and VirtualLabs [92], respectively. I am grateful to Caltech for its hospitality as I completed this article, and I gratefully acknowledge financial support under NSF Grant SES-0617753.

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

VII.5	Global Convergence	22
VII.6	Nonconvergence	29
VIII	Stochastic Dynamics	36
VIII.1	Deterministic Approximation	36
VIII.2	Convergence to Equilibria and to Better-Reply Closed Sets	37
VIII.3	Stochastic Stability and Equilibrium Selection	38
IX	Local Interaction	44
IX.1	Stochastic Stability and Equilibrium Selection Revisited	45
IX.2	Complex Spatial Dynamics	46
X	Applications	49
XI	Future Directions	50
XII	Bibliography	51

Glossary

Deterministic evolutionary dynamic: A deterministic evolutionary dynamic is a rule for assigning population games to ordinary differential equations describing the evolution of behavior in the game. Deterministic evolutionary dynamics can be derived from *revision protocols*, which describe choices (in economic settings) or births and deaths (in biological settings) on an agent-by-agent basis.

Evolutionarily stable strategy (ESS): In a symmetric normal form game, an evolutionarily stable strategy is a (possibly mixed) strategy with the following property: a population in which all members play this strategy is resistant to invasion by a small group of mutants who play an alternative mixed strategy.

Normal form game: A normal form game is a strategic interaction in which each of n players chooses a strategy and then receives a payoff that depends on all agents' choices of strategy. In a *symmetric two-player normal form game*, the two players choose from the same set of strategies, and payoffs only depend on own and opponent's choices, not on a player's identity.

Population game: A population game is a strategic interaction among one or more large populations of agents. Each agent's payoff depends on his own choice of strategy and the distribution of others' choices of strategies. One can generate a population game from a

normal form game by introducing random matching; however, many population games of interest, including congestion games, do not take this form.

Replicator dynamic: The replicator dynamic is a fundamental deterministic evolutionary dynamic for games. Under this dynamic, the percentage growth rate of the mass of agents using each strategy is proportional to the excess of the strategy's payoff over the population's average payoff. The replicator dynamic can be interpreted biologically as a model of natural selection, and economically as a model of imitation.

Revision protocol: A revision protocol describes both the timing and the results of agents' decisions about how to behave in a repeated strategic interaction. Revision protocols are used to derive both deterministic and stochastic evolutionary dynamics for games.

Stochastically stable state: In models of stochastic evolution in games described by irreducible Markov processes, a population state is stochastically stable if it retains positive weight in the process's stationary distribution as the level of noise in agents' choices approaches zero, or as the population size approaches infinity.

I. Definition and Importance

Evolutionary game theory studies the behavior of large populations of agents who repeatedly engage in strategic interactions. Changes in behavior in these populations are driven either by natural selection via differences in birth and death rates, or by the application of myopic decision rules by individual agents.

The birth of evolutionary game theory is marked by the publication of a series of papers by mathematical biologist John Maynard Smith [137, 138, 140]. Maynard Smith adapted the methods of traditional game theory [151, 215], which were created to model the behavior of rational economic agents, to the context of biological natural selection. He proposed his notion of an *evolutionarily stable strategy (ESS)* as a way of explaining the existence of ritualized animal conflict.

Maynard Smith's equilibrium concept was provided with an explicit dynamic foundation through a differential equation model introduced by Taylor and Jonker [205]. Schuster and Sigmund [189], following Dawkins [58], dubbed this model the *replicator dynamic*, and recognized the close links between this game-theoretic dynamic and dynamics studied much earlier in population ecology [132, 214] and population genetics [73]. By the 1980s, evolutionary game theory was a well-developed and firmly established modeling framework in biology [106].

Towards the end of this period, economists realized the value of the evolutionary approach to game theory in social science contexts, both as a method of providing foundations for the equilibrium concepts of traditional game theory, and as a tool for selecting among equilibria in games that admit more than one. Especially in its early stages, work by economists in evolutionary game theory hewed closely to the interpretation set out by biologists, with the notion of ESS and the replicator dynamic understood as modeling natural selection in populations of agents genetically programmed to behave in specific ways. But it soon became clear that models of essentially the same form could be used to study the behavior of populations of active decision makers [50, 76, 133, 149, 167, 191]. Indeed, the two approaches sometimes lead to identical models: the replicator dynamic itself can be understood not only as a model of natural selection, but also as one of imitation of successful opponents [35, 188, 216].

While the majority of work in evolutionary game theory has been undertaken by biologists and economists, closely related models have been applied to questions in a variety of fields, including transportation science [143, 150, 173, 175, 177, 197], computer science [72, 173, 177], and sociology [34, 62, 126, 225, 226]. Some paradigms from evolutionary game theory are close relatives of certain models from physics, and so have attracted the attention of workers in this field [141, 201–203]. All told, evolutionary game theory provides a common ground for workers from a wide range of disciplines.

II. Introduction

This article offers a broad survey of the theory of evolution in games. Section III introduces normal form games, a simple and commonly studied model of strategic interaction. Section IV presents the notion of an evolutionarily stable strategy, a static definition of stability proposed for this normal form context.

Section V defines population games, a general model of strategic interaction in large populations. Section VI offers the notion of a revision protocol, an individual-level description of behavior used to define the population-level processes of central concern.

Most of the article concentrates on these population-level processes: Section VII considers deterministic differential equation models of game dynamics; Section VIII studies stochastic models of evolution based on Markov processes; and Section IX presents deterministic and stochastic models of local interaction. Section X records a range of applications of evolutionary game theory, and Section XI suggests directions for future research. Finally, Section XII offers an extensive list of primary references.

III. Normal Form Games

In this section, we introduce a very simple model of strategic interaction: the symmetric two-player normal form game. We then define some of the standard solution concepts used to analyze this model, and provide some examples of games and their equilibria. With this background in place, we turn in subsequent sections to evolutionary analysis of behavior in games.

In a *symmetric two-player normal form game*, each of the two players chooses a (pure) strategy from the finite set S ; which we write generically as $S = \{1, \dots, n\}$. The game's *payoffs* are described by the matrix $A \in \mathbf{R}^{n \times n}$. 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.

The fundamental solution concept of noncooperative game theory is Nash equilibrium [151]. We say that the pure strategy $i \in S$ is a *symmetric Nash equilibrium* of A if

$$(1) \quad A_{ii} \geq A_{ji} \text{ for all } j \in S.$$

Thus, if his opponent chooses a symmetric Nash equilibrium strategy i , a player can do no better than to choose i himself.

A stronger requirement on strategy i demands that it be superior to all other strategies regardless of the opponent's choice:

$$(2) \quad A_{ik} > A_{jk} \text{ for all } j, k \in S.$$

When condition (2) holds, we say that strategy i is *strictly dominant* in A .

Example III.1. The game below, with strategies C ("cooperate") and D ("defect"), is an instance of a *Prisoner's Dilemma*:

	C	D	
C	2	0	
D	3	1	.

(To interpret this game, note that $A_{CD} = 0$ is the payoff to cooperating when one's opponent defects.) Since $1 > 0$, defecting is a symmetric Nash equilibrium of this game. In fact, since $3 > 2$ and $1 > 0$, defecting is even a strictly dominant strategy. But since $2 > 1$, both players are better off when both cooperate than when both defect. §

In many instances, it is natural to allow players to choose *mixed* (or *randomized*) strategies. When a player chooses mixed strategy from the simplex $X = \{x \in \mathbf{R}_+^n : \sum_{i \in S} x_i = 1\}$, his behavior is stochastic: he commits to playing pure strategy $i \in S$ with probability x_i .

When either player makes a randomized choice, we evaluate payoffs by taking expectations: a player choosing mixed strategy x against an opponent choosing mixed strategy y garners an expected payoff of

$$(3) \quad x' Ay = \sum_{i \in S} \sum_{j \in S} x_i A_{ij} y_j.$$

In biological contexts, payoffs are *fitnesses*, and represent levels of reproductive success relative to some baseline level; equation (3) reflects the idea that in a large population, expected reproductive success is what matters. In economic contexts, payoffs are *utilities*: a numerical representation of players' preferences under which equation (3) captures players' choices between uncertain outcomes [215].

The notion of Nash equilibrium extends easily to allow for mixed strategies. Mixed strategy x is a *symmetric Nash equilibrium* of A if

$$(4) \quad x' Ax \geq y' Ax \text{ for all } y \in X.$$

In words, x is a symmetric Nash equilibrium if its expected payoff against itself is at least as high as the expected payoff obtainable by any other strategy y against x . Note that we can represent the pure strategy $i \in S$ using the mixed strategy $e_i \in X$, the i th standard basis vector in \mathbf{R}^n . If we do so, then definition (4) restricted to such strategies is equivalent to definition (1).

We illustrate these ideas with a few examples.

Example III.2. Consider the *Stag Hunt* game:

	H	S	
H	h	h	
S	0	s	.

Each player in the Stag Hunt game chooses between hunting hare (H) and hunting stag (S). A player who hunts hare always catches one, obtaining a payoff of $h > 0$. But hunting stag is only successful if both players do so, in which case each obtains a payoff of $s > h$. Hunting stag is potentially more profitable than hunting hare, but requires a coordinated effort.

In the Stag Hunt game, H and S (or, equivalently, e_H and e_S) are symmetric pure Nash equilibria. This game also has a symmetric mixed Nash equilibrium, namely $x^* = (x_H^*, x_S^*) = (\frac{s-h}{s}, \frac{h}{s})$. If a player's opponent chooses this mixed strategy, the player's expected payoff is h whether he chooses H, S, or any mixture between the two; in particular, x^* is a best response against itself.

To distinguish between the two pure equilibria, we might focus on the one that is *payoff dominant*, in that it achieves the higher joint payoff. Alternatively, we can concentrate on the *risk dominant* equilibrium [89], which utilizes the strategy preferred by a player who thinks his opponent is equally likely to choose either option (that is, against an opponent playing mixed strategy $(x_H, x_S) = (\frac{1}{2}, \frac{1}{2})$). In the present case, since $s > h$, equilibrium S is payoff dominant. Which strategy is risk dominant depends on further information about the value of s . If $s > 2h$, then S is risk dominant. But if $s < 2h$, H is risk dominant: evidently, payoff dominance and risk dominance need not agree. §

Example III.3. In the *Hawk-Dove game* [139], the two players are animals contesting a resource of value $v > 0$. The players choose between two strategies: display (D) or escalate (E). If both display, the resource is split; if one escalates and the other displays, the escalator claims the entire resource; if both escalate, then each player is equally likely to claim the entire resource or to be injured, suffering a cost of $c > v$ in the latter case.

The payoff matrix for the Hawk-Dove game is therefore

$$\begin{array}{cc} & \begin{array}{cc} \text{D} & \text{E} \end{array} \\ \begin{array}{c} \text{D} \\ \text{E} \end{array} & \begin{array}{|cc|} \hline \frac{1}{2}v & 0 \\ v & \frac{1}{2}(v - c) \\ \hline \end{array} . \end{array}$$

This game has no symmetric Nash equilibrium in pure strategies. It does, however, admit the symmetric mixed equilibrium $x^* = (x_D^*, x_E^*) = (\frac{c-v}{c}, \frac{v}{c})$. (In fact, it can be shown that every symmetric normal form game admits at least one symmetric mixed Nash equilibrium [151].)

In this example, our focus on symmetric behavior may seem odd: rather than randomizing symmetrically, it seems more natural for players to follow an asymmetric Nash equilibrium in which one player escalates and the other displays. But the symmetric equilibrium is the most relevant one for understanding natural selection in a populations that are randomly matched in pairwise contests—see Section IV. §

Example III.4. Consider the class of *Rock-Paper-Scissors* games:

	R	P	S	
R	0	-l	w	
P	w	0	-l	
S	-l	w	0	.

Here $w > 0$ is the benefit of winning the match and $l > 0$ the cost of losing; ties are worth 0 to both players. We call this game *good RPS* if $w > l$, so that the benefit of winning the match exceeds the cost of losing, *standard RPS* if $w = l$, and *bad RPS* if $w < l$. Regardless of the values of w and l , the unique symmetric Nash equilibrium of this game, $x^* = (x_R^*, x_P^*, x_S^*) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, requires uniform randomization over the three strategies. §

IV. Static Notions of Evolutionary Stability

In introducing game-theoretic ideas to the study of animal behavior, Maynard Smith promoted a fundamental principle: that the evolutionary success of (the genes underlying) a given behavioral trait can depend on the prevalences of all traits. It follows that natural selection among the traits can be modeled as random matching of animals to play normal form games [137–140]. Working in this vein, Maynard Smith offered a stability concept for populations of animals sharing a common behavioral trait—that of playing a particular mixed strategy in the game at hand. Maynard Smith’s concept of evolutionary stability, influenced by the work of Hamilton [87] on the evolution of sex ratios, defines such a population as stable if it is resistant to invasion by a small group of mutants carrying a different trait.

Suppose that a large population of animals is randomly matched to play the symmetric normal form game A . We call mixed strategy $x \in X$ an *evolutionarily stable strategy (ESS)* if

$$(5) \quad x'A((1 - \varepsilon)x + \varepsilon y) > y'A((1 - \varepsilon)x + \varepsilon y) \text{ for all } \varepsilon \leq \bar{\varepsilon}(y) \text{ and } y \neq x.$$

To interpret condition (5), imagine that a population of animals programmed to play mixed strategy x is invaded by a group of mutants programmed to play the alternative mixed strategy y . Equation (5) requires that regardless of the choice of y , an incumbent’s expected payoff from a random match in the post-entry population exceeds that of a mutant so long as the size of the invading group is sufficiently small.

The definition of ESS above can also be expressed as a combination of two conditions:

$$(4) \quad x'Ax \geq y'Ax \text{ for all } y \in X;$$

$$(6) \quad \text{For all } y \neq x, [x'Ax = y'Ax] \text{ implies that } [x'Ay > y'Ay].$$

Condition (4) is familiar: it requires that the incumbent strategy x be a best response to itself, and so is none other than our definition of symmetric Nash equilibrium. Condition (6) requires that if a mutant strategy y is an alternative best response against the incumbent strategy x , then the incumbent earns a higher payoff against the mutant than the mutant earns against itself.

A less demanding notion of stability can be obtained by allowing the incumbent and the mutant in condition (6) to perform equally well against the mutant:

$$(7) \quad \text{For all } y \in X, [x'Ax = y'Ax] \text{ implies that } [x'Ay \geq y'Ay].$$

If x satisfies conditions (4) and (7), it is called a *neutrally stable strategy (NSS)* [139].

Let us apply these stability notions to the games introduced in Section III. Since every ESS and NSS must be a Nash equilibrium, we need only consider whether the Nash equilibria of these games satisfy the additional stability conditions, (6) and (7).

Example IV.1. In the Prisoner's Dilemma game (Example III.1), the dominant strategy D is an ESS. §

Example IV.2. In the Stag Hunt game (Example III.2), each pure Nash equilibrium is an ESS. But the mixed equilibrium $(x_H^*, x_S^*) = (\frac{s-h}{s}, \frac{h}{s})$ is not an ESS: if mutants playing either pure strategy enter the population, they earn a higher payoff than the incumbents in the post-entry population. §

Example IV.3. In the Hawk-Dove game (Example III.3), the mixed equilibrium $(x_D^*, x_E^*) = (\frac{c-v}{c}, \frac{v}{c})$ is an ESS. Maynard Smith used this and other examples to explain the existence of ritualized fighting in animals. While an animal who escalates always obtains the resource when matched with an animal who merely displays, a population of escalators is unstable: it can be invaded by a group of mutants who display, or who merely escalate less often. §

Example IV.4. In Rock-Paper-Scissors games (Example III.4), whether the mixed equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is evolutionarily stable depends on the relative payoffs to winning and losing a match. In good RPS ($w > l$), x^* is an ESS; in standard RPS ($w = l$), x^* is a NSS but not an ESS, while in bad RPS ($w < l$), x^* is neither an ESS nor an NSS. The last case shows that neither evolutionary nor neutrally stable strategies need exist in a given game. §

The definition of an evolutionarily stable strategy has been extended to cover a wide range of strategic settings, and has been generalized in a variety of directions. Prominent among these developments are set-valued versions of ESS: in rough terms, these concepts consider a set of mixed strategies $Y \subset X$ to be stable if the no population playing a strategy

in the set can be invaded successfully by a population of mutants playing a strategy outside the set. Hines [95] provides a thorough survey of the first 15 years of research on ESS and related notions of stability; key references on set-valued evolutionary solution concepts include [15, 199, 206].

Maynard Smith's notion of ESS attempts to capture the dynamic process of natural selection using a static definition. The advantage of this approach is that his definition is often easy to check in applications. Still, more convincing models of natural selection should be explicitly dynamic models of natural selection, building on techniques from the theories of dynamical systems and stochastic processes. Indeed, this thoroughgoing approach can help us understand whether and when the ESS concept captures the notion of robustness to invasion in a satisfactory way.

The remainder of this article concerns explicitly dynamic models of behavior. These models will differ from approach taken in this section in two other important ways as well. First, rather than looking at populations whose members all play a particular mixed strategy, these models consider populations in which different members play different pure strategies. Second, instead of maintaining a purely biological point of view, our dynamic models will be equally well-suited to studying behavior in animal and human populations.

V. Population Games

Population games provide a simple and general framework for studying strategic interactions in large populations whose members play pure strategies. The simplest population games are generated by random matching in normal form games, but the population game framework allows for interactions of a more intricate nature.

We focus here on games played by a single population (i.e., games in which all agents play equivalent roles). We suppose that there is 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* $x \in X$, with x_j representing the proportion of agents choosing pure strategy j . We identify a *population game* with a continuous vector-valued payoff function $F : X \rightarrow \mathbf{R}^n$. The scalar $F_i(x)$ represents the payoff to strategy i when the population state is x .

Population state x^* is a *Nash equilibrium* of F if no agent can improve his payoff by unilaterally switching strategies. More explicitly, x^* is a Nash equilibrium if

$$(8) \quad x_i^* > 0 \text{ implies that } F_i(x) \geq F_j(x) \text{ for all } j \in S.$$

Example V.1. Suppose that the unit mass of agents are randomly matched to play the symmetric normal form game A . At population state x , the (expected) payoff to strategy i is the linear function $F_i(x) = \sum_{j \in S} A_{ij}x_j$; the payoffs to all strategies can be expressed concisely as $F(x) = Ax$. It is easy to verify that x^* is a Nash equilibrium of the population game F if and only if x^* is a symmetric Nash equilibrium of the symmetric normal form game A . §

While population games generated by random matching are especially simple, many games that arise in applications are not of this form. In the biology literature, games outside the random matching paradigm are known as *playing the field* models [139].

Example V.2. Consider the following model of highway congestion [17, 143, 166, 173]. 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 constituent links, while the delay on a link is a function of the number of agents who use that link.

Population games embodying this description are known as a *congestion games*. To define a congestion game, let Φ be the collection of links in the highway network. Each strategy $i \in S$ is a route from Home to Work, and so is identified with a set of links $\Phi_i \subseteq \Phi$. Each link ϕ is assigned a *cost function* $c_\phi : \mathbf{R}_+ \rightarrow \mathbf{R}$, whose argument is link ϕ 's *utilization level* u_ϕ :

$$u_\phi(x) = \sum_{i \in \rho(\phi)} x_i, \text{ where } \rho(\phi) = \{i \in S : \phi \in \Phi_i\}$$

The payoff of choosing route i is the negation of the total delays on the links in this route:

$$F_i(x) = - \sum_{\phi \in \Phi_i} c_\phi(u_\phi(x)).$$

Since driving on a link increases the delays experienced by other drivers on that link (i.e., since highway congestion involves *negative externalities*), cost functions in models of highway congestion are increasing; they are typically convex as well. Congestion games can also be used to model positive externalities, like the choice between different technological standards; in this case, the cost functions are decreasing in the utilization levels. §

VI. Revision Protocols

We now introduce foundations for our models of evolutionary dynamics. These foundations are built on the notion of a revision protocol, which describes both the timing and results of agents' myopic decisions about how to continue playing the game at hand [24, 35, 96, 175, 217]. Revision protocols will be used to derive both the deterministic dynamics studied in Section VII and the stochastic dynamics studied in VIII; very similar ideas underlie the local interaction models introduced in Section IX.

VI.1 Definition

Formally, a *revision protocol* is a map $\rho : \mathbf{R}^n \times X \rightarrow \mathbf{R}_+^{n \times n}$ that takes the payoff vectors π and population states x as arguments, and returns nonnegative matrices as outputs. For reasons to be made clear below, scalar $\rho_{ij}(\pi, x)$ is called the *conditional switch rate* from strategy i to strategy j .

To move from this notion to an explicit model of evolution, let us consider a population consisting of $N < \infty$ members. (A number of the analyses to follow will consider the limit of the present model as the population size N approaches infinity—see Sections VII.1, VIII.1, and VIII.3.4.) In this case, the set of feasible social states is the finite set $\mathcal{X}^N = X \cap \frac{1}{N}\mathbf{Z}^n = \{x \in X : Nx \in \mathbf{Z}^n\}$, a grid embedded in the simplex X .

A revision protocol ρ , a population game F , and a population size N define a continuous-time evolutionary process—a Markov process $\{X_t^N\}$ —on the finite state space \mathcal{X}^N . A one-size-fits-all description of this process is 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. 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 ρ_{ij}/R . If a switch occurs, the population state changes accordingly, from the old state x to a new state y that accounts for the agent's change in strategy.

While this interpretation of the evolutionary process can be applied to any revision protocol, simpler interpretations are sometimes available for protocols with additional structure. The examples to follow illustrate this point.

VI.2 Examples

VI.2.1 Imitative Protocols and Natural Selection Protocols

In economic contexts, revision protocols of the form

$$(9) \quad \rho_{ij}(\pi, x) = x_j \hat{\rho}_{ij}(\pi, x)$$

are called *imitative protocols* [35, 96, 216]. These protocols can be given a very simple interpretation: when an agent receives a revision opportunity, he chooses an opponent at random and observes her strategy. If our agent is playing strategy i and the opponent strategy j , the agent switches from i to j with probability proportional to $\hat{\rho}_{ij}$. Notice that the value of the population share x_j is not something the agent need know; this term in (9) accounts for the agent's observing a randomly chosen opponent.

Example VI.1. Suppose that after selecting an opponent, the agent imitates the opponent only if the opponent's payoff is higher than his own, doing so in this case with probability proportional to the payoff difference:

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

This protocol is known as *pairwise proportional imitation* [188]. §

Protocols of form (9) also appear in biological contexts, starting with the work of Moran [144], and revisited more recently in [153, 158]; in these cases we refer to them as *natural selection protocols*. The biological interpretation of (9) supposes that each agent is programmed to play a single pure strategy. An agent who receives a revision opportunity dies, and is replaced through asexual reproduction. The reproducing agent is a strategy j player with probability $\rho_{ij}(\pi, x) = x_j \hat{\rho}_{ij}(\pi, x)$, which is proportional both to the number of strategy j players and to some function of the prevalences and fitnesses of all strategies. Note that this interpretation requires the restriction

$$\sum_{j \in S} \rho_{ij}(\pi, x) \equiv 1.$$

Example VI.2. Suppose that payoffs are always positive, and let

$$(10) \quad \rho_{ij}(\pi, x) = \frac{x_j \pi_j}{\sum_{k \in S} x_k \pi_k}.$$

Understood as a natural selection protocol, (10) says that the probability that the repro-

ducing agent is a strategy j player is proportional to $x_j\pi_j$, the aggregate fitness of strategy j players.

In economic contexts, we can interpret (10) as an imitative protocol based on repeated sampling. When an agent's clock rings he chooses an opponent at random. If the opponent is playing strategy j , the agent imitates him with probability proportional to π_j . If the agent does not imitate this opponent, he draws a new opponent at random and repeats the procedure. §

VI.2.2 Evaluative Protocols

In the previous examples, only strategies currently in use have any chance of being chosen by a revising agent (or of being the programmed strategy of the newborn agent). Under other protocols, agents' choices are not mediated through the population's current behavior, except indirectly via the effect of behavior on payoffs. These *evaluative protocols* require agents to directly evaluate the payoffs of each strategy, rather than to indirectly evaluate them as under an imitative procedure.

Example VI.3. Suppose that choices are made according to the *logit choice rule*:

$$(11) \quad \rho_{ij}(\pi, x) = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}.$$

The interpretation of this protocol is simple. Revision opportunities arrive at unit rate. When an opportunity is received by an i player, he switches to strategy j with probability $\rho_{ij}(\pi, x)$, which is proportional to an exponential function of strategy j 's payoffs. The parameter $\eta > 0$ is called the *noise level*. If η is large, choice probabilities under the logit rule are nearly uniform. But if η is near zero, choices are optimal with probability close to one, at least when the difference between the best and second best payoff is not too small. §

Additional examples of revision protocols can be found in the next section, and one can construct new revision protocols by taking linear combinations of old ones; see [183] for further discussion.

VII. Deterministic Dynamics

Although antecedents of this approach date back to the early work of Brown and von Neumann [45], the use of differential equations to model evolution in games took root

with the introduction of the replicator dynamic by Taylor and Jonker [205], and remains an vibrant area of research; Hofbauer and Sigmund [108] and Sandholm [183] offer recent surveys. In this section, we derive a deterministic model of evolution: the *mean dynamic* generated by a revision protocol and a population game. We study this deterministic model from various angles, focusing in particular on local stability of rest points, global convergence to equilibrium, and nonconvergent limit behavior.

While the bulk of the literature on deterministic evolutionary dynamics is consistent with the approach we take here, we should mention that other specifications exist, including discrete time dynamics [5, 59, 131, 218], and dynamics for games with continuous strategy sets [41, 42, 77, 100, 159, 160] and for Bayesian population games [62, 70, 179]. Also, deterministic dynamics for extensive form games introduce new conceptual issues; see [28, 30, 51, 53, 55] and the monograph of Cressman [54].

VII.1 Mean Dynamics

As described in Section VI.1, a revision protocol ρ , a population game F , and a population size N define a Markov process $\{X_t^N\}$ on the finite state space \mathcal{X}^N . We now derive a deterministic process—the *mean dynamic*—that describes the expected motion of $\{X_t^N\}$. In Section VIII.1, we will describe formally the sense in which this deterministic process provides a very good approximation of the behavior of the stochastic process $\{X_t^N\}$, at least over finite time horizons and for large population sizes. But having noted this result, we will focus in this section on the deterministic process itself.

To compute the expected increment of $\{X_t^N\}$ over the next dt time units, recall first that each of the N agents receives revision opportunities via a rate R exponential distribution, and so expects to receive $R dt$ opportunities during the next dt time units. If the current state is x , the expected number of revision opportunities received by agents currently playing strategy i is approximately $Nx_i R dt$. Since an i player who receives a revision opportunity switches to strategy j with probability ρ_{ij}/R , the expected number of such switches during the next dt time units is approximately $Nx_i \rho_{ij} dt$. Therefore, the expected change in the number of agents choosing strategy i during the next dt time units is approximately

$$(12) \quad N \left(\sum_{j \in \mathcal{S}} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in \mathcal{S}} \rho_{ij}(F(x), x) \right) dt.$$

Dividing expression (12) by N and eliminating the time differential dt yields a differential

equation for the rate of change in the *proportion* of agents choosing strategy i :

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

Equation (M) is the *mean dynamic* (or *mean field*) generated by revision protocol ρ in population game F . The first term in (M) captures the inflow of agents to strategy i from other strategies, while the second captures the outflow of agents to other strategies from strategy i .

VII.2 Examples

We now describe some examples of mean dynamics, starting with ones generated by the revision protocols from Section VI.2. To do so, we let

$$\bar{F}(x) = \sum_{i \in S} x_i F_i(x)$$

denote the *average payoff* obtained by the members of the population, and define the *excess payoff* to strategy i , $\hat{F}_i(x) = F_i(x) - \bar{F}(x)$, to be the difference between strategy i 's payoff and the population's average payoff.

Example VII.1. In Example VI.1, we introduced the pairwise proportional imitation protocol $\rho_{ij}(\pi, x) = x_j[\pi_j - \pi_i]_+$. This protocol generates the mean dynamic

$$(13) \quad \dot{x}_i = x_i \hat{F}_i(x).$$

Equation (13) is the *replicator dynamic* [205], 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 that strategy's current excess payoff; unused strategies always remain so. There are a variety of revision protocols other than pairwise proportional imitation that generate the replicator dynamic as their mean dynamics; see [35, 96, 108, 217]. §

Example VII.2. In Example VI.2, we assumed that payoffs are always positive, and introduced the protocol $\rho_{ij} \propto x_j \pi_j$, which we interpreted both as a model of biological natural selection and as a model of imitation with repeated sampling. The resulting mean dynamic,

$$(14) \quad \dot{x}_i = \frac{x_i F_i(x)}{\sum_{k \in S} x_k F_k(x)} - x_i = \frac{x_i \hat{F}_i(x)}{\bar{F}(x)},$$

is the *Maynard Smith replicator dynamic* [139]. This dynamic only differs from the standard replicator dynamic (13) by a change of speed, with motion under (14) being relatively fast when average payoffs are relatively low. (In multipopulation models, the two dynamics are less similar, and convergence under one does not imply convergence under the other—see [183,216].) §

Example VII.3. In Example VI.3 we introduced the logit choice rule $\rho_{ij}(\pi, x) \propto \exp(\eta^{-1}\pi_j)$. The corresponding mean dynamic,

$$(15) \quad \dot{x}_i = \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} - x_i,$$

is called the *logit dynamic* [82].

If we take the noise level η to zero, then the probability with which a revising agent chooses the best response approaches one whenever the best response is unique. At such points, the logit dynamic approaches the *best response dynamic* [84]:

$$(16) \quad \dot{x} \in B^F(x) - x,$$

where

$$B^F(x) = \operatorname{argmax}_{y \in X} y'F(x)$$

defines the (*mixed*) *best response correspondence* for game F . Note that unlike the other dynamics we consider here, (16) is defined not by an ordinary differential equation, but by a differential inclusion, a formulation proposed in [97]. §

Example VII.4. Consider the evaluative protocol

$$\rho_{ij}(\pi, x) = \left[\pi_j - \sum_{k \in S} x_k \pi_k \right]_+.$$

When an agent's clock rings, he chooses a strategy at random; if that strategy's payoff is above average, the agent switches to it with probability proportional to its excess payoff. The resulting mean dynamic,

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

is called the *Brown-von Neumann-Nash (BNN) dynamic* [45]; see also [98, 176, 194, 200, 217]. §

Revision protocol	Mean dynamic	Name and source
$\rho_{ij} = x_j[\pi_j - \pi_i]_+$	$\dot{x}_i = x_i \hat{F}_i(x)$	replicator [205]
$\rho_{ij} = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}$	$\dot{x}_i = \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} - x_i$	logit [82]
$\rho_{ij} = 1_{\{j = \operatorname{argmax}_{k \in S} \pi_k\}}$	$\dot{x} \in B^F(x) - x$	best response [84]
$\rho_{ij} = [\pi_j - \sum_{k \in S} x_k \pi_k]_+$	$\dot{x}_i = [\hat{F}_i(x)]_+ - x_i \sum_{j \in S} [\hat{F}_j(x)]_+$	BNN [45]
$\rho_{ij} = [\pi_j - \pi_i]_+$	$\dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+$	Smith [197]

Table I: Five basic deterministic dynamics.

Example VII.5. Consider the evaluative revision protocol

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

When an agent's clock rings, he selects a strategy at random. If the new strategy's payoff is higher than his current strategy's payoff, he switches strategies with probability proportional to the difference between the two payoffs. The resulting mean dynamic,

$$(17) \quad \dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+,$$

is called the *Smith dynamic* [197]; see also [178]. §

We summarize these examples of revision protocols and mean dynamics in Table I.

Figure 1 presents phase diagrams for the five basic dynamics when the population is randomly matched to play standard Rock-Paper-Scissors (Example III.4). 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 reveals 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 pictured are based on evaluative protocols, which allow agents to select strategies that are currently unused. In these cases, the Nash equilibrium

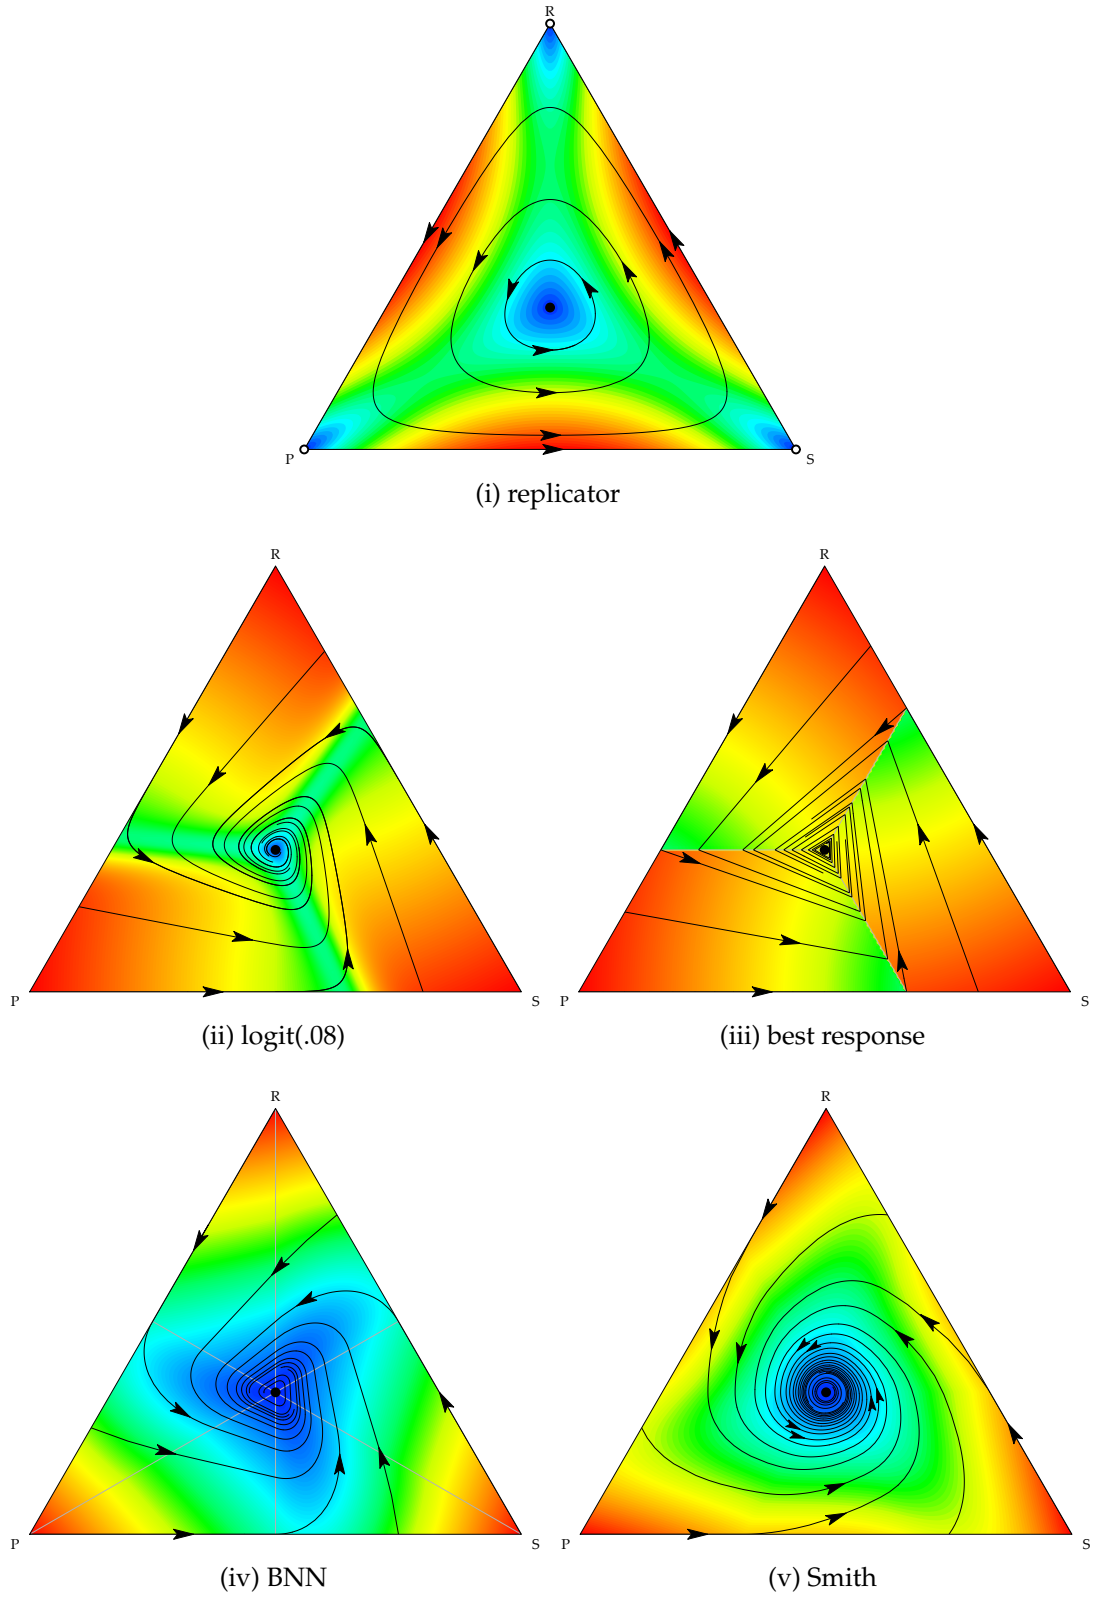


Figure 1: Five basic deterministic dynamics in standard Rock-Paper-Scissors. Colors represent speeds: red is fastest, blue is slowest

is the sole rest point, and attracts solutions from all initial conditions. (In 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 [101, 104].) Under the logit and best response dynamics, solution trajectories quickly change direction and then accelerate when the best response to the population state changes; under the BNN and especially the Smith dynamic, solutions approach the Nash equilibrium in a less angular fashion.

VII.3 Evolutionary Justification of Nash Equilibrium

One of the goals of evolutionary game theory is to justify the prediction of Nash equilibrium play. For this justification to be convincing, it must be based on a model that makes only mild assumptions about agents' knowledge about one another's behavior. This sentiment can be captured by introducing two desiderata for revision protocols:

- (C) *Continuity*: ρ is Lipschitz continuous.
- (SD) *Scarcity of data*: ρ_{ij} only depends on π_i, π_j , and x_j .

Continuity (C) asks that revision protocols depend continuously on their inputs, so that small changes in aggregate behavior do not lead to large changes in players' responses. *Scarcity of data* (SD) demands that the conditional switch rate from strategy i to strategy j only depend on the payoffs of these two strategies, so that agents need only know those facts that are most germane to the decision at hand [178]. (The dependence of ρ_{ij} on x_j is included to allow for dynamics based on imitation.) Protocols that respect these two properties do not make unrealistic demands on the amount of information that agents in an evolutionary model possess.

Our two remaining desiderata impose restrictions on mean dynamics $\dot{x} = V^F(x)$, linking the evolution of aggregate behavior to incentives in the underlying game.

- (NS) *Nash stationarity*: $V^F(x) = \mathbf{0}$ if and only if $x \in NE(F)$.
- (PC) *Positive correlation*: $V^F(x) \neq \mathbf{0}$ implies that $V^F(x)'F(x) > 0$.

Nash stationarity (NS) is a restriction on stationary states: it asks that the rest points of the mean dynamic be precisely the Nash equilibria of the game being played. *Positive correlation* (PC) is a restriction on disequilibrium adjustment: it requires that away from rest points, strategies' growth rates be positively correlated with their payoffs. Condition (PC) is among the weakest of the many conditions linking growth rates of evolutionary dynamics and payoffs in the underlying game; for alternatives, see [76, 110, 149, 162, 170, 173, 200].

Dynamic	Family	(C)	(SD)	(NS)	(PC)
replicator	imitation	yes	yes	no	yes
best response		no	yes*	yes*	yes*
logit	perturbed best response	yes	yes*	no	no
BNN	excess payoff	yes	no	yes	yes
Smith	pairwise comparison	yes	yes	yes	yes

Table II: Families of deterministic evolutionary dynamics and their properties; yes* indicates that a weaker or alternate form of the property is satisfied.

In Table II, we report how the the five basic dynamics fare under these criteria. For the purposes of justifying the Nash prediction, the most important row in the table is the last one, which reveals that the Smith dynamic satisfies all four desiderata at once: while the revision protocol for the Smith dynamic (see Example VII.5) requires only limited information on the part of the agents who employ it, this information is enough to ensure that rest points of the dynamic and Nash equilibria coincide.

In fact, the dynamics introduced above can be viewed as members of families of dynamics that are based on similar revision protocols and that have similar qualitative properties. For instance, the Smith dynamic is a member of the family of *pairwise comparison* dynamics [178], under which agents only switch to strategies that outperform their current choice. For this reason, the exact functional forms of the previous examples are not essential to establishing the properties noted above.

In interpreting these results, it is important to remember that Nash stationarity only concerns the rest points of a dynamic; it says nothing about whether a dynamic will converge to Nash equilibrium from an arbitrary initial state. The question of convergence is addressed in Sections VII.5 and VII.6. There we will see that in some classes of games, general guarantees of convergence can be obtained, but that there are some games in which no reasonable dynamic converges to equilibrium.

VII.4 Local Stability

Before turning to the global behavior of evolutionary dynamics, we address the question of local stability. As we noted at the onset, an original motivation for introducing game dynamics was to provide an explicitly dynamic foundation for Maynard Smith's notion of ESS [205]. Some of the earliest papers on evolutionary game dynamics [105,224] established that being an ESS is a sufficient condition for asymptotically stability under

the replicator dynamic, but that it is not a necessary condition. It is curious that this connection obtains despite the fact that ESS is a stability condition for a population whose members all play the same mixed strategy, while (the usual version of) the replicator dynamic looks at populations of agents choosing among different pure strategies.

In fact, the implications of ESS for local stability are not limited to the replicator dynamic. Suppose that the symmetric normal form game A admits a symmetric Nash equilibrium that places positive probability on each strategy in S . One can show that this equilibrium is an ESS if and only if the payoff matrix A is negative definite with respect to the tangent space of the simplex:

$$(18) \quad z'Az < 0 \text{ for all } z \in TX = \left\{ \hat{z} \in \mathbf{R}^n : \sum_{i \in S} \hat{z}_i = 0 \right\}.$$

Condition (18) and its generalizations imply local stability of equilibrium not only under the replicator dynamic, but also under a wide range of other evolutionary dynamics: see [52, 98, 99, 102, 111, 179] for further details.

The papers cited above use linearization and Lyapunov function arguments to establish local stability. An alternative approach to local stability analysis, via index theory, allows one to establish restrictions on the stability properties of all rest points at once—see [60].

VII.5 Global Convergence

While analyses of local stability reveal whether a population will return to equilibrium after a small disturbance, they do not tell us whether the population will approach equilibrium from an arbitrary disequilibrium state. To establish such global convergence results, we must restrict attention to classes of games defined by certain interesting payoff structures. These structures appear in applications, lending strong support for the Nash prediction in the settings where they arise.

VII.5.1 Potential Games

A *potential game* [17, 106, 143, 166, 173, 181] is a game that admits a *potential function*: a scalar valued function whose gradient describes the game's payoffs. In a full potential game $F : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ (see [181]), all information about incentives is captured by the potential function $f : \mathbf{R}_+^n \rightarrow \mathbf{R}$, in the sense that

$$(19) \quad \nabla f(x) = F(x) \text{ for all } x \in \mathbf{R}_+^n.$$

If F is smooth, then it is a full potential game if and only if it satisfies *full externality symmetry*:

$$(20) \quad \frac{\partial F_i}{\partial x_j}(x) = \frac{\partial F_j}{\partial x_i}(x) \text{ for all } i, j \in S \text{ and } x \in \mathbf{R}_+^n.$$

That is, the effect on the payoff to strategy i of adding new strategy j players always equals the effect on the payoff to strategy j of adding new strategy i players.

Example VII.6. Suppose a single population is randomly matched to play the symmetric normal form game $A \in \mathbf{R}^{n \times n}$, generating the population game $F(x) = Ax$. We say that A exhibits *common interests* if the two players in a match always receive the same payoff. This means that $A_{ij} = A_{ji}$ for all i and j , or, equivalently, that the matrix A is symmetric. Since $DF(x) = A$, this is precisely what we need for F to be a full potential game. The full potential function for F is $f(x) = \frac{1}{2}x'Ax$, which is one-half of the average payoff function $\bar{F}(x) = \sum_{i \in S} x_i F_i(x) = x'Ax$. The common interest assumption defines a fundamental model from population genetics, where the common interest assumption reflects the shared fate of two genes that inhabit the same organism [73, 106, 107]. §

Example VII.7. In Example V.2, we introduced congestion games, a basic model of network congestion. To see that these games are potential games, observe that an agent taking path $j \in S$ affects the payoffs of agents choosing path $i \in S$ through the marginal increases in congestion on the links $\phi \in \Phi_i \cap \Phi_j$ that the two paths have in common. But since the marginal effect of an agent taking path i on the payoffs of agents choosing path j is identical, full externality symmetry (20) holds:

$$\frac{\partial F_i}{\partial x_j}(x) = - \sum_{\phi \in \Phi_i \cap \Phi_j} c'_\phi(u_\phi(x)) = \frac{\partial F_j}{\partial x_i}(x).$$

In congestion games, the potential function takes the form

$$f(x) = - \sum_{\phi \in \Phi} \int_0^{u_\phi(x)} c_\phi(z) dz,$$

and so is typically unrelated to aggregate payoffs,

$$\bar{F}(x) = \sum_{i \in S} x_i F_i(x) = - \sum_{\phi \in \Phi} u_\phi(x) c_\phi(u_\phi(x)).$$

However, potential is proportional to aggregate payoffs if the cost functions c_ϕ are all

monomials of the same degree [56,173]. §

Population state x is a Nash equilibrium of the potential game F if and only if it satisfies the Kuhn-Tucker first order conditions for maximizing the potential function f on the simplex X [17,173]. Furthermore, it is simple to verify that any dynamic $\dot{x} = V^F(x)$ satisfying positive correlation (PC) ascends the potential function:

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

It then follows from classical results on Lyapunov functions that any dynamic satisfying (PC) converges to a connected set of rest points. If the dynamic also satisfies Nash stationarity (NS), these sets consist entirely of Nash equilibria. Thus, in potential games, very mild conditions on agents' adjustment rules are sufficient to justify the prediction of Nash equilibrium play.

In the case of the replicator dynamic, one can say more. On the interior of the simplex X , the replicator dynamic for the potential game F is a *gradient system* for the potential function f (i.e., it always ascends f in the direction of maximum increase). However, this is only true after one introduces an appropriate Riemannian metric on X [123,192]. An equivalent statement of this result, due to Akin [2], is that the replicator dynamic is the gradient system for f under the usual Euclidean metric if we stretch the state space X onto the radius 2 sphere. This stretching is accomplished using the *Akin transformation* $H_i(x) = 2\sqrt{x_i}$, which emphasizes changes in the use of rare strategies relative to changes in the use of common ones [2,4,185]. (There is also a dynamic that generates the gradient system for f on X under the usual metric: the so-called *projection dynamic* [130,150,185].)

Example VII.8. Consider evolution in 123 Coordination:

$$\begin{array}{c} \begin{array}{ccc} & 1 & 2 & 3 \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \boxed{\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{array}} & . \end{array}$$

Figure 2(i) presents a phase diagram of the replicator dynamic on its natural state space X , drawn atop of a contour plot of the potential function $f(x) = \frac{1}{2}((x_1)^2 + 2(x_2)^2 + 3(x_3)^2)$. Evidently, all solution trajectories ascend this function and converge to one of the seven symmetric Nash equilibria, with trajectories from all but a measure zero set of initial conditions converging to one of the three pure equilibria.

Figure 2(ii) presents another phase diagram for the replicator dynamic, this time after the solution trajectories and the potential function have been transported to the surface of

the radius 2 sphere using the Akin transformation. In this case, solutions cross the level sets of the potential function orthogonally, moving in the direction that increases potential most quickly. §

VII.5.2 Stable Games

A population game F is a *stable game* [102] if

$$(21) \quad (y - x)'(F(y) - F(x)) \leq 0 \text{ for all } x, y \in X.$$

If the inequality in (21) always holds strictly, then F is a *strictly stable game*.

If F is smooth, then F is a stable game if and only if it satisfies *self-defeating externalities*:


$$(22) \quad z'DF(x)z \leq 0 \text{ for all } z \in TX \text{ and } x \in X,$$

where $DF(x)$ is the derivative of $F : X \rightarrow \mathbf{R}^n$ at x . This condition requires that the improvements in the payoffs of strategies to which revising agents are switching are always exceeded by the improvements in the payoffs of strategies which revising agents are abandoning.


Example VII.9. The symmetric normal form game A is *symmetric zero-sum* if A is skew-symmetric (i.e., if $A = -A'$), so that the payoffs of the matched players always sum to zero. (An example is provided by the standard Rock-Paper-Scissors game (Example III.4).) Under this assumption, $z'Az = 0$ for all $z \in \mathbf{R}^n$; thus, the population game generated by random matching in A , $F(x) = Ax$, is a stable game that is not strictly stable. §

Example VII.10. Suppose that A satisfies the interior ESS condition (18). Then (22) holds strictly, so $F(x) = Ax$ is a strictly stable game. Examples satisfying this condition include the Hawk-Dove game (Example III.3) and any good Rock-Paper-Scissors game (Example III.4). §

Example VII.11. A *war of attrition* [33] is a symmetric normal form game in which strategies represent amounts of time committed to waiting for a scarce resource. If the two players choose times i and $j > i$, then the j player obtains the resource, worth v , while both players pay a cost of c_i : once the first player leaves, the other seizes the resource immediately. If both players choose time i , the resource is split, so payoffs are $\frac{v}{2} - c_i$ each. It can be shown that for any resource value $v \in \mathbf{R}$ and any increasing cost vector $c \in \mathbf{R}^n$, random matching in a war of attrition generates a stable game [102]. §



(i) on the original state space



(ii) on the transformed state space

Figure 2: The replicator dynamic in 123 Coordination. Colors represent the value of the game's potential function.

Dynamic	Lyapunov function for stable games
replicator	$H_{x^*}(x) = \sum_{i \in S(x^*)} x_i^* \log \frac{x_i^*}{x_i}$
logit	$\tilde{G}(x) = \max_{y \in \text{int}(X)} (y' \hat{F}(x) - \eta \sum_{i \in S} y_i \log y_i) + \eta \sum_{i \in S} x_i \log x_i$
best response	$G(x) = \max_{i \in S} \hat{F}_i(x)$
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 III: Lyapunov functions for five basic deterministic dynamics in stable games.

The flavor of the self-defeating externalities condition (22) suggests that obedience of incentives will push the population toward some “central” equilibrium state. In fact, the set of Nash equilibria of a stable game is always convex, and in the case of strictly stable games, equilibrium is unique. Moreover, it can be shown that the replicator dynamic converges to Nash equilibrium from all interior initial conditions in any strictly stable game [4, 105, 224], and that the evaluative dynamics introduced above converge to Nash equilibrium from all initial conditions in all stable games, strictly stable or not [98, 102, 104, 197]. In each case, the proof of convergence is based on the construction of a Lyapunov function that solutions of the relevant dynamic descend. The Lyapunov functions for the five basic dynamics are presented in Table III.

Interestingly, the convergence results for evaluative dynamics are not restricted to the dynamics listed in Table III, but extend to other dynamics in the same families (cf Table II). But compared to the conditions for convergence in potential games, the conditions for convergence in stable games demand additional structure on the adjustment process [102].

VII.5.3 Perturbed Best Response Dynamics in Supermodular Games

Supermodular games are defined by the property that higher choices by one’s opponents (with respect to the natural ordering on $S = \{1, \dots, n\}$) make one’s own higher strategies look relatively more desirable. Let $\Sigma \in \mathbf{R}^{(n-1) \times n}$ satisfy Σ_{ij} equal 1 if $j > i$ and equal 0 otherwise, so that $\Sigma x \in \mathbf{R}^{n-1}$ is the “decumulative distribution function” corresponding to the “density function” x . The population game F is a *supermodular game* if it

exhibits *strategic complementarities*:

$$(23) \quad \text{If } \Sigma y \geq \Sigma x, \text{ then } F_{i+1}(y) - F_i(y) \geq F_{i+1}(x) - F_i(x) \text{ for all } i < n \text{ and } x \in X.$$

If F is smooth, condition (23) is equivalent to

$$(24) \quad \frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) \geq 0 \text{ for all } i, j < n \text{ and } x \in X.$$

Example VII.12. Consider this model of *search with positive externalities*. A population of agents choose levels of search effort in $S = \{1, \dots, n\}$. The payoff to choosing effort i is

$$F_i(x) = m(i) b(a(x)) - c(i),$$

where $a(x) = \sum_{k \leq n} kx_k$ is the aggregate search effort, b is some increasing benefit function, m is an increasing multiplier function, and c is an arbitrary cost function. Notice that the benefits from searching are increasing in both own search effort and in the aggregate search effort. It is easy to check that F is a supermodular game. §

Complementarity condition (23) implies that the agents' best response correspondence is monotone in the stochastic dominance order, which in turn ensures the existence of minimal and maximal Nash equilibria [207]. One can take advantage of the monotonicity of best responses in studying evolutionary dynamics by appealing to the theory of monotone dynamical systems [196]. To do so, one needs to focus on dynamics that respect the monotonicity of best responses and that also are smooth, so that the theory of monotone dynamics can be applied. It turns out that the logit dynamic satisfies these criteria; so does any perturbed best response dynamic defined in terms of stochastic payoff perturbations. In supermodular games, these dynamics define cooperative differential equations; consequently, solutions of these dynamics from almost every initial condition converge to an approximate Nash equilibrium [104].

VII.5.4 Imitation Dynamics in Dominance Solvable Games

Suppose that in the population game F , strategy i is strictly dominated by strategy j : $F_i(x) < F_j(x)$ for all $x \in X$. Consider the evolution of behavior under the replicator dynamic (13). Since for this dynamic we have

$$\frac{d x_i}{d t x_j} = \frac{\dot{x}_i x_j - \dot{x}_j x_i}{(x_j)^2} = \frac{x_i \hat{F}_i(x) x_j - x_j \hat{F}_j(x) x_i}{(x_j)^2} = \frac{x_i}{x_j} (\hat{F}_i(x) - \hat{F}_j(x)),$$

solutions from every interior initial condition converge to the face of the simplex where the dominated strategy is unplayed [3]. It follows that the replicator dynamic converges in games with a strictly dominant strategy, and by iterating this argument, one can show that this dynamic converges to equilibrium in any game that can be solved by iterative deletion of strictly dominated strategies. In fact, this argument is not specific to the replicator dynamic, but can be shown to apply to a range of dynamics based on imitation [110,170]. Even in games which are not dominance solvable, arguments of a similar flavor can be used to restrict the long run behavior of imitative dynamics to better-reply closed sets [162]; see Section VIII.2 for a related discussion.

While the analysis here has focused on imitative dynamics, it is natural to expect the results described above to extend to any reasonable evolutionary dynamic. But we will see in Section VII.6.4 that this is not the case: the elimination of dominated strategies that obtains under imitative dynamics is the exception, not the rule.

VII.6 Nonconvergence

The previous section revealed that when certain global structural conditions on payoffs are satisfied, one can establish global convergence to equilibrium under various classes of evolutionary dynamics. Of course, if these conditions are not met, convergence cannot be guaranteed. In this section, we offer examples to illustrate some of the possibilities for nonconvergent limit behavior.

VII.6.1 Conservative Properties of the Replicator Dynamic in Zero-Sum Games

In Section VII.5.2, we noted that in strictly stable games, the replicator dynamic converges to Nash equilibrium from all interior initial conditions. To prove this, one shows that interior solutions descend the function

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

until converging to its minimizer, the unique Nash equilibrium x^* .

Now, random matching in a symmetric zero-sum game generates a population game that is stable, but not strictly stable (Example VII.9). In this case, for each Nash equilibrium x^* , the function H_{x^*} is a constant of motion for the replicator dynamic: its value is fixed along every interior solution trajectory.

Example VII.13. Suppose that agents are randomly matched to play the symmetric zero-sum game A , given by




Figure 3: Solutions of the replicator dynamic in a zero-sum game. The solutions pictured lie on the level set $H_{x^*}(x) = .58$.

$$\begin{array}{c}
 \begin{array}{cccc}
 & 1 & 2 & 3 & 4 \\
 1 & 0 & -1 & 0 & 1 \\
 2 & 1 & 0 & -1 & 0 \\
 3 & 0 & 1 & 0 & -1 \\
 4 & -1 & 0 & 1 & 0
 \end{array}
 \end{array}
 .$$

The Nash equilibria of $F(x) = Ax$ are the points on the line segment NE connecting states $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $(0, \frac{1}{2}, 0, \frac{1}{2})$, a segment that passes through the barycenter $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Figure 3 shows solutions to the replicator dynamic that lie on the level set $H_{x^*}(x) = .58$. Evidently, each of these solutions forms a closed orbit. §

Although solution trajectories of the replicator dynamic do not converge in zero-sum games, it can be proved that the time average of each solution trajectory converges to Nash equilibrium [190].

The existence of a constant of motion is not the only conservative property enjoyed by replicator dynamics for symmetric zero-sum games: these dynamics are also volume preserving after an appropriate change of speed or change of measure [5,96].

VII.6.2 Games with Nonconvergent Dynamics

The conservative properties described in the previous section have been established only for the replicator dynamic (and its distant relative, the projection dynamic [185]).

Starting with Shapley [193], many researchers have sought to construct games in which large classes of evolutionary dynamics fail to converge to equilibrium.

Example VII.14. Suppose that players are randomly matched to play the following symmetric normal form game [107, 109]:

	1	2	3	4	
1	0	0	-1	ε	
2	ε	0	0	-1	
3	-1	ε	0	0	
4	0	-1	ε	0	.

When $\varepsilon = 0$, the payoff matrix $A^\varepsilon = A^0$ is symmetric, so F^0 is a potential game with potential function $f(x) = \frac{1}{2}x'A^0x = -x_1x_3 - x_2x_4$. The function f attains its minimum of $-\frac{1}{4}$ at states $v = (\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $w = (0, \frac{1}{2}, 0, \frac{1}{2})$, has a saddle point with value $-\frac{1}{8}$ at the Nash equilibrium $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, and attains its maximum of 0 along the closed path of Nash equilibria γ consisting of edges $\overline{e_1e_2}$, $\overline{e_2e_3}$, $\overline{e_3e_4}$, and $\overline{e_4e_1}$.


Let $\dot{x} = V^F(x)$ be an evolutionary dynamic that satisfies Nash stationarity (NS) and positive correlation (PC), and that is based on a revision protocol that is continuous (C). If we apply this dynamic to game F^0 , then the foregoing discussion implies that all solutions to $\dot{x} = V^{F^0}(x)$ whose initial conditions ξ satisfy $f(\xi) > -\frac{1}{8}$ converge to γ . The Smith dynamic for F^0 is illustrated in Figure 4(i).

Now consider the same dynamic for the game F^ε , where $\varepsilon > 0$. By continuity (C), the attractor γ of V^{F^0} continues to an attractor γ^ε of V^{F^ε} whose basin of attraction approximates that of γ under $\dot{x} = V^{F^0}(x)$ (Figure 4(ii)). But since the unique Nash equilibrium of F^ε is the barycenter x^* , this proves that solutions from most initial conditions converge to an attractor far from any Nash equilibrium. §


Other examples of games in which many dynamics fail to converge include monocyclic games [22, 83, 97, 106], Mismatching Pennies [91, 116], and the hypnodisk game [103]. These examples demonstrate that there is no evolutionary dynamic that converges to Nash equilibrium regardless of the game at hand, and so that in general, analyses of long run behavior should not restrict attention to equilibria alone.

VII.6.3 Chaotic Dynamics

We have seen that deterministic evolutionary game dynamics can follow closed orbits and approach limit cycles. We now show that they also can behave chaotically.



(i)



(ii)

Figure 4: Solutions of the Smith dynamic in (i) the potential game F^0 ; (ii) the perturbed potential game F^ϵ , $\epsilon = \frac{1}{10}$.




Figure 5: Chaotic behavior under the replicator dynamic.

Example VII.15. Consider evolution under the replicator dynamic when agents are randomly matched to play the symmetric normal form game below [13, 195], whose lone interior Nash equilibrium is the barycenter $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$:

	1	2	3	4	
1	0	-12	0	22	
2	20	0	0	-10	
3	-21	-4	0	35	
4	10	-2	2	0	.

Figure 5 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 the game in Example VII.15 has a complicated payoff structure, in multipopulation contexts one can find chaotic evolutionary dynamics in very simple games [187].

VII.6.4 Survival of Dominated Strategies

In Section VII.5.4, we saw that dynamics based on imitation eliminate strictly dominated strategies along solutions from interior initial conditions. While this result seems unsurprising, it is actually extremely fragile: [25, 103] prove that dynamics that satisfy continuity (C), Nash stationarity (NS), and positive correlation (PC) and that are not based exclusively on imitation must fail to eliminate strictly dominated strategies in some games. Thus, evolutionary support for a basic rationality criterion is more tenuous than the results for imitative dynamics suggest.

Example VII.16. Figure 6(i) presents the Smith dynamic for “bad RPS with a twin”:


	R	P	S	T	
R	0	-2	1	1	
P	1	0	-2	-2	
S	-2	1	0	0	
T	-2	1	0	0	.

The Nash equilibria of this game are the states on line segment $NE = \{x^* \in X : x^* = (\frac{1}{3}, \frac{1}{3}, c, \frac{1}{3} - c)\}$, which is a repellor under the Smith dynamic. Under this dynamic, strategies gain players at rates that depend on their payoffs, but lose players at rates proportional to their current usage levels. It follows that when the dynamics are not at rest, the proportions of players choosing strategies 3 and 4 become equal, so that the dynamic approaches the plane $P = \{x \in X : x_3 = x_4\}$ on which the twins receive equal weight. Since A is based on bad RPS, solutions on the plane P approach a closed orbit away from any Nash equilibrium.


Figure 6(ii) presents the Smith dynamic in “bad RPS with a feeble twin”,

	R	P	S	T	
R	0	-2	1	1	
P	1	0	-2	-2	
S	-2	1	0	0	
T	$-2 - \varepsilon$	$1 - \varepsilon$	$-\varepsilon$	$-\varepsilon$,

with $\varepsilon = \frac{1}{10}$. Evidently, the attractor from Figure 6(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



(i) bad RPS with a twin



(ii) bad RPS with a feeble twin

Figure 6: The Smith dynamic in two games.

of X , the strictly dominated strategy $Twin$ is always played with probabilities bounded far away from zero. §

VIII. Stochastic Dynamics

In Section VI we defined the stochastic evolutionary process $\{X_t^N\}$ in terms of a simple model of myopic individual choice. We then turned in Section VII to the study of deterministic dynamics, which we claimed could be used to approximate the stochastic process $\{X_t^N\}$ over finite time spans and for large population sizes. In this section, we turn our attention to the stochastic process $\{X_t^N\}$ itself. After offering a formal version of the deterministic approximation result, we investigate the long run behavior of $\{X_t^N\}$, focusing on the questions of convergence to equilibrium and selection among multiple stable equilibria.

VIII.1 Deterministic Approximation

In Section VI, we defined the Markovian evolutionary process $\{X_t^N\}$ from a revision protocol ρ , a population game F , and a finite population size N . In Section VII.1, we argued that the expected motion of this process is captured by the mean dynamic

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

The basic link between the Markov process $\{X_t^N\}$ and its mean dynamic (M) is provided by *Kurtz's Theorem* [127], variations and extensions of which have been offered in a number of game-theoretic contexts [24,29,43,44,175,204]. Consider the sequence of Markov processes $\{\{X_t^N\}_{t \geq 0}\}_{N=N_0}^\infty$, supposing that the initial conditions X_0^N converge to $x_0 \in X$. Let $\{x_t\}_{t \geq 0}$ be the solution to the mean dynamic (M) starting from x_0 . Kurtz's Theorem tells us that for each finite time horizon $T < \infty$ and error bound $\varepsilon > 0$, we have that

$$(25) \quad \lim_{N \rightarrow \infty} P \left(\sup_{t \in [0, T]} |X_t^N - x_t| < \varepsilon \right) = 1.$$

Thus, 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 (M) through time T . 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 for all t between 0 and T (Figure 7).




Figure 7: Deterministic approximation of the Markov process $\{X_t^N\}$.

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. Still, at most population states the expected number of revision opportunities that arrive during the brief time interval $I = [t, t + dt]$ is large—in particular, of order $N dt$. Since each opportunity leads to an increment of the state of size $\frac{1}{N}$, the size of the overall change in the state during time interval I is of order dt . Thus, during this interval there are a large number of revision opportunities, each following nearly the same transition probabilities, and hence having nearly the same expected increments. 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 (M).

VIII.2 Convergence to Equilibria and to Better-Reply Closed Sets

Stochastic models of evolution can also be used to address directly the question of convergence to equilibrium [61, 78, 117, 118, 125, 143, 172, 219]. Suppose that a society of agents is randomly matched to play an (asymmetric) normal form game that is *weakly acyclic in better replies*: from each strategy profile, there exists a sequence of profitable unilateral deviations leading to a Nash equilibrium. If agents switch to strategies that do at least as well as their current one against the choices of random samples of opponents, then the society will eventually escape any better-response cycle, ultimately settling upon a Nash equilibrium.

Importantly, many classes of normal form games are weakly acyclic in better replies: these include potential games, dominance solvable games, certain supermodular games, and certain *aggregative games*, in which each agent's payoffs only depend on opponents' behavior through a scalar aggregate statistic. Thus, in all of these cases, simple stochastic better-reply procedures are certain to lead to Nash equilibrium play.

Outside these classes of games, one can narrow down the possibilities for long run behavior by looking at *better-reply closed sets*: that is, subsets of the set of strategy profiles that cannot be escaped without a player switching to an inferior strategy (cf [16,162]). Stochastic better-reply procedures must lead to a cluster of population states corresponding to a better-reply closed set; once the society enters such a cluster, it never departs.

VIII.3 Stochastic Stability and Equilibrium Selection

To this point, we used stochastic evolutionary dynamics to provide foundations for deterministic dynamics and to address the question of convergence to equilibrium. But stochastic evolutionary dynamics introduce an entirely new possibility: that of obtaining unique long-run predictions of play, even in games with multiple locally stable equilibria. This form of analysis, which we consider next, was pioneered by Foster and Young [74], Kandori, Mailath, and Rob [119], and Young [219], building on mathematical techniques due to Freidlin and Wentzell [75].

VIII.3.1 Stochastic Stability

To minimize notation, let us describe the evolution of behavior using a discrete-time Markov chain $\{X_k^{N,\varepsilon}\}_{k=0}^\infty$ on \mathcal{X}^N , where the parameter $\varepsilon > 0$ represents the level of "noise" in agents' decision procedures. The noise ensures that the Markov chain is irreducible and aperiodic: any state in \mathcal{X}^N can be reached from any other, and there is positive probability that a period passes without a change in the state.

Under these conditions, the Markov chain $\{X_k^{N,\varepsilon}\}$ admits a unique *stationary distribution*, $\mu^{N,\varepsilon}$, a measure on the state space \mathcal{X}^N that is invariant under the Markov chain:

$$\sum_{x \in \mathcal{X}^N} \mu^{N,\varepsilon}(x) P(X_{k+1}^{N,\varepsilon} = y \mid X_k^{N,\varepsilon} = x) = \mu^{N,\varepsilon}(y) \text{ for all } y \in \mathcal{X}^N.$$

The stationary distribution describes the long run behavior of the process $\{X_t^{N,\varepsilon}\}$ in two distinct ways. First, $\mu^{N,\varepsilon}$ is the *limiting distribution* of $\{X_t^{N,\varepsilon}\}$:

$$\lim_{k \rightarrow \infty} P(X_k^{N,\varepsilon} = y \mid X_0^{N,\varepsilon} = x) = \mu^{N,\varepsilon}(y) \text{ for all } x, y \in \mathcal{X}^N.$$

Second, $\mu^{N,\varepsilon}$ almost surely describes the *limiting empirical distribution* of $\{X_t^{N,\varepsilon}\}$:

$$P\left(\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} 1_{\{X_k^{N,\varepsilon} \in A\}} = \mu^{N,\varepsilon}(A)\right) = 1 \text{ for any } A \subseteq \mathcal{X}^N.$$

Thus, if most of the mass in the stationary distribution $\mu^{N,\varepsilon}$ were placed on a single state, then this state would provide a unique prediction of long run behavior.

With this motivation, consider a sequence of Markov chains $\{\{X_k^{N,\varepsilon}\}_{k=0}^{\infty}\}_{\varepsilon \in (0,\varepsilon)}$ parameterized by noise levels ε that approach zero. Population state $x \in \mathcal{X}^N$ is said to be *stochastically stable* if it retains positive weight in the stationary distributions of these Markov chains as ε becomes arbitrarily small:

$$\lim_{\varepsilon \rightarrow 0} \mu^{N,\varepsilon}(x) > 0.$$

When the stochastically stable state is unique, it offers a unique prediction of play over sufficiently long time spans.

VIII.3.2 Bernoulli Arrivals and Mutations

Following the work of Kandori, Mailath, and Rob [119–121] and other early contributors to the literature, let us consider a model of stochastic evolution based on *Bernoulli arrivals of revision opportunities* and *best responses with mutations*. The former assumption means that during each discrete time period, each agent has probability $\theta \in (0, 1]$ of receiving an opportunity to update his strategy. This assumption differs than the one we proposed in Section VI (and revisit below). The key new implication is that all agents may receive revision opportunities simultaneously. (Models that assume this directly generate similar results.) The latter assumption posits that when an agent receives a revision opportunity, he plays a best response to the current strategy distribution with probability $1 - \varepsilon$, and chooses a strategy at random with probability ε .

Example VIII.1. Suppose that a population of N agents is randomly matched to play the Stag Hunt game (Example III.2):

$$\begin{array}{cc} & \begin{array}{cc} \text{H} & \text{S} \end{array} \\ \begin{array}{c} \text{H} \\ \text{S} \end{array} & \begin{array}{|cc|} \hline h & h \\ \hline 0 & s \\ \hline \end{array} . \end{array}$$

Since $s > h > 0$, hunting hare and hunting stag are both symmetric pure equilibria; the game also admits the symmetric mixed equilibrium $x^* = (x_H^*, x_S^*) = (\frac{s-h}{s}, \frac{h}{s})$.




Figure 8: Equilibrium selection via mutation counting in Stag Hunt games.

If more than fraction x_H^* of the agents hunt hare, then hare is the unique best response, while if more than fraction x_S^* of the agents hunt stag, then stag is the unique best response. Thus, under any deterministic dynamic that respects payoffs, the mixed equilibrium x^* divides the state space into two basins of attraction, one for each of the two pure equilibria.

Now consider our stochastic evolutionary process. If the noise level ε is small, this process typically behaves like a deterministic process, moving quickly toward one of the two pure states, $e_H = (1, 0)$ or $e_S = (0, 1)$, and remaining there for some time. But since the process is ergodic, it will eventually leave the pure state it reaches first, and in fact will switch from one pure state to the other infinitely often.

To determine the stochastically stable state, we must compute and compare the “improbabilities” of these transitions. If the current state is e_H , a transition to e_S requires mutations to cause roughly Nx_S^* agents to switch to the suboptimal strategy S, sending the population into the basin of attraction of e_S ; the probability of this event is of order $\varepsilon^{Nx_S^*}$. Similarly, to transit from e_S to e_H , mutations must cause roughly $Nx_H^* = N(1 - x_S^*)$ to switch from S to H; this probability of this event is of order $\varepsilon^{N(1-x_S^*)}$.

Which of these rare events is more likely ones depends on whether x_S^* is greater than or less than $\frac{1}{2}$. If $s > 2h$, so that $x_S^* < \frac{1}{2}$, then $\varepsilon^{Nx_S^*}$ is much smaller than $\varepsilon^{N(1-x_S^*)}$ when ε is small; thus, state e_S is stochastically stable (Figure 8(i)). If instead $s < 2h$, so that $x_S^* > \frac{1}{2}$, then $\varepsilon^{N(1-x_S^*)} < \varepsilon^{Nx_S^*}$, so e_H is stochastically stable (Figure 8(ii)).

These calculations show that *risk dominance*—being the optimal response against a uniformly randomizing opponent—drives stochastic stability. In particular, when $s < 2h$, so that risk dominance and payoff dominance disagree, stochastic stability favors the former over the latter. §

This example illustrates how under Bernoulli arrivals and mutations, stochastic stability analysis is based on *mutation counting*: that is, on determining how many simultaneous mutations are required to move from each equilibrium into the basin of attraction of each other equilibrium. In games with more than two strategies, completing the argument becomes more complicated than in the example above: the analysis, typically based on the tree-analysis techniques of Freidlin and Wentzell [75], requires one to account for the relative difficulties of transitions between all pairs of equilibria. Ellison [68] develops a streamlined method of computing the stochastically stable state based on radius-coradius calculations; while this approach is not always sufficiently fine to yield a complete analysis, in the cases where it works it can be considerably simpler to apply than the Freidlin-Wentzell method.

These techniques have been employed successfully to variety of classes of games, including pure coordination games, supermodular games, games satisfying “bandwagon” properties, and games with equilibria that satisfy generalizations of risk dominance [68, 120, 121, 134]. A closely related literature uses stochastic stability as a basis for evaluating traditional solution concepts for extensive form games [90, 115, 122, 128, 152, 168, 169].

A number of authors have shown that variations on the Bernoulli arrivals and mutations model can lead to different equilibrium selection results. For instance, [165, 211] show that if choices are determined from the payoffs from a single round of matching (rather than from expected payoffs), the payoff dominant equilibrium rather than the risk dominant equilibrium is selected. If choices depend on strategies’ relative performances rather than their absolute performances, then long run behavior need not resemble a Nash equilibrium at all [26, 161, 171, 198]. Finally, if the probability of mutation depends on the current population state, then any recurrent set of the unperturbed process (e.g., any pure equilibrium of a coordination game) can be selected in the long run if the mutation rates are specified in an appropriate way [27]. This last result suggests that mistake probabilities should be provided with an explicit foundation, a topic we take up in Section VIII.3.3.

Another important criticism of the stochastic stability literature concerns the length of time needed for its predictions to become relevant [31, 67]. If the population size N is large and the mutation rate ε is small, then the probability ε^{cN} that a transition between equilibria occurs during given period is miniscule; the waiting time between transitions is thus enormous. Indeed, if the mutation rate falls over time, or if the population size grows over time, then ergodicity may fail, abrogating equilibrium selection entirely [163, 186]. These analyses suggest that except in applications with very long time horizons, the unique predictions generated by analyses of stochastic stability may be inappropriate, and that modelers would do better to focus on history-dependent predictions of the sort

provided by deterministic models. At the same time, there are frameworks in which stochastic stability becomes relevant much more quickly. The most important of these are local interaction models, which we discuss in Section IX.

VIII.3.3 Poisson Arrivals and Payoff Noise

Combining the assumption of Bernoulli arrivals of revision opportunities with that of best responses with mutations creates a model in which the probabilities of transitions between equilibria are easy to compute: one can focus on events in which large numbers of agents switch to a suboptimal strategy at once, each doing so with the same probability. But the simplicity of this argument also highlights the potency of the assumptions behind it.

An appealing alternative approach is to model stochastic evolution using *Poisson arrivals of revision opportunities* and *payoff noise* [29,31,38,39,63,135,145,209,210,222]. (One can achieve similar effects by looking at models defined in terms of stochastic differential equations; see [18,48,74,79,113].) By allowing revision opportunities to arrive in continuous time, as we did in Section VI, we ensure that agents do not receive opportunities simultaneously, ruling out the simultaneous mass revisions that drive the Bernoulli arrival model. (One can accomplish this using a discrete time model by assuming that one agent updates during each period; the resulting process is a random time change away from the Poisson arrivals model.)

Under Poisson arrivals, transitions between equilibria occur gradually, as the population works its way out of basins of attraction one agent at a time. In this context, the mutation assumption becomes particularly potent, ensuring that the probabilities of suboptimal choices do not vary with their payoff consequences. Under the alternative assumption of payoff noise, one supposes that agents play best responses to payoffs that are subject to random perturbations drawn from a fixed multivariate distribution. In this case, suboptimal choices are much more likely near basin boundaries, where the payoffs of second-best strategies are not much less than those of optimal ones, than they are at stable equilibria, where payoff differences are larger.

Evidently, assuming Poisson arrivals and payoff noise means that stochastic stability cannot be assessed by way of mutation counting. To determine the unlikelihood of escaping from an equilibrium's basin of attraction, one must not only account for the "width" of the basin of attraction (i.e., the number of suboptimal choices needed to escape it), but also for its "depth" (the unlikelihood of each of these choices). In two-strategy games this is not difficult to accomplish: since in this case the evolutionary process is a birth-and-death chain, its stationary distribution can be expressed using an explicit

formula. Beyond this case, one can employ the Freidlin and Wentzell [75] machinery, although doing so tends to be computationally demanding.

This computational burden is less in models that retain Poisson arrivals, but replace perturbed optimization with decision rules based on imitation and mutation [80]. Because agents imitate successful opponents, the population spends the vast majority of periods on the edges of the simplex, implying that the probabilities of transitions between vertices can be determined using birth-and-death chain methods [158]. As a consequence, one can reduce the problem of finding the stochastically stable state in an n strategy coordination game to that of computing the limiting stationary distribution of an n state Markov chain.

VIII.3.4 Stochastic Stability via Large Population Limits

The approach to stochastic stability followed thus far relies on small noise limits: that is, on evaluating the limit of the stationary distributions $\mu^{N,\varepsilon}$ as the noise level ε approaches zero. Binmore and Samuelson [29] argue that in the contexts where evolutionary models are appropriate, the amount of noise in agents decisions is not negligible, so that taking the low noise limit may not be desirable. At the same time, evolutionary models are intended to describe behavior in large populations, suggesting an alternative approach: that of evaluating the limit of the stationary distributions $\mu^{N,\varepsilon}$ as the population size N grows large.

In one respect, this approach complicates the analysis. When N is fixed and ε varies, each stationary distribution $\mu^{N,\varepsilon}$ is a measure on the fixed state space $\mathcal{X}^N = \{x \in X : Nx \in \mathbf{Z}^n\}$. But when ε is fixed and N varies, the state space \mathcal{X}^N varies as well, and one must introduce notions of weak convergence of probability measures in order to define stochastic stability.

But in other respects taking large population limits can make analysis simpler. We saw in Section VIII.1 that by taking the large population limit, we can approximate the finite-horizon sample paths of the stochastic evolutionary process $\{X_t^{N,\varepsilon}\}$ by solutions to the mean dynamic (M). Now we are concerned with infinite horizon behavior, but it is still reasonable to hope that the large population limit will again reduce some of our computations to a calculus problems.

As one might expect, this approach is easiest to follow in the two-strategy case, where for each fixed population size N , the evolutionary process $\{X_t^{N,\varepsilon}\}$ is a birth-and-death chain. When one takes the large population limit, the formulas for waiting times and for the stationary distribution can be evaluated using integral approximations [24,29,39,222]. Indeed, the approximations so obtained take an appealing simple form [182].

The analysis becomes more complicated beyond the two-strategy case, but certain

models have proved amenable to analysis. Building on their work for fixed population size models [80], Fudenberg and Imhof [81] characterize large population stochastic stability in models based on imitation and mutation. The population again spends nearly all periods on the edges of the simplex X , and the large population limit makes evaluating the probabilities of transitions along these edges relatively simple.

If one supposes that agents play best responses to noisy payoffs, then one must account directly for the behavior of the process $\{X_t^{N,\varepsilon}\}$ in the interior of the simplex. One possibility is to combine the deterministic approximation results from Section VIII.1 with techniques from the theory of stochastic approximation [20,21] to show that the large N limiting stationary distribution is concentrated on attractors of the mean dynamic. By combining this idea with convergence results for deterministic dynamics from Section VII.5, reference [104] shows that the limiting stationary distribution must be concentrated around equilibrium states in potential games, stable games, and supermodular games.

The results in [104] do not address the question of equilibrium selection. However, for the specific case of logit evolution in potential games, a complete characterization of the large population limit of the process $\{X_t^{N,\varepsilon}\}$ has been obtained [23]. By combining deterministic approximation results, which describe the usual behavior of the process within basins of attraction, with a large deviations analysis, which characterizes the rare escapes from basins of attraction, one can obtain a precise asymptotic formula for the large N limiting stationary distribution. This formula accounts both for the typical procession of the process along solutions of the mean dynamic, and for the rare sojourns of the process against this deterministic flow.

IX. Local Interaction

All of the game dynamics considered so far have been based implicitly on the assumption of *global interaction*: each agent's payoffs depend directly on all agents' actions. In many contexts, one expects to the contrary that interactions will be local in nature: for instance, agents may live in fixed locations and interact only with neighbors. In addition to providing a natural fit for these applications, *local interaction* models respond to some of the criticisms of the stochastic stability literature. At the same time, once one moves beyond relatively simple cases, local interaction models become exceptionally complicated, and so lend themselves to methods of analysis very different from those considered thus far.

IX.1 Stochastic Stability and Equilibrium Selection Revisited

In Section VIII.3, we saw the prediction of risk dominant equilibrium play provided by stochastic stability models subverted by the waiting-time critique: namely, that the length of time required before this equilibrium is reached may be extremely long. Ellison [67,68] shows that if interactions are local, then selection of the risk dominant equilibrium persists, and waiting times are no longer an issue.

Example IX.1. In the simplest version of Ellison's model, a population of N agents are located at N distinct positions around a circle. During each period of play, each agent plays the Stag Hunt game (Examples III.2 and VIII.1) with his two nearest neighbors, following the same action against both of his opponents. If we suppose that $s \in (h, 2h)$, so that hunting hare is the risk dominant strategy, then by definition, an agent whose neighbors play different strategies finds it optimal to choose H himself.

Now suppose that there are Bernoulli arrivals of revision opportunities, and that decisions are based on best responses with occasional rare mutations. To move from the all S state to the all H state, it is enough that a single agent mutates from S to H. This one mutation begins a chain reaction: the mutating agent's neighbors respond optimally by switching to H themselves; they are followed in this by their own neighbors; and the contagion continues until all agents choose H. Since a single mutation is always enough to spur the transition from all S to all H, the expected wait before this transition is small, even when the population is large.

In contrast, the transition back from all H to all S is extremely unlikely. Even if all but one of the agents simultaneously mutate to S, the contagion process described above will return the population to the all-H state. Thus, while the transition from all-S to all-H occurs quickly, the reverse transition takes even longer than in the global interaction setting. §

The local interaction approach to equilibrium selection has been advanced in a variety of directions: by allowing agents to choose their locations [69], or to pay a cost to choose different strategies against different opponents [86], and by basing agents' decisions on the attainment of aspiration levels [11] or on imitation of successful opponents [9,10]. A portion of this literature initiated by Blume develops connections between local interaction models in evolutionary game theory with models from statistical mechanics [36–38,124,141]. These models provide a point of departure for research on complex spatial dynamics in games, which we consider next.

IX.2 Complex Spatial Dynamics

The local interaction models described above address the questions of convergence to equilibrium and selection among multiple equilibria. In the cases where convergence and selection results obtain, behavior in these models is relatively simple, as most periods are spent with most agents coordinating on a single strategy. A distinct branch of the literature on evolution and local interaction focuses on cases with complex dynamics, where instead of settling quickly into a homogeneous, static configuration, behavior remains in flux, with multiple strategies coexisting for long periods of time.

Example IX.2. Cooperating is a dominated strategy in the Prisoner's Dilemma, and is not played in equilibrium in finitely repeated versions of this game. Nevertheless, a pair of Prisoner's Dilemma tournaments conducted by Axelrod [14] were won by the strategy Tit-for-Tat, which cooperates against cooperative opponents and defects against defectors. Axelrod's work spawned a vast literature aiming to understand the persistence of individually irrational but socially beneficial behavior.

To address this question, Nowak and May [153–157] consider a population of agents who are repeatedly matched to play the Prisoner's Dilemma

$$\begin{array}{c} \text{C} \quad \text{D} \\ \text{C} \begin{array}{|c|c|} \hline 1 & -\varepsilon \\ \hline \end{array} \\ \text{D} \begin{array}{|c|c|} \hline g & 0 \\ \hline \end{array} \end{array},$$

where the greedy payoff g exceeds 1 and $\varepsilon > 0$ is small. The agents are positioned on a two-dimensional grid. During each period, each agent plays the Prisoner's Dilemma with the eight agents in his (Moore) neighborhood. In the simplest version of the model, all agents simultaneously update their strategies at the end of each period. If an agent's total payoff that period is as high as that of any of neighbor, he continues to play the same strategy; otherwise, he switches to the strategy of the neighbor who obtained the highest payoff.

Since defecting is a dominant strategy in the Prisoner's Dilemma, one might expect the local interaction process to converge to a state at which all agents defect, as would be the case in nearly any model of global interaction. But while an agent is always better off defecting himself, he also is better off the more of his neighbors cooperate; and since evolution is based on imitation, cooperators tend to have more cooperators as neighbors than do defectors.

In Figures 9–11, we present snapshots of the local interaction process for choices of the greedy payoff g from each of three distinct parameter regions. If $g > \frac{5}{3}$ (Figure 9),




Figure 9: Local interaction in a Prisoner’s Dilemma; greedy payoff $g = 1.7$. In Figures 9-11, agents are arrayed on a 100×100 grid with periodic boundaries (i.e., a torus). Initial conditions are random with 75% cooperators and 25% defectors. Agents update simultaneously, imitating the neighbor who earned the highest payoff Blue cells represent cooperators who also cooperated last period, green cells represent new cooperators; red cells represent defectors who also defected last period, yellow cells represent new defectors. (Figures 9-11 created using VirtualLabs [92].)







Figure 10: Local interaction in a Prisoner’s Dilemma; greedy payoff $g = 1.55$.




(i) $t = 0$




(ii) $t = 5$




(iii) $t = 30$



(iv) $t = 60$



(v) $t = 90$



(vi) $t = 120$

Figure 11: Local interaction in a Prisoner's Dilemma; greedy payoff $g = 1.65$.

the process quickly converges to a configuration containing a few rectangular islands of cooperators in a sea of defectors; the exact configuration depending on the initial conditions. If instead $g < \frac{8}{5}$ (Figure 10), the process moves towards a configuration in which agents other than those in a “web” of defectors cooperate. But for $g \in (\frac{8}{5}, \frac{5}{3})$ (Figure 11), the system evolves in a complicated fashion, with clusters of cooperators and of defectors forming, expanding, disappearing, and reforming. But while the configuration of behavior never stabilizes, the proportion of cooperators appears to settle down to about .30.

The specification of the dynamics considered above, based on simultaneous updating and certain imitation of the most successful neighbor, presents a relatively favorable environment for cooperative behavior. Nevertheless, under Poisson arrivals of revision opportunities, or probabilistic decision rules, or both, cooperation can persist for very long periods of time for values of g significantly larger than 1 [154, 155]. §

The literature on complex spatial dynamics in evolutionary game models is large and rapidly growing, with the evolution of behavior in the spatial Prisoners’ Dilemma being the single most-studied environment. While analyses are typically based on simulations, analytical results have been obtained in some relatively simple settings [71, 94].

Recent work on complex spatial dynamics has considered games with three or more strategies, including Rock-Paper-Scissors games, as well as public good contribution games and Prisoner’s Dilemmas with voluntary participation. Introducing more than two strategies can lead to qualitatively novel dynamic phenomena, including large-scale spatial cycles and traveling waves [93, 202, 203]. In addition to simulations, the analysis of complex spatial dynamics is often based on approximation techniques from non-equilibrium statistical physics, and much of the research on these dynamics has appeared in the physics literature. Szabó and Fáth [201] offer a comprehensive survey of work on this topic.

X. Applications

Evolutionary game theory was created with biological applications squarely in mind: in the prehistory of the field, Fisher [73] and Hamilton [87] used game-theoretic ideas to understand the evolution of sex ratios. Maynard Smith [137–140] introduced his definition of ESS as a way of understanding ritualized animal conflicts. Since these early contributions, evolutionary game theory has been used to study a diverse array of biological questions, including mate choice, parental investment, parent-offspring conflict,

social foraging, and predator-prey systems. For overviews of research on these and other topics in biology, see [65,88].

The early development of evolutionary game theory in economics was motivated primarily by theoretical concerns: the justification of traditional game-theoretic solution concepts, and the development of methods for equilibrium selection in games with multiple stable equilibria. More recently, evolutionary game theory has been applied to concrete economic environments, in some instances as a means of contending with equilibrium selection problems, and in others to obtain an explicitly dynamic model of the phenomena of interest. Of course, these applications are most successful when the behavioral assumptions that underlie the evolutionary approach are appropriate, and when the time horizon needed for the results to become relevant corresponds to the one germane to the application at hand.

Topics in economics topics studied using the methods of evolutionary game theory range from behavior in markets [1,6–8,12,19,64,112,129,212], to bargaining and hold-up problems [32,46,57,66,164,208,220–222], to externality and implementation problems [47,49,136,174,177,180], to questions of public good provision and collective action [146–148]. The techniques described here are being applied with increasing frequency to problems of broader social science interest, including residential segregation [40,62,142,222,223,225,226] and cultural evolution [34,126], and to the study of behavior in transportation and computer networks [72,143,150,173,175,177,197]. A proliferating branch of research extends the approaches described in this article to address the evolution of structure and behavior in social networks; a number of recent books [85,114,213] offer detailed treatments of work in this domain.

XI. Future Directions

Evolutionary game theory is a maturing field; many basic theoretical issues are well understood, but many difficult questions remain. It is tempting to say that stochastic and local interaction models offer the more open terrain for further explorations. But while it is true that we know less about these models than about deterministic evolutionary dynamics, even our knowledge of the latter is limited: while dynamics on one and two dimensional state spaces, and for games satisfying a few interesting structural assumptions, are well-understood, the dynamics of behavior in the vast majority of many-strategy games are not.

The prospects for further applications of the tools of evolutionary game theory are brighter still. In economics, and in other social sciences, the analysis of mathematical

models has too often been synonymous with the computation and evaluation of equilibrium behavior. The questions of whether and how equilibrium will come to be are often ignored, and the possibility of long-term disequilibrium behavior left unmentioned. For settings in which its assumptions are tenable, evolutionary game theory offers a host of techniques for modeling the dynamics of economic behavior. The exploitation of the possibilities for a deeper understanding of human social interactions has hardly begun.

XII. Bibliography

Primary Literature

- [1] Murali Agastya. Stochastic stability in a double auction. *Games and Economic Behavior*, 48:203–222, 2004.
- [2] Ethan Akin. *The Geometry of Population Genetics*. Springer, Berlin, 1979.
- [3] Ethan Akin. Domination or equilibrium. *Mathematical Biosciences*, 50:239–250, 1980.
- [4] Ethan Akin. The differential geometry of population genetics and evolutionary games. In S. Lessard, editor, *Mathematical and Statistical Developments of Evolutionary Theory*, pages 1–93. Kluwer, Dordrecht, 1990.
- [5] Ethan Akin and Viktor Losert. Evolutionary dynamics of zero-sum games. *Journal of Mathematical Biology*, 20:231–258, 1984.
- [6] Carlos Alós-Ferrer and Ana B. Ania. The evolutionary stability of perfectly competitive behavior. *Economic Theory*, 26:497–516, 2005.
- [7] Carlos Alós-Ferrer, Ana B. Ania, and Klaus Reiner Schenk-Hoppé. An evolutionary model of Bertrand oligopoly. *Games and Economic Behavior*, 33:1–19, 2000.
- [8] Carlos Alós-Ferrer, Georg Kirchsteiger, and Markus Walzl. On the evolution of market institutions: The platform design paradox. Unpublished manuscript, University of Konstanz, 2006.
- [9] Carlos Alós-Ferrer and Simon Weidenholzer. Contagion and efficiency. Unpublished manuscript, University of Konstanz and University of Vienna, 2006.
- [10] Carlos Alós-Ferrer and Simon Weidenholzer. Imitation, local interactions, and efficiency. *Economics Letters*, 93:163–168, 2006.
- [11] Luca Anderlini and Antonella Ianni. Path dependence and learning from neighbors. *Games and Economic Behavior*, 13:141–177, 1996.

- [12] Ana B. Ania, Thomas Tröger, and Achim Wambach. An evolutionary analysis of insurance markets with adverse selection. *Games and Economic Behavior*, 40:153–184, 2002.
- [13] A. Arneodo, P. Couillet, and C. Tresser. Occurrence of strange attractors in three-dimensional Volterra equations. *Physics Letters*, 79A:259–263, 1980.
- [14] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
- [15] Dieter Balkenborg and Karl H. Schlag. Evolutionarily stable sets. *International Journal of Game Theory*, 29:571–595, 2001.
- [16] Kaushik Basu and Jörgen W. Weibull. Strategy sets closed under rational behavior. *Economics Letters*, 36:141–146, 1991.
- [17] Martin Beckmann, C. B. McGuire, and Christopher B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, 1956.
- [18] Alan W. Beggs. Stochastic evolution with slow learning. *Economic Theory*, 19:379–405, 2002.
- [19] Assaf Ben-Shoham, Roberto Serrano, and Oscar Volij. The evolution of exchange. *Journal of Economic Theory*, 114:310–328, 2004.
- [20] Michel Benaïm. Recursive algorithms, urn processes, and the chaining number of chain recurrent sets. *Ergodic Theory and Dynamical Systems*, 18:53–87, 1998.
- [21] Michel Benaïm and Morris W. Hirsch. On stochastic approximation algorithms with constant step size whose average is cooperative. *Annals of Applied Probability*, 30:850–869, 1999.
- [22] Michel Benaïm, Josef Hofbauer, and Ed Hopkins. Learning in games with unstable equilibria. Unpublished manuscript, Université de Neuchâtel, University of Vienna, and University of Edinburgh, 2006.
- [23] Michel Benaïm and William H. Sandholm. Logit evolution in potential games: Reversibility, rates of convergence, large deviations, and equilibrium selection. Unpublished manuscript, Université de Neuchâtel and University of Wisconsin, 2007.
- [24] Michel Benaïm and Jörgen W. Weibull. Deterministic approximation of stochastic evolution in games. *Econometrica*, 71:873–903, 2003.
- [25] Ulrich Berger and Josef Hofbauer. Irrational behavior in the Brown-von Neumann-Nash dynamics. *Games and Economic Behavior*, 56:1–6, 2006.
- [26] James Bergin and Dan Bernhardt. Comparative learning dynamics. *International Economic Review*, 45:431–465, 2004.

- [27] James Bergin and Barton L. Lipman. Evolution with state-dependent mutations. *Econometrica*, 64:943–956, 1996.
- [28] Ken Binmore, John Gale, and Larry Samuelson. Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, 8:56–90, 1995.
- [29] Ken Binmore and Larry Samuelson. Muddling through: Noisy equilibrium selection. *Journal of Economic Theory*, 74:235–265, 1997.
- [30] Ken Binmore and Larry Samuelson. Evolutionary drift and equilibrium selection. *Review of Economic Studies*, 66:363–393, 1999.
- [31] Ken Binmore, Larry Samuelson, and Richard Vaughan. Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35, 1995.
- [32] Ken Binmore, Larry Samuelson, and H. Peyton Young. Equilibrium selection in bargaining models. *Games and Economic Behavior*, 45:296–328, 2003.
- [33] D. T. Bishop and C. Cannings. A generalised war of attrition. *Journal of Theoretical Biology*, 70:85–124, 1978.
- [34] Alberto Bisin and Thierry Verdier. The economics of cultural transmission and the dynamics of preferences. *Journal of Economic Theory*, 97:298–319, 2001.
- [35] Jonas Björnerstedt and Jörgen W. Weibull. Nash equilibrium and evolution by imitation. In Kenneth J. Arrow et al., editors, *The Rational Foundations of Economic Behavior*, pages 155–181. St. Martin’s Press, New York, 1996.
- [36] Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [37] Lawrence E. Blume. The statistical mechanics of best reponse strategy revision. *Games and Economic Behavior*, 11:111–145, 1995.
- [38] Lawrence E. Blume. Population games. In W. B. Arthur, S. N. Durlauf, and D. A. Lane, editors, *The Economy as an Evolving Complex System II*, pages 425–460. Addison-Wesley, 1997.
- [39] Lawrence E. Blume. How noise matters. *Games and Economic Behavior*, 44:251–271, 2003.
- [40] Martin Bøg. Is segregation robust? Unpublished manuscript, Stockholm School of Economics, 2006.
- [41] Immanuel M. Bomze. Dynamical aspects of evolutionary stability. *Monatshefte für Mathematik*, 110:189–206, 1990.
- [42] Immanuel M. Bomze. Cross entropy minimization in uninvadable states of complex populations. *Journal of Mathematical Biology*, 30:73–87, 1991.

- [43] Tilman Börgers and Rajiv Sarin. Learning through reinforcement and the replicator dynamics. *Journal of Economic Theory*, 77:1–14, 1997.
- [44] Richard T. Boylan. Continuous approximation of dynamical systems with randomly matched individuals. *Journal of Economic Theory*, 66:615–625, 1995.
- [45] George W. Brown and John von Neumann. Solutions of games by differential equations. In Harold W. Kuhn and Albert W. Tucker, editors, *Contributions to the Theory of Games I*, volume 24 of *Annals of Mathematics Studies*, pages 73–79. Princeton University Press, Princeton, 1950.
- [46] Mary A. Burke and H. Peyton Young. Competition and custom in economic contracts: A case study of Illinois agriculture. *American Economic Review*, 91:559–573, 2001.
- [47] Antonio Cabrales. Adaptive dynamics and the implementation problem with complete information. *Journal of Economic Theory*, 86:159–184, 1999.
- [48] Antonio Cabrales. Stochastic replicator dynamics. *International Economic Review*, 41:451–481, 2000.
- [49] Antonio Cabrales and Giovanni Ponti. Implementation, elimination of weakly dominated strategies, and evolutionary dynamics. *Review of Economic Dynamics*, 3:247–282, 2000.
- [50] Vincent P. Crawford. An “evolutionary” interpretation of Van Huyck, Battalio, and Beil’s experimental results on coordination. *Games and Economic Behavior*, 3:25–59, 1991.
- [51] Ross Cressman. Evolutionary stability in the finitely repeated prisoner’s dilemma game. *Journal of Economic Theory*, 68:234–248, 1996.
- [52] Ross Cressman. Local stability of smooth selection dynamics for normal form games. *Mathematical Social Sciences*, 34:1–19, 1997.
- [53] Ross Cressman. Subgame monotonicity in extensive form evolutionary games. *Games and Economic Behavior*, 32:183–205, 2000.
- [54] Ross Cressman. *Evolutionary Dynamics and Extensive Form Games*. MIT Press, Cambridge, 2003.
- [55] Ross Cressman and Karl H. Schlag. On the dynamic (in)stability of backwards induction. *Journal of Economic Theory*, 83:260–285, 1998.
- [56] Stella Dafermos and Frederick T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards B*, 73:91–118, 1969.

- [57] Herbert Dawid and W. Bentley MacLeod. Hold-up and the evolution of investment and bargaining norms. *Games and Economic Behavior*, forthcoming.
- [58] Richard Dawkins. *The Selfish Gene*. Oxford University Press, Oxford, 1976.
- [59] Eddie Dekel and Suzanne Scotchmer. On the evolution of optimizing behavior. *Journal of Economic Theory*, 57:392–407, 1992.
- [60] Stefano Demichelis and Klaus Ritzberger. From evolutionary to strategic stability. *Journal of Economic Theory*, 113:51–75, 2003.
- [61] Martin Dindoš and Claudio Mezzetti. Better-reply dynamics and global convergence to Nash equilibrium in aggregative games. *Games and Economic Behavior*, 54:261–292, 2006.
- [62] Emin Dokumacı and William H. Sandholm. Schelling redux: An evolutionary model of residential segregation. Unpublished manuscript, University of Wisconsin, 2007.
- [63] Emin Dokumacı and William H. Sandholm. Stochastic evolution with perturbed payoffs and rapid play. Unpublished manuscript, University of Wisconsin, 2007.
- [64] Edward Droste, Cars Hommes, and Jan Tuinstra. Endogenous fluctuations under evolutionary pressure in Cournot competition. *Games and Economic Behavior*, 40:232–269, 2002.
- [65] Lee Alan Dugatkin and Hudson Kern Reeve, editors. *Game Theory and Animal Behavior*. Oxford University Press, Oxford, 1998.
- [66] Tore Ellingsen and Jack Robles. Does evolution solve the hold-up problem? *Games and Economic Behavior*, 39:28–53, 2002.
- [67] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61:1047–1071, 1993.
- [68] Glenn Ellison. Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–45, 2000.
- [69] Jeffrey C. Ely. Local conventions. *Advances in Economic Theory*, 2:1 (30), 2002.
- [70] Jeffrey C. Ely and William H. Sandholm. Evolution in Bayesian games I: Theory. *Games and Economic Behavior*, 53:83–109, 2005.
- [71] Ilan Eshel, Larry Samuelson, and Avner Shaked. Altruists, egoists, and hooligans in a local interaction model. *American Economic Review*, 88:157–179, 1998.
- [72] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing. Unpublished manuscript, RWTH Aachen, 2006.

- [73] Ronald A. Fisher. *The Genetical Theory of Natural Selection*. Clarendon Press, Oxford, 1930.
- [74] Dean P. Foster and H. Peyton Young. Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232, 1990. Corrigendum, 51:77–78, 1997.
- [75] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer, New York, second edition, 1998.
- [76] Daniel Friedman. Evolutionary games in economics. *Econometrica*, 59:637–666, 1991.
- [77] Daniel Friedman and Joel Yellin. Evolving landscapes for population games. Unpublished manuscript, UC Santa Cruz, 1997.
- [78] James W. Friedman and Claudio Mezzetti. Learning in games by random sampling. *Journal of Economic Theory*, 98:55–84, 2001.
- [79] Drew Fudenberg and Christopher Harris. Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57:420–441, 1992.
- [80] Drew Fudenberg and Loren A. Imhof. Imitation processes with small mutations. *Journal of Economic Theory*, 131:251–262, 2006.
- [81] Drew Fudenberg and Loren A. Imhof. Monotone imitation dynamics in large populations. *Journal of Economic Theory*, forthcoming.
- [82] Drew Fudenberg and David K. Levine. *Theory of Learning in Games*. MIT Press, Cambridge, 1998.
- [83] Andrea Gaunersdorfer and Josef Hofbauer. Fictitious play, Shapley polygons, and the replicator equation. *Games and Economic Behavior*, 11:279–303, 1995.
- [84] Itzhak Gilboa and Akihiko Matsui. Social stability and equilibrium. *Econometrica*, 59:859–867, 1991.
- [85] Sanjeev Goyal. *Connections: An Introduction to the Economics of Networks*. Princeton University Press, Princeton, forthcoming.
- [86] Sanjeev Goyal and Maarten C. W. Janssen. Non-exclusive conventions and social coordination. *Journal of Economic Theory*, 77:34–57, 1997.
- [87] William D. Hamilton. Extraordinary sex ratios. *Science*, 156:477–488, 1967.
- [88] Peter Hammerstein and Reinhard Selten. Game theory and evolutionary biology. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory*, volume 2, chapter 28, pages 929–993. Elsevier, Amsterdam, 1994.
- [89] John C. Harsanyi and Reinhard Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, 1988.

- [90] Sergiu Hart. Evolutionary dynamics and backward induction. *Games and Economic Behavior*, 41:227–264, 2002.
- [91] Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93:1830–1836, 2003.
- [92] Christoph Hauert. VirtualLabs in evolutionary game theory. Software. <http://www.univie.ac.at/virtuallabs>, 2007.
- [93] Christoph Hauert, Silvia De Monte, Josef Hofbauer, and Karl Sigmund. Volunteering as Red Queen mechanism for cooperation in public goods games. *Science*, 296:1129–1132, 2002.
- [94] A. V. M. Herz. Collective phenomena in spatially extended evolutionary games. *Journal of Theoretical Biology*, 169:65–87, 1994.
- [95] W. G. S. Hines. Evolutionary stable strategies: A review of basic theory. *Theoretical Population Biology*, 31:195–272, 1987.
- [96] Josef Hofbauer. Imitation dynamics for games. Unpublished manuscript, University of Vienna, 1995.
- [97] Josef Hofbauer. Stability for the best response dynamics. Unpublished manuscript, University of Vienna, 1995.
- [98] Josef Hofbauer. From Nash and Brown to Maynard Smith: Equilibria, dynamics, and ESS. *Selection*, 1:81–88, 2000.
- [99] Josef Hofbauer and Ed Hopkins. Learning in perturbed asymmetric games. *Games and Economic Behavior*, 52:133–152, 2005.
- [100] Josef Hofbauer, Jörg Oechssler, and Frank Riedel. Brown-von Neumann-Nash dynamics: The continuous strategy case. Unpublished manuscript, University of Vienna, 2005.
- [101] Josef Hofbauer and William H. Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70:2265–2294, 2002.
- [102] Josef Hofbauer and William H. Sandholm. Stable games. Unpublished manuscript, University of Vienna and University of Wisconsin, 2006.
- [103] Josef Hofbauer and William H. Sandholm. Survival of dominated strategies under evolutionary dynamics. Unpublished manuscript, University of Vienna and University of Wisconsin, 2006.
- [104] Josef Hofbauer and William H. Sandholm. Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, 132:47–69, 2007.

- [105] Josef Hofbauer, Peter Schuster, and Karl Sigmund. A note on evolutionarily stable strategies and game dynamics. *Journal of Theoretical Biology*, 81:609–612, 1979.
- [106] Josef Hofbauer and Karl Sigmund. *Theory of Evolution and Dynamical Systems*. Cambridge University Press, Cambridge, 1988.
- [107] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, 1998.
- [108] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society (New Series)*, 40:479–519, 2003.
- [109] Josef Hofbauer and Jeroen M. Swinkels. A universal Shapley example. Unpublished manuscript, University of Vienna and Northwestern University, 1996.
- [110] Josef Hofbauer and Jörgen W. Weibull. Evolutionary selection against dominated strategies. *Journal of Economic Theory*, 71:558–573, 1996.
- [111] Ed Hopkins. A note on best response dynamics. *Games and Economic Behavior*, 29:138–150, 1999.
- [112] Ed Hopkins and Robert M. Seymour. The stability of price dispersion under seller and consumer learning. *International Economic Review*, 43:1157–1190, 2002.
- [113] Lorens A. Imhof. The long-run behavior of the stochastic replicator dynamics. *Annals of Applied Probability*, 15:1019–1045, 2005.
- [114] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, forthcoming.
- [115] Hans Jørgen Jacobsen, Mogens Jensen, and Birgitte Sloth. Evolutionary learning in signalling games. *Games and Economic Behavior*, 34:34–63, 2001.
- [116] James S. Jordan. Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior*, 5:368–386, 1993.
- [117] Jens Josephson. Stochastic better reply dynamics in finite games. *Economic Theory*, forthcoming.
- [118] Jens Josephson and Alexander Matros. Stochastic imitation in finite games. *Games and Economic Behavior*, 49:244–259, 2004.
- [119] Michihiro Kandori, George J. Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.
- [120] Michihiro Kandori and Rafael Rob. Evolution of equilibria in the long run: A general theory and applications. *Journal of Economic Theory*, 65:383–414, 1995.

- [121] Michihiro Kandori and Rafael Rob. Bandwagon effects and long run technology choice. *Games and Economic Behavior*, 22:84–120, 1998.
- [122] Yong-Gwan Kim and Joel Sobel. An evolutionary approach to pre-play communication. *Econometrica*, 63:1181–1193, 1995.
- [123] Motoo Kimura. On the change of population fitness by natural selection. *Heredity*, 12:145–167, 1958.
- [124] Michael Kosfeld. Stochastic strategy adjustment in coordination games. *Economic Theory*, 20:321–339, 2002.
- [125] Nikolai S. Kukushkin. Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior*, 48:94–110, 2004.
- [126] Timur Kuran and William H. Sandholm. Cultural integration and its discontents. *Review of Economic Studies*, forthcoming, 2008.
- [127] Thomas G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7:49–58, 1970.
- [128] Christoph Kuzmics. Stochastic evolutionary stability in extensive form games of perfect information. *Games and Economic Behavior*, 48:321–336, 2004.
- [129] Ratul Lahkar. The dynamic instability of dispersed price equilibria. Unpublished manuscript, University College London, 2007.
- [130] Ratul Lahkar and William H. Sandholm. The projection dynamic and the geometry of population games. Unpublished manuscript, University of Wisconsin and University College London, 2007.
- [131] Viktor Losert and Ethan Akin. Dynamics of games and genes: Discrete versus continuous time. *Journal of Mathematical Biology*, 17:241–251, 1983.
- [132] Alfred J. Lotka. Undamped oscillation derived from the law of mass action. *Journal of the American Chemical Society*, 42:1595–1598, 1920.
- [133] George J. Mailath. Introduction: Symposium on evolutionary game theory. *Journal of Economic Theory*, 57:259–277, 1992.
- [134] Toshimasa Maruta. On the relationship between risk-dominance and stochastic stability. *Games and Economic Behavior*, 19:221–234, 1997.
- [135] Toshimasa Maruta. Binary games with state dependent stochastic choice. *Journal of Economic Theory*, 103:351–376, 2002.
- [136] Laurent Mathevet. Supermodular Bayesian implementation: Learning and incentive design. Unpublished manuscript, Caltech, 2007.

- [137] John Maynard Smith. Game theory and the evolution of fighting. In *On Evolution*, pages 8–28. Edinburgh University Press, Edinburgh, 1972.
- [138] John Maynard Smith. The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 47:209–221, 1974.
- [139] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, 1982.
- [140] John Maynard Smith and George R. Price. The logic of animal conflict. *Nature*, 246:15–18, 1973.
- [141] Jacek Miękisz. Statistical mechanics of spatial evolutionary games. *Journal of Physics A*, 37:9891–9906, 2004.
- [142] Markus M. Möbius. The formation of ghettos as a local interaction phenomenon. Unpublished manuscript, MIT, 2000.
- [143] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [144] P. A. P. Moran. *The Statistical Processes of Evolutionary Theory*. Clarendon Press, Oxford, 1962.
- [145] David P. Myatt and Chris C. Wallace. A multinomial probit model of stochastic evolution. *Journal of Economic Theory*, 113:286–301, 2003.
- [146] David P. Myatt and Chris C. Wallace. An evolutionary justification for thresholds in collective-action problems. Unpublished manuscript, Oxford University, 2007.
- [147] David P. Myatt and Chris C. Wallace. An evolutionary analysis of the volunteer’s dilemma. *Games and Economic Behavior*, forthcoming.
- [148] David P. Myatt and Chris C. Wallace. When does one bad apple spoil the barrel? An evolutionary analysis of collective action. *Review of Economic Studies*, forthcoming.
- [149] John H. Nachbar. ‘Evolutionary’ selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory*, 19:59–89, 1990.
- [150] Anna Nagurney and Ding Zhang. Projected dynamical systems in the formulation, stability analysis, and computation of fixed demand traffic network equilibria. *Transportation Science*, 31:147–158, 1997.
- [151] John F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:287–295, 1951.
- [152] Georg Nöldeke and Larry Samuelson. An evolutionary analysis of backward and forward induction. *Games and Economic Behavior*, 5:425–454, 1993.

- [153] Martin A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap/Harvard, Cambridge, 2006.
- [154] Martin A. Nowak, Sebastian Bonhoeffer, and Robert M. May. More spatial games. *International Journal of Bifurcation and Chaos*, 4:33–56, 1994.
- [155] Martin A. Nowak, Sebastian Bonhoeffer, and Robert M. May. Spatial games and the maintenance of cooperation. *Proceedings of the National Academy of Sciences*, 91:4877–4881, 1994.
- [156] Martin A. Nowak and Robert M. May. Evolutionary games and spatial chaos. *Nature*, 359:826–829, 1992.
- [157] Martin A. Nowak and Robert M. May. The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos*, 3:35–78, 1993.
- [158] Martin A. Nowak, Akira Sasaki, Christine Taylor, and Drew Fudenberg. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428:646–650, 2004.
- [159] Jörg Oechssler and Frank Riedel. Evolutionary dynamics on infinite strategy spaces. *Economic Theory*, 17:141–162, 2001.
- [160] Jörg Oechssler and Frank Riedel. On the dynamic foundation of evolutionary stability in continuous models. *Journal of Economic Theory*, 107:141–162, 2002.
- [161] Paul Rhode and Mark Stegeman. A comment on “Learning, mutation, and long run equilibria in games”. *Econometrica*, 64:443–449, 1996.
- [162] Klaus Ritzberger and Jörgen W. Weibull. Evolutionary selection in normal form games. *Econometrica*, 63:1371–1399, 1995.
- [163] Jack Robles. Evolution with changing mutation rates. *Journal of Economic Theory*, 79:207–223, 1998.
- [164] Jack Robles. Evolution, bargaining, and time preferences. *Economic Theory*, forthcoming.
- [165] Arthur Robson and Fernando Vega-Redondo. Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory*, 70:65–92, 1996.
- [166] Robert W. Rosenthal. A class of games possessing pure strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [167] Larry Samuelson. Evolutionary foundations of solution concepts for finite, two-player, normal-form games. In Moshe Y. Vardi, editor, *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge (Pacific Grove, CA, 1988)*, pages 211–225. Morgan Kaufmann Publishers, Los Altos, CA, 1988.

- [168] Larry Samuelson. Stochastic stability in games with alternative best replies. *Journal of Economic Theory*, 64:35–65, 1994.
- [169] Larry Samuelson. *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, 1997.
- [170] Larry Samuelson and Jianbo Zhang. Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391, 1992.
- [171] William H. Sandholm. Simple and clever decision rules in a model of evolution. *Economics Letters*, 61:165–170, 1998.
- [172] William H. Sandholm. Almost global convergence to p -dominant equilibrium. *International Journal of Game Theory*, 30:107–116, 2001.
- [173] William H. Sandholm. Potential games with continuous player sets. *Journal of Economic Theory*, 97:81–108, 2001.
- [174] William H. Sandholm. Evolutionary implementation and congestion pricing. *Review of Economic Studies*, 69:81–108, 2002.
- [175] William H. Sandholm. Evolution and equilibrium under inexact information. *Games and Economic Behavior*, 44:343–378, 2003.
- [176] William H. Sandholm. Excess payoff dynamics and other well-behaved evolutionary dynamics. *Journal of Economic Theory*, 124:149–170, 2005.
- [177] William H. Sandholm. Negative externalities and evolutionary implementation. *Review of Economic Studies*, 72:885–915, 2005.
- [178] William H. Sandholm. Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. Unpublished manuscript, University of Wisconsin, 2006.
- [179] William H. Sandholm. Evolution in Bayesian games II: Stability of purified equilibria. *Journal of Economic Theory*, 136:641–667, 2007.
- [180] William H. Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132:367–382, 2007.
- [181] William H. Sandholm. Potential functions for population games and normal form games. Unpublished manuscript, University of Wisconsin, 2007.
- [182] William H. Sandholm. Simple formulas for stationary distributions and stochastically stable states. *Games and Economic Behavior*, 59:154–162, 2007.
- [183] William H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, forthcoming.

- [184] William H. Sandholm and Emin Dokumacı. *Dynamo: Phase diagrams for evolutionary dynamics*. Software. <http://www.ssc.wisc.edu/~whs/dynamo>, 2007.
- [185] William H. Sandholm, Emin Dokumacı, and Ratul Lahkar. The projection dynamic and the replicator dynamic. Unpublished manuscript, University of Wisconsin and University College London, 2007.
- [186] William H. Sandholm and Ady Pauzner. Evolution, population growth, and history dependence. *Games and Economic Behavior*, 22:84–120, 1998.
- [187] Yuzuru Sato, Eizo Akiyama, and J. Doyne Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99:4748–4751, 2002.
- [188] Karl H. Schlag. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78:130–156, 1998.
- [189] Peter Schuster and Karl Sigmund. Replicator dynamics. *Journal of Theoretical Biology*, 100:533–538, 1983.
- [190] Peter Schuster, Karl Sigmund, Josef Hofbauer, and Robert Wolff. Selfregulation of behaviour in animal societies I: Symmetric contests. *Biological Cybernetics*, 40:1–8, 1981.
- [191] Reinhard Selten. Evolution, learning, and economic behavior. *Games and Economic Behavior*, 3:3–24, 1991.
- [192] S. Shahshahani. A new mathematical framework for the study of linkage and selection. *Memoirs of the American Mathematical Society*, 211, 1979.
- [193] Lloyd S. Shapley. Some topics in two person games. In M. Dresher, Lloyd S. Shapley, and Albert W. Tucker, editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–28. Princeton University Press, Princeton, 1964.
- [194] Brian Skyrms. *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge, 1990.
- [195] Brian Skyrms. Chaos in game dynamics. *Journal of Logic, Language, and Information*, 1:111–130, 1992.
- [196] Hal L. Smith. *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, Providence, RI, 1995.
- [197] Michael J. Smith. The stability of a dynamic model of traffic assignment—an application of a method of Lyapunov. *Transportation Science*, 18:245–252, 1984.
- [198] Mark Stegeman and Paul Rhode. Stochastic Darwinian equilibria in small and large populations. *Games and Economic Behavior*, 49:171–214, 2004.

- [199] Jeroen M. Swinkels. Evolutionary stability with equilibrium entrants. *Journal of Economic Theory*, 57:306–332, 1992.
- [200] Jeroen M. Swinkels. Adjustment dynamics and rational play in games. *Games and Economic Behavior*, 5:455–484, 1993.
- [201] György Szabó and Gábor Fáth. Evolutionary games on graphs. *Physics Reports*, 446:97–216, 2007.
- [202] György Szabó and Christoph Hauert. Phase transitions and volunteering in spatial public goods games. *Physical Review Letters*, 89:11801(4), 2002.
- [203] Kei-ichi Tainaka. Physics and ecology of rock-paper-scissors game. In T. A. Marsland and I. Frank, editors, *Computers and Games, Second International Conference (Hammatsu 2000)*, number 2063 in Lecture Notes in Computer Science, pages 384–395, Berlin, 2001. Springer.
- [204] Yasuo Tanabe. The propagation of chaos for interacting individuals in a large population. *Mathematical Social Sciences*, 51:125–152, 2006.
- [205] Peter D. Taylor and Leo Jonker. Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156, 1978.
- [206] Bernhard Thomas. On evolutionarily stable sets. *Journal of Mathematical Biology*, 22:105–115, 1985.
- [207] Donald Topkis. Equilibrium points in nonzero-sum n -person submodular games. *SIAM Journal on Control and Optimization*, 17:773–787, 1979.
- [208] Thomas Tröger. Why sunk costs matter for bargaining outcomes: An evolutionary approach. *Journal of Economic Theory*, 102:28–53, 2002.
- [209] Takashi Ui. Robustness of stochastic stability. Unpublished manuscript, Bank of Japan, 1998.
- [210] Eric van Damme and Jörgen W. Weibull. Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory*, 106:296–315, 2002.
- [211] Fernando Vega-Redondo. *Evolution, Games, and Economic Behaviour*. Oxford University Press, Oxford, 1996.
- [212] Fernando Vega-Redondo. The evolution of Walrasian behavior. *Econometrica*, 65:375–384, 1997.
- [213] Fernando Vega-Redondo. *Complex Social Networks*. Cambridge University Press, Cambridge, 2007.
- [214] Vito Volterra. *Lecons sur la Theorie Mathematique de la Lutte pour la Vie*. Gauthier-Villars, Paris, 1931.

- [215] John von Neumann and Oscar Morgenstern. *Theory of Games and Economic Behavior*. Prentice-Hall, Princeton, 1944.
- [216] Jörgen W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, 1995.
- [217] Jörgen W. Weibull. 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, 1996.
- [218] Franz J. Weissing. Evolutionary stability and dynamic stability in a class of evolutionary normal form games. In Reinhard Selten, editor, *Game Equilibrium Models I*, pages 29–97. Springer, Berlin, 1991.
- [219] H. Peyton Young. The evolution of conventions. *Econometrica*, 61:57–84, 1993.
- [220] H. Peyton Young. An evolutionary model of bargaining. *Journal of Economic Theory*, 59:145–168, 1993.
- [221] H. Peyton Young. Conventional contracts. *Review of Economic Studies*, 65:773–792, 1998.
- [222] H. Peyton Young. *Individual Strategy and Social Structure*. Princeton University Press, Princeton, 1998.
- [223] H. Peyton Young. The dynamics of conformity. In Steven N. Durlauf and H. Peyton Young, editors, *Social Dynamics*, pages 133–153. Brookings Institution Press/MIT Press, Washington/Cambridge, 2001.
- [224] E. Christopher Zeeman. Population dynamics from game theory. In Zbigniew Nitecki and Clark Robinson, editors, *Global Theory of Dynamical Systems (Evanston, 1979)*, number 819 in *Lecture Notes in Mathematics*, pages 472–497, Berlin, 1980. Springer.
- [225] Junfu Zhang. A dynamic model of residential segregation. *Journal of Mathematical Sociology*, 28:147–170, 2004.
- [226] Junfu Zhang. Residential segregation in an all-integrationist world. *Journal of Economic Behavior and Organization*, 24:533–550, 2004.