Population Games and Evolutionary Dynamics

William H. Sandholm

September 9, 2009

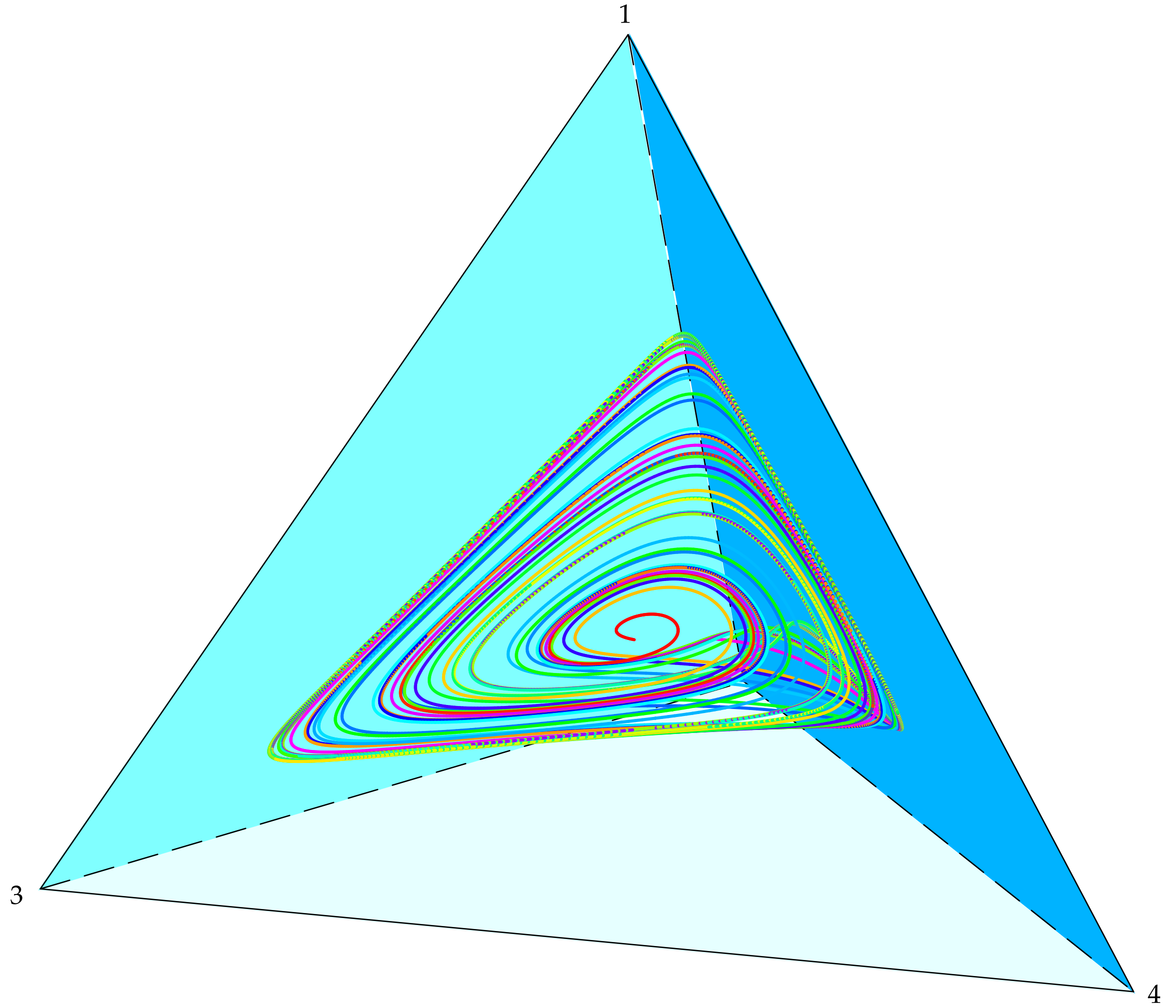

Chaos under the replicator dynamic

CONTENTS IN BRIEF

CONTENTS

CHAPTER

ONE

Introduction

This book describes an approach to modeling recurring strategic interactions in large populations of small, anonymous agents. The approach is built upon two basic elements. The first, called a *population game*, describes the strategic interaction that is to occur repeatedly. The second, called a *revision protocol*, specifies the myopic procedure that agents employ to decide when and how to choose new strategies. Starting with a population game and a revision protocol, one can derive dynamic processes, both deterministic and stochastic, that describe how the agents' aggregate behavior changes over time. These processes are known as *evolutionary game dynamics*.

This introductory chapter begins the work of adding substance to this austere account of evolutionary game theory, providing motivations for and overviews of the analyses to come. For the most part, the chapter is written with an eye toward modeling in economics and other social sciences, though it also discusses the biological origins of the field. But these perspectives should not be viewed as constraints, as the methods presented in this book have ready applications in other disciplines that require models of interacting populations of humans, animals, or machines.

Section 1.1 introduces the notion of a population game by presenting applications, offering informal definitions, and discussing connections with normal form games. It then previews our treatment of population games in Chapters 2–3. Section 1.2 describes dynamic models of behavior in recurrent play of population games, and contrasts this dynamic approach with the equilibrium approach traditionally used in game theory. This section also offers an overview of our presentation of evolutionary dynamics in Chapters 4–12. Section 1.3 concludes with some remarks on motivations for and interpretations of evolutionary game theory. References relevant to the discussions in the text can be found in the Notes at the end of the chapter, where the references from the Preface can also be found.

1.1 Population Games

> [O]nly after the theory for moderate numbers of participants has been satisfactorily developed will it be possible to decide whether extremely great numbers of participants simplify the situation... We share the hope...that such simplifications will indeed occur...
>
> von Neumann and Morgenstern (1944, p. 14)

> We shall now take up the "mass-action" interpretation of equilibrium points... It is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal.
>
> To be more detailed, we assume that there is a population (in the sense of statistics) of participants for each position of the game. Let us also assume that the "average playing" of the game involves n participants selected at random from the n populations, and that there is a stable average frequency with which each pure strategy is employed by the "average member" of the appropriate population.
>
> Nash (1950b, p. 21)

> There are many situations, however, in which an individual is, in effect, competing not against an individual opponent but against the population as a whole, or some section of it. Such cases can loosely be described as "playing the field"... [S]uch contests against the field are probably more widespread and important than pairwise contests.
>
> Maynard Smith (1982, p. 23)

1.1.1 Modeling Interactions in Large Populations

One can imagine many economic, social, and technological environments in which large collections of small agents make strategically interdependent decisions:

- Network congestion. Drivers commute over a highway network. The delay each driver experiences depends not only on the route he selects, but also on the congestion created by other agents along this route.
- Public goods and externalities. A local government maintains a collection of public recreation facilities. The benefit that a family obtains from using a facility depends on the quality of the facility and the number of other families that use it.
- Industrial organization. Software developers choose whether to copyright their products or make them freely available under a public license; they can also choose to work within an existing open-source framework. The latter options entail a loss of control, but allow products to improve through the accumulation of uncoordinated individual efforts.

- The emergence of conventions, norms, and institutions. Firms in a developing economy choose among various business practices: whether to accept credit or require cash, whether to fight or acquiesce to corrupt officials, whether to reward merit or engage in nepotism. Through historical precedent and individual decisions, conventions about business conduct are formed. Whether or not these conventions are efficient, they enable firms to form accurate expectations about how trading partners will act.
- Cultural integration and assimilation. The behavior of immigrants settling in a new country is influenced by traditions imported from their home country, and, to the extent that interactions with the incumbent population require coordination, by the practices of these incumbents as well. At the same time, the choices of the incumbents are influenced by the need to coordinate with the immigrants. The interplay of these forces determines how behavior in the society as a whole evolves.
- Language and communication. Agents without a common language interact repeatedly, attempting to communicate their intentions to the others on each occasion. Whether these attempts are ultimately understood, and whether they enable the agents to coordinate on mutually beneficial behavior, depends on meanings determined by the aggregate communication pattern.
- Task allocation and decentralized control. The employees of a large firm provide services to customers in a number of distinct locations. To take advantage of the ground-level information possessed by the employees, the firm has them allocate themselves among the customers requiring service, providing incentives to ensure that the employees' choices further the firm's objectives.
- Markets and bargaining. Large numbers of buyers and sellers participate in a centralized exchange. Each individual specifies acceptable terms of trade in the hopes of obtaining the greatest benefit from his initial endowment.

While the environments listed above are quite varied, they have certain basic features in common. First, each environment contains a large number of agents capable of making independent decisions. Second, each agent is small, in that his choices have only a minor impact on other agents' outcomes. Third, agents are anonymous: an agent's outcome from the interaction depends on his own strategy and the distribution of others' strategies; further individuation of the opponents is not required.

Simultaneous interactions exhibiting these three properties can be modeled using *population games*. The participants in a population game form a *society* consisting of one or more *populations* of agents. Agents in a given population are assumed to be identical: an agent's population determines his role in the game, the strategies available to him, and his preferences. These preferences are described by a payoff function that conditions on the agent's own strategy and the distribution of strategies in each population. The populations may have either finite numbers of agents or continua of agents, whichever is more convenient. When populations are continuous, the payoffs to each strategy are assumed to depend on the society's aggregate behavior in a continuous fashion, reflecting the idea that very small changes in aggregate behavior do not lead to large changes in the consequences of playing any given strategy.

Aggregate behavior in a population game is described by a *social state*, which specifies the empirical distribution of strategy choices in each population. For simplicity, we assume throughout this book that there are a finite number of populations, and that members of each population choose from a finite strategy set. Doing so ensures that the social state is finite-dimensional, expressible as a vector with a finite number of components; if populations are continuous the set of social states is a polytope.

Except when they are revising (see the next section), agents in population games are assumed to play pure strategies. One of the main reasons for introducing randomized strategies is moot here: when the populations are continuous, and payoffs are continuous in the social state, pure strategy Nash equilibria always exist. This guarantee may be one of the "simplifications" that von Neumann and Morgenstern had in mind when they looked ahead to the study of games with large numbers of players. But this fact is not as essential to our theory as it is to traditional approaches to analyzing games: while traditional approaches are grounded on the assumption of equilibrium play, we emphasize the process through which agents adjust their behavior in response to their current strategic environment.

From our point of view, a more important simplification provided by the population games framework is the description of behavior using social states—that is, by distributions of agents' choices. To understand the advantages this approach brings, let us contrast it with the standard framework for modeling simultaneous-move interactions: normal form games. To define a normal form game, we must specify a (finite or infinite) set of players, and to describe behavior in such a game, we must stipulate each player's strategy choice. If there are many players, these tasks can be laborious.

One can view a population game as a normal form game that satisfies certain restrictions on the diversity and anonymity of the players. But it is nevertheless preferable to work directly with population games, and so to avoid the extra work involved in individuating the players. Moreover, describing behavior using empirical strategy distributions shifts our attention from questions like "who chose strategy i?" to ones like "how many chose strategy i?" and "what happens to the payoff of strategy j players if some agents switch to strategy *i*?" This is just where our attention should be: as we will see in Chapter 3, the answers to the last question determine the incentive structure of a population game.

Much research in evolutionary game theory has focused on population games of a particularly simple sort: those generated when populations of agents are matched to play a normal form game. Indeed, Nash informally introduced population games of this sort in 1949 in proposing the "mass action" interpretation of his equilibrium concept, a development that seems to have gone unnoticed for the next 45 years (see the Notes). This population-matching interpretation of normal form games has a clear appeal, and we will use normal form games as a ready source of simple examples throughout the book.

At the same time, if our main interest is in large population interactions rather than in normal form games, then focusing only on matching in normal form games is quite restrictive. Maynard Smith observes that matching is a rather special sort of interaction in large populations. Instead, interactions in which each agent's payoffs are determined directly from all agents' behavior—what Maynard Smith terms "playing the field"—seem to us to be the rule rather than the exception. Only a few of the applications listed at the onset are most naturally modeled using matching; some, like congestion in highway networks, require payoffs to depend nonlinearly on the population state, and so are mathematically inconsistent with a random matching approach. While one might expect that moving from linear to nonlinear payoffs might lead to intractable models, we will see that it does not: the dynamics we study are nonlinear even when payoffs in the underlying game are not, so allowing nonlinear payoffs does not lead to a qualitative increase in the complexity of the analysis.

1.1.2 Definitions and Classes of Population Games

Our formal introduction to population games begins in Chapter 2, which offers definitions of population states, payoff functions, and various other items mentioned and unmentioned above. To reduce the degree of abstraction, and to illustrate the range of possible applications, the chapter presents a number of basic examples of population games; among these are *congestion games*, which provide a general but tractable model of network congestion and of related sorts of multilateral externalities. Finally, by showing how low-dimensional population games can be represented in pictures, Chapter 2 ushers in the geometric methods of analysis that are emphasized throughout the book.

From a purely formal point of view, a population game is defined by an arbitrary collection of continuous, real-valued functions on an appropriate domain. While some basic results, including existence of Nash equilibrium, can be proved at this level of generality, obtaining more specific conclusions requires us to focus on classes of games defined by certain structural properties. In Chapter 3, we introduce three important classes of population games: *potential games*, *stable games*, and *supermodular games*. Each of these classes of games includes a number of important examples, and each is characterized by restrictions on the nature of the externalities that users of each strategy impose on one another. The structure imposed by these restrictions makes analyses of games in these classes relatively simple. For instance, in games from all three classes, one can prove existence of Nash equilibrium without recourse to fixed point theorems. But more important to us is the impact of this structure on disequilibrium behavior: we will see in Chapter 7 that in potential games, stable games, and supermodular games, broad classes of evolutionary dynamics are assured of converging to Nash equilibrium from arbitrary initial conditions.

1.2 Evolutionary Dynamics

> The state of equilibrium. . . is therefore *stable*; i.e., if either of the producers, misled as to his true interest, leaves it temporarily, he will be brought back to it by a series of reactions, constantly declining in amplitude, and of which the dotted lines of the figure give a representation by their arrangement in steps.
>
> Cournot (1838)

> We repeat most emphatically that our theory is thoroughly static. A dynamic theory would unquestionably be more complete and preferable.
>
> von Neumann and Morgenstern (1944, p. 44–45)

> An equilibrium would be just an extreme state of rare occurrence if it were not stable—that is, if there were no forces which tended to restore equilibrium as soon as small deviations from it occurred.
>
> Besides this stability "in the small", one may consider stability "in the large"—that is, the ability of the system to reach an equilibrium from any initial position.
>
> Beckmann, McGuire, and Winsten (1956, p. 70)

> An obvious weakness of the game-theoretic approach to evolution is that it places great emphasis on equilibrium states, whereas evolution is a process of continuous, or at least periodic, change. It is, of course, mathematically easier to analyse equilibria than trajectories of change. There are, however, two situations in which game theory models force us to think about change as well as constancy. The first is that a game may not have an ESS, and hence the population cycles indefinitely. . .
>
> The second situation. . . is when, as is often the case, a game has more than one ESS. Then, in order to account for the present state of the population, one has to allow for initial conditions—that is, for the state of the ancestral population.
>
> Maynard Smith (1982, p. 8)

After modeling a strategic interaction using a population game, one would like to use the game as the basis for predicting how agents in the interaction will behave. Traditionally, most predictions in game theory have been based on equilibrium analysis: one introduces some notion of equilibrium play, and then finds all behaviors in the game that agree with the equilibrium notion. While the equilibrium approach is standard practice in applications of game theory, the quotations above, drawn from seminal works in economic theory, game theory, transportation science, and theoretical biology, all emphasize that this approach is incomplete, and should be complimented by an analysis of dynamics. The latter quotes go further: they point out that local stability analysis, which checks whether equilibrium play will be restored after small disturbances in behavior, is only a first step, as it begs the question of how equilibrium is established in the first place. These concerns are most pronounced in settings with large numbers of players, where the interplayer coordination of beliefs and actions associated with equilibrium seems most difficult to achieve.

The majority of this book studies dynamic models of behavior in large population games, and so is an attempt to provide some answers to the questions raised above.

1.2.1 Knowledge, Rationality, and Large Games

The fundamental solution concept of noncooperative game theory is *Nash equilibrium*: the requirement that each agent choose a strategy that is optimal given the choices of the others. There are many other solution concepts for games, but most of them are refinements of Nash's definition, and are called upon to reduce the set of predictions in games with multiple Nash equilibria.

Despite the central role of the Nash equilibrium concept, the traditional, rationalistic justification for applying this concept is not especially convincing. This justification is based on three assumptions about the players in the game. First, each player is assumed to be *rational*, acting to maximize his payoffs given what he knows. Second, players have *knowledge of the game* they are playing: they know what strategies are available, and what payoffs result from every strategy profile. Third, the players have *equilibrium knowledge*: they are able to anticipate correctly what their opponents will do. If all players expect a certain strategy profile to be played, and if each player is rational and understands the payoff consequences of switching strategies, then each player is content to play his part in the strategy profile if and only if that profile is a Nash equilibrium.

Of the three assumptions listed above, the equilibrium knowledge assumption is the hardest to accept. Certainly, shared expectations about play can be put into place by a disinterested moderator who guides the players to a particular strategy profile. But without such guidance, it is hard to explain how players can introspectively anticipate how others will act, particularly in games with large numbers of participants.

In fact, when we consider games with many players or strategies, even apparently innocuous conditions for equilibrium play may be called into question. Under the traditional interpretation of equilibrium play in a traffic network, a driver choosing a route to work has a complete mental account of all of the routes he could take, and he is able to anticipate the delay that would arise on each route for any possible profile of choices by his fellow drivers. Evidently, the assumption of knowledge of the game, while seemingly innocent, may actually be quite bold when the game is large.

This discussion suggests that in large games, even the force of seemingly weak solution concepts, ones that do not require equilibrium knowledge, should not be taken for granted. For instance, a basic tenet of traditional game-theoretic analysis holds that a strictly dominated strategy—a strategy that performs worse than some single alternative strategy regardless of how opponents behave—should not be chosen. This requirement is uncontroversial when players have full knowledge of the game. But if players are unable or unwilling to keep the entire game in mind, they may well not notice that one strategy is dominated by another. While one might expect that an accumulation of experience might ensure that players eventually avoid dominated strategies, we will see that this is not necessarily so: in Section 9.4, we will present a set of seemingly mild conditions on players' updating rules that are enough to ensure that strictly dominated strategies must survive in perpetuity in some games.

Obtaining a convincing rationalistic justification for equilibrium play in games seems an impossible task. But in settings where the same game is played many times, the possibilities become brighter, as one can replace introspection with repetition as the basis for coordination of behavior. In large population settings, repetition may be enough to coordinate behavior even when agents' information and abilities are quite limited. But while dynamic approaches can support and even refine traditional game-theoretic predictions, they also can also lead to predictions of cyclical or more complex nonstationary behavior, possibilities that are ignored by traditional analyses.

1.2.2 Foundations for Evolutionary Dynamics

There are a variety of approaches one could take to studying disequilibrium dynamics in games, depending on the number of players involved, the information the players are expected to possess, and the importance of the interaction to the players. *Evolutionary game theory*, the approach studied in this book, considers the dynamics of behavior in large, strategically interacting populations. This approach posits that agents only occasionally switch strategies, and then use simple myopic rules to decide how to act. While these assumptions are certainly not appropriate for every application, they seem natural when the interaction in question is just one among many the agent faces, so that the sporadic application of a rule of thumb is a reasonable way for the agent to proceed.

While it is possible to proceed directly with a description of aggregate behavior dynamics, we find it preferable to begin by specifying when and how individual agents make decisions. We accomplish these tasks using a modeling device called a *revision protocol*. A revision protocol is a function that takes the strategies' payoffs and utilization levels as inputs; it returns as outputs the overall rate of switching strategies, and the probabilities with which each alternative strategy will be chosen.

In defining a revision protocol, we implicitly specify the informational burden that the agents must bear. Starting in Chapter 4, we show that revision protocols come in many different varieties, from ones that embody exact myopic optimization, to others that require each agent to know nothing more than his own current payoff. In all cases, though, the protocol only relies on information about current strategic conditions: historical information, as well as counterfactual information about strategies' performances under other conditions, are not considered.

Also implicit in the definition of a revision protocol is the method to be used to identify alternative strategies. One can place protocols into two broad categories according to this criterion. Under *imitative protocols*, an agent obtains a candidate strategy by observing the strategy of a randomly chosen member of his population. Under *direct protocols*, agents are assumed to choose candidate strategies directly; a strategy's popularity does not directly influence the probability with which it is considered. (Agents may also meander among different protocols as time passes, in which case they are said to employ a *hybrid protocol*.) After obtaining a candidate strategy, an agent can evaluate its current payoff by briefly experimenting with it. If the strategy is currently in use, its current payoff can also be determined by observing the outcomes of an opponent who employs it.

We will see in Chapter 5 that the aggregate behavior dynamics generated by imitative protocols are very similar to those studied in mathematical biology. Thus, the *replicator dynamic*, introduced in the mathematical biology literature to model natural selection, also describes the aggregate behavior of agents who use certain imitative protocols. In contrast, dynamics based on direct selection—for instance, the *best response dynamic*, which is based on optimal myopic choices—behave rather differently than those studied in biology. Direct selection allows unused strategies to be introduced to the population, which is impossible under pure imitation (or under biological reproduction without mutations). For its part, imitation generates dynamics with relatively simple functional forms and behavior. But the simple forms and special properties of imitative dynamics are rather special: we will find that when agents use hybrid protocols, their aggregate behavior agrees in broad qualitative terms with what one would see under direct selection of alternative strategies.

1.2.3 Deterministic Evolutionary Dynamics

Suppose that one or more large populations of agents recurrently play a population game, with each agent occasionally updating his choice of strategies using a fixed revision protocol. Since there are many agents, and since the stochastic elements of the agents' updating procedures are idiosyncratic, one expects these stochastic influences to be averaged away, leaving aggregate behavior to evolve in an essentially deterministic fashion. In Chapter 4, we explain how such a deterministic evolutionary process can be described by an ordinary differential equation, and how this equation can be derived from the population game and revision protocol. A formal justification for our study of this differential equation, which we call the *mean dynamic*, is deferred until Chapter 10; see Section 1.2.4 below.

The mean dynamic specifies the rate of change in the use of each strategy i. It is the difference of two terms: an inflow term, which captures agents' switches from other strategies to strategy i, and an outflow term, which captures agents' switches from strategy i to other strategies. Of course, the exact specification of the dynamic depends on the primitives of the model: namely, the protocol the agents employ and the game they play.

Taking a slightly different point of view, we note that any fixed revision protocol defines a map that takes population games as inputs, and returns specific instances of mean dynamics as outputs. We call this map from population games to ordinary differential equations an *evolutionary dynamic*.

Chapters 5 and 6 introduce a variety of families of evolutionary dynamics and investigate their properties. Each family of dynamics is generated from a set of revision protocols sharing certain qualitative features: for instance, being based on imitation, or on myopic optimization. The main results in Chapters 5 and 6 establish properties of evolutionary dynamics that hold regardless of the population game at hand.

One of the basic issues considered in these chapters is the relationship between the rest points of the dynamics and the payoff structure of the underlying game. We show in Chapter 5 that the rest points of imitative dynamics include all Nash equilibria of the game being played. Further analyses reveal that under many direct and hybrid dynamics, the sets of rest points and Nash equilibria are identical. This latter property, which we call *Nash stationarity*, provides a first link between the population dynamics and traditional equilibrium analyses. A second issue we address is the connection between out-of-equilibrium dynamics and incentives in the underlying game. Many of the dynamics we study satisfy a monotonicity property called *positive correlation*, which requires strategies' growth rates to be positively correlated with their payoffs. This property and its relatives are a basic ingredients in our later analyses of dynamic stability.

Properties like Nash stationarity only provide a weak justification for the prediction of Nash equilibrium play. To obtain a more convincing defense of this prediction, one must address the questions of stablilty raised at the start of this section.

Chapter 8 considers local stability of equilibrium: whether equilibrium will be reached if play begins at a nearby social state. It is easy to see that some Nash equilibria are unlikely to be locally stable. For instance, if play begins near the mixed equilibrium of a coordination game, then myopic adjustment will lead the population away from this equilibrium. Additional restrictions beyond Nash's condition are thus needed to obtain general stability results. A main finding in Chapter 8 is that the notion of an *evolutionarily stable state* (*ESS*), introduced by Maynard Smith and Price for model of evolution in populations of mixed strategists, provides a general sufficient condition for local stability under the pure-strategist dynamics studied here.

Local stability results only offer a partial justification of equilibrium predictions. To be convinced that equilibrium will be reached, one must look instead for global convergence results, which establish that equilibrium is attained from any initial state. We offer such results in Chapter 7, where we prove that in certain classes of games—namely, the classes of potential games, stable games, and supermodular games, introduced in Chapter 3—there are classes of dynamics that converge to equilibrium from all or almost all initial conditions.

These stability results provide strong support for the Nash prediction in some settings, but they say little about behavior in games outside of the classes the results cover. While one might hope that these results are nevertheless representative of behavior in most games, this seems not to be so. In Chapter 9, we present a variety of examples in which deterministic dynamics enter cycles far from Nash equilibrium, and others in which the dynamics display chaotic behavior. Thus, if we take the model of myopic adjustment dynamics seriously, we must sometimes accept these more complicated limit behaviors as more credible predictions than equilibrium play. Moreover, as we hinted earlier, the possibility of nonconvergent behavior has some counterintuitive consequences: we show in Section 9.4 that under "typical" evolutionary dynamics, we can always find simple games in which strictly dominated strategies survive.

1.2.4 Orders of Limits for Stochastic Evolutionary Models

The deterministic dynamics studied in Chapters 4–9 form the largest portion of the literature on evolutionary game dynamics. But there is another major branch of the literature that focuses on the infinite horizon behavior of stochastic evolutionary dynamics, and whose central aim is to obtain unique predictions of play in games with multiple equilibria. In Chapter 10, we bring together these two literatures by showing how deterministic and stochastic evolutionary game dynamics can be derived from a single foundation.

To accomplish this, we consider population games played by large but finite populations of agents. Together, a finite-population game and a revision protocol define a stochastic evolutionary process—in particular, a Markov process—on the now finite set of population states. What methods are most useful for studying this process depends on the parameter values—the population size N and the time horizon T—that are relevant to the application at hand.

This point is made most clearly by taking limits. To begin, suppose that we fix the time horizon T and take the population size N to infinity. The main result of Chapter 10 shows that once N is large enough, the stochastic evolutionary process is very likely to behave in a nearly deterministic way, mirroring a solution trajectory of the relevant mean dynamic through time T. For intuition, observe that the stochastic aspects of the evolutionary process—the random arrivals of revision opportunities, and the randomizations among candidate strategies during these opportunities—are idiosyncratic. When the population size becomes very large, the idiosyncratic noise is averaged away. Thus, the behavior of the process is driven by the expected changes in the state, which are precisely what the mean dynamic captures.

Alternatively, suppose we fix the population size N and take the time horizon T to infinity. In this case we are studying the infinite horizon behavior of a finite-state Markov process. If this process is irreducible—that is, if all states are mutually accessible, as is true, for instance, under any revision protocol that always places positive probability on each alternative strategy—then results from probability theory tell us that every social state will be visited infinitely often, and that the proportion of time spent in each state over the infinite horizon does not depend on the initial state. This invariance property opens the door to obtaining unique predictions of play.

1.2.5 Stationary Distributions and Stochastic Stability

The proportion of time that an irreducible Markov process spends in each state is described by the process's *stationary distribution*. By computing this distribution, we can obtain an exact description of the infinite horizon behavior of our evolutionary process. Chapter 11 presents some basic results on stationary distributions of stochastic evolutionary models, and studies in detail those cases—namely, the models that generate *reversible* Markov processes—in which the stationary distribution can be computed exactly.

These analyses provide a first illustration of a fundamental idea. Even if the underlying game has multiple equilibria, it is nevertheless the case that in the infinite horizon, a stochastic evolutionary dynamic may spend the vast majority of periods in the vicinity of a single state. But we also argue that the amount of time necessary for this analysis to become relevant can be exceedingly long, limiting the range of possible applications.

Once one moves beyond reversible cases, obtaining an exact expression for the stationary distribution is typically impossible. One way of circumventing this difficulty is to take certain parameters of the evolutionary process to their limiting values. If one takes the noise level that parameterizes the agents' revision protocol to zero, or the population size to infinity, one can sometimes describe the limiting stationary distribution, even though the stationary distributions for fixed parameter values cannot be computed explicitly. There are many interesting cases in which the limiting stationary distribution places all of its mass on a single state, providing an especially tidy prediction of infinite horizon play.

States that retain mass in the limiting stationary distribution are said to be *stochastically stable*. In Chapter 12 we offer a complete presentation of stochastic stability theory. We describe the main analytical techniques used to determine the stochastically stable states, and we present key selection results. We also establish formal connections between the states that are stochastically stable in the large population limit and the recurrent states of the relevant mean dynamic. In doing so, we provide another link between stochastic stability analysis and the deterministic dynamics studied in earlier chapters.

1.3 Remarks on History, Motivation, and Interpretation

Evolutionary game theory dates from the early 1970s, when John Maynard Smith introduced the ESS concept as a way of understanding ritualized animal conflict (see the Notes). In pursuing this line of research, Maynard Smith and other biologists borrowed the standard game-theoretic definitions of mathematicians and economists, but reinterpreted these definitions to suit their own purposes. In a biological model, a strategy is not something that an animal chooses, but a behavior that an animal is hardwired to perform by virtue of its genetic endowment. Similarly, payoffs are not a numerical representation of a preference relation, but instead represent "Darwinian fitnesses": they describe how the use of the strategy improves an animal's prospects for survival and reproduction.

Thus, the development of a biological approach to game theory was marked not only by the invention of new solution concepts and methods of analysis, but also by a recasting of existing definitions.

In the late 1980s, economists surmised that this biological paradigm could be used to address problems with the foundations of traditional game theory. During the heyday of the equilibrium refinements literature, game theorists introduced solution concepts that seemed to impose ever larger demands on the reasoning abilities of the players whose behavior they purported to describe. This development raised the criticism that the solution concepts could only describe the behavior of hyperrational players. If, though, it could be shown that the behavior of populations of preprogrammed creatures could be described by means of these same solution concepts, then this would seem to provide a way of dispelling the hyperrationality critique.

In view of this motivation, it is not surprising that many early efforts by economists in evolutionary game theory adopted the biological paradigm wholesale, with populations of agents hardwired to choose certain strategies, and payoffs interpreted as Darwinian fitnesses. These analyses often found that behavior in these populations, as described by ESS and other biological solution concepts, agreed in a formal sense with the most demanding of the equilibrium refinements introduced by economists.

While these links between the biological and economic approaches to game theory might provide some sense of reassurance, it is important to be mindful of the limitations of this line of reasoning. In many economic applications of game theory, the game in question is played by, say, two or three players. Nash equilibrium and its refinements offer predictions about the mixed strategies these players will employ. For their part, the evolutionary models introduced by biologists concern interactions among large populations of animals, and their solution concepts describe the effects of natural selection on population shares of different genetic types. The fact that economists' predictions about rational agents' choices of mixed strategies formally agree with biologists' predictions about population shares under natural selection is certainly intriguing. But there is no obvious logic by which this agreement enables either analysis to justify the other.

Economists working on population dynamics soon realized that to obtain a modeling tool useful for typical social science environments, the Darwinian assumptions and interpretations of evolutionary game theory would need to be replaced with ones reflecting human decision making. This new vision of evolutionary game theory brought back the original interpretation of the game itself: strategies returned to being objects of choice, and payoffs resumed being descriptions of individual preferences. But other parts of the biologists' approach—in particular, the very idea of population dynamics, with changes in strategies' population shares driven by differences in current payoffs—were retained. Of course, some reinterpretation was necessary here, as the "evolution" of population shares would no longer be driven by relative rates of reproduction, but rather by conscious decisions to switch from one strategy to another. More intriguing agreements were discovered between the biological and economic approaches. For instance, Taylor and Jonker introduced the replicator dynamic in 1978 to provide a dynamic foundation for Maynard Smith's static approach to natural selection. But in the 1990s, economists found that this dynamic could also be derived from models of imitation in populations of economic agents—see Section 5.4.

This history has created some understandable confusion about the meaning of the term "evolutionary game theory". From the origins of the field through the earliest contributions of economists, this phrase referred exclusively to game-theoretic models of natural selection. But since the mid 1990s, it has also encompassed models of myopic behavior in games played by large populations of active decision makers. In this interpretation, which predominates in this book, the word "evolution" should not be understood as a direct reference to Darwinian survival of the fittest, but instead should be taken in its broader sense, referring to a process of gradual change.

As the interpretation of evolutionary game theory has changed, so too has its role within economic modeling. As we noted earlier, the initial flush of interest in evolutionary game theory in the economics community stemmed from the theory's potential to provide low-rationality foundations for high-rationality solution concepts. This motivation helps explain why economists working in evolutionary game theory have devoted so much effort to the study of random matching in normal form games. If the point of the theory is to motivate solution concepts for such games, then there is not much reason to move beyond the random matching setting.

As we have emphasized here, predictions about the aggregate behavior of large populations do not directly address how two players would play a two-player normal form game. For such predictions to play more than a metaphorical role, they must do so in applications that specifically concern the behavior of populations. But once such applications are brought to the fore, restricting attention to matching in normal form games begins to feel artificial: if one is interested in modeling behavior in large populations, the games one should write down are those one believes populations actually play.

1.N Notes

Preface: *General references on evolutionary game theory*. Books: Maynard Smith (1982), Hofbauer and Sigmund (1988, 1998), Bomze and Pötscher (1989), Cressman (1992, 2003), Weibull (1995), Vega-Redondo (1996), Samuelson (1997), Fudenberg and Levine (1998), Young (1998b). Surveys: van Damme (1991, Chapter 9), Hines (1987), Kandori (1997), Mailath (1998), Weibull (2002), Hofbauer and Sigmund (2003), Sandholm (2008b),

References on omitted topics. Evolution in extensive form games. Cressman (1996a, 2000, 2003), Cressman and Schlag (1998), Chamberland and Cressman (2000), Binmore et al. (1995a), Binmore and Samuelson (1999), Ponti (2000) (deterministic models); Samuelson (1994, 1997), Nöldeke and Samuelson (1993), Hart (2002), Kuzmics (2004) (stochastic models). Local interaction models: Blume (1993, 1995, 1997), Ellison (1993, 2000), Kosfeld (2002), Miękisz (2004); Herz (1994), Ely (2002), Eshel et al. (1998), Anderlini and Ianni (1996), Goyal and Janssen (1997), Alós-Ferrer and Weidenholzer (2006, 2007, 2008), Goyal (2007), Vega-Redondo (2007), Jackson (2008); Nowak and May (1992, 1993), Nowak (2006), Szabó and Fáth (2007), Hauert (2007). Continuous strategy sets: Eshel (1983), Bomze (1990, 1991), Friedman and Yellin (1997), Oechssler and Riedel (2001, 2002), Eshel and Sansone (2003), Cressman (2005, 2009), Cressman and Hofbauer (2005), Cressman et al. (2006), Friedman and Ostrov (2008), Norman (2008), Hofbauer et al. (2009). Stochastic evolution via diffusion processes: Foster and Young (1990), Fudenberg and Harris (1992), Cabrales (2000), Imhof (2005), Benaïm et al. (2008), Hofbauer and Imhof (2009).

References on heuristic learning. Original references: Brown (1949, 1951), Robinson (1951). Books: Fudenberg and Levine (1998), Young (2004), Cesa-Bianchi and Lugosi (2006). Survey: Hart (2005). Brief overview: Sandholm (2008d).

References on applications. Markets: Hopkins and Seymour (2002), Lahkar (2007), Droste et al. (2002), Sethi (1999), Ben-Shoham et al. (2004), Agastya (2004), Vega-Redondo (1997), Alós-Ferrer et al. (2000), Ania et al. (2002), Alós-Ferrer and Ania (2005), Alós-Ferrer et al. (2006), Kandori et al. (2008), Friedman and Ostrov (2008). Bargaining and hold-up problems: Young (1993b, 1998b,a), Burke and Young (2001), Ellingsen and Robles (2002), Tröger (2002), Binmore et al. (2003), Abreu and Sethi (2003), Dawid and MacLeod (2008), Robles (2008). Signaling and cheap talk: Nöldeke and Samuelson (1993, 1997), Jacobsen et al. (2001), Robson (1990), Matsui (1991), Wärneryd (1993), Blume et al. (1993), Kim and Sobel (1995), Bhaskar (1998), Banerjee and Weibull (2000), Trapa and Nowak (2000), Skyrms (2002), Pawlowitsch (2008), Demichelis and Weibull (2008). Public good provision: Sethi and Somanathan (1996), Myatt and Wallace (2008b,a, 2009). Implementation and decentralized control: Cabrales (1999), Cabrales and Ponti (2000), Sandholm (2002, 2005b, 2007b), Cabrales and Serrano (2007), Mathevet (2007), Arslan et al. (2007), Marden et al.

(2009a). Residential segregation: Young (1998b, 2001), Möbius (2000), Zhang (2004a,b), Dokumacı and Sandholm (2007a). Preference evolution: Güth and Yaari (1992), Güth (1995), Huck and Oechssler (1999), Koçkesen et al. (2000), Sethi and Somanathan (2001), Ok and Vega-Redondo (2001), Ely and Yilankaya (2001), Dekel et al. (2007), Heifetz et al. (2007), Herold and Kuzmics (2008). Cultural evolution: Bisin and Verdier (2001), Sandholm (2001b), Kuran and Sandholm (2008) Montgomery (2009). Applications in biology: Maynard Smith (1982), Hofbauer and Sigmund (1988) (books); Hammerstein and Selten (1994), Dugatkin and Reeve (1998) (surveys).

Section 1.1. Noncooperative games with a continuum of agents were introduced by Schmeidler (1973), who proves that pure strategy Nash equilibria exist when each player's payoffs only depend on his own action and on the aggregate behavior of all players. Schmeidler (1973) uses a normal form approach, in that he names and describes the behavior of each agent separately. Mas-Colell (1984) introduces a simpler formulation based on empirical distributions of strategies rather than on strategy profiles; his formulation can be understood as a population game in which both the set of populations and the set of strategies may be continuous. Mas-Colell (1984) also points out the overlap between his model and the distributional approach to Bayesian games with a continuum of types (Milgrom and Weber (1985)). This connection can be used as the basis for studying evolution in Bayesian games; see Ely and Sandholm (2005). The literature on noncooperative games with large numbers of players is surveyed in Khan and Sun (2002); see Balder (2002) and Carmona and Podczeck (2009) for recent results.

Nash's "mass action" interpretation of his equilibrium concept appears in his dissertation (Nash (1950b)), but not in either of its published versions (Nash (1950a, 1951)). It was later rediscovered by Leonard (1994), Weibull (1994, 1996), and Björnerstedt and Weibull (1996); see also Hofbauer (2000).

Section 1.2. There are many good general references on traditional game-theoretic analysis: see Osborne (2004) at the undergraduate level, and Fudenberg and Tirole (1991), Myerson (1991), Ritzberger (2002), and van Damme (1991) at the graduate level. For an overview of the use of rationality and knowledge assumptions in game theory, see Dekel and Gul (1997).

While evolutionary game theory models assume that agents are myopic, one can imagine large population interactions whose participants are better modeled as forward-looking dynamic optimizers, who attempt to maximize discounted flows of future payoffs. The *perfect foresight dynamics* of Matsui and Matsuyama (1995) introduce forward-looking behavior to large population models while retaining the assumption of inertia from the evolutionary approach. Recent work on perfect foresight dynamics includes Hofbauer and Sorger (1999), Oyama (2002), Matsui and Oyama (2006), Kojima (2006a), Kojima and Takahashi (2008), Takahashi (2008), Oyama et al. (2008), and Oyama and Tercieux (2009).

Section 1.3. While papers by John Maynard Smith in the early 1970s (Maynard Smith (1972, 1974), Maynard Smith and Price (1973)), are usually viewed as marking the birth of evolutionary game theory, Maynard Smith (1982, p. 2) himself notes the appearance of game-theoretic ideas in earlier studies of the evolution of sex ratios: they appeared implicitly in the work on Fisher (1930), and explicitly in the work of Hamilton (1967). But it was only with Maynard Smith's work that game theory began to take on its role as a standard framework for biological modeling.