

Penn
Institute
For
Economic
Research

Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19103-6297
pier@ssc.upenn.edu
<http://www.econ.upenn.edu/pier>

PIER Working Paper 98-001

“Do People Play Nash Equilibrium? Lessons From
Evolutionary Game Theory”

by

George J. Mailath

Do People Play Nash Equilibrium? Lessons From Evolutionary Game Theory*

George J. Mailath
University of Pennsylvania

January 24, 1998
Forthcoming, *Journal of Economic Literature*

Abstract

Evolutionary game theory provides an answer to two of the central questions in economic modeling: When is it reasonable to assume that people are rational? And, when is it reasonable to assume that behavior is part of a Nash equilibrium (and if it is reasonable, which equilibrium)? The traditional answers are not compelling, and much of evolutionary modeling is motivated by the need for a better answer. Evolutionary game theory suggests that, in a range of settings, agents do (eventually) play a Nash equilibrium. Moreover, evolutionary modeling has shed light on the relative plausibility of different Nash equilibria.

1. Introduction

At the same time that noncooperative game theory has become a standard tool in economics, it has also come under increasingly critical scrutiny from theorists and experimentalists. Noncooperative game theory, like neoclassical economics, is built on two heroic assumptions: **Maximization**—every economic agent is a rational decision maker with a clear understanding of the world. And **Consistency**—the agent's understanding, in particular, expectations of other agents' behavior, is correct (i.e., the overall pattern of individual optimizing behavior forms a Nash equilibrium). These assumptions are no less controversial in the context of noncooperative game theory than they are in neoclassical economics.

*I thank Steven Matthews, Loretta Mester, John Pencavel, three referees, and especially Larry Samuelson for their comments.

A major challenge facing noncooperative game theorists today is that of providing a compelling justification for these two assumptions. As I will argue here, many of the traditional justifications are *not* compelling. But without such justification, the use of game theory in applications is problematic. The appropriate use of game theory requires understanding when its assumptions make sense and when they do not.

In some ways, the challenge of providing a compelling justification is not a new one. A major complaint other social scientists (and some economists) have about economic methodology is the central role of the maximization hypothesis. A common informal argument is that any agent not optimizing, in particular, any firm not maximizing profits, will be driven out by market forces. This is an evolutionary argument, and as is well known, Charles Darwin was led to the idea of natural selection from reading Thomas Malthus.¹ But does such a justification work? Is Nash equilibrium (or some related concept) a good predictor of behavior?

While the parallel between noncooperative game theory and neoclassical economics is close, it is not perfect. Certainly, the question of whether agents maximize is essentially the same in both. Moreover, the consistency assumption also appears in neoclassical economics as the assumption that prices clear markets. However, a fundamental distinction between neoclassical economics and noncooperative game theory is that, while the many equilibria of a competitive economy almost always share many of the same properties (such as efficiency or its lack),² the many equilibria of games often have dramatically different properties. While neoclassical economics does not address the question of equilibrium selection, game theory must.

Much of the work in evolutionary game theory is motivated by two basic questions:

1. Do agents play Nash equilibrium?
2. Given that agents play Nash equilibrium, which equilibrium do they play?

Evolutionary game theory formalizes and generalizes the evolutionary argument given above by assuming that more successful behavior tends to be more

¹“In October 1838, that is, fifteen months after I had begun my systematic enquiry, I happened to read for amusement ‘Malthus on Population,’ and being well prepared to appreciate the struggle for existence which everywhere goes on from long-continued observation of the habits of animals and plants, it at once struck me that under these circumstances favourable variations would tend to be preserved, and unfavourable ones to be destroyed. The results of this would be the formation of new species. Here, then I had at last got a theory by which to work;” Charles Darwin (1887, p. 83).

²Or perhaps economists have chosen to study only those properties that are shared by all equilibria. For example, different competitive equilibria have different income distributions.

prevalent. The canonical model has a population of players interacting over time, with their behavior adjusting over time in response to the payoffs (utilities, profits) that various choices have historically received. These players could be workers, consumers, firms, etc. The focus of study is the dynamic behavior of the system. The crucial assumptions are that there is a population of players, these players are interacting, and the behavior is naive (in two senses: players do not believe—understand—that their own behavior potentially affects future play of their opponents, and players typically do not take into account the possibility that their opponents are similarly engaged in adjusting their own behavior). It is important to note that successful behavior becomes more prevalent not just because market forces select against unsuccessful behavior, but also because agents imitate successful behavior.

Since evolutionary game theory studies populations playing games, it is also useful for studying social norms and conventions. Indeed, many of the motivating ideas are the same.³ The evolution of conventions and social norms is an instance of players learning to play an equilibrium. A convention can be thought of as a symmetric equilibrium of a coordination game. Examples include a population of consumers who must decide which type of good to purchase (in a world of competing standards); a population of workers who must decide how much effort to exert; a population of traders at a fair (market) who must decide how aggressively to bargain; and a population of drivers randomly meeting at intersections who must decide who gives way to whom.

Evolutionary game theory has provided a qualified affirmative answer to the first question: In a range of settings, agents do (eventually) play Nash. There is thus support for equilibrium analysis in environments where evolutionary arguments make sense. Equilibrium is best viewed as the steady state of a community whose members are myopically groping toward maximizing behavior. This is in marked contrast to the earlier view (which, as I said, lacks satisfactory foundation), according to which game theory and equilibrium analysis are the study of the interaction of (ultra-) rational agents with a large amount of (common) knowledge.⁴

The question of which equilibrium is played has received a lot of attention, most explicitly in the refinements literature. The two most influential ideas in that literature are backward and forward induction. Backward induction and its extensions—subgame perfection and sequentiality—capture notions of credibility and sequential rationality. Forward induction captures the idea that a player's

³See, for example, Elster (1989), Skyrms (1996), Sugden (1989), and Young (1996).

⁴Even in environments where an evolutionary analysis would not be appropriate, equilibrium analysis *is* valuable in illuminating the strategic structure of the game.

choice of current action can be informative about his future play. The concerns about adequate foundations extend to these refinement ideas. While evolutionary game theory does discriminate between equilibria, backward induction receives little support from evolutionary game theory. Forward induction receives more support. One new important principle for selecting an equilibrium, based on stochastic stability, does emerge from evolutionary game theory, and this principle discriminates between strict equilibria (something backward and forward induction cannot do).

The next section outlines the major justifications for Nash equilibrium, and the difficulties with them. In that section, I identify learning as the best available justification for Nash equilibrium. Section 3 introduces evolutionary game theory from a learning perspective. The idea that Nash equilibrium can be usefully thought of as an evolutionary stable state is described in Section 4. The question of which Nash equilibrium is played is then discussed in Section 5. As much as possible, I have used simple examples. Very few theorems are stated (and then only informally). Recent surveys of evolutionary game theory include van Damme (1987, Chapter 9), Kandori (1997), Mailath (1992), and Weibull (1995).

2. The Question

Economics and game theory typically assume that agents are “rational” in the sense of pursuing their own welfare, as they see it. This hypothesis is tautological without further structure, and it is usually further assumed that agents understand the world as well as (if not better than) the researcher studying the world inhabited by these agents. This often requires an implausible degree of computational and conceptual ability on the part of the agents. For example, while chess is strategically trivial,⁵ it is computationally impossible to solve (at least in the foreseeable future).

Computational limitations, however, are in many ways less important than conceptual limitations of agents. The typical agent is not like Gary Kasparov, the world champion chess player who knows the rules of chess, but also knows that he doesn’t know the winning strategy. In most situations, people do not know they are playing a game. Rather, people have some (perhaps imprecise) notion of the environment they are in, their possible opponents, the actions they and their opponents have available, and the possible payoff implications of different actions. These people use heuristics and rules of thumb (generated from experience) to

⁵Since chess is a finite game of perfect information, it has been known since 1912 that either White can force a win, Black can force a win, or either player can force a draw (Zermelo (1912)).

guide behavior; sometimes these heuristics work well and sometimes they don't.⁶ These heuristics can generate behavior that is inconsistent with straightforward maximization. In some settings, the behavior can appear as if it was generated by concerns of equity, fairness, or revenge.

I turn now to the question of consistency. It is useful to first consider situations that have aspects of coordination, i.e., where an agent (firm, consumer, worker, etc.) maximizes his welfare by choosing the same action as the majority. For example, in choosing between computers, consumers have a choice between PCs based on Microsoft's operating systems and Apple-compatible computers. There is significantly more software available for Microsoft-compatible computers, due to the market share of the Microsoft computers, and this increases the value of the Microsoft-compatible computers. Firms must often choose between different possible standards.

The first example concerns a team of workers in a modern version of Rousseau (1950)'s Stag Hunt.⁷ In the example, each worker can put in low or high effort, the team's total output (and so each worker's compensation) is determined by the minimum effort of all the workers, and effort is privately costly. Suppose that if all workers put in low effort, the team produces a per capita output of 3, while if all workers put in high effort, per capita output is 7. Suppose, moreover, the disutility of high effort is 2 (valued in the same units as output). We can thus represent the possibilities as in Figure 1.⁸ It is worth emphasizing at this point that the characteristics that make the stag hunt game interesting are pervasive. In most organizations, the value of a worker's effort is increasing in the effort levels of the other workers.

What should we predict to be the outcome? Consider a typical worker, whom I call Bruce for definiteness. If all the other workers are only putting in low effort, then the best choice for Bruce is also low effort: high effort is costly, and choosing high effort cannot increase output (since output is determined by the minimum effort of all the workers). Thus, if Bruce expects the other workers to put in low effort, then Bruce will also put in low effort. Since all workers find themselves

⁶Both Andrew Postlewaite and Larry Samuelson have made the observation that in life there are no one-shot games and no "last and final offers." Thus, if an experimental subject is placed in an artificial environment with these properties, the subject's heuristic will not work well until it has adjusted to this environment. I return to this in my discussion of the ultimatum game.

⁷Rousseau's (1950) stag hunt describes several hunters in the wilderness. Individually, each hunter can catch rabbits and survive. Acting together, the hunters can catch a stag and have a feast. However, in order to catch a stag, *every* hunter must cooperate in the stag hunt. If even one hunter does not cooperate (by catching rabbits), the stag escapes.

⁸The stag hunt game is an example of a *coordination game*. A *pure coordination game* differs from the game in Figure 1 by having zeroes in the off-diagonal elements. The game in Figure 13 is a pure coordination game.

		minimum of other workers' efforts	
		high	low
worker's effort	high	5	0
	low	3	3

Figure 1: A “stag-hunt” played by workers in a team.

in the same situation, we see that all workers choosing to put in low effort is a *Nash equilibrium*: each worker is behaving in his own best interest, given the behavior of the others. Now suppose the workers (other than Bruce) are putting in high effort. In this case, the best choice for Bruce is now high effort. While high effort is costly, Bruce’s choice of high rather than low effort now does affect output (since Bruce’s choice is the minimum) and so the increase in output (+4) is more than enough to justify the increase in effort. Thus, if Bruce expects all the other workers to be putting in high effort, then Bruce will also put in high effort. As for the low effort case, a description of behavior in which all workers choose high effort constitutes a Nash equilibrium.

These two descriptions of worker behavior (all choose low effort and all choose high effort) are internally consistent; they are also *strict*: Bruce strictly prefers to choose the same effort level as the other workers.⁹ This implies that even if Bruce is somewhat unsure about the minimum effort choice (in particular, as long as Bruce assigns a probability of no more than 0.4 to some worker choosing the other effort choice), this does not affect his behavior.

But are these two descriptions good predictions of behavior? Should we, as outside observers, be confident in a prediction that *all* the workers in Bruce’s team will play one of the two Nash equilibrium? And if so, why and which one? Note that this is not the same as asking if Bruce will choose an effort that is consistent with equilibrium. After all, both choices (low and high effort) are consistent with equilibrium, and so Bruce necessarily chooses *an* equilibrium effort.

Rather, the question concerns the behavior of the group as a whole. How do we rule out Bruce choosing high effort because he believes everyone else will, while Sheila (another worker on Bruce’s team) chooses low effort because she believes everyone else will.¹⁰ This is, of course, ruled out by equilibrium considerations. But what does that mean? The critical feature of the scenario just described is

⁹A *strict* Nash equilibrium is a Nash equilibrium in which, given the play of the opponents, each player has a unique best reply.

¹⁰While the term “coordination failure” would seem to be an apt one to describe this scenario, that term is commonly understood (particularly in macroeconomics) to refer to coordination on an inefficient equilibrium (such as low effort chosen by *all* workers).

that the expectations of Bruce or Sheila about the behavior of the other members of the team are *incorrect*, something that Nash equilibrium *by definition* does not allow.

As I said earlier, providing a compelling argument for Nash equilibrium is a major challenge facing noncooperative game theory today.¹¹ The consistency in Nash equilibrium seems to require that players know what the other players are doing. But where does this knowledge come from? When or why is this a plausible assumption? There are several justifications typically given for Nash equilibria: preplay communication, self-fulfilling prophecies (consistent predictions), focal points, and learning.¹²

The idea underlying *preplay communication* is straightforward. Suppose the workers in Bruce's team meet before they must choose their effort levels and discuss how much effort they each will exert. If the workers reach an agreement that they all believe will be followed, it must be a Nash equilibrium (otherwise at least one worker has an incentive to deviate). This justification certainly has appeal and some range of applicability. However, it does not cover all possible applications. It also assumes that an agreement is reached. While it seems clear that an agreement will be reached (and which one) in our stag hunt example (at least if the team is small!), this is not true in general. Suppose, for example, the interaction has the characteristics of a battle-of-the-sexes game (Figure 2). Such a game, which may describe a bargaining interaction, has several Pareto noncomparable Nash equilibria: there are several profitable agreements that can be reached, but the bargainers have opposed preferences over which agreement is reached. In this case, it is not clear that an agreement will be reached. Moreover, if the game does have multiple Pareto noncomparable Nash equilibria, then the

¹¹The best introspective (i.e., knowledge or epistemic) foundations for Nash equilibrium assume that each player's conjectures about the behavior of other players is *known* by all the players, see Aumann and Brandenburger (1995). This assumption does not appear to be a significant improvement over the original assumption of Nash behavior.

More generally, there is no compelling introspective argument for any useful equilibrium notion. The least controversial are those do not require the imposition of a consistency condition, such as rationalizability (introduced by Bernheim (1984) and Pearce (1984), for two players it is equivalent to the iterated deletion of strictly dominated strategies), the iterated deletion of weakly dominated strategies, and backward induction. However, in most games, rationalizability does little to constrain players' behavior, and as we will see below, the iterated deletion of weakly dominated strategies and backward induction are both controversial.

¹²What if there is only one equilibrium? Does this by itself give us a reason to believe that the unique equilibrium will be played? The answer is no. It is possible that the unique Nash equilibrium yields each player their maximin values, while at the same time being riskier (in the sense that the Nash equilibrium strategy does not *guarantee* the maximin value). This is discussed by, for example, Harsanyi (1977, p. 125) and Aumann (1985). Kreps (1990b, p. 135) describes a complicated game with a unique equilibrium that is also unlikely to be played.

		employer	
		high wage	low wage
employee	high wage	3, 2	0, 0
	low wage	0, 0	2, 3

Figure 2: A “Battle-of-the-sexes” between an employer and a potential employee bargaining over wages. Each simultaneously makes a wage demand or offer. The worker is only hired if they agree.

preplay communication stage is itself a bargaining game and so perhaps should be explicitly modelled (at which point, the equilibrium problem resurfaces).¹³ Finally, there may be no possibility of preplay communication.

The second justification of *self-fulfilling prophecy* runs as follows: If a theory uniquely predicting players’ behaviors is known by the players in the game, then it must predict Nash equilibria (see Myerson (1991, pp. 105-108) for an extended discussion of this argument). The difficulty, of course, is that the justification requires a theory that uniquely predicts player behavior, and that is precisely what is at issue.

The *focal point* justification, due to Schelling (1960), can be phrased as “if there is an obvious way to play in a game (derived from either the structure of the game itself or from the setting), then players will know what other players are doing.” There are many different aspects of a game that can single out an “obvious way to play.” For example, considerations of fairness may make equal divisions of a surplus particularly salient in a bargaining game. Previous experience suggests that stopping at red lights and going through green is a good strategy, while another possible strategy (go at red and stop at green) is not (even though it is part of another equilibrium). It is sometimes argued that efficiency is such an aspect: if an equilibrium gives a higher payoff to every player than any other equilibrium, then players “should not have any trouble coordinating their expectations at the commonly preferred equilibrium point” (Harsanyi and Selten (1988, p. 81)). In our earlier stag hunt example, this principle (called *payoff dominance* by Harsanyi and Selten (1988)) suggests that the high effort equilibrium is the obvious way to play. On the other hand, the low effort equilibrium is less risky, with Bruce receiving a payoff of 3, no matter what the other members of his team do. In contrast, it is possible that a choice

¹³The preplay communication stage might also involve a *correlating device*, like a coin. For example, the players might agree to flip a coin: if heads, then the worker receives a high wage, while if tails, the worker receives a low wage.

of high effort yields a payoff of only 0. As we will see, evolutionary game theory has been particularly important in addressing this issue of riskiness and payoff dominance. See Kreps (1990b) for an excellent extended discussion and further examples of focal points.¹⁴

Finally, agents may be able to *learn* to play an equilibrium. In order to learn to play an equilibrium, players must be playing the same game repeatedly, or at least, similar games that can provide valuable experience. Once all players have learned how their opponents are playing, and if all players are maximizing, then we must be at a Nash equilibrium. There are two elements to this learning story. The first is that, given maximizing behavior of players, players can learn the behavior of their opponents.¹⁵ The second is that players are maximizing. This involves, as I discussed earlier, *additional* considerations of learning. Even if a player knows how his or her opponents have played (for example, the player may be the “last mover”), he may not know what the best action is. A player will use his or her past experience, as well as the experience of other players (if that is available and relevant), to make forecasts as to the current or future behavior of opponents, as well as the payoff implications of different actions. Since learning itself changes the environment that the agents are attempting to learn (as other agents change their behavior in response to their own learning), the process of learning is quite subtle. Note that theories of learning are focal point theories, since histories (observed patterns of play of other agents) can serve as coordinating devices that make certain patterns of play “the obvious way to play,” as in the traffic example from the previous paragraph.

The discussion so far points out some of the problems with many of the standard justifications for Nash equilibrium and its two assumptions (maximization and consistency). Besides learning, another approach is to concede the lack of a sound foundation for consistency, but maintain the hypothesis that agents maximize. The question then is whether rationality and knowledge of the game (including the rationality of opponents) is enough, at least in some interesting cases, to yield usable predictions. The two (related) principles that people have typically used in applications are backward induction and the iterated deletion of weakly dominated strategies.¹⁶ In many games, these procedures identify a unique outcome. The difficulty is that they require an implausible degree of ra-

¹⁴Kreps (1990a, Chapter 12) is a formal version of Kreps (1990b).

¹⁵Non-evolutionary game theory work on learning has focused on the question of when maximizing players can in fact learn the behavior of their opponents. Examples of this work include Fudenberg and Kreps (1989), Fudenberg and Levine (1993), Gul (1996), and Kalai and Lehrer (1993).

¹⁶Backward induction is also the basic principle underlying the elimination of equilibria relying on “incredible” threats.




Figure 3: A short centipede game.

tionality and knowledge of other players' rationality. The two key examples here are Rosenthal's centipede game (so-called because of the appearance of its extensive form—Figure 3 is a short centipede) and the finitely-repeated Prisoner's dilemma. The centipede is conceptually simpler, since it is a game of perfect information. The crucial feature of the game is that each player, when it is his or her turn to move, strictly prefers to end the game immediately (i.e., choose E_n), rather than have the opponent end the game on the next move by choosing E_{n+1} . Moreover, at each move each player strictly prefers to have play proceed for a further two moves, rather than end immediately. For the game in Figure 3, if play reaches the last possible move, player I surely ends the game by choosing E_3 rather than C_3 . Knowing this, player II should choose E_2 . The induction argument then leads to the conclusion that player I necessarily ends the game on the first move.¹⁷ I suspect everyone is comfortable with the prediction that in a two move game, player I stops the game immediately. However, while this logic is the same in longer versions, many researchers are no longer comfortable with the same prediction.¹⁸ It only requires one player to think that there is some chance that the other player is willing to play C initially to support playing C in early moves. Similarly, if we consider the repeated prisoner's dilemma, the logic of backward induction (together with the property that the unique one period dominant strategy equilibrium yields each player their maximin payoff) implies that even in early periods cooperation is not possible.

3. Evolutionary Game Theory and Learning

The previous section argued that, of the various justifications that have been advanced for equilibrium analysis, learning is the least problematic. Evolutionary game theory is a particularly attractive approach to learning. In the typical evolutionary game theoretic model, there is a population of agents, each of whose payoff is a function of not only how they behave, but also how the agents they interact with behave. At any point in time, behavior within the population is distributed over the different possible strategies, or behaviors. If the population

¹⁷This argument is not special to the extensive form. If the centipede is represented as a normal form game, this backward induction is mimicked by the iterated deletion of weakly dominated strategies.

While this game has many Nash equilibria, they all involve the same behavior on the equilibrium path: player I chooses E_1 . The equilibria only differ in the behavior of players off-the-equilibrium-path.

¹⁸Indeed, more general knowledge-based considerations have led some researchers to focus on the procedure of one round of weak domination followed by iterated rounds of strict domination. A nice (although technical) discussion is Dekel and Gul (1997).

is finite, a state (of the population) is a description of which agents are choosing which behavior. If the population is infinite, a state is a description of the fractions of the population that are playing each strategy. If a player can maximize, and knows the state, then he can choose a best reply. If he does not know the state of the population, then he must draw inferences about the state from whatever information he has. In addition, even given knowledge of the state, the player may not be able to calculate a best reply. Calculating a best reply requires that a player know all the strategies available and the payoff implications of all these strategies. The observed history of play is now valuable for two reasons. First, the history conveys information about how the opponents are expected to play. Second, the observed success or failure of various choices helps players determine what might be good strategies in the future. Imitation is often an important part of learning; successful behavior tends to be imitated. In addition, successful behavior will be taught. To the extent that players are imitating successful behavior and not explicitly calculating best replies, it is not necessary for players to distinguish between knowledge of the game being played and knowledge of how opponents are playing. Players need know only what was successful, not why it was successful.

Evolutionary game theoretic models are either static or dynamic, and the dynamics are either in discrete or continuous time. A *discrete time dynamic* is a function that specifies the state of the population in period $t + 1$ as a function of the state in period t , i.e., given a distribution of behavior in the population, the dynamic specifies next period's distribution. A *continuous time dynamic* specifies the relative rates of change of the fractions of the population playing each strategy as a function of the current state. *Evolutionary* (also known as *selection* or *learning*) dynamics specify that behavior that is successful this period will be played by a larger fraction in the immediate future. Static models study concepts that are intended to capture stability ideas motivated by dynamic stories without explicitly analyzing dynamics.

An important point is that evolutionary dynamics do not build in any assumptions on behavior or knowledge, other than the basic principle of differential selection—apparently successful behavior increases its representation in the population, while unsuccessful behavior does not.

Evolutionary models are *not* structural models of learning or bounded rationality. While the motivation for the basic principle of differential selection involves an appeal to learning and bounded rationality, individuals are not explicitly modelled. The feature that successful behavior last period is an attractive choice this period does seem to require that agents are naive learners. They do not believe, or understand, that their own behavior potentially affects future play

of their opponents, and they do not take into account the possibility that their opponents are similarly engaged in adjusting their own behavior. Agents do not look for patterns in historical data. They behave as if the world is stationary, even though their own behavior should suggest to them it is not.¹⁹ Moreover, agents behave as if they believe that other agents' experience is relevant for them. Imitation then seems reasonable. Note that the context here is important. This style of modelling does not lend itself to small numbers of agents. If there is only a small population, is it plausible to believe that the agents are not aware of this? And if agents are aware, then imitation is not a good strategy. As we will see, evolutionary dynamics have the property that, in large populations, if they converge then they converge to a Nash equilibrium.²⁰ This property is a necessary condition for any reasonable model of social learning. For suppose we had a model in which behavior converged to something that is not a Nash equilibrium. Since the environment is eventually stationary and there is a behavior (strategy) available to some agent that yields a higher payoff, then that agent should eventually figure this out and so deviate.²¹

One concern sometimes raised about evolutionary game theory is that its agents are implausibly naive. This concern is misplaced. If an agent is boundedly rational, then he does not understand the model as written. Typically, this model is very simple (so that complex dynamic issues can be studied) and so the bounds of the rationality of the agents are often quite extreme. For example, agents are usually not able to detect any cycles generated by the dynamics. Why then are the agents not able to figure out what the modeler can? As in most of economic theory, the role of models is to improve our intuition and to deepen our understanding of how particular economic or strategic forces interact. For this literature to progress, we must analyze (certainly now, and perhaps forever) simple and tractable games. The games are intended as examples, experiments, and allegories. Modelers do not make assumptions of bounded rationality because they believe players are stupid, but rather that players are not as sophisticated as our models generally assume. In an ideal world, modelers would study very

¹⁹It is difficult to build models of boundedly rational agents who look for patterns. Aoyagi (1996) and Sonsino (1997) are rare examples of models of boundedly rational agents who can detect cycles.

²⁰More accurately, they converge to a Nash equilibrium of the game determined by the strategies that are played along the dynamic path. It is possible that the limit point fails to be a Nash equilibrium because a strategy that is not played along the path has a higher payoff than any strategy played along the path.

Even if there is "drift" (see the ultimatum game below), the limit point will be close to a Nash equilibrium.

²¹Of course, this assumes that this superior strategy is something the agent could have thought of. If the strategy is never played, then the agent might never think of it.

	Bold	Cautious
Bold	0	3
Cautious	1	2

Figure 4: Payoff to a trader following the row strategy against a trader following the column strategy. The gains from trade are 4.

complicated games and understand how agents who are boundedly rational in some way behave and interact. But the world is not ideal and these models are intractable. In order to better understand these issues, we need to study models that can be solved. Put differently, the bounds on rationality are to be understood relative to the complexity of the environment.

4. Nash Equilibrium as an Evolutionary Stable State

Consider a population of traders that engage in randomly determined pairwise meetings. As is usual, I will treat the large population as being infinite. Suppose that when two traders meet, each can choose one of two strategies, “bold” and “cautious.” If a trader has chosen to be “bold,” then he will bargain aggressively, even to the point of losing a profitable trade; on the other hand, if a trader has chosen to be “cautious,” then he will never lose a profitable trade. If a bold trader bargains with a cautious trader, a bargain will be struck that leaves the majority of gains from trade with the bold trader. If two cautious traders bargain, they equally divide the gains from trade. If two bold traders bargain, no agreement is reached. One meeting between two traders is depicted as the symmetric game in Figure 4.²²

Behaviors with higher payoffs are more likely to be followed in the future. Suppose that the population originally consists only of cautious traders. If no trader changes his behavior, the population will stay this way. Now suppose there is a perturbation that results in the introduction of some bold traders into the population. This perturbation may be the result of entry, perhaps traveling traders from another community have arrived, or experimentation, perhaps some of the traders are not sure that they are behaving optimally and try something different.²³ In a population of cautious traders, bold traders also consummate their deals and receive a higher payoff than the cautious traders. So over time the

²²This is the Hawk-Dove game traditionally used to introduce the concept of an evolutionary stable strategy.

²³In the biological context, this perturbation is referred to as a mutation or an invasion.

fraction of cautious traders in the population will decrease and the fraction of bold traders will increase. However, once there are enough bold traders in the population, bold traders no longer have an advantage (on average) over cautious traders (since two bold traders cannot reach an agreement), and so the fraction of bold traders will always be strictly less than one. Moreover, if the population consists entirely of bold traders, a cautious trader can successfully invade the population. The only stable population is divided between bold and cautious traders, with the precise fraction determined by payoffs. In our example, the stable population is equally divided between bold and cautious traders. This is the distribution with the property that bold and cautious traders have equal payoffs (and so also describes a mixed strategy Nash equilibrium). At that distribution, if the population is perturbed, so that, for example, slightly more than half the population are now bold while slightly less than half are cautious, cautious traders have a higher payoff and so learning will lead to an increase in the number of cautious traders at the expense of bold traders, until balance is once again restored.

It is worth emphasizing that the final state is independent of the original distribution of behavior in the population, and that this state corresponds to the symmetric Nash equilibrium. Moreover, this Nash equilibrium is dynamically stable: any perturbation from this state is always eliminated.

This parable illustrates the basics of an evolutionary game theory model, in particular, the interest in the dynamic behavior of the population. The next section describes the well-known notion of an evolutionary stable strategy, a static notion that attempts to capture dynamic stability. Section 4.2 then describes explicit dynamics, while Section 4.3 discusses asymmetric games.

4.1. Evolutionary stable strategies

In the biological setting, the idea that a stable pattern of behavior in a population should be able to eliminate any invasion by a “mutant” motivated Maynard Smith and Price (1973) to define an evolutionary stable strategy (ESS).²⁴ If a population pattern of behavior is to eliminate invading mutations, it must have a higher fitness than the mutant in the population that results from the invasion. In biology, animals are programmed (perhaps genetically) to play particular strategies and the payoff is interpreted as “fitness,” with fitter strategies having higher reproductive rates (reproduction is asexual).

It will be helpful to use some notation at this point. The collection of available behaviors (strategies) is S and the payoff to the agent choosing i when his

²⁴Good references on biological dynamics and evolution are Maynard Smith (1982) and Hofbauer and Sigmund (1988).

opponent chooses j is $\pi(i, j)$. We will follow most of the literature in assuming that there is only a finite number of available strategies. Any behavior in S is called *pure*. A *mixed* strategy is a probability distribution over pure strategies. While any pure strategy can be viewed as the mixed strategy that places probability one on that pure strategy, it will be useful to follow the convention that the term “mixed strategy” always refers to a mixed strategy that places strictly positive probability on at least two strategies. Mixed strategies have two leading interpretations as a description of behavior in a population: either the population is *monomorphic*, in which every member of the population plays the same mixed strategy, or the population is *polymorphic*, in which each member plays a pure strategy and the fraction of the population playing any particular pure strategy equals the probability assigned to that pure strategy by the mixed strategy.²⁵ As will be clear, the notion of an evolutionary stable strategy is best understood by assuming that each agent can choose a mixed strategy and the population is originally monomorphic.

Definition 1. A (potentially mixed) strategy p is an Evolutionary Stable Strategy (ESS) if:

1. the payoff from playing p against p is at least as large as the payoff from playing any other strategy against p ; and
2. for any other strategy q that has the same payoff as p against p , the payoff from playing p against q is at least as large as the payoff from playing q against q .²⁶

Thus, p is an evolutionary stable strategy if it is a symmetric Nash equilibrium, and if, in addition, when q is also a best reply to p then p does better against q than q does. For example, a is an ESS in the game in Figure 5. ESS is a static notion that attempts to capture dynamic stability.

There are two cases to consider: The first is that p is a strict Nash equilibrium (see footnote 9). Then, p is the only best reply to itself (so p must be pure), and any agent playing q against a population whose members mostly play p receives a lower payoff (on average) than a p player. As a result, the fraction of the population playing q shrinks.

²⁵The crucial distinction is whether agents can play and inherit (learn) mixed strategies. If not, then any mixed strategy state is necessarily the result of a polymorphic population. On the other hand, even if agents can play mixed strategies, the population may be polymorphic with different agents playing different strategies.

²⁶The payoff to p against q is $\pi(p, q) = \sum_{i,j} \pi(i, j)p_i q_j$. Formally, the strategy p is an ESS if, for all q , $\pi(p, p) \geq \pi(q, p)$, and if there exists $q \neq p$ such that $\pi(p, p) = \pi(q, p)$, then $\pi(p, q) > \pi(q, q)$.

	a	b
a	2	1
b	2	0

Figure 5: The numbers in the matrix are the row player’s payoff. The strategy a is an ESS in this game.

The other possibility is that p is not the only best reply to itself (if p is a mixed strategy, then any other mixture with the same or smaller support is also a best reply to p). Suppose q is a best reply to p . Then both an agent playing p and an agent playing q earn the same payoff against a population of p players. After the perturbation of the population by the entry of q players, however, the population is not simply a population of p players. There is also a small fraction of q players, and their presence will determine whether the q players are eliminated. The second condition in the definition of ESS guarantees that in the perturbed population, it is the p players who do better than the q players when they play against a q player.

4.2. The replicator and other more general dynamics

While plausible, the story underlying ESS suffers from its reliance on the assumption that agents can learn (in the biological context, inherit) mixed strategies. ESS is only useful to the extent that it appropriately captures some notion of dynamic stability. Suppose individuals now choose only pure strategies. Define p_i^t as the proportion of the population choosing strategy i at time t . The state of the population at time t is then $p^t \equiv (p_1^t, \dots, p_n^t)$, where n is the number of strategies (of course, p^t is in the $n - 1$ dimensional simplex). The simplest evolutionary dynamic one could use to investigate the dynamic properties of ESS is the *replicator dynamic*. In its simplest form, this dynamic specifies that the proportional rate of growth in a strategy’s representation in the population, p_i^t , is given by the extent to which that strategy does better than the population average.²⁷ The payoff to strategy i when the state of the population is p^t is $\pi(i, p^t) = \sum_j \pi(i, j)p_j^t$, while the population average payoff is $\pi(p^t, p^t) = \sum_{ij} \pi(i, j)p_i^t p_j^t$. The continuous time replicator dynamic is then:

$$\frac{dp_i^t}{dt} = p_i^t \times \left(\pi(i, p^t) - \pi(p^t, p^t) \right). \tag{1}$$

²⁷The replicator dynamic can also be derived from more basic biological arguments.

Thus, if strategy i does better than average, its representation in the population grows ($dp_i^t/dt > 0$), and if another strategy i' is even better, then its growth rate is also higher than that of strategy i . Equation (1) is a differential equation that, together with an initial condition, uniquely determines a path for the population that describes, for any time t , the state of the population.

A state is a *rest point* of the dynamics if the dynamics leave the state unchanged (i.e., $dp_i^t/dt = 0$ for all i). A rest point is *Liapunov stable* if the dynamics do not take states close to the rest point far away. A rest point is *asymptotically stable* if, in addition, any path (implied by the dynamics) that starts sufficiently close to the rest point converges to that rest point.

There are several features of the replicator dynamic to note. First, if a pure strategy is extinct (i.e., no fraction of the population plays that pure strategy) at any point of time, then it is never played. In particular, any state in which the same pure strategy is played by every agent (and so every other strategy is extinct) is a rest point of the replicator dynamic. So, being a rest point is *not* a sufficient condition for Nash equilibrium. This is a natural feature that we already saw in our discussion of the game in Figure 4—if everyone is bargaining cautiously and if traders are *not aware* of the possibilities of bold play, then there is no reason for traders to change their behavior (even though a rational agent who understood the payoff implications of the different strategies available would choose bold behavior rather than cautious).

Second, the replicator dynamic is not a *best reply dynamic*: strategies that are not best replies to the current population will still increase their representation in the population if they do better than average (this feature only becomes apparent when there are at least three available strategies). This again is consistent with the view that this is a model of boundedly rational learning, where agents do not understand the full payoff implications of the different strategies.

Finally, the dynamics can have multiple asymptotically stable rest points. The asymptotic distribution of behavior in the population can depend upon the starting point. Returning to the stag hunt game of Figure 1, if a high fraction of workers has chosen high effort historically, then those workers who had previously chosen low effort would be expected to switch to high effort, and so the fraction playing high effort would increase. On the other hand, if workers have observed low effort, perhaps low effort will continue to be observed. Under the replicator dynamic (or any deterministic dynamic that rewards actions with higher payoffs), if $p_{\text{high}}^0 > 3/5$, then $p_{\text{high}}^t \rightarrow 1$, while if $p_{\text{high}}^0 < 3/5$, then $p_{\text{high}}^t \rightarrow 0$. The equilibrium that players eventually learn is determined by the original distribution of players across high and low effort. If the original distribution is random (e.g., p_{high}^0 is determined as a realization of a uniform random variable), then the low

effort equilibrium is 3/5's as likely to arise as the high effort equilibrium. This notion of path dependence—that history matters—is important and attractive.

Zeeman (1980) and Taylor and Jonker (1978) have shown that if p is an ESS, it is asymptotically stable under the continuous time replicator dynamic, and that there are examples of asymptotically stable rest points of the replicator dynamic that are not ESS. If dynamics are specified that allow for mixed strategy inheritance, then p is an ESS if and only if it is asymptotically stable (see Hines (1980), Robson (1992), and Zeeman (1981)). A point I will come back to is that both asymptotic stability and ESS are concerned with the stability of the system after a once and for all perturbation. They do not address the consequences of continual perturbations. As we will see, depending upon how they are modelled, continual perturbations can profoundly change the nature of learning.

While the results on the replicator dynamic are suggestive, the dynamics are somewhat restrictive, and there has been some interest in extending the analysis to more general dynamics.²⁸ Interest has focused on two classes of dynamics. The first, *monotone dynamics*, roughly requires that on average, players switch from worse to better (not necessarily the best) pure strategies. The second, more restrictive, class, *aggregate monotone dynamics*, requires that, in addition, the switching of strategies have the property that the induced distribution over strategies in the population have a higher average payoff. It is worth noting that the extension to aggregate monotone dynamics is not that substantial: aggregate monotone dynamics are essentially multiples of the replicator dynamic (Samuelson and Zhang (1992)).

Since, by definition, a Nash equilibrium is a strategy profile with the property that every player is playing a best reply to the behavior of the other players, every Nash equilibrium is a rest point of any monotone dynamic. However, since the dynamics may not introduce behavior that is not already present in the population, not all rest points are Nash equilibria. If a rest point is asymptotically stable, then the learning dynamics, starting from a point close to the rest point but with all strategies being played by a positive fraction of the population, converge to the rest point. Thus, if the rest point is not a Nash equilibrium, some agents are not optimizing and the dynamics will take the system away from that rest point. This is the first major message from evolutionary game theory: if the

²⁸There has also been recent work exploring the link between the replicator dynamic and explicit models of learning. Börgers and Sarin (1997) consider a single boundedly rational decision maker using a version of the Bush and Mosteller (1951, 1955) model of positive reinforcement learning, and show that the equation describing individual behavior looks like the replicator dynamic. Gale, Binmore and Samuelson (1995) derive the replicator dynamic from a behavioral model of aspirations. Schlag (1994) derives the replicator dynamic in a bandit setting, where agents learn from others.

		visiting trader	
		Bold	Cautious
owner	Bold	0, 0	3, 1
	Cautious	1, 3	2, 2

Figure 6: Payoffs to two traders bargaining in the row trader’s establishment. The first number is the owner’s payoff, the second is the visitor’s. The gains from trade are 4.

state is asymptotically stable, then it describes a Nash equilibrium.

4.3. Asymmetric games and nonexistence of ESS

The assumption that the traders are drawn from a single population and that the two traders in any bargaining encounter are symmetric is important. In order for two traders in a trading encounter to be symmetric in this way, their encounter must be on “neutral” ground, and not in one of their establishments.²⁹ Another possibility is that each encounter occurs in one of the traders’ establishments, and the behavior of the traders depends upon whether he is the visitor or the owner of the establishment. This is illustrated in Figure 6. This game has three Nash equilibria: The stable profile of a fifty-fifty mixture between cautious and bold behavior is still a mixed strategy equilibrium (in fact, it is the only symmetric equilibrium). There are also two pure strategy asymmetric equilibria (the owner is bold, while the visitor is cautious; and the owner is cautious, while the visitor is bold). Moreover, these two pure strategy equilibria are strict.

Symmetric games, like that in Figure 6, have the property that the strategies available to the players do not depend upon their *role*, i.e., row (owner) or column (visitor). The assumption that in a symmetric game, agents cannot condition on their role is called *no role identification*. In most games of interest, the strategies available to a player also depend upon his role. Even if not, as in the example just discussed, there is often some distinguishing feature that allows players to identify their role (such as the row player being the incumbent and the column player being an entrant in a contest for a location). Such games are called *asymmetric*.

The notion of an ESS can be applied to such games by either changing the definition to allow for asymmetric mutants (as in Swinkels (1992)), or, equivalently, by symmetrizing the game. The symmetrized game is the game obtained by assuming that, *ex ante*, players do not know which role they will have, so that

²⁹More accurately, the traders’ behavior cannot depend on the location.

	α	β
a	2, 2	1, 2
b	2, 1	0, 0

Figure 7: The asymmetric version of the game in Figure 5. The strategy (a, α) is not an ESS of the symmetrized game, since the payoff to (a, α) when playing against (b, α) equals the payoff to (b, α) when playing against (b, α) (which equals $\frac{1}{2}\pi(a, \alpha) + \frac{1}{2}\pi(\alpha, b) = \frac{3}{2}$).

players' strategies specify behavior conditional on different roles, and first having a move of nature that allocates each player to a role. (In the trader example, a coin flip for each trader determines whether the trader stays "home" this period, or visits another establishment.) However, every ESS in such a symmetrized game is a strict equilibrium (Selten (1980)). This is an important negative result, since most games (in particular, non-trivial extensive form games) do not have strict equilibria and so ESSs do not exist for most asymmetric games.

The intuition for the nonexistence is helpful in what comes later, and is most easily conveyed if we take the monomorphic interpretation of mixed strategies. Fix a non-strict Nash equilibrium of an asymmetric game, and suppose it is the row player that has available other best replies. Recall that a strategy specifies behavior for the agent in his role as the row player and also in his role as the column player. The mutant of interest is given by the strategy that specifies one of these other best replies for the row role and the existing behavior for the column role. This mutant is not selected against, since its expected payoff is the same as that of the remainder of the population. First note that all agents in their column role still behave the same as before the invasion. Then, in his row role, the mutant is (by assumption) playing an alternative best reply, and in his column role he receives the same payoff as every other column player (since he is playing the same as them). See Figure 7. The idea that evolutionary (or selection) pressures may not be effective against alternative best replies plays an important role in subsequent work on cheap talk and forward and backward induction.

It is also helpful to consider the behavior of dynamics in a two-population world playing the trader game. There are two populations, owners and visitors. A state now consists of the pair (p, q) , where p is the fraction of owners who are bold, while q is the fraction of visitors who are bold.³⁰ There is a separate replicator dynamic for each population, with the payoff to a strategy followed

³⁰This is equivalent to considering dynamics in the one-population model where the game is symmetrized and p is the fraction of the population who are bold in the owner role.

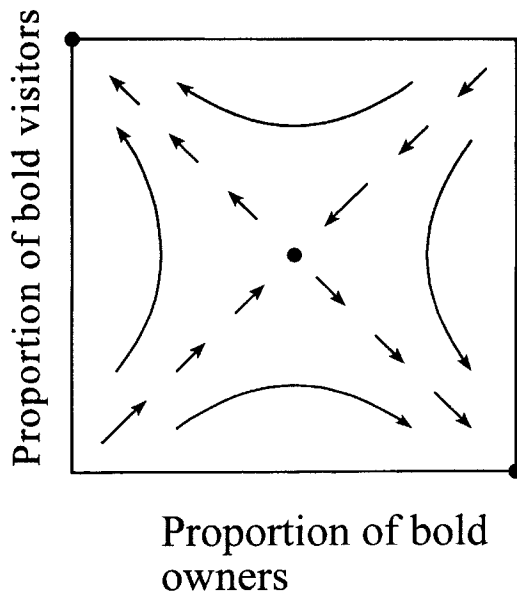


Figure 8: The phase diagram for the two population trader game.

by an owner depending only on q , and not p (this is the observation from the previous paragraph that row players do not interact with row players). While (p^*, q^*) , where $p^* = q^* = \frac{1}{2}$, is still a rest point of the two dimensional dynamical system describing the evolution of trader behavior, it is no longer asymptotically stable. The phase diagram is illustrated in Figure 8. If there is a perturbation, then owners and traders move toward one of the two strict pure equilibria. Both of the asymmetric equilibria are asymptotically stable.

If the game is asymmetric, then we have already seen that the only ESS are strict equilibria. There is a similar lack of power in considering asymptotic stability for general dynamics in asymmetric games. In particular, asymptotic stability in asymmetric games implies “almost” strict Nash equilibria, and if the profile is pure, it is strict.³¹ For the special case of replicator dynamics, a Nash

³¹Samuelson and Zhang (1992, p. 377) has the precise statement. See also Friedman (1991)

		Player 2	
		<i>L</i>	<i>R</i>
Player 1	<i>T</i>	1, 1	1, 0
	<i>B</i>	1, 1	0, 0

Figure 9: Any state in which all agents in population 2's play *L* is a stable rest point.

equilibrium is asymptotically stable if and only if it is strict (Ritzberger and Weibull (1995)).

5. Which Nash equilibrium?

Beyond excluding non-Nash behavior as stable outcomes, evolutionary game theory has provided substantial insight into the types of behavior that are consistent with evolutionary models, and the types of behavior that are not.

5.1. Domination

The strongest positive results concern the behavior of the dynamics with respect to strict domination: If a strategy is strictly dominated by another (that is, it yields a lower payoff than the other strategy for all possible choices of the opponents), then over time that strictly dominated strategy will disappear (a smaller fraction of the population will play that strategy). Once that strategy has (effectively) disappeared, any strategy that is now strictly dominated (given the deletion of the original dominated strategy) will also now disappear.

There is an important distinction here between strict and weak domination. It is not true that weakly dominated strategies are similarly eliminated. Consider the game taken from Samuelson and Zhang (1992, p. 382) in Figure 9. It is worth noting that this is the normal form of the extensive form game with perfect information given in Figure 10.

There are two populations, with agents in population 1 playing the role of player 1 and agents in population 2 playing the role of player 2. Any state in which all agents in population 2 play *L* is Liapunov stable: Suppose the state starts with almost all agents in population 2 playing *L*. There is then very little incentive for agents in the first population playing *B* to change behavior, since *T* is only marginally better (moreover, if the game played is in fact the extensive

and Weibull (1995) for general results on continuous time dynamics.




Figure 10: An extensive form game with the normal form in Figure 9.

form of Figure 10, then agents in population 1 only have a choice to make if they are matched with one of the few agents choosing R). So, dynamics will not move the state far from its starting point, if the starting point has mostly L -playing agents in population 2. If we model the dynamics for this game as a two-dimensional continuous time replicator dynamic, we have³²

$$\frac{dp^t}{dt} = p^t \times (1 - p^t)(1 - q^t), \text{ and}$$

$$\frac{dq^t}{dt} = q^t \times (1 - q^t),$$

where p^t is the fraction of population 1 playing T and q^t is the fraction of population 2 playing L . The adjustment of the fraction of population 2 playing L reflects the strict dominance of L over R : Since L always does better than R , if

³²The two-population version of (1) is:

$$\begin{aligned} dp_i^t/dt &= p_i^t \cdot (\pi_1(i, q^t) - \pi_1(p^t, q^t)), \text{ and} \\ dq_j^t/dt &= q_j^t \cdot (\pi_2(p^t, j) - \pi_2(p^t, q^t)), \end{aligned}$$

where $\pi_k(i, j)$ is player k 's payoff from the strategy pair (i, j) and $\pi_k(p, q) = \sum_{i,j} \pi_k(i, j)p_iq_j$.

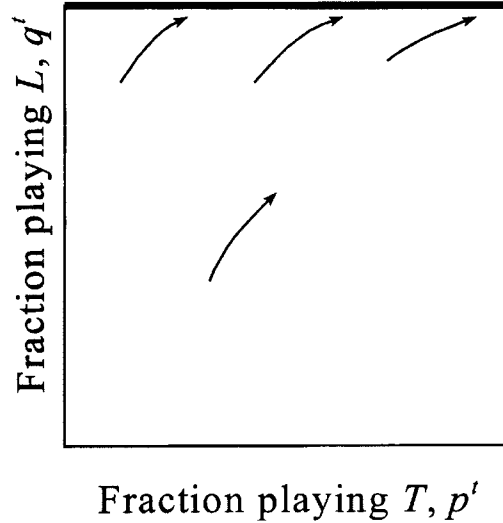


Figure 11: The phase diagram for the game with weak dominance.

q^t is interior (that is, there are both agents playing L and R in the population) the fraction playing L increases ($dq^t/dt > 0$), independently of the fraction in population 1 playing T , with the adjustment only disappearing as q^t approaches 1. The adjustment of the fraction of population 1 playing T , on the other hand depends on the fraction of population 2 playing R : if almost all of population 2 is playing L (q^t is close to 1), then the adjustment in p^t is small (in fact, arbitrarily small for q^t arbitrarily close to 1), no matter what the value of p^t . The phase diagram is illustrated in Figure 11.

It is also important to note that no rest point is asymptotically stable. Even the state with all agents in population 1 playing T and all agents in population 2 playing L is not asymptotically stable, because a perturbation that increases the fraction of population 2 playing R while leaving the entire population 1 playing T will not disappear: the system has been perturbed towards another rest point. Nothing prevents the system from “drifting” along the heavy lined horizontal

segment in Figure 11. Moreover, each time the system is perturbed from a rest point toward the interior of the state space (so that there is a positive fraction of population 1 playing B and of population 2 playing R), it typically returns to a different rest point (and for many perturbations, to a rest point with a higher fraction of population 2 playing L). This logic is reminiscent of the intuition given earlier on for the nonexistence of ESS in asymmetric games. It also suggests that *sets of states* will have better stability properties than individual states.

Recall that if a single strategy profile is “evolutionarily stable” then behavior within the population, once near that profile, converges to it and never leaves it. A single strategy profile describes the aggregate pattern of behavior within the population. A set of strategy profiles is a *collection* of such descriptions. Loosely, we can think of a set of strategy profiles as being “evolutionarily stable” if behavior within the population, once near any profile in the set, converges to some profile within it and never leaves the set. The important feature is that behavior within the population need not settle down to a steady state, rather it can “drift” between the different patterns within the “evolutionarily stable” set.

5.2. The ultimatum game and backward induction

The ultimatum game is a simple game with multiple Nash equilibria and a unique backward induction outcome. There is a \$1 to divide. The proposer proposes a division to the responder. The responder either accepts the division, in which case it is implemented, or rejects it, in which case both players receive nothing. If the proposer can make any proposal, the only backward induction solution has the receiver accepting the proposal in which the proposer receives the entire dollar. If the dollar must be divided into whole pennies, there is another solution in which the responder rejects the proposal in which he receives nothing and accepts the proposal of 99 cents to the proposer and 1 cent to the responder. This prediction is uniformly rejected in experiments!

How do Gale et al. (1995) explain this? The critical issue is the relative speeds of convergence. Both proposers and responders are “learning.” The proposers are learning which offers will be rejected (this is learning as we have been discussing it). In principle, if the environment (terms of the proposal) is sufficiently complicated, responders may have difficulty evaluating offers. In experiments, responders do reject as much as 30 cents. However, it is difficult to imagine that responders do not understand that 30 cents is better than zero. There are at least two possibilities that still allow the behavior of the responders to be viewed as learning. The first is that responders do not believe the rules of the game as described and take time to learn that the proposer really was making a take-

it-or-leave-it offer. As mentioned in footnote 6, most people may not be used to take-it-or-leave-it offers, and it may take time for the responders to properly appreciate what this means. The second is that the monetary reward is only one ingredient in responders' utility functions, and that responders must learn what the "fair" offer is.

If proposers learn sufficiently fast relative to responders, then there can be convergence to a Nash equilibrium that is not the backward induction solution. In the backward induction solution, the responder gets almost nothing, so that the cost of making an error is low, while the proposer loses significantly if he misjudges the acceptance threshold of the responder. In fact, in a simplified version of the ultimatum game, Nash equilibria in which the responder gets a substantial share are stable (although not asymptotically stable). In a non-backward induction outcome, the proposer never learns that he could have offered less to the responder (since he never observes such behavior). If he is sufficiently pessimistic about the responder's acceptance threshold, then he will not offer less, since a large share is better than nothing. Consider a simplified ultimatum game that gives the proposer and responder two choices: the proposer either offers even division or a small positive payment, and the responder only responds to the small positive payment (he must accept the equal division). Figure 12 is the phase diagram for this simplified game. While this game has some similarities to that in Figure 9, there is also an important difference. All the rest points inconsistent with the backward induction solution (with the crucial exception of the point labelled *A*) are Liapunov stable (but not asymptotically so). Moreover, in response to some infrequent perturbations, the system will effectively "move along" these rest points toward *A*. But once at *A* (unlike the corresponding stable rest point in Figure 11), the system will move far away. The rest point labelled *A* is not stable: Perturbations near can move the system onto a trajectory that converges to the backward induction solution.

While this seems to suggest that non-backward induction solutions are fragile, such a conclusion is premature. Since any model is necessarily an approximation, it is important to allow for drift. Binmore and Samuelson (1996) use the term *drift* to refer to unmodelled small changes in behavior. One way of allowing for drift would be to add to the learning dynamic an additional term reflecting a deterministic flow from one strategy to another that was independent of payoffs. In general, this drift would be small, and in the presence of strong incentives to learn, irrelevant. However, if players (such as the responders above) have little incentive to learn, then the drift term becomes more important. In fact, adding an arbitrarily small uniform drift term to the replicator dynamic changes the dynamic properties in a fundamental way. With no drift, the only asymptotically

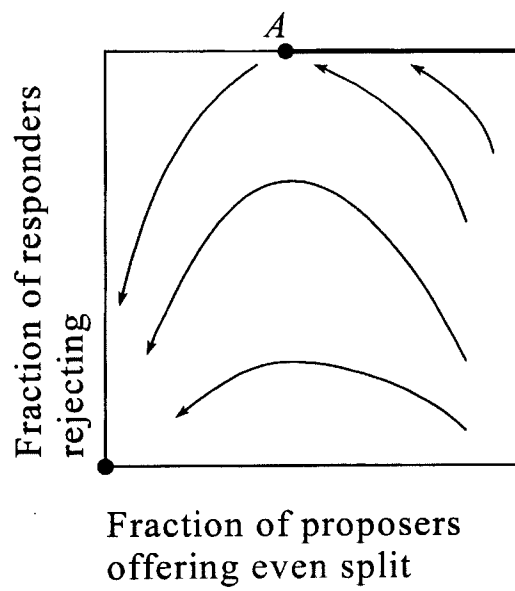


Figure 12: A suggestive phase diagram for a simplified version of the ultimatum game.

stable rest point is the backward induction solution. With drift, there can be another asymptotically stable rest point near the non backward-induction Nash equilibria.

The ultimatum game is special in its simplicity. Ideas of backward induction and sequential rationality have been influential in more complicated games (like the centipede game, the repeated prisoner's dilemma, and alternating offer bargaining games). In general, backward induction has received little support from evolutionary game theory for more complicated games (see, for example, Cressman and Schlag (1995), Nöldeke and Samuelson (1993), and Ponti (1996)).

5.3. Forward induction and efficiency

In addition to backward induction, the other major refinement idea is forward induction. In general, forward induction receives more support from evolutionary arguments than does backward induction. The best examples of this are in the context of cheap talk (starting with Matsui (1991)). Some recent papers are Bhaskar (1995), Blume, Kim and Sobel (1993), Kim and Sobel (1995), and Wärneryd (1993)—see Sobel (1993) and Samuelson (1993) for a discussion.

Forward induction is the idea that actions can convey information about the future intentions of players even off the equilibrium path. Cheap talk games (signaling games in which the messages are costless) are ideal to illustrate these ideas. Cheap talk games have both revealing equilibria (cheap talk can convey information) and the so-called babbling equilibria (messages do not convey any information because neither the sender nor the receiver expects them to), and forward induction has been used to eliminate the babbling equilibria.

Consider the following cheap talk game: There are two states of the world, *rain* and *sun*. The sender knows the state and announces a message, *rain* or *sun*. On the basis of the message, the receiver chooses an action, *picnic* or *movie*. Both players receive 1 if the receiver's action agrees with the state of the world (i.e., *picnic* if *sun*, and *movie* if *rain*), and 0 otherwise. Thus, the sender's message is payoff irrelevant and so is "cheap talk". The obvious pattern of behavior is for the sender to signal the state by making a truthful announcement in each state and for the receiver to then choose the action that agrees with the state. In fact, since the receiver can infer the state from the announcement (if the announcement differs across states), there are two *separating* equilibrium profiles (truthful announcing, where the sender announces *rain* if *rain* and *sun* if *sun*; and false announcing, where sender announces *sun* if *rain* and *rain* if *sun*).

A challenge for traditional non-cooperative theory, however, is that babbling is also an equilibrium: The sender places equal probability on *rain* and *sun*, in-

dependent of the state. The receiver, learning nothing from the message, places equal probability on *rain* and *sun*. Consider ESS in the symmetrized game (where each player has equal probability of being the sender or receiver). It turns out that only separating equilibria are ESS. The intuition is in two parts. First, the babbling equilibrium is not an ESS: Consider the truthful entrant (who announces truthfully and responds to announcements by choosing the same action as the announcement). The payoff to this entrant is strictly greater than that of the babbling strategy against the perturbed population (both receive the same payoff when matched with the babbling strategy, but the truthful strategy does strictly better when matched with the truthful strategy). Moreover, the separating equilibria are strict equilibria, and so, ESSs. Suppose, for example, all players are playing the truthful strategy. Then any other strategy must yield strictly lower payoffs: Either, as a sender, the strategy specifies an action conditional on a state that does not correspond to that state, leading to an incorrect action choice by the truthful receiver, or, as a responder, the strategy specifies an incorrect action after a truthful announcement).

This simple example is driven by the crucial assumption that the number of messages equals the number of states. If there are more messages than states, then there are no strict equilibria, and so no ESSs. To obtain similar efficiency results for a larger class of games, we need to use set-valued solution concepts, such as cyclical stability (Gilboa and Matsui (1991)) used by Matsui (1991) and equilibrium evolutionary stability (Swinkels (1992)) used by Blume et al. (1993). Matsui (1991) and Kim and Sobel (1995) study coordination games with a pre-play round of communication. In such games, communication can allow players to coordinate their actions. However, as in the above example, there are also babbling equilibria, so that communication appears not to guarantee coordination. Evolutionary pressures, on the other hand, destabilize the babbling equilibria. Blume et al. (1993) study cheap talk signaling games like that of the example. Bhaskar (1995) obtains efficiency with noisy pre-play communication, and shows that, in his context at least, the relative importance of noise and mutations is irrelevant.

The results in this area strongly suggest that evolutionary pressures can destabilize inefficient outcomes. The key intuition is that suggested by the example above. If an outcome is inefficient, then there is an entrant that is equally successful against the current population, but that can achieve the efficient outcome when playing against a suitable opponent. A crucial aspect is that the model allow the entrant to have sufficient flexibility to achieve this, and that is the role of cheap talk above.

		column player	
		H	L
row player	H	100, 100	0, 0
	L	0, 0	1, 1

Figure 13: (H, H) seems more likely than (L, L)

5.4. Multiple strict equilibria

Multiple best replies for a player raise the question of determining which of these best replies are “plausible” or “sensible.” The refinements literature (of which backward and forward induction are a part) attempted to answer this question, and by so doing eliminate some equilibria as being uninteresting. Multiple strict equilibria raise a completely new set of issues. It is also worth recalling that any strict equilibrium is asymptotically stable under any monotonic dynamic. I argued earlier that this led to the desirable feature of history dependence. However, even at an intuitive level, some strict equilibria are more likely. For example, in the game described by Figure 13, (H, H) seems more likely than (L, L) . There are several ways this can be phrased. Certainly, (H, H) seems more “focal,” and if asked to play this game, I would play H , as would (I suspect) most people. Another way of phrasing this is to compare the basins of attraction of the two equilibria under monotone dynamics (they will all agree in this case, since the game is so simple),³³ and observe that the basin of attraction of (H, H) is 100-times the size of (L, L) . If we imagine that the initial condition is chosen randomly, then the H pattern of behavior is 100 times as likely to arise. I now describe a more recent perspective that makes the last idea more precise by eliminating the need to specify an initial condition.

The motivation for ESS and (asymptotic) stability was a desire for robustness to a *single* episode of perturbation. It might seem that if learning operates at a sufficiently higher rate than the rate at which new behavior is introduced into the population, focusing on the dynamic implications of a single perturbation is reasonable. Foster and Young (1990) have argued that the notions of an ESS and attractor of the replicator dynamic do not adequately capture long-run stability when there are continual small stochastic shocks. Young and Foster (1991) describe simulations and discuss this issue in the context of Axelrod (1984)’s computer tournaments.

³³A state is in the basin of attraction of an equilibrium if, starting at that state and applying the dynamic, eventually the system is taken to the state in which all players play the equilibrium strategy.

There is a difficulty that must be confronted when explicitly modelling randomness. As I mentioned above, the standard replicator dynamic story is an idealization for large populations (specifically, a continuum). If the mutations-experimentation occurs at the individual level, there should be no aggregate impact; the resulting evolution of the system is deterministic and there are no “invasion events.” There are two ways to approach this. One is to consider aggregate shocks (this is the approach of Foster and Young (1990) and Fudenberg and Harris (1992)). The other is to consider a finite population and analyze the impact of individual experimentation. Kandori, Mailath and Rob (1993) study the implications of randomness on the individual level. In addition to emphasizing individual decision making, this paper analyzes the simplest model that illustrates the role of perpetual randomness.

Consider again the stag hunt game, and suppose each work-team consists of two workers. Suppose, moreover, that the firm has N workers. The relevant state variable is z , the number of workers who choose high effort. Learning implies that if high effort is a better strategy than low effort, then at least some workers currently choosing low effort will switch to high effort, i.e., $z^{t+1} > z^t$ if $z^t < N$. A similar property holds if low effort is a better strategy than high. The learning or selection dynamics describe, as before, a dynamic on the set of population states, which is now the number of workers choosing high effort. Since we are dealing with a finite set of workers, this can also be thought of as a Markov process on a finite state space. The process is Markov, because, by assumption workers only learn from last period’s experience. Moreover, both all workers choosing high effort and all workers choosing low effort are absorbing states of this Markov process.³⁴ This is just a restatement of the observation that both states correspond to Nash equilibria.

Perpetual randomness is incorporated by assuming that, in each period, each worker independently switches his effort choice (i.e., experiments) with probability ε , where $\varepsilon > 0$ is to be thought of as small. In each period, there are two phases: the learning phase and the experimentation phase. Note that after the experimentation phase (in contrast to the learning phase), with positive probability, fewer workers may be choosing a best reply.

Attention now focuses on the behavior of the Markov chain with the perpetual randomness. Because of the experimentation, every state is reached with positive probability from any other state (including the states in which all workers choose high effort and all workers choose low effort). Thus, the Markov chain is irreducible and aperiodic. It is a standard result that such a Markov chain has a unique stationary distribution. Let $\mu(\varepsilon)$ denote the stationary distribution. The

³⁴An absorbing state is a state that the process, once in, never leaves.

goal has been to characterize the limit of $\mu(\varepsilon)$ as ε becomes small. This, if it exists, is called the *stochastically stable distribution* (Foster and Young (1990)) or the *limit distribution*. States in the support of the limit distribution are sometimes called *long-run equilibria* (Kandori et al. (1993)). The limit distribution is informative about the behavior of the system for positive but very small ε . Thus, for small degrees of experimentation, the system will spend almost all of its time with all players choosing the same action. However, every so often (but infrequently) enough players will switch action, which will then switch play of the population to the other action, until again enough players switch.

It is straightforward to show that the stochastically stable distribution exists; characterizing it is more difficult. Kandori et al. (1993) is mostly concerned with the case of 2×2 symmetric games with two strict symmetric Nash equilibria. Any monotone dynamic divides the state space into the same two basins of attraction of the equilibria. The *risk dominant equilibrium* (Harsanyi and Selten (1988)) is the equilibrium with the larger basin of attraction. The risk dominant equilibrium is “less risky” and may be Pareto dominated by the other equilibrium (the risk dominant equilibrium results when players choose best replies to beliefs that assign equal probability to the two possible actions of the opponent). In the stag hunt example, the risk dominant equilibrium is low effort and it is Pareto dominated by high effort. In Figure 13, the risk dominant equilibrium is H and it Pareto dominates L . Kandori et al. (1993) shows that the limit distribution puts probability one on the risk dominant equilibrium. The non-risk dominant equilibrium is upset because the probability of a sufficiently large number of simultaneous mutations that leave society in the basin of attraction of the risk dominant equilibrium is of higher order than that of a sufficiently large number of simultaneous mutations that cause society to leave the basin of attraction.

This style of analysis allows us to formally describe strategic uncertainty. To make this point, consider again the workers involved in team production as a stag hunt game. For the payoffs as in Figure 1, risk dominance leads to the low effort outcome even if each work team has only two workers. Suppose, though, the payoffs are as in Figure 14, with V being the value of high effort (if reciprocated). For $V > 6$ and two-worker teams, the principles of risk dominance and payoff dominance agree: high effort. But now suppose each team requires n workers. While this is no longer a two player game, it is still true that (for large populations) the size of the basin of attraction is the determining feature of the example. For example if $V = 8$, high effort has a smaller basin of attraction than low effort for all $n \geq 3$. As V increases, the size of the team for which high effort becomes too risky increases.³⁵ This attractively captures the idea that co-

³⁵High effort has the larger basin of attraction if $(1/2)^n > 3/V$.

		minimum of other workers' efforts	
		high	low
worker's effort	high	V	0
	low	3	3

Figure 14: A new “stag-hunt” played by workers in a team.

operation in large groups can be harder than cooperation in small groups just due to the uncertainty that everyone will cooperate. While a small possibility of noncooperation by any one agent is not destabilizing in small groups (since cooperation is a strict equilibrium), it is in large ones.

In contrast to both ESS and replicator dynamics, there is a unique outcome. History does not matter. This is the result of taking two limits: first, time is taken to infinity (which justifies looking at the stationary distribution), and *then* the probability of mutation is taken to zero (looking at small rates). The rate at which the stationary distribution is approached from an arbitrary starting point is decreasing in population size (since the driving force is the probability of a simultaneous mutation by a fraction of the population). Motivated by this observation, Ellison (1993) studied a model with local interactions that has substantially faster rates of convergence. The key idea is that, rather than playing against a randomly drawn opponent from the entire population, each player plays only against a small number of neighbors. The neighborhoods are overlapping, however, so that a change of behavior in one neighborhood can (eventually) influence the entire population.

Since it is only for the case of 2×2 symmetric games that the precise modelling of the learning dynamic is irrelevant, extensions to larger games require specific assumptions about the dynamics. Kandori and Rob (1995), Nöldeke and Samuelson (1993), and Young (1993) generalize the best reply dynamic in various directions.

Kandori et al. (1993), Young (1993), and Kandori and Rob (1995) study games with strict equilibria, and (as the example above illustrates) the relative magnitudes of probabilities of simultaneous mutations are important. In contrast, Samuelson (1994) studies normal form games with nonstrict equilibria and Nöldeke and Samuelson (1993) studies extensive form games. In these cases, since the equilibria are not strict, states that correspond to equilibria can be upset by a single mutation. This leads to the limit distribution having nonsingleton support. This is the reflection in the context of stochastic dynamics of the issues illustrated by discussion of Figures 9, 10, and 12. In general, the support will contain “con-

nected” components, in the sense that there is a sequence of single mutations from one state to another state that will not leave the support. Moreover, each such state will be a rest point of the selection dynamic. The results on extensive forms are particularly suggestive, since different points in a connected component of the support correspond to different specifications of off-the-equilibrium path behavior.

The introduction of stochastic dynamics does not, by itself, provide a general theory of equilibrium selection. Bergin and Lipman (1996) show that allowing the limiting behavior of the mutation probabilities to depend on the state gives a general possibility theorem: Any strict equilibrium of any game can be selected by an appropriate choice of state-dependent mutation probabilities. In particular, in 2×2 games, if the risk dominant strategy is “harder” to learn than the other (in the sense that the limiting behavior of the mutation probabilities favors the non-risk dominant strategy), then the risk dominant equilibrium will not be selected. On the other hand, if the state dependence of the mutation probabilities only arises because the probabilities depend on the difference in payoffs from the two strategies, the risk dominant equilibrium is selected (Blume (1994)). This latter state dependence can be thought of as being strategy neutral in that it only depends on the payoffs generated by the strategies, and not on the strategies themselves. The state dependence that Bergin and Lipman (1996) need for their general possibility result is perhaps best thought of as strategy dependence, since the selection of some strict equilibria only occurs if players find it easier to switch to (learn) certain strategies (perhaps for complexity reasons).

Binmore, Samuelson and Vaughan (1995), who study the result of the selection process itself being the source of the randomness, also obtain different selection results. Finally, the matching process itself can also be an important source of randomness, see Young and Foster (1991) and Robson and Vega-Redondo (1996).

6. Conclusion

The result that any asymptotically stable rest point of an evolutionary dynamic is a Nash equilibrium is an important result. It shows that there are primitive foundations for equilibrium analysis. However, for asymmetric games, asymptotic stability is effectively equivalent to strict equilibria (which do not exist for many games of interest). To a large extent, this is due to the focus on individual states. If we instead consider sets of states (strategy profiles), as I discussed at the end of Section 5.1, there is hope for more positive results.³⁶

³⁶Sets of strategy profiles that are asymptotically stable under plausible deterministic dynamics turn out also to have strong Kohlberg and Mertens (1986) type stability properties (Swinkels

The lack of support for the standard refinement of backward induction is in some ways a success. Backward induction has always been a problematic principle, with some examples (like the centipede game) casting doubt on its universal applicability. The reasons for the lack of support improve our understanding of when backward induction is an appropriate principle to apply.

The ability to discriminate between different strict equilibria and provide a formalization of the intuition of strategic uncertainty is also a major contribution of the area.

I suspect that the current evolutionary modelling is still too stylized to be used directly in applications. Rather, applied researchers need to be aware of what they are implicitly assuming when they do equilibrium analysis.

In many ways, there is an important parallel with the refinements literature. Originally, this literature was driven by the hope that theorists could identify the unique “right” equilibrium. If that original hope had been met, applied researchers need never worry about a multiplicity problem. Of course, that hope was not met, and we now understand that that hope, in principle, could never be met. The refinements literature still served the useful role of providing a language to describe the properties of different equilibria. Applied researchers find the refinements literature of value for this reason, even though they cannot rely on it mechanically to eliminate “uninteresting” equilibria. The refinements literature is currently out of fashion because there were too many papers in which one example suggested a minor modification of an existing refinement and no persuasive general refinement theory emerged.

There is a danger that evolutionary game theory could end up like refinements. It is similar in that there was a lot of early hope and enthusiasm. And, again, there have been many perturbations of models and dynamic processes, not always well motivated. As yet, the overall picture is still somewhat unclear.

However, on the positive side, important insights are still emerging from evolutionary game theory (for example, the improving understanding of when backward induction is appropriate and the formalization of strategic uncertainty). Interesting games have many equilibria, and evolutionary game theory is an important tool in understanding which equilibria are particularly relevant in different environments.

(1993)), in particular, the property of robustness to deletion of never weak best replies. This latter property implies many of the refinements that have played an important role in the refinements literature and signaling games, such as the intuitive criterion, the test of equilibrium domination, and D1 (Cho and Kreps (1987)). A similar result under different conditions was subsequently proved by Ritzberger and Weibull (1995), who also characterize the sets of profiles that can be asymptotically stable under certain conditions.

References

- Aoyagi, Masaki**, "Evolution of Beliefs and the Nash Equilibrium of Normal Form Games," *Journal of Economic Theory*, 1996, 70, 444–469.
- Aumann, Robert J.**, "On the Non-Transferable Utility Value: A Comment on the Roth-Shafer Examples," *Econometrica*, 1985, 53, 667–677.
- and **Adam Brandenburger**, "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 1995, 63, 1161–1180.
- Axelrod, Robert**, *The Evolution of Cooperation*, New York: Basic Books, 1984.
- Bergin, James and Barton L. Lipman**, "Evolution with State-Dependent Mutations," *Econometrica*, 1996, 64, 943–956.
- Bernheim, B. Douglas**, "Rationalizable Strategic Behavior," *Econometrica*, 1984, 52, 1007–1028.
- Bhaskar, V.**, "Noisy Communication and the Evolution of Cooperation," 1995. University of St. Andrews.
- Binmore, Kenneth G. and Larry Samuelson**, "Evolutionary Drift and Equilibrium Selection," 1996. Mimeo, University of Wisconsin.
- , — , and **Richard Vaughan**, "Musical Chairs: Modeling Noisy Evolution," *Games and Economic Behavior*, 1995, 11, 1–35.
- Blume, Andreas, Yong-Gwan Kim, and Joel Sobel**, "Evolutionary Stability in Games of Communication," *Games and Economic Behavior*, 1993, 5, 547–575.
- Blume, Lawrence**, "How Noise Matters," 1994. Mimeo, Cornell University.
- Börgers, Tilman and Rajiv Sarin**, "Learning Through Reinforcement and Replicator Dynamics," *Journal of Economic Theory*, 1997, 77, 1–14.
- Bush, Robert R. and Frederick Mosteller**, "A Mathematical Model for Simple Learning," *Psychological Review*, 1951, 58, 313–323.
- and — , *Stochastic Models of Learning*, New York: Wiley, 1955.
- Cho, In-Koo and David Kreps**, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 1987, 102, 179–221.

- Cressman, Ross and Karl H. Schlag**, “The Dynamic (In)Stability of Backwards Induction,” Technical Report, Wilfred Laurier University and Bonn University 1995.
- Darwin, Charles**, *The Life and Letters of Charles Darwin, Including an Autobiographical Chapter. Edited by his Son, Francis Darwin*, second ed., Vol. 1, Albemarle Street, London: John Murray, 1887.
- Dekel, Eddie and Faruk Gul**, “Rationality and Knowledge in Game Theory,” in David M. Kreps and Kenneth F. Wallis, eds., *Advances in Economics and Econometrics: Theory and Applications - Seventh World Congress of the Econometric Society, Vol. 1*, Cambridge University Press Cambridge 1997, pp. 87–172.
- Ellison, Glenn**, “Learning, Local Interaction, and Coordination,” *Econometrica*, 1993, *61*, 1047–1071.
- Elster, Jon**, “Social Norms and Economic Theory,” *Journal of Economic Perspectives*, 1989, *3*, 99–117.
- Foster, Dean and H. Peyton Young**, “Stochastic evolutionary game dynamics,” *Theoretical Population Biology*, 1990, *38*, 219–232.
- Friedman, Daniel**, “Evolutionary Games in Economics,” *Econometrica*, 1991, *59*, 637–666.
- Fudenberg, Drew and Christopher Harris**, “Evolutionary Dynamics with Aggregate Shocks,” *Journal of Economic Theory*, 1992, *57*, 420–441.
- and **David Kreps**, “A Theory of Learning, Experimentation, and Equilibrium,” Mimeo, MIT and Stanford 1989.
- and **David Levine**, “Steady state learning and Nash equilibrium,” *Econometrica*, 1993, *61*, 547–573.
- Gale, John, Kenneth G. Binmore, and Larry Samuelson**, “Learning to be Imperfect: The Ultimatum Game,” *Games and Economic Behavior*, 1995, *8*, 56–90.
- Gilboa, Itzhak and Akihiko Matsui**, “Social stability and equilibrium,” *Econometrica*, 1991, *59*, 859–867.
- Gul, Faruk**, “Rationality and Coherent Theories of Strategic Behavior,” *Journal of Economic Theory*, 1996, *70*, 1–31.

- Harsanyi, John C.**, *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge, UK: Cambridge University Press, 1977.
- and **Reinhard Selten**, *A General Theory of Equilibrium in Games*, Cambridge: MIT Press, 1988.
- Hines, W. G. S.**, “Three characterizations of population strategy stability,” *Journal of Applied Probability*, 1980, 17, 333–340.
- Hofbauer, Josef and Karl Sigmund**, *The Theory of Evolution and Dynamical Systems*, Cambridge: Cambridge University Press, 1988.
- Kalai, Ehud and Ehud Lehrer**, “Rational Learning leads to Nash equilibrium,” *Econometrica*, 1993, 61, 1019–1045.
- Kandori, Michihiro**, “Evolutionary Game Theory in Economics,” in David M. Kreps and Kenneth F. Wallis, eds., *Advances in Economics and Econometrics: Theory and Applications - Seventh World Congress of the Econometric Society, Vol. 1*, Cambridge University Press Cambridge 1997, pp. 243–277.
- and **Rafael Rob**, “Evolution of Equilibria in the Long Run: A General Theory and Applications,” *Journal of Economic Theory*, 1995, 65, 383–414.
- , **George J. Mailath**, and **Rafael Rob**, “Learning, Mutation, and Long Run Equilibria in Games,” *Econometrica*, 1993, 61, 29–56.
- Kim, Yong-Gwan and Joel Sobel**, “An Evolutionary Approach to Pre-Play Communication,” *Econometrica*, 1995, 63, 1181–1193.
- Kohlberg, Elon and Jean-Francois Mertens**, “On the Strategic Stability of Equilibria,” *Econometrica*, 1986, 54, 1003–1037.
- Kreps, David M.**, *A Course in Microeconomic Theory*, Princeton, NJ: Princeton University Press, 1990.
- , *Game Theory and Economic Modelling*, Oxford: Clarendon Press, 1990.
- Mailath, George J.**, “Introduction: Symposium on Evolutionary Game Theory,” *Journal of Economic Theory*, 1992, 57, 259–277.
- Matsui, Akihiko**, “Cheap-Talk and Cooperation in Society,” *Journal of Economic Theory*, 1991, 54, 245–258.
- Maynard Smith, John**, *Evolution and the Theory of Games*, Cambridge: Cambridge Univ. Press, 1982.

- and **G. R. Price**, “The Logic of Animal Conflict,” *Nature*, 1973, *246*, 15–18.
- Myerson, Roger B.**, *Game Theory: Analysis of Conflict*, Cambridge, MA: Harvard Univ. Press., 1991.
- Nöldeke, Georg and Larry Samuelson**, “An Evolutionary Analysis of Backward and Forward Induction,” *Games and Economic Behavior*, 1993, *5*, 425–454.
- Pearce, David**, “Rationalizable strategic behavior and the problem of perfection,” *Econometrica*, 1984, *52*, 1029–50.
- Ponti, Giovanni**, “Cycles of Learning in the Centipede Game,” Technical Report, University College London 1996.
- Ritzberger, Klaus and Jörgen W. Weibull**, “Evolutionary Selection in Normal-Form Games,” *Econometrica*, 1995, *63*, 1371–1399.
- Robson, Arthur J.**, “Evolutionary Game Theory,” in John Creedy, Jeff Borland, and Jürgen Eichberger, eds., *Recent developments in game theory*, Hants, England: Edward Elgar Publishing Limited, 1992.
- and **Fernando Vega-Redondo**, “Efficient Equilibrium Selection in Evolutionary Games with Random Matching,” *Journal of Economic Theory*, 1996, *70*, 65–92.
- Rousseau, Jean-Jacques**, “A Discourse on the Origin of Inequality,” in “The Social Contract and Discourses,” New York: Dutton, 1950. Translated by G.D.H. Cole.
- Samuelson, Larry**, “Recent Advances in Evolutionary Economics: Comments,” *Economics Letters*, 1993, *42*, 313–319.
- , “Stochastic Stability in Games with Alternative Best Replies,” *Journal of Economic Theory*, 1994, *64*, 35–65.
- and **Jianbo Zhang**, “Evolutionary Stability in Asymmetric Games,” *Journal of Economic Theory*, 1992, *57*, 363–391.
- Schelling, Thomas**, *The Strategy of Conflict*, Cambridge, MA: Harvard Univ. Press, 1960.

- Schlag, Karl H.**, "Why Imitate, and If So, How?," Mimeo, University of Bonn 1994.
- Selten, Reinhard**, "A Note on Evolutionary Stable Strategies in Asymmetric Animal Conflicts," *Journal of Theoretical Biology*, 1980, *84*, 93–101.
- Skyrms, Brian**, *Evolution of the Social Contract*, Cambridge: Cambridge University Press, 1996.
- Sobel, Joel**, "Evolutionary Stability and Efficiency," *Economics Letters*, 1993, *42*, 301–312.
- Sonsino, Doron**, "Learning to Learn, Pattern Recognition, and Nash Equilibrium," *Games and Economic Behavior*, 1997, *18*, 286–331.
- Sugden, Robert**, "Spontaneous Order," *Journal of Economic Perspectives*, 1989, *3*, 85–97.
- Swinkels, Jeroen M.**, "Evolutionary Stability with Equilibrium Entrants," *Journal of Economic Theory*, 1992, *57*, 306–332.
- , "Adjustment Dynamics and Rational Play in Games," *Games and Economic Behavior*, 1993, *5*, 455–484.
- Taylor, Peter D. and Leo B. Jonker**, "Evolutionary stable strategies and game dynamics," *Mathematical Biosciences*, 1978, *40*, 145–156.
- van Damme, Eric**, *Stability and Perfection of Nash Equilibria*, Berlin: Springer-Verlag, 1987.
- Wärneryd, Karl**, "Cheap Talk, Coordination, and Evolutionary Stability," *Games and Economic Behavior*, 1993, *5*, 532–546.
- Weibull, Jörgen W.**, *Evolutionary Game Theory*, Cambridge: MIT Press, 1995.
- Young, H. Peyton**, "The Evolution of Conventions," *Econometrica*, 1993, *61*, 57–84.
- , "The Economics of Conventions," *Journal of Economic Perspectives*, 1996, *10*, 105–122.
- and **Dean Foster**, "Cooperation in the Short and in the Long Run," *Games and Economic Behavior*, 1991, *3*, 145–156.

- Zeeman, E.**, “Population Dynamics from Game Theory,” in Zbigniew Nitecki and Clark Robinson, eds., *Global Theory of Dynamical Systems*, Vol. 819 of *Lecture Notes in Mathematics*, Berlin: Springer-Verlag, 1980, pp. 471–497.
- , “Dynamics of the evolution of animal conflicts,” *Journal of Theoretical Biology*, 1981, *89*, 249–270.
- Zermelo, Ernst**, “Über eine Anwendung der Mengenlehre auf die Theorie des Schachspeils,” *Proceedings of the Fifth International Congress of Mathematicians*, 1912, *II*, 501–504.