Lecture Notes in Economics and Mathematical Systems Vol. 324

GAME THEORETICAL FOUNDATIONS OF EVOLUTIONARY STABILITY by Bomze, I. M. and Pötscher, B. M., Springer-Verlag, New York, 1988.

THE THEORY OF EVOLUTION AND DYNAMICAL SYSTEMS by Hofbauer, J. and Sigmund, K., Cambridge University Press, Cambridge, 1988.

INTRODUCTION

The use of noncooperative game theory has now spread throughout economics. In the process, the standard equilibrium concept in noncooperative game theory, the Nash equilibrium, has been complemented by a host of relaxations and refinements. Recently, the concept of an 'evolutionarily stable strategy' has joined the ranks of these refinements. At the same time, evolutionary game theory, or the study of evolutionary games, has been growing in popularity.

Since 1982, the standard book-length reference in evolutionary game theory has been Maynard Smith. The flurry of recent activity in the field has created an acute need for another examination of evolutionary game theory. The books by Bomze and Pötscher (1989) and Hofbauer and Sigmund (1988) are among the first attempts to systematically and rigorously examine the theory of evolutionary games.

In this review, I shall first introduce the basic notions of evolutionary game theory. The third section then identifies the issues encountered in extending these notions, which arose in biology as models of evolution, to human interactions. This provides the background for a consideration of the two books in the fourth section. The final section comments on open problems in evolutionary game theory.

ORIGINS

The terms 'evolutionary games' and 'evolutionary game theory' are misnomers, especially when one is considering economic applications. It is more descriptive to speak of the evolutionary *approach* to games. Any game can be examined with evolutionary techniques, with the label 'evolutionary game' being attached to a game that is so examined. The question of when a game is an evolutionary game then depends upon how the analyst treats the game rather than on characteristics of the game.

Evolutionary game theory has its origins in biology.[1] Here, the analysis is better described as the 'game theory of evolution' rather than the 'evolutionary theory of games'. This section provides a brief look at this biological foundation.

Consider a population of animals who are territorial. Occasionally, pairs of these animals meet in a contest in which each vies for possession of a particularly advantageous territory, perhaps because the territory has plentiful sources of food. Assume that there are a number of alternative courses of action or strategies available in these contests, represented by the elements of a finite set S. Let $\pi(s_1, s_2)$ be the payoff to an animal who plays s_1 against an opponent who plays s_2. A strategy here represents not a conscious course of action but rather a behavior that is genetically determined. 'Payoff' refers to increases in Darwinian fitness, or number of offspring, where reproduction is asexual and offspring are assumed to play the same strategy as their parents. Assuming that the behavior of the animals is shaped by evolution, what type of behavior is likely to appear?

Two approaches to this question arise. First, let ΔS be the set of probability mixtures over S. We can then interpret elements of ΔS as mixed strategies. Suppose that there exists a (possibly mixed) strategy $s^* \in \Delta S$ such that

$$\pi(s^*, s^*) \geq \pi(s, s^*) \quad \forall s \in \Delta S \tag{1}$$

$$\pi(s^*, s^*) = \pi(s, s^*) \Rightarrow \pi(s^*, s) > \pi(s, s) \quad \forall s \in \Delta S,\ s \neq s^* \tag{2}$$

Then s^* is said to be an *evolutionary stable strategy*. Condition (1) states that s^* is a best reply to itself, and is hence a Nash equilibrium. Condition (2) is interpreted as a stability condition. In particular, suppose that all of the animals in the population are initially playing strategy s^*. Then suppose that a collection of mutants playing strategy s appears, for any $s \in S$. Then conditions (1) and (2) ensure that as long as the fraction of the population playing s is not too large, the average payoff to s will fall short of the average payoff to s^*. Animals playing s^* will then produce larger numbers of offspring than those playing s, and evolution will cause the proportion of the population playing s^* to increase and will cause the proportion playing s to decrease. This process will continue until the animals playing s have been eliminated from the population.

The interpretation of this result is that an evolutionarily stable strategy is robust to the appearance of mutants, where this is interpreted to mean that any mutant strategy that appears will be eliminated from the population. This result requires that mutants appear in relatively small numbers and that mutations appear relatively rarely. Both assumptions appear reasonable in a biological context.

It follows immediately from condition (1) that any ESS in a Nash equilibrium, and the presence of condition (2) ensures that ESS is a refinement of Nash equilibrium. An ESS can in fact be shown to be a relatively powerful refinement. For example, an ESS is a proper equilibrium (where a proper equilibrium in a normal form game induces a sequential equilibrium on any extensive form version of that game). A host of studies have appeared in the biological literature that use the concept of an ESS to explain various types of natural behavior.

The dynamic process by which the ESS drives mutants out of the population is not explicitly modeled. An alternative approach begins with a model of this process. Suppose again that the payoffs in a game represent rates of (asexual) reproduction, as is conventionally assumed in biological applications. Let the players in the population be repeatedly and randomly matched to play the game. Suppose further that no error occurs in genetic transmission between matchings, so that an offspring plays the same strategy as its parent. The proportions of the population playing various strategies will then adjust over time in response to payoff differences. Let us take the limit of this adjustment process as the size of the population increases to infinity and let us further use a continuous-time process to represent the evolutionary system. Then one can show that the adjustment process will be given by:

$$\frac{dx_{it}}{dt}\frac{1}{x_{it}} = \pi_{it} - \bar{\pi}_{it} \tag{3}$$

where x_{it} is the proportion of the population playing the ith strategy at time t; π_{it} is the expected or average payoff to the ith strategy at time t, and $\bar{\pi}_{it}$ is the average of such expected payoffs over strategies. The growth rate of the population proportion playing strategy i is thus given by the difference between i's payoff and the average payoff.

The dynamic system given by Eqn (3) is commonly referred to as the replicator dynamics. The attraction of these dynamics is that they can be traced to a precise biological model of reproduction and genetic transfer. It may initially appear as if the replicator dynamics makes no provision for mutation, and is hence a questionable model of evolution. This point is usually addressed by assuming that x_{i0} is positive for every pure strategy, so that every strategy is initially represented in the population. This initial condition combines with the continuity of the process specified in Eqn (3) to ensure that at each time t, every strategy is played by at least some proportion of the population, though this proportion may approach zero for some strategies.

Players are generally restricted to playing pure strategies in replicator models. Mixtures appear not as mixed strategies played by particular players but as population polymorphisms, or limiting outcomes in which a number of pure strategies are played by a positive proportion of the population.

Interest now turns to the limiting outcomes of the replicator dynamics. It can be shown that if a strategy is an ESS, then that strategy is asymptotically stable under the replicator dynamics. Hence, for any ESS s^* there exists a neighborhood of s^* with the property that if the population initially falls within that neighborhood, then the replicator dynamics for that population will converge to s^*. An ESS is thus a local attractor under the replicator dynamics. Under certain (though not all) conditions, the converse of this statement holds. For example, if S contains just two strategies, then a limiting outcome under the replicator dynamics will be an ESS. If the replicator dynamics converge monotonically, then the limiting outcome will be an ESS.

Together these results form a convenient evolutionary theory. The replicator dynamics provides a model of the evolutionary process. The concept of an ESS is closely related to the limiting outcomes of the replicator dynamics, and is both a convenient equilibrium concept with which to work and a relatively powerful refinement of the Nash equilibrium concept.

This theory of evolutionary games is attractive to both biologists and economists, though for different reasons.

For biologists, evolutionary game theory is a tool with which to examine the effects of evolution. For economists, the concept of an ESS provides another potentially useful equilibrium concept.

Perhaps more importantly, evolutionary game theory presents economists with the prospect of constructing a consistent foundation for common game-theoretic solution concepts. Economists have recently turned attention to constructing decision-theoretic foundations for solution concepts such as Nash equilibrium and its refinements. The initial approach to this problem involved assuming that players are perfectly rational and that this perfect rationality is common knowledge. The results of this initial approach have been disappointing. On the one hand, it has been found that common knowledge of rationality motivates relatively weak solution concepts. For example, Tan and Werlang (1988) (see also Bernheim, 1984; Pearce, 1984) have shown that common knowledge of rationality is equivalent to assuming that players will not play strategies that are eliminated by the iterated elimination of strictly dominated strategies, or iterated strict dominance. Iterated strict dominance is a quite weak solution concept, being much less restrictive than Nash equilibrium, and the common-knowledge-of-rationality approach thus provides little help in motivating such equilibrium concepts as Nash equilibrium and its refinements. At the same time, work by Binmore (1987–8), Reny (1985) and Basu (1990) indicates that perfect rationality is a logically inconsistent concept. Together, these results suggest that economists will be unlikely to find a decision-theoretic foundation for common game-theoretic practice.

Evolutionary game theory potentially provides an alternative justification for game-theoretic outcomes. In an evolutionary model, one no longer must assert that the players are perfectly rational, and the paradoxes of common-knowledge-of-rationality arguments can thus be avoided. Instead, evolutionary game theory allows one to implicitly argue that agents' decision processes may be quite far from perfectly rational. The strategies that emerge from these decision processes, however, are directed by an evolutionary selection process toward certain outcomes. These outcomes will often exhibit characteristics, such as the playing of best replies, that might be interpreted as an implication of rationality.

This type of argument should be a familiar one to economists. Economists typically make assumptions concerning maximizing behavior, such as the assumption that firms maximize profits. When pressed, the usual defense of this assumption is that firms may not consciously maximize profits, but market forces cause those firms that come closest to maximizing profits to prosper and cause those that do not perform well to fail. This is implicitly an evolutionary argument, and evolutionary game theory simply makes this type of argument explicit for abstract games.

EVOLUTIONARY GAMES AND ECONOMICS

The model of evolutionary games presented above reflects the biological origins of the theory. If evolutionary game theory is to be applied in economics, several modifications in the theory or in its interpretation must be made. The most important of these modifications can be divided into two groups.

First, it appears quite reasonable in a biological context to assume that the agents in a population are genetically coded to play their strategies and to assume that the proportions of agents playing the various strategies are adjusted by differential rates of reproduction. These assumptions are less appropriate in an economic context. We may be unwilling to assume that players in an economic context are perfectly rational, but may also be unwilling to assume that players are unconscious automata that are genetically destined to take certain actions. A middle ground is presumably appropriate, in which players are either boundedly rational or are utility maximizers subject to the information they have available, which may be incomplete or incorrect. Players will then potentially switch strategies over time, and the proportions of the population playing various strategies will adjust, as players receive new information or learn to choose better strategies. Learning thus replaces selection as the driving force. In such a model, the replicator dynamics may no longer be appropriate.

There are two possible responses to the potential difficulties with the replicator dynamics. First, one can work with more general dynamic processes. For example, one might examine the class of dynamic processes that exhibit a monotonicity property which calls for the proportion of the population playing strategy s to grow faster than that playing strategy s' if the average payoff to s exceeds that to s'. At best, however, this is a stop-gap measure. At some point, explicit models must be constructed of the process by which players learn about and choose strategies, and then a dynamic process extracted from this model. Work is currently progressing on such models. The important observation is that the useful application of evolutionary game theory to economics will require that attention be devoted to modeling the origins and characteristics of the dynamic process by which players move between strategies.

The second modification required to apply evolutionary game theory to economics arises out of the fact that most results in evolutionary game theory hold for symmetric, normal form games. A normal form game $(\{1, 2\}, S_1, S_2, \pi_1, \pi_2)$ is symmetric if $S_1 = S_2$, if

$$\pi_1(s_1, s_2) = \pi_2(s_2, s_1) \quad (4)$$

and if players are unable to tell if they are players 1 or 2 when choosing strategies. These symmetry assumptions are implicitly built into the definition of an ESS given in conditions (1) and (2), where the fact that players 1 and 2 have identical strategy sets and cannot identify which role in the game they occupy allows us to model the game as if it were played by a single population of players choosing from a single strategy set.

It may be appropriate in many biological applications to limit attention to symmetric games, though even here asymmetries can arise. In economics, it seems clear that asymmetric games are often important. The conventional method for dealing with an asymmetric game is to create a larger, symmetric game in which each player is potentially ready to play either of the roles 1 or 2 in the original game. In any play of the game, nature first assigns players to roles and then players choose the action specified for that role by their strategy. This larger game with nature is a symmetric game and the conventional definition of an ESS can be applied.

This may appear to eliminate the symmetry problem. It can be shown, however, that an ESS in the larger, symmetric game is equivalent to a strict Nash equilibrium in the original game (Selten, 1980). Strict Nash equilibria are well known to satisfy virtually every refinement of Nash equilibrium, but they also often fail to exist. A theory that prescribes that attention be limited to strict Nash equilibria is of limited interest. Alternative means of dealing with asymmetric games must then be found.

One apparently promising alternative approach to asymmetric games is to explicitly model the behavior of dynamic evolutionary processes on such games. It appears as if this type of approach is necessary if progress is to be made on asymmetric games. Work in this area is in its early stages.

In addition to addressing primarily symmetric games, existing work on evolutionary games has focused on normal form games. This again may be appropriate for biological applications, but many interesting economic issues appear to be best modeled as extensive form games. Selten (1983, 1988) provides the first step in extending evolutionary game theory to extensive form games. In the process, Selten shows that this extension encounters formidable difficulties. One of the key features of Selten's work on extensive form games is the introduction of an equilibrium concept designed for extensive form games and denoted the 'limit ESS'. In order to assess this equilibrium concept and make further progress on extensive form games, it appears again as if work with dynamic processes is required.

Each time a difficulty is encountered in extending evolutionary game theory to economics, one possible solution is offered by pursuing more work with dynamic evolutionary models. With the help of this background, we now turn to the two volumes under consideration.

TWO TREATMENTS OF EVOLUTIONARY GAME THEORY

Given the rapidity with which articles on evolutionary game theory have been appearing, both in biology and economics, comprehensive treatments that could provide some unification are badly needed. One then approaches these volumes with high hopes. These hopes are all the more encouraged by the observation that these volumes have an impressive collection of authors. Their previous works would appear repeatedly in any useful bibliography of recent work in evolutionary game theory.

Four features of these volumes are immediately apparent. First, they are not mere compilations of previously existing results. Each volume contains significant original results. For this reason alone, they will become indispensable parts of the literature. Second, the two volumes are complements and not substitutes. The focus of the two volumes is quite different and there is remarkably little overlap in coverage. Third, each book involves some sophisticated mathematics. This is not to say that one must be a mathematician to follow the analysis, but one must be comfortable with rigorous mathematical argument. Finally, some of those working in evolutionary games will find one of the books useful and others will find the other useful, depending upon the problems examined and techniques involved, but virtually everyone working in evolutionary games will find at least one of the volumes useful.

It is convenient to start with the differences between the two volumes. Hofbauer and Sigmund is the longer of the two books by a factor of more than two, and consequently contains significantly more material. Bomze and Pötscher is much narrower in focus. Hofbauer and Sigmund offer a heavy dose of the biological foundations of evolutionary game theory and of dynamic modeling. Bomze and Pötscher offer only a mention of the biological foundations of the subject and only a very brief treatment of dynamic models. Bomze and Pötscher's analysis proceeds by defining solution concepts, observing that these concepts have an evolutionary flavor, and then investigating the properties of these concepts. Hofbauer and Sigmund, in contrast, discuss solution concepts only as one way of describing the limiting outcomes of the dynamic systems they examine. Much of the material in Bomze and Pötscher appears because of a desire to treat infinite games and games in which expected payoffs may not be linear functions of strategies. Each of these is somewhat of a departure from conventional game theory, and Hofbauer and Sigmund do not examine such issues.

Hofbauer and Sigmund devote considerable attention to biological foundations of evolution. Much of the first third of their book concerns these foundations. From these biological foundations, they construct dynamic models of the evolutionary process and examine the limiting outcomes of these processes. The phrase 'evolutionary game' does not appear in Hofbauer and Sigmund's title. Their interest is in dynamic models of evolution, and results for game theory appear only as a byproduct.

The beauty of Hofbauer and Sigmund's analysis arises out of the interweaving of biological and dynamic considerations. This appears early in the book. Chapter 3 is devoted to the 'fundamental theorem', which states that if a dynamic system represents a process of biological selection, then the average fitness of a population governed by that system must increase over time. Variations on this theme, involving alternative specifications of the model generating the dynamic process and the characteristics of the latter, recur throughout the book. These results constitute an important step in confronting the criticism that evolutionary models are *ad hoc*. Hofbauer and Sigmund masterfully demonstrate that they need not be *ad hoc*, and that one can work with models with rigorously specified foundations. Their specific results as well as the example they set in pursuing this type of analysis will have a profound effect on evolutionary game theory.

There are some minor shortcomings with the book. The index is short and will often present an obstacle to those using the book for reference. While the volume deals in mathematically sophisticated analysis, the presentation could benefit in places from being more precise.

For example, it is often difficult to identify when proofs begin and end. The concept of an orbit first appears on page 14, but is defined on page 42. It is often difficult to keep track of which assumptions, especially which symmetry assumptions, are maintained. Some of the remarks concerning the especially interesting concept of a Nash–Pareto pair (e.g. page 287) are especially hard to follow.

These are matters of style. If one is forced to criticize the book on matters of substance, one can only resort to the observation that it would be nice to have the book cover even more topics than it does. For example, very little attention is devoted to establishing the relationship between the outcome of the evolutionary models constructed in this volume and traditional game-theoretic solution concepts. Such links are especially interesting for those who have turned to evolutionary game theory out of an interest in foundations of game theory. In addition, much of the book is concerned with symmetric games, or games satisfying Eqn (4), and many results hold for the even more restrictive case of symmetric games in which it is also the case that

$$\pi_1(s, s') = \pi_1(s', s) \qquad (5)$$

Both assumptions impose limitations when considering economic applications.

This last paragraph constitutes not so much a criticism of Hofbauer and Sigmund as a recognition that it would always be nice to have a good analysis cover even more topics than it does and to have effective researchers share one's own concerns and interests. In particular, it appears that the application of evolutionary techniques to economics will require consideration of asymmetric games, causing one to wish more attention were devoted to this topic. At the same time, it appears as if progress in asymmetric games will require dynamic models. Hofbauer and Sigmund lay the groundwork for the construction of such models.

Hofbauer and Sigmund's text is complemented by an extensive collection of exercises. These range from simple extensions or verifications of material in the text to difficult problems, some of which develop important results. These exercises will prove invaluable for those using the book as part of a class on evolution.

The title of Bomze and Pötscher suggests a primary interest in evolutionary stability. Bomze and Pötscher also introduce a restriction of evolutionary stability which they refer to as 'uninvadability'. To present the notion of uninvadability, some notation is required. Following Bomze and Pötscher, let $s_\varepsilon(s')$ denote a population in which proportion ε of the agents play strategy s and proportion $1-\varepsilon$ play strategy s'. Then let $E(s, s_\varepsilon(s'))$ denote the average payoff to strategy s in a population characterized by $s_\varepsilon(s')$ and let $E(s', s_\varepsilon(s'))$ denote the average payoff to strategy s' in such a population. Let us now replace the definition of an ESS given in conditions (1) and (2) with the following. A strategy s^* is an evolutionarily stable strategy (ESS) if, for every $s \in S$ with $s \neq s^*$, there exists an $\varepsilon(s)$ such that

$$E(s^*, s_{\varepsilon(s)}(s^*)) > E(s, s_{\varepsilon(s)}(s^*)) \qquad (6)$$

Condition (6) thus requires that for every potential mutant strategy s, there exists an $\varepsilon(s)$ with the property that if the mutant appears as a proportion of the population smaller than $\varepsilon(s)$, then the average payoff to the mutant will fall short of that of the ESS. Selection will then presumably remove the mutant form of the population. Bomze and Pötscher show that under mild continuity conditions on the payoff function, condition (6) is equivalent to conditions (1) and (2). If strategy s satisfies conditions (1) and (2), then an $\varepsilon(s)$ can be found such that condition (6) holds; while condition (6) can hold for some $\varepsilon(s)$ only if condition (1) and (2) hold. This justifies the application of the term ESS to both conditions (1) and (2) and condition (6).

The value of ε involved in condition (6) is allowed to depend upon the mutant strategy s. If we place a uniformity condition on the value of ε, then we have the definition of uninvadability. In particular, s^* is uninvadable if there exists an ε such that for any $s \in S$ with $s \neq s^*$, we have

$$E(s^*, s_\varepsilon(s^*)) > E(s, s_\varepsilon(s^*)) \qquad (7)$$

It is clear that any uninvadable strategy is an ESS. In games in which S is finite, an ESS will be uninvadable. Bomze and Pötscher, however, present examples to show that an ESS need not be uninvadable in infinite games.

When the notions of ESS and uninvadability diverge, the latter appears to be the more interesting of the two because it establishes a stronger stability property. Bomze and Pötscher identify their primary objective as being the formulation of necessary and sufficient conditions for a strategy to be either an ESS or to be uninvadable. A variety of such conditions are offered, some of which rest on analytical arguments and some are geometric in nature. Some brief remarks are offered on the relationship between evolutionary stability and conventional solution concepts. Similar links for uninvadability are not established, though invadability is primarily of interest in infinite games, while many conventional solution concepts are most useful in finite games.

The innovation of Bomze and Pötscher is their extension of the theory of evolutionary games beyond finite games with bilinear payoff functions. This extension requires the concept of uninvadability, and the examination of the latter exposes the implicit assumptions that have been lurking in conventional studies of evolutionary stability. The authors exceed admirably in their quest to characterize uninvadability.

We can again note some difficulties with the exposition. The concept of a population 'state' appears to be unnecessary, and to simply introduce needless complication. The concept of a metagame is also somewhat awkward, and could be eliminated from the exposition. The book contains no material on biological origins and relatively little on dynamic processes. This is presumably forgivable in light of the author's interest in solution concepts. Given the latter, however, it would be helpful to have more attention devoted to the relationship between the concepts examined here, especially uninvadability, and standard solution concepts. Finally, this volume contains no index, making it all the harder to use as a reference work.

What can we say about the impact that these volumes are likely to have on evolutionary game theory, and especially on the use of evolutionary game theory in economics? This impact will be shaped by the fact that the books take quite different approaches to evolutionary game theory. The two volumes will accordingly be of use to researchers with different interests. Hofbauer and Sigmund will become a standard reference for those interested in the biological foundations of evolutionary games. It will also be a standard reference for those interested in serious examinations of dynamic evolutionary models. As indicated in the previous section, many of the current difficulties in evolutionary game theory are likely to be addressed by turning to dynamic models. Hofbauer and Sigmund is thus assured of playing a crucial role in the further development of evolutionary game theory. Its material will be an essential part of the standard techniques of an evolutionary game theorist, and the volume is destined to join Maynard Smith (1982) as one of the works with which one must be familiar if one is to work in evolutionary games.

The book by Bomze and Pötscher is harder to judge. This will be important if one is interested in treating infinities and nonlinearities. This is somewhat more speculative material than appears in Hofbauer and Sigmund and hence more difficult to assess. Bomze and Pötscher have laid the groundwork for extending evolutionary game theory into new areas; it remains to be seen what uses will be found for these new techniques.

FUTURE DIRECTIONS

The two volumes under consideration make contributions to our understanding of evolutionary game theory. Each will play an important role in the continued development of the theory. At the same time, these volumes leave unaddressed many of the issues raised in this review. This should not be taken as a criticism of these volumes but rather as an indication of the magnitude of the task that remains in order to effectively use evolutionary game theory in economics.

It is convenient to close by discussing one further issue that must be examined before evolutionary game theory becomes operational. What is the purview of evolutionary game theory, or what is the class of games to which it applies? The formal theory of evolutionary games directs our attention to the limiting outcomes of a sequence of repeated plays of the game. In practice, one does not often observe agents repeatedly playing a game. In light of this, to what games does evolutionary game theory apply?

One possible response is that applications of evolutionary game theory should be confined to cases in which games are literally repeatedly played. An opportunity for examining such games is provided by laboratory experiments, and Crawford (1990) and Miller and Andreoni (1990) have recently invoked evolutionary game theory to explain experimental outcomes. On the one hand, the success of evolutionary game theory in explaining these outcomes is encouraging. On the other, unless this success can be repeated in cases in which one does not have the luxury of arranging for a sequence of repeated plays, the importance of evolutionary game theory for economics is likely to be limited.

The other possible response is to recognize that economic agents may seldom literally play a game repeatedly, but these agents do constantly play games. It may be that these games can be sorted by their players into a small number of categories, with the games in a given category exhibiting common strategic features and with evolutionary processes acting on this sequence of similar games. The games in the evolutionary process would then not be identical but would be strategically equivalent. This interpretation of evolutionary game theory holds the promise for significantly expanding the purview of evolutionary theory. Constructing a more precise theory along these lines is an important priority for future work. If such a theory can be formulated, then evolutionary game theory promises to provide economists with a powerful analytical tool. The two volumes under consideration will serve as a useful point of departure for this investigation.

NOTE

1. A concise summary of the basic results in evolutionary game theory appears in van Damme (1987), from which this section occasionally borrows.

REFERENCES

K. Basu (1990). On the non-existence of a rationality definition for extensive games. *International Journal of Game Theory* **19**, 33–44.

B. D. Bernheim (1984). Rationalizable strategic behavior. *Econometrica* **52**, 1007–28.

K. G. Binmore (1987–8). Modelling rational players: I and II. *Economics and Philosophy* **3**, 179–215; **4**, 9–55.

V. P. Crawford (1990). An 'evolutionary' interpretation of Van Huyck, Battalio and Beil's experimental results on coordination. *Games and Economic Behavior* (forthcoming).

J. Maynard Smith (1982). *Evolution and the Theory of Games*, Cambridge: Cambridge University Press.

J. H. Miller and J. Andreoni (1990). A coevolutionary model of free riding behavior: replicator dynamics as an explanation of theoretical results. SSRI working paper 9010, University of Wisconsin.

D. G. Pearce (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* **52**, 1029–50.

P. Reny (1985). Rationality, common knowledge, and the theory of games. Mimeo, University of Western Ontario.

R. Selten (1980). A note on evolutionarily stable strategies in asymmetric animal conflicts. *Journal of Theoretical Biology* **84**, 93–101.

R. Selten (1983). Evolutionary stability in extensive two-person games. *Mathematical Social Sciences* **5**, 269–363.

R. Selten (1988). Evolutionary stability in extensive two-person games—correction and further development. *Mathematical Social Sciences* **16**, 223–66.

T. C.-C. Tan and S. R. da Costa Werlang (1988). The Bayesian foundations of solution concepts of games. *Journal of Economic Theory* **45**, 370–91.

E. van Damme (1987). *Stability and Perfection of Nash Equilibria*, New York: Springer-Verlag.

LARRY SAMUELSON
Department of Economics
University of Wisconsin
Madison, WIS 53706

–