Maynard Smith, J. and G.R. Price, 1973, The logic of animal conflict, Nature 246, 15–18.
Myerson, R.B., G.B. Pollock and J.M. Swinkels, 1991, Viscous population equilibria, Games and Economic Behavior 3, 101–109.
Nash, J.F., 1950, Non-cooperative games, Ph.D. Dissertation (Princeton University, Princeton, NJ).
Penrose, E.T., 1952, Biological analogies in the theory of the firm, American Economic Review 42, 804–819.
Selten, R., 1980, A note on evolutionary stable strategies in asymmetrical animal conflicts, Journal of Theoretical Biology 84, 93–101.
Selten, R., 1983, Evolutionary stability in extensive 2-person games, Mathematical Social Sciences 5, 269–363.
Selten, R., 1991, Evolution, learning and economic behavior, Games and Economic Behavior 3, 3–24.
Shapley, L.S., 1964, Some topics in two-person games, in: M. Dresher, L.S. Shapley and A.W. Tucker, eds., Advances in game theory (Princeton University Press, Princeton, NJ).
Swinkels, J.M., 1992, Evolutionary stability with equilibrium entrants, Journal of Economic Theory 57, 306–332.
Taylor, P.D. and L.B. Jonker, 1978, Evolutionary stable strategies and game dynamics, Mathematical Bioscience 40, 145–156.
Thomas, B., 1985, On evolutionary stable sets, Journal of Mathematical Biology 22, 105–115.
Van Damme, E., 1987, Stability and perfection of Nash equilibria. Second edition (Springer-Verlag, Berlin).
Vega-Redondo, F., 1993, On the evolution of cooperation in general games of common interest, Mimeo. (University of Alicante, Alicante).
Wärneryd, K.M., 1992, Evolutionary stability in unanimity games with cheap talk, Economics Letters 39, 295–300.
Wärneryd, K.M., 1993, Communication, complexity, and evolutionary stability, CentER DP 9313 (CentER for Economic Research, Tilburg).
Weibull, J., 1993, The 'As If' approach to game theory: 3 positive results, and 4 obstacles, Mimeo. (Stockholm University, Stockholm).
Young, P., 1993a, An evolutionary model of bargaining, Journal of Economic Theory 59, 145–168.
Young, P, 1993b, The evolution of conventions, Econometrica 61, 57–84.
Zeeman, E.C., 1981, Dynamics of the evolution of animal conflict, Journal of Theoretical Biology 89, 249–270.

ELSEVIER
European Economic Review 38 (1994) 859–867
EUROPEAN ECONOMIC REVIEW

Drift

Ken Binmore[a], Larry Samuelson[b]*

[a]Department of Economics, University College London, Gower Street, London WC1E 6BT, UK
[b]Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA

Abstract

Which equilibria should command our attention in a game with multiple equilibria? Attempts to solve this equilibrium selection problem have recently turned from the invention of equilibrium refinements to evolutionary models of games. Unfortunately, the basic equilibrium concept for evolutionary games, an evolutionarily stable strategy, frequently fails to exist. In this paper, we examine some of the modified and alternative stability concepts that have been applied when evolutionarily stable strategies do not exist.

Key words: Drift; Evolutionary stability; Evolutionarily stable strategy; Equilibrium
JEL classification: C72

1. Evolutionarily stable strategies

The Game. We shall conduct our discussion in the context of 'Game G', shown in Fig. 1. We refer to the participants in Game G as 'player one' and 'player two' or as 'role one' and 'role two'. Game G has three pure-strategy Nash equilibria, given by $(L\ell, L)$, (Lr, L) and $(R\ell, R)$. The first two are contained in a single component in which player one plays any mixture of $L\ell$ and Lr and player two plays L with probability at least 1/2. The third is contained in a component in which two plays R and player one mixes between $R\ell$ and Rr, with probability at least 1/3 on $R\ell$. Game G has a unique subgame perfect equilibrium, given by $(R\ell, R)$.

* Corresponding author. Financial support from National Science Foundation grant SES-9122176 and the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at the University of Bonn, is gratefully acknowledged. We thank Georg Nöldeke for helpful discussions, and thank Avner Shaked and the University of Bonn for their hospitality.

0014-2921/94/$07.00 © 1994 Elsevier Science B.V. All rights reserved
SSDI 0014-2921(93)E0111-W

Maynard Smith, J. and G.R. Price, 1973, The logic of animal conflict, Nature 246, 15–18.
Myerson, R.B., G.B. Pollock and J.M. Swinkels, 1991, Viscous population equilibria, Games and Economic Behavior 3, 101–109.
Nash, J.F., 1950, Non-cooperative games, Ph.D. Dissertation (Princeton University, Princeton, NJ).
Penrose, E.T., 1952, Biological analogies in the theory of the firm, American Economic Review 42, 804–819.
Selten, R., 1980, A note on evolutionary stable strategies in asymmetrical animal conflicts, Journal of Theoretical Biology 84, 93–101.
Selten, R., 1983, Evolutionary stability in extensive 2-person games, Mathematical Social Sciences 5, 269–363.
Selten, R., 1991, Evolution, learning and economic behavior, Games and Economic Behavior 3, 3–24.
Shapley, L.S., 1964, Some topics in two-person games, in: M. Dresher, L.S. Shapley and A.W. Tucker, eds., Advances in game theory (Princeton University Press, Princeton, NJ).
Swinkels, J.M., 1992, Evolutionary stability with equilibrium entrants, Journal of Economic Theory 57, 306–332.
Taylor, P.D. and L.B. Jonker, 1978, Evolutionary stable strategies and game dynamics, Mathematical Bioscience 40, 145–156.
Thomas, B., 1985, On evolutionary stable sets, Journal of Mathematical Biology 22, 105–115.
Van Damme, E., 1987, Stability and perfection of Nash equilibria. Second edition (Springer-Verlag, Berlin).
Vega-Redondo, F., 1993, On the evolution of cooperation in general games of common interest, Mimeo. (University of Alicante, Alicante).
Wärneryd, K.M., 1992, Evolutionary stability in unanimity games with cheap talk, Economics Letters 39, 295–300.
Wärneryd, K.M., 1993, Communication, complexity, and evolutionary stability, CentER DP 9313 (CentER for Economic Research, Tilburg).
Weibull, J., 1993, The 'As If' approach to game theory: 3 positive results, and 4 obstacles, Mimeo. (Stockholm University, Stockholm).
Young, P., 1993a, An evolutionary model of bargaining, Journal of Economic Theory 59, 145–168.
Young, P, 1993b, The evolution of conventions, Econometrica 61, 57–84.
Zeeman, E.C., 1981, Dynamics of the evolution of animal conflict, Journal of Theoretical Biology 89, 249–270.

ELSEVIER European Economic Review 38 (1994) 859–867 EUROPEAN ECONOMIC REVIEW

Drift

Ken Binmore[a], Larry Samuelson[b*]

[a]*Department of Economics, University College London, Gower Street, London WC1E 6BT, UK*
[b]*Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706, USA*

Abstract

Which equilibria should command our attention in a game with multiple equilibria? Attempts to solve this equilibrium selection problem have recently turned from the invention of equilibrium refinements to evolutionary models of games. Unfortunately, the basic equilibrium concept for evolutionary games, an evolutionarily stable strategy, frequently fails to exist. In this paper, we examine some of the modified and alternative stability concepts that have been applied when evolutionarily stable strategies do not exist.

Key words: Drift; Evolutionary stability; Evolutionarily stable strategy; Equilibrium
JEL classification: C72

1. Evolutionarily stable strategies

The Game. We shall conduct our discussion in the context of 'Game G', shown in Fig. 1. We refer to the participants in Game G as 'player one' and 'player two' or as 'role one' and 'role two'. Game G has three pure-strategy Nash equilibria, given by $(L\ell, L)$, (Lr, L) and $(R\ell, R)$. The first two are contained in a single component in which player one plays any mixture of $L\ell$ and Lr and player two plays L with probability at least 1/2. The third is contained in a component in which two plays R and player one mixes between $R\ell$ and Rr, with probability at least 1/3 on $R\ell$. Game G has a unique subgame perfect equilibrium, given by $(R\ell, R)$.

* Corresponding author. Financial support from National Science Foundation grant SES-9122176 and the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at the University of Bonn, is gratefully acknowledged. We thank Georg Nöldeke for helpful discussions, and thank Avner Shaked and the University of Bonn for their hospitality.

0014-2921/94/$07.00 © 1994 Elsevier Science B.V. All rights reserved
SSDI 0014-2921(93)E0111-W

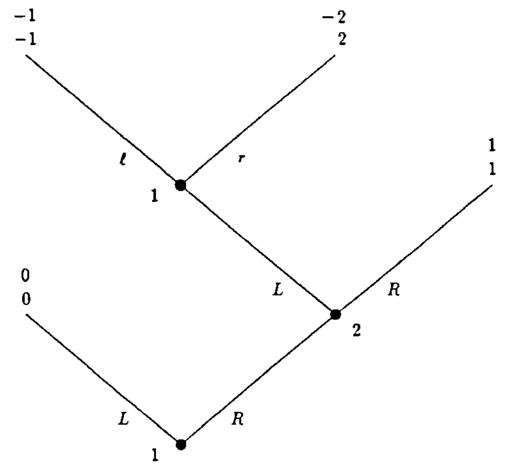

Fig. 1. Game G

Evolutionarily stable strategies. Maynard Smith (1982) and Maynard Smith and Price (1973) introduced the idea of an evolutionarily stable strategy for two-player, symmetric normal form games. Let S be the strategy set of both players one and two, and let $\pi(s,t)$: $\Delta S \times \Delta S \rightarrow \Re$ be the expected payoff to strategy $s \in \Delta S$ (the set of mixed strategies) when played against strategy $t \in \Delta S$. Then strategy $s \in S$ is an evolutionarily stable strategy if, for all $t \in \Delta S$,

$$\pi(s,s) \geqslant \pi(t,s), \tag{1}$$

$$\pi(s,s) = \pi(s,t) \Rightarrow \pi(s,t) > \pi(t,t). \tag{2}$$

These conditions ensure that if a population in which all agents play s is invaded by a small collection of mutants playing t, then the s agents receive a higher average payoff. This in turn presumably causes an evolutionary process to eliminate t.

To apply the ESS concept to the asymmetric Game G, we imagine a larger, symmetric game in which each agent in a single population is characterized by a strategy for each role in the game. The players are randomly matched and assigned to roles. The characterization of evolutionary stability given by (1)–(2) then still holds, where the strategies are interpreted as pairs of strategies from the original game, one for each role, and payoffs are calculated in the larger, 'symmetrized' game.

An evolutionarily stable strategy does not exist for Game G. For example, the subgame perfect equilibrium $(R\ell,R)$ is not an ESS because the strategy (Rr,R) can invade. More precisely, taking s to be $(R\ell,R)$ and t to be (Rr,R) yields equality throughout (1)–(2). In general, Selten (1980) has shown that only strict Nash equilibria in an underlying asymmetric game can be evolutionarily stable strategies.

Accommodating drift. There are two responses to the nonexistence of an evolutionarily stable strategy in Game G: to assimilate the problematic 'drift' into the equilibrium concept or to eradicate the drift. The result of a successful mutant invasion of a pure strategy Nash equilibrium in Game G is the peaceful coexistence of the original and invading strategies, with identical observed play and payoffs. Why worry about such invasions, or, equivalently, why worry about 'drift' in behavior off the equilibrium path? Why not follow set-valued equilibrium refinements, such as strategic stability (Kohlberg and Mertens, 1986), and decline to make sharp predictions about behavior off the equilibrium path?

Toward this end, Maynard Smith (1982) suggests the concept of a *neutrally stable strategy*, which is defined by (1)–(2) with the strict inequality in (2) replaced by a weak inequality. The *weak evolutionarily stable strategy* of Fudenberg and Maskin (1990) and the *modified evolutionarily stable strategy* of Binmore and Samuelson (1992) are similar concepts. In Game G, every Nash equilibrium is a neutrally stable strategy except that in which player two plays R and player one plays $R\ell$ with probability 1/3 and Rr with probability 2/3; and that in which player one mixes between $L\ell$ and Lr while player two plays L with probability 1/2.

Eliminating drift. Maynard Smith (1982) suggests that small perturbations in either the game or the evolutionary process will eliminate drift off the equilibrium path. Hofbauer and Sigmund (1988, p. 288) capture this by offering the concept of a weak evolutionarily stable strategy for asymmetric normal form games. The *limit evolutionarily stable strategy* of Selten (1983, 1988) is similarly motivated. If S_1 and S_2 are player one and two's strategy sets, then $(s_1^*,s_2^*) \in S_2 \times S_2$ is a *weak evolutionarily stable strategy* if it is a Nash equilibrium and for $i \in \{1,2\}$ and $j \neq i$,

$$\pi_i(s_i^*,s_j^*) = \pi_i(s_i,s_j^*) \Rightarrow \pi_i(s_i^*,s_j) > \pi_i(s_i,s_j) \quad \forall s_j \in S_j \setminus \{s_j^*\}. \tag{3}$$

Hence, alternative best replies against the equilibrium strategy must fare strictly worse against all other strategies. Game G has a unique weak evolutionarily stable strategy, given by the subgame perfect equilibrium of $(R\ell,R)$.

These two responses to drift yield quite different results; giving in one case all Nash equilibria (except certain boundary equilibria) and in the other case giving only the subgame perfect equilibrium. In order to evaluate the concepts, two issues must be addressed: set-valued solution concepts and the dynamic processes underlying these solution concepts.

Set valued-solution concepts. Neutrally stable strategies generally occur in the form of connected components of strategies, corresponding to a variety of behavior off the equilibrium path. The inability to repel mutants raises the possibility that the system can drift between the neutrally stable strategies in such a component. We are thus forced to consider set-valued solution concepts.

One possibility is to simply take the set of all neutrally stable strategies. However, consider the component of neutrally stable strategies in Game G that includes the equilibrium $(L\ell,L)$. Occasional mutations may cause the underlying dynamic system to drift between the elements of this component, until R is played with a probability very close to 1/2. At this point, a mutant that plays $(R\ell,R)$ can successfully invade, leading the system away from the component. The neutral stability of strategies in the interior of a component is of little comfort if the system can drift to states from which mutants can conquer. Instead, the stability of the entire component must be assessed, and hence the set-valued nature of the equilibrium concept must be built into the definition of the concept.

We will consider two set-valued equilibrium concepts, due to Swinkels (1992) and Sobel (1993). The equilibrium concepts of Thomas (1985a, b), Gilboa and Matsui (1991), and Ritzberger and Weibull (1993) are similarly motivated. Let ΔS be the joint (mixed) strategy set of a normal form game, let $C(s)$ be the carrier of s, and let $B(s)$ be the set of all strategies that are best replies to s. Then a set of strategies $X \subset \Delta S$ is *equilibrium evolutionarily stable* (Swinkels, 1993) if X is closed, each $s \in X$ is a Nash equilibrium, and $\exists \epsilon^* > 0$ such that

$$\forall \epsilon \in (0,\epsilon^*), \forall s \in X, \forall s' \in \Delta S, \quad C(s') \subset B((1-\epsilon)s + \epsilon s') \Rightarrow (1-\epsilon)s + \epsilon s' \in X.$$

Hence, any strategy profile that can be created from a member of an equilibrium evolutionarily stable set by adding mutants who play best responses to their post-entry environment must also be contained in the set.

Game G has a unique equilibrium evolutionarily stable set in which player two always plays R, yielding the subgame perfect equilibrium outcome. In particular, the strategy profile $(L\ell,L)$ cannot be contained in an equilibrium evolutionarily stable set because the set would also have to contain the strategy profile in which player one chooses $L\ell$ and player two mixes equally between L and R, at which point there are mutant strategies (placing higher probabilities on R for player 2) that pass the test of being best replies to their post-entry environment, but take the system outside the set of Nash equilibria.

We can contrast this with the nonequilibrium evolutionarily stable set of Sobel (1993). Consider two finite populations. Say that strategy profile s, identifying a pure strategy for each agent of each population, can 'replace' s' if s and s' differ only in the strategy of a single agent, with that agent's expected payoff against a randomly selected agent of the opposing population being at least as high under s as it is under s'. Then a set X is a *nonequilibrium evolutionarily stable* set if it is nonempty and minimal with respect to the property that $\{s' \in X$ and s able to replace $s'\}$ implies that $s \in X$. Hence, the set must contain any population profile to which it can move via a process of having one player at a time shift in a payoff increasing direction.

In Game G, there is a single nonequilibrium evolutionarily stable set, containing (among other strategy profiles) every strategy profile in which player one chooses any mixture between $L\ell$ and Lr as well as every strategy profile in which player one mixes between $R\ell$ and Rr while player two chooses R. Hence, both Nash equilibrium outcomes are included. In addition, as the name suggests, nonequilibrium evolutionarily stable sets can contain nonequilibrium strategy profiles.

To see the difference between equilibrium and nonequilibrium evolutionarily stable sets, consider the subgame perfect equilibrium $(R\ell,R)$ in Game G, and consider the dynamic stories that implicitly lie behind the two concepts. Both concepts allow player one to drift toward Rr, to the point that L can become a weak best reply for player two. The implicit assumption behind an equilibrium evolutionarily stable set is that mutants adjust faster than incumbents. Further drift toward Rr then prompts player two to adjust toward L, which *immediately* pushes the mutant agent ones back to $R\ell$ and restores the Nash equilibrium. A nonequilibrium evolutionarily stable set, in contrast, admits the possibility that once sufficient drift toward Rr has occurred to make L a best reply for player two, the primary adjustments may be for incumbent player-two agents to switch to L and player-one agents to switch to either $L\ell$ *or* Lr, leading away from the subgame perfect equilibrium.

Dynamic models. We thus have contending set-valued equilibrium concepts that differ in their assumptions about the relative rates at which various agents adjust strategies. We can see no way to gain insight into such matters other than to model the dynamic processes by which agents choose strategies. The intuition that perturbations will eliminate drift off the equilibrium path can also be asssessed only by constructing a dynamic model incorporating the perturbations.

We will consider a model offered by Nöldeke and Samuelson (1993). (Other dynamic analyses include Binmore and Samuelson, 1993b; Binmore, Samuelson and Vaughn, 1993; Swinkels, 1993; and Ritzberger and Weibull, 1993). The game is played by two finite populations of agents. At the beginning of each discrete time period, each agent is characterized by a pure strategy and a conjecture over the pure strategies of agents in the opposing population. A round-robin tournament is played. Agents are then independently and randomly selected to 'learn', in which case they adjust their conjectures to match their most recent experience and choose a best reply to these new conjectures. After this learning process, agents are again independently and randomly drawn, this time to be mutants, in which case they are assigned a new conjecture and strategy from an exogenously specified, completely mixed distribution. Attention is then directed to the unique stationary distribution of this Markov process, in the limit as the probability of a mutation gets arbitrarily small, which we refer to as the 'limiting' distribution.

The techniques for examining this limiting distribution have been developed by Kandori et al. (1993), Young (1993), Samuelson (1992), and Nöldeke and Samuelson (1993). The limiting distribution will attach support only to absorbing sets under the learning process, i.e., sets that the learning process can cause the system to enter but not exit. In Game G, all absorbing sets are singletons, and correspond to *self-confirming* equilibria, where we define a self-confirming equilibrium to be a collection of conjectures and strategies such that the strategies are best replies to the conjectures and the conjectures are confirmed along the path of play (cf. Fudenberg and Levine, 1993). A self-confirming equilibrium need not be a Nash equilibrium because conjectures need not be correct off the equilibrium path. For example, the strategies (Lr,R) can be part of a self-confirming equilibrium in Game G supported in part by conjectures on the part of player-one agents that player-two agents are playing L.

The self-confirming equilibria that appear in the limiting distribution are those that 'require the fewest mutations to reach'. Nöldeke and Samuelson (1993) show that the limiting distribution always includes the subgame perfect equilibrium. They also show that if a single mutation can move the system from one self-confirming equilibrium to the basin of attraction (under the learning dynamics) of another, then the first appears in the limiting distribution only if the second also appears. Since mutations can cause the system to drift between self-confirming equilibria that differ only in behavior off the equilibria path and since the subgame perfect equilibrium appears in the limiting distribution, then every other self-confirming equilibrium supporting the subgame perfect outcome, including those in which all player-one agents play Rr, must also appear in the limiting distribution. From here, however, a single mutation causing a player-two agent to play L can lead to a learning sequence in which first player-two agents happen to get the learn draw and switch to L and then player-one agents happen to get the learn draw and switch to $L\ell$ or Lr, giving a Nash equilibrium that is not subgame perfect. Since only one mutation was required before the learning dynamics took over, the latter equilibrium also appears in the limiting distribution. The limiting outcome thus includes all of the Nash equilibrium outcomes, including those that are not subgame perfect. Notice that the outcome and the process by which it is supported resemble Sobel's nonequilibrium

evolutionarily stable sets. Notice also that in this model, perturbations do not suffice to eliminate drift.

Learning 'symmetrically'. The learning and mutation dynamics in Nöldeke and Samuelson's model allow quite asymmetric realizations, such as cases where one population learns or mutates much faster than another. For example, the path we described, leading away from the subgame perfect equilibrium, begins with a sequence of mutations that hits population one much more often than population two. What if the two populations are restricted to adjust at similar speeds? For example, we might work with infinite populations and some variant of the replicator dynamics. Relative speeds of learning would then be fixed by the payoffs of the game. Might this suffice to restrict the limiting distribution to the subgame perfect equilibrium?

Samuelson and Zhang (1992), Binmore and Samuelson (1993a) and Binmore et al. (1993) investigate models based on noisy replicator dynamics. In each case, the finding is that perturbations do not suffice to eliminate drift and the models do not restrict attention to subgame perfect equilibrium. A similar result emerges from Binmore and Samuelson (1993b), Ritzberger and Weibull (1993) and some variants of the dynamic process in Swinkels (1993).

Binmore and Samuelson (1993a) and Binmore et al. (1993) examine the Ultimatum Game, which is even simpler than Game G because it has only two moves: player one makes a proposal of how the two players are to split a dollar, which player two can either accept, implementing the proposal, or refuse, yielding zero payoffs for both. The subgame perfect equilibrium gives the entire dollar (or all but one penny, if offers must be made in pennies) to player one. The replicator dynamics have a sink at the subgame perfect equilibrium, but have a host of other rest points at Nash equilibria that are not subgame perfect. The replicator dynamics can thus cause the system to come to rest on equilibria that are not subgame perfect.

Equilibria that are not subgame perfect in the Ultimatum Game can be rest points under the replicator dynamics only because some strategies are extinct (i.e., played by no agents), and the replicator dynamics cannot increase the proportion of agents playing such a strategy. In particular, outcomes can survive in which player two receives part of the dollar only because all agents in the player one population who offer a smaller share to player two have disappeared. One's intuition is again that if perturbations in the system allowed extinct strategies to be continually revived, then there would be constant pressure pushing the system toward the subgame perfect equilibrium. In order to investigate this conjecture, Binmore et al. examine 'noisy' versions of the replicator dynamics that continually inject all strategies into the system. Not only do rest points that are not the subgame perfect equilibrium persist, but they can now be asymptotically stable. Simulations show that from a wide variety of initial conditions, the system converges to an equilibrium that gives player two slightly more than twenty percent of the dollar, rather than the subgame perfect equilibrium payoff of zero. 'Symmetric' learning models then do not allow us to restrict attention to subgame perfect equilibria.

Swinkels (1993) shows that there is some hope for obtaining backward induction results from the limiting outcome of a dynamic process if the dynamic process has rest points at Nash equilibria and *only* at Nash equilibria. In particular, the self-confirming equilibrium of Game G with behavior (Lr,R) (supported by player-two conjectures of $L\ell$) must not be a rest point, instead creating pressure (without mutations) for player-two agents to switch to L. Such pressure may arise if players have access to information concerning not only the outcome of play but also the strategies played by their opponents. However, if players can observe only outcomes and not strategies, as will often be the case in extensive form games, then it is problematic to assume that self-confirming equilibria are not rest points and hence less likely that backward induction can be rescued.

2. Conclusion

The void left by the nonexistence of an ESS in most asymmetric games has been filled by a host of alternative concepts. The most important message of our discussion is that these concepts can be evaluated only by examining the dynamic processes that produce them. On the one hand, this is a formidable task, as current models of how agents actually learn to play games are both preliminary and primitive. On the other hand, by transferring the debate into one concerning individual learning behavior, a topic amenable to both theoretical and empirical (or experimental) analysis, this investigation holds out the promise of making important new progress on the equilibrium selection problem. This progress may well require fundamental changes in our views of game theory. In particular, it appears as if backward induction will lose its currently exalted position.

References

Binmore, Kenneth G. and Larry Samuelson, 1992, Evolutionary stability in repeated games played by finite automata, Journal of Economic Theory 57, 278–305.

Binmore, Kenneth G. and Larry Samuelson, 1993a, An economist's perspective on the evolution of norms, Mimeo. (University College London, London and University of Wisconsin, Madison, WI).

Binmore, Kenneth G. and Larry Samuelson, 1993b, Muddling through: Noisy equilibrium selection, Mimeo. (University College London, London and University of Wisconsin, Madison, WI).

Binmore, Kenneth G. and Larry Samuelson, 1993c, Musical chairs: The evolutionary mechanics of equilibrium selection, Mimeo. (University College London, London and University of Wisconsin, Madison, WI).

Binmore, Kenneth G., John Gale and Larry Samuelson, 1993, Learning to be imperfect: The ultimatum game, Mimeo. (University College London, London and University of Wisconsin, Madison, WI).

Fudenberg, Drew and David K. Levine, 1993, Self-confirming equilibrium, Econometrica 61, 523–446.

Fudenberg, Drew and Eric Maskin, 1990, Evolution and cooperation in noisy repeated games, American Economic Review 80, 274–279.

Gilboa, Itzhak and Akihiko Matsui, 1991, Social stability and equilibrium, Econometrica 59, 859–868.

Hofbauer, J. and K. Sigmund, 1988, The theory of evolution and dynamical systems (Cambridge University Press, Cambridge).

Kandori, Michihiro, George J. Mailath and Rafael Rob, 1993, Learning, mutation, and long run equilibria in games, Econometrica 61, 29–56.

Kohlberg, Elon and Jean-François Mertens, 1986, On the strategic stability of equilibria, Econometrica 54, 1003–1038.

Maynard Smith, John, 1982, Evolution and the theory of games (Cambridge University Press, Cambridge).

Maynard Smith, John and G.R. Price, 1973, The logic of animal conflict, Nature 273, 15–18.

Nöldeke, Georg and Larry Samuelson,,1993, An evolutionary analysis of backward and forward induction. Games and Economic Behavior 5, 425–454.

Ritzberger, Klaus and Jörgen Weibull, 1993, Evolutionary selection in normal form games, Mimeo. (Institute for Advanced Studies, Vienna and Stockholm University, Stockholm).

Samuelson, Larry, 1992, Dominated strategies and common knowledge, Games and Economic Behavior 4, 284–313.

Samuelson, Larry and Jianbo Zhang, 1992, Evolutionary stability in asymmetric games, Journal of Economic Theory 57, 363–391.

Selten, Reinhard, 1980, A note on evolutionarily stable strategies in asymmetric animal contests, Journal of Theoretical Biology 84, 93–101.

Selten, Reinhard, 1983, Evolutionary stability in extensive two-person games, Mathematical Social Sciences 5, 269–363.

Selten, Reinhard, 1988, Evolutionary stability in extensive two-person games-correction and further development, Mathematical Social Sciences 16, 223–266.

Sobel, Joel, 1993, Evolutionary stability in communication games, Economics Letters, Forthcoming.

Swinkels, Jeroen, 1992, Evolutionary stability with equilibrium entrants, Journal of Economic Theory 57, 306–332.

Swinkels, Jeroen, 1993, Adjustment dynamics and rational play in games, Games and Economic Behavior 5, 455–484.

Thomas, B., 1985a, Evolutionarily stable sets in mixed-strategist models, Theoretical Population Biology 28, 332–341.

Thomas, B., 1985b, On evolutionarily stable sets, Journal of Mathematical Biology 22, 105–115.

Young, Peyton, 1993, The evolution of conventions, Econometrica 61, 57–84.