Evolutionary Dynamics and Extensive Form Games[†] by Ross Cressman

Reviewed by William H. Sandholm[*]

Noncooperative game theory is one of a handful of fundamental frameworks used for economic modeling. It is therefore troubling that the solution concepts on which the theory's predictions are based are not as firmly grounded as one might desire. For example, while Nash equilibrium is the starting point for most game theoretic analyses, the conditions on players' rationality and knowledge needed to ensure that a group of players will play a Nash equilibrium are disconcertingly strong. If game theoretic solution concepts can only be derived from heroic assumptions about individuals' abilities, predictions based on these concepts are cast into doubt.

Evolutionary game theory offers a novel approach for generating and justifying predictions of behavior in strategic situations. Under the evolutionary paradigm, one associates a large population of agents with each player role in the underlying game. Rather than impose strong requirements on individual agents' abilities, one assumes instead that agents dynamically adjust their choices in response to the choices made by others. Among other things, this framework allows one to ask whether and when the empirical distributions of agents' choices will come to resemble a Nash equilibrium of the underlying game.

Evolutionary game theorists have succeeded in characterizing a range of contexts in which populations can learn to behave as the standard theory predicts. Still, from the point of view of applied research these developments leave a significant gap. The vast majority of work on evolution in games focuses on normal form games. In contrast, most economic applications of game theory are based on extensive form games. Of course, one can analyze the latter games in terms of their reduced normal forms, but evolutionary analyses based on the reduced normal form ignore the original game's extensive form structure. This structure underlies the basic refinements of Nash equilibrium, and can also considerably simplify equilibrium computations—compare the fixed point methods required to find Nash equilibria with the backward induction techniques used to compute subgame perfect equilibria. For these reasons, the development of new techniques for studying extensive form evolution is a topic of obvious importance.

Ross Cressman's *Evolutionary Dynamics for Extensive Form Games* is the first book length study of evolution in games with sequential moves. Cressman shows how the

---

[†] MIT Press, 2003.

[*] Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706. e-mail: whs@ssc.wisc.edu. website: http://www.ssc.wisc.edu/~whs. July 20, 2004.

notion of backward induction can be incorporated into the study of evolutionary game dynamics, demonstrating that backward induction is useful not only for finding refined equilibria, but also for understanding out-of-equilibrium play. Cressman provides complete evolutionary analyses of many canonical extensive form examples. Perhaps most importantly, he takes the first steps toward a general theory of evolution based on backward induction methods. I expect that Cressman's book will serve as the principal starting point for future research on extensive form evolution.

The heart of Cressman's book uses techniques based on subgame decompositions to prove convergence results for the replicator dynamic.[1] The main results begin in Chapter 7, which focuses on two player "simultaneity games". This class of games includes finitely repeated games, two stage signaling games, and the War of Attrition as special cases. The strongest results concern subgame perfect equilibria (SPE) that are "pervasive", in the sense of reaching each information set with positive probability. Cressman proves that pervasive SPE are asymptotically stable if and only if they can derived by applying asymptotic stability inductively, starting at the terminal subgames and working backward through the game tree. The chapter concludes with an evolutionary analysis of the finitely repeated Prisoner's Dilemma: introducing new analytical methods, Cressman establishes convergence of the replicator dynamic to Nash equilibrium, thereby showing that cooperative behavior cannot persist.

Chapter 8 analyzes two player games of perfect information. Despite the fact that classical analysis of these games via backward induction is very simple, evolutionary analysis can be rather difficult, and can lead to unexpected results. For example, suppose we apply the replicator dynamic to a two node game of Entry Deterrence. In the unique SPE of this game, entry occurs. But while some solutions of the replicator dynamic converge to the SPE, others converge to the other component of Nash equilibria in which entry is deterred by a noncredible threat. However, only the equilibrium component containing the SPE is asymptotically stable.

Cressman offers general conditions on perfect information games that imply the final statement above: namely, that the unique asymptotically stable set is the component of Nash equilibria containing the SPE. But he also shows that these sufficient conditions can be quite demanding. For instance, in the Centipede game the unique component of equilibria is asymptotically stable when the number of decision nodes is 2 or 3, but fails to be even Lyapunov stable otherwise.

Weaved throughout the book and culminating in Chapter 9 are the beginnings of a general theory of evolution in extensive form games. This theory is built on two basic concepts: the Wright manifold and subgame monotonicity.

---

[1] The *replicator dynamic* is one of the fundamental dynamics of evolutionary game theory. It is defined by the property that the percentage growth rate of each strategy is equal to that strategy's *excess payoff*: that is, the difference between the strategy's actual payoff and the population's average payoff.

Let me introduce these concepts in the context of a symmetric extensive form game,[2] Γ. A pure strategy in the game Γ specifies a player's behavior at each of information sets. Thus, if S_h is the set of choices available to the player at information set h, his pure strategies are elements of the product set $\prod_h S_h$, where the product is taken over all of his information sets. By definition, a mixed strategy is an element of $X = \Delta(\prod_h S_h)$, the set of probability distributions over pure strategies.

In evolutionary contexts, the set X plays a different role. If a population of agents are randomly matched to play the symmetric game Γ, then a population state describes what proportion of agents are choosing each pure strategy. Thus, population states are elements of $X = \Delta(\prod_h S_h)$ as well.

If we view X as the set of mixed strategies, then *Wright manifold* $W \subseteq X$ consists of those mixed strategies with the property that *choices made at temporally unordered pairs of information sets are stochastically independent*. In evolutionary contexts, the population state $x \in X$ is in the Wright manifold if individuals' choices at unordered pairs of information sets are *statistically independent*. In other words, for x to be in W, learning that a randomly chosen member of the population would chose strategy s at information set h must provide no information about whether this agent would have chosen strategy s' at information set h', at least when h' neither precedes nor follows h in the game tree.

In games like Centipede in which all information sets are ordered, the Wright manifold is simply all of X (i.e., the set of all probability measures on the product set $\prod_h S_h$). More generally, the Wright manifold can be much smaller than X: when all of a player's pairs of information sets in Γ are unordered, W need only contain the *product distributions* on $\prod_h S_h$. The restriction to product distributions reflects the fact that choices made at one information set provide no information about the choices made at the others.

Cressman uses a simple example to illustrate a fundamental fact about evolution in extensive form games: away from the Wright manifold, the replicator dynamic can behave in very peculiar ways. The example, presented in Section 4.6.1, is a Rock-Paper-Scissors (RPS) game that is augmented by a sunspot: before play of RPS begins, the two players commonly observe a toss of a fair coin. While the payoffs of the RPS game are specified so that its unique equilibrium is globally stable, Cressman shows that in the sunspot game the unique Nash equilibrium component is *unstable*.

The importance of the Wright manifold W for understanding evolutionary dynamics now becomes evident. In the sunspot game, each player has two information sets, one in each subgame. Since the subgames are unordered, the Wright manifold consists of the product distributions on $X = \prod_{h \in \{h_1, h_2\}} S_h$. Cressman shows that the replicator

[2] For a formal definition, see Cressman's Section 6.4.

dynamic is invariant on W: if the individuals' behaviors in the two subgames are initially statistically independent, so that the initial state is a product distribution on X, the replicator dynamic preserves this independence over time. Moreover, if we restrict attention to the action of the replicator dynamic on W, we find that the unique Nash equilibrium that lies in W is asymptotically stable.[3] The intuition for this result is straightforward: since the agents' behaviors in the two subgames always remain independent, evolution is able to act separately in each. On the other hand, if the agents' behaviors in the two subgames are correlated, then evolutionary forces in one subgame can interfere with the selection process in the other, preventing convergence to Nash equilibrium.

This example is a primary motivation for the book's final chapter, which introduces the concept of subgame monotonicity. A trajectory through the state space X is called *subgame monotone* if strategy's growth rates reflect relative payoffs not only in the game as a whole, but also in the subgames of Γ and in the truncations of Γ obtained by collapsing subgames. Cressman shows that subgame monotonicity has a number of important implications: for instance, in generic perfect information games, all interior subgame monotone trajectories converge to components of Nash equilibria.

Are there dynamics whose solution trajectories are subgame monotone? Cressman shows that under the replicator dynamic, solution trajectories that lie *on the Wright manifold* are subgame monotone. This fact can be used to prove convergence results for the Centipede game: since in this game $W = X$, all solutions to the replicator dynamic are subgame monotone, and hence all interior solutions converge to Nash equilibrium. But when W and X differ, the conclusions we can draw about the replicator dynamic are more limited: either we must arbitrarily impose independence assumptions to ensure that the initial state lies on the Wright manifold, or we must accept the fact that the replicator dynamic will sometimes exhibit seemingly strange behavior, as in the sunspot example above.

Still, there is a third alternative: we can abandon the replicator dynamic and instead seek dynamics whose solution trajectories are *all* subgame monotone. While the existence of such dynamics is an open question, Cressman suggests how they might be constructed. In Sections 2.10, 4.6, and 8.4, he reports on joint work with Karl Schlag on evolution in one-player games in which payoffs are determined by multiarmed bandits—that is, by draws from unknown but fixed probability distributions. One can construct rules for playing these bandit games that perform well for large classes of payoff distributions and game trees. When these rules are used by many agents, their behavior within each subgame can resemble the replicator dynamic defined directly for

[3] The Nash equilibrium in W is the product distribution whose marginal distributions both equal the Nash equilibrium of the original RPS game.

that subgame. By employing similar rules in multiplayer contexts, it may be possible to construct dynamics for general extensive form games whose solutions are always subgame monotone.

While I have attempted to summarize Cressman's main line of argument, his book covers an array of other distinct but related topics. Chapters 2 and 3 review evolution in matrix and bimatrix games, covering not only the replicator dynamic, but also more general classes of imitative dynamics as well as the best response dynamic and the fictitious play process. These chapters cover topics not treated elsewhere, including the use of center manifold theory for evaluating the local stability of Nash equilibria, and the relationship between solutions of bimatrix games and solutions of their symmetrized counterparts. Chapter 3 also presents the phase diagrams for the replicator and best response dynamics in all nine distinct classes of 2 x 2 bimatrix games, including six degenerate classes. Since nondegenerate extensive form games often induce degenerate normal form games, the phase diagrams for the degenerate cases are crucial building blocks for the study of extensive form evolution. Cressman's book may be the first to collect all of these diagrams in one place.

While many of the motivations for studying evolution in extensive form games are drawn from economics, the origins of evolutionary game theory actually lie in biological models of intraspecies competition. Indeed, the replicator dynamic first appeared in precisely this context.[4] Researchers later realized that the replicator dynamic is equivalent to the fundamental selection equation of theoretical population genetics, an equation with a much longer history in the mathematical biology literature.

In Chapter 5, Cressman develops this connection by presenting results from theoretical population genetics using his extensive form framework. He shows that a basic model of natural selection of trait profiles (i.e., of gametes containing genes at multiple loci) can be described by applying the replicator dynamic to a certain class of symmetric extensive form games; these games consist of an initial move by Nature followed by the play of a matrix game with symmetric payoffs. By applying his results on extensive form evolution, Cressman shows why natural selection increases evolutionary fitness and ultimately leads to locally efficient gamete distributions.

One can obtain more realistic models of genetic evolution by introducing *recombination*: the random reshuffling of genes situated at different loci during cell division (meiosis). If selection pressures are absent, recombination causes individuals' genes at different loci to become statistically independent: in our earlier terminology, the Wright manifold is a global attractor of the recombination process. If we allow selection and recombination to act at once, we obtain a powerful integrated

[4] Peter D. Taylor and Leo B. Jonker, "Evolutionarily Stable Strategies and Game Dynamics," *Mathematical Biosciences* 40 (1978), 145-156.

convergence result: regardless of the initial specification of the gamete distribution, natural selection leads this distribution to a Nash equilibrium that lies on the Wright manifold.

Thus, the Wright manifold, which proves so important in studying extensive form evolution, has its origins elsewhere; it is actually the namesake of Sewall Wright, a founding father of population genetics. This connection illustrates how applications of game theory in economics and biology can build on one another in unexpected ways. By addressing his book to both economists and biologists, and in exploring the links between the contributions of both groups, Cressman ably promotes future investment in cross-fertilization.