Deterministic Evolutionary Dynamics

Prepared for the *New Palgrave Dictionary of Economics*, 2nd edition
William H. Sandholm[1]
University of Wisconsin
1 November 2005

1. Introduction

Deterministic evolutionary dynamics for games first appeared in the mathematical biology literature, where Taylor and Jonker (1978) introduced the *replicator dynamic* to provide an explicitly dynamic foundation for the static evolutionary stability concept of Maynard Smith and Price (1973) (see the entry on "Evolutionarily Stable Strategies" in this dictionary). But one can find precursors to this approach in the beginnings of game theory: Brown and von Neumann (1950) introduced differential equations as a tool for computing equilibria of zero-sum games. In fact, the replicator dynamic appeared in the mathematical biology literature long before game theory itself: while Maynard Smith and Price (1973) and Taylor and Jonker (1978) studied game theoretic models of animal conflict, the replicator equation is equivalent to much older models from population ecology and population genetics. These connections are explained by Schuster and Sigmund (1983), who also introduced the name "replicator dynamic", borrowing the term "replicator" from Dawkins (1982).

In economics, the initial phase of research on deterministic evolutionary dynamics in the late 1980s and early 1990s focused on populations of agents who are randomly matched to play normal form games, with evolution described by the replicator dynamic or other closely related dynamics. The motivation behind the dynamics continued to be essentially biological: individual agents are preprogrammed to play specific strategies, and the dynamics themselves are driven by differences in birth and death rates. Since that time, the purview of the literature has broadened considerably, allowing more general sorts of large population interactions, and admitting dynamics derived from explicit models of active myopic decision making.

This article provides a brief overview of deterministic evolutionary dynamics in game theory. More detailed treatments of topics introduced here can be found in the recent survey article by Hofbauer and Sigmund (2003), and in books by Maynard Smith (1982), Hofbauer and Sigmund (1988, 1998), Weibull (1995), Vega-Redondo (1996), Samuelson (1997), Fudenberg and Levine (1998), Cressman (2003), and Sandholm (2006).

2. Population Games

Population games provide a general model of strategic interactions among large numbers of anonymous agents. For simplicity, we focus on games played by a single population, in which agents are not differentiated by roles; allowing for multiple populations is mostly a matter of introducing more elaborate notation.

In a one-population game, each agent from a unit-mass population chooses a strategy from the finite set $S = \{1, \dots, n\}$, with typical elements i and j. The distribution of strategy choices at a given moment in time is described by a *population state* $x \in X = \{x \in \mathbf{R}_+^n: \sum_{i \in S} x_i = 1\}$. The *payoff* to strategy i, denoted $F_i: X \to \mathbf{R}$, is a continuous function of the population state; we use the notation $F: X \to \mathbf{R}^n$ to refer to all strategies' payoffs at once. By taking the set of strategies S as fixed, we can refer to F itself as a *population game*.

The simplest example of a population game is the most commonly studied one: random matching to play a symmetric normal form game $A \in \mathbf{R}^{n \times n}$, where A_{ij} is the payoff obtained by an agent choosing strategy i when his opponent chooses strategy j. When the population state is $x \in X$, the expected payoff to strategy i is simply the weighted average of the elements of the ith row of the payoff matrix: $F_i(x) = \sum_{j \in S} A_{ij} x_j = (Ax)_i$. Thus, the population game generated by random matching in A is the linear population game $F(x) = Ax$.

Many models of strategic interactions in large populations that arise in applications do not take this simple linear form. Consider, for example, in models of highway congestion, payoff functions are convex: increases in traffic when traffic levels are low have virtually no effect on delays, while increases in traffic when traffic levels are high increase delays substantially (see Beckmann, McGuire and Winsten (1956) and Sandholm (2001)). Happily, allowing nonlinear payoffs extends the range of possible applications of population games without making evolutionary dynamics especially more difficult to analyze, since the dynamics themselves are nonlinear even when the underlying payoffs are not.

3. Foundations of Evolutionary Dynamics

Formally, an *evolutionary dynamic* is a map that assigns to each population game F a differential equation $\dot{x} = V^F(x)$ on the state space X. While one can define evolutionary dynamics directly, it is preferable to derive them from explicit models of myopic individual choice.

We can accomplish this by introducing the notion of a *revision protocol* $\rho: \mathbf{R}^n \times X \to \mathbf{R}_+^{n \times n}$. Given a payoff vector $F(x)$ and a population state x, the revision protocol specifies for each pair of strategies i and j a nonnegative number $\rho_{ij}(F(x), x)$, representing the rate

at which agents currently playing strategy i switch to strategy j.

A given revision protocol can admit a variety of interpretations. For one all-purpose interpretation, suppose each agent is equipped with an exponential alarm clock. When the clock belonging to an agent playing strategy i rings, he selects a strategy $j \in S$ at random, and then switches to this strategy with probability proportional to $\rho_{ij}(F(x), x)$. While this interpretation is always available, others may be simpler in certain instances. For example, if the revision protocol is of the imitative form $\rho_{ij} = x_j \times \hat{\rho}_{ij}$, we can incorporate the x_j term into our story by supposing that the revising agent selects his candidate strategy j not by drawing a *strategy* at random, but by drawing an *opponent* at random and observing this opponent's strategy. Revision protocols that are most consistent with the evolutionary paradigm only require agents to possess limited information: for example, a revising agent might only know the current payoffs of his own strategy i and his candidate strategy j.

A population game F and a revision protocol ρ together generate an ordinary differential equation $\dot{x} = V^F(x)$ on the state space X. This equation, which captures the population's *expected* motion under F and ρ, is known as the *mean dynamic* or *mean field* for F and ρ:

(M) $$\dot{x}_i = V_i^F(x) = \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x).$$

The form of the mean dynamic is easy to explain. The first term describes the "inflow" into strategy i from other strategies; it is obtained by multiplying the mass of agents playing each strategy j by the rate at which such agents switch to strategy i, and then summing over j. Similarly, the second term describes the "outflow" from strategy i to other strategies. The difference between these terms is the net rate of change in the use of strategy i.

To obtain a formal link between the mean dynamic (M) and our model of individual choice, imagine that the population game F is not played by a continuous mass of agents, but rather by a large, finite population with N members. Then the model described above defines a Markov process $\{X_t^N\}$ on a fine but discrete grid in the state space X. The foundations for deterministic evolutionary dynamics are provided by the *Finite Horizon Deterministic Approximation Theorem*: Fix a time horizon $T < \infty$. Then the behavior of the stochastic process $\{X_t^N\}$ through time T is approximated by a solution of the mean dynamic (M); the approximation is uniformly good with probability close to one once the population size N is large enough. (For a formal statement of this result, see Benaïm and Weibull (2003).)

In cases where one is interested in phenomena that occur over very long time horizons, it may be more appropriate to consider the infinite horizon behavior of the

stochastic process $\{X_t^N\}$. Over this infinite time horizon, the deterministic approximation fails, as a correct analysis must explicitly account for the stochastic nature of the evolutionary process. For more on the distinction between the two time scales, see Benaïm and Weibull (2003), as well as the entry on "Stochastic Adaptive Dynamics" in this dictionary.

4. Examples and Families of Evolutionary Dynamics

We now describe revision protocols that generate some of the most commonly studied evolutionary dynamics. In the table below, $\bar{F}(x) = \sum_{i \in S} x_i F_i(x)$ represents the population's average payoff at state x, and $B^F(x) = \mathrm{argmax}_{y \in X} \; y'F(x)$ is the best response correspondence for the game F.

Revision protocol	Evolutionary Dynamic	Name	Origin
$\rho_{ij} = x_j\,(K - F_i)$, or $\rho_{ij} = x_j\,(K + F_j)$, or $\rho_{ij} = x_j\,[F_j - F_i]_+$	$\dot{x}_i = x_i(F_i(x) - \bar{F}(x))$	replicator	Taylor and Jonker (1978)
$\rho_{ij} = [F_j - \bar{F}]_+$	$\dot{x}_i = [F_i(x) - \bar{F}(x)]_+$ $-x_i \sum_{j \in S}[F_j(x) - \bar{F}(x)]_+$	Brown-von Neumann-Nash (BNN)	Brown and von Neumann (1950)
$\rho_{ij} = [F_j - F_i]_+$	$\dot{x}_i = \sum_{j \in S} x_j[F_i(x) - F_j(x)]_+$ $-x_i \sum_{j \in S}[F_i(x) - F_j(x)]_+$	pairwise difference (PD)	Smith (1984)
$\rho_{ij} = \dfrac{\exp(\eta^{-1}F_j)}{\sum_{k \in S}\exp(\eta^{-1}F_k)}$	$\dot{x}_i = \dfrac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S}\exp(\eta^{-1}F_k(x))} - x_i$	logit	Fudenberg and Levine (1998)
$\rho_{ij} = B_i^F(x)$	$\dot{x} = B_i^F(x) - x_i$	best response	Gilboa and Matsui (1991)

A common critique of evolutionary analysis of games is that the choice of a specific revision protocol, and hence the evolutionary analysis that follows, is necessarily arbitrary. There is surely some truth to this criticism: to the extent that one's analysis is sensitive to the fine details of the choice of protocol, the conclusions of the analysis are cast into doubt. But much of the force of this critique is dispelled by this important observation: *evolutionary dynamics based on qualitatively similar revision protocols lead to qualitatively similar aggregate behavior*. We call a collection of dynamics generated by similar revision protocols a *family* of evolutionary dynamics.

To take one example, many properties that hold for the replicator dynamic also hold for dynamics based on revision protocols of the form $\rho_{ij} = x_j\,\hat{\rho}_{ij}$ where $\hat{\rho}_{ij}$ satisfies

$$\text{sgn}\left((\hat{\rho}_{ki} - \hat{\rho}_{ik}) - (\hat{\rho}_{kj} - \hat{\rho}_{jk})\right) = \text{sgn}(F_i - F_j) \text{ for all } k \in S.$$

(In words: if i earns a higher payoff than j, then the net conditional switch rate from k to i is higher than that from k to j for all $k \in S$.) For reasons described in the previous section, dynamics generated in this way are called *imitative dynamics*. (See Björnerstedt and Weibull (1996) for a related formulation.) For another example, most properties of the PD dynamic remain true for dynamics are based on protocols of the form $\rho_{ij} = \phi(F_i - F_j)$, where $\phi\colon \mathbf{R} \to \mathbf{R}_+$ satisfies sign-preservation:

$$\text{sgn}(\phi(d)) = \text{sgn}([d]_+).$$

Dynamics in this family are called *pairwise comparison dynamics*. For more on these and other families of dynamics, see Sandholm (2006, Ch. 5).

5. Rest Points and Local Stability

Having introduced families of evolutionary dynamics, we now turn to questions of prediction: if agents playing game F follow the revision protocol ρ (or, more broadly, a revision protocol from a given family), what predictions can we make about how they will play the game? To what extent do these predictions accord with those provided by traditional game theory?

A natural first question to ask concerns the relationship between the rest points of an evolutionary dynamic V^F and the Nash equilibria of the underlying game F. In fact, one can prove for a very wide range of evolutionary dynamics that if a state $x^* \in X$ is a Nash equilibrium (that is, if $x \in B(x)$), then x^* is a rest point as well.

One way to show that $NE(F) \subseteq RP(V^F)$ is to first establish a *monotonicity* property for V^F: that is, a property that relates strategies' growth rates under V^F with their payoffs in the underlying game (see, e.g., Nachbar (1990), Friedman (1991), and Weibull (1995)). The most general such property, first studied by Friedman (1991) and Swinkels (1993), we call *positive correlation*:

(PC)   If $x \notin RP(V^F)$, then $F(x)'V^F(x) > 0$.

Property (PC) is equivalent to requiring a positive correlation between strategies' growth rates $V_i^F(x)$ and payoffs $F_i(x)$ (where the underlying probability measure is the uniform measure on the strategy set S). This property is satisfied by the first three dynamics in the table above, and modifications of it hold for the remaining two as well. Moreover, it is not difficult to show that if V^F satisfies (PC), then all Nash equilibria of F

are rest points of V^F: that is, $NE(F) \subseteq RP(V^F)$, as desired (see Sandholm (2006, Ch. 5)).

In many cases, one can also prove that every rest point of V^F is a Nash equilibrium of F, and hence that $NE(F)$ equals $RP(V^F)$. In fact, versions of this statement are true for all of the dynamics introduced above, with the notable exception of the replicator dynamic and other imitative dynamics. The reason for this failure is easy to see: when revisions are based on imitation, unused strategies, even ones that are optimal, are never chosen. On the other hand, if we introduce a small number of agents playing the unused optimal strategy, then these agents will be imitated. Developing this logic, Bomze (1986) and Nachbar (1990) show that under many imitative dynamics, every Lyapunov stable rest point is a Nash equilibrium.

As we noted at the onset, the original motivation for the replicator dynamic was to provide a foundation for Maynard Smith and Price's (1973) notion of an evolutionarily stable strategy. Hofbauer, Schuster, and Sigmund (1979) and Zeeman (1980) show that an ESS is asymptotically stable under the replicator dynamic, but that an asymptotically state need not be an ESS.

More generally, when is a Nash equilibrium a dynamically stable rest point, and under which dynamics? Under differentiable dynamics, stability of isolated equilibria can often be determined by linearizing the dynamic around the equilibrium. In many cases, the question of the stability of the rest point x^* reduces to a question of the negativity of certain eigenvalues of the Jacobian matrix $DF(x^*)$ of the payoff vector field. In nondifferentiable cases, and in cases where the equilibria in question form a connected component, stability can sometimes be established using another standard approach: the construction of suitable Lyapunov functions. For an overview of work in these directions, see Sandholm (2006, Ch. 6).

In the context of random matching in normal form games, it is natural to ask whether an equilibrium that is stable under an evolutionary dynamic also satisfies the restrictions proposed in the equilibrium refinements literature (see the entry on "Refinements of Nash Equilibrium" in this dictionary). Swinkels (1993) and Demichelis and Ritzberger (2003) show that this is true in great generality under even the most demanding refinements: in particular, any component of rest points that is asymptotically stable under a dynamic that respects condition (PC) contains a strategically stable set in the sense of Kohlberg and Mertens (1986). While proving this result is difficult, the idea behind the result is simple. If a component is asymptotically stable under an evolutionary dynamic, then this dynamic stability ought not to be affected by slight perturbations of the payoffs of the game. *A fortiori*, the *existence* of the component ought not to be affected by the payoff perturbations either. But this preservation of existence is precisely what strategic stability demands.

This argument also shows that asymptotic stability under evolutionary dynamics is

a qualitatively stronger requirement than strategic stability: while strategic stability requires equilibria not to vanish after payoff perturbations, it does not demand that they be attracting under a disequilibrium adjustment process. For example, while all Nash equilibria of simple coordination games are strategically stable, only the pure Nash equilibria are stable under evolutionary dynamics.

Demichelis and Ritzberger (2003) establish their results using tools from index theory. Given an evolutionary dynamic V^F for a game F, one can assign each component of rest points an integer, called the *index*, that is determined by the behavior of the dynamic in a neighborhood of the rest point; for instance, regular, stable rest points are assigned an index of 1. The set of all indices for the dynamic V^F is constrained by the *Poincaré-Hopf Theorem*, which tells us that the sum of the indices of the equilibrium components of V^F must equal one. As a consequence of this deep topological result, one can sometimes determine the local stability of one component of rest points by evaluating the local stability of the others.

6. Global Convergence: Positive and Negative Results

To provide the most satisfying evolutionary justification for the prediction of Nash equilibrium play, it is not enough to link the rest points of a dynamic and the Nash equilibria of the underlying game, or to prove local stability results. Rather, one must establish convergence to Nash equilibrium from *arbitrary* initial conditions.

One way to proceed is to focus on a class of games defined by some noteworthy payoff structure, and then to ask whether global convergence can be established for games in this class under certain families of evolutionary dynamics. As it turns out, general global convergence results can be proved for a number of classes of games. Among these classes are *potential games*, which include common interest games, congestion games, and games generated by externality pricing schemes; *stable games*, which include zero-sum games, games with an interior ESS, and (perturbed) concave potential games; and *supermodular games*, which include models of Bertrand oligopoly, arms races, and macroeconomic search. A fundamental paper on global convergence of evolutionary dynamics is Hofbauer (2000); for a full treatment of these results, see Sandholm (2006).

Once we move beyond specific classes of games, global convergence to Nash equilibrium cannot be guaranteed; cycling and chaotic behavior become possible. Indeed, Hofbauer and Swinkels (1996) and Hart and Mas-Colell (2003) construct examples of games in which all reasonable deterministic evolutionary dynamics fail to converge to Nash equilibrium from most initial conditions. These results tell us that general guarantees of converge to Nash equilibrium are impossible to obtain.

In light of this fact, we might instead consider the extent to which solution concepts simpler than Nash equilibrium are supported by evolutionary dynamics. Cressman and Schlag (1998) and Cressman (2003) investigate whether imitative dynamics lead to subgame perfect equilibria in reduced normal forms of extensive form games—in particular, generic games of perfect information. In these games, interior solution trajectories do converge to Nash equilibrium components, and only subgame perfect components can be interior asymptotically stable. But even in very simple games, interior asymptotically stable components need not exist, so the dynamic analysis may fail to select subgame perfect equilibria. For a full treatment of these issues, see Cressman (2003).

What about games with strictly dominated strategies? Early results on this question were positive: Akin (1980), Nachbar (1990), Samuelson and Zhang (1992), and Hofbauer and Weibull (1996) prove that dominated strategies are eliminated under certain classes of imitative dynamics. However, Berger and Hofbauer (2005) show that dominated strategies need not be eliminated under the BNN dynamic. Pushing this argument further, Hofbauer and Sandholm (2005) find that dominated strategies can survive under any continuous evolutionary dynamic that satisfies positive correlation and *innovation*. The later condition requires that whenever there is an unused best response, there is an unused best response that is growing. Thus, when revision is based on direct evaluation of payoffs rather than imitation of successful opponents, evolutionary dynamics may violate even the mildest rationality criteria.

7. Conclusion

Because the literature on evolutionary dynamics came to prominence shortly after the literature on equilibrium refinements, it is tempting to view the former literature as a branch of the latter. But while it is certainly true that evolutionary models have something to say about selection among multiple equilibria, viewing them simply as equilibrium selection devices can be misleading. As we have seen, evolutionary dynamics capture the behavior of large numbers of myopic, imperfectly informed decision makers. Using evolutionary models to predict behavior in interactions between, say, two well informed players is daring at best.

The negative results described in the previous section should be understood in this light. If we view evolutionary dynamics as an equilibrium selection device, the fact that they need not eliminate strictly dominated strategies might be viewed with disappointment. But if we take the result at face value, it should not come as a surprise: if agents switch to strategies that perform reasonably well at the moment of choice, the fact that a strategy is never optimal is something they need not even know, and so will

not deter them from playing it.

A related point can be made about failures of convergence to equilibrium. From a traditional point of view, persistence of disequlibrium behavior might seem to undermine the very possibility of a satisfactory economic analysis. But the work described in this entry suggests that in large populations, this possibility is not only real, but is also one that game theorists are well equipped to analyze.

References

Akin, E. (1980). "Domination or Equilibrium," *Mathematical Biosciences* 50, 239-250.

Beckmann, M., C. B. McGuire, and C. B. Winsten (1956). *Studies in the Economics of Transportation*. New Haven: Yale University Press.

Benaïm, M., and J. W. Weibull (2003). "Deterministic Approximation of Stochastic Evolution in Games," *Econometrica* 71, 873-903.

Berger, U., and Hofbauer J. (2005). "Irrational Behavior in the Brown-von Neumann-Nash Dynamics." Forthcoming, *Games Econ. Behav*.

Björnerstedt, J., and J. W. Weibull (1996). "Nash Equilibrium and Evolution by Imitation," in K. J. Arrow *et. al*. (eds.), *The Rational Foundations of Economic Behavior*, 155-181. New York: St. Martin's Press.

Bomze, I. M. (1986). "Non-Cooperative Two-Person Games in Biology: A Classification," *Int. J. Game Theory* 15, 31-57.

Brown, G. W. and J. von Neumann (1950). "Solutions of Games by Differential Equations." In H. W. Kuhn and A. W. Tucker, eds., *Contributions to the Theory of Games I*, 73-79. Annals of Mathematics Studies 24. Princeton: Princeton University Press.

Cressman, R. (2003). *Evolutionary Dynamics and Extensive Form Games*. Cambridge: MIT Press.

Cressman, R. and K. Schlag (1998). "The Dynamic (In)Stability of Backwards Induction," *J. Econ. Theory* 83, 260-285.

Dawkins, R. (1982). *The Extended Phenotype*. San Francisco: Freeman.

Demichelis, S. and K. Ritzberger (2003). "From Evolutionary to Strategic Stability," *J. Econ. Theory* 113, 51-75.

Friedman, D. (1991). "Evolutionary Games in Economics," *Econometrica* 59, 637-666.

Fudenberg, D., and D. K. Levine (1998). *Theory of Learning in Games*. MIT.

Gilboa, I., and A. Matsui (1991). "Social Stability and Equilibrium," *Econometrica* 59, 859-867.

Hart, S., and A. Mas-Colell (2003). "Uncoupled Dynamics Do Not Lead to Nash Equilibrium," *Amer. Econ. Rev*. 93, 1830-1836.

Hofbauer, J. (2000). "From Nash and Brown to Maynard Smith: Equilibria, Dynamics, and ESS," *Selection* 1, 81-88.

Hofbauer, J., and W. H. Sandholm (2005). "Survival of Dominated Strategies under Evolutionary Dynamics," unpublished manuscript, University College London and University of Wisconsin.

Hofbauer, J., P. Schuster, and K. Sigmund (1979). "A Note on Evolutionarily Stable Strategies and Game Dynamics," *J. Theoretical Biology*. 27, 537-548.

Hofbauer, J., and K. Sigmund (1988). *Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge University Press.

Hofbauer, J., and K. Sigmund (1998). *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.

Hofbauer, J., and K. Sigmund (2003). "Evolutionary Game Dynamics," *Bull. Amer. Math. Soc.* (*N.S.*) 40, 479-519.

Hofbauer, J., and J. Swinkels (1996). "A Universal Shapley Example." Unpublished manuscript, University of Vienna and Northwestern University.

Hofbauer, J., and J. Weibull (1996). "Evolutionary Selection against Dominated Strategies," *J. Econ. Theory* 71, 558-573.

Kohlberg, E. and J.-F. Mertens (1986). "On the Strategic Stability of Equilibria," *Econometrica* 54, 1003-1038.

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

Maynard Smith, J., and G. R. Price (1973). "The Logic of Animal Conflict," *Nature* 246, 15-18.

Nachbar, J. H. (1990). "'Evolutionary" Selection Dynamics in Games: Convergence and Limit Properties," *Int. J. Game Theory* 19, 59-89.

Samuelson, L. (1997). "Evolutionary Games and Equilibrium Selection," Cambridge: MIT Press.

Samuelson, L., and J. Zhang (1992). "Evolutionary Stability in Asymmetric Games," *J. Econ. Theory* 57, 363-391.

Sandholm, W. H. (2001). "Potential Games with Continuous Player Sets," *J. Econ. Theory* 97, 81-108.

Sandholm, W. H. (2006). *Population Games and Evolutionary Dynamics*. Unpublished manuscript, University of Wisconsin.

Schuster, P., and K. Sigmund (1983). "Replicator Dynamics," *J. Theoretical Biology* 100, 533-538.

Smith, M. J. (1984). "The Stability of a Dynamic Model of Traffic Assignment—An Application of a Method of Lyapunov," *Trans. Sci.* 18, 245-252.

Swinkels, J. M. (1993). "Adjustment Dynamics and Rational Play in Games," *Games*

Econ. Behav. 5, 455-484.

Taylor, P. D., and L. Jonker (1978). "Evolutionarily Stable Strategies and Game Dynamics," *Mathematical Biosciences* 40, 145-156.

Vega-Redondo, F. (1996). *Evolution, Games, and Economic Behavior*. Oxford: Oxford University Press.

Weibull, J. W. (1995). *Evolutionary Game Theory*. Cambridge: MIT Press.

Zeeman, E. C. (1980). "Population Dynamics from Game Theory." In Z. Nitecki and C. Robinson, eds., *Global Theory of Dynamical Systems*, 471-497. Berlin: Springer.