History Independent Prediction in Evolutionary Game Theory[*]

William H. Sandholm[†]
MEDS – KGSM
Northwestern University
Evanston, IL 60208, U.S.A.
e-mail: whs@nwu.edu

January 28, 1998
Forthcoming in *Rationality and Society*.

[*] I thank Eddie Dekel, Nolan Miller, Jeroen Swinkels, and an anonymous referee for their comments. Financial support from a State Farm Companies Foundation Dissertation Fellowship is gratefully acknowledged.

[†] William H. Sandholm is a Ph.D. candidate in the Department of Managerial Economics and Decision Sciences of the J. L. Kellogg Graduate School of Management, Northwestern University. His research interests include evolution and learning in games and the economics of technological standardization. His most recent publication, entitled "Evolution, Population Growth, and History Dependence", was coauthored by Ady Pauzner and appeared in *Games and Economic Behavior* (1998).

Abstract

We survey three classes of models from evolutionary game theory which make history independent predictions: stochastic stability, stochastic stability with local interaction, and cheap talk. We argue that of the three, only local interaction models yield credible history independent predictions.

Keywords: evolutionary game theory, equilibrium selection, history dependence, conventions.

1. Introduction

Game theory is the mathematical study of interactive decision making. At first, research in game theory focused exclusively on choices made by agents with both knowledge and reasoning ability adequate to completely and unfailingly analyze any strategic encounter. While this eductive approach has proved extremely fruitful, it is nevertheless subject to the criticism that it demands more knowledge and ability of its players than is typically reasonable to assume.

The pioneering work of Maynard Smith and Price (1973) and Maynard Smith (1974) showed that game theory can be applied profitably in biological contexts, in which the agents are animals whose behaviors are determined by their genetic endowments. Rather than conscious decision, natural selection drives the population towards equilibrium behavior.

Evolutionary game theory uses the structures originally developed for biological models to study the evolution of behavior in both the biological and social sciences. In the social sciences, the evolutionary approach studies the behavior of large human populations. Each individual's success depends on the interplay between his behavior and that of his society. In contrast to the eductive approach, in the evolutionary approach players are assumed to change their behavior myopically. Actions which are currently well rewarded are played with greater frequency in the population. Despite the agents' lack of foresight, equilibrium behavior in evolutionary models satisfies many of the strongest rationality conditions defined in the eductive literature.[1] Hence, despite the simple approach to decisions that members of the population follow, their aggregate behavior is as if the agents satisfied stringent rationality postulates.

While early evolutionary game theory models attempted to capture the process of evolution within a static framework, beginning with the work of Taylor and Jonker (1978) attention has turned to explicitly dynamic models of evolution. Most of these models utilize deterministic dynamics: for every possible distribution of behaviors, the results of the myopic adjustment process are uniquely specified. Consequently, knowledge of the initial play of the population is enough to determine the complete path of its behavior. This is particularly useful when more than one stable social outcome exists. While a static analysis often provides no clear

---

1 See Weibull (1995) for a thorough survey of these results.

basis for prediction among multiple stable outcomes, dynamic analysis can provide a unique prediction provided that one knows the relevant historical data.

One might expect that historical influences necessarily play a leading role in determining which conventions emerge, and hence that the prediction of which conventions arise in a particular population must entail a study of the population's past. However, game theorists have recently focused attention on formal frameworks in which only the mathematical structure of the interactions determines which conventions are established. This research program has proposed models of interaction in which the ultimate social outcome depends only on the payoff structure of individual exchanges. The main motivation for this program is to simplify the prediction of social behavior: in settings in which the models apply, rather than depending on history, prediction can be based solely on the payoffs of the underlying game.

We survey recent work on history independent prediction in evolutionary game theory. Underlying each of the models we consider is a simple coordination game. In this game, players choose between two strategies, Up and Down. Taking his opponent's strategy as given, a player prefers mimicking his opponent to choosing a different strategy. For this reason, coordination by both players on the same strategy constitutes equilibrium play: if both players choose Up or both choose Down, neither desires to deviate.

A large population of players is repeatedly randomly matched to play the coordination game. If players myopically choose best responses to the play of their fellows, all players eventually coordinate on the same strategy. Whether coordination occurs on Up or Down depends on the population's initial behavior. Each arrangement is a social convention: players prefer to emulate the action played by a wide majority of their fellows, whatever that action happens to be.

We consider three models which add to the structure described above in order to generate history independent predictions. In the basic *stochastic stability* model, players infrequently experiment with suboptimal actions. The model of *stochastic stability with local interaction* is similar, but incorporates a neighborhood structure: players live in fixed locations, and are only matched with opponents who live close by. Finally, the *communication* or *cheap talk* model assumes that players are able to signal each other before playing the coordination game.

Each of these models incorporates randomness: in the stochastic stability models, it is introduced through occasional experimentation, while in the communication model, it takes the form of drift between combinations of signals and actions yielding the same outcome. The inclusion of a stochastic element in all three models is not accidental. Under most reasonable specifications of deterministic dynamics, either convention is stable. Once a convention is established, no player can improve his payoffs through an individual effort, and hence no strategy changes occur. Adding a stochastic element to the dynamic adjustment process disrupts the stability of the conventions, making unique prediction possible.

All three models guarantee that one of the conventions will eventually predominate regardless of the population's initial behavior. Which convention is selected depends only on the payoff structure of the underlying coordination game. In the communication model, the efficient convention is chosen. In both stochastic stability models, coordination is on the risk dominant convention, which uses the strategy that players prefer when they are completely uncertain about how opponents will act. In all cases, a result describing the limiting behavior of the model provides a formal justification for the history independent prediction.

Each evolutionary model is an attempt to capture the critical details of a class of human interactions in order to draw conclusions about how they will proceed. If one does not view the model as simply a technical device used to generate unique predictions, one must interpret the results of the model in light of the situations from which it abstracts. For a model's history independent predictions to be credible, the amount of time required before the model's results become meaningful must be measurable on a human rather than an astronomical scale. In particular, if play begins at the convention which the model does not predict, the population must switch to the predicted convention within a reasonable time span if the conclusions of the model are to be believed.

In both the basic stochastic stability model and the communication model, the amount of time before such a switch occurs grows exponentially in the size of the population. Consequently, for populations of even moderate size, the relevant prediction in applications is the first convention played. In contrast, in the local interaction model, the delay before a shift to the predicted risk dominant convention is small and essentially independent of the population size. Moreover, once established, coordination on the risk dominant convention is for all practical purposes impossible to break. Hence, when local interaction is an appropriate modeling assumption, unique prediction of outcomes based on the risk structure of payoffs is possible.

The paper proceeds as follows. Sections 2 and 3 introduce coordination games and deterministic evolutionary dynamics. Sections 4, 5, and 6 present examples of the three classes of models described above and provide interpretations of the models' formal results. Section 7 concludes.

2. Coordination Games

Underlying all of the models we consider is a class of two player, one-shot games. In these games each player chooses one of two strategies, Up and Down. The payoffs a player receives depend on both his action and that of his opponent, as specified in Figure 1.

$$
\begin{array}{ccc}
 & & \text{Opponent chooses} \\
 & & \begin{array}{cc} \text{Up} & \text{Down} \end{array} \\
\text{Player chooses} & \begin{array}{c} \text{Up} \\ \text{Down} \end{array} & \begin{bmatrix} a & b \\ c & d \end{bmatrix}
\end{array}
$$

Figure 1: A Game

We call any 2 x 2 matrix of payoffs a *game*.[2] The interpretation is always that given in Figure 1. For concreteness, we give three examples of games which we revisit throughout the paper:

$$A = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}, \qquad B = \begin{bmatrix} 5 & 0 \\ 4 & 2 \end{bmatrix}, \qquad C = \begin{bmatrix} 3 & 0 \\ 4 & 1 \end{bmatrix}.$$

Games A and B, which satisfy the inequalities $a > c$ and $d > b$, are *coordination games*. They are so named because it is in the players' best interests to coordinate their actions. If John knows that Mary intends to play, say, Up, he maximizes his payoffs by also playing Up. The choice of the same strategy by each player constitutes a (*Nash*) *equilibrium* of the game: given the choice of the opponent, no player has an incentive to switch.[3]

---

[2] More precisely, it is a two player symmetric normal form game.

[3] There is a third Nash equilibrium in these games in which the players base their choices on the result of randomizing devices; see footnote 6.

The (*Pareto*) *efficient* payoff in a coordination game can always be achieved at an equilibrium. In games *A* and *B* efficiency is achieved when both players choose Up. Efficiency seems a natural history independent prediction; indeed, some of the models we discuss make this prediction.

Following Harsanyi and Selten (1988), we can also classify the equilibria according to their riskiness. Suppose that John has no idea how Mary intends to play the game, and in his ignorance assumes that she is equally likely to play either strategy. The strategy that John would then prefer to play is called the *risk dominant* strategy, and the corresponding equilibrium the risk dominant equilibrium. The other strategy and equilibrium are called *risk dominated.* We demonstrate below that in game *A*, the Up equilibrium is risk dominant, while in *B* it is the Down equilibrium. Hence, in game *B*, the efficient equilibrium and the risk dominant equilibrium fail to coincide; indeed, the latter minimizes the sum of the players' payoffs. Nevertheless, the risk dominance criterion forms the basis for the history independent prediction in two of the models we consider below.

Game *C* is the celebrated *prisoner's dilemma.*[4] In this game, each player is better off playing Down regardless of the behavior of her opponent; it is therefore not a coordination game. The only equilibrium has both players choose Down, despite the fact that efficiency is achieved when both players choose Up. In the prisoner's dilemma, individuals' incentives prevent the achievement of social efficiency.

Evolutionary game theory studies the behavior of a large population of agents who play a game repeatedly and anonymously. The basic model assumes that pairs of members of the population are chosen at random to play a game which is fixed in advance. Players select one of the available strategies, occasionally updating their selection if it performs inadequately.

We can describe the behavior of the population with a number x between zero and one, representing to the proportion of players current selecting Up. We call x as the *state* of the population. The *expected payoffs* from a random match when the state is x are

$$\pi_U(x) = ax + b(1 - x)$$

for a player who chooses Up, and

$$\pi_D(x) = cx + d(1 - x)$$

[4] For an entertaining (and the original) account of this game, see Luce and Raiffa (1957).

for one who chooses Down.[5] Players prefer actions which maximize their expected payoffs.

A state is a *population equilibrium* if no player can improve his payoffs by switching strategies at that state. In coordination games, both All Up ($x = 1$) and All Down ($x = 0$) are equilibria. If all players in the population are choosing, say, Up, then no player has an incentive to switch to Down, since he will be worse off in matches with his Up playing opponents. Formally, we check that state 1 is an equilibrium by observing that the expected payoff to Up at state 1, $\pi_U(1) = a$, exceeds the payoff to down, $\pi_D(1) = c$. Hence, no player would benefit by switching from his current choice of Up to Down.

The All Up and All Down equilibria can be viewed as *conventions*. A player prefers to play Up so long as most of his opponents play Up; the same is true of Down. Each convention is a self-reinforcing pattern of social conduct. Absent other information about the population's behavior, each constitutes a reasonable prediction about play.

Every coordination game possesses a third population equilibrium, namely

$$x^* = \frac{d-b}{(a-c)+(d-b)}.$$

At this state, a fraction of x^* of the players play Up, while the remainder play Down. At state x^*, payoffs to the two strategies are equal: $\pi_U(x^*) = \pi_D(x^*) = \frac{ad-bc}{(a-c)+(d-b)}$. Thus, no player can benefit from switching strategies. As we shall see, this equilibrium is inherently unstable and hence unlikely to arise. Nevertheless, it will play an important role in the analyses which follow.[6]

---

5 These definitions assume an infinite population; for large but finite populations they are close approximations.

6 The population equilibrium x^* corresponds to a *mixed strategy Nash equilibrium* of the one-shot game. In this equilibrium, players use randomizing devices to determine which strategy to play. Each player flips a coin which comes up heads with probability x^* and plays Up if she sees a head. Since $\pi_U(x^*) = \pi_D(x^*)$, each player's randomization makes her opponent precisely indifferent between playing Up and Down. This indifference justifies the opponent's willingness to randomize between the two strategies. For more on the interpretation of mixed strategy equilibrium in one-shot games, see Osborne and Rubinstein (1994, Chapter 3).

3. Deterministic Dynamics

Equilibrium play is a sensible prediction because it is self-reinforcing. Once an equilibrium is reached, no player has reason to deviate. A fundamental difficulty in game theory is the problem of multiple equilibria: when a game possesses more than one equilibrium, which is the appropriate prediction? Evolutionary game theory addresses this issue through dynamic analysis. Knowledge of the adjustment process and initial behavior together yield the basis for a unique prediction about future play.

Evolutionary dynamics are based on myopic adjustment. They are an application of the principle that when the population is not in equilibrium, players switch to those strategies which currently perform well. They do so myopically, simply choosing those which perform better given the current state. By basing evolutionary dynamics on myopic adjustment, we make only weak demands on the abilities of the players. We need only assume that they know the payoff structure of the game and the current population state and are able to compute their expected payoffs. In fact, under certain assumptions myopic dynamics can be derived from a process of imitation: players simply copy the behavior of opponents whose payoffs are higher than their own. Either way, the demands on players are far less than under the eductive paradigm.

The simplest dynamics to write down are deterministic with a discrete time parameter. They are defined by a function which maps each state x_t to the state x_{t+1} which arises after a round of adjustment occurs. In two strategy coordination games, any dynamics based on myopic optimization have essentially the same character. Myopic adjustment implies that the representation of the better performing strategy increases over time. That is, x increases when $\pi_U(x) > \pi_D(x)$ (and $x \neq 1$), x decreases when $\pi_D(x) > \pi_U(x)$ and (and $x \neq 0$), and x stays fixed when $\pi_U(x) = \pi_D(x)$. In games A, B, and C, this property suffices to describe evolution. The *phase diagrams* of the evolutionary dynamics in these games are presented in Figure 2.

Recall that in game C, the prisoner's dilemma, players prefer to choose Down regardless of the behavior of their opponents. Therefore, evolution always increases the representation of Down, and the All Down equilibrium eventually is reached regardless of the initial state. Hence, in the prisoner's dilemma, a simple dynamic analysis generates a history independent prediction: we should eventually expect all players to choose Down. Of course, since this game has only one equilibrium, this prediction is not particularly surprising.[7]

In contrast, coordination games have three equilibria. The dynamics divide the state space into two *basins of attraction*, one for each convention. For example, in game A, if the initial state has more than one third of the population playing Up, the proportion of players selecting Up increases until all choose Up. Similarly, if more than two thirds of the players initially choose Down, everyone eventually plays Down. Thus, knowledge of the initial behavior of the population allows us to make a unique prediction about the course of play.

$$A = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \qquad B = \begin{bmatrix} 5 & 0 \\ 4 & 2 \end{bmatrix} \qquad C = \begin{bmatrix} 3 & 0 \\ 4 & 1 \end{bmatrix}$$

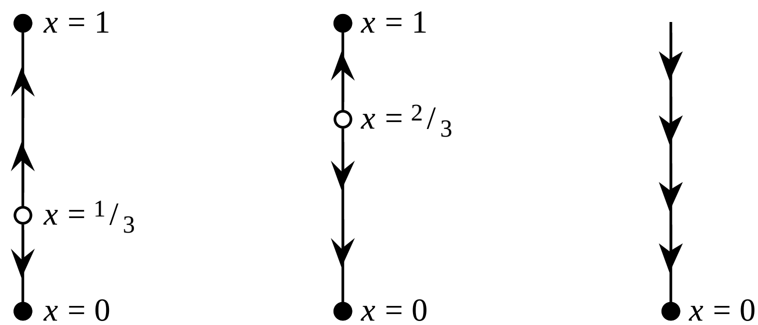

Figure 2: Deterministic Dynamics

We observe that the state x^* which separates the two basins of attraction is the mixed strategy equilibrium of the game. If in game A, exactly one third of the population plays Up, both strategies receive the same payoff ($\pi_U(\frac{1}{3}) = \pi_D(\frac{1}{3}) = \frac{2}{3}$), and no adjustment occurs. However, this equilibrium is *dynamically unstable*: any slight change in the state will lead the population towards either All Up or All Down. Because of the instability of the mixed equilibrium, only the two pure strategy equilibria constitute reasonable predictions about play.

While the mixed equilibrium is unlikely to persist, it does help us characterize risk dominance. Recall that the risk dominant strategy is the one a player chooses in

---

7 On the other hand, the existence of a unique equilibrium does not necessarily imply that it is the only possible result of evolution. For example, if one allows more than two strategies, one can find games and evolutionary dynamics such that the unique equilibrium of the game is dynamically unstable and evolution leads to cyclical behavior. See Gilboa and Matsui (1991) for an example.

a one-shot game when he is completely uncertain about the play of his opponent. Equivalently, it is the strategy he would prefer to choose in the population game when the state is $\frac{1}{2}$. Examining Figure 2, we see that in game *A*, this state is in the basin of attraction of All Up, which is therefore the risk dominant equilibrium. In game *B*, All Down is risk dominant.

Since the mixed equilibrium separates the two basins of attraction, it characterizes risk dominance. In particular, if $x^* < \frac{1}{2}$, All Up is risk dominant, while if $x^* > \frac{1}{2}$, All Down is risk dominant. We can also characterize the risk dominant equilibrium by noting that it always has the larger basin of attraction. It is this property which drives the history independent predictions of stochastic stability models.

4. Stochastic Stability

Under deterministic dynamics, initial conditions dictate the population's entire course of play. Hence, when multiple conventions are possible, the dynamics alone do not foretell which will eventually be established. The remainder of this survey considers models which yield unique, history independent predictions. Unique prediction is clearly at odds with the stability of multiple conventions which exists under deterministic dynamics. Consequently, the models which we now consider employ stochastic dynamics. As we shall see, slight but persistent randomness, even in minute amounts, can radically effect the long run behavior of a dynamical system, making history independent prediction possible.

In the next two models we consider, randomness takes the form of *choice trembles*: occasionally, players fail to play an optimal response to the behavior of the population. This may be a result of experimentation, the entry of newcomers unfamiliar with established conventions, or simple errors in judgment. That players might sometimes make mistakes is uncontroversial. The surprising finding of the stochastic stability literature is that the introduction of trembles can generate unique predictions. Important stochastic stability models include Foster and Young (1990), Kandori, Mailath, and Rob (1993), Nöldeke and Samuelson (1993), Young (1993), Samuelson (1994), and Ellison (1995). Our treatment follows Kandori, Mailath, and Rob (1993).

For every state of a system, a deterministic dynamic specifies the state which must follow. Consequently, the initial state preordains the entire course of evolution. Adding even slight random influences can lead to radically different behavior, especially in the long run. *Markov chains* are the closest stochastic analog to a deterministic system. Rather than specify particular successors, a Markov chain instead specifies a probability distribution of possible transitions from each state. Hence, any initial condition can lead to an infinite variety of evolutionary paths. Crucially, in contrast with the deterministic case, the long run behavior of the system can be completely independent of the initial state. It is this property which underlies history independent prediction.

We introduce Markov chains in the context of a simple stochastic stability model. As we saw in the previous section, deterministic dynamics force the population to one of the two stable conventions. Now suppose there is a possibility of choice trembles: even when society is coordinated upon a convention, players occasionally experiment with the unused strategy. If the initial state is All Up, and enough players experiment simultaneously, it is possible that the proportion of players choosing Up will leap from one to something less than x^*, and hence into the basin of attraction of All Down. If this happens, all players begin switching to Down; barring another focused episode of experimentation, All Down will eventually be established. Jumps of this kind are rare, but regardless of the infrequency of experimentation they are always possible; this fact has surprising implications for the population's long term behavior.

We now describe a simple version of the stochastic stability model. Assume that the population size is some finite number N, and that the population initially plays All Up or All Down. Each period begins with the possibility of experimentation. Each player independently with some small probability ε switches strategies: that is, each flips a coin which comes up heads with probability ε and switches strategies if the result of the toss is heads. After this, all players switch to a best response to the resulting distribution of strategies.[8] Thus, at the end of each period, the population coordinates on a convention; the convention only changes if a rash of experimentation sends the distribution of strategies in the population to the other side of the mixed strategy equilibrium.

We now define a Markov chain based on this verbal description of the model. A Markov chain is defined by its state space and the probabilities of transitions between any pair of states. In this simple model, the state space is the two element

---

8 It should be noted that the qualitative results described above continue to hold if rather than switch immediately to the best response, players only switch to the best response in any particular period with some small, fixed probability.

set {All Up, All Down}. We determine the transition probabilities by computing the likelihood that experimentation upsets a convention. Let P_{DU} denote the probability that a population coordinated upon All Down switches to All Up. For this to occur, at least fraction x^* of the population must experiment simultaneously. The probability of this event is given in terms of the binomial distribution:

$$P_{DU} = \sum_{z > Nx^*} B_{N,\varepsilon}(z),$$

where

$$B_{N,\varepsilon}(z) = \binom{N}{z} \varepsilon^z (1-\varepsilon)^{(N-z)} \text{ for } z = 0, 1, ..., N.$$ [9]

If less than the critical amount of experimentation occurs, there is no jump; hence $P_{DD} = 1 - P_{DU}$. Similarly, more than fraction $1 - x^*$ must experiment for a jump to occur from All Up. Hence,

$$P_{UD} = \sum_{z > N(1-x^*)} B_{N,\varepsilon}(z),$$

and $P_{UU} = 1 - P_{UD}$. Using these data it is possible to calculate the probability that at any future time t the population coordinates on All Up or All Down. Hence, the initial state and the transition probabilities give us a complete probabilistic description of the course of evolution.

Since all transition probabilities are positive, there is a unique pair of numbers (m_U, m_D) which satisfy the following three equations:

$$m_U + m_D = 1,$$
$$m_U P_{UU} + m_D P_{DU} = m_U,$$
$$m_U P_{UD} + m_D P_{DD} = m_{D.}$$

Since m_U and m_D sum to one, we can think of them as the probabilities that the population is at All Up or All Down at some time t. The interpretation of the latter two equations is that if (m_U, m_D) describes the likelihoods of the two conventions at time t, the probabilities of jumps "cancel each other out", and (m_U, m_D) continues to

9 Suppose one has a coin which comes up heads ε of the time. Then $B_{N,\varepsilon}(z)$ is the probability that exactly z out of N tosses of the coin turn up heads.

describe the likelihoods at time $t + 1$, and, by induction, at all future times. For this reason, (m_U, m_D) is called the *stationary distribution* of the Markov chain.

If as in our model the stationary distribution is unique, it gives important characterizations of long run behavior. Most important for this model is the *time average property*: in the long run, the proportion of time spent at each state is given by the stationary distribution, *regardless of the initial state of the system*. Hence, although the population randomly jumps between states forever, this randomness engenders a regularity in the average behavior of the system.[10]

Solving the simultaneous equations, we see that the long run behavior of the Markov chain is described by

$$(m_U, m_D) = \left(\frac{P_{DU}}{P_{UD} + P_{DU}}, \frac{P_{UD}}{P_{UD} + P_{DU}} \right).$$

Thus, the long run behavior of the system depends only on the relative likelihoods of jumps.

When ε is small and N large, it can be shown that either $m_U \approx 1$ or $m_D \approx 1$, depending on whether All Up or All Down is risk dominant equilibrium. Therefore, in the long run, nearly all time is spent at the risk dominant equilibrium. Hence, we can predict that regardless of the initial behavior of the population, the risk dominant convention will eventually have been in force for a vast majority of play. This is the history independent prediction of the stochastic stability model.

The intuition behind this result can be illustrated using the phase diagrams in Figure 2. For a jump between equilibria to occur, enough experimentation must occur at once to move the population to the other basin of attraction. Clearly, the larger the basin of attraction, the more jumps that are required. When the rate of mutation is small, a jump requiring more mutations is far less likely than one requiring fewer. But the risk dominant equilibrium is the equilibrium with the larger basin of attraction and hence the equilibrium from which it is harder to jump. Therefore, in the long run, the population will spend far more time at the risk dominant equilibrium than at the risk dominated equilibrium.[11]

---

[10] A special case of the time average property is the Strong Law of Large Numbers, which guarantees, for example, that in the long run, the proportion of tosses of a fair coin which come up heads will equal one half.

[11] To see this more formally, suppose that $x^* < 1/2$, so that All Up is the risk dominant equilibrium. A jump from All Up from all Down takes at least $N(1 - x^*)$ experiments, while a jump in the other direction takes at least Nx^* experiments. Since $x^* < 1/2$, when N is large, the former jump requires

What can be said if the underlying game has more than two strategies? Ellison (1995) shows that equilibrium selection results still hold in such games so long as one of the strategies satisfies $\frac{1}{2}$*-dominance*, a notion due to Morris, Rob, and Shin (1995). A strategy is $\frac{1}{2}$-dominant if it is a best response whenever at least half of a player's opponents are playing it. By definition, any risk dominant strategy in a two strategy game is $\frac{1}{2}$-dominant. Unfortunately, not all games possess a $\frac{1}{2}$-dominant strategy; in these cases, determining the stochastically stable outcome is a somewhat more complicated task. See Kandori and Rob (1995) and Ellison (1995) for further discussion.

Stochastic stability analysis provides a mathematically rigorous history independent prediction. How believable is this prediction over the lengths of time relevant in social science models? As we discussed in the Introduction, a history independent prediction is only credible if the predicted outcome will arise with a very high probability over a time span which is meaningful given the object of study. A number of authors, including Ellison (1993), Binmore, Samuelson, and Vaughan (1995), and Sandholm and Pauzner (1998), have observed that the prediction of the basic stochastic stability model does not satisfy this criterion.

To illustrate this, suppose that a population playing game A initially coordinates on All Down. The prediction of the stochastic stability model is All Up. However, in order for All Up to arise, a jump must occur: more than one third of the population must experiment simultaneously. When the experimentation rate is small, even moderate population sizes make this jump extremely unlikely. For example, if $\varepsilon = .1$ and $N = 200$, the probability of this jump is $P_{DU} = 1.52 \text{ x } 10^{-19}$; the expected number of periods before the first jump occurs is therefore $\frac{1}{P_{DU}} = 6.57 \text{ x } 10^{18}$ periods.[12] Regardless of our interpretation of a period, this number is enormous. Even if a period lasts only a second, we should not expect to see the first jump for 200 billion years.

Sandholm and Pauzner (1998) extend this criticism by considering a model in which the population grows over time. They show that even extremely slow growth overturns the history independence results described above. If experimentation is sufficiently rare, a growing population is virtually guaranteed

---

many more experiments than the latter. So, when ε, the probability of experimentation, is small, the probability of the former, P_{UD}, is orders of magnitude smaller than the probability of the latter, P_{DU}. Thus, applying the formula above, we see that $m_U = P_{DU}/(P_{UD} + P_{DU})$ is very close to one.

[12] This follows because the number of periods before the first jump is geometrically distributed with parameter P_{DU}.

never to jump. In stark contrast to the fixed population model, the first equilibrium played is quite likely to be played forever. Hence, even if the population is initially quite small, if it is growing the history independent prediction is unlikely to be relevant. In light of these criticisms, it appears that in settings in which the stochastic stability model applies, unless the population size is both small and fixed in size, the relevant prediction is not the risk dominant equilibrium, but rather the equilibrium predicted by the deterministic dynamics.

5. Stochastic Stability with Local Interaction

The predictions of the basic stochastic stability model require an unreasonably long time span to become relevant for social science modeling. A substantial proportion of the population must experiment simultaneously to induce any change in convention. When the population is large, these changes are too rare to form the basis for a believable prediction.

The basic stochastic stability model is predicated on uniform random matching: all matches between players are equally likely. However, it is often natural to think that certain matches are more likely than others. If the players each reside in some particular physical location, we should expect that most encounters will take place between players who live close by. Rather than worry about the behavior of the entire society, such players are mainly concerned with the actions of their neighbors. If your neighbors all play Up, you too should play Up, regardless of the behavior of the population at large. Hence, a small, localized set of choice trembles can induce changes in neighbors' behavior. How does local interaction affect the way that conventions evolve? Work in this area includes papers by Blume (1993, 1995) and Ellison (1993, 1995); our treatment follows Ellison (1993).

Imagine a society of N players whose locations are specified by positions on a circle (see Figure 3), and suppose for simplicity that each only interacts with his immediate neighbor on either side. There are three possible distributions of one's neighbors' strategies: both Up, both Down, or one of each. Clearly, if both of one's neighbors are playing the same strategy, the optimal response is to follow along. If the neighbors are split, the expected payoffs to each strategy are $\pi_U(\frac{1}{2})$ and $\pi_D(\frac{1}{2})$. By definition, the higher value will accrue to the risk dominant strategy. Hence, if at least one neighbor plays the risk dominant strategy, it is optimal to play the risk dominant strategy in response.

Now suppose that in each period, each player chooses a best response to his two neighbors, except on occasions on which she experiments. Formally, this is modeled via a Markov chain with 2^N states: one for each possible specification of the N players' actions. As before, the stationary distribution describes the long run time averaged behavior; once again, nearly all weight in the stationary distribution is placed on the state at which all players choose Up. However, in contrast to the basic stochastic stability model, under local interaction the All Up convention is a credible prediction regardless of the initial state.

Consider a population of twelve players located around a circle. They repeatedly play game A, in which All Up is risk dominant, but are initially coordinated upon All Down. Assume that in each period, each player adjusts to play the best response given the behavior of his neighbors. Now, suppose that two adjacent players simultaneously tremble, switching to Up. In the following period, these two players continue to play Up, as do their two immediate neighbors. In each of the periods which follow, barring further trembles, two more players switch to Up, until all players have switched, at which point strategy choices stabilize. Thus, two trembles to Up are enough to effect the institution of the risk dominant strategy throughout the population. The spread of the risk dominant strategy through local interaction is known as *contagion*, and is illustrated in Figure 3.

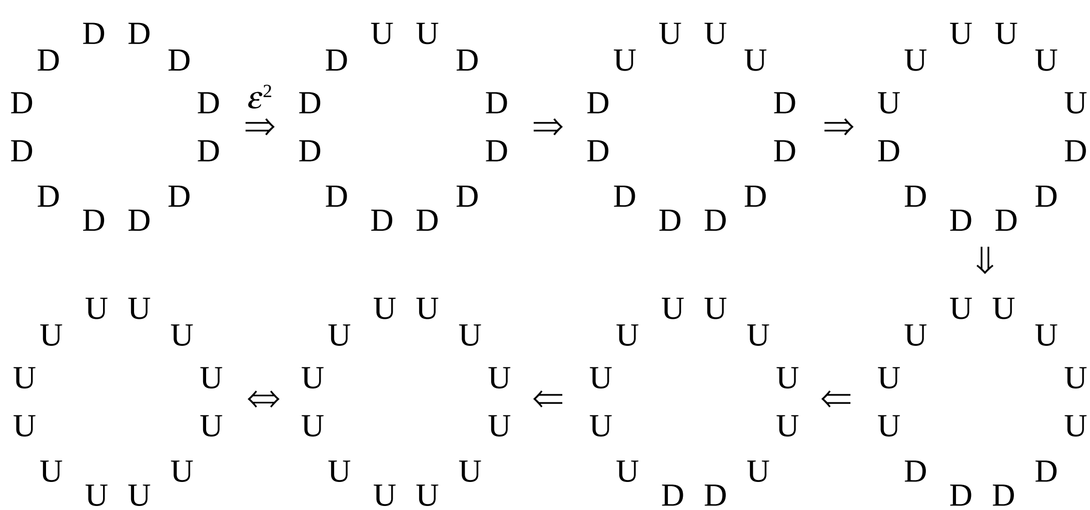

Figure 3: Contagion

Contagion is irreversible: coordination on Up, once established, is very difficult to break. As long as two adjacent players remain who play Up, myopic best responses will quickly return Up to predominance. No less than half the population must simultaneously switch to Down in order for the All Up convention to be disturbed. Hence, coordination on All Up is as stable here as it is in the basic stochastic stability model. Moreover, even if the risk dominated convention is established first, the amount of time before it is disrupted is not very large. In the two neighbor matching model described above, only two mutations are needed to break coordination on All Down. This is true regardless of the size of the population: the behavior of a million member society can be radically changed by a minute spark of experimentation.[13] Thus, the local interaction makes credible history independent predictions: regardless of the initial conditions, we can expect the risk dominant convention to emerge with very high probability reasonably soon; once established, the time span before it is dislodged is astronomical.[14]

While in Ellison's (1993) model the risk dominant equilibrium is selected, local interaction can lead to other outcomes if we change the assumptions governing individual players' behavior. For example, consider a situation in which Up and Down correspond to two languages. Through some extra effort, an agent could learn both languages, easing communication in all encounters.

With this motivation, we consider a variant of a local interaction model of Goyal and Janssen (1997) which admits the possibility of *non-exclusive conventions*. That is, in addition to playing Up or Down, a player can choose Both. In doing so, he plays Up against opponents playing Up and Down against those playing Down. Moreover, when meeting another player who chooses Both, they can choose to follow the efficient convention, which in games A and B is Up. However, the flexibility of playing Both carries with it a fixed cost which must be paid each time Both is chosen.

Applying tools for stochastic stability analysis developed by Ellison (1995), one can show that if the cost of playing Both is low (in games A and B, less than $\frac{1}{2}$), then the only stable long run outcome is coordination on the Pareto dominant equilibrium. Thus, in contrast with the stochastic stability models considered above,

---

[13] While we have focused on two neighbor matching, this property holds more generally: the time needed to disrupt the risk dominated convention depends only on the size of the neighborhood, while the time needed to disrupt the risk dominant convention depends on the size of the whole population.

[14] In analogue with his global interaction results, Ellison (1995) shows that under local interaction, when the underlying game has more than two strategies, any 1/2-dominant equilibrium will be stochastically stable.

evolution in this model always leads to efficient play.[15] To see why, suppose that the underlying game is *B*, and consider a player who faces one neighbor playing Down and another playing Up. If playing Both was not an option, the best response would be to play the risk dominant strategy, Down. However, so long as playing Both is not too costly, it is worthwhile to play Both and so to coordinate with both neighbors in the succeeding periods. But once both of the player's neighbors are playing either Up or Both, by playing Up he can coordinate with both neighbors without paying the flexibility fee. Through such strategy adjustments, the possibility of flexibility becomes a mechanism which allows efficient behavior to emerge.

Is it possible for different, interconnected regions to follow different conventions? A number of authors, including Sugden (1995), Anderlini and Ianni (1996), and Goyal and Janssen (1997) consider the possibility of the coexistence of conventions in the absence of repeated choice trembles. Under certain assumptions on the neighborhood structure of the local interaction model, it becomes possible for different conventions to exist in different regions of the interaction space. Moreover, this coexistence of conventions can be stable in the face of one-time choice trembles. Whether the coexistence of conventions can be stochastically stable (i.e., whether it persists under repeated choice trembles) is a question requiring further investigation.

6. Communication

Without non-exclusive conventions, the stochastic stability models presented above predict that society will coordinate on the risk dominant equilibrium. While there is good reason for a player who is uncertain about the behavior of his opponents to play a risk dominant strategy, this can lead to undesirable outcomes: in game *B*, the risk dominant equilibrium is the point at which societal welfare is minimized.

In both games *A* and *B*, the outcome which arises when all players choose Up is every player's favorite. Suppose that before playing the game, each player can express his intention to play Up. Can the opportunity to communicate guarantee efficient play?

---

[15] Non-exclusive conventions can also be considered in the global interaction framework of Section 4. In this setting, the Pareto dominant equilibrium is selected, but once again the waiting time before selection occurs is extremely long.

There is only weak support for the notion that non-binding communication, also known as *cheap talk*, can lead fully rational players to play efficient equilibria.[16] First, there exist "babbling equilibria" in which players ignore each others' signals to play the efficient equilibrium. Second, even if we assume that players listen to credible signals, efficiency is not guaranteed. For example, in game B, regardless of his own action, each player always prefers that his opponent play Up. Hence, signaling that opponent to play Up hardly constitutes a credible commitment.[17]

In contrast to the eductive approach, the evolutionary approach to games with communication assumes that signals have no intrinsic meaning. Rather, any meaning that the signals acquire is built up endogenously. More importantly, through evolution the efficient equilibrium can become the unique prediction for long run play. A number of authors, among them Robson (1990), Matsui (1991), Wärneryd (1991), Kim and Sobel (1992, 1995), Banerjee and Weibull (1993), and Schlag (1993) have studied evolution in games with communication. The model we present is adapted from Kim and Sobel (1995).

For concreteness, we consider the evolution of play of game A; the analysis is identical for other two strategy coordination games. A population of N players is repeatedly randomly matched to play this game. Each match consists of the play of a two stage communication game. In the first stage, each player sends one of two messages, numbered 1 and 2. Then, after hearing his opponents' message, each player chooses Up or Down. Thus, each communication game strategy consists of three elements: a message, and the base game strategy to be played in response to messages 1 and 2. For example, the strategy 1UD is one in which message 1 is sent, Up is played as a response to the play of message 1 by an opponent, and Down is played in response to message 2.

The payoffs generated in a match only depend on the coordination game strategies that end up being played. For instance, if a player choosing 1UD meets another playing 2DU, both players play Down, resulting in a payoff of 1. Many strategy combinations yielding the same outcome: for example, 1UD vs. 2DD and 1DD vs. 2DD also lead to a payoff of 1. As we shall see, it is precisely this redundancy which makes a unique prediction possible.

There are 2 x 2 x 2 = 8 strategies available in the communication game. As in the basic stochastic stability model, players are randomly matched, with a match with

[16] For treatments of cheap talk within the standard game theoretic paradigm, see, for example, Farrell (1988), Myerson (1989), and Rabin (1994).

[17] This observation is due to Aumann (1990).

any other member of the population equally likely. Hence, the performance of each communication game strategy depends on the current distribution of communication game strategies within the population.

In the stochastic stability models, a stochastic component was introduced to the evolutionary process via choice trembles. Here, chance elements are not due to suboptimal choices. In the communication model, randomness instead takes the form of *drift* between equally profitable strategies. In particular, we assume that in each period, one player gets the opportunity to switch strategies, and chooses randomly from among his *acceptable* communication game strategies: those that perform at least as well as his current choice against the current distribution of strategies in the population.

The model predicts that all players eventually play Up. However, unlike in the previous models, the All Up convention is associated with numerous combinations of communication game strategies, each of which leads to play of All Up in the base game. It is clear that if all players choose 1UD, then all matches will result in the play of Up in the base game and hence a payoff of 2; this is also true if all players choose 1UU, or 2UU, or 2DU. Furthermore, if players are divided between 1UD and 1UU, the second message is never played, and so Up is always played. Similarly, if the population is divided between 2UU and 2DU, or between 1UU and 2UU, the efficient payoff is achieved. The population will drift perpetually through all of these sorts of states, but to none at which an inefficient payoff is possible: once a state at which all players receive a payoff of 2 is reached, no player will ever receive less, as this would require a switch to a suboptimal strategy. Thus, once the efficient return is achieved, disruption is not just extremely unlikely; it is impossible.

Is there a stable collection of states at which everyone plays Down? The best candidate uses the communication game strategies 1DD and 2DD in combination. If all players choose one of these two strategies, and more than one is playing each, then the only acceptable switches are to 1DD and 2DD: since both messages are in use, and since players always choose Down in the base game, an optimal response must always play Down to obtain the payoff of 1. At these states players simply ignore the messages and play Down regardless of what they see.

Since play drifts among the acceptable strategies, the distribution between 1DD and 2DD fluctuates over time. Eventually, drift will lead the population to all coordinate on a single communication game strategy, say 2DD. At this point, since no player is playing message 1, strategy 1UD becomes acceptable. In the context created by the strategy distribution, a player who chooses 1UD can be imagined telling his opponents, "Everyone else sends message 2 and always plays Down. If you meet a person who sends message 1, it's me; if you signal me by sending message 1 yourself, we can both play Up and receive the good payoff."

Once one person switches to 1UD, the best alternative for other players receiving the opportunity to switch is to also play 1UD; when two such players meet, their "secret handshake" allows them to coordinate on the good outcome without being punished when playing other members of the population.[18] While some players who receive the opportunity to switch may linger playing 2DD, once two players play 1UD, it becomes inevitable that all players will eventually switch to this strategy. Once 1UD is established in this way, all players receive the efficient payoff of 2; as we argued above the efficient payoff is never abandoned. Thus, drift alone can disrupt the All Down convention.[19]

To test whether this prediction is credible, we must consider whether a population initially coordinated upon the All Down convention will egress from it in a reasonable time span. Suppose that initially all players choose 1DD or 2DD. Until only one message is used, the only acceptable switches are to 1DD and 2DD. Hence, the population continues to be divided between 1DD and 2DD until one message goes unused. How much time elapses before this occurs?

We can speed up the drift process by assuming that players always switch strategies: that is, the chosen player always decides to alter his message. Suppose that the population size N is even, and for the moment, suppress the possibility that strategies besides 1DD and 2DD will be chosen at the extreme states. In this case, the drift process takes the form of an Ehrenfest chain, a Markov chain which has been studied extensively. If we express the state of the Ehrenfest chain as the number of players choosing message 1, its stationary distribution is binomial with parameters N and $\frac{1}{2}$. Thus, rescaling the state space to reflect the proportion of players choosing message 1, the Central Limit Theorem implies that the stationary distribution is approximately normal with mean $\frac{1}{2}$ and variance $\frac{1}{4N}$. Hence, as N grows large, the

18 The secret handshake metaphor was first proposed by Robson (1990).

19 It turns out that there is another stable state in which half of the players play 1DU, and half play 2UD. At this distribution, players have a nearly even chance of meeting an opponent sending the same message. When two opponents who send the same message meet, they both play Down and receive the payoff of 1; if the players send different messages, they both play down and receive 2. Thus, the average payoff from a match is approximately 3/2; any player who switches strategies receives a strictly lower payoff. However, the stability of this state can be destroyed by altering a modeling detail: by splitting the population into two groups, say, males and females, and then assuming that each match involves exactly one male and one female, the stability of the equilibrium described above is destroyed. For details, see Kim and Sobel (1995).

stationary distribution, and so the long run time average of play, is almost completely contained in a neighborhood of $\frac{1}{2}$: the distribution of messages is nearly always very close to even. Furthermore, if the initial state has half of the players using each message, the expected wait before the state at which all players choose message 1 is reached is approximately $\frac{1}{N}2^N$ periods. For a population size as small as forty, the expected wait is about thirty trillion periods.[20]

Hence, if the population is large, it will take an extraordinarily long time before it drifts to a strategy distribution with an unused message. Inefficient states in which messages are ignored can persist almost indefinitely. Thus, we conclude that the mechanism which generates the history independent prediction is too slow to be of practical relevance.

7. Conclusion

We have considered three models of evolution which generate history independent predictions. The predictions of the basic stochastic stability model and the communication model require far longer time spans to gain relevance than seems sensible in most applications concerning human behavior. In contrast, local interaction models, whose results only depend on small, concentrated choice disturbances, yield credible history independent predictions.

Further study of local interaction models is needed to more thoroughly assess the extent to which they make unique predictions possible.[21] On the other hand, when interactions are of a more global nature, the results described here indicate that predictions of play may not easily be disentangled from historical conditions. While game theoretic analysis can rule out most non-equilibrium societal

---

[20] One could obtain very different results than the ones we have described by assuming that we choose a *message* at random and increase its representation by one. This results in a standard random walk, which arrives at the boundaries considerably faster than does the Ehrenfest chain. However, selecting a message (rather than a player) at random requires that we not treat the players in an anonymous fashion, which seems an inappropriate modeling assumption.

The Ehrenfest chain was originally introduced as a probabilistic model of physical equilibrium. Interestingly, it was created in order to argue that although all states are possible, only those with a nearly even distribution of messages are ever likely to be observed. For further discussion of the Ehrenfest chain, as well as derivations of the formulas used above, see Bhattacharya and Waymire (1990, Chapter III.5).

[21] Recent work on local interaction models includes Mailath, Samuelson, and Shaked (1994), Ely (1995), and Morris (1996, 1997).

outcomes, selection among those that remain may not always be reducible to a question of payoff structure, but only to one of precedence.

References

Anderlini, L., and A. Ianni (1996). "Path Dependence and Learning from Neighbors," *Games and Economic Behavior* **13**, 141-177.

Aumann, R. (1990). "Communication Need Not Lead to Nash Equilibrium," mimeo, Hebrew University.

Banerjee, A., and J. W. Weibull (1993). "Evolutionary Selection with Discriminating Players," mimeo, Harvard University and Stockholm University.

Bergin, J., and B. L. Lipman (1996). "Evolution with State Dependent Mutations," *Econometrica* **64**, 943-956.

Bhattacharya, R. N., and E. C. Waymire (1990). *Stochastic Processes with Applications*. New York: Wiley.

Binmore, K. G., and L. Samuelson (1994). "An Economist's Perspective on the Evolution of Norms," *Journal of Institutional and Theoretical Economics* **150**, 45-63.

Binmore, K. G., L. Samuelson, and R. Vaughan (1995). "Musical Chairs: Modeling Noisy Evolution," *Games and Economic Behavior* **11**, 1-35.

Blume, L. (1993). "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior* **5**, 387-424.

Blume, L. (1995). "The Statistical Mechanics of Best Response Strategy Revision," *Games and Economic Behavior* **11**, 111-145.

Ellison, G. (1993). "Learning, Local Interaction, and Coordination," *Econometrica* **61**, 1047-1071.

Ellison, G. (1995). "Basins of Attraction and Long Run Equilibria," mimeo, M.I.T.

Ellison, G. (1997). "Learning from Personal Experience: One Rational Guy and the Justification of Myopia," *Games and Economic Behavior* **19**, 180-210.

Ely, J. (1995). "Local Conventions," mimeo, Northwestern University.

Farrell, J. (1988). "Communication, Coordination, and Nash Equilibrium," *Economics Letters* **27**, 209-214.

Foster, D., and H. P. Young (1990). "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology* **38**, 219-232.

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

Goyal, S. and Janssen, M. C. W. (1997). "Non-Exclusive Conventions and Social Coordination," *Journal of Economic Theory* **77**, 34-57.

Harsanyi, J. C., and R. Selten (1988). *A General Theory of Equilibrium in Games*. Cambridge: MIT Press.

Kandori, M., G. J. Mailath, and R. Rob (1993). "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* **61**, 29-56.

Kandori, M., and R. Rob (1995). "Evolution of Equilibria in the Long Run: A General Theory and Applications." *Journal of Economic Theory* **65**, 383-414.

Kim, Y.-G., and J. Sobel (1992). "An Evolutionary Approach to Pre-Play Communication," mimeo, University of Iowa and UCSD.

Kim, Y.-G., and J. Sobel (1995). "An Evolutionary Approach to Pre-Play Communication," *Econometrica* **63**, 1181-1193.

Luce, R. D., and H. Raiffa (1957). *Games and Decisions*. New York: Wiley.

Mailath, G., L. Samuelson, and A. Shaked (1994). "Evolution and Endogenous Interactions," CARESS Working Paper #94-13, University of Pennsylvania.

Matsui, A. (1991). "Cheap-Talk and Cooperation in a Society," *Journal of Economic Theory* **54**, 245-258.

Maynard Smith, J. (1974). "The Theory of Games and The Evolution of Animal Conflict," *Journal of Theoretical Biology* **47**, 209-221.

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

Morris, S. (1996). "Contagion," mimeo, University of Pennsylvania.

Morris, S. (1997). "Interaction Games: A Unified Analysis of Incomplete Information, Local Interaction, and Random Matching," mimeo, University of Pennsylvania.

Morris, S., R. Rob, and H. S. Shin (1995). "*p*-dominance and Belief Potential." *Econometrica* **63**, 145-157.

Myerson, R. (1989). "Credible Negotiation Statements and Coherent Plans," *Journal of Economic Theory* **48**, 264-291.

Nöldeke, G., and L. Samuelson (1993). "An Evolutionary Analysis of Backward and Forward Induction," *Games and Economic Behavior* **5**, 424-454.

Osborne, M. J., and A. Rubinstein (1994). *A Course in Game Theory*. Cambridge: MIT Press.

Rabin, M. (1994). "A Model of Pre-Game Communication," *Journal of Economic Theory* **63**, 370-391.

Robson, A. J. (1990). "Efficiency in Evolutionary Games: Darwin, Nash, and the Secret Handshake," *Journal of Theoretical Biology* **144**, 379-396.

Samuelson, L. (1994). "Stochastic Stability in Games with Alternative Best Replies," *Journal of Economic Theory* **64**, 35-65.

Sandholm, W. H., and A. Pauzner (1998). "Evolution, Population Growth, and History Dependence," *Games and Economic Behavior* **22**, 84-120.

Schlag, K. H. (1993). "Cheap Talk and Evolutionary Dynamics," mimeo, Universität Bonn.

Sugden, R. (1995). "The Coexistence of Conventions," *Journal of Economic Behavior and Organization* **28**, 241-256.

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

Wärneryd, K. (1991). "Evolutionary Stability in Unanimity Games with Cheap Talk," *Economics Letters* **36**, 375-378.

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

Young, H. P. (1993). "The Evolution of Conventions," *Econometrica* **61**, 57-84.