

## An Evolutionary Analysis of Backward and Forward Induction\*

GEORG NÖLDEKE

*Department of Economics, Princeton University, Princeton, New Jersey 08544*

AND

LARRY SAMUELSON

*Department of Economics, University of Wisconsin, Madison, Wisconsin 53706*

Received July 16, 1992

We examine the limiting outcomes of a dynamic evolutionary process driven by stochastic learning and rare mutations. We first show that *locally stable outcomes* are subgame perfect and satisfy a forward induction property. To address cases in which locally stable outcomes fail to exist, we turn to a dynamic analysis. The limiting distribution of the dynamic process in a class of extensive form games with perfect information always includes the subgame perfect equilibrium outcome, but consists exclusively of that outcome only under stringent conditions. The limiting distribution in a class of outside option games satisfies a forward induction requirement. *Journal of Economic Literature* Classification Numbers C70, C72. © 1993 Academic Press, Inc.

### 1. INTRODUCTION

This paper examines the limiting behavior of a dynamic evolutionary process driven by stochastic learning and rare mutations. The analysis is focused on extensive form games. We are especially interested in whether

\* This paper originated from discussions with Eric van Damme, for whose help, insight, and encouragement we are grateful. We thank Dieter Balkenborg for helpful discussions and thank Vince Crawford, Joel Sobel, and two referees for their detailed comments. Financial support from the National Science Foundation and the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at the University of Bonn, is gratefully acknowledged.

the process yields outcomes that exhibit backward induction properties (such as subgame perfection) and forward induction properties [such as those examined by van Damme (1989)].

We first examine what we call *locally stable outcomes*. Intuitively, our definition of local stability requires that *any* strategy combination yielding a locally stable outcome is surrounded by learning dynamics that at least eventually lead back to that outcome. Our interest in these outcomes not only has an evolutionary motivation, as do static evolutionary stability concepts, but emerges from our dynamic model. If our evolutionary process selects a unique outcome, this outcome must be locally stable.

Locally stable outcomes exhibit both backward and forward induction properties. In extensive form games in which each player moves at most once along any path, every locally stable outcome is a subgame perfect equilibrium outcome. Furthermore, every locally stable outcome must satisfy a forward induction property. In two-player games from our class, this property implies the Never-Weak-Best-Response property.

In many games, locally stable outcomes fail to exist.<sup>1</sup> In such games, the limiting distribution of the dynamic process assigns strictly positive probability to multiple outcomes that are contained in *locally stable components* of absorbing sets of the learning process. A component is locally stable if it is a minimal subset of absorbing sets with the property that the learning dynamics at all nearby strategies lead back to that component.

To address the question of whether the limiting distribution will satisfy backward and forward induction properties in games where it does not generate a locally stable outcome, we turn to an analysis of the dynamic process. We consider two simple classes of games, allowing us to deal with backward and forward induction one at a time.

We first examine backward induction in games of perfect information in which each player moves at most once along any path. In these games, there is a unique locally stable component, containing the subgame perfect equilibrium outcome. Local stability of the subgame perfect equilibrium outcome is thus both necessary and sufficient for the limiting distribution to consist entirely of this outcome. However, the subgame perfect equilibrium outcome will be locally stable only under stringent conditions. If these conditions are not met, the subgame perfect equilibrium still appears in the limiting distribution (since it is contained in the unique locally stable component), but is accompanied by other self-confirming equilibria that yield different outcomes.

<sup>1</sup> It is also possible for locally stable outcomes to exist without appearing in the limiting distribution of the dynamic process, which must then be multivalued.

We next examine forward induction in games where player 1 has a choice between exercising an outside option or playing a normal form game with player 2. We find that if there is *any* strict Nash equilibrium of the subgame yielding a higher payoff to player 1 than the outside option, then equilibria in which the latter is played cannot appear in the limiting distribution.

Van Damme's (1987b, 1989) notion of forward induction, applied to generic games in this class, is that if there is one and only one Nash equilibrium of the subgame that offers player 1 a higher payoff than the outside option (with other Nash equilibria of the subgame offering lower payoffs), then equilibria in which player 1 chooses the outside option will not appear. The uniqueness requirement appears in van Damme's definition because if there are multiple equilibria in the subgame that dominate the outside option for player 1, then entering the subgame provides an ambiguous signal and van Damme's forward induction argument loses its force. In the evolutionary model, however, the evolutionary process "assigns" a meaning to the event of entering the subgame, even if this meaning is *a priori* ambiguous, allowing us to obtain a stronger result.

Our results build upon a large literature that investigates the connection between evolutionary models and equilibrium refinement properties. The point of departure for this literature is van Damme's (1987a) demonstration that an evolutionarily stable strategy (ESS) in a symmetric game must be a proper equilibrium. Coupled with the result that proper equilibria in normal form games correspond to sequential equilibria in extensive form games (van Damme, 1984; Kohlberg and Mertens, 1986), this associates a backward induction property with evolutionarily stable strategies.

Unfortunately, evolutionarily stable strategies do not exist in many games of interest in economics.<sup>2</sup> A collection of alternative concepts has been created to address this difficulty. These concepts generally involve some provision for unreached information sets or alternative best replies that are reminiscent of equilibrium refinements. For example, Selten's (1983, 1988) limit ESS concept incorporates a normal form perfection requirement (cf. Samuelson, 1991b). A strong result is obtained by Swinkels (1992a), who shows that his equilibrium evolutionarily stable sets contain Kohlberg-and-Mertens stable sets.

Rather than work with refinements of static equilibrium concepts, our

<sup>2</sup> The reasons for this nonexistence are closely related to the issues that drive the equilibrium refinement literature. For example, an equilibrium in an extensive form game that does not reach every information set with positive probability will fail to be an ESS (Selten, 1983, 1988), and in an asymmetric normal form game, any equilibrium to which there are alternative best replies will not be an ESS (Selten, 1980).

approach is to examine an explicitly dynamic model of the evolutionary process. Swinkels (1993) also examines refinement ideas in the context of a dynamic model. He introduces a stability notion for sets of outcomes that is analogous to our local stability and shows that for a broad class of deterministic dynamics, sets satisfying this assumption must contain a Kohlberg-and-Mertens stable and hyperstable set.

Our analysis is similar to that of Swinkels in several respects. First, we both obtain set-valued outcomes. In each case, this set-valuedness is driven by the inability of the evolutionary process to prevent the system from drifting to alternative best replies. Second, in each case, "good" outcomes, such as subgame perfect or stable equilibria, are contained in the limiting outcome. However, each model encounters difficulties in restricting the limiting outcome to include only "good" outcomes of this type. In our case, if any Nash equilibrium strategy profile appears in the limiting distribution, then so do any other strategy profiles created by drift at unreached subgames. These can include strategies that are not Nash equilibria and that can be surrounded by learning dynamics that lead to quite different outcomes, causing the latter to also appear in the limiting distribution. In Swinkels' model, similar forces appear. Every Nash equilibrium appearing in the outcome is accompanied by a host of strategy profiles that give the same outcome, in some cases including profiles surrounded by dynamics that preclude asymptotic stability. Swinkels observes that his model holds some hope of yielding asymptotically stable sets that contain only subgame perfect outcomes if the dynamics have rest points only at *Nash* equilibria (as opposed, for example, to having rest point at all self-confirming equilibria, as in our model). Such an assumption appears to be inappropriate in our context because it requires players to react to changes in actions at information sets that are not reached during play. In an effort to respect the extensive forms of our games, we assume that players can observe (and hence react to) only changes in actions that are revealed in the course of play.

The most difficult aspect of a dynamic analysis is generally establishing conditions under which the dynamic process will converge. This problem is often avoided, with results taking the form of characterizing the limiting behavior of the process if it does converge. Our work with locally stable outcomes is in this vein. Crawford (1992), Milgrom and Roberts (1990, 1991), and Marimon and McGrattan (1992) examine conditions under which certain dynamic systems converge.

Section 2 presents the basic framework for our analysis. Section 3 examines locally stable outcomes. Section 4 presents the analysis of backward induction. Section 5 examines our class of forward induction games. Section 6 concludes.

## 2. THE MODEL

We work with a finite, stochastic evolutionary model. The study of stochastic evolutionary models in games was pioneered by Foster and Young (1990) and Young (1993) and pursued by Kandori *et al.* (1993) and then Samuelson (1991a). Our model borrows heavily from Samuelson (1991a), which in turn relies heavily on the work of Kandori *et al.*

Let  $G$  be a finite extensive form game with perfect recall and without moves by nature. Let  $I = 1, \dots, n$  denote the set of players and  $Z$  the set of terminal nodes (outcomes). An assignment of payoffs to terminal nodes is given by the function  $\pi: Z \rightarrow R^n$ . We restrict attention to games  $G$  that satisfy the condition that every path through the game tree intersects at most one information set of every player.<sup>3</sup>

### 2.1. Learning

Given an extensive form game  $G$ , we assume that for every player  $i$  there is a finite population of size  $\Delta > 1$ . A typical member of a population is referred to as an *agent*. At each time  $t \in \{0, 1, 2, \dots\}$ , every possible combination of agents capable of playing the game meets and plays.

At time  $t$ , each agent of each population is described by a *characteristic*, consisting of a pure behavior strategy and a conjecture. A conjecture for an agent specifies, for each population representing an opponent and for each action such an opponent may take at one of his information sets, the number of agents in the population taking that action.<sup>4</sup> A *state* of the system is a specification of how many agents in each population have each possible characteristic. We let  $\Theta$  denote the set of possible states of the system and let  $\theta$  denote an element of  $\Theta$ . Associated with every state is a distribution over terminal nodes, denoted by  $z(\theta)$ , that results from the matching process.

In each period, after agents have been matched, each agent of each population independently takes a random draw from a Bernoulli trial. With probability  $1 - \mu \in (0, 1)$ , the agent's characteristic does not change. With probability  $\mu$ , the draw produces the outcome "learn." We assume that an agent who learns is able to observe the *outcomes* of all matches of the current round of play. Note that we take the extensive form of the

<sup>3</sup> In addition to being a natural class of games in which to examine subgame perfection, because questions of rational behavior after evidence of irrationality do not arise (cf. Binmore, 1987, 1988; Reny, 1993), this assumption allows us to specify very straightforward rules for updating strategies and conjectures. See footnotes 6 and 8 for details.

<sup>4</sup> Note that our conjectures describe opponents' actions rather than beliefs about nodes in information sets, as in Kreps and Wilson (1982).

game seriously here, in that we do not allow players to observe actions at information sets that are not reached during the course of play.<sup>5</sup>

Given this information, the agent first updates his conjectures to match the observed frequency of actions at all information sets that were reached during period  $t$ . Updated conjectures about play at information sets that are reached are uniquely specified<sup>6</sup> and are the same for all agents who learn. Conjectures at unreached information sets are unchanged.<sup>7</sup>

Given his new conjecture, the agent updates his behavior strategy. At all information sets where his current action is a best response, his action remains unchanged.<sup>8</sup> At all information sets where his current action is not a best response against his conjecture, the agent changes to an action that is chosen according to some probability distribution (which depends only on the agent's information about the current state) that puts positive probability on all actions that are best responses.

The learning mechanism defines a collection of probabilities of the form  $p_{ij}$  for all  $i$  and  $j$  in  $\Theta$ . These probabilities in turn constitute a Markov process on the state space  $\Theta$ .

## 2.2. Mutations

We now add "mutations" to our evolutionary model. At the end of each time  $t$ , each agent takes another independent draw from a Bernoulli trial. With probability  $1 - \lambda \in (0, 1)$ , this draw produces no change in this agent's characteristic. With probability  $\lambda$ , this draw produces the

<sup>5</sup> Canning (1992a,b) also examines evolutionary processes that are explicitly tailored to the structure of extensive form games.

<sup>6</sup> Here we use our assumptions that every player moves at most once along each path and that learning players observe the outcomes of all matches: Suppose player  $i$ 's information set  $h$  is reached in some match. Whether  $h$  is reached depends only on the choices of player  $i$ 's opponents, so it must be that for all agents from population  $i$  there is a matching that allows their behavior at  $h$  to be observed. All learning agents then observe this behavior and update their conjectures to match the actual distribution of actions by player  $i$  agents at information set  $h$ . The assumption that learning agents observe the current state is strong, but allows a convenient characterization of conjectures. We suspect that our results will hold as long as learning agents have "good enough" information, but have not investigated how good is good enough.

<sup>7</sup> Because agents update conjectures to match the most recently observed play, we take an agent's conjecture about another agent to be a strategy rather than a probability distribution over strategies. Assuming that learning agents observe the most recent play and change their conjectures to match this play allows us to work with a finite space state, which simplifies the analysis.

<sup>8</sup> Again, whether one of the agent's own information sets is reached does not depend on the agent's behavior strategy. Therefore his conjectures and hence action remains unchanged at any information set that was not reached in the previous state. In particular, changes in an agent's actions at an information set  $h$  cannot cause that agent to now potentially reach and hence choose a new action at the previously unreached information set  $h'$ .

outcome “mutate.” In this case, the agent changes to a characteristic that is randomly determined according to a probability distribution that puts positive probability on each of the characteristics that are possible for this agent.

Let  $q_{ij}$  be the probability that mutations change the state of the system from  $i$  to  $j$ . Note that for all  $i$  and  $j$ ,  $q_{ij} > 0$ . Let

$$\gamma_{ij} \equiv \sum_{k \in \Theta} p_{ik} q_{kj}. \quad (1)$$

Then  $\gamma_{ij}$  is the probability that the combination of learning and mutation moves the system from state  $i$  to  $j$ . We are interested in the Markov process given by  $\{\gamma_{ij}\}_{i,j \in \Theta}$ . We let  $\Gamma(G)$  denote this process.

Because  $q_{ij} > 0$  for all  $i$  and  $j$  in  $\Theta$ , we have  $\gamma_{ij} > 0$  for all  $i$  and  $j$  in  $\Theta$ . The following results are standard (cf. Billingsley, 1986; Freidlin and Wentzell, 1984) and are given here without proof:

**LEMMA 1.** *Given  $\lambda$ , the Markov process  $\{\gamma_{ij}\}_{i,j \in \Theta}$  has a unique stationary distribution  $\zeta^*(\lambda)$ , where  $\zeta^*(\lambda)$  is a probability measure on  $\Theta$  such that  $\zeta^*(\lambda)\Gamma(G) = \zeta^*(\lambda)$ . The system converges to  $\zeta^*(\lambda)$  from any initial condition.*

We are interested in the limit of the stationary distribution  $\zeta^*(\lambda)$  as mutations become rare, i.e., as  $\lambda$  becomes small. We refer to this limit as the *limiting distribution*. This limit exists (and hence is unique):

**LEMMA 2.**  *$\lim_{\lambda \rightarrow 0} \zeta^*(\lambda)$  exists.*

The limiting distribution is thus unique and independent of the initial conditions, allowing us to avoid the nonexistence problems of static evolutionary equilibrium concepts. The properties of the model driving this result are that mutations are completely mixed, so that any strategy may be introduced by a mutation; the probability of a mutation occurring does not vary with time; and the probability that a mutation introduces a given strategy (contingent on a mutation occurring) does not change as the probability of a mutation decreases.<sup>9</sup> If these properties hold, then the remaining details of the mutation process are irrelevant.

An alternative source of noise would be for learning agents to observe the outcomes of only some rather than all matches (perhaps meeting

<sup>9</sup> Crawford (1992) argues that time invariance of mutations may not be applicable. In particular, if mutations are taken to represent either players' experimentation or mistakes, then one might expect the incidence of mutation to decrease over time, as players learn the game and hence are less inclined to experiment and less prone to make mistakes. It may be more plausible to think of mutations that do not vary over time if mutations represent new entrants into the game.

only a subset of the agents in the opposing population and observing the outcomes only of their matches. We suspect that the results would be similar to ours if populations of agents are sufficiently large (and the proportion of agents matched in each period remained constant) and the matching scheme is uniform across agents. Quite different results might appear if the matching scheme is not uniform, so that there are neighborhood effects. We might also add "trembling hand" trembles of the kind studied by Selten (1975) to the model.<sup>10</sup> This would cause all information sets to be reached with positive probability and would preclude our constructing proofs that rely on the freedom of actions at unreached information sets to drift. However, we would be interested in the case of small trembles, which we would capture by examining the limit as tremble probabilities approach zero. The results of Samuelson (1991a) suggest that the addition of trembles will affect the results only if (in the limit) they are arbitrarily more likely than mutations.

### 2.3. Limiting Distribution

The basic tool we use to examine the support of the limiting distribution is the concept of an *absorbing set*. Let  $p_{ij}^n$  be the probability that the learning mechanism leads the system from state  $i$  to  $j$  in  $n$  steps. Then:

**DEFINITION 1.** The set of states  $Q \subseteq \Theta$  is absorbing if we have  $p_{ij} = 0$  for all  $i \in Q$  and  $j \notin Q$  and if no proper subset of  $Q$  has this property. The *basin of attraction* of an absorbing set  $Q$  is the set  $B(Q) = \{\theta : \exists n, \exists \theta' \in Q \text{ s.t. } p_{\theta\theta'}^n > 0\}$ .

An absorbing set is then a minimal set with the property that the *learning mechanism* cannot take the system out of this set (although mutations may still move the system out of an absorbing set). Note that an absorbing set may contain more than one state. The basin of attraction of an absorbing set is the collection of states from which there is a positive probability that the learning scheme leads to the absorbing set.

The following lemma is immediate from Samuelson (1991a) (as are Lemmas 4 and 5 below) and is given here without proof:

**LEMMA 3.** *The support of the limiting distribution consists only of elements of absorbing sets. If a state  $\theta$  appears in the limiting distribution, then all states in the absorbing set containing  $\theta$  are in the support of the limiting distribution.*

<sup>10</sup> Canning (1992a,b) examines evolutionary models with trembling hand trembles but without mutations.

The intuitive interpretation of Lemma 3 is that the selection mechanism moves the system more easily into absorbing sets than out of them, and hence in the limit (as  $\lambda$  becomes small) is concentrated on absorbing sets.

Not all absorbing sets will appear in the limiting distribution. Young (1993), Kandori *et al.* (1993), and Samuelson (1991a) show that the absorbing sets appearing in the limit are those that are easiest to reach from all other absorbing sets, in the sense that it takes the fewest mutations to get to their basins of attraction from other absorbing sets. For example, Young (1993) and Kandori *et al.* (1993) show that in a  $2 \times 2$  symmetric normal form game with two strict Nash equilibria (denoted  $A$  and  $B$ ), the limiting distribution will consist entirely of equilibrium  $A$  if it takes fewer mutations to transform the system from equilibrium  $B$  to  $A$  than from  $A$  to  $B$ .

This characterization of the limiting distribution in terms of the number of mutations required to move between absorbing sets may demand the comparison of very large numbers of very unlikely mutations. It is important to note that our results are obtained by asking whether absorbing sets are robust against a *single* mutation. We think that such results are the most robust to emerge from the model.

**DEFINITION 2.** States  $\theta$  and  $\theta'$  are adjacent if one mutation can change the state from  $\theta$  to  $\theta'$  (and hence from  $\theta'$  to  $\theta$ ). The *single-mutation neighborhood*,  $M(Q)$ , of an absorbing set  $Q$  is the set of all  $\theta'$  that are adjacent to some  $\theta \in Q$ .

Our next lemma states that if the basin of attraction of an absorbing set  $Q'$  is only "one mutation away" from an absorbing set  $Q$  that appears in the limiting distribution, then  $Q'$  is also contained in the limiting distribution.

**LEMMA 4.** Suppose absorbing set  $Q$  is contained in the support of the limiting distribution and absorbing set  $Q'$  satisfies

$$M(Q) \cap B(Q') \neq \emptyset. \quad (2)$$

Then  $Q'$  is also contained in the support of the limiting distribution.

We can use this result to identify collections of absorbing sets that have the property that either all or none of them appear in the support of the limiting distribution.

**DEFINITION 3.** A collection of absorbing sets  $\Phi$  is a cycle if for all  $Q, Q' \in \Phi$  there exists absorbing sets  $Q_1, \dots, Q_n \in \Phi$  such that  $Q = Q_1$ ,  $Q' = Q_n$  and for  $i = 1, \dots, n - 1$

$$M(Q_i) \cap B(Q_{i+1}) \neq \emptyset.$$

A cycle  $\Phi$  is a *component* if there exists no cycle  $\Phi' \neq \Phi$  containing  $\Phi$ . A component  $\Phi$  is *locally stable* if there is no pair of absorbing sets  $Q \in \Phi$  and  $Q' \notin \Phi$  for which condition 2 holds.

Note that the set of components is a partition of the sets of absorbing sets. Throughout the following we use  $C$  to denote components and write  $C(Q)$  [or  $C(\theta)$ ] to denote the unique component containing absorbing set  $Q$  (or state  $\theta$ ). A component is locally stable if more than one mutation is required to go from  $C$  to *any* other component. Note that we could also define a locally stable component as a minimal collection of absorbing sets having this stability property.<sup>11</sup>

*Remark 1.* We find it convenient to speak of the single mutation neighborhood and basin of attraction of a component (defined in the obvious way as the union of the corresponding sets over the absorbing sets in the component). We denote the case when the single mutation neighborhood of a component  $C$  and the basin of attraction of a component  $C'$  intersect (i.e., there exists  $Q \in C$  and  $Q' \in C'$  such that condition 2 holds) by writing  $\sigma^*(C, C') = 1$ . If they do not intersect, we write  $\sigma^*(C, C') > 1$ .

From Lemma 4 it is easy to see that either all or none of the states appearing in a component are in the support of the limiting distribution. The following lemma states that a sufficient condition for a component *not* to appear in the limiting distribution is that it is only "one mutation away" from a locally stable component.

**LEMMA 5.** *Suppose  $C'$  is a locally stable component and let  $Q$  be an absorbing set not contained in  $C'$ . If there exists an absorbing set  $Q'$  in  $C'$  such that condition 2 holds, then the states in  $C(Q)$  are not contained in the support of the limiting distribution.*

Combining Lemmas 4 and 5 we can now prove the following result, which provides our basic tool for examining the limiting distribution.

**PROPOSITION 1.** *State  $\theta$  is in the support of the limiting distribution only if  $\theta$  is contained in a locally stable component. If  $\theta$  is in the limiting distribution, so are all states in the locally stable component  $C(\theta)$ .*

*Proof.* First we show that a locally stable component must exist. Consider a component, denoted by  $C_1$ , and suppose  $C_1$  is not locally stable. Then there must exist a component  $C_2$  such that  $\sigma^*(C_1, C_2) = 1$ . It is

<sup>11</sup> Our definition of local stability allows for the possibility that, starting in the single mutation neighborhood of a locally stable component  $C$ , the learning dynamics initially lead away from this component. It is only required that learning must ultimately lead back to some state in  $C$ . Because our results are established by finding situations where a *single* mutation suffices to destabilize a component, they would continue to hold if we defined local stability in terms of larger numbers of mutations.

easy to verify that  $\sigma^*(C_2, C_1) > 1$  must hold, since otherwise the union of the absorbing sets contained in  $C_1$  and  $C_2$  would form a cycle, contradicting the fact that  $C_1$  is a component. Hence, either  $C_2$  is locally stable or there exists a component  $C_3 \neq C_1$  such that  $\sigma^*(C_2, C_3) = 1$ . As before, since a component cannot be contained in a cycle, we must have  $\sigma^*(C_3, C_1) > 1$  and  $\sigma^*(C_3, C_2) > 1$ , implying that either  $C_3$  is locally stable or that we can continue the argument to find a new component  $C_4$  such that  $M(C_4)$  does not intersect the basins of attraction of  $C_1$ ,  $C_2$ , and  $C_3$ . Since the state space of our process is finite, there can only be a finite number of components. Continuing this construction must thus ultimately lead to a locally stable component. Next, it is immediate from Lemma 4 that either all or none of the states appearing in a component are contained in the limiting distribution. Now suppose a component, again denoted by  $C_1$ , is contained in the support of the limiting distribution. If  $C_1$  is not locally stable, then we can use the same construction as above to find a sequence  $C_1, \dots, C_k$  such that  $C_k$  is locally stable and  $\sigma^*(C_i, C_{i+1}) = 1$ . But then Lemma 4 implies that  $C_{k-1}$  must be in the support of the limiting distribution, whereas Lemma 5 implies that  $C_{k-1}$  cannot be in the support of the limiting distribution. This contradiction proves that  $C_1$  must be locally stable, proving the proposition. ■

If there is a unique locally stable component, Proposition 1 implies that the support of the limiting distribution coincides with the states appearing in that component. If there are multiple locally stable components and one wishes to determine which of these components constitute the support of the limiting distribution, then one must resort to the more general mutation-counting arguments of the type employed by Young (1993) and Kandori *et al.* (1993).

We are especially interested in cases in which the limit consists entirely of singleton absorbing sets because the latter exhibit considerable structure. First:

**DEFINITION 4.** A state is a *self-confirming equilibrium* if each agent's strategy is a best response to that agent's conjecture and if each agent's conjecture about opponents' strategies matches the opponents' choices at information sets that are reached in the play of some matches.

This is similar to Fudenberg and Levine's more general notion of a self-confirming equilibrium (1991), except that we find it convenient to define this concept directly in the state space of the Markov process rather than in the strategy space of the game.<sup>12</sup>

<sup>12</sup> See Kalai and Lehrer (1991) for a similar equilibrium concept. There may be self-confirming equilibria in the game  $G$  that we cannot support as self-confirming equilibria in the state space of our Markov process because the former require mixed strategies that

It then follows immediately from the definitions that:

*Remark 2.* For  $\theta \in \Theta$ ,  $\{\theta\}$  is a singleton absorbing set if and only if  $\theta$  corresponds to a self-confirming equilibrium.

Note that we get equivalence between singleton absorbing sets and self-confirming equilibria rather than Nash equilibria.<sup>13</sup>

When discussing singleton absorbing sets, we often let  $\theta$  denote both an element in the state space and the singleton absorbing set containing that element. When discussing a component  $C$ , we also find it helpful to abuse notation by writing  $\theta \in C$  for any state  $\theta$  contained in any absorbing set contained in  $C$ .

Consider a self-confirming equilibrium. The component containing this state contains states corresponding to the same path through the extensive form game but a variety of out-of-equilibrium behavior and out-of-equilibrium conjectures. Some additional terminology is helpful in making this precise. Given a self-confirming equilibrium outcome, we say that *player  $i$  can force entry into a subgame  $G(a_i)$*  if there exists an information set  $h$  of player  $i$  that is reached during the play of the self-confirming equilibrium and an action  $a_i$  for player  $i$  at  $h$  that is not chosen by any player  $i$  agent and such that the decision node resulting from the choice of  $a_i$  is the root of a subgame of  $G$ . This subgame is denoted by  $G(a_i)$ . Then (the proof is in the Appendix):

**PROPOSITION 2.** (2.1) *Let  $\{\theta\}$  be a singleton absorbing set. Let  $\theta'$  differ from  $\theta$  only in actions prescribed at information sets that are not reached in any matchings in state  $\theta$  (i.e., every agent in  $\theta$  can be paired with a distinct agent in  $\theta'$  such that the two agents hold identical conjectures (at all information sets) and play identical actions at all information sets reached in any matching under  $\theta$ ). Then  $\{\theta'\}$  is also a singleton absorbing set and  $\theta' \in C(\theta)$ .*

(2.2) *Suppose that, given  $\theta'$ , player  $i$  can force entry into a subgame  $G(a_i)$ . If  $\theta''$  differs from  $\theta'$  only in the conjectures that agents from populations other than  $i$  hold over choices at information sets in  $G(a_i)$ , then  $\theta'' \in C(\theta)$ .*

---

cannot be duplicated as population proportions in our finite population. A similar difficulty arises when we consider subgame perfect equilibria. In the following analysis, we do not distinguish between equilibria in the game  $G$  and in the state space of our Markov process. Hence, we implicitly assume that the population size is such that we can achieve any desired equilibrium on every subgame in the state space. Similar results can be obtained for pure strategy versions of the equilibrium concepts without this assumption. It would be useful to extend the model in order to address the question of when population proportions are "close enough" to effectively allow mixed strategy equilibria to be achieved.

<sup>13</sup> Nöldeke and Samuelson (1992) provide an example in which a self-confirming equilibrium that is not a Nash equilibrium appears in the limiting distribution of the evolutionary process.

The implication of this result is that the evolutionary process is not effective at imposing discipline at unreached information sets. If a self-confirming equilibrium appears in the support of the limiting distribution, so does any state that can be reached by allowing actions at unreached information sets to drift.

### 3. LOCALLY STABLE OUTCOMES

This section examines locally stable components of singleton absorbing sets. From Proposition 1, we know that the limiting distribution contains only locally stable components and contains all of the absorbing sets in any component that contributes to the limiting distribution. Our only hope for single-valued equilibrium predictions then lies with locally stable components of absorbing sets that all correspond to the same self-confirming equilibrium outcome. In this case, we call the outcome involved a *locally stable outcome*. More formally:

**DEFINITION 5.** An outcome  $z^*$  is *locally stable* if there exists a locally stable component  $C$  such that for all  $\theta \in C$ ,  $z(\theta) = z^*$ .

Studying such outcomes allows us to characterize the limiting distribution when the latter is “nicely behaved” and also provides some clue as to when it will be so behaved.

Locally stable outcomes are analogs of static refinements of evolutionary stability concepts. In each case, one couples an equilibrium condition with a stability requirement and shows that the resulting outcome possess desirable properties. In each case, however, existence and uniqueness is problematic. In the following sections, we turn to dynamic considerations because locally stable outcomes may not exist and may not be unique, and even the existence of a unique locally stable outcome may not ensure that it is contained in the support of the limiting distribution.<sup>14</sup>

We first establish a backward induction property for locally stable outcomes:

**PROPOSITION 3.** *Let  $G$  be an extensive form game with each player moving at most once along each path. Then every locally stable outcome is a subgame perfect equilibrium outcome.*

<sup>14</sup> It is possible to construct games that have two locally stable components, one corresponding to a locally stable outcome and one containing a nonsingleton-absorbing set. In such a case, the support of the limiting distribution may consist solely of the states in the second component. The existence of a locally stable outcome thus does not suffice to imply that the evolutionary process results in convergence to equilibrium behavior.

*Proof.* Let  $z^*$  be a locally stable outcome generated by the locally stable component  $C$ .  $C$  contains only singleton-absorbing sets.<sup>15</sup> Suppose  $z^*$  does not correspond to the subgame perfect equilibrium outcome. Then every state  $\theta \in C$  satisfies the condition that there exists a player  $i$  who can force entry into a subgame  $G(a_i)$  such that every subgame perfect equilibrium of  $G(a_i)$  yields  $i$  a higher payoff than does  $z^*$ . Proposition 2 ensures that the component  $C$  contains a state  $\theta$  in which the actions and conjectures of all players who have an information set in  $G(a_i)$  correspond to a subgame perfect equilibrium of  $G(a_i)$ .<sup>16</sup> Now consider a mutation in the actions of a player  $i$  agent that causes  $G(a_i)$  to be reached. Then with positive probability, all player  $i$  agents learn. These agents will switch their actions so as to enter  $G(a_i)$ . The learning mechanism then cannot further adjust actions or conjectures in  $G(a_i)$ , since these constitute a self-confirming equilibrium on  $G(a_i)$ . This in turn ensures that the learning mechanism leads to an absorbing set, say  $Q$ , that contains at least one state not resulting in the outcome  $z$ . We then have  $Q \notin C$  and  $M(\theta) \cap B(Q) \neq \emptyset$ , contradicting the local stability of  $C$ .

We now turn to forward induction properties. Proposition 3 showed that a self-confirming but *not* subgame perfect equilibrium will fail to be locally stable because there must be a player who can force entry into a subgame with a subgame perfect equilibrium that makes the player better off than the original equilibrium. The following proposition extends this argument to show that *subgame perfect* equilibria also may fail to be locally stable. This occurs because states supporting the subgame perfect equilibrium outcome again allow a variety of behavior off the equilibrium path, including behavior that, once revealed, will tempt agents away from the subgame perfect equilibrium in quest of a higher payoff. This requires only that an agent has the ability to cause a subgame that has a self-confirming equilibrium promising the player a higher payoff than the subgame perfect equilibrium to be reached. It may be impossible to support this higher payoff as an equilibrium in the entire game, so that subsequent adjustments must lead away from this payoff, but these adjustments cannot lead back to the original equilibrium. This potential instability of subgame perfect equilibria provides the basis for our discussion of forward induction.

**PROPOSITION 4.** *Let  $G$  be an extensive form game with each player moving at most once along each path. Suppose outcome  $z^*$  is locally stable. If, given the locally stable component supporting  $z^*$ , player  $i$  can*

<sup>15</sup> It is easy to see that every nonsingleton-absorbing set must contain states  $\theta, \theta'$ , such that  $z(\theta) \neq z(\theta')$ .

<sup>16</sup> Here we use our assumption that each player moves at most once along every path, which implies that player  $i$  has no information set in  $G(a_i)$ . Proposition 2 then implies that we can choose the desired conjectures for all players having an information set in  $G(a_i)$ .

force entry into a subgame  $G(a_i)$ , then no self-confirming equilibrium of  $G(a_i)$  can give player  $i$  a higher payoff than does  $z^*$ .

*Proof.* This follows by the same argument as the proof of Proposition 3. The only additional observation required is that if the actions and conjectures of the agents who have information sets in  $G(a_i)$  agree with a self-confirming equilibrium on this game, then once player  $i$  enters the subgame  $G(a_i)$ , the learning process cannot cause any adjustments on this subgame. ■

The result in Proposition 4 associates a strong forward induction property with locally stable outcomes. In many games this property will be too stringent to be satisfied by any of the subgame perfect equilibrium outcomes. Proposition 4 hence identifies one possible source for the nonexistence of locally stable outcomes (Example 2 in the next section illustrates another possibility for the nonexistence of locally stable outcomes).

In two-player games (with each player moving at most once along any path), these results can be sharpened. We think of a game  $G$  as proceeding by having player 1 first choose which of player 2's information sets  $h$  in the game  $G$  to reach, at which point a normal form game is played with strategy sets given by 2's behavior strategies (in  $G$ ) at  $h$  and 1's strategies (in  $G$ ) that cause  $h$  to be reached. We refer to this latter representation of the game as the *extended form of  $G$* . It is a straightforward variation on Proposition 3 that any locally stable outcome of  $G$  must correspond to a subgame perfect equilibrium outcome of the extended form of  $G$ .<sup>17</sup>

In addition, for generic games of this class, Proposition 4 implies that a locally stable outcome of  $G$  corresponds to a subgame perfect equilibrium outcome of the extended form satisfying the Never-Weak-Best-Response (NWBR) property.<sup>18</sup> This follows from the fact that in generic games from this class, a subgame perfect equilibrium outcome  $s^*$  fails the NWBR property only if there is an unreached subgame with a Nash equilibrium giving player 1 a higher payoff than  $s^*$ .

We illustrate the forward induction properties of locally stable outcomes in two-player games with the following example.

EXAMPLE 1. Consider the game whose extended form is given by

	$L_1$	$R_1$		$L_2$	$R_2$
$T_1$	4, 1	-4, 0	$T_2$	3, 3	-5, -5
$B_1$	0, 0	2, 1	$B_2$	-5, -5	1, 1

<sup>17</sup> We require here that the population size is such that any (possibly mixed) Nash equilibrium in any subgame of the extended form of  $G$  can be achieved by appropriate population proportions.

<sup>18</sup> A subgame perfect equilibrium outcome  $s^*$  of the extended form satisfies NWBR if it is also a subgame perfect equilibrium outcome of the extended form obtained by removing strategies that fail to be best responses to any elements of the Nash component containing  $s^*$ . See Kohlberg and Mertens (1986).

where player 1 first chooses one of the matrices, and then players 1 and 2 play the game represented by that matrix, with player 1 choosing rows and player 2 choosing columns. There are no dominated strategies in this game, so that iterated weak dominance fails to eliminate any strategies. NWBR eliminates the subgame perfect equilibrium  $(T_2, L_2)$ . In particular,  $B_1$  is a best response to no element in the component supporting outcome  $(T_2, L_2)$ . Once this is eliminated, subgame perfection requires player 2 to choose  $L_1$  at his first information set, at which point  $(T_2, L_2)$  ceases to be an equilibrium. Note that NWBR cannot eliminate the subgame perfect equilibrium given by  $(B_2, R_2)$ , since every strategy is a best reply to some element in the component supporting this outcome.

In contrast, the only locally stable outcome in this game is  $(T_1, L_1)$ . To see this, consider a singleton-absorbing set  $\theta$  supporting the outcome  $(B_2, R_2)$  or  $(T_2, L_2)$ . By Proposition 2, the component  $C(\theta)$  contains a state in which all player 2 agents choose  $L_1$  at their first information set. Now let a mutation that causes a player 1 agent to play  $T_1$  occur. Then with positive probability, all player 1 agents learn and switch to  $T_1$ . This yields a new singleton-absorbing set  $\theta'$  supporting outcome  $(T_1, L_1)$ . Since the latter equilibrium gives player 1 the highest possible payoff in the game and action  $L_1$  is a strict best reply for player 2 at the relevant information set, the component  $C(\theta')$  contains only states yielding the outcome  $(T_1, L_1)$  and is easily seen to be locally stable for all sufficiently large population sizes.  $(T_1, L_1)$  is the *unique* locally stable outcome, since an argument similar to the one just given shows that neither the outcome  $(B_1, R_1)$  nor a mixed strategy equilibrium outcome is locally stable.

#### 4. BACKWARD INDUCTION

Even if a game admits a unique subgame perfect equilibrium  $s^*$ , the results of the previous section do not allow us to conclude that  $s^*$  appears in the limiting distribution, much less that  $s^*$  is the only outcome in that distribution. The only possible locally stable outcome is  $s^*$ , but a locally stable outcome may fail to exist or may not comprise the entire limiting distribution. To address these issues we must go beyond the study of locally stable outcomes to examine the dynamic process. The complexity of this task forces us to consider a restricted class of games. In this section we consider extensive form games with each player moving at most once along each path, as before, but now restrict attention to generic games of perfect information. We find that the subgame perfect equilibrium must appear in the limiting distribution, but only under stringent conditions will it be the only outcome in the limiting distribution.<sup>19</sup>

<sup>19</sup> These games are dominance solvable, so that the subgame perfect equilibrium is the outcome of iterated elimination of weakly dominated strategies in the normal form. Our

Our genericity condition is that for every player  $i \in I$ ,  $z \neq z' \in Z \Rightarrow \pi_i(z) \neq \pi_i(z')$ . It is well known that such games have a unique subgame perfect equilibrium.<sup>20</sup>

We begin by establishing that such games have only singleton absorbing sets, which in turn implies that the evolutionary process converges to a locally stable component containing only states that yield self-confirming equilibrium outcomes. The Appendix proves (all subsequent proofs are in the Appendix):<sup>21</sup>

**PROPOSITION 5.** *Let  $G$  be a generic extensive form game of perfect information with each player moving at most once along each path. Then all absorbing sets are singletons.*

Next, we show that there is a unique locally stable component that contains the subgame perfect equilibrium outcome. From Proposition 1 this implies that the limiting distribution satisfies a backward induction property in the sense that it must assign strictly positive probability to the subgame perfect equilibrium. Furthermore, if the subgame perfect equilibrium outcome is locally stable, it must be the unique outcome appearing in the support of the limiting distribution. The proof of this result, which combines ideas from the proof of Propositions 1 and 3, exploits the fact that absorbing sets are singletons.

**PROPOSITION 6.** *Let  $G$  be a generic extensive form game of perfect information with each player moving at most once along each path. Then there is a unique locally stable component containing the subgame perfect equilibrium outcome.*

Proposition 4 showed that the subgame perfect equilibrium outcome will fail to be locally stable if some player can be tempted away from the subgame perfect equilibrium path by a self-confirming equilibrium that promises that player a higher payoff.<sup>22</sup> The following example shows that the necessary conditions for a locally stable outcome established in

---

finding that the subgame perfect equilibrium will often not be the only outcome in the limiting distribution is then related to Samuelson's (1991a) finding that a similar model does not in general eliminate weakly dominated strategies in the normal form.

<sup>20</sup> Note that the unique subgame perfect equilibrium will always be in pure strategies, so that the question of whether mixed strategies can be duplicated by population proportions does not arise.

<sup>21</sup> Canning (1992b) shows that all the absorbing sets for a fictitious play process in extensive form games of perfect information are singletons. In a model with trembling hand trembles of the type introduced by Selten (1975), but without mutations, Canning finds that the subgame perfect equilibrium is the unique limiting outcome (as trembles become small).

<sup>22</sup> It is easy to establish that for all two-player games (from the class of games under consideration) the subgame perfect equilibrium outcome is locally stable for every sufficiently large population size and is thus the only outcome in the limiting distribution of our evolutionary process.

Proposition 4 are not sufficient to imply that the subgame perfect equilibrium outcome is locally stable.

EXAMPLE 2. Consider the game shown in Fig. 1.

The unique subgame perfect equilibrium is given by the strategy combination  $(R, R, R)$ , for a payoff of  $(1, 1, 1)$ . Let  $\theta_R$  be the state in which all agents from all populations play  $R$  and conjectures match these actions. Let  $\theta_L$  be a state in which all agents in populations 1 and 2 play  $L$  and agents from population 3 play  $R$ , with conjectures matching these actions. Note that  $\theta_L$  is a self-confirming equilibrium and gives payoffs  $(0, 0, 0)$ . We now show that  $\theta_L$  must be contained in the support of the limiting distribution.

By Proposition 6,  $\theta_R$  is in the support of the limiting distribution and so is (by Proposition 2) the state in which all agents from populations 1 and 2 play  $R$  and all agents from population 3 play  $L$ , with conjectures matching the behavior of agents at reached decision nodes and agents from population 1 and 2 conjecturing that all agents from population 3 play  $R$ . Now let a single mutation cause an agent from population 2 to play  $L$ . With positive probability, all agents from population 2 (but no other agents) now receive the learn draw, update their conjecture to match the observation that all agents from population 3 play  $L$ , and switch to playing a best reply of  $L$ . Suppose in the next stage of the evolutionary process all agents from population 1 (but no other agents) learn. They update their conjectures to match the observed behavior of other agents,

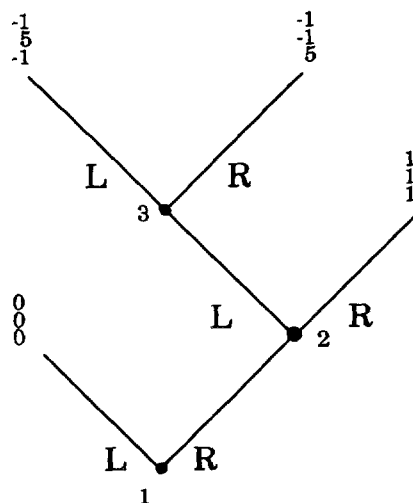


FIG. 1.

viz. that agents from population 2 are playing  $L$ , and then switch to their best reply of playing  $L$ . It is easily checked that the resulting state is a singleton-absorbing set yielding the same outcome as  $\theta_L$ . Using Proposition 2 again, it thus follows that  $\theta_L \in C(\theta_R)$  and consequently this self-confirming equilibrium must appear in the support of the limiting distribution.

This example allows us to illustrate the differences between our outcomes and Swinkels' (1992b) equilibrium evolutionarily stable sets. It is easy to show that in the game of Example 2, the (unique) EES set and our limiting distribution both contain a set of strategies corresponding to the payoff  $(1, 1, 1)$ .<sup>23</sup> In our case, this set contains states corresponding to any possible mixture on the part of player 3, including mixtures that violate the conditions for a Nash equilibrium. It is the presence of these states that causes the subgame perfect equilibrium outcome to fail local stability, because single mutations can lead from these states to learning dynamics that draw the system away from the subgame perfect outcome. The corresponding EES set allows only Nash equilibria and hence includes mixtures on the part of player 3 only if less than one third probability is placed on  $L$ . The dynamic intuition behind the EES set is that deviations from Nash equilibrium, even at unreached subgames, prompt immediate responses that push the system back to the EES set, so that the system is not free to drift away from Nash equilibria via movements at unreached information sets.

Proposition 4 and Example 2 suggest that the subgame perfect equilibrium outcome will correspond to the unique outcome in the limiting distribution of the evolutionary process if one precludes the possibility that some player is "tempted" away from the subgame perfect equilibrium by a (possibly nonequilibrium) higher payoff. To make this intuition precise, we introduce a bit more terminology. We say that the subgame perfect equilibrium outcome  $z^*$  is a *strict outcome* (cf. Balkenborg, 1992) if no player can force entry into a subgame that has a terminal node with a higher payoff than the subgame perfect equilibrium, or more formally: Suppose that given the subgame perfect equilibrium outcome, player  $i$  can force entry into the subgame  $G(a_i)$  and let  $Z(a_i)$  denote the set of terminal nodes contained in  $G(a_i)$ . For  $z^*$  to be a strict outcome

$$\forall z \in Z(a_i): \pi_i(z^*) > \pi_i(z)$$

<sup>23</sup> Thomas (1985a,b) offers the set-valued concept of an evolutionarily stable (ES) set. This differs from Swinkels' EES set in that potential entrants are not required to be best responses to their immediate postentry environment. As a result, ES sets require robustness against more entrants than do EES sets, causing the former to be subsets of the latter. For reasons similar to those causing our limiting distribution to contain more states than the EES set in Example 2, there is no ES in this game.

must hold. We then have:<sup>24</sup>

**PROPOSITION 7.** *Let the game be as in Proposition 6. If the subgame perfect equilibrium outcome is strict, it is the unique outcome in the limiting distribution for all sufficiently large population sizes.*

Our results provide both good and bad news for backward induction. Subgame perfection presents the only possibility for a locally stable outcome. In many games, however, the limiting distribution contains states that do not yield the subgame perfect outcome path. These results suggest a reinterpretation of the role of backward induction in finite extensive form games of perfect information, with subgame perfection being a sufficient but not necessary conditions for an outcome to be interesting.

## 5. FORWARD INDUCTION

In the spirit of the previous section, we examine a class of games simple enough to investigate whether the forward induction properties of locally stable components carry over to characterize the limiting distributions of our evolutionary process. We consider two player games whose extended form calls for player 1 to first choose between an outside option, yielding payoff vector  $x$ , and entering a normal form game  $\tilde{G}$ .

The analysis of these games is complicated by the fact that they may have nonsingleton-absorbing sets. We can make some progress in spite of this, but the following assumption allows us to strengthen the results.

**Assumption 1.** (1.1) If a player  $i$  agent, called agent  $A$ , receives the learn draw and (after updating conjectures) finds that he is not playing a best reply, but other player  $i$  agents are playing best replies, then agent  $A$  switches to a best reply played by one of the other player  $i$  agents.

(1.2) No Nash equilibrium of game  $\tilde{G}$  gives player 1 the same payoff as the outside option.

Assumption (1.1) introduces some inertia into the learning process, indicating that when switching strategies, agents will choose best replies that are already played by other members of their population if such strategies exist. This assumption is likely to be reasonable in contexts where learning is at least partly driven by imitation. Assumption (1.2) is a mild genericity assumption.

<sup>24</sup> Note that Propositions 4 and 7 are not converses of one another, since Proposition 7 requires that no terminal node in  $G(a_i)$  give player  $i$  a higher payoff than  $\pi_i(z^*)$ , while Proposition 4 requires that a self-confirming equilibrium outcome in  $G(a_i)$  give a higher payoff than  $\pi_i(z^*)$ .

**PROPOSITION 8.** *Suppose that there exists a strict Nash equilibrium of the game  $\tilde{G}$  providing player 1 with a payoff higher than  $x_1$ .*

*(8.1) Then for all sufficiently large  $\Delta$  the limiting distribution does not contain a state in which all player 1 agents choose the outside option.*

*(8.2) If Assumption 1 holds, then for all sufficiently large  $\Delta$ , the limiting distribution does not contain a state in which any player 1 agent chooses the outside option.*

Van Damme (1987b, 1989) formulates a notion of forward induction for generic extensive form games. For generic games from our class, van Damme's formulation takes the following form: if one Nash equilibrium in the game  $\tilde{G}$  provides player 1 with a payoff higher than  $x_1$  and all other Nash equilibria of  $\tilde{G}$  provide player 1 a lower payoff, then any equilibrium in which player 1 plays the outside option fails forward induction. The idea here is that if  $\tilde{G}$  contains a unique equilibrium offering a higher payoff to player 1 than  $x_1$ , then player 1's act of entering  $\tilde{G}$  provides an unambiguous signal that 1 expects this equilibrium to appear in  $\tilde{G}$  and will play 1's part of this equilibrium. Player 2 then finds it optimal to play 2's part of the equilibrium. This makes the equilibrium available to player 1, so that 1 will not choose the outside option.

Proposition 8 yields a forward induction notion that differs in two important respects from van Damme's. First, we do not require that the equilibrium in the game  $G$  providing  $i$  with a higher payoff than adhering to the outside option be unique.<sup>25</sup> Van Damme requires uniqueness to ensure that the signal provided by entry into the subgame will be unambiguous. Ambiguity is not a difficulty in the evolutionary approach, as evolution will assign an unambiguous meaning to the action of entering the subgame.

Second, although the existence of a Nash equilibrium in the subgame exhibiting certain properties appears as a condition of the proposition, our result does not require an assumption of equilibrium play in the subgame and hence requires no assumption of equilibrium play after nonequilibrium actions. In particular, while our result establishes that the outside option cannot appear in the limiting distribution, it does *not* follow that the limiting distribution must correspond to an equilibrium of  $\tilde{G}$ . We cannot rule out the possibility that the limiting distribution corresponds to a locally stable component containing some nonsingleton-absorbing set (that contains only states in which player 1 always enters  $\tilde{G}$  and receives a higher payoff than the outside option).

We now turn to an example (where we do not require Assumption 1) that allows us to explore the relationship between the outside option and

<sup>25</sup> Our requirement that this equilibrium be strict plays a role similar to van Damme's assumption that the game is generic.

the equilibrium that is selected in game  $\tilde{G}$ . Unlike our earlier results, this example exploits arguments based on comparisons of large numbers of mutations. These are required to ascertain which of several locally stable outcomes will appear in the limiting distribution.

EXAMPLE 3. Let the game  $\tilde{G}$  be given by, where letters represent payoffs:

	L	R
T	a,a	b,c
B	c,d	d,d

We assume that  $d > a > b > c$  with  $(a - c) > (d - b)$ . This ensures that  $(T, L)$  and  $(B, R)$  are pure strategy Nash equilibria, with  $(B, R)$  being Pareto dominant and  $(T, L)$  being risk dominant, and with both yielding a higher payoff than the mixed strategy equilibrium payoff (denoted  $m$ ).

Now consider the game  $G$  in which player 1 chooses between an outside option  $x$ , giving payoff  $x$  to player 1, and game  $\tilde{G}$ . The outcome here will depend upon the relative magnitudes of  $x$  and the payoffs in the subgame. We assume that the population size is large enough that any strict Nash equilibrium is locally stable.

Two cases have been addressed by our previous results. If  $x > d$ , then there is only one component, which is composed of singleton-absorbing sets and gives outcome  $x$ . All states in the support of the limiting distribution thus yield outcome  $x$ . If  $x < m$ , then there are three components, corresponding to the three equilibria in  $\tilde{G}$ . The component consisting of the mixed strategy equilibrium fails to be locally stable, whereas the other two components (for sufficiently large population size) yield locally stable outcomes. An argument analogous to that of Kandori *et al.* (1993) (extending their analysis to our state space and learning mechanism) shows that the limiting distribution will consist entirely of the risk dominant equilibrium  $(T, L)$  in this case.

If  $m < x < d$ , we must examine the dynamics of this model to derive the support of the limiting distribution:

PROPOSITION 9. Let  $m < x < d$  and

$$x^* = \frac{bd - ca}{(d - c) - (a - b)}. \quad (3)$$

Then for  $\Delta$  sufficiently large, the limiting distribution yields the following locally stable outcomes as a function of  $x$ :

<i>Case 1</i>	$m < x < x^*$	$(T, L)$
<i>Case 2</i>	$x^* < x < a$	$(B, R)$
<i>Case 3</i>	$a < x < d$	$(B, R)$

In Case 1, the outside option is so unprofitable as to be irrelevant. In Case 3, the outside option eliminates  $a$  but preserves  $d$  as an equilibrium outcome. There are then two components, corresponding to the outcomes  $x$  and  $(B, R)$ . The latter is locally stable but the former is not, since it takes only one mutation to reach the basin of attraction of outcome  $(B, R)$  from the component supporting  $x$ . The limiting distribution then selects only  $(B, R)$ .

Case 2 is more intriguing. In the absence of the outside option, the limit consists entirely of  $(T, L)$ , while the addition of the outside option causes the limit to be entirely  $(B, R)$ . The interesting aspect of this result is that the outside option disrupts neither of the pure strategy Nash equilibria of the game  $\bar{G}$  (since  $x < a < d$ ) and never appears in the limiting distribution, but still affects that distribution. To see why this result appears, note that in the absence of the outside option, the outcome is  $(T, L)$  because it takes fewer mutations to move the system from  $(B, R)$  to the basin of attraction of  $(T, L)$  than to accomplish the reverse movement. With the addition of the outside option, the relevant calculation now concerns how many mutations it takes to move the system first from  $(B, R)$  to the outside option and then from the outside option to  $(T, L)$ . If the outside option is sufficiently attractive, this can now require more mutations than the reverse path, yielding  $(B, R)$ .<sup>26</sup> It is easy to show that  $x^* < a$ , ensuring that the proposition is not vacuous.

## 6. CONCLUSION

We have examined the evolutionary foundations of backward and forward induction, the two basic building blocks of equilibrium refinements.

We find that our evolutionary model offers mixed support for backward induction. We have examined the class of games in which backward induction is most likely to “work”: finite extensive form games of perfect information in which each player moves at most once along any equilibrium path. We find that the limiting distribution of the evolutionary process will include the unique subgame perfect equilibrium of such a game and that this equilibrium presents the only possibility for a unique limiting

<sup>26</sup> This occurs because  $(B, R)$  is Pareto dominant, making it easier in terms of mutations to move from  $(T, L)$  to the outside option than from  $(B, R)$  to the outside option.

distribution. However, in the absence of strong conditions, the limiting distribution will also include self-confirming equilibria that are not subgame perfect.

Self-confirming equilibria that are not subgame perfect can appear in the limiting distribution because the evolutionary system can be tempted away from the subgame perfect equilibrium outcome by the lure of higher payoffs. In games with multiple subgame perfect equilibria, these same forces, by allowing the system to be drawn away from some of the equilibria, can produce forward induction properties. In outside option games, we find a forward induction property that strengthens van Damme's (1987b, 1989) notion of forward induction.

We conclude that if one embraces the evolutionary approach to games, then backward induction is likely to play a somewhat different role than is often the case, being sufficient but not necessary to identify an outcome as interesting, while forward induction is likely to play a larger role than is often the case.

Our analysis can be extended to examine communication in games where messages do not have exogenously given or intrinsic meanings. The inferences attached to out-of-equilibrium moves in our forward induction games arise endogenously as part of the evolutionary process. In a similar way, meaning can be attached to unused signals in cheap talk and signaling games. This allows these unused messages to destabilize some equilibria, thus providing evolutionary foundations for the arguments that appear in the equilibrium refinements literature.<sup>27</sup> In Nöldeke and Samuelson (1992), we develop this argument for "small" cheap talk and signaling games. Further work is required to extend the analysis to more general games and to refine the connection between our limiting distribution and equilibrium refinements.

## APPENDIX: PROOFS

*Proof of Proposition 2.* Let  $\theta$  and  $\theta'$  have the specified property. Then  $\theta$  and  $\theta'$  feature identical conjectures at all information sets and identical actions at every information set

---

<sup>27</sup> For example, Farrell's credible neologisms (1993) and Matthews *et al.*'s announcement proof equilibria (1991) appear to involve forces similar to those that appear in our evolutionary model.

reached during the play of any match. Because all of the actions taken in state  $\theta$  are best replies to agents' conjectures, the same is true of  $\theta'$ , ensuring that  $\theta'$  is a singleton-absorbing set. To show that  $\theta'$  lies in  $C(\theta)$ , we need only note that we can construct a sequence of singleton-absorbing sets beginning with  $\theta$  and leading to  $\theta'$ , each successive pair of which is adjacent, by altering, one agent at a time, the actions at unreached information sets the prevail in state  $\theta$  to match those of state  $\theta'$ . Similarly,  $\theta''$  is in  $C(\theta)$  because altering the conjectures of agents from populations other than  $i$  in a subgame that they cannot cause to be reached does not affect the optimality of their actions at reached information sets.

*Proof of Proposition 5.* Let  $\theta$  be a state that is not a self-confirming equilibrium. It suffices to show that with positive probability, the learning mechanism leads from  $\theta$  to a state that yields a self-confirming equilibrium (and is hence a singleton-absorbing set). This implies that every state that is not a self-confirming equilibrium lies in the basin of attraction of a self-confirming equilibrium, so that the only absorbing sets are singletons consisting of self-confirming equilibria. To show this, fix  $\theta$  and consider the evolutionary sequence in which, in each period, all agents receive the learn draw. Let  $X_1, \dots, X_m$  be a sequence of sets of nodes constructed by letting  $X_1$  consist of all nodes that are followed only by terminal nodes and letting  $X_i$  for  $i > 1$  consist of all nodes not contained in  $X_1, \dots, X_{i-1}$  and followed only by nodes contained in  $X_1, \dots, X_{i-1}$ . Let  $t_1$  be large enough that if  $x_1$  is a node of  $X_1$  and if there are at least two periods in which some match reaches node  $x_1$ , then  $x_1$  has already been reached in at least two periods prior to  $t_1$  (since  $X_1$  is finite,  $t_1$  exists). Then no conjectures about actions at any node  $x_1 \in X_1$  can change after period  $t_1$ ; because either node  $x_1$  has already been twice reached, in which case actions (after the first time the node is reached) and conjectures (after the second time) must match the unique optimal action at this node and cannot subsequently be changed (because the actions in question can never be suboptimal),<sup>28</sup> or the node is never again reached, in which case conjectures cannot be subsequently changed by the learning process. Now let  $X_2$  be the collection of nodes that are followed only by either terminal nodes or nodes in  $X_1$ . A similar argument shows that there exists a time  $t_2$  after which conjectures at nodes in  $X_2$  cannot change. Continuing this argument establishes the existence of a (finite) period  $t_m$  after which no conjectures can change. In period  $t_m + 1$ , all agents will switch to best responses to these conjectures, with these best responses confirming the conjectures (since there is no further conjecture adjustment), yielding a self-confirming equilibrium. Hence, the learning process leads to a self-confirming equilibrium with strictly positive probability, establishing the desired result. ■

---

<sup>28</sup> Note that we use here the assumption that each player moves only once along each path, so that if player  $i$  moves at this node, then conjectures about  $i$ 's behavior at this node cannot subsequently be altered by having  $i$  change an action at a prior node that previously made this node unreachable.

*Proof of Proposition 6.* Let  $C$  be a locally stable component. We show that  $C$  must contain a state yielding the subgame perfect equilibrium outcome. A simple extension of the argument showing Proposition 2 then implies that  $C$  must contain all singleton-absorbing states corresponding to the subgame perfect equilibrium outcome, thus yielding the result. To show that  $C$  must contain a state yielding the subgame perfect equilibrium outcome, consider an arbitrary singleton-absorbing set  $\theta$  contained in  $C$ . Let  $X_1, \dots, X_m$  be the sequence of nodes constructed in the proof of Proposition 5. Let  $n_\theta$  be the smallest integer such that  $\theta$  prescribes play at a node  $x_{n_\theta}$  that is reached in the course of play, that is contained in the set  $X_{n_\theta}$ , and that does not match play prescribed by the subgame perfect equilibrium. (If there is no such  $n_\theta$ , then  $\theta$  yields the subgame perfect equilibrium outcome and we are done.) By Proposition 2,  $C$  contains an element, say  $\theta'$ , with  $n_{\theta'} = n_\theta$ , with  $\theta$  and  $\theta'$  yielding identical play and conjectures, and with  $\theta'$  prescribing actions corresponding to the subgame perfect equilibrium at every node that is not reached. Consider  $\theta'$ . Let a single mutation that causes an agent at node  $x_{n_\theta}$  to switch to the subgame perfect equilibrium action occur. Then with positive probability, all agents who move at node  $x_{n_\theta}$  learn and switch to the subgame perfect equilibrium action at node  $x_{n_\theta}$ . Any subsequent learning, because it cannot alter actions following node  $x_{n_\theta}$  (or any other nodes followed by subgame perfect equilibrium actions) will then lead to a state  $\theta_1$ , which is either the subgame perfect equilibrium or satisfies  $n_{\theta_1} > n_\theta$ . Since  $C$  is locally stable, we have  $\theta_1$  in  $C$ . As this argument can be applied to any  $\theta \in C$  and  $m$  is finite, this ensures that the subgame perfect equilibrium is in  $C$ . ■

*Proof of Proposition 7.* Let  $\bar{\pi}_i$  be the maximal payoff that could result for player  $i$  if he deviates from the subgame perfect equilibrium path. Since the subgame perfect equilibrium outcome is a strict outcome, we have  $\bar{\pi}_i < \pi_i(z^*)$ . Let  $\underline{\pi}_i = \min_{z \in Z} \pi_i(z)$ ; i.e.,  $\underline{\pi}_i$  is the worst possible payoff player  $i$  could receive. Define  $\bar{\Delta}_i = (\pi_i(z^*) - \underline{\pi}_i) / (\pi_i(z^*) - \bar{\pi}_i)$  and let  $\bar{\Delta} = \max_i \bar{\Delta}_i$ . Consider  $\Delta > \bar{\Delta}$  and a singleton-absorbing set  $\theta^*$  yielding the subgame perfect equilibrium outcome. By Propositions 5 and 6, it suffices to show that, given  $\theta^*$ , a single mutation cannot yield a state that lies in the basin of attraction of a self-confirming equilibrium that does not yield the subgame perfect equilibrium outcome. Suppose that a single mutation occurs, and call the resulting state  $\theta_1$ . If this mutation changes the characteristic of an agent of player  $i$  who cannot force entry into a subgame, then  $\theta_1$  yields the subgame perfect equilibrium outcome. Hence we may assume that the mutation changes the characteristic of an agent, who we refer to as the affected agent, of player  $i$  who could force entry into a subgame  $G(a_i)$ . Suppose this mutation does not cause the affected agent to change his strategy on the equilibrium path. Then the current state remains unchanged until the affected agent receives the learn draw, and if he does so he will update his conjecture to match the observed behavior of other agents. Since the subgame perfect equilibrium outcome is a strict outcome, the resulting state must be a self-confirming equilibrium yielding the subgame perfect equilibrium path. Finally, suppose the mutation causes the affected agent to force entry into a subgame  $G(a_i)$ . By construction it is the case that for any  $\Delta > \bar{\Delta}$ , all agents but the affected agent will, upon receiving the learn draw, not change their action. Hence, the only adjustment in actions that can be caused by the learning process is that the affected agent eventually switches back to the subgame perfect equilibrium path, completing the proof of local stability of the subgame perfect equilibrium outcome. ■

*Proof of Proposition 8.* (8.1) Let  $\theta'$  be a singleton-absorbing set corresponding to the Nash equilibrium of the game  $G$  that yields player 1 a higher payoff than  $x$ . Because this Nash equilibrium is assumed to be strict,  $\{\theta'\}$  is a locally stable component for sufficiently

large  $\Delta$ . To show that the limiting distribution excludes states in which all player 1 agents play the outside option, it suffices to show that any absorbing set  $Q$  containing such a state satisfies the condition of Lemma 5, i.e.,  $\sigma(C(Q), C(\theta')) = 1$ . Hence, consider a state  $\theta$  in which all player 1 agents play the outside option. If  $\theta$  is a singleton-absorbing set, then it follows by a simple variation of the argument in Proposition 6 that  $\sigma^*(C(\theta), C(\theta')) = 1$ . [In particular,  $C(\theta)$  contains a state in which all player 2 agents play their part of the equilibrium of  $\tilde{G}$  whose player 1 payoff exceeds  $x_1$ , a single mutation causing a player 1 agent to enter  $\tilde{G}$  then yielding a state in the basin of attraction of  $\theta'$ .] Hence a singleton-absorbing set in which all agents from population 1 choose the outside option cannot appear in the limiting distribution. To complete the proof, it suffices to show that a state  $\theta$  in which all player 1 agents choose the outside option cannot be part of a nonsingleton-absorbing set. Suppose that it is. Then this absorbing set, denoted  $Q$ , must contain a state  $\theta'$  in which not all player 1 agents play the outside option (since otherwise  $Q$  would be a singleton), with  $p_{\theta'\theta} > 0$  (since the learning mechanism leads with positive probability from any state in a nonsingleton-absorbing set to any other).  $p_{\theta'\theta} > 0$  can hold only if the outside option is a strict best reply in state  $\theta'$  or is a weak best reply and no player 1 agents currently play best replies. Let either of these conditions hold. Then there is a positive probability that all (and only) player 1 agents learn in state  $\theta'$ , with these agents switching to the outside option. This yields a singleton-absorbing set (since the outside option is by hypothesis a best reply for player 1 agents and any strategies for player 2 agents are best replies to the outside option). Because  $\theta$  is a singleton-absorbing set, it cannot be that  $Q$  is a nonsingleton-absorbing set (because  $\theta \in Q$ ), yielding a contradiction.

(8.2) Let Assumption 1 hold. Then we show that there cannot exist a nonsingleton-absorbing set containing a state in which some (although possibly not all) player 1 agents play the outside option. Suppose such a set exists, denoted  $Q$ . Then  $Q$  must contain two states, say  $\theta$  and  $\theta'$ , such that no player 1 agent plays the outside option in state  $\theta'$ , some player 1 agents play the outside option in state  $\theta$ , and  $p_{\theta'\theta} > 0$ .<sup>29</sup> Then the outside option must be a best reply for player 1 agents in state  $\theta'$ . If it is a unique best reply (or is a weak best reply but no other best reply is currently played), then with positive probability all (and only) player 1 agents learn and the learning mechanism leads to a self-confirming equilibrium and hence singleton-absorbing set in which all player 1 agents choose the outside option, contradicting the assumption that  $Q$  is a nonsingleton-absorbing set. If some other best reply is currently played, then under Assumption 1.1, no player 1 agents can switch to the outside option, contradicting  $p_{\theta'\theta} > 0$ . Hence, no nonsingleton-absorbing set can contain states in which player 1 agents play the outside option. The proof is then completed by noting that,

---

<sup>29</sup> If not, then the outside option must be a best reply in every state in  $Q$ . However, there cannot exist a set of strategies for a player that are best replies in every state of a nonsingleton-absorbing set. If such a set  $S$  existed, then there must exist a state in the absorbing set  $Q$  in which all player 1 agents play strategies in  $S$ . However, no subsequent adjustments in player 1 strategies can occur, since the strategies in  $S$  are best replies in every state in  $Q$ . With positive probability, the other player's agents then all learn and switch to best replies, yielding a self-confirming equilibrium and hence a singleton-absorbing set, contradicting the fact that  $Q$  is an absorbing set.

given Assumption 1.2, there cannot exist a singleton-absorbing set in which some but not all player 1 agents play the outside option. ■

*Proof of Proposition 9.* We only consider Case 2; Cases 1 and 3 are straightforward. For sufficiently large population size, the game has three components of singleton-absorbing sets, corresponding to the three Nash equilibrium outcomes of the game:  $x$  (supported by every state in which all player 1 agents have conjectures that make choosing the outside option a best reply),  $(T, L)$ , and  $(B, R)$ . Each of the latter two components consists of a single state. Let  $n_x$  be the number of states in  $C(x)$ . Let  $n_1 = \sigma^*((B, R), C(x))$ , and let  $n_2 = \sigma^*((T, L), C(x))$ , where  $n_1, n_2 > 1$  by local stability of the strict Nash equilibria. Note that  $\sigma^*(C(x), (T, L)) = 1 = \sigma^*(C(x), (B, R))$ . Let  $D(\theta_x) = n_1 + n_2 + n_x - 1$ ,  $D((T, L)) = n_1 + n_x$ , and  $D((B, R)) = n_2 + n_x$ , where  $\theta_x$  represents any state in the component  $C(x)$ .  $D(Q)$  can intuitively be interpreted as the least number of mutations required to construct a path from every other absorbing set to  $Q$ . It is shown in Samuelson (1991a) that the limiting distribution will consist of those absorbing sets that minimize  $D(\cdot)$ . So the limit will consist entirely of  $(T, L)$  if  $n_2 > n_1$  and will consist entirely of  $(B, R)$  if  $n_2 < n_1$ . Now note that  $n_2 > n_1$  if and only if  $p > q$ , where  $(1 - p)a + pb = x$  and  $qc + (1 - q)d = x$ . Recalling that  $a > b$  and  $d > c$ , we see that  $p < q$ , and the limiting distribution will consist entirely of  $(B, R)$ , if  $(a - x)/(a - b) < (d - x)/(d - c)$  or, from Eq. (3), if  $x > x^*$ , which is the desired result. ■

## REFERENCES

- BALKENBORG, D. (1992). "Repeated Games with Common Interests and the Notion of a Strict Outcome Path," Mimeo, University of Bonn, 1992.
- BILLINGSLEY, P. (1986). *Probability and Measure*. New York: Wiley.
- BINMORE, K. G. (1987). "Modelling Rational Players. I," *Econ. Philos.* **3**, 179–214.
- BINMORE, K. G. (1988). "Modelling Rational Players. II," *Econ. Philos.* **4**, 9–55.
- CANNING, D. (1992a). "Learning Language Conventions in Common Interest Signaling Games," Department of Economics Discussion Paper Series 607, Columbia University.
- CANNING, D. (1992b). "Learning the Subgame Perfect Equilibrium," Department of Economics Discussion Paper Series 608, Columbia University.
- CRAWFORD, V. P. (1992). "Adaptive Dynamics in Coordination Games," Department of Economics Working Paper 92-02R, University of California, San Diego.
- FARRELL, J. (1993). "Meaning and Credibility in Cheap Talk Games," *Games Econ. Behav.*, in press.
- FOSTER, D., AND YOUNG, P. (1990). "Stochastic Evolutionary Game Dynamics," *J. Theor. Biol.* **38**, 219–232.
- FREIDLIN, M. I., AND WENTZELL, A. D. (1984). *Random Perturbations of Dynamical Systems*. New York: Springer-Verlag.
- FUDENBERG, D., AND LEVINE, D. K. (1991). "Self-confirming Equilibrium," Department of Economics Working Paper 581, Massachusetts Institute of Technology.

- KALAI, E., AND LEHRER, E. (1991). "Private-Beliefs Equilibrium," Center for Mathematical Studies in Economics and Management Science Discussion Paper 926, Northwestern University.
- KANDORI, M., MAILATH, G. J., AND ROB, R. (1993). "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* **61**, 29–56.
- KOHLBERG, E., AND MERTENS, J. F. (1986). "On the Strategic Stability of Equilibria," *Econometrica* **54**, 1003–1038.
- KREPS, D. M., AND WILSON, R. J. (1982). "Sequential Equilibrium," *Econometrica* **50**, 863–894.
- MARIMON, R., AND MCGRATTEN, E. (1992). "On Adaptive Learning in Strategic Games," in *Learning and Rationality in Economics* (A. Kirman and M. Salmon, Eds.). Basil Blackwell, in press.
- MATTHEWS, S. A., OKUNO-FUJIWARA, M., AND POSTLEWAITE, A. (1991). "Refining Cheap-Talk Equilibria," *J. Econ. Theory* **55**, 247–273.
- MILGROM, P., AND ROBERTS, J. (1990). "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica* **58**, 1255–1278.
- MILGROM, P., AND ROBERTS, J. (1991). "Adaptive and Sophisticated Learning in Normal Form Games," *Games Econ. Behav.* **3**, 82–100.
- NÖLDEKE, G., AND SAMUELSON, L. (1992). "The Evolutionary Foundations of Backward and Forward Induction," SFB Discussion Paper B-216, University of Bonn.
- RENY, P. J. (1993). "Common Belief and the Theory of Games with Perfect Information," *J. Econ. Theory* **59**, 257–274.
- SAMUELSON, L. (1991a). "How to Tremble if You Must," Social Systems Research Institute Working Paper 9122, University of Wisconsin.
- SAMUELSON, L. (1991b). "Limit Evolutionarily Stable Strategies in Two-Player, Normal Form Games," *Games Econ. Behav.* **3**, 110–119.
- SELTEN, R. (1975). "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive-form Games," *Int. J. Game Theory* **4**, 25–55.
- SELTEN, R. (1980). "A Note on Evolutionarily Stable Strategies in Asymmetric Animal Contests," *J. Theor. Biol.* **84**, 93–101.
- SELTEN, R. (1983). "Evolutionary Stability in Extensive Two-person Games," *Math. Soc. Sci.* **5**, 269–363.
- SELTEN, R. (1988). "Evolutionary Stability in Extensive Two-person Games-Correction and Further Development," *Math. Soc. Sci.* **16**, 223–266.
- SWINKELS, J. (1992a). "Evolution and Strategic Stability: From Maynard Smith to Kohlberg-Mertens," *J. Econ. Theory* **57**, 333–342.
- SWINKELS, J. (1992b). "Evolutionary Stability with Equilibrium Entrants," *J. Econ. Theory* **57**, 306–332.
- SWINKELS, J. (1993). "Adjustment Dynamics and Rational Play in Games," *Games Econ. Behav.*, in press.
- THOMAS, B. (1985a). "Evolutionarily Stable Sets in Mixed-strategist Models," *Theor. Pop. Biol.* **28**, 332–341.
- THOMAS, B. (1988b). "On Evolutionarily Stable Sets," *J. Math. Biol.* **22**, 105–115.
- VAN DAMME, E. (1984). "A Relation Between Perfect Equilibria in Extensive Form Games and Proper Equilibria in Normal Form Games," *Int. J. Game Theory* **13**, 1–13.

- VAN DAMME, E. (1987a). *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag.
- VAN DAMME, E. (1987b). "Stable Equilibria and Forward Induction," SFB Discussion Paper A-128/128, University of Bonn.
- VAN DAMME, E. (1989). "Stable Equilibria and Forward Induction," *J. Econ. Theory* **48**, 476–509.
- YOUNG, P. (1993). "The Evolution of Conventions," *Econometrica* **61**, 57–84.