

The Simple Geometry of Perfect Information Games

Stefano Demichelis*, Klaus Ritzberger[†]
and Jeroen M. Swinkels[‡]

January 2002

Abstract

Perfect information games have a particularly simple structure of equilibria in the associated normal form. For generic such games each of the finitely many connected components of Nash equilibria is contractible. For every perfect information game there is a unique connected and contractible component of subgame perfect equilibria. Finally, the graph of the subgame perfect equilibrium correspondence, after a very mild deformation, looks like the space of perfect information extensive form games.

*CORE, 34, Voie du Roman Pays, B-1348 Louvain-la-Neuve, Belgium. e-mail. stefano@core.ucl.ac.be

[†]Institute for Advanced Studies, Department of Economics and Finance, Stumpergasse 56, A-1060 Vienna, Austria. Tel. (+43-1) 599 91-153, Fax. (+43-1) 599 91-163, e-mail. ritzbe@ihs.ac.at

[‡]John M. Olin School of Business, Washington University at St. Louis, St. Louis, MO 63005, USA. e-mail. swinkels@mail.olin.wustl.edu

1 Introduction

Perfect information games have played an important role in the development of game theory. One of the earliest existence proofs for Nash equilibrium was given for this class of games (Kuhn (1953)). Perfect information games have guided equilibrium selection in applied models with multiple equilibria well before the refinement debate came into swing (Selten (1965)). They have provided the playground for the idea of backwards induction in its various forms. Last, but not least, they have been applied to a large number of important economic problems: bilateral bargaining, Stackelberg duopoly, wage and employment determination (Leontief (1946)), monetary policy (Barro and Gordon (1983)), and numerous other models in industrial organization.

The reason for this prominence is because of the simplicity of perfect information games. Of course, this is also the reason that perfect information is considered such a restrictive assumption in modern game theory. Such games have the highest possible degree of decomposability. Every move is the root of a subgame that can in principle be analyzed separately. This invites the application of a backwards induction or “dynamic programming” procedure. (Incidentally, perfect information games have also been used to show that such procedures do not necessarily yield intuitive results; see Rosenthal (1981) on the chain-store paradox.)

Recently, a further application for perfect information games has surfaced at the intersection of equilibrium selection theory and evolutionary game theory. As an alternative to rationality-based refinements of Nash equilibrium, evolutionary game theorists argue that equilibria should be selected that have certain dynamic (or stochastic) stability properties in evolutionary selection dynamics. (For an overview on evolutionary game theory see Weibull (1995) or Samuelson (1997).)

Quite a number of results are already available in this field. Though most refer to normal form games, many use the normal forms of perfect information games as their benchmarks and for illustrations. The spirit of the exercise is that an evolutionary selection of the subgame perfect equilibrium outcome in perfect information games provides a justification for backwards induction independent of issues of rationality. Moreover, perfect information games are sufficiently simple to allow for first advances in the theory of evolution on extensive form games. (See Nöldeke and Samuelson (1993) for an interesting analysis in more general extensive form games.)

Cressman and Schlag (1998) and Hart (1999) identify conditions for when

evolution will select the backwards induction solution in (the normal forms of) perfect information games. In other approaches the results on backwards induction are more implicit. For example, Marx (1999) shows conditions under which iterated weak dominance is implied by an adaptive dynamic. Of course, iterated weak dominance selects the backward induction outcome in games of perfect information.

Frequently, results about backward induction are implied by results on strategic stability (see Swinkels (1992) and (1993), Ritzberger and Weibull (1995), and Demichelis and Ritzberger (2000)). But such indirect conclusions depend heavily on the structure of the equilibrium set. The well-known uniqueness of subgame perfect equilibrium in generic perfect information games is helpful here, but often not sufficient, because in the normal form the whole component of Nash equilibria that induces the backwards induction outcome has to be considered.

In this paper, we aim at a complete description of the structure of the equilibrium set for perfect information games. Three results are presented. First, for generic perfect information games all components of Nash equilibria are contractible. For such games with only two players they are convex. Moreover, each component will contain a pure equilibrium, whatever the number of players. Second, every perfect information game, even if degenerate, has only one component of equilibria that contains subgame perfect ones; and this component is also contractible. This, together with uniqueness for generic games, already suggests that the subgame perfect equilibrium correspondence on the space of perfect information games is rather simple. Our third result shows that an arbitrary small deformation of the graph makes it “look like” (viz. project homeomorphically onto) the space of games.

There are important applications of these insights. In a recent paper Demichelis and Ritzberger ((2000), Theorem 1) show that for a component of Nash equilibria to be asymptotically stable in an evolutionary (deterministic continuous-time) selection dynamics (acting on the normal form) it is necessary that its index (Ritzberger (1994)) coincides with its Euler characteristic. (This is a substantial weakening of an analogous condition used by Swinkels (1993).) They apply this to two-player outside-option games (as introduced by van Damme (1989)) to show that, if it selects an outcome at all, evolution will select the forward induction outcome. (This result on dynamics is analogous to the result obtained by Hauk and Hurkens (forthcoming) in a static context.) But there are outside-option games where for no component the index does agree with the Euler characteristic.

The present result implies that for generic perfect information games the situation is better: for the backwards induction component the index will always agree with the Euler characteristic. This is because for generic perfect information games we show that all components are contractible and thus have Euler characteristic $+1$. The index assignment is very easy for such games: since every component with nonzero index contains a Mertens-stable set (Demichelis and Ritzberger (2000), Theorem 2), the backwards induction component has index $+1$ and all other components of equilibria have index 0 . To see this, note that if any other than the backwards induction component had nonzero index, it would contain a Mertens-stable set (for a definition see Mertens (1989) and (1991)). Since any Mertens-stable set contains a proper equilibrium (Myerson (1978)) and any proper equilibrium induces a sequential (hence, subgame perfect) equilibrium in any compatible normal form (van Damme (1984); Kohlberg and Mertens (1986)), such an other component would contain a subgame perfect equilibrium, in contradiction to uniqueness of the latter. That the backwards induction component has index $+1$ then follows from the property that the sum of the index across all components is $+1$.

That components of equilibria are contractible does not extend to degenerate perfect information games. A counterexample is given below. But when attention is restricted to subgame perfect equilibria, the result is resurrected. Even for degenerate perfect information games the unique component that consists of subgame perfect equilibria is contractible.

Degenerate perfect information games are of considerable interest in considering the properties of the subgame perfect equilibrium correspondence. At generic points, the subgame perfect equilibrium correspondence is a function. Hence, the behavior of the correspondence at degenerate games determines its properties. To clarify this issue, we prove a result analogous to the structure-theorem by Kohlberg and Mertens ((1986), Theorem 1): the graph of the subgame perfect equilibrium correspondence can be continuously deformed into the space of perfect information games. In particular, the subgame perfect equilibrium correspondence is upper hemi-continuous.

The rest of the paper is organized as follows. Section 2 gives basic definitions. In Section 3 it is shown that all components are contractible in the generic case. Section 4 focuses on subgame perfect equilibria and contains the other two results. Section 5 concludes. An appendix contains a discussion of the relation between mixed and behavior strategies.

2 Definitions and Notation

The following basic definitions for extensive form games are used throughout. A *tree* T is a finite connected directed graph without loops and with a distinguished node, called the *root*, that comes before all other nodes. Nodes that have no successors are called *terminal*, and all other nodes are called *moves*. The set of all nodes is denoted by N , the set of moves by X , and the root by $x_0 \in X$.

On a tree T define the *immediate predecessor* function $P : N \setminus \{x_0\} \rightarrow X$ by the condition that $P(x) \in X \setminus \{x\}$ comes before $x \in N$ and all nodes that come before x come before $P(x)$ or coincide with it, for all $x \in N$. That is, the immediate predecessor $P(x)$ of x is the “latest” node that comes before x . By convention, extend P to N , setting $P(x_0) = x_0$. By finiteness, for every $x \in N$ there is some $t = 1, 2, \dots$ such that $P^t(x) = P(P^{t-1}(x)) = x_0$, where P^0 is the identity. Say that a node x *comes before* node y if $x = P^t(y)$ for some $t \in \{1, 2, 3, \dots\}$.

Definition 1 *An n -player extensive form with perfect information is a triple $F = (T, \mathcal{X}, p)$, where*

- T is a tree,
- $\mathcal{X} = (X_0, X_1, \dots, X_n)$ is a partition of X into decision points,
- $p : P^{-1}(X_0) \rightarrow \mathbb{R}_{++}$ is a function such that

$$\sum_{y \in P^{-1}(x)} p(y) = 1 \text{ for all } x \in X_0 \tag{1}$$

which assigns probabilities to (immediate successors of) chance moves $x \in X_0$, where $P^{-1}(x) = \{y \in N \mid P(y) = x\}$ for all $x \in X$ and $P^{-1}(X_0) \equiv \cup_{x \in X_0} P^{-1}(x)$.

Moves in $X_{++} \equiv \cup_{i=1}^n X_i$ are decision points of personal players, moves in X_0 are chance moves. A *choice* for player i is a node $y \in P^{-1}(x)$ for some move $x \in X_i$. A *play* is a maximal chain (completely ordered subset) of nodes that starts with the root and ends with a terminal node. Let W denote the set of plays for the tree T . By finiteness, terminal nodes and plays are one-to-one.

An n -player perfect information extensive form *game* is a pair $G = (F, v)$, where F is an n -player extensive form with perfect information and $v = (v_1, \dots, v_n) : W \rightarrow \mathbb{R}^n$ is the *payoff function*. A *subgame* G_x of a perfect information game G is the extensive form game obtained by restricting G to the tree rising at $x \in X$.

A *pure strategy* for player $i = 1, \dots, n$ in an n -player perfect information extensive form game G is a function $s_i : X_i \rightarrow N$ such that

$$P(s_i(x)) = x \text{ for all } x \in X_i \quad (2)$$

That is, a pure strategy assigns a “next” node for each move belonging to i . The set of all pure strategies of player i is denoted by S_i . The product $S = S_1 \times \dots \times S_n$ is the set of all *pure strategy combinations*.

A *mixed strategy* for player i is a probability distribution σ_i on S_i , and Δ_i denotes the simplex of all mixed strategies for player i . The product $\Theta = \Delta_1 \times \dots \times \Delta_n$ is the set of all *mixed strategy combinations*.

A *behavior strategy* for player i is a function $b_i : P^{-1}(X_i) \rightarrow \mathbb{R}_+$, where $P^{-1}(X_i) \equiv \cup_{x \in X_i} P^{-1}(x)$, such that

$$\sum_{y \in P^{-1}(x)} b_i(y) = 1 \text{ for all } x \in X_i \quad (3)$$

and B_i denotes the set of all behavior strategies for player i - a product of $|X_i|$ simplices. The subvector $(b_i(y))_{y \in P^{-1}(x)}$ for some move $x \in X_i$ will be referred to as *behavior* of player i at x . The product $B = B_1 \times \dots \times B_n$ is the set of all *behavior strategy combinations*.

Every pure or behavior strategy combination induces (together with p) a unique transition probability to a node from its immediate predecessor. Multiplying transition probabilities over all predecessors of a node induces a nonnegative real-valued function π on nodes that assigns the probability $\pi(x)$ of node $x \in N$ being reached (from the root). From this, a unique probability distribution $\pi : W \rightarrow \mathbb{R}_+$ on plays is obtained, by selecting the corresponding terminal nodes. This distribution on plays is the *outcome* associated with the pure or behavior strategy combination.

Due to the structure of a tree, the basic relation between the outcome $\pi = (\pi(w))_{w \in W}$ and the probability $\pi(x)$ of a node $x \in N$ being reached is given by

$$\pi(x) = \sum_{x \in w} \pi(w) \quad (4)$$

for all $x \in N$. (Recall that $w \in W$ refers both to a terminal node and to a path. So, $x \in w$ means simply that x is on the path to w .) Likewise, under a mixed strategy combination $\sigma \in \Theta$ the probability $\pi_\sigma(x)$ of node $x \in N$ being reached is given by

$$\pi_\sigma(x) = \sum_{s \in S} \prod_{i=1}^n \sigma_i(s_i) \pi_s(x) \quad (5)$$

where $\pi_s(x)$ denotes the probability of node x being reached under the pure strategy combination $s \in S$. Selecting the corresponding terminal nodes, this yields the *outcome* associated with the mixed strategy combination $\sigma \in \Theta$.

The outcome is used to extend the payoff function v to pure, mixed, or behavior strategy combinations by taking the expectation of v with respect to the outcome over all plays. To distinguish, we denote by $u = (u_1, \dots, u_n)$ the payoff function on pure strategy combinations $s \in S$, and by $U = (U_1, \dots, U_n)$ the payoff functions on behavior ($b \in B$) resp. mixed ($\sigma \in \Theta$) strategy combinations.

Since for a given behavior strategy combination $b \in B$ and every move $x \in X$ the function π_b induces a unique probability distribution $\pi_b(\cdot | x)$ over plays passing through x , the conditional payoff $U_i(b|x)$ from strategy combination b given x is well defined for every $x \in X$, all $b \in B$, and all $i = 1, \dots, n$. The same, of course, holds for pure strategies.

Associated with each extensive form game G is its *normal form* (S, u) . Allowing randomized strategies yields two further normal form games associated with G : the mixed extension (Θ, U) of (S, u) , and the normal form game (B, U) played with behavior strategies. Here, reference to (Θ, U) will be expressed by referring to “equilibria in mixed strategies” (or “mixed equilibria”); reference to (B, U) will be expressed by referring to “equilibria in behavior strategies”.

Furthermore, every extensive form game gives rise to a reduced normal form. Two pure strategies of the same player are here called *strategically equivalent* (in the extensive form) if they induce the same outcomes¹ for all (pure) strategy combinations among the opponents. The *pure-strategy reduced normal form* (or the reduced normal form, for short) is the (mixed

¹Sometimes strategic equivalence is defined in the normal form, i.e. by payoff ties rather than outcome ties. If two strategies are strategically equivalent in the extensive form, then they are in the normal form. Yet, in degenerate cases there may be strategically equivalent strategies in the normal form that are not equivalent in the extensive form. Generically the two notions coincide.

extension of the) normal form game that arises when all strategically equivalent strategies are identified.

Here, the analysis will be conducted in behavior strategies. Yet, what we will have to say about the topological structure of equilibrium components carries over to the reduced normal form. (This is shown in the appendix at the end.) That behavior strategies entail no loss of generality with regard to attainable outcomes follows from Kuhn's theorem (Kuhn (1953)) and that every perfect information game automatically satisfies perfect recall.

A *Nash equilibrium* for an extensive form game is a strategy combination such that no player can gain by a unilateral deviation. A Nash equilibrium is *subgame perfect* (Selten (1965)) if it induces a Nash equilibrium in every subgame.

It is well known that for every finite normal form game the set of mixed Nash equilibria consists of finitely many (closed) connected components (see Kohlberg and Mertens (1986)). Applying this to the agent normal form of an extensive form game shows that the same is true in behavior strategies. Furthermore, for generic extensive form games there are only finitely many Nash equilibrium outcomes (Kreps and Wilson (1982)). Therefore, for almost all extensive form games the outcome is constant across every component of equilibria, both in mixed and in behavior strategies. Accordingly, we call a perfect information extensive form game $G = (F, v)$ *generic* if the outcome is constant across every connected component of Nash equilibria.

3 Generic Perfect Information Games

In this section generic perfect information games are considered. First, consider such a game with only *two* players. Let C be a component of Nash equilibria, π_C the associated outcome induced by (all) equilibria in C , and $b^1, b^2 \in C$ two Nash equilibria in the same component. Because the outcome is constant across C , the equilibrium b^1 must induce the same choices as b^2 at all decision points that are reached, i.e. for which $\pi_C(x) > 0$.

At unreached (i.e. $\pi_C(x) = 0$) decision points b^1 may induce other choices than b^2 . But, because (by perfect information) every choice at a decision point leads to the root of a subgame, the fact that b^1 is a Nash equilibrium with the same outcome as b^2 implies that the choices induced by b^1_i at unreached decision points of player i cannot make it profitable for either player to choose differently at reached decision points than under b^2 , for $i = 1, 2$.

Hence, the replacing b_i^2 by b_i^1 is irrelevant to the incentives at reached decision points. In other words, $U_{3-i}(b_i^1, b_{3-i}^2 | x) \geq U_{3-i}(b_i^1, b_{3-i}^1 | x)$ for all reached decision points $x \in X_{3-i}$ and all $b_{3-i} \in B_{3-i}$, for $i = 1, 2$. Therefore, b_{3-i}^2 is a best reply against b_i^1 for $i = 1, 2$.

Since b_{3-i}^2 is a best reply against b_i^2 by hypothesis, under the assumption that there are only two players, linearity of the payoff function implies that b_{3-i}^2 is a best reply against $\lambda b_i^1 + (1 - \lambda) b_i^2$ for all $\lambda \in [0, 1]$, for $i = 1, 2$.

Interchanging the roles of b^1 and b^2 , an analogous argument shows that b_{3-i}^1 is a best reply against $\lambda b_i^1 + (1 - \lambda) b_i^2$ for all $\lambda \in [0, 1]$, for $i = 1, 2$. But then, under the assumption of only two players, linearity of the payoff function implies that $\mu b_{3-i}^1 + (1 - \mu) b_{3-i}^2$ is a best reply against $\lambda b_i^1 + (1 - \lambda) b_i^2$ for all $\mu \in [0, 1]$ and all $\lambda \in [0, 1]$, for $i = 1, 2$. In particular, $\lambda b^1 + (1 - \lambda) b^2 \in C$ is a Nash equilibrium with outcome π_C for all $\lambda \in [0, 1]$. Thus, we have shown:

Proposition 1 *For a finite generic two-player perfect information extensive form game, every component of Nash equilibria is convex.*

Unfortunately, this conclusion is peculiar to two-player perfect information games. With more players only a weaker property holds: it will be shown that, for a generic perfect information game, all components of Nash equilibria are contractible. (A subset of a Euclidean space is *contractible* if the identity is homotopic to a constant.) But first, it is illustrated that convexity may fail for more than two players.

Example 1 *Consider the three-player extensive form game in Figure 1, where first player 1 can either terminate, yielding $v(w_1) = (5, 0, 0)$, or give the move to player 2. If player 2 is reached, she can terminate, yielding $v(w_2) = (0, 3, 3)$, or give the move to player 3. If player 3 is reached, she chooses between payoff vectors $v(w_3) = (9, 1, 2)$ and $v(w_4) = (1, 2, 1)$. (The first entry in payoff vectors is player 1's, the second player 2's, and the third player 3's*




Figure 1: A three-player perfect information game

payoff.) This yields the following normal form.

		s_2^1	s_2^2
s_1^1	5	0	0
	0	3	3
s_1^2	5	0	0
	9	1	2
		s_3^1	s_3^2

		s_2^1	s_2^2
s_1^1	5	0	0
	0	3	3
s_1^2	5	0	0
	1	2	1
		s_3^1	s_3^2

The component \bar{C} that contains the subgame perfect equilibrium (the “backwards induction component”, henceforth) is given by $\sigma_1(s_1^1) = 1$ and

$$1 \geq \sigma_2(s_2^1) \geq (8\sigma_3(s_3^1) - 4) / (1 + 8\sigma_3(s_3^1))$$

In the face $\sigma_1(s_1^1) = 1$ of the cube Θ (in which \bar{C} is contained) this is depicted by all points to the northwest of the concave curve in Figure 2, where $\sigma_3(s_3^1)$ is on the horizontal and $\sigma_2(s_2^1)$ on the vertical axis. (The subgame perfect equilibrium is located at the point (1,1) in Figure 2.) Therefore, \bar{C} fails to be convex.

The example also illustrates another point. It is tempting to prove that at least the backwards induction component \bar{C} is contractible by starting “at the

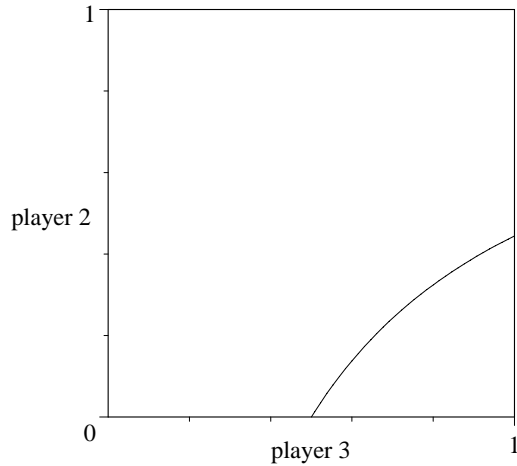


Figure 2: The backwards induction component is not convex

end” and contracting behavior strategies towards subgame perfect behavior. But, in the example, consider the equilibrium where player 1 terminates at the beginning, player 2 continues, and player 3 chooses $v(w_4) = (1, 2, 1)$ (point $(0, 0)$ in Figure 2), rather than $v(w_3) = (9, 1, 2)$, as in the subgame perfect equilibrium. This is an equilibrium in \bar{C} , because players 2 and 3 are not reached and their behavior does not induce 1 to deviate. Yet, if player 3’s choice at her (final) decision point is switched to subgame perfect equilibrium behavior (keeping strategies of other players fixed), the resulting strategy combination is not an equilibrium. (It corresponds to the bottom-right point $(1, 0)$ in Figure 2.)

Therefore, and because the argument will apply to all components, the focus cannot be on subgame perfect behavior. Rather, for a component C with associated outcome π (constant across the component), consider the strategy combination obtained by working backwards in the tree, and modifying the behavior of agents at unreached nodes to minimize the payoff of the last player who is reached along the path to this node. (Note that if x and y are two unreached nodes along the same path, the “last player” is the same, and so there is no inconsistency in this construction.) Iteratively modifying all strategies towards this strategy combination preserves equilibrium and the outcome π and, thus, remains inside the component C . This

modification results in a homotopy between the identity on C and a constant function that maps into a particular strategy combination.

In the example, the relevant player at both unreached nodes is player 1, and the actions that minimize 1's payoff are s_3^2 and s_2^1 . Following this order, all strategy combinations in Figure 2 are first moved horizontally to the left edge of the square, and then vertically up to the top *left* corner. Note that in this example, even though the component in question was the backward induction one, the resultant strategy combination is not the backward induction one. However, it is, of course, the case that if a set can be contracted to one point in itself, then it can also be contracted to any other, and so for backward induction components, the contraction toward the non-backward induction outcome is only for technical convenience.

Theorem 1 *For a finite generic perfect information extensive form game, every connected component of Nash equilibria in behavior strategies is contractible.*

Proof. Fix a generic perfect information extensive form game and order decision points as follows. Define X_{++}^0 as the union of chance moves X_0 with all terminal nodes and for any $t = 1, 2, \dots$ define X_{++}^t recursively as the set

$$\{x \in X_{++}^t \mid \text{if } x \text{ comes before } y \in N \setminus \{x\}, \text{ then } y \in \cup_{\tau=0}^{t-1} X_{++}^\tau\} \quad (6)$$

That is, X_{++}^1 contains all decision points which come before terminal nodes or chance moves only; from moves in X_{++}^2 only terminal nodes, chance moves, or moves in X_{++}^1 can be reached, and so forth.

Now, consider a particular component $C \subseteq B$ of Nash equilibria and let π_C denote the unique outcome associated with equilibria in C . A move $x \in X$ is *reached* if $\pi_C(x) = \sum_{x \in w} \pi_C(w) > 0$, otherwise it is *unreached* (see (4)).

Because the game is finite, for any unreached decision point $x \in X_{++}$ there is a unique smallest $t(x) = 1, 2, \dots$ such that $\pi_C(P^{t(x)}(x)) > 0$. Denote this by $\xi(x) = P^{t(x)}(x)$ and let $\iota(x) \in \{1, \dots, n\}$ be the player such that $\xi(x) \in X_{\iota(x)}$. That is, $\xi(x)$ is the last node reached on the path to x , and $\iota(x)$ is the player to whom $\xi(x)$ belongs. Note that $\iota(x)$ must be a personal player, because $\xi(x) \in X_0$ would imply $\pi_C(y) > 0$ for all $y \in P^{-1}(\xi(x))$ in contradiction to the construction of $\xi(x)$. If $x \in X_{++}$ is unreached, then $\pi_C(y) = 0$ for all $y \in N$ that come (properly) after $\xi(x)$.

We define now a recursive procedure for how to modify a given equilibrium $b \in C$. Consider any $b = b^0 \in C$ and an unreached decision point $x \in X_{++}^t$. Assume that behavior at all decision points (of personal players) that come after x has been adjusted to b^{t-1} in accordance with the procedure, where $b^{t-1} \in C$ is an equilibrium with outcome π_C . (If $t = 1$ this assumption is void, providing a starting point for the recursive procedure.) Let i be such that $x \in X_i$ and choose a successor $y \in P^{-1}(x)$ such that

$$U_{\iota(x)}(b^{t-1} | y) \leq U_{\iota(x)}(b^{t-1} | z) \text{ for all } z \in P^{-1}(x)$$

Then, modify i 's behavior at x such that $y \in P^{-1}(x)$ is chosen with certainty and denote the resulting behavior strategy combination by $b^x = (b_{-i}^{t-1}, b_i^x) \in B$. (That is, b_i^x is identical to b_i^{t-1} , except possibly at x .)

We claim that $\lambda b^{t-1} + (1 - \lambda) b^x \in C$ for all $\lambda \in [0, 1]$: First, since behavior under b^x differs from that under b^{t-1} only at the unreached move x , both b^x and b^{t-1} must induce the same outcome, π_C . Next, by the construction of b^x , for any unreached move $y \in X_{++}$ player $\iota(y)$ cannot gain more by deviating at $\xi(y)$ under b^x than she could have gained under b^{t-1} . Therefore, if b^{t-1} is an equilibrium, so is b^x . Moreover, for any unreached move $y \in X_{++}$ player $\iota(y)$'s conditional payoff $U_{\iota(y)}(\lambda b^{t-1} + (1 - \lambda) b^x | \xi(y))$ is linearly increasing (or constant) in λ , implying that $\lambda b^{t-1} + (1 - \lambda) b^x \in C$ for all $\lambda \in [0, 1]$, as desired.

Repeat this modification for all $x \in X_{++}^t$ to obtain $b^t \in C$. Since each $x \in X_{++}^t$ is the root of a separate subgame (by the perfect information assumption), all these modifications can be done independently.

Now, let $\tau \in \{1, 2, \dots\}$ be the maximum over unreached decision points $x \in X_{++}$ such that $P^\tau(x) = \xi(x)$. We define for every $t = 1, \dots, \tau$ recursively a continuous function $h_t : C \times [0, 1] \rightarrow C$ such that $h_t(b, 0) = b$ for all $b \in C$. For $t = \tau$ this will yield the desired homotopy.

Define $h_0 : C \times [0, 1] \rightarrow C$ as the identity, $h_0(b, \lambda) = b$ for all $b \in C$ and all $\lambda \in [0, 1]$. For $t > 0$, enumerate by x_1, \dots, x_k all unreached moves in X_{++}^t , denote by $j(m)$ the player for whom $x_m \in X_{++}^t \cap X_{j(m)}$, let $b^0(b) = h_{t-1}(b, 1)$, and define recursively

$$b^m(b) = \left(b_{-j(m)}^{m-1}(b), b_{j(m)}^{x_m} \right) \text{ for all } m = 1, \dots, k$$

where $b_{j(m)}^{x_m}$ is constructed as above. Then, define for every $t = 1, \dots, \tau$ recursively the functions $h_t : C \times [0, 1] \rightarrow C$ by $h_t(b, \lambda) = h_{t-1}(b, 2\lambda)$ for all

$\lambda \in [0, 1/2]$ and

$$h_t(b, \lambda) = (2k\lambda - k - m + 1) b^m(b) + (k + m - 2k\lambda) b^{m-1}(b)$$

for all $\lambda \in ((k + m - 1) / 2k, (k + m) / 2k]$ and all $m = 1, \dots, k$.

Choosing $t = \tau$, a continuous piecewise linear function $h_\tau : C \times [0, 1] \rightarrow C$ is obtained such that $h_\tau(b, 0) = b$ for all $b \in C$ and under $h_\tau(b, 1)$ for any unreached move $x \in \cup_{t=1}^\tau X_{++}^t$ the conditional payoff of player $\iota(x)$ given $\xi(x)$ is minimized with respect to behavior at x . Since $h_\tau(b, 1)$ does not depend on b , the desired homotopy has been constructed. ■

Since mixed strategies of the normal form induce behavior at all decision points (see Appendix), the logic of the proof carries over to mixed strategies as well.

Corollary 1 *For a finite generic perfect information extensive form game, every connected component of Nash equilibria in mixed strategies is contractible, both in the normal form and the reduced normal form.*

Proof. For the reduced normal form the statement follows from Proposition 2 in the Appendix. For the (unreduced) normal form, the proof is identical to the one of Theorem 1, except that any given mixed equilibrium $\sigma = \sigma^0 \in C \subseteq \Theta$, for an unreached decision point $x \in X_{++}^t \cap X_i$ and a node $y \in \arg \min_{z \in P^{-1}(x)} U_{\iota(x)}(\sigma^{t-1} | z)$, is modified as follows.

Define $S_i(y) = \{s_i \in S_i | s_i(x) = y\}$ as the set of i 's pure strategies that choose y at x , the mapping $\varphi_x : S_i \rightarrow S_i(y)$ by $\varphi_x(s_i)(z) = s_i(z)$ for all $z \in X_i \setminus \{x\}$, and let $\sigma_i^x \in \Delta_i$ be the unique mixed strategy that satisfies

$$\sigma_i^x(s_i) = \sum_{r_i \in \varphi_x^{-1}(s_i)} \sigma_i(r_i) \text{ for all } s_i \in S_i(y)$$

Then the behavior induced by σ_i^x agrees with the one induced by σ_i at all reached (according to π_C) moves of i .² Therefore, $\sigma \in C$ and $\sigma^x = (\sigma_{-i}, \sigma_i^x) \in \Theta$ induce the same outcome (π_C). Moreover, player $\iota(x)$ cannot gain more by deviating at $\xi(x)$ under σ^x than she could have gained under σ , by the construction of σ_i^x . Hence, if σ is an equilibrium, so is σ^x . Linearity of the

²In fact, the behavior induced by σ_i^x agrees with the one induced by σ_i at all moves of i that can be reached under both strategies (each combined with some strategy combination among the opponents), with the possible exception of x .

payoff function in σ_i then implies that $\lambda\sigma + (1 - \lambda)\sigma^x$ is an equilibrium for all $\lambda \in [0, 1]$. Since C is connected, it follows that $\lambda\sigma + (1 - \lambda)\sigma^x \in C$ for all $\lambda \in [0, 1]$.

Consequently, σ^x can play the same role as b^x in the proof of Theorem 1, and the homotopy can be constructed analogously. ■

As pointed out earlier, this result has important applications in evolutionary game theory. It implies that, if evolution selects any equilibrium component at all (by asymptotic stability in a selection dynamics), then for generic perfect information games it will be the backwards induction component (Demichelis and Ritzberger (2000), Theorem 1). This is, because the backwards induction component is the only one for which its index agrees with its Euler characteristic.

This, of course, does not imply that there is always a sensible selection dynamics for which the backwards induction component is asymptotically stable. In fact, Cressman and Schlag (1998) give an example of a perfect information game for which the backwards induction component cannot be asymptotically stable in the replicator dynamics. Since this is a game where each player has only two strategies, it is easy to show that the backwards induction component cannot be asymptotically stable in any “payoff consistent” (for a definition see Demichelis and Ritzberger (2000)) selection dynamics either.

Still, Theorem 1 shows that for perfect information games the situation is better than for two-player outside option games (van Damme (1989), Hauk and Hurkens (forthcoming)). In the latter class there are games which have no component for which the index agrees with the Euler characteristic. So, the present insight represents at least some support for backwards induction.

The preceding result has an additional interesting implication. The homotopy from Theorem 1 yields pure choices at unreached moves. Since at reached moves a backwards induction argument shows that generically all choices are pure, it follows that the constant strategy combination $h_\tau(\cdot, 1)$ is a pure strategy combination $s \in C$ which is such that the incentives for agents at the equilibrium path to deviate into unreached subgames are minimized. This demonstrates the following.

Corollary 2 *For almost all finite perfect information extensive form games, every component of Nash equilibria contains a pure strategy combination.*

Topologically Theorem 1 says that the equilibrium set of a generic perfect information game is equivalent to a finite collection of points, both in mixed and behavior strategies. Precisely one of these corresponds to the (unique) subgame perfect equilibrium.

4 Subgame Perfect Equilibria

Perfect information games constitute the prime case where subgame perfect equilibrium appears to be the natural equilibrium refinement concept. This is so because such games have the highest possible degree of decomposability. Every player, when called upon to move, knows the entire history that led to her move, there are no simultaneous decisions, and every choice leads into a subgame. This makes perfect information game the ideal domain for backwards induction. Accordingly, we now turn to the set of subgame perfect equilibria for perfect information games.

First, it is well known that in generic perfect information games the subgame perfect equilibrium is unique (in behavior strategies). So, the interesting cases are degenerate perfect information games. But for those, the conclusion from Theorem 1 on Nash equilibrium components fails, as the following example shows.

Example 2 Consider again the three-player perfect information game from Figure 1, but now with the degenerate payoff function $v(w_1) = (0, 0, 0)$, $v(w_2) = (0, 0, 0)$, $v(w_3) = (2, -1, 0)$, and $v(w_4) = (-1, 2, 0)$. (Player 3's payoffs at w_1 and w_2 do not matter for the argument.) In this game player 3 is always indifferent when she is called upon to move. Yet, depending on what player 3 chooses, the other players will want to either take their outside options (s_i^1 for $i = 1, 2$) or to pass the move on (s_i^2 for $i = 1, 2$). The set of Nash equilibria consists of a single connected component that is homeomorphic to a circle with two rectangles glued to it, and hence, is homotopy-equivalent to a circle. It is given by the union of

$$\begin{aligned} & \text{the segment } \sigma_1(s_1^2) = 1, 0 \leq \sigma_2(s_2^1) \leq 1, \sigma_3(s_3^1) = 2/3, \\ & \text{the rectangle } 0 \leq \sigma_1(s_1^1) \leq 1, \sigma_2(s_2^1) = 1, 2/3 \leq \sigma_3(s_3^1) \leq 1, \\ & \text{the segment } \sigma_1(s_1^1) = \sigma_2(s_2^1) = 1, 1/3 \leq \sigma_3(s_3^1) \leq 2/3, \\ & \text{the rectangle } \sigma_1(s_1^1) = 1, 0 \leq \sigma_2(s_2^1) \leq 1, 0 \leq \sigma_3(s_3^1) \leq 1/3, \\ & \text{the segment } 0 \leq \sigma_1(s_1^1) \leq 1, \sigma_2(s_2^2) = 1, \sigma_3(s_3^1) = 1/3, \end{aligned}$$




Figure 3: The set of Nash equilibria is homotopy-equivalent to a circle

and the segment $\sigma_1(s_1^2) = \sigma_2(s_2^2) = 1, 1/3 \leq \sigma_3(s_3^1) \leq 2/3$.

Hence, the only component of Nash equilibria in this game is not contractible. Figure 3 illustrates with $\sigma_1(s_1^1)$ on the horizontal axis, $\sigma_2(s_2^1)$ on the vertical axis, and $\sigma_3(s_3^1)$ on the third axis.

But not all the Nash equilibria of this game are subgame perfect. More precisely, all Nash equilibria belonging to the segment (third piece), $\sigma_1(s_1^1) = \sigma_2(s_2^1) = 1, 1/3 \leq \sigma_3(s_3^1) \leq 2/3$, and the rectangle (fourth piece), $\sigma_1(s_1^1) = 1, 0 \leq \sigma_2(s_2^1) \leq 1, 0 \leq \sigma_3(s_3^1) \leq 1/3$, (the rightward “vertical” rectangle and the upper line segment connecting it to the other rectangle, in Figure 3) fail subgame perfection. (All other equilibria are subgame perfect.) Therefore, if only those Nash equilibria, whose outcome corresponds to a subgame perfect equilibrium, are considered, these do form a contractible component.

The example suggests that, for degenerate perfect information games, the set of Nash equilibria may have a complicated structure. Yet, the set of subgame perfect equilibria appears to be simpler. Theorem 2 below shows that this is indeed so. Intuitively, it says that the set of subgame perfect equilibria “has no holes”, even for degenerate perfect information games.




Figure 4: Perturbing the subgame perfect equilibrium correspondence

Theorem 3 extends this observation to a statement about the graph of the subgame perfect equilibrium correspondence (from perfect information games with a fixed tree to, say, behavior strategies). It says that an arbitrary small perturbation makes the graph of this correspondence into the graph of a (continuous) function. Figure 4 illustrates this: The bold graph depicts the correspondence mapping perfect information games (on the horizontal axis) into subgame perfect equilibria (i.e. strategies on the vertical axis); the broken graph depicts an appropriate small perturbation.

Before stating the theorems, however, it has to be clarified what “subgame perfect” means for mixed strategies. For the (unreduced) normal form, subgame perfect equilibria are well defined, provided an extension of Bayes’ rule is adopted for decision points that are not reached (see Appendix). Such an extension, however, cannot be continuous.

For the reduced normal form mixed strategies only induce behavior at decision points that can be reached - but they do so continuously. Thus, for the reduced normal form subgame perfection has to be defined more loosely. Accordingly, define a mixed equilibrium of the (mixed extension of the) reduced normal form as *subgame perfect* if its image in behavior strategies contains a subgame perfect equilibrium (see Appendix). Since behavior strategies map continuously into mixed strategies (both for the normal form and the

reduced normal form; see Lemma 2 in the Appendix), we continue to work with behavior strategies.

Fix an extensive form with perfect information F with n players. Let $B = \times_{i=1}^n B_i$ be the associated space of behavior strategy combinations and identify the set of all payoff functions $v : W \rightarrow \mathbb{R}^n$ with a Euclidean space \mathbb{R}^K of appropriate dimension, i.e. with $K = n |W|$. The latter is then the space of all perfect information games with extensive form F . The *subgame perfect equilibrium correspondence* $E : \mathbb{R}^K \rightarrow B$ maps each perfect information game $G = (F, v)$ into the set $E(v)$ of its subgame perfect equilibria (in behavior strategies). Denote by $\mathcal{G} = \{(v, b) \in \mathbb{R}^K \times B \mid b \in E(v)\}$ the *graph* of the subgame perfect equilibrium correspondence.

For most of the space of perfect information games E is a function. At degenerate games the structure of $E(v)$ is yet unknown. The next result clarifies the geometry of $E(v)$ for all perfect information games.

Theorem 2 *For every perfect information extensive form game the set $E(v)$ of subgame perfect equilibria is contractible.*

Proof. For a given extensive form F denote by $\tau(F) \geq 1$ the unique integer such that $X = \cup_{t=0}^{\tau(F)} X_{++}^t$. We proceed by induction over the “size” $\tau(F)$ of the tree. If $\tau(F) = 1$, then all moves come before terminal nodes or chance moves only; hence, there is only a single player, who decides once and for all at the root. The set $E(v)$ is then the convex hull of this player’s choices, that maximize the single player’s expected payoff, and, therefore, a contractible polyhedron.

Now, suppose the statement of the proposition holds true for all $\tau(F) = 1, \dots, k - 1$ and consider an extensive form F for which $\tau(F) = k$. Let i be the player, who decides first (at the root). Each of player i ’s choices leads to the root of a subgame $G_x = (F_x, v_x)$ with $x \in P^{-1}(x_0)$ for which $\tau(F_x)$ is at most $k - 1$ (where v_x denotes the restriction of v to the plays passing through $x \in X$). Hence, by the induction hypothesis, the set of subgame perfect equilibria of each of the subgames G_x with $x \in P^{-1}(x_0)$ is contractible and polyhedral (in fact, a simplex).

Consider the map ϕ from $E(v)$ into the product $\times_{x \in P^{-1}(x_0)} E(v_x)$ of subgame perfect equilibria of subgames starting immediately after the (decision at the) root that assigns to each subgame perfect equilibrium of G the equilibria that it induces in the subgames G_x with $x \in P^{-1}(x_0)$. Since the game is finite, this map is surjective, because every (subgame perfect) equilibrium

of a subgame G_x is part of a subgame perfect equilibrium of G by Kuhn's Lemma (Kuhn (1953)).

Now consider the preimage $\phi^{-1}(b)$ of a point

$$b = (b_x)_{x \in P^{-1}(x_0)} \in \times_{x \in P^{-1}(x_0)} E(v_x)$$

where b_x denotes an equilibrium of the subgame G_x for all $x \in P^{-1}(x_0)$. This preimage is nonempty, because the map is surjective. Since the behavior strategy combination b assigns a unique probability distribution to each set of plays passing through a move $x \in P^{-1}(x_0)$, every such move is associated with a unique expected payoff for player i . Therefore, the preimage $\phi^{-1}(b)$ is the face (subsimplex) of player i 's behaviors simplex spanned by the choices at the root, that assign payoff maximizing choices at the root, and behavior consistent with b in later parts of the tree, together with b . It follows that $\phi^{-1}(b)$ is contractible (in fact, convex).

In other words, ϕ is a surjection for which the preimage of any point is contractible. Since, by the induction hypothesis, each $E(v_x)$ is also polyhedral, ϕ is a cell-like map (for a definition see Lacher (1969), p. 718). For cell-like maps on polyhedra Corollary 1.3 of Lacher ((1969), p. 720) implies that the map is a homotopy equivalence.³ It follows that the whole set $E(v)$ is contractible also for $\tau(F) = k$. ■

The proof of Theorem 2 could have been stated for mixed strategies of the normal form without changing the argument. This observation together with Proposition 2 (in the Appendix) yields the following.

Corollary 3 *For every perfect information extensive form game, the set of subgame perfect equilibria is contractible, both in the normal form and the reduced normal form.*

Furthermore, since every contractible set is connected, a perfect information game cannot have two distinct components of Nash equilibria both of which contain subgame perfect equilibria. Again, this also holds for all types of strategies.

Corollary 4 *Every perfect information extensive form game has precisely one connected component of subgame perfect equilibria, both in behavior and mixed strategies.*

³We are grateful to Steve Ferry for providing the adequate reference.

At this point we know two things: First, on an open dense subset of \mathbb{R}^K the correspondence E is a function. Second, at nongeneric points, where it is not, $E(v)$ is a contractible set. If the branches of E over generic v 's hang nicely together by contractible pieces, then there is a good chance that the whole graph \mathcal{G} “looks like” the space \mathbb{R}^K of games, at least after some mild deformation. And, indeed, a slight modification of the mapping used by Kohlberg and Mertens (1986) serves to show precisely this.

As a first step, we characterize the mapping introduced by Kohlberg and Mertens (1986) in a geometrically transparent way.

Lemma 1 *Let $\alpha : \mathbb{R}^l \rightarrow \Delta^{l-1}$ be defined by $\alpha(v) = \arg \min_{a \in \Delta^{l-1}} \|v - a\|$, where Δ^{l-1} denotes the $(l-1)$ -dimensional unit simplex, for some integer $l \geq 1$. Then:*

(a) *α is a continuous function such that $a = \alpha(v)$ is the unique fixed point of the mapping $f_v : \Delta^{l-1} \rightarrow \Delta^{l-1}$ defined by*

$$f_v(b) = \arg \max_{a \in \Delta^{l-1}} a \cdot (v - b), \text{ for all } v \in \mathbb{R}^l$$

(b) *if $u = v + re \in \mathbb{R}^l$ for some $r \in \mathbb{R}$ with $e = (1, 1, \dots, 1) \in \mathbb{R}^l$, then $\alpha(u) = \alpha(v)$.*

Proof. First, note that $\|v - a\|$ is minimized at $a \in \Delta^{l-1}$ if and only if $-\|v - a\|^2$ is maximized at $a \in \Delta^{l-1}$. Since for all $a, b \in \Delta^{l-1}$ and all $\lambda \in (0, 1)$

$$\|v - \lambda a - (1 - \lambda)b\|^2 \leq \lambda \|v - a\|^2 + (1 - \lambda) \|v - b\|^2$$

where equality implies $a = b$, that Δ^{l-1} is convex implies that $\alpha(v)$ is unique for all $v \in \mathbb{R}^l$. Hence, α is a function that is continuous by the maximum theorem (that yields upper hemi-continuity).

(a) That f_v has a fixed point, follows from Kakutani's fixed point theorem by observing that $f_v(b)$ is a face of Δ^{l-1} (and, therefore, convex) for all $b \in \Delta^{l-1}$ and applying the maximum theorem (to deduce upper hemi-continuity).

To see that $b \in f_v(b)$ if and only if $b = \alpha(v)$, let $b \in \Delta^{l-1}$ and $a = \alpha(v)$. By definition,

$$\|v - b\|^2 \geq \|v - a\|^2 = \|v - b\|^2 + 2(b - a) \cdot (v - b) + \|b - a\|^2 \quad (7)$$

$$\Leftrightarrow \|b - a\|^2 \leq 2a \cdot (v - b) - 2b \cdot (v - b) \quad (8)$$

Therefore, if $b \in f_v(b)$ and there is some $a \in \Delta^{l-1}$ such that $\|v - a\| < \|v - b\|$, then inequality (7) would be strict, so that

$$b \cdot (v - b) = \max_{c \in \Delta^{l-1}} c \cdot (v - b) \geq a \cdot (v - b) \quad (9)$$

would imply the contradiction $\|b - a\| < 0$ by (8). Hence, $b \in f_v(b)$ implies $b = \alpha(v)$. Conversely, if $b \in f_v(b)$ and $a = \alpha(v)$, then (9) implies from (8) that $\|b - a\|^2 \leq 0$, i.e. $b = a$. Because α is a function, the fixed point of f_v is unique for all $v \in \mathbb{R}^l$.

(b) Suppose $u = v + re$ for $r \in \mathbb{R}$ and let $a = \alpha(v)$. Then, using $e \cdot a = 1 = e \cdot b$ for all $b \in \Delta^{l-1}$,

$$\begin{aligned} \|u - a\|^2 &= \|v - a\|^2 + 2re \cdot v - 2r + r^2l \leq \\ \|v - b\|^2 + 2re \cdot v - 2r + r^2l &= \|u - b\|^2 \end{aligned}$$

for all $b \in \Delta^{l-1}$ verifies that $a = \alpha(u)$. ■

With this intermediate step a proof that \mathcal{G} is homeomorphic to \mathbb{R}^K is straightforward.

Theorem 3 *The graph \mathcal{G} of the subgame perfect equilibrium correspondence is homeomorphic to the space \mathbb{R}^K of perfect information games.*

Proof. Let $v \in \mathbb{R}^K$ be the payoff vector for a perfect information extensive form game $G = (F, v)$. For any player i and a decision point $x \in X_i$ denote by $H_i(x)$ the set of nodes belonging to $P^{-1}(X_i) \equiv \cup_{y \in X_i} P^{-1}(y)$ that come before x or agree with x . Likewise, for any player i and a node $y \in P^{-1}(X_i)$ denote by $L_i(y)$ the set of nodes belonging to $P^{-1}(X_i)$ that come after y , but do not agree with y .

Consider a move $x \in X_{++}^t$, for some $t \geq 1$, and let i be the player, who decides at x . Assume that to all moves $z \in X_{++}$ (of any player) that come after x (but do not agree with x) probability distributions $a_z(v) = (a_z(y|v))_{y \in P^{-1}(z)}$ have already been assigned. (If $t = 1$ this hypothesis is void, thereby providing a starting point for the recursive construction.) Associate to each play $w_y \in W$ passing through some $y \in P^{-1}(x)$ the ‘‘payoff’’

$$q_{ix}(w_y | v) = v_i(w_y) - \sum_{z \in L_i(y) \cap w_y} a_{P(z)}(z | v)$$

and denote by $Q_{ix}(v) = (Q_{ix}(y|v))_{y \in P^{-1}(x)}$ the corresponding vector of expected “payoffs”, where the expectation is taken with respect to the probability distribution induced on plays (passing through x) by the $a_z(v)$ ’s for $z \in X_{++}^k$ with $k < t$ and p (for chance moves). Set $a_x(v) = \alpha(Q_{ix}(v))$, i.e.

$$a_x(v) = (a_x(y|v))_{y \in P^{-1}(x)} = \arg \min_{a \in \Delta^{l-1}} \|Q_{ix}(v) - a\| \quad (10)$$

where $l = |P^{-1}(x)|$. Do this for all moves in X_{++}^t and repeat this procedure for all t , until a probability distribution $a_x(v) = (a_x(y|v))_{y \in P^{-1}(x)}$ has been assigned to all moves of personal players. Finally, set for all players $i = 1, \dots, n$ and all plays $w \in W$

$$\begin{aligned} v'_i(w) &= v_i(w) - \sum_{y \in P^{-1}(X_i) \cap w} a_{P(y)}(y|v) \text{ and} \\ b_i &= \left((b_i(y))_{y \in P^{-1}(x)} \right)_{x \in X_i} = (a_x(v))_{x \in X_i} \end{aligned}$$

Then, $b = (b_i)_{i=1}^n \in B$ is a behavior strategy combination and $v' \in \mathbb{R}^K$ a payoff vector for the extensive form F . This defines a continuous mapping ψ from \mathbb{R}^K to $\mathbb{R}^K \times B$ by $v \mapsto (v', b)$.

We claim that for any $v \in \mathbb{R}^K$ one has $\psi(v) = (v', b) \in \mathcal{G}$, i.e., b is a subgame perfect equilibrium for v' . To see this, let $x \in X_{++}^t \cap X_i$ and assume that $b \in B$ assigns optimal behavior at all $z \in \cup_{k=1}^{t-1} X_{++}^k$. (If $t = 1$, then the latter hypothesis is void and, thus, provides the starting point for an induction argument.) For every play $w_y \in W$ passing through some $y \in P^{-1}(x)$ the corresponding payoff of i can be written as

$$\begin{aligned} v'_i(w_y) &= v_i(w_y) - \sum_{z \in L_i(y) \cap w_y} a_{P(z)}(z|v) - b_i(y) - \sum_{z \in H_i(x) \cap w_y} b_i(z) = \\ & q_{ix}(w_y|v) - \sum_{z \in H_i(x) \cap w_y} b_i(z) - b_i(y) \end{aligned}$$

Since the set $H_i(x)$ contains only nodes that come before x (or agree with it), the sum $\gamma_x = \sum_{z \in H_i(x) \cap w_y} b_i(z)$ is a constant with respect to $y \in P^{-1}(x)$ (i.e. does not depend on y). Taking expectations over plays passing through all $y \in P^{-1}(x)$ with respect to the probability distribution induced on plays by b and p , therefore, yields the vector

$$(Q_{ix}(y|v) - \gamma_x - b_i(y))_{y \in P^{-1}(x)}$$

Since by (10) $(b_i(y))_{y \in P^{-1}(x)} = a_x(v) = \alpha(Q_{ix}(v)) = \alpha(Q_{ix}(v) - \gamma_x e)$ (the latter by Lemma 1(b)),

$$a_x(v) = (b_i(y))_{y \in P^{-1}(x)} \in \arg \max_{a \in \Delta^{l-1}} a \cdot (Q_{ix}(v) - \gamma_x e - a_x(v))$$

by Lemma 1(a), where $e = (1, 1, \dots, 1) \in \mathbb{R}^l$ and $l = |P^{-1}(x)|$. Because $Q_{ix}(v) - \gamma_x e - a_x(v)$ is precisely the vector of expected payoffs v'_i with respect to the probability distribution induced on plays (passing through any $y \in P^{-1}(x)$) by b and p , this shows that b_i also induces optimal behavior at x . By induction it follows that $b \in B$ is a subgame perfect equilibrium for $v' \in \mathbb{R}^K$.

Hence, ψ is a continuous mapping from \mathbb{R}^K to \mathcal{G} . To construct an inverse, simply set for any $(v', b) \in \mathbb{R}^K \times \mathcal{G}$

$$v''_i(w) = v'_i(w) + \sum_{y \in P^{-1}(X_i) \cap w} b_i(y)$$

for every player i and all plays $w \in W$. Then, if $\psi(v) = (v', b)$ it follows that $v'' = v$. Because continuity of ψ is evident from the construction and ψ^{-1} is linear, ψ is the desired homeomorphism. ■

This shows that the subgame perfect equilibrium correspondence over the space of perfect information extensive form games is topologically trivial. It is essentially a function (topologically a constant) that occasionally makes a few steps that add “vertical” pieces to its graph.

Precisely as in Kohlberg and Mertens ((1986), Theorem 1) it can now be shown that the projection of the homeomorphism onto the space of (perfect information) games is properly homotopic to the identity. Hence, the global degree of the subgame perfect equilibrium correspondence is +1.

Corollary 5 *There is a homeomorphism $\psi : \mathbb{R}^K \rightarrow \mathcal{G}$ and a proper homotopy $H : \mathbb{R}^K \times [0, 1] \rightarrow \mathbb{R}^K$ such that $H(\rho \circ \psi, 0) = \rho \circ \psi$, $H(\rho \circ \psi, 1) = \text{id}$, and $H(\rho \circ \psi, \lambda)$ is a homeomorphism for all $\lambda > 0$, where $\rho : \mathcal{G} \rightarrow \mathbb{R}^K$ denotes the projection onto the space of games.*

The last part of the corollary points out one important difference between the graph of the subgame perfect equilibrium correspondence and the graph of the Nash equilibrium correspondence on the space of normal form games. The latter has to be deformed quite substantially to project homeomorphically onto the space of games. Figure 5 illustrates with (normal form) games




Figure 5: Deformation of the equilibrium correspondence for the normal form

on the horizontal and strategies on the vertical axis: The bold graph depicts the Nash equilibrium correspondence and the broken one its deformation that projects homeomorphically onto the space of games.

For the subgame perfect equilibrium correspondence there is an arbitrarily small deformation that makes its graph project homeomorphically (compare Figures 4 and 5). This is, because almost everywhere E is a continuous function.

5 Conclusions

This paper contains three results on perfect information extensive form games. First, for generic such games all components of Nash equilibria are contractible. (With only two players they are even convex.) Second, even degenerate perfect information games have only one component of subgame perfect equilibria, and this component is contractible. Third, the graph of the subgame perfect equilibrium correspondence, mapping perfect information games into behavior strategy combinations, “looks like” the space of games, at least after a very small deformation.

These results give a complete understanding of the geometry of equilibria for perfect information games. The first result has important applications in

evolutionary game theory. It implies that, for this class of games, only the backwards induction component can be asymptotically stable in a wide class of (deterministic continuous time) selection dynamics. The second result says that, even for degenerate perfect information games, the set of subgame perfect equilibria has a topologically trivial structure: a single contractible component. The third result shows that, on all perfect information games, subgame perfect equilibrium behaves approximately like a continuous function.

6 Appendix

In this appendix the relation between behavior and mixed strategies is studied. The main purpose is to show that what we have said about the topological structure of equilibrium components in behavior strategies carries over to mixed strategies in the reduced normal form. (The discussion, in fact, applies to all extensive form games. We phrase it in terms of perfect information games to save on notation for general extensive form games.)

A first observation is that behavior strategies map nicely into mixed strategies.

Lemma 2 *For a perfect information extensive form, the mapping from B to Θ defined by*

$$b_i = \left((b_i(y))_{y \in P^{-1}(x)} \right)_{x \in X_i} \mapsto \left(\prod_{\substack{x \in X_i \\ s_i \in S_i}} b_i(s_i(x)) \right) \in \Delta_i \text{ for all } i = 1, \dots, n$$

*is an embedding.*⁴

Proof. Because the mapping (introduced by Kuhn (1953)) under scrutiny is continuous by construction and both B and Θ are compact subsets of Euclidean spaces, it is enough to show that the mapping $\Xi : B \rightarrow \Theta$ is one-to-one. (This is, because any continuous injection from a compact space to a Hausdorff space is an embedding.) For any player $i = 1, \dots, n$ and

⁴An *embedding* is an injection between topological spaces that maps homeomorphically onto its image.

$y \in P^{-1}(X_i) \equiv \cup_{x \in X_i} P^{-1}(x)$ define $S_i(y) = \{s_i \in S_i \mid s_i(P(y)) = y\}$ as the set of strategies of i that choose y . Since for any $x \in X_i$ and $y \in P^{-1}(x)$

$$1 = \sum_{y \in P^{-1}(x)} \sum_{s_i \in S_i(y)} \prod_{z \in X_i} b_i(s_i(z)) = \sum_{y \in P^{-1}(x)} b_i(y) \sum_{s_i \in S_i} \prod_{z \in X_i \setminus \{x\}} b_i(s_i(z)) =$$

$$\sum_{s_i \in S_i} \prod_{z \in X_i \setminus \{x\}} b_i(s_i(z)) \left[\sum_{y \in P^{-1}(x)} b_i(y) \right] = \sum_{s_i \in S_i} \prod_{z \in X_i \setminus \{x\}} b_i(s_i(z))$$

it follows that $\sum_{s_i \in S_i(y)} \prod_{x \in X_i} b_i(s_i(x)) = b_i(y)$ for all $y \in P^{-1}(x)$ and all $x \in X_i$. Therefore, that the images of $b_i \in B_i$ and $b'_i \in B_i$ in Θ agree, implies that $b_i = b'_i$, i.e. the mapping is one-to-one (injective). ■

This implies that the image of behavior strategies in mixed strategies is a submanifold of the mixed strategy space with the dimension of the behavior strategy space. Therefore, unless the player moves only once, there are many mixed strategies that are not images of behavior strategies. Consequently, there can be mixed equilibria that induce subgame perfect behavior strategy combinations, without being images of behavior strategy combinations.

A mixed strategy $\sigma_i \in \Delta_i$ for player i in the normal form *induces* a behavior strategy $b_i \in B_i$ as follows (see Kuhn (1953)). For every decision point $x \in X_i$, let

$$R_i(x) = \{s_i \in S_i \mid \exists s_{-i} \in S_{-i} \equiv \times_{j \neq i} S_j : \pi_s(x) > 0, \text{ when } s = (s_{-i}, s_i) \in S\}$$

be the strategies of i that *can reach* x , and $S_i(y) = \{s_i \in S_i \mid s_i(x) = y\}$ the strategies that choose y at x , for all $y \in P^{-1}(x)$. Then, for any $x \in X_i$ and $y \in P^{-1}(x)$,

$$b_i(y) = \frac{\sum_{s_i \in S_i(y) \cap R_i(x)} \sigma_i(s_i)}{\sum_{s_i \in R_i(x)} \sigma_i(s_i)} \quad (11)$$

whenever this quantity is defined, and $b_i = \sum_{s_i \in S_i(y)} \sigma_i(s_i)$ otherwise. Since the mapping from (11) is indeterminate, when $\sum_{s_i \in R_i(x)} \sigma_i(s_i)$ becomes zero, it is not a continuous function. Nor is its composition with the embedding from Lemma 2 continuous.

Example 3 Consider a player i in a perfect information extensive form game, who has two choices at the root x_0 , say, T and B , and then after B another decision point x_1 , where she again has two choices, say, L and R .

The player has no other decision points. Then, she has four pure strategies, say, $s_i^1 = TL$, $s_i^2 = TR$, $s_i^3 = BL$, and $s_i^4 = BR$. Consider the mixed strategy $\sigma_i = ((1 - \varepsilon - \delta)/2, (1 - \varepsilon - \delta)/2, \varepsilon, \delta)$ with $\varepsilon, \delta > 0$ small. This induces the behavior strategy $b_i(T) = 1 - \varepsilon - \delta$, $b_i(B) = \varepsilon + \delta$, $b_i(L) = \varepsilon/(\varepsilon + \delta)$, and $b_i(R) = \delta/(\varepsilon + \delta)$. The image of this behavior strategy under the embedding Ξ from Lemma 2 is

$$\Xi_i(b_i) = \left(\frac{(1 - \varepsilon - \delta)\varepsilon}{\varepsilon + \delta}, \frac{(1 - \varepsilon - \delta)\delta}{\varepsilon + \delta}, \varepsilon, \delta \right)$$

At $\delta = 0$ this evaluates to $(1 - \varepsilon, 0, \varepsilon, 0)$; at $\varepsilon = 0$ it evaluates to $(0, 1 - \delta, 0, \delta)$. But at $\varepsilon = \delta = 0$ the mixed strategy $\sigma_i^0 = (1/2, 1/2, 0, 0)$ induces the behavior strategy $b_i^0(T) = 1$, $b_i^0(B) = 0$, $b_i^0(L) = 1/2$, and $b_i^0(R) = 1/2$ that maps into $\Xi_i(b_i^0) = (1/2, 1/2, 0, 0) = \sigma_i^0$.

Matters are slightly different for the (pure-strategy) reduced normal form. Since, there, a pure strategy induces choices only at decision points that can be reached under this strategy, the embedding from Lemma 2 has to be redefined. Moreover, for the reduced normal form, the convention, how to extend (11) to decision points that cannot be reached, does not apply, because $S_i(y) = R_i(x) \cap S_i(y)$ for all $y \in P^{-1}(x)$ and all $x \in X_i$.

For, if two strategies $s_i, s'_i \in S_i$ (of the unreduced normal form) differ only at decision points $x \in X_i$ that cannot be reached under either, i.e. $s_i(x) \neq s'_i(x) \Rightarrow \{s_i, s'_i\} \subseteq S_i \setminus R_i(x)$, then they must clearly be strategically equivalent. Conversely, if there is a decision point $x \in X_i$ such that $s_i(x) \neq s'_i(x)$, but, say, $s_i \in R_i(x)$, then there is a play $w \in W$ such that $s_i(x) \in w$ and $\pi_s(w) > 0$, where $s = (s_{-i}, s_i) \in S$ for some strategy combination $s_{-i} \in S_{-i}$ among the opponents, but $\pi_{s'}(w) = 0$ (because $s'_i(x) \notin w$) whenever $s' = (s_{-i}, s'_i) \in S$; hence, s_i and s'_i cannot be strategically equivalent. In other words, two strategies are strategically equivalent if and only if they differ only at decision points that cannot be reached under either of the two.

Therefore, a pure strategy s_i of the reduced normal form specifies choices precisely at the decision points in $R_i^{-1}(s_i) = \{x \in X_i \mid s_i \in R_i(x)\}$ that are relevant for s_i , i.e. $S_i(y) \subseteq R_i(x)$ for all $y \in P^{-1}(x)$ and all $x \in X_i$. Accordingly, the embedding from Lemma 2 has to be redefined as

$$b_i = \left((b_i(y))_{y \in P^{-1}(x)} \right)_{x \in X_i} \mapsto \left(\prod_{x \in R_i^{-1}(s_i)} b_i(s_i(x)) \right)_{s_i \in S_i} \in \Delta_i \text{ for all } i \quad (12)$$

Since, for the reduced normal form, $S_i(y) \subseteq R_i(x)$ for all $y \in P^{-1}(x)$ and all $x \in X_i$, however, the proof of Lemma 2 stays the same.

The advantage of the reduced normal form is that behavior at decision points that cannot be reached is irrelevant for the embedding. Therefore, the discontinuity illustrated in Example 3 cannot occur. To make this precise, let Θ^* be the mixed strategy space of the reduced normal form and define the correspondence $\beta : \Theta^* \rightarrow B$ by the requirement that $\beta(\sigma)$ satisfies (11) whenever $\sum_{s_i \in R_i(x)} \sigma_i(s_i) > 0$, and assigns the whole simplex $\left\{ b_i : P^{-1}(x) \rightarrow \mathbb{R}_+ \mid \sum_{y \in P^{-1}(x)} b_i(y) = 1 \right\}$ otherwise, for all $x \in X_i$ and all $i = 1, \dots, n$.

Proposition 2 *For a perfect information extensive form game, let $C \subseteq B$ be a connected component of equilibria in behavior strategies, $C^* \subseteq \Theta^*$ the set of mixed strategy combinations $\sigma \in \Theta^*$ such that $\beta(\sigma) \cap C \neq \emptyset$, and $\Xi(C)$ the image of C in Θ^* under the embedding from (12). Then, C^* deformation retracts onto $\Xi(C)$.*

Proof. By definition, $\Xi(C) \subseteq C^*$. We wish to show that there is a mapping $r : C^* \times [0, 1] \rightarrow C^*$ such that (a) $r(\sigma, 0) = \sigma$ for all $\sigma \in C^*$, (b) $r(\sigma, 1) \in \Xi(C)$ for all $\sigma \in C^*$, and (c) $r(\sigma, \lambda) = \sigma$ for all $\sigma \in \Xi(C)$ and all $\lambda \in [0, 1]$. The composition of β with the embedding Ξ from (12) will serve this purpose.

First, β is an upper hemi-continuous correspondence with convex values (products of simplices with points). Since Ξ is a continuous function, the composition $\Xi \circ \beta : \Theta^* \rightarrow \Theta^*$ is also an upper hemi-continuous correspondence. If we can show that $\Xi \circ \beta$ is single-valued, it would follow that it is a continuous function.

For fixed $\sigma \in \Theta^*$ let $\theta, \theta' \in \Xi(\beta(\sigma))$ and denote by β^σ the values of β at σ . Suppose there is a move $x \in X_i$ of player i such that $\sum_{s_i \in R_i(x)} \sigma_i(s_i) = 0$. Then, there must be a move $x' \in X_i$ and a choice $y' \in P^{-1}(x')$ that both come before x such that $\sum_{s_i \in R_i(x')} \sigma_i(s_i) > 0$, but $\sum_{s_i \in S_i(y')} \sigma_i(s_i) = 0$, i.e. $\beta_i^\sigma(y') = 0$.

To see this, first note that if $x' \in X_i$ comes before $x \in X_i$, then $R_i(x) \subseteq R_i(x')$ by definition. Second, observe that if for some $x' \in X_i$ the choice $y' \in P^{-1}(x')$ comes before $x \in X_i$, then $R_i(x) \subseteq S_i(y')$, because a strategy $s_i \in S_i$ for which $s_i(x') \neq y'$ cannot reach x . Third, if there is *no* decision point $x'' \in X_i \setminus \{x\}$ that comes before x such that $y' \in P^{-1}(x')$ comes before x'' (i.e. x' is the “latest” move of i , other than x , that comes before x), then

$R_i(x) = S_i(y') \subseteq R_i(x')$, because then any strategy that can reach x' and chooses y' at x' can reach x . Now, let $x_j \in X_i$ for $j = 1, \dots, k$ be all the decision points of i that come before x such that x_1 comes before x_2, \dots , and x_{k-1} comes before $x_k = x$. By finiteness there is a largest $j = 1, \dots, k-1$ such that $\sum_{s_i \in R_i(x_j)} \sigma_i(s_i) > 0$. If there would be some $y_j \in P^{-1}(x_j)$ that comes before x such that $\sum_{s_i \in S_i(y_j)} \sigma_i(s_i) > 0$, then $\sum_{s_i \in R_i(x_{j+1})} \sigma_i(s_i) > 0$ (because y_j comes “directly” before x_{j+1}) contradicts the definition of x_j .

It follows that $\theta_i(s_i) = \prod_{x \in R_i^{-1}(s_i)} \beta_i^\sigma(s_i(x)) = 0 = \theta'_i(s_i)$, because $s_i(x') = y'$, for all $s_i \in R_i(x)$. Hence, both $\theta_i(s_i) > 0$ or $\theta'_i(s_i) > 0$ imply that $\sum_{s_i \in R_i(x)} \sigma_i(s_i) > 0$ for all $x \in R_i^{-1}(s_i)$, for all pure strategies s_i in the reduced normal form. But at such decision points β assigns a single value. Hence, because Ξ is a continuous function, $\theta = \theta'$, as desired.

Thus, the map $\theta \equiv \Xi \circ \beta : \Theta^* \rightarrow \Theta^*$ is a continuous function. By construction, for each $\sigma \in C^*$ and every player i , both σ_i and $\theta_i(\sigma)$ induce the same behavior at all decision points of i that are reached by σ , for all $i = 1, \dots, n$. Therefore, if σ is a Nash equilibrium, so is $(\sigma_{-i}, \theta_i(\sigma)) \in \Theta^*$; and $(\sigma_{-i}, \theta_i(\sigma))$ induces the same outcome as $\sigma \in C^*$, so $(\sigma_{-i}, \theta_i(\sigma)) \in C^*$. Then, linearity of the payoff function in player i 's mixed strategy implies that $(\sigma_{-i}, \lambda\theta_i(\sigma) + (1-\lambda)\sigma_i) \in C^*$ is a Nash equilibrium that, by (5), induces the same outcome as $\sigma \in C^*$, for all $\lambda \in [0, 1]$.

Define for each player $i = 1, \dots, n$ and all $\lambda \in [0, 1]$ the function $\lambda_i(\lambda) = \max\{0, \min\{1, 1 - i + n\lambda\}\}$ which is continuous, piecewise linear, zero for all $\lambda \leq (i-1)/n$, and equal to 1 for all $\lambda \geq i/n$. Then, $r : C^* \times [0, 1] \rightarrow C^*$ is given by

$$r_i(\sigma, \lambda) = \lambda_i(\lambda) \theta_i(\sigma) + [1 - \lambda_i(\lambda)] \sigma_i \text{ for all } i = 1, \dots, n$$

Since $\lambda_i(0) = 0$ for all i , $r(\sigma, 0) = \sigma$ for all $\sigma \in C^*$, verifying (a). Since $\lambda_i(1) = 1$ for all i , $r(\sigma, 1) = \theta(\sigma) \in \Xi(C)$ for all $\sigma \in C^*$, verifying (b). Since θ is the identity on $\Xi(C)$, so is $r(\cdot, \lambda)$ on $\Xi(C)$ for all $\lambda \in [0, 1]$, verifying (c). Thus, r is the desired retraction. ■

Proposition 2 can be applied regardless of whether subgame perfect or merely Nash equilibria are considered. This is, because the definition of C^* requires that all behavior strategy combinations $b \in \beta(C^*)$ induce the same outcome as some equilibrium $b' \in C$ in behavior strategies. Hence, if a mixed strategy combination $\sigma \in \Theta^*$ for the reduced normal form is declared a (mixed) subgame perfect equilibrium if $\beta(\sigma)$ contains a subgame perfect

equilibrium in behavior strategies, then C^* is a component of mixed subgame perfect equilibria.

Proposition 2 does not say that C^* and $\Xi(C)$ are homeomorphic. Indeed, they can have different dimensions. The simplest example of this is a trivial one-player game, where first chance chooses a state, then the player learns the state and gets to choose one of two alternatives at each of her decision points; the payoff is constant across all plays. Then, every strategy is a (subgame perfect) equilibrium. And all mixed strategies induce (continuously) behavior strategies as conditional probability distributions. Yet, the space of mixed strategies is three-dimensional, but the space of behavior strategies (and its image under the embedding from Lemma 2) is only two-dimensional.

References

- [1] Barro, R., and D. Gordon (1983), Rules, Discretion, and Reputation in a Model of Monetary Policy, *Journal of Monetary Economics* **12**, 101-121.
- [2] Cressman, R., and K. H. Schlag (1998), The Dynamic (In)Stability of Backwards Induction, *Journal of Economic Theory* **83**, 260-285.
- [3] van Damme, E. (1984), A Relation between Perfect Equilibria in Extensive-form Games and Proper Equilibria in Normal-form Games, *International Journal of Game Theory* **13**, 1-13.
- [4] van Damme, E. (1989), Stable Equilibria and Forward Induction, *Journal of Economic Theory* **48**, 476-496.
- [5] Demichelis, S., and K. Ritzberger, From Evolutionary to Strategic Stability, unpublished, CORE Discussion Paper 2000/59.
- [6] Hart, S. (1999), Evolutionary Dynamics and Backward Induction, unpublished, The Hebrew University of Jerusalem, Center for Rationality DP-195.
- [7] Hauk, E., and S. Hurkens (forthcoming), On Forward Induction and Evolutionary and Strategic Stability, *Journal of Economic Theory*.
- [8] Kohlberg, E., and J.-F. Mertens (1986), On the Strategic Stability of Equilibria, *Econometrica* **54**, 1003-1037.

- [9] Kreps, D. M., and R. Wilson (1982), Sequential Equilibrium, *Econometrica* **50**, 863-894.
- [10] Kuhn, H. W. (1953), Extensive Games and the Problem of Information, in: Contributions to the Theory of Games II, ed. by H. W. Kuhn and A. W. Tucker, Princeton, Princeton University Press, 193-216.
- [11] Lacher, R. C. (1969), Cell-like mappings, I, *Pacific Journal of Mathematics* **30**, 717-731.
- [12] Leontief, W. (1946), The Pure Theory of the Guaranteed Annual Wage Contract, *Journal of Political Economy* **54**, 76-79.
- [13] Marx, L. (1999), Adaptive Learning and Iterative Weak Dominance, *Games and Economic Behavior* **26**, 253-278.
- [14] Mertens, J.-F. (1989), Stable Equilibria - A Reformulation, Part I: Definition and Basic Properties, *Mathematics of Operations Research* **14**, 575-624.
- [15] Mertens, J.-F. (1991), Stable Equilibria - A Reformulation, Part II: Discussion of the Definition and Further Results, *Mathematics of Operations Research* **16**, 694-753.
- [16] Myerson, R. B. (1978), Refinements of the Nash Equilibrium Concept, *International Journal of Game Theory* **7**, 73-80.
- [17] Nöldeke, G., and L. Samuelson (1993), An Evolutionary Analysis of Backward and Forward Induction, *Games and Economic Behavior* **5**, 425-454.
- [18] Ritzberger, K. (1994), The Theory of Normal Form Games from the Differentiable Viewpoint, *International Journal of Game Theory* **23**, 201-236.
- [19] Ritzberger, K., and J. W. Weibull (1995), Evolutionary Selection in Normal Form Games, *Econometrica* **63**, 1371-1399.
- [20] Rosenthal, R. (1981), Games of Perfect Information, Predatory Pricing, and the Chain-Store Paradox, *Journal of Economic Theory* **25**, 92-100.

- [21] Samuelson, L. (1997), "Evolutionary Games and Equilibrium Selection," Cambridge, Mass., MIT Press.
- [22] Selten, R. (1965), Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit, *Zeitschrift für die gesamte Staatswissenschaft* **121**, 301-324 and 667-689.
- [23] Swinkels, J. (1992), Evolution and Strategic Stability: From Maynard Smith to Kohlberg and Mertens, *Journal of Economic Theory* **57**, 333-342.
- [24] Swinkels, J. (1993), Adjustment Dynamics and Rational Play in Games, *Games and Economic Behavior* **5**, 455-484.
- [25] Weibull, J. W. (1995), "Evolutionary Game Theory," Cambridge, Mass., MIT Press.