

Does Backwards Induction Imply Subgame Perfection?

Carlos Alós-Ferrer^a, Klaus Ritzberger^b

^a*University of Cologne, Department of Economics. Albertus-Magnus Platz, D-50923
Cologne, Germany. Phone: (+49) 221 470 8303.*

^b*Vienna Graduate School of Finance and Institute for Advanced Studies, A-1060 Vienna
(Austria). Phone: (+43) 1 59991 153.*

Abstract

In finite games subgame perfect equilibria are precisely those that are obtained by a backwards induction procedure. In large extensive form games with perfect information this equivalence does not hold: Strategy combinations fulfilling the backwards induction criterion may not be subgame perfect in general. The full equivalence is restored only under additional (topological) assumptions. This equivalence is in the form of a one-shot deviation principle for large games, which requires lower semi-continuous preferences. As corollaries we obtain one-shot deviation principles for particular classes of games, when each player moves only finitely often or when preferences are representable by payoff functions that are continuous at infinity.

Keywords: Backwards Induction, Subgame Perfection, Large Extensive Form Games, Perfect Information, One-Shot Deviation Principle

JEL: C72

1. Introduction

Subgame perfect equilibrium is a natural solution concept for extensive form games of perfect information (Selten, 1965, 1975). In the finite case subgame perfect equilibria are precisely those delivered by the backwards induction algorithm due to Kuhn (1953). The equivalence is essentially a “one-shot deviation principle” (see, e.g., Perea, 2001, Chapter 3, or Osborne and Rubinstein, 1994, Lemma 98.2, for the finite case, and Blackwell, 1965

Email addresses: `carlos.alos-ferrer@uni-koeln.de` (Carlos Alós-Ferrer),
`ritzbe@ihs.ac.at` (Klaus Ritzberger)

for the principle in dynamic programming). In order to check whether a strategy combination satisfies backwards induction the algorithm starts at the end of the tree and proceed backwards, verifying at every decision point (node) that the player controlling it has no profitable deviation. Suppose, however, that a player i plays at least twice, once at a given node x and once further “down” the tree, i.e., at some other node y that comes after his or her first decision. Backwards induction will eventually fix the player’s decision at y , and then proceed up until the first node, x , is reached. There, the algorithm will simply check whether player i can profitably deviate *at that node* (x), taking the behavior of all players in the future, including player i at y , as given.

In principle, such a strategy combination may satisfy backwards induction and still fail to constitute a subgame perfect equilibrium, indeed even a Nash equilibrium. For, a priori, even if there is no profitable deviation for i at x , taking i ’s behavior at y (and later) as given, there may still be *multiple* profitable deviations, where i changes his behavior both at x and at y , and possibly even at further nodes following y . This, of course, cannot happen in finite games, and in that sense the equivalence of backwards induction and subgame perfection amounts to the claim that checking for “one-shot deviations” is enough to establish subgame perfection.

We go beyond the class of finite games and consider general extensive form games with perfect information allowing for infinite horizon, infinite action spaces, and even infinitely many players. Moreover, we consider ordinal preferences rather than cardinal payoff functions. The purpose of this paper is to show that, in general, in infinite games backwards induction may not be equivalent to subgame perfection, but that the equivalence (in the form of a one-shot deviation principle) holds under additional assumptions.

Somewhat surprisingly, those assumptions are topological in nature. The one-shot deviation principle holds as long as all the players’ preferences are lower semi-continuous; in particular, full continuity is sufficient. Of course, this begs the question of which topology is used, and the even more primitive question of which space is endowed with a topology. Following the original formulation by von Neumann and Morgenstern (1944), we consider the space of ultimate outcomes (on which preferences are defined) as the appropriate primitive on which a topology is defined. This approach has the convenient feature that all elements of the tree, e.g. the nodes, become sets of outcomes rather than, say, abstract elements of a graph. With this approach, the one-shot deviation principle holds for *any* topology which provides the

minimal conditions necessary for players to actually be able to solve optimization problems at every node, in particular that the topology is compact and every node is a closed set of outcomes—call such topologies *admissible*. Of course, these conditions are trivial in the finite case for the discrete topology (which makes every payoff function continuous), and one obtains the standard equivalence as a straightforward corollary.

Why are topological conditions needed? Our counterexample is an infinite-horizon game where players play infinitely often, as it is the case in e.g. infinitely repeated games or Rubinstein’s (1982) alternating-offers bargaining game. The failure of equivalence arises from the fact that it is not possible to identify a “last decision” for a player. If this were possible, the one-shot deviation principle would hold without continuity assumptions. Indeed, we show that this is precisely the case for the class of games where it is always possible to identify the last decision of a player. This class includes all finite-horizon games independently of the cardinality of the action sets.

A second application of the main result concerns games with payoff functions that are continuous at infinity (Fudenberg and Levine, 1983). We show that the concept of continuity at infinity is equivalent to full continuity for a particular admissible topology. Since our result only requires lower semi-continuity (of preferences) with respect to some admissible topology, a one-shot deviation principle for games satisfying continuity at infinity follows directly.

The present results are conceptually related to the literature on dynamic programming, because a dynamic optimization problem is a one-player game. For this case the necessity of appropriate continuity assumptions is known. Subgame perfection corresponds to the concept of policy optimality in that literature, while backwards induction reduces to the statement that a policy cannot be improved through a single deviation; the latter property is sometimes referred to as “unimprovability”. Blair (1984, Example 1) provides an example showing that without what amounts to a lower semi-continuity at infinity (his axiom A2’) a one-shot deviation principle may *not* hold. That is, backwards induction for the one-player case does not imply optimality. An example to that effect was also later provided by Streufert (1993, Section 5.2). Under A2’ (plus a weak axiom of monotonicity), however, the implication does hold (Blair, 1984, Theorem 4). Thus, for dynamic (single-player) optimization problems it is known that lower semi-continuity is needed for a one-shot deviation principle to hold. Our result extends this logic in two directions: First, to subgame perfection in multi-player games and, second,

to arbitrary (admissible) topologies. The topologies that we consider allow greater flexibility, as will be illustrated below by an example (see Example 2). The result on continuity at infinity is also conceptually related to the work on consumer patience and myopia as initiated by Koopmans (1960) (see also Brown and Lewis, 1981).

Section 2 defines extensive form games with perfect information without finiteness assumptions and introduces the notation necessary for the analysis. Section 3 introduces the concept of backwards induction and shows by means of a counterexample that it does not imply subgame perfection in general. Section 4 introduces the topological framework and proves a one-shot deviation principle, which establishes that subgame perfection and backwards induction are equivalent if players' preferences are lower-semicontinuous. Section 5 contains the applications. Section 6 concludes.

2. Perfect Information Games

Many alternative definitions of extensive form games have been introduced in the literature. Since we will ultimately speak of topological properties of preferences, and preferences are defined on ultimate outcomes of the game, it is convenient to follow the original approach of von Neumann and Morgenstern (1944, Section 8). In that approach extensive form games are defined on trees, but the latter are not seen as graphs. Rather, trees are viewed as collections of (nonempty) subsets of an underlying set of outcomes (also called plays). Although Kuhn (1953) later popularized the “graphical approach” where trees are viewed as graphs on abstract nodes, in the original approach by von Neumann and Morgenstern (1944) a node is simply a collection of outcomes. The relation between both approaches is simple: A node should be seen as the set of outcomes which are still available when a player decides at that node. In that way, a node precedes another node if and only if the latter node is properly contained in the former. Intuitively, decisions discard possible outcomes and hence reduce the size of the nodes.

We have followed this approach in our previous work on large extensive form games, in particular Alós-Ferrer and Ritzberger (2005, henceforth AR1), Alós-Ferrer and Ritzberger (2008, henceforth AR2), and Alós-Ferrer and Ritzberger (2013, henceforth AR3). These papers develop the concept of a game tree, show that viewing nodes as sets of plays ordered by set inclusion is without loss of generality (AR1), characterize the class of game trees for which every pure strategy combination induces an outcome and does so

uniquely (AR2; see also Alós-Ferrer et al., 2011), and characterize discrete game trees and the associated extensive forms (AR3). In this paper, however, we need a relatively small part of the formalism, because we will deal with games of perfect information only, and those are relatively simple.

Formally, we work with discrete game trees as introduced in AR3 (see also Ritzberger, 2001, Definition 3.1), which are defined as follows.¹

Definition 1. A **discrete** (rooted and complete) **game tree** (N, \supseteq) is a collection of nonempty subsets $x \in N$ (the nodes) of a given set W (of plays or outcomes) partially ordered by set inclusion such that $W \in N$, $\{w\} \in N$ for all $w \in W$, and

(GT1) $h \subseteq N$ is a chain² if and only if there exists $w \in W$ such that $w \in \bigcap_{x \in h} x$,

(GT2) every chain in the set $X = N \setminus \{\{w\}\}_{w \in W}$ (of moves) has a maximum and either an infimum in the set $E = \{\{w\}\}_{w \in W}$ (of terminal nodes) or a minimum.³

Property (GT1) requires that if two nodes have a nonempty intersection then one contains the other (“Trivial Intersection”), and that all chains have a nonempty intersection (“Boundedness”). Property (GT2) is discreteness. Its first part (the existence of maxima) essentially states that nodes have immediate successors, which are the maxima of maximal chains of successors (AR2, p. 235). This property is called “up-discreteness” and is actually necessary for every pure strategy combination to induce a unique play/outcome (AR2, Theorem 6 and Corollary 5; see also Alós-Ferrer et al., 2011). Its second part (the existence of minima or infima) implies that every node $x \in X$ is reached after finitely many decisions (AR3, Theorem 1). This property, called “down-discreteness,” is satisfied in all standard applications and greatly simplifies the formalism, without excluding large action spaces

¹ Definition 1 is equivalent to Definition 5 of AR3 plus the added property that $\{w\} \in N$ for all $w \in W$, which is called *completeness* in that work and means that (the singleton sets made of) outcomes are included in the formalism as nodes, even if they are “reached” only after infinitely many decisions. AR3 (Proposition 4) shows that assuming completeness is without loss of generality.

² A nonempty subset of a partially ordered set is called a *chain* if it is completely ordered.

³ Maximum, minimum, and infimum of a chain are with respect to set inclusion. Henceforth \subseteq denotes weak set inclusion and \subset denotes proper inclusion.

or an infinite horizon.⁴

For each node $x \in N$ define the *up-set* (“the past”) $\uparrow x$ and the *down-set* (“the future”) $\downarrow x$ by

$$\uparrow x = \{y \in N \mid y \supseteq x\} \text{ and } \downarrow x = \{y \in N \mid x \supseteq y\}. \quad (1)$$

By the if-part of (GT1) $\uparrow x$ is a chain for all $x \in N$. A *play* is a chain of nodes $h \subseteq N$ that is maximal in N , i.e., there is no $x \in N \setminus h$ such that $h \cup \{x\}$ is a chain. Intuitively, a play is a complete history of all events along the tree, from the beginning (the root $W \in N$) to the “end”—if there is an end: Since infinite histories are allowed, plays need not be finite.

The advantage of game trees is that the set of plays can be one-to-one identified with the underlying set W (AR1, Theorem 3(c)). That is, for every ultimate outcome $w \in W$, the set $\uparrow \{w\}$ is a play, and for every play h there exists a unique outcome w such that $h = \uparrow \{w\}$, or, equivalently, $\bigcap_{x \in h} x = \{w\}$. Hence, abusing notation, one can identify outcomes and plays. That is, an element $w \in W$ can be seen either as a possible outcome (element of some node) or as a play (maximal chain of nodes). Henceforth we will not distinguish between plays and outcomes. A node can then be identified with the set of plays passing through it, and the underlying set W represents all plays.

For a discrete game tree (N, \supseteq) , let (by a slight abuse of notation) $W : N \rightarrow W$ denote the correspondence⁵ that assigns to every node, viewed as an element of the tree, the set of its constituent plays, that is, the node itself viewed as a set of plays, i.e. $W(x) = x$ for all $x \in N$. For a set $Y \subseteq N$ of nodes write $W(Y) = \bigcup_{x \in Y} x \subseteq W$ for the union, and refer to $W(Y)$ as the *plays passing through* Y .

In an extensive form game every node that has proper successors corresponds to a decision point where some player is active. In contrast, nodes (if any) “at the end” of the tree do not capture decisions. Formally, for a discrete game tree (N, \supseteq) a node $x \in N$ is *terminal* if $\downarrow x = \{x\}$. A node is a *move* if it is not terminal. It can be shown (see AR2, Lemma 1, and AR3, Lemma 3(b)) that a node $x \in N$ in a discrete game tree is terminal if and only if there is $w \in W$ such that $x = \{w\}$. Hence the set $E = \{\{w\}\}_{w \in W}$

⁴ (GT2) excludes e.g. games in continuous time.

⁵ Even though the same symbol serves for the map and its range, no confusion can arise, because the argument will always be specified.

introduced in (GT2) coincides with the set of terminal nodes. Likewise, the set $X = N \setminus E$ is the set of moves.

A more fundamental classification of nodes in a discrete game tree is as follows. Say that $x \in N \setminus \{W\}$ is *finite* if $\uparrow x \setminus \{x\}$ has a minimum and *infinite* if $x = \inf \uparrow x \setminus \{x\}$. By AR3 (Proposition 3) in a discrete game tree every node is either finite or infinite. Intuitively, a node is finite if it is reached after finitely many decisions, while infinite nodes are those which are only reached “in the limit,” e.g., in the case of infinite horizon games. Again by Proposition 3 of AR3 all infinite nodes are terminal. Denote by $F(N)$ the set of finite nodes together with the root $W \in N$. On this set a function $p : F(N) \rightarrow X$ can be defined that assigns to every finite node its *immediate predecessor*. Namely, for each $x \in F(N) \setminus \{W\}$ let

$$p(x) = \min \uparrow x \setminus \{x\} \quad (2)$$

and $p(W) = W$ by convention. Hence, $x \subset p(x) = \cap \{y \mid y \in \uparrow x \setminus \{x\}\}$ for all $x \in F(N) \setminus \{W\}$.

In order to define general extensive form games, one needs the concept of choices and the derived concept of information sets. That framework has been developed in AR1 and AR2, and specialized to discrete trees in AR3. However, in this work we deal with perfect information games only, and hence a game will be fully characterized by the tree, the set of players, their decision points, and their preferences on W .

Definition 2. An **extensive form game with perfect information** (EFPI) is a tuple $(T, I, \mathcal{X}, \succsim)$, where $T = (N, \supseteq)$ is a discrete game tree, I is the (possibly infinite) set of players, $\mathcal{X} = (X_i)_{i \in I}$ is a partition of the set of moves X of the tree T into the sets of decision points for players $i \in I$, and $\succsim = (\succsim_i)_{i \in I}$ is a profile of reflexive, complete, and transitive binary relations on the set W of plays, one *preference relation* for each player $i \in I$.

In an EFPI, at every move $x \in X$, the player for whom $x \in X_i$ makes a decision and selects an immediate successor⁶ $y \in p^{-1}(x)$. Moves of nature are ruled out. In the present setting this is natural, because in games of perfect

⁶ A node y is an immediate successor of x if and only if x is the immediate predecessor of y . The set of immediate successors of y is hence given by $p^{-1}(x) = \{y \in N \mid x = p(y)\}$. Note that immediate successors might include terminal nodes.

information with chance moves SPE may not exist (Luttmer and Mariotti, 2003).

A *pure strategy* for player $i \in I$ in an EFPI $(T, I, \mathcal{X}, \succsim)$ is a function $s_i : X_i \rightarrow N$ such that $p(s_i(x)) = x$ for all $x \in X_i$, i.e., a function that specifies which of the immediate successors of each $x \in X_i$ is chosen by player i . The set S_i of all such functions is player i 's pure strategy space. A (pure) *strategy combination* is a tuple $s = (s_i)_{i \in I} \in S \equiv \times_{i \in I} S_i$ of pure strategies, one for each player.

A *subgame* of an EFPI as above is the perfect information game that starts at some move $x \in X$. Under perfect information every move $x \in X$ is the root of a subgame. For every $x \in X$ there is a function $\phi_x : S \rightarrow x$ which assigns to the subgame starting at $x \in X$ the play $\phi_x(s) \in x$ that the strategy combination $s \in S$ (uniquely) induces in the subgame starting at x . Formally this follows from AR3, Theorem 3, and AR2, Theorems 2 and 6, and Corollary 5. This translates preferences on plays into preferences on strategy combinations.

A (pure) *Nash equilibrium* is a pure strategy combination $s \in S$ such that no player has an incentive to deviate from it. A (pure) *subgame perfect equilibrium* (SPE; Selten, 1965) is a pure strategy combination $s \in S = \times_{i \in I} S_i$ that induces a Nash equilibrium in every subgame.

3. Backwards induction

In the finite case an SPE is a backwards induction equilibrium obtained by pasting together solutions of subgames. In the infinite case we first need to formalize the notion of backwards induction without reference to an endpoint.

First, for a fixed move $x \in X$ in a discrete game tree (N, \supseteq) define the immediate successor that contains a particular play $w \in x$ by

$$\gamma(x, w) = \cup \{y \in N \mid w \in y \subset x\}. \quad (3)$$

That $\gamma(x, w)$ is indeed a node in $p^{-1}(x)$ follows from the first part of (GT2).

Call a function $f : N \rightarrow W$ that satisfies $f(x) \in x$ for all $x \in N$ a *selection* and let \mathcal{F} denote the set of all selections. Then, combining the functions ϕ_x for all $x \in X$, a strategy combination $s \in S$ uniquely defines a selection $f_s \in \mathcal{F}$ by $f_s(x) = \phi_x(s)$ for all $x \in X$ and $f_s(\{w\}) = w$ for all $w \in W$. Call this f_s the selection *induced* by $s \in S$. Conversely, in an EFPI every selection $f \in \mathcal{F}$ uniquely defines a strategy combination $s \in S$

by $s_i(x) = \gamma(x, f(x))$ for all $x \in X_i$ and all $i \in I$. In particular, f_s yields the strategy combination s .

Definition 3. For an EFPI $(T, I, \mathcal{X}, \succsim)$, a strategy combination $s \in S$ satisfies **backwards induction** if the induced selection $f_s \in \mathcal{F}$ is such that, for all players $i \in I$ and all $x \in X_i$,

(BI1) if $f_s(p(x)) \in x$ then $f_s(x) = f_s(p(x))$, and

(BI2) $f_s(y) \succsim_i f_s(x)$ for all $y \in p^{-1}(x)$.

Property (BI1) states that a player at move $p(x)$ correctly foresees the decisions at the successor x and thereafter, i.e., every player holds the correct expectations about all future decisions, including those of other players. (BI2) is states that every individual choice of every player is optimal. These two properties clearly have to hold in every SPE.

Proposition 1. *Every SPE of an EFPI satisfies backwards induction.*

Proof. If (BI2) fails, then $f_s \in \mathcal{F}$ cannot be the selection induced by an SPE, because then some player i has a profitable deviation at some node $x \in X_i$. If (BI1) fails, then the player i who controls the immediate predecessor $p(x)$ of x expects the play to continue differently than the player j who decides at x , which can also not be the case in an SPE. \square

In a finite game the converse of this result would also hold. In an infinite game, though, a strategy combination may satisfy backwards induction, but not subgame perfection, as the following example illustrates.

Example 1. Let $W = \{1, 2, \dots, \infty\}$ be the set of natural numbers, ordered by the standard relation \geq , together with an additional element denoted ∞ , taken to be strictly larger than any other element. Let the set of nodes be $N = \{(x_w)_{w \in W}, (\{w\})_{w \in W}\}$, where $x_w = \{w' \in W \mid w' \geq w\}$ for all $w \in W$, which is easily seen to form a discrete game tree (an “infinite centipede”). Consider two players, $I = \{1, 2\}$, and let $X_1 = \{x_w \in X \mid w \text{ is odd}\}$ and $X_2 = \{x_w \in X \mid w \text{ is even}\}$. Preferences are represented by the payoff functions $u_i(w) = (-1)^w / w$ for all $w = 1, 2, \dots$ and $u_i(\infty) = 1$ for $i = 1, 2$. See Figure 1 for a graphical representation.

Consider a strategy combination $s \in S$ such that player 1 always “continues” (chooses x_{w+1} at any $x_w \in X_1$) and player 2 always “stops” (chooses $\{w\} \in E$ at any $x_w \in X_2$). This strategy combination induces the selection




Figure 1: Graphical representation of Example 1. Player 1 plays at odd-numbered moves, player 2 at even-numbered ones. Both players have the same payoffs. The dotted-lines profile where player 1 always continues and player 2 always stops fulfills backwards induction but is not subgame-perfect.

$f_s \in \mathcal{F}$ given by

$$f_s(x) = \begin{cases} w + 1 & \text{if } x = x_w \in X \text{ and } w \text{ is odd} \\ w & \text{if } x = x_w \in X \text{ and } w \text{ is even or } x = \{w\} \end{cases}$$

which satisfies backwards induction, but is not an SPE. To see that (BI1) in the definition of backwards induction holds, let $x \in X_i$ for some $i \in I$ and suppose $f_s(p(x)) \in x$. By definition of f_s , it follows that $p(x) = x_w$ with w odd, so that $f_s(x) = f_s(x_{w+1}) = w + 1 = f_s(p(x))$. Therefore, condition (BI1) holds. To see (BI2), let first $x = x_w \in X$ for w odd. Then $x \in X_1$ and

$$u_1(f_s(x_w)) = u_1(w + 1) = \frac{1}{w + 1} = u_1(f_s(x_{w+1})) > u_1(f_s(\{w\})) = \frac{-1}{w}.$$

If, on the other hand, $x = x_w \in X$ for w even, then $x \in X_2$ and

$$u_2(f_s(x_w)) = u_2(w) = \frac{1}{w} = u_2(f_s(\{w\})) > u_2(f_s(x_{w+1})) = \frac{1}{w + 2}.$$

Since this completes the verification of (BI2), backwards induction holds. The resulting payoffs are $(1/2, 1/2)$. But if player 2 deviates to “always continue” (choosing x_{w+1} at each $x_w \in X_2$), she obtains $u_2(\infty) = 1$, showing that the deviation is profitable. Hence, $s \in S$ is not an SPE.

In this example backwards induction fails to detect the profitable deviation of player 2, because the deviation is, in a sense, “infinite.” That 2’s infinite deviation is profitable, even though 2 has no profitable finite devia-

tion, is due to a failure of lower semi-continuity of preferences.

To gain intuition on this point, consider the natural topology in this case. Note that W as defined is just the set of Alexandroff-compactified natural numbers. Take the corresponding topology, i.e. the Alexandroff compactification based on the discrete topology of the natural numbers.⁷ Then (W, τ) is a well-behaved topological space, e.g. it is compact and separated (Hausdorff)⁸ and all nodes are closed sets.

Payoffs in the example above are easily seen to be upper semi-continuous, but they are not lower semi-continuous and “jump upwards” at ∞ .⁹

4. A General One-Shot Deviation Principle

For finite games with perfect information a *one-shot deviation principle* holds that backwards induction is equivalent to subgame perfection.¹⁰ The last example shows that, without further assumptions, this is incorrect for large extensive form games. That is, in general backwards induction is necessary, but not sufficient for an SPE. Our main result will show that the one-shot deviation principle in infinite games takes lower semi-continuous preferences. This, however, requires us to clarify what is an acceptable topology for an extensive form game.

Definition 4. Given a discrete game tree $T = (N, \supseteq)$ with set of plays W , a topology τ on W is **T -admissible** if W is compact in τ and x is closed for every node $x \in N$.

Since closed subsets of compact sets are compact, an admissible topology is one where all nodes are compact. Intuitively, if some node were not compact, it would be easy to specify profiles of continuous preferences such that

⁷ The *Alexandroff compactification* of a topological space adds the element ∞ to the space and defines the topology on the space to consist of the open sets of the original space together with all subsets that contain ∞ and are such that their complements (in the original space) are closed and compact (Kelley, 1975, p. 150).

⁸ A topological space is *Hausdorff* or *separated* if any two distinct points can be separated by disjoint neighborhoods.

⁹ A real valued function $f : W \rightarrow \mathbb{R}$ is *upper* resp. *lower semi-continuous* if for each $r \in \mathbb{R}$ the upper resp. lower contour set $\{w \in W \mid f(w) \geq r\}$ resp. $\{w \in W \mid f(w) \leq r\}$ is closed.

¹⁰ Hendon et al. (1996) and Perea (2002) consider one-shot deviation principles for finite games with imperfect information.

the player active at that node would not be able to maximize, and hence SPE would not exist in general. The definition above is actually weaker than the conditions necessary to guarantee existence of SPE for continuous preferences (for details, see Alós-Ferrer and Ritzberger, 2016 and forthcoming).

The following lemma is needed to establish our main result. It states that, given a T -admissible topology and an open neighborhood of a play $w \in W$, there exists a finite node containing w which is fully contained in the neighborhood.

Lemma 1. *Let $T = (N, \supseteq)$ be a discrete game tree with set W of plays and τ a T -admissible topology on W . If $w \in u \in \tau$, then there is $x \in \uparrow\{w\} \cap F(N)$ such that $x \subseteq u$.*

Proof. Let $w \in u \in \tau$. If $\{w\} \in F(N)$, there is nothing to prove. Otherwise, because the tree is discrete, it follows that $\{w\} \in E$ is an infinite node. Then by Lemma 2 of AR2 $\bigcap \{y \mid y \in \uparrow\{w\} \setminus \{\{w\}\}\} = \{w\}$. If there is no y such that $\{w\} \subset y \subseteq u$, then for every $y \in \uparrow\{w\} \setminus \{\{w\}\}$ there is $w_y \in y \setminus u$. The set $\{w_y \in W \setminus u \mid y \in \uparrow\{w\} \setminus \{\{w\}\}\}$, where the set of indices is ordered by $y \leq z \Leftrightarrow z \subseteq y$, forms a sequence (by discreteness, Theorem 1(d) of AR3) in $W \setminus u$. Since $W \setminus u$ is closed, hence, compact, this sequence contains a subsequence that converges to some $w' \in W \setminus u$. Since $\uparrow\{w\} \setminus \{\{w\}\}$ is a chain, this subsequence is eventually contained in every $y \in \uparrow\{w\} \setminus \{\{w\}\}$. Hence w' is a cluster point of every such node y . Since nodes are closed by hypothesis, $w' \in y$ for all $y \in \uparrow\{w\} \setminus \{\{w\}\}$. But then $w' \in \bigcap \{y \mid y \in \uparrow\{w\} \setminus \{\{w\}\}\} = \{w\}$ implies $w' = w$, in contradiction to $w' \in W \setminus u$ and $w \in u$. Hence, there is $x \in \uparrow\{w\} \setminus \{\{w\}\} \subseteq F(N)$, the latter by Theorem 1(b) of AR3, such that $x \subseteq u$. This completes the proof. \square

With this auxiliary step at hand we are now able to prove the main result.

Theorem 1. *Let $(T, I, \mathcal{X}, \lesssim)$ be an EFPI on a discrete game tree $T = (N, \supseteq)$ with set W of plays and τ a T -admissible topology on W . If \lesssim_i is lower semi-continuous with respect to τ for all $i \in I$,¹¹ then a strategy combination $s \in S$ constitutes an SPE if and only if $s \in S$ satisfies backwards induction.*

¹¹ A binary relation \lesssim on W is *upper* resp. *lower semi-continuous* if for each $w \in W$ the upper resp. lower contour set $\{w' \in W \mid w' \lesssim w\}$ resp. $\{w' \in W \mid w' \lesssim w\}$ is closed in τ .

Proof. By Proposition 1 it is enough to show the if part. Let $s \in S$ satisfy backwards induction and let $f \in \mathcal{F}$ be the selection induced by s . Hence, $s_i(x) = \gamma(x, f(x))$ for all $x \in X_i$ and all $i \in I$.

If $s \in S$ is not an SPE, then there is some player $i \in I$ and a move $x \in X_i$ for which there exists a function $g : \downarrow x \rightarrow x$ such that $g(y) \in y$, $g(p(y)) \in y \Rightarrow g(y) = g(p(y))$, and $y \notin X_i \Rightarrow \gamma(y, g(y)) = \gamma(y, f(y))$ for all $y \in \downarrow x$, and $f(x) \prec_i g(x)$. The function g captures the deviation of player i in the subgame starting at x . The play induced by this deviation in the subgame is $g(x)$.

Consider first the case where player i deviates only finitely often along the play $g(x)$, i.e., there are only finitely many $y \in X_i \cap \downarrow x \cap \uparrow \{g(x)\}$ such that $\gamma(y, g(y)) \neq \gamma(y, f(y))$. Let y_0 be the minimum of this chain, i.e., the last deviation by player i along the play $g(x)$. Note that, for every $z \in X_i \cap \downarrow x \cap \uparrow \{g(x)\}$, the play $g(z)$ coincides with the play $g(x)$. By (BI2) $f(y_0) \succsim_i f(\gamma(y_0, g(x))) = g(x)$. Let y_1 be the minimum of the set of $y \in X_i \cap \downarrow x \cap (\uparrow y_0 \setminus \{y_0\})$ such that $\gamma(y, g(y)) \neq \gamma(y, f(y))$; that is, y_1 is the last deviation of player i before y_0 . For nodes z such that $y_0 \subset z \subset y_1$ the active player at z does not deviate from the original profile. This implies that for all these nodes $f(z) = f(y_0)$, applying (BI1) inductively. It now follows, again from (BI2), that $f(y_1) \succsim_i f(\gamma(y_1, g(x))) = f(y_0) \succsim_i g(x)$. Proceeding inductively, it follows that $f(x) \succsim_i g(x)$, contradicting $f(x) \prec_i g(x)$.

Next, consider the case where player i deviates infinitely often along the play $g(x)$, i.e. there are infinitely many $y \in X_i \cap \downarrow x \cap \uparrow \{g(x)\}$ such that $\gamma(y, g(y)) \neq \gamma(y, f(y))$. We claim that then there is another profitable deviation that deviates from f only at a finite chain of moves—which will then lead to a contradiction.

Because \succsim_i is lower semi-continuous on W , the strict upper contour set $\{w \in x \mid f(x) \prec_i w\}$ is open and $g(x) \in W$ belongs to it. By Lemma 1 there is $y' \in \uparrow \{g(x)\} \cap F(N)$ such that $y' \subseteq \{w \in x \mid f(x) \prec_i w\}$. Therefore, $f(x) \prec_i f(z)$ for all $z \in \downarrow y'$. Since player i deviates infinitely often along $g(x)$, there exists $y \in F(N)$ with $y \in X_i \cap \downarrow y' \cap \uparrow \{g(x)\}$ such that $y \subseteq \{w \in x \mid f(x) \prec_i w\}$.

Let $w^* = f(\gamma(y, g(y)))$ and define the function $g' : \downarrow x \rightarrow x$ by

$$g'(z) = \begin{cases} w^* & \text{if } z \in \downarrow x \cap \uparrow y \\ f(z) & \text{otherwise} \end{cases}$$

Note that $\gamma(z, g'(z)) = \gamma(z, f(z))$ for any player $j \neq i$ and all $z \in X_j$. This

is because for any move of player j before y the successor that contains w^* coincides with the one that contains $g(x)$.

Because the tree is discrete and $y \in F(N)$, the chain $\downarrow x \cap \uparrow y$ is finite (AR3, Theorem 1(c)), thus, g' deviates from f only at finitely many moves. That $\downarrow x \cap \uparrow y$ is a chain and the definition of w^* imply that (BI1) holds. Since $w^* \in y$, it follows that $w^* \in \{w \in x \mid f(x) \prec_i w\}$. Therefore, g' describes a finite and profitable deviation after x , which according to the previous argument cannot exist. This completes the proof. \square

It is important to realize that Theorem 1 applies for any T -admissible topology. Topological approaches to infinite horizon games have often relied on a product construction, where the set of actions at each step is assumed to be compact and compactness of the space of plays is then derived from Tychonoff's theorem (e.g., Fudenberg and Levine, 1983; Harris, 1985; Hellwig and Leininger, 1987). This imposes a particular choice of topology, which Theorem 1 does not rely on. The following example illustrates that the freedom to choose a natural topology can be advantageous.

Example 2. Consider the following two-player, infinite-horizon game. The outcome of the game is a real number $w \in [0, 1]$, to be determined jointly by two players in an iterative fashion. In period 1 player 1 can either choose some $w \in [0, 1/2)$ (in which case the game ends), or pass the decision to player 2. In the latter case in period 2 player 2 can either choose some $w \in [1/2, 3/4)$ (in which case the game ends), or return the decision to player 1, and so on. That is, player 1 moves at odd periods $t = 1, 3, 5, \dots$, and player 2 at even periods, $t = 2, 4, 6, \dots$. Let $q_t = 1 - \frac{1}{2^t}$, $t = 0, 1, \dots$, and $I_t = [q_{t-1}, q_t)$, $t = 1, 2, \dots$. At period t , the player who moves can either fix some $w \in I_t$, in which case the game ends with that choice of w , or pass the decision to the other player. If no player ever picks a number, the outcome is $w = 1$. Payoffs are functions $u_i : [0, 1] \rightarrow \mathbb{R}$, for $i = 1, 2$.

Our modeling approach for this game is simple. The natural set of plays is obviously $W = [0, 1]$, and the most natural choice of topology is the relative Euclidean topology on this interval. The tree is formed by all singleton nodes $\{w\}$, $w \in W$, and the “continue” nodes $x_t = [q_t, 1]$ for $t \geq 0$. At period t , the active player decides at move x_{t-1} and can choose among the successors $\{w\}$, with $q_{t-1} \leq w < q_t$, and the “continue” successor x_t . Since W is compact and all nodes are closed, the relative Euclidean topology is T -admissible.

In contrast, the product construction would go as follows. At period t , let the set of actions be $A_t = [q_{t-1}, q_t]$, where the action q_t stands for “continue”.




Figure 2: Graphical representation of player 1's payoff in Example 2. Horizontal scale is stretched to increase visibility.

Hence A_t is compact in the relative Euclidean topology and the product set $A = \times_{t=1}^{\infty} A_t$ is compact. As noted elsewhere (Alós-Ferrer and Ritzberger, 2016 and forthcoming), the Tychonoff construction blows up the space of plays, in this case generating an infinite-dimensional space even though the set of outcomes is simply a compact interval of real numbers.

Let us complete the description of the game by specifying payoffs. For simplicity, let u_2 be constant, $u_2(w) = k$ for all $w \in [0, 1]$. The payoff function for player 1 will be specified in a piecewise way as follows. For each $t = 1, 2, \dots$,

$$u_1(w) = (2^t - 1)(w - q_{t-1}) \quad \text{for all } w \in I_t.$$

Note that $u_1(q_{t-1}) = 0$ and $\lim_{w \rightarrow q_t^-} u_1(q_t) = q_t$ for all $t \geq 1$. Further, set $u_1(1) = 0$. Figure 2 represents this payoff function.

Consider any payoff π with $0 \leq \pi < 1$. For each $t \geq 0$, there exists a unique $w_t(\pi) \in [q_{t-1}, q_t]$ such that, for $w \in I_t$, $u_1(w) > \pi$ if and only if $w \in (w_t(\pi), q_t)$. Further, the interval $(w_t(\pi), q_t)$ can be empty ($w_t(\pi) = q_t$) only for finitely many values of t . That is, the strict upper contour set for payoff π is given by

$$U^1(\pi) = \{w \in W \mid u_1(w) > \pi\} = \bigcup_{t=1}^{\infty} (w_t(\pi), q_t)$$

which is open in the relative Euclidean topology. Further, $U^1(1) = \emptyset$. Hence,

the payoff function $u_1 : [0, 1] \rightarrow \mathbb{R}$ is lower semi-continuous in the relative Euclidean topology, and, since the constant function u_2 is also lower semicontinuous, Theorem 1 applies and every strategy profile satisfying backwards induction is a subgame-perfect equilibrium.

This game has many SPE. Consider a strategy profile as follows. Player 1 always continues, i.e. $s_1(x_t) = x_{t+1}$ for all $t \geq 0$. For each $t = 2, 4, \dots$, choose any w_t such that $u_1(w_t) > q_{t-1}$ (recall Figure 2). Let player 2 stop and choose one such w_t at period t , i.e., $s_2(x_t) = w_t$. Player 2's choice is always optimal since u_2 is constant. Player 1's choice to continue is always optimal, since player 2 "promises" him a strictly larger payoff in the following period than the supremum of the payoffs he can obtain by stopping. Hence backwards induction holds, and the profile is an SPE by Theorem 1. Other, qualitatively different SPE exist, for instance where player 2 also continues for finitely many periods and then behaves as described above.¹²

Analyzing this simple example with the Tychonoff construction gives rise to a number of difficulties. First, the payoff function u_1 is not lower semi-continuous. To see this, recall that numbers $w \in [0, 1]$ giving a strictly larger payoff than $\pi \in [0, 1)$ are those in $\bigcup_{t=1}^{\infty} (w_t(\pi), q_t)$. If plays are taken to be elements of the infinite product $A = \times_{t=1}^{\infty} [q_{t-1}, q_t]$, essentially the strict upper contour set would map into the product set $\times_{t=1}^{\infty} (w_t(\pi), q_t)$ (see below for the precise argument). But this is not an open set in the product topology. Any open set has to contain a basic open neighborhood of each of its points, where basic open neighborhoods are product sets such that finitely many coordinates are open subsets of the coordinate (action) spaces, but the remaining infinitely many coordinates are "unrestricted", i.e., they contain the whole coordinate space. It follows that, for the strict upper contour set to be open, we would need to have that $[q_{t-1}, q_t] = (w_t(\pi), q_t)$ for all but finitely many coordinates, a contradiction.

The precise argument showing that u_1 is not lower semi-continuous is slightly more involved, due to the technicalities involved in the product approach. Specifically, the function u_1 is not defined on the full product space A , because this space is far too large. Rather, as observed e.g. by Harris (1985), one needs to reconstruct the set of plays as a strict subset of A . (For the

¹² Hence the two-player game has infinitely many SPE. It is interesting to note that, if the example is recast as a one-player game (with player 1 making all choices), there exist no subgame perfect equilibrium. The example then becomes a continuous-action version of "Denardo's counterexample" as reported in Sobel (1975, Example 4).

analysis to be fruitful the set of plays also needs to be a closed subset of the product set, but we will ignore this added difficulty here; for details, see Alós-Ferrer and Ritzberger, 2016 and forthcoming.) In the present case the set of plays can be represented as a subset of the product space by recalling the convention that “continuing” corresponds to the choice $q_t \in A_t = [q_{t-1}, q_t]$. A play where players continue for $t - 1$ periods and then stop by choosing the number $w_t \in I_t$ at period t can be represented by the infinite sequence $(q_1, q_2, \dots, q_{t-1}, w_t, q_{t+1}, q_{t+2}, \dots)$. With this convention the set of plays is given by

$$H = \bigcup_{t=1}^{\infty} H_t \quad \text{where } H_t = \left\{ (w^t)_{t=1}^{\infty} \mid w^t \in A_t, w^r = q_r \forall r \neq t \right\}.$$

For a payoff $\pi \in [0, 1)$, the strict upper contour set is then

$$U^1(\pi) = \bigcup_{t=1}^{\infty} B_t \quad \text{where } B_t = \left\{ (w^t)_{t=1}^{\infty} \mid w^t \in (w_t(\pi), q_t), w^r = q_r \forall r \neq t \right\}.$$

If B were open in the relative topology on H inherited from the product topology, there would be a basic open set of the product topology whose intersection with H would be contained in B . This would imply $(w_t(\pi), q_t) = A_t$ for all but finitely many t , a contradiction. Hence, u_1 is not lower semi-continuous if one insists on the product topology. And Theorem 1 does not apply, simply because the “wrong” topology has been selected.

5. Applications

In this section we apply Theorem 1 to particular classes of games. The first class concerns games where players decide only finitely often along each play. The second allows players to decide infinitely often along a play, but assumes finitely many actions at each move and representable preferences.

5.1. Games Where Players Move Finitely Often

Consider the case of finite games. The set of outcomes, W , is finite, and one can simply take the discrete topology on W , which is trivially T -admissible. All payoff functions are automatically continuous in this topology, and the standard equivalence between backwards induction and subgame perfection follows as a trivial corollary of Theorem 1.

Corollary 1. *In finite extensive form games with perfect information a strategy combination constitutes an SPE if and only if it satisfies backwards induction.*

In the infinite case, however, the discrete topology will in general not be T -admissible, e.g. by virtue of failing compactness. In Example 1 the set W of plays is not compact in the discrete topology. It follows from Theorem 1 that in this example there exist no T -admissible topology such that the payoff functions are lower semi-continuous.

There are classes of games where the equivalence can be established without explicit reference to a topology. One such class is obtained as follows. Given an EFPI, a player i *plays finitely often* if for every play w , viewed as a maximal chain of nodes $\uparrow\{w\}$, i decides only finitely often along the play, i.e. $\uparrow\{w\} \cap X_i$ is finite for all plays $w \in W$.

The class of extensive form games where all players play finitely often includes many applications. For instance, in every finite-horizon game, all plays are finite, and hence every player can play only finitely often. A different subclass is the one where every player has only finitely many moves, i.e. X_i is finite for each $i \in I$. For instance, games with infinite horizon and infinitely many players where each player only has finitely many decisions belong to this category.¹³

The following result shows that the one-deviation principle holds for this class of games without topological assumptions. The key is that the proof of Theorem 1 does not rely on continuity if it is always possible to identify a last decision for every player along any play.

Corollary 2. *Let $(T, I, \mathcal{X}, \succsim)$ be an EFPI on a discrete game tree $T = (N, \supseteq)$ with set W of plays such that every player plays finitely often. Then, a strategy combination constitutes an SPE if and only if it satisfies backwards induction.*

Proof. Let $s \in S$ satisfy backwards induction. Proceeding as in the proof of Theorem 1, there is a player $i \in I$ who can profitably deviate in the

¹³ This does *not* mean that all games in this class are such that every player is active at finitely many moves only. Consider, for example, the classical Stackelberg duopoly (von Stackelberg, 1934). The game has a finite horizon and hence every player plays finitely often. However, the follower has a continuum of moves (one for each possible production decision of the leader).

subgame starting at some move $x \in X_i$. Since every player plays finitely often, it follows that i deviates only finitely often along the play induced by the deviation. The proof of Theorem 1 in that case reaches a contradiction without using any topological assumption. \square

5.2. Continuity at Infinity

Previous results beyond the finite case relied on a notion of “continuity at infinity,” introduced by Fudenberg and Levine (1983). For instance, Fudenberg and Tirole (1991, Theorem 4.2) use it to prove a one-stage deviation principle for infinite horizon multi-stage games with observed actions. Hendon et al. (1996, p. 281) also rely on this condition to prove a one-shot deviation principle for games with finite information sets and finite choice sets at each information set, but for an infinite horizon. For the case of perfect information this implies that backwards induction is equivalent to subgame perfection in this particular setting.¹⁴

Here we will show that the concept of continuity at infinity coincides with full continuity in a particular T -admissible topology. The insight from that is that the assumptions of Theorem 1 are strictly weaker than continuity at infinity. Hence, Theorem 1 will imply a one-shot deviation principle for games with payoff functions that are continuous at infinity.

We first construct the appropriate topology and demonstrate that it is T -admissible. Suppose that at every move only finitely many options are available, that is, $p^{-1}(x)$ is finite for each $x \in X$. For $t \geq 0$ let Y_t be the set of nodes which can be reached after t steps from the root, that is, $x \in Y_t$ if and only if $\uparrow x$ contains exactly $t + 1$ nodes. Let A_t be the set containing all possible options available at moves in Y_t plus a dummy option $*$, that is, $A_t = (\cup_{x \in Y_t} p^{-1}(x)) \cup \{*\}$. The set of plays W can be identified with a subset of the product space $\times_{t \geq 0} A_t$. For instance, a finite play would correspond to an infinite sequence where all entries after some finite t are the dummy option $*$.

Endow each A_t with the discrete topology. Then $A = \times_{t \geq 0} A_t$ can be endowed with the product topology and W inherits the relative topology.

¹⁴Sobel (1975) assumes “countable transitivity”, which amounts to an upper semi-continuity at infinity. Blair (1984) shows that what is needed for a one-shot deviation principle to hold for one-player settings is the “dual” of this assumption, that is, a lower semi-continuity at infinity.

Call this topology the *product-discrete topology* on plays. This topology is T -admissible, that is, it is compact and all nodes are closed sets.

Lemma 2. *Consider an EFPI on a discrete game tree T such that every move has finitely many immediate successors. Then the product-discrete topology on plays is T -admissible.*

Proof. For a given sequence $a \in A$ and a number k define

$$V(a, k) = \left(\prod_{t \leq k} \{a(t)\} \right) \times \left(\prod_{t \geq k+1} A_t \right). \quad (4)$$

This set is open in the product topology on A , because each $\{a(t)\}$ is open in the discrete topology on A_t and only finitely many coordinates are restricted. The set $V(a, k)$ is also closed (i.e., it is clopen) in the product topology, because it is a product of closed sets.

Now, notice that each A_t is compact by virtue of being finite, hence by Tychonoff's Theorem A is compact with the product topology. We now show that W is a closed subset of A , hence compact with the relative topology. Equivalently we show that $A \setminus W$ is open. Let a be a sequence in $A \setminus W$. This is the case if and only if some coordinate $a(t)$ is followed by a coordinate $a(t+1)$ which is incompatible with the tree, that is, one of the following three conditions holds: either (i) $a(t) \in X$ but $a(t+1) \notin p^{-1}(a(t))$, (ii) $a(t)$ is a terminal node but $a(t+1) \neq *$, or (iii) $a(t) = *$ but $a(t+1) \neq *$. Let k be the first index at which one of the three violations occurs. In all three cases the open set $V(a, k+1)$ contains a and is contained in $A \setminus W$. That is, every point in $A \setminus W$ has an open neighborhood contained in $A \setminus W$. It follows that $A \setminus W$ is open, therefore W is closed, hence compact.

Furthermore, every node, identified with the corresponding set of plays in $W \subseteq A$, is closed. To see this, notice that a node is the intersection of W with a set of the form $V(a, k)$ for some $a \in A$ and some fixed k . As observed above, the latter set is closed in the product topology of A . Hence, the node is closed in the product-discrete topology on plays. \square

To define continuity at infinity requires representable preferences. Hence, assume that for each players $i \in I$ there is a payoff function u_i that represents i 's preferences. Then, a game is *continuous at infinity* (Fudenberg and Levine, 1983) if for each $\varepsilon > 0$ there is a K such that whenever $s, s' \in S$ are identical on all Y_t for any $t \leq K$, then $|u_i(\phi(s)) - u_i(\phi(s'))| < \varepsilon$, for all

$i \in I$. Since the payoffs only depend on the induced play $\phi(s)$, this definition can be stated equivalently on plays. That is, the game is continuous at infinity if for each $\varepsilon > 0$ there is a K such that whenever $w, w' \in W$ coincide up to K , then $|u_i(w) - u_i(w')| < \varepsilon$, for all $i \in I$.¹⁵

It follows from Theorem 1 that backwards induction and subgame perfection are equivalent whenever the preferences represented by the payoff functions u_i are lower semi-continuous with respect to some T -admissible topology. Continuity at infinity is a stronger criterion, though. It is equivalent to full continuity with respect to the product-discrete topology on plays.

Lemma 3. *Consider an EFPI on a discrete game tree T such that every move has finitely many immediate successors and preferences are representable by payoff functions. It satisfies continuity at infinity if and only if payoff functions are continuous with respect to the product-discrete topology on plays.*

Proof. “if:” Let $\varepsilon > 0$ and $w \in W$. By continuity of u_i in the product-discrete topology the preimage $u_i^{-1}((u_i(w) - \varepsilon, u_i(w) + \varepsilon))$ is open. By definition of the product-discrete topology there exists a $t(w, \varepsilon)$ such that $w \in V(w, t(w, \varepsilon)) \subseteq u_i^{-1}((u_i(w) - \varepsilon, u_i(w) + \varepsilon))$ for $V(w, t(w, \varepsilon))$ as in (4). That is, $|u_i(w') - u_i(w)| < \varepsilon$ for every $w' \in W$ that coincides with w until $t(w, \varepsilon)$. However, continuity at infinity entails a uniformity in the sense that $t(w, \varepsilon)$ has to be independent of w (see also Fudenberg and Levine, 1983, Lemma 3.1). To show this, we use the fact that W is compact by Lemma 2.

Fix $\varepsilon > 0$. The collection of sets $V(w, t(w, \varepsilon/2)) \cap W$ for all $w \in W$ form an open covering of W . By compactness there is a finite subcovering with indices w_1, \dots, w_L for some finite L . Let $\bar{t} = \max_{\ell=1, \dots, L} t(w_\ell, \varepsilon/2)$ and $w, w' \in W$ be such that they coincide up to \bar{t} . There is a w_ℓ such that $w \in V(w_\ell, t(w_\ell, \varepsilon/2))$. Hence w and w_ℓ coincide up to $t(w_\ell, \varepsilon/2)$. By the previous observation this implies that $|u_i(w_\ell) - u_i(w)| < \varepsilon/2$. Since w and w' coincide up to $\bar{t} \geq t(w_\ell, \varepsilon/2)$, it follows that w' also coincides with w_ℓ up to $t(w_\ell, \varepsilon/2)$. Hence, $|u_i(w_\ell) - u_i(w')| < \varepsilon/2$. We conclude that $|u_i(w') - u_i(w)| \leq |u_i(w_\ell) - u_i(w)| + |u_i(w_\ell) - u_i(w')| < \varepsilon$. This proves continuity at infinity.

“only if:” Let (b_0, b_1) be an open interval of real numbers and let $w \in W$ be such that $u_i(w)$ lies in this interval. By continuity at infinity there is a

¹⁵ The equivalence requires the fact that for every play there is a strategy combination that induces it, see Theorem 4 of AR2.

$t \geq 0$ and a set $V(w, t)$ (as in (4)), which is open in the product-discrete topology on A , that contains w and such that every play in this set yields a payoff in (b_0, b_1) . This shows that $u_i^{-1}((b_0, b_1))$ is open. It follows that u_i is continuous in the product-discrete topology on plays. \square

Putting together the last two lemmata with Theorem 1 we obtain the following result.

Proposition 2. *In any extensive form game of perfect information where players have finitely many options available at every move and payoff functions are continuous at infinity, a strategy combination constitutes an SPE if and only if it satisfies backwards induction.*

The discussion above shows that Theorem 1 is strictly stronger than results based on continuity at infinity, because it relies on weaker assumptions. First, it relies on lower semi-continuity rather than on full continuity. Second, continuity can be with respect to any (admissible) topology, rather than a particular one based on a product structure. Third, our result relies on ordinal preferences only, not on cardinal payoff functions.

6. Conclusion

The main message of this paper is that, for infinite games, the equivalence between backwards induction and subgame perfection can be seen as a topological property. Backwards induction guarantees subgame perfection as long as the preference relations are lower semi-continuous. In particular, continuity always suffices to establish the equivalence. Continuity, however, is a topological property, and hence in order to be able to state the equivalence, one needs to specify the topology in advance. This is hardly too much to ask for in game theory, since most equilibrium existence theorems rely on continuity and compactness anyway. It makes explicit, however, that the choice of topology is not inconsequential.

In particular, we have assumed topologies such that every node is closed. In two related papers (Alós-Ferrer and Ritzberger, 2016 and forthcoming) we provide a characterization of the class of (compact) topologies such that every perfect information game with continuous preferences has a subgame perfect equilibrium (for a given discrete game tree). This characterization implies that all nodes are closed in the topology. Hence, this requirement is actually necessary for conducting equilibrium analysis in extensive form games.

Therefore, our assumption in Theorem 1 that topologies are T -admissible is a rather weak one.

Corollary 2, however, clarifies when topological assumptions are actually needed. If every player plays only finitely often, as is always the case in finite-horizon games, the one-shot deviation principle always holds without any reference to payoff continuity. The topological framework is hence needed when at least some players may play infinitely often along some plays. That class, however, includes many interesting examples, from infinitely repeated games and examples as e.g. alternating-offers bargaining (Rubinstein, 1982), all the way to games with one long-run player and infinitely many short-run players (Fudenberg and Levine, 2006, 2012), which is used as a workhorse to study self-control problems. The condition that is known from the literature on such games, continuity at infinity, turns out to be stronger than our assumptions.

In retrospect, it is natural that the discussion of the one-shot deviation principle leads to topological concepts in large games. This is precisely the class of games where there is in general no “last decision” for a given player. Hence, it becomes necessary to, explicitly or implicitly, consider behavior “in the limit.” Indeed, limits are, in themselves, topological concepts.

Acknowledgments

The authors gratefully acknowledge helpful comments and suggestions by Larry Blume, Johannes Kern, and Andrés Perea. We also thank the German Research Foundation (DFG) and the Austrian Science Fund (FWF) for financial support under projects Al-1169/1 and I1242-G16, respectively.

References

- Alós-Ferrer, C., Kern, J., Ritzberger, K., 2011. Comment on ‘Trees and Extensive Forms’. *Journal of Economic Theory* 146, 2165–2168.
- Alós-Ferrer, C., Ritzberger, K., 2005. Trees and Decisions. *Economic Theory* 25(4), 763–798.
- Alós-Ferrer, C., Ritzberger, K., 2008. Trees and Extensive Forms. *Journal of Economic Theory* 43(1), 216–250.
- Alós-Ferrer, C., Ritzberger, K., 2013. Large Extensive Form Games. *Economic Theory* 52(1), 75–102.

- Alós-Ferrer, C., Ritzberger, K., 2016. Equilibrium Existence for Large Perfect Information Games. *Journal of Mathematical Economics* 62, 5–18.
- Alós-Ferrer, C., Ritzberger, K., forthcoming. Characterizing Existence of Equilibrium for Large Extensive Form Games: A Necessity Result. *Economic Theory* , doi: 10.1007/s00199-015-0937-0.
- Blackwell, D., 1965. Discounted Dynamic Programming. *Annals of Mathematical Statistics* 36, 226–235.
- Blair, C.E., 1984. Axioms and Examples Related to Ordinal Dynamic Programming. *Mathematics of Operations Research* 9, 345–347.
- Brown, D.J., Lewis, L.M., 1981. Myopic Economic Agents. *Econometrica* 49, 359–368.
- Fudenberg, D., Levine, D.K., 1983. Subgame-Perfect Equilibria of Finite and Infinite Horizon Games. *Journal of Economic Theory* 31, 251–268.
- Fudenberg, D., Levine, D.K., 2006. A Dual-Self Model of Impulse Control. *The American Economic Review* 96, 1449–1476.
- Fudenberg, D., Levine, D.K., 2012. Timing and Self-Control. *Econometrica* 80, 1–42.
- Fudenberg, D., Tirole, J., 1991. *Game Theory*. The MIT Press, Cambridge, Massachusetts.
- Harris, C., 1985. Existence and Characterization of Perfect Equilibrium in Games of Perfect Information. *Econometrica* 53, 613–628.
- Hellwig, M., Leininger, W., 1987. On the Existence of Subgame-Perfect Equilibrium in Infinite-Action Games of Perfect Information. *Journal of Economic Theory* 43, 55–75.
- Hendon, E., Jacobsen, H.J., Sloth, B., 1996. The One-Shot-Deviation Principle for Sequential Rationality. *Games and Economic Behavior* 12, 274–282.
- Kelley, J., 1975. *General Topology*. volume 27 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin/Heidelberg/New York.

- Koopmans, T.C., 1960. Stationary Ordinal Utility and Impatience. *Econometrica* 28, 287–309.
- Kuhn, H., 1953. Extensive Games and the Problem of Information, in: Kuhn, H., Tucker, A. (Eds.), *Contributions to the Theory of Games*, vol. II. Princeton University Press, Princeton, NJ.
- Luttmer, E.G.J., Mariotti, T., 2003. The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: A Comment. *Econometrica* 71, 1909–1911.
- von Neumann, J., Morgenstern, O., 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.
- Osborne, M.J., Rubinstein, A., 1994. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts.
- Perea, A., 2001. *Rationality in Extensive Form Games*. volume 29 of *Theory and Decision Library, Series C*. Kluwer Academic Publishers, Boston/Dordrecht/London.
- Perea, A., 2002. A Note on the One-Deviation Property in Extensive Games. *Games and Economic Behavior* 40, 322–338.
- Ritzberger, K., 2001. *Foundations of Non-Cooperative Game Theory*. Oxford University Press, Oxford, UK.
- Rubinstein, A., 1982. Perfect Equilibrium in a Bargaining Model. *Econometrica* 50, 97–109.
- Selten, R., 1965. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit. *Zeitschrift für die gesamte Staatswissenschaft* 121, 301–324 and 667–689.
- Selten, R., 1975. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory* 4, 25–55.
- Sobel, M.J., 1975. Ordinal Dynamic Programming. *Management Science* 21, 967–975.

von Stackelberg, H., 1934. Marktform und Gleichgewicht. Springer-Verlag, Heidelberg.

Streufer, P.A., 1993. Markov-Perfect Equilibria in Intergenerational Games with Consistent Preferences. *Journal of Economic Dynamics and Control* 17, 929–951.