



HAL
open science

Infinite Subgame Perfect Equilibrium in the Hausdorff Difference Hierarchy

Stéphane Le Roux

► **To cite this version:**

Stéphane Le Roux. Infinite Subgame Perfect Equilibrium in the Hausdorff Difference Hierarchy. 1st International Conference on Theoretical Computer Science (TTCS), Aug 2015, Tehran, Iran. pp.147-163, 10.1007/978-3-319-28678-5_11 . hal-01446258

HAL Id: hal-01446258

<https://inria.hal.science/hal-01446258v1>

Submitted on 25 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Infinite subgame perfect equilibrium in the Hausdorff difference hierarchy

Stéphane Le Roux

Université Libre de Bruxelles

Abstract. Subgame perfect equilibria are specific Nash equilibria in perfect information games in extensive form. They are important because they relate to the rationality of the players. They always exist in infinite games with continuous real-valued payoffs, but may fail to exist even in simple games with slightly discontinuous payoffs. This article considers only games whose outcome functions are measurable in the Hausdorff difference hierarchy of the open sets (*i.e.* Δ_2^0 when in the Baire space), and it characterizes the families of linear preferences such that every game using these preferences has a subgame perfect equilibrium: the preferences without infinite ascending chains (of course), and such that for all players a and b and outcomes x, y, z we have $\neg(z <_a y <_a x \wedge x <_b z <_b y)$. Moreover at each node of the game, the equilibrium constructed for the proof is Pareto-optimal among all the outcomes occurring in the subgame. Additional results for non-linear preferences are presented.

Keywords: infinite multi-player games in extensive form, subgame perfection, Borel hierarchy, preference characterization, Pareto-optimality.

1 Introduction

Game theory is the theory of competitive interactions between agents having different interests. Until the late 1960's an agent would usually represent a human or group of humans, when game theory was mainly meant for economics and political science. Then game theory was also applied to evolutionary biology [24] and to theoretical computer science [4], especially to system verification and system synthesis (against given specifications). Classically, the verification or synthesis problem is represented as a game with two players: the system trying to win the game by meeting the specifications, and the environment trying to win the game by preventing the system from doing so. The two players play this game in turn and deterministically on a finite or infinite directed graph, and the key notion is that of winning strategy. For a decade, though, computer scientists such as Ummels [26] have been considering multi-player games to represent more complex verification or synthesis problems, *e.g.* relating to distributed systems. The notion of winning strategy is specific to two-player win-lose games, but in a multi-player setting it may be replaced with a faithful extension, namely the famous notion of (pure) Nash equilibrium. It does not only accommodate more than two players, it also allows for refined/quantitative objectives.

The deterministic turn-based games on graphs may be unfolded, usually without much loss of information, into deterministic turn-based games on finite or infinite trees, which have been widely studied in game theory. It is one reason why this article focuses on perfect information games in extensive form (*i.e.* played on trees) and their deterministic strategies, unless otherwise stated.


Kuhn [12] proved the existence of Nash equilibrium (NE) in finite games with real-valued payoffs. His proofs uses backward induction and constructs a special kind of NE that was later called subgame perfect equilibrium (SPE) by Selten [23]. An extension of Kuhn's result [13] characterizes the preferences that always yield finite games with NE/SPE: the acyclic preferences. Also, Escardó and Oliva [5] studied generalizations of backward induction in possibly infinite yet well-founded game trees, *i.e.* without infinite plays. The SPE have nice extra properties and are usually preferred over the more general NE: for psychology in a broad sense an SPE amounts to the absence of empty threats, and for system engineering in a broad sense it amounts to stability of a system regardless of the initial state.

The concept of infinite horizon is convenient in economics and also in computer science, *e.g.*, for liveness. Gale and Stewart [9] studied infinite two-player win-lose games, but backward induction is no longer applicable since there may not be leaves to start the induction at. Nevertheless, they proved that if the winning set of each player is open or closed (with the usual topology), one player has a winning strategy. This result was extended by Wolfe [27] for Σ_2^0 and Π_2^0 sets, then by other people to more complex sets, and eventually by Martin for Borel [19] and even quasi-Borel [20] sets. This is called (quasi-)Borel determinacy.

Mertens and Neymann [21, p. 1567] found that Borel determinacy can be used to show existence of ϵ -NE in infinite games with bounded Borel-measurable real-valued payoffs. By generalizing their technique, an abstract result about point-classes and determinacy [14] implies a characterization of the preferences that always yield games with NE, in games with (quasi-)Borel-measurable outcome functions, countably many outcomes, and an arbitrary cardinality of players: the preferences whose inverses are well-founded. Then it was shown [16] that two-player antagonistic games (*i.e.* abstract zero-sum games) with finitely many outcomes and (quasi-)Borel-measurable outcome function have SPE.

When the outcome function is a continuous real-valued payoff function, Fudenberg and Levine [8] showed that there is always an SPE in multi-player games. Similar results were obtained recently via an abstract and uniform density argument [16]. The continuity assumption may be slightly relaxed if one is willing to accept approximate SPE. Indeed existence of ϵ -SPE was proved for lower-semicontinuous [6] and upper-semicontinuous [22] payoffs.

However, when the real-valued payoff function is discontinuous enough and the preferences are not antagonistic, there may be no (ϵ -)SPE, as in the following example which is similar to [25, Example 3]. Let a and b be two players with preferences $z <_a y <_a x$ and $x <_b z <_b y$. They are alternatively given the possibility to stop and yield outcomes y and z , respectively, but the outcome is x if no one ever stops.



In addition, a real-valued two-player game [7] was recently designed with the following features: it has a similar preference pattern as in the example above, and it has no ϵ -SPE for small enough ϵ even when the players are allowed to use mixed strategies at every node of the game tree. All this shows that Mertens [21] was right when writing that "Subgame perfectness is a completely different issue" from NE. This article solves the issue partially, and the contribution is two-fold. First, it characterizes the linear preferences that always yield SPE in games with outcome functions in the Hausdorff difference hierarchy of the open sets: the preferences that are void of infinite ascending chains (of course) and of the *SPE killer* from the example above. Said otherwise, if $\neg(z <_a y <_a x \wedge x <_b z <_b y)$ holds for all players a and b and all outcomes x , y and z of a multi-player game with outcome function in the difference hierarchy (and preferences without infinite ascending chains), the game has an SPE, and even a *global-Pareto* one, as in Definition 7. Second contribution, the characterization still holds for two-player games with strict weak order preferences, but no longer for three-player games. (Strict weak orders are important since they are an abstraction of the usual preference order over the real-valued payoff functions.)

Section 2 consists of definitions; Section 3 proves the characterization for many players and linear preferences; Section 4 proves the characterization for two players and strict weak order preferences.

Related works Characterizing the preferences that guarantee existence of NE in interesting classes of games is not a new idea: earlier than the two examples above ([13], [14]), Gimbert and Zielonka [10] "characterise the family of payoff mappings for which there always exist optimal positional strategies for both players" in some win-lose games played on finite graphs. Also, [15] characterizes the preferences that guarantee existence of NE in determined, countable two-player game forms.

The notion of SPE has also been studied in connection with system verification and synthesis, at low levels of the Borel hierarchy: in [26] with qualitative objectives, in [2] with quantitative objective for reachability, and in [3] with quantitative objectives and a weak variant of SPE.

Finally, some specific infinite games in extensive form (such as the dollar auction) and especially their SPE have been studied using co-algebraic methods in [18] and [1].

2 Technical background

The games in this article are built on infinite trees, which may be defined as prefix-closed sets of finite sequences. The elements of a tree are called nodes. Intuitively, a node represents both a position in the tree and the only path from the root to this position.

Definition 1 (Tree). Let Σ be a set.

- Σ^* (Σ^ω) is the set of finite (infinite) sequences over Σ , and a tree over Σ is a subset T of Σ^* such that $\gamma \sqsubseteq \delta$ (i.e. γ is a prefix of δ) and $\delta \in T$ implies $\gamma \in T$.
- For a node γ in a tree T , let $\text{succ}(T, \gamma) := \{\delta \in T \mid \gamma \sqsubseteq \delta \wedge |\delta| = |\gamma| + 1\}$, where $|\gamma|$ is the length of γ .
- A tree T is pruned if $\text{succ}(T, \gamma) \neq \emptyset$ for all $\gamma \in T$.
- Let T be a tree over Σ . The set $[T]$ is made of the infinite paths of T , namely the elements of Σ^ω whose every finite prefix is in T .
- Let T be a tree over Σ . For $\gamma \in T$ let $T_\gamma := \{\delta \in \Sigma^* \mid \gamma\delta \in T\}$.

In this article the outcomes of a game correspond to some partition of the infinite paths of the game tree, and the subsets of the partition are restricted to the Hausdorff difference hierarchy of the open sets. This hierarchy is defined in, e.g., [11, 22.E], but a probably-folklore result [17, Section 2.4] gives an equivalent presentation, which is in turn rephrased in Definition 4 below. These new definitions facilitate the proofs by induction in this article. Then, the Hausdorff-Kuratowski theorem (see, e.g., [11, Theorem 22.27]) implies that, in the Baire space (i.e. $\mathbb{N}^\mathbb{N}$), the difference hierarchy is equal to Δ_2^0 , the sets that are both countable unions of closed sets and countable intersections of open sets. This equality tells how low the difference hierarchy lies in the Borel hierarchy.

For every pruned tree T , let $\{\gamma[T_\gamma] \mid \gamma \in T\}$ be a basis for the open subsets of $[T]$. Definition 2 below is a special case of a more general definition that can be found, e.g., in [11, 22.E].

Definition 2 (Difference hierarchy). Let T be a pruned tree, $\theta > 0$ be a countable ordinal, and $(A_\eta)_{\eta < \theta}$ be an increasing sequence of open subsets of $[T]$. $D_\theta((A_\eta)_{\eta < \theta})$ is defined as below.

$x \in D_\theta((A_\eta)_{\eta < \theta}) :\Leftrightarrow x \in \cup_{\eta < \theta} A_\eta$ and the least $\eta < \theta$ with $x \in A_\eta$ has parity opposite to that of θ .

And $D_\theta([T]) := \{D_\theta((A_\eta)_{\eta < \theta}) \mid \forall \eta < \theta, A_\eta \text{ is an open subset of } [T]\}$.

Observation 3 $D_{\theta+1}((A_\eta)_{\eta < \theta+1}) = A_\theta \setminus D_\theta((A_\eta)_{\eta < \theta})$

Definition 4, Lemma 1, and Proposition 1 relate to [17, Section 2.4].

Definition 4 (Quasi-difference sets). Let T be a pruned tree. The set $\mathcal{D}([T])$ is defined by transfinite induction below.

- Every open set of $[T]$ is in $\mathcal{D}([T])$.
- Let $(\gamma_i)_{i \in I}$ be pairwise non-comparable nodes in T , and for all $i \in I$ let $D_i \in \mathcal{D}([T])$ be such that $D_i \subseteq \gamma_i[T_{\gamma_i}]$. Then $\cup_{i \in I} D_i \in \mathcal{D}([T])$.
- The complement in $[T]$ of a set in $\mathcal{D}([T])$ is also in $\mathcal{D}([T])$.

One can prove Observation 5 below by induction along Definition 4.

Observation 5 Given a node γ of a pruned tree T , and $A \in \mathcal{D}([T])$, $\gamma[T_\gamma] \cap A \in \mathcal{D}([T])$.

Lemma 1. $D_\theta \subseteq \mathcal{D}([T])$ for all non-zero countable ordinal θ and all pruned tree T .

Proof. By induction on θ , which holds for $\theta = 1$ since $D_1([T])$ is made of the open subsets of $[T]$. Let $\theta > 1$ be an ordinal and let $A \in D_\theta$, so $A = D_\theta((A_\eta)_{\eta < \theta})$ for some family $(A_\eta)_{\eta < \theta}$ of open subsets of $[T]$ that is increasing for the inclusion. Every A_η is open so it can be written $\cup_{i \in I} \gamma_{\eta,i}[T_{\gamma_{\eta,i}}]$ where $\gamma_{\eta,i}$ and $\gamma_{\eta,j}$ are not proper prefixes of one another (but are possibly equal) for all $\eta < \theta$ and $i, j \in I$. We can also require some minimality for all $\gamma_{\eta,i}$, more specifically $\gamma \sqsubset \gamma_{\eta,i} \Rightarrow \neg(\gamma[T_\gamma] \subseteq A_\eta)$. Let F consists of the minimal prefixes among $\{\gamma_{\eta,i} \mid i \in I \wedge \eta < \theta\}$, and for all $\gamma \in F$ let $f(\gamma)$ be the least η such that $\gamma = \gamma_{\eta,i}$ for some $i \in I$. So $f(\gamma) < \theta$ for all $\gamma \in F$, and $A \subseteq \cup_{\eta < \theta} A_\eta = \cup_{\gamma \in F} \gamma[T_\gamma]$, so $A = \cup_{\gamma \in F} A \cap \gamma[T_\gamma]$. Let $\gamma \in F$ and let us make a case disjunction to show that $A \cap \gamma[T_\gamma] \in \mathcal{D}([T])$. First case, $f(\gamma)$ and θ have the same parity, so $A \cap \gamma[T_\gamma] = D_{f(\gamma)}((A_\eta \cap \gamma[T_\gamma])_{\eta < f(\gamma)}) \in D_{f(\gamma)}([T])$, so $A \cap \gamma[T_\gamma] \in \mathcal{D}([T])$ by induction hypothesis. Second case, $f(\gamma)$ and θ have opposite parity, so

$$\begin{aligned} A \cap \gamma[T_\gamma] &= D_{f(\gamma)+1}((A_\eta \cap \gamma[T_\gamma])_{\eta < f(\gamma)+1}) \text{ by Definition 2,} \\ &= (A_{f(\gamma)} \cap \gamma[T_\gamma]) \setminus D_{f(\gamma)}((A_\eta \cap \gamma[T_\gamma])_{\eta < f(\gamma)}) \text{ by Observation 3,} \\ &= \gamma[T_\gamma] \setminus D_{f(\gamma)}((A_\eta \cap \gamma[T_\gamma])_{\eta < f(\gamma)}) \text{ by definition of } F \ni \gamma \text{ and } f(\gamma), \\ &\in \mathcal{D}([T]) \text{ by induction hypothesis and Observation 5.} \end{aligned}$$

Therefore $A \in \mathcal{D}([T])$.

Proposition 1 below shows that the quasi-difference sets coincide with the sets in the difference hierarchy for countable trees, just like quasi-Borel sets [20] and Borel sets coincide on Polish spaces.

Proposition 1. $\mathcal{D}([T]) = \cup_{\theta < \omega_1} D_\theta([T])$ for all countable pruned tree T .

Proof. $\cup_{\theta < \omega_1} D_\theta \subseteq \mathcal{D}([T])$ was already proved in Lemma 1, so let $A \in \mathcal{D}([T])$. Let us prove that $A \in D_\theta([T])$ for some θ , by induction on the definition of $\mathcal{D}([T])$. Base case, if A is open, $A \in D_1([T])$.

Second case, let $(\gamma_i)_{i \in I}$ be pairwise non-comparable nodes in T , and let $A_i \in \mathcal{D}([T]) \cap \gamma_i[T_{\gamma_i}]$ for all $i \in I$, such that $A = \cup_{i \in I} A_i$. By induction hypothesis $A_i \in D_{\theta_i}$ for some $\theta_i < \omega_1$. Let $\theta := \sup_{i \in I} \theta_i$, so $\theta < \omega_1$ by countability of T , and $A_i \in D_\theta$ for all $i \in I$. For all $i \in I$ let $(A_{i,\eta})_{\eta < \theta}$ be an increasing sequence of open sets such that $A_i = D_\theta((A_{i,\eta})_{\eta < \theta})$. So $A = \cup_{i \in I} D_\theta((A_{i,\eta})_{\eta < \theta}) = D_\theta((\cup_{i \in I} A_{i,\eta})_{\eta < \theta})$, which shows that $A \in D_\theta([T])$.

Third case, A is the complement of $B \in D_\theta([T])$ for some $\theta < \omega_1$. So $B = D_\theta((B_\eta)_{\eta < \theta})$ for some increasing sequence of open sets of $[T]$. Let $B_\theta = [T]$, so $A = D_{\theta+1}((B_\eta)_{\eta < \theta+1})$.

Informally, a play of a game starts at the root of an infinite game tree and at each stage of the game the unique owner of the current node chooses a child of the node. The articles [14] and [16] used game trees of the form C^* because it was more convenient and done without much loss of generality, but this article

works with general pruned trees because they will be cut in a non-uniform way. Moreover pruned trees are general enough, since leaves in infinite games can be simulated by the pseudo-leaves defined below.

Definition 6 (Game, subgame, pseudo-leaf). *An infinite game g is a tuple $\langle A, T, d, O, v, (\prec_a)_{a \in A} \rangle$ complying with the following.*

- A is a non-empty set (of players).
- T is a non-empty pruned tree (of possible finite plays).
- $d : T \rightarrow A$ (assigns a decision maker to each stage of the game).
- O is a non-empty set (of possible outcomes of the game).
- $v : [T] \rightarrow O$ (uses outcomes to value the infinite plays in the tree).
- Each \prec_a is a binary relation over O (modelling the preference of player a).

For $\gamma \in T$, the subgame g_γ is defined by $d_\gamma : T_\gamma \rightarrow A$ such that $d_\gamma(\delta) := d(\gamma\delta)$ and by $v_\gamma : [T_\gamma] \rightarrow O$ such that $v_\gamma(p) := v(\gamma p)$. A x -pseudo-leaf of g is a shortest node $\gamma \in T$ such that only the outcome x occurs in g_γ .

Definition 7 (strategy profile, induced play, global-Pareto equilibrium).

Let $g = \langle A, T, d, O, v, (\prec_a)_{a \in A} \rangle$ be a game.

- A strategy profile is a function $s : T \rightarrow T$ such that $s(\gamma) \in \text{succ}(T, \gamma)$ for all $\gamma \in T$. Let S_g be the set of the strategy profiles for g . For $\gamma \in T$ and $s \in S_g$ the subprofile $s_\gamma : T_\gamma \rightarrow T_\gamma$ is defined by the equality $\gamma s_\gamma(\delta) = s(\gamma\delta)$.
- For $\gamma \in T$ and $s \in S_g$, the play $p = p^\gamma(s)$ induced by s at γ is defined inductively by $p_0 \dots p_{|\gamma|-1} := \gamma$ and $p_n := s(p_0 \dots p_{n-1})$ for all $n > |\gamma|$.
- A Nash equilibrium is a profile $s \in S_g$ such that

$$NE_g(s) := \forall s' \in S_g, \forall a \in A, \neg(v \circ p^\epsilon(s) \prec_a v \circ p^\epsilon(s') \wedge (\forall \gamma \in T, s(\gamma) \neq s'(\gamma) \Rightarrow d(\gamma) = a))$$

A subgame perfect equilibrium is a profile $s \in S_g$ such that $NE_{g_\gamma}(s_\gamma)$ for all $\gamma \in T$.

- Let $O' \subseteq O$. One says that $x \in O'$ is Pareto-optimal in O' if for all $y \in O'$ and $a \in A$ such that $x \prec_a y$ there exists $b \in B$ such that $y \prec_b x$. A global-Pareto Nash equilibrium (GP-NE) is an NE whose induced outcome is Pareto-optimal in the outcomes occurring in the underlying game. A GP-SPE is a profile that induces a GP-NE in every subgame.

The proofs in this article do not build SPE by mere backward induction, but more generally by recursively refining rational behavioral promises. At each stage the refinement is optimal given the existing promises and regardless of the future ones. Since a promise not to choose a specific successor of a given node cannot be represent by a strategy profile, the more general notion of quasi-profile is defined below.

Definition 8 (quasi profile). Let $g = \langle A, T, d, O, v, (\prec_a)_{a \in A} \rangle$ be a game.

- A quasi profile is a multivalued function $q : T \multimap T$ such that $\emptyset \neq q(\gamma) \subseteq \text{succ}(\gamma)$ for all $\gamma \in T$. Let Q_g be the set of the quasi profiles for g . For $\gamma \in T$ and $q \in Q_g$ the sub-quasi-profile $q_\gamma : T_\gamma \multimap T_\gamma$ is defined by the equality $\gamma q_\gamma(\delta) = q(\gamma\delta)$.

- For $\gamma \in T$ and $q \in Q_g$, the tree induced by q starting at γ is defined inductively by $\epsilon \in T_\gamma(q)$, where ϵ is the empty sequence, and $\delta \in T_\gamma(q) \Rightarrow q_\gamma(\delta) \subseteq T_\gamma(q)$.
- Let $q \in Q_g$. Let $(\gamma_i)_{i \in I}$ be the nodes of T such that $\gamma_i \notin q(\gamma)$ for all $\gamma \in T$, and let $G(g, q) := \{g_{\gamma_i} \mid_{T_{\gamma_i}(q)}\}_{i \in I}$.

Making a promise in a game g by defining a quasi-profile q splits the game into "smaller" games, formally *via* $G(g, q)$. If the promise is rational, these "smaller" games can be processed independently since gluing any of their respective SPE will yield an SPE for g . Towards this, Observation 9 below suggests that the recursive refinement will lead to a fully defined strategy profile of g , if performed a sufficiently great (ordinal) number of times.

Observation 9 *Let g be a game on a tree T , let q be a quasi profile for g , and let $G(g, q) = \{g_{\gamma_i} \mid_{T_{\gamma_i}(q)}\}_{i \in I}$. Then $\{\gamma_i T_{\gamma_i}(q)\}_{i \in I}$ is a partition of T .*

3 Many players with linearly ordered preferences

This section characterizes the linear preferences that always yield SPE in games with \mathcal{D}_{ω_1} -measurable outcome functions: the families of preferences without infinite ascending chains and without the SPE killer, *i.e.* the pattern $z <_a y <_a x \wedge x <_b z <_b y$ for some players a and b and outcomes x, y, z . The main difficulty is tackled by Lemma 2 and corollary 1 below. It consists in slightly generalizing an existing result [16] stating that two-player Borel games with antagonist preferences have SPE by considering preferences that are almost antagonist, but in addition there is an outcome y that is the worst one for both players, and the set of plays with outcome y is a closed set (union an open set). This is then generalized for a set in \mathcal{D}_{ω_1} by induction, and eventually to multi-player games without the SPE killer thanks to a combinatorial result.

Lemma 2. *Let a game involve two players a and b , preferences $y <_a x_n <_a \dots <_a x_1$ and $y <_b x_1 <_b \dots <_b x_n$ for some n , such that all plays without pseudo-leaves have outcome y . The game has a global-Pareto subgame perfect equilibrium.*

Proof. Let us consider only infinite games involving two players a and b , preferences $y <_a x_n <_a \dots <_a x_1$ and $y <_b x_1 <_b \dots <_b x_n$ for some n . Let us call a game *weak-stop* if every play without pseudo-leaves has outcome y , and *strong-stop* if in addition every node that does not lie on a play with outcome y has a prefix that is a pseudo-leaf. Note the following: in every weak-stop game the plays with outcome y form a closed set; for every quasi profile q for a weak-stop game g , the set $G(g, q)$ contains only weak-stop games; and modifying the outcome function of a weak-stop game such that it is constant on given subgames yields a weak-stop game. (But the same does not hold for strong-stop games.) Let us call a node of a strong-stop game an a_k (b_k) node if it is owned by player a (b), and if it is the parent of a x_k -pseudo-leaf. Let us call every a_k (b_k) node

a a -stop (b -stop) node, and furthermore let us call every a -stop or b -stop node a stop node.

Let us prove the claim by induction on n , which holds for $n = 0$ and $n = 1$, so let us assume that $1 < n$. Five transformations on games are defined below, and they are meant to be applied recursively to a weak-stop game.

1. "Weak-stop towards strong-stop": Let g be a weak-stop game on tree T . Let γ be a node such that g_γ involves more than one outcome but not y . By construction g_γ is an antagonist game, and it amounts, when seeing a pseudo-leaf as a leaf, to a game without infinite plays, but possibly without uniform bound on the length of the plays. By [5] it has a GP-SPE s_γ nonetheless, which induces some x_k . Let us derive a weak-stop game g' from g by modification of the outcome function: for all $p \in [T]$ let $v'(p) := x_k$ if $\gamma \sqsubseteq p$ and $v'(p) := v(p)$ otherwise. Pasting s_γ at node γ on a GP-SPE s' for g' yields a GP-SPE s for g . Formally $s(\gamma\delta) := s_\gamma(\delta)$ for all $\delta \in T_\gamma$ and $s(\delta) := s'(\delta)$ for all $\delta \in T$ such that $\gamma \not\sqsubseteq \delta$.
2. "Emptying the interior of y ": Let g be a strong-stop game on tree T and let γ be a y -pseudo-leaf. If $\gamma = \epsilon$ all profiles for g are GP-SPE. If γ is not the root of g let us define a quasi profile q for g by letting the owner of the parent of γ ignore γ . Formally, $q(\delta) := \text{succ}(T, \delta) \setminus \{\gamma\}$ for all $\delta \in T$. Since y is the worst outcome for both players, none will have an incentive to deviate from this promise, regardless of the future choices at the other nodes. $G(g, q)$ contains g_γ and a weak-stop game g' . Combining any profile s_γ for g_γ and a GP-SPE for g' yields a GP-SPE for g .
3. " b chooses x_n ": Let g be a strong-stop game on tree T , let γ be a b_n node and let $\delta \in \text{succ}(T, \gamma)$ be an x_n -pseudo-leaf. Let us define a quasi profile q for g by letting b choose δ at γ . Formally, $q(\gamma) := \{\delta\}$ and $q(\alpha) := \text{succ}(T, \alpha)$ for all $\alpha \in T \setminus \{\gamma\}$. Since x_n is b 's preferred outcome, she will have no incentive to deviate from this choice, regardless of the choices at the other nodes. So, finding a GP-SPE for every weak-stop game in $G(g, q)$ will complete the definition of a GP-SPE for g .
4. " a ignores x_n ": Let g be a strong-stop game. Let γ be a node in g such that g_γ involves outcome y but no b -stop nodes, and such that every subgame of g_γ involving outcome y has an a_k node for some $k < n$. Let us define a quasi profile q for g by letting a ignore all x_n -pseudo-leaves at all a_n nodes below γ . The set $G(g, q)$ is made of games involving only outcome x_n and of one g' such that g'_γ does not involve x_n . By induction hypothesis g'_γ has an GP-SPE s'_γ , which induces some x_k . Let us define g'' by modification of the outcome function of g : for all $p \in \gamma[T_\gamma]$ let $v''(p) := x_k$ and for all $p \in [T] \setminus \gamma[T_\gamma]$ let $v''(p) := v(p)$. Pasting s'_γ on a GP-SPE s'' for g'' yields a GP-SPE for g .
5. " a chooses x_n ": Let g be a strong-stop game. Let γ be such that g_γ involves outcome y and every subgame of g_γ involving y has some a_n nodes but no a_k node for all $k < n$. To build for g_γ a GP-SPE s_γ inducing x_n on all of its subprofiles, it suffices, first, to choose arbitrary profiles for the subgames rooted at the x_n -pseudo-leaves of g_γ , and second, to fix consistently paths from each node to an x_n -pseudo-leaf. This second step can be done by letting

player a choose an x_n -pseudo-leaf at some node, which defines a quasi profile q , and by repeating it recursively for the games in $G(g_\gamma, q)$. This s_γ is a GP-SPE because every subprofile induces outcome x_n , which is b 's preferred outcome, and the only alternative for a in every subprofile is outcome y since all a -stop nodes are a_n nodes. Let us define g' by modification of the outcome function of g : for all $p \in \gamma[T_\gamma]$ let $v'(p) := x_n$ and for all $p \in [T] \setminus \gamma[T_\gamma]$ let $v'(p) := v(p)$. Pasting s_γ on a GP-SPE s' for g' yields a GP-SPE for g .

Given a weak-stop game g , let us apply to it the five transformations above, sequentially, non-deterministically whenever they are applicable, and until none of them is applicable, *i.e.* possibly an ordinal number of times. This yields a set G of strong-stop games (otherwise Transformation 1 could be applied) whose subgames that involve outcome y all have stop nodes (otherwise Transformation 2 could be applied), without b_n nodes (otherwise Transformation 3 could be applied), such that every subgame that involves y but no b -stop nodes has a subgame without a_k nodes for all $k < n$ (otherwise Transformation 4 could be applied), and such that every subgame h' of every game in G has the following property (otherwise Transformation 5 could be applied): if every subgame of h' has a a -stop node, h' has an a_k node for some $k < n$.

Let h' be a subgame of $h \in G$, and that involves y . If h' has no b -stop nodes, combining the properties above shows that all of its subgames have a -stop nodes, so one of them has only a_n nodes, contradiction, so every subgame of h that involves y has a b -stop node.

For every $h \in G$ let us define the quasi profile q by letting a ignore all the x_n -pseudo-leaves. $G(h, q)$ is made of games involving the outcome x_n only and of one h' void of x_n . Since every subgame of h that involves y has a b -stop node, it also holds for h' . By induction hypothesis, h' has a GP-SPE, which is easily extended to a GP-SPE for h , thus completing the definition of a GP-SPE for the original g .

Corollary 1. *Given a game g with two players a and b , a quasi-Borel measurable outcome function, and preferences $y <_a x_n <_a \dots <_a x_1$ and $y <_b x_1 <_b \dots <_b x_n$ for some n . If the plays with outcome y form the union of an open set and a closed set, the game has a global-Pareto subgame perfect equilibrium.*

Proof. Let the plays with outcome y be the union $Y = Y_o \cup Y_c$ of an open set and a closed set. Wlog Y_o and Y_c are disjoint. Let us derive g' from g by removing the plays in Y_o . So the plays of g' that do not yield outcome y form an open set, *i.e.* a disjoint union of clopen balls with defined by the prefixes $(\gamma_i)_{i \in I}$. Every game g_{γ_i} is antagonist and quasi-Borel, so it has an SPE s_i by [16]. Let us define g'' by modification of the outcome function of g' : every play going through γ_i yields the outcome induced by s_i . This g'' has a GP-SPE s'' by Lemma 2, and together with the s_i it can be used to build a GP-SPE for g .

Let us extend Corollary 1 from open union closed to the difference hierarchy.

Lemma 3. *Let be a game on a tree T , with two players a and b and preferences $y <_a x_n <_a \dots <_a x_1$ and $y <_b x_1 <_b \dots <_b x_n$, and let us assume that each set*

of plays with outcome x_i is quasi-Borel and that Y the set of plays with outcome y is in $\mathcal{D}_{\omega_1}([T])$. Then the game has a global-Pareto subgame perfect equilibrium.

Proof. By transfinite induction on the level of Y in the difference hierarchy. The base case where Y is open or closed is solved by Corollary 1.

For the inductive case, let us make a case disjunction depending on the last step of the construction of Y . For the union case, $Y = \cup_{i \in \mathbb{N}} \gamma_i Y_i$ for some Y_i that have lower levels than Y in the difference hierarchy, and where the γ_i are not prefixes of one another. By induction hypothesis each g_{γ_i} has a GP-SPE s_i inducing either y or an outcome $x_{k(i)}$. Let us start the construction of a profile s for g by fixing the s_i as the respective subprofiles for the $g|_{\gamma_i T_{\gamma_i}}$. Let us define g' by modification of the outcome function of g : let each play going through γ_i yield the outcome induced by s_i . This is a quasi-Borel game and the plays with outcome y form an open set, so it has a GP-SPE s' by Corollary 1, which we use to complete the definition of s . It is easy to check that s is a GP-SPE for g .

For the complementation case, $Y = [T] \setminus ([T] \setminus Y)$, where $[T] \setminus Y$ is equal to $\cup_{i \in \mathbb{N}} \gamma_i X_i$ for some X_i that have lower levels than $[T] \setminus Y$ (and Y) in the difference hierarchy, and where the γ_i are not prefixes of one another. Since all $Y \cap [\gamma_i T_{\gamma_i}] = [\gamma_i T_{\gamma_i}] \setminus \gamma_i X_i$ have lower levels than Y in the difference hierarchy, by induction hypothesis each g_{γ_i} has a GP-SPE s_i inducing either y or an outcome $x_{k(i)}$. Let us start the construction of a profile s for g by fixing the s_i as the respective subprofiles for the $g|_{\gamma_i T_{\gamma_i}}$. Let us define g' by modification of the outcome function of g : let each play going through γ_i yield the outcome induced by s_i . This is a quasi-Borel game and the plays with outcome y form the union of an open set and the closed set $[T] \setminus \cup_{i \in I} [\gamma_i T_{\gamma_i}]$, so it has a GP-SPE s' by Corollary 1, which we use to complete the definition of s . It is easy to check that s is a GP-SPE for g .

The combinatorial Lemma 4 below shows that the "local" absence of the SPE-killer amounts to a very simple "global" structure.

Lemma 4. *Let A be a non-empty set and for all $a \in A$ let $<_a$ be a strict linear order over some non-empty set O . The following assertions are equivalent.*

1. $\forall a, b \in A, \forall x, y, z \in O, \neg(z <_a y <_a x \wedge x <_b z <_b y)$.
2. *There exists a partition $\{O_i\}_{i \in I}$ of O and a linear order $<$ over I such that:*
 - (a) $i < j$ implies $x <_a y$ for all $a \in A$ and $x \in O_i$ and $y \in O_j$.
 - (b) $<_b|_{O_i} = <_a|_{O_i}$ or $<_b|_{O_i} = <_a|_{O_i}^{-1}$ for all $a, b \in A$.

If 1. and 2. hold, we may also assume that $<_b|_{O_i} = <_a|_{O_i}^{-1}$ is witnessed for all $i \in I$. Also, O_i is always a $<_a$ -interval for all $(i, a) \in I \times A$, as implied by 2a.

Proof. $2 \Rightarrow 1$ is straightforward, so let us assume 1. Let $x \sim y$ stand for $\exists a, b \in A, x \leq_a y \leq_b x$, which defines a reflexive and symmetric relation, and note that due to the SPE killer $z <_a y <_a x \wedge x <_b z <_b y$ the following holds: if $x <_a y$ and $y <_b x$, then $x <_a z <_a y$ iff $y <_b z <_b x$. To show that \sim is transitive too, let us assume that $x \sim y \sim z$. If x, y, z are not pairwise distinct, $x \sim z$ follows

directly, so let us assume that they are pairwise distinct, so by assumption there exist $a, b, c, d \in A$ such that $y <_a x <_b y <_c z <_d y$. To show that $x \sim z$ there are three cases depending on where z lies with respect to $y <_a x$, all cases invoking the (above-mentioned) forbidden-pattern argument: if $y <_a z <_a x$ then $x <_b z <_b y$, and $x \sim z$ follows; if $y <_a x <_a z$ then $z <_d x <_d y$ and subsequently $y <_c x <_c z$, by invoking twice the forbidden-pattern argument, and $x \sim z$ follows; third case, let us assume that $z <_a y <_a x$. If $x <_b z$ then $x \sim z$ follows, and if $z <_b x$ then $z <_b x <_b y$, so $y <_c x <_c z$, and $x \sim z$ follows. Therefore \sim is an equivalence relation; let $\{O_i\}_{i \in I}$ be the corresponding partition of O .

Now let us show that the \sim -classes are $<_a$ -intervals for all a , so let $x \sim y$ and $x <_a z <_a y$. By definition of \sim , there exists b such that $y <_b x$, in which case $y <_b z <_b x$ by the forbidden-pattern argument, so $x \sim z$ by definition.

Let $x \in O_i$ and $y \in O_j$ be such that $x <_a y$. Since O_i and O_j are intervals, $x' <_a y'$ for all $x' \in O_i$ and $y' \in O_j$. Since $\neg(x' \sim y')$ by assumption, $x' <_b y'$ for all $b \in A$, by definition of \sim . In this case defining $i < j$ meets the requirements.

Before proving 2b let us prove that if $x <_a y$ and $y <_b x$ and $z \sim y$, then $z <_a y$ iff $y <_b z$: this is trivial if z equals x or y , so let us assume that $x \neq z \neq y$, and also that $z <_a y$. If $z <_b y$, then $y <_c z$ for some c since $z \sim y$, and wherever x may lie with respect to $y <_c z$, it always yields a SPE killer using $<_a$ or $<_b$, so $y <_b z$. The converse is similar, it follows actually from the application of this partial result using $<_b^{-1}$ and $<_a^{-1}$ instead of $<_a$ and $<_b$.

Now assume that $<_b|_{O_i} \neq <_a|_{O_i}$ for some O_i , so $x <_a y$ and $y <_b x$ for some $x, y \in O_i$. Let $z, t \in O_i$. By the claim just above $z <_a y$ iff $y <_b z$, so by the same claim again $z <_a t$ iff $t <_b z$, which shows that $<_b|_{O_i} = <_a|_{O_i}^{-1}$. This proves the equivalence.

Finally, let us assume that the assertions hold. By definition of \sim , if O_i is not a singleton, $x <_a y$ and $y <_b x$ for some $a, b \in A$ and $x, y \in O_i$, so $<_b|_{O_i} = <_a|_{O_i}^{-1}$ is witnessed.

Theorem 1 extends Lemma 3 to many players and more complex preferences.

Theorem 1. *Let g be a quasi-Borel game with players in A , outcomes in O , and linear preferences $<_a$ for all $a \in A$. Let us assume that the inverses of the $<_a$ are well-ordered, and that there exists a partition $\{O_i\}_{i \in I}$ of O and a linear order $<$ of I such that:*

- $i < j$ implies $x <_a y$ for all $a \in A$ and $x \in O_i$ and $y \in O_j$.
- $<_b|_{O_i} = <_a|_{O_i}$ or $<_b|_{O_i} = <_a|_{O_i}^{-1}$ for all $a, b \in A$ and $i \in I$.

Let us further assume that for all $i \in I$ the plays with outcome in $\cup_{j < i} O_j$ form a \mathcal{D}_{ω_1} set. Then g has a global-Pareto subgame perfect equilibrium.


Proof. By Lemma 4 let us further assume wlog that $<_b|_{O_j} = <_a|_{O_j}^{-1}$ is witnessed for all $j \in I$, so the O_j are finite by well-ordering. Let us build a GP-SPE for g as the limit of a recursive procedure: Let O_i be such that some outcome of O_i occurs in g and such that for all $j > i$ no outcome from O_j occurs in g . Let

$(\gamma_k)_{k \in K}$ be the shortest nodes of g such that the outcomes occurring in g_{γ_k} are in $\cup_{j < i} O_j$ only. Let us define a quasi profile q for g by having the γ_k ignored by their parents. $G(g, q)$ consists of the g_{γ_k} and of a game g' . By Lemma 3 there is a GP-SPE s' for g' . (To see this, replace $\cup_{j < i} O_j$ with one single fresh outcome y , i.e. $y \notin O$ and set $y <_a x$ for all $a \in A$ and $x \in O_i$.) Combining s' with GP-SPE for the g_{γ_k} (obtained recursively) yields a GP-SPE for g , since the choices made in g' hold regardless of the choices made in the g_{γ_k} .

Corollary 2. *Let A and O be non-empty finite sets (of players and of outcomes) and for all $a \in A$ let $<_a$ be a linear preference. The following are equivalent.*

1. $\forall a, b \in A, \forall x, y, z \in O, \neg(z <_a y <_a x \wedge x <_b z <_b y)$.
2. Every \mathcal{D}_{ω_1} -Gale-Stewart game using A, O and the $<_a$ has a GP-SPE.

Proof. For 1. \Rightarrow 2. invoke Lemma 4 and Theorem 1, and prove 2. \Rightarrow 1. by contraposition with the following folklore example which is detailed, e.g., in [16].



Corollary 2 and the results that lead to it considers linear preference only. Proposition 2 below show that this restriction incurs a loss of generality, which is partly solved in Section 4.

Proposition 2. *Let us define two binary relations by $z, t \prec_a x, y$ and $y \prec_b z \prec_b x \prec_b t$.*

1. \mathcal{D}_{ω_1} infinite games with players a and b and preferences \prec_a and \prec_b have SPE.
2. The SPE killer occurs in any strict linear extensions of \prec_a and \prec_b .

Proof. 1. follows Theorem 2. For 2. let \prec'_a be a linear extension of \prec_a . If $x \prec'_a y$ then $z \prec'_a x \prec'_a y$ and $y \prec_b z \prec_b x$. If $y \prec'_a x$ then $t \prec'_a y \prec'_a x$ and $y \prec_b x \prec_b t$.

4 Two players with strict weak order preferences

The preferences considered in Proposition 2 are strict weak orders. Informally, strict weak orders are strict partial orders that can be seen as strict linear orders from afar, i.e. up to an equivalence relation. Traditionally in game theory the outcomes are real-valued payoff functions $f, g : A \rightarrow \mathbb{R}$ and the preferences are defined by $f \prec_a g$ iff $f(a) < g(a)$. These preferences are not strict linear orders but they are strict weak orders, so the results from Section 3 are worth generalizing. Strict weak orders are defined below.

Definition 10 (Strict weak order). *A strict weak order is a strict partial order whose complement is transitive, i.e. it satisfies $\neg(x \prec x)$ and $x \prec y \wedge y \prec z \Rightarrow x \prec z$ and $\neg(x \prec y) \wedge \neg(y \prec z) \Rightarrow \neg(x \prec z)$.*

Lemma 4 above describes the structure of strict linear orders void of the SPE killer. A similar result for strict weak orders will be useful. Lemma 5 below is part to it, and the other part appears directly in the proof of Lemma 6.

Lemma 5. *Let \prec_a and \prec_b be two strict weak orders over some finite O .*

1. *If \prec_a and \prec_b are void of the SPE killer, if there exists a \prec_a -non-extremal element and a \prec_b -non-extremal element, and if there is no partition $\{O_u, O_l\}$ of O such that $\neg(x \prec_a y)$ and $\neg(x \prec_b y)$ for all $(x, y) \in O_u \times O_l$, then $\prec_a \cap \prec_b = \emptyset$.*
2. *If $\prec_a \cap \prec_b = \emptyset$, there exists a linear extension $<$ of \prec_a such that $\prec_b \subseteq <^{-1}$.*

Proof. 1. Let x be \prec_b -minimal among the \prec_a -maximal elements, and let y be \prec_a -minimal among the \prec_b -maximal elements, so $x \neq y$ by the partition assumption. Towards a contradiction let us assume that, e.g., x is not \prec_b -minimal, and let z be \prec_b -minimal. For all $t \prec_b x$, it follows that t is not \prec_a -maximal by definition of x , and $\neg(y \prec_a t)$ by absence of the SPE killer. So y is not \prec_a -minimal, otherwise z is also \prec_a -minimal, thus contradicting the partition assumption. So likewise, for all $t \prec_a y$, it follows that t is not \prec_b -maximal by definition of y , and $\neg(x \prec_b t)$ by absence of the SPE killer. So the two-element partition induced by $\{t \in O \mid t \prec_a y \vee t \prec_b x\}$ contradicts the partition assumption. This shows that x is \prec_b -minimal and y is \prec_a -minimal. Towards a contradiction let us assume that $t(\prec_a \cap \prec_b)z$ for some $t, z \in O$. So $\{x, y\} \cap \{z, t\} = \emptyset$. By the partition assumption z is not both \prec_a and \prec_b -maximal, so, e.g., $z \prec_a x$, and t is not both \prec_a and \prec_b -minimal. By absence of the SPE killer t is \prec_b -minimal, so $y \prec_a t$ by the partition assumption., and subsequently z is \prec_b -maximal. By assumption there exists γ that is neither \prec_b -maximal nor \prec_b -minimal. Wherever γ lies wrt \prec_a , the SPE killer occurs.

2. *By induction on the cardinality of O , which holds for $|O| = 0$. Let x be \prec_b -minimal among the \prec_a -maximal elements, so x is also \prec_b -minimal since $\prec_a \cap \prec_b = \emptyset$. By induction hypothesis let $<_x$ witness the claim for $\prec_a|_{O \setminus \{x\}}$ and $\prec_b|_{O \setminus \{x\}}$. The linear order $< := <_x \cup \{(y, x) \mid y \in O \setminus \{x\}\}$ witnesses the claim.*

Lemma 6 below is a generalization of Corollary 1 from an order-theoretic point of view and a special case thereof from a topological point of view. Due lack of space and strong similarities with the proof of Lemma 2, the proof of Lemma 6 is in appendix.

Lemma 6. *Let g be a game with two players a and b , finitely many outcomes O , and strict weak order preferences void of the SPE killer. If each outcome corresponds to the union of an open set and a closed set, the game has an SPE.*

Furthermore, for every node γ of g let $\{O_1^\gamma, \dots, O_{n_\gamma}^\gamma\}$ be a partition of the outcomes of g_γ such that $\neg(x \prec_a y)$ and $\neg(x \prec_b y)$ for all $1 \leq k < n_\gamma$ and $(x, y) \in O_{k+1}^\gamma \times O_k^\gamma$. There exists an SPE for g such that the outcome induced at every node γ belongs to $O_{n_\gamma}^\gamma$.

Proof. By induction on the number of outcomes, which holds up to two outcomes. Let us make a case disjunction for the inductive case. First main case, there is no partition $\{O_u, O_l\}$ of O such that $\forall(x, y) \in O_u \times O_l, \neg(x \prec_a y) \wedge \neg(x \prec_b y)$. Let us make a nested case distinction. First nested case, there exists a \prec_a -non-extremal element and a \prec_b -non-extremal element, so by Lemma 5 let $<$ be a linear extension of \prec_a such that $\prec_b \subseteq <^{-1}$. By [16] the antagonist game with preference $<$ has an SPE, which is also an SPE for \prec_a and \prec_b .

Second nested case, one preference, e.g., \prec_a has only extremal elements. By the partition assumption let y be \prec_b -maximal and \prec_a -minimal. Let Y be the set of plays with outcome y , and let us define a quasi-profile q as follows. Let γ be the parent of a y -pseudo-leaf. If γ is owned by a , let a ignore the y -pseudo-leaves at γ ; otherwise let b choose a y -pseudo-leaf at γ . Let us apply this construction recursively (an ordinal number of times) to the games in $G(g, q)$ that do not involve y , until the Y' of each remaining game has empty interior. It is easy to check that Y' is also closed since Y is the union of an open set and a closed set, by assumption. So let $(\gamma_i)_{i \in I}$ be the shortest nodes that are not on any play with outcome y . The g_{γ_i} do not involve y , so by induction hypothesis they have suitable SPE s_i inducing some $x_i \in O$, which allow us to start the definition of a suitable SPE for g . Let us define g' by modification of the outcome function of g : let $v'(\gamma_i p) := x_i$ for all $i \in I$ and let $v'(p) := v(p)$ when $\gamma_i \not\sqsubseteq p$ for all $i \in I$. Let M_a be the \prec_a -maximal outcomes, and let us define a quasi-profile q' for g' as follows. Let γ be the parent of a M_a -pseudo-leaf, i.e. a shortest node involving outcome in M_a only. If γ is owned by a , let a choose a M_a -pseudo-leaf at γ ; otherwise let b ignore the M_a -pseudo-leave at γ . Let us apply this recursively to the games in $G(g', q)$ that involve y . In the remaining games a can deviate from a play with outcome y only to reach a $(O \setminus M_a)$ -pseudo-leaf, so every profile that follows plays with outcome y whenever possible is a suitable SPE.

Second main case, there exist partitions $\{O_u, O_l\}$ of O such that $\forall(x, y) \in O_u \times O_l, \neg(x \prec_a y) \wedge \neg(x \prec_b y)$. Among these partitions let us consider the one with the smallest possible O_u . Let us make a further case disjunction. First case, $|O_l| > 1$. As is now customary, let us start defining a suitable SPE for g by using the induction hypothesis on the maximal subgames involving only outcomes in O_l , and on the game derived from g by replacing outcomes in O_l with a fresh outcome y that is the new \prec_a and \prec_b -minimum.

Second case, $|O_l| = \{y\}$. By minimality of $|O_u|$, there is no partition $\{O_{uu}, O_{ul}\}$ of O_u such that $\forall(x, y) \in O_{uu} \times O_{ul}, \neg(x \prec_a y) \wedge \neg(x \prec_b y)$. Therefore the situation is reminiscent of the first main case above, but for $\prec_a|_{O_u}$ and $\prec_b|_{O_u}$ instead of \prec_a and \prec_b . In both nested cases from the first main case, there exists some x_n that is, e.g., $\prec_a|_{O_u}$ -minimal and $\prec_b|_{O_u}$ -maximal. Applying the proof of Lemma 2 almost verbatim yields a suitable SPE for g .

Theorem 2. Let g be a game with two players a and b , finitely many outcomes, a \mathcal{D}_{ω_1} -measurable outcome function, strict weak order preferences such that $\neg(z \prec_a y \prec_a x \wedge x \prec_b z \prec_b y)$ for all outcomes x, y and z . Then the game has an SPE.

Furthermore, for every node γ of g let $\{O_1^\gamma, \dots, O_{n_\gamma}^\gamma\}$ be a partition of the outcomes of g_γ such that $\neg(x \prec_a y)$ and $\neg(x \prec_b y)$ for all $1 \leq k < n_\gamma$ and $(x, y) \in O_{k+1}^\gamma \times O_k^\gamma$. There exists an SPE for g such that the outcome induced at every node γ belongs to $O_{n_\gamma}^\gamma$.


Proof. By induction on the levels in the difference hierarchy of the sets of plays corresponding to the outcomes. The base case holds by Lemma 6.

For the inductive case, let y be an outcome whose corresponding set Y has level more than one in the difference hierarchy, and let us make a case disjunction depending on the last step of the construction of Y . The remainder of the proof can be taken almost verbatim from the proof of Lemma 3, but by replacing "GP-SPE" with "suitable SPE", and by invoking Lemma 6 or the induction hypothesis instead of Corollary 1.

Theorem 2 considers two-player games only. Observation 11 shows that absence of the SPE killer is no longer a sufficient condition for a three-player game with strict weak order preferences to have an SPE.

Observation 11 *Let three players a, b and c have preferences $z \prec_a y \prec_a x$ and $t \prec_b z \prec_b y$ and $x \prec_c t \prec_c y$. (and, e.g., $y \sim_a t$, $z \sim_b x$, and $y \sim_c z$ or $x \sim_c z$)*

1. The SPE killer does not occur in the strict weak orders \prec_a and \prec_b and \prec_c .
2. The following game with \prec_a and \prec_b and \prec_c has no SPE.



Proof. For 2. Towards a contradiction let us assume that there exists an SPE for the game. Let us consider a node where player a chooses y . Then at the node right above it c chooses to continue to benefit from y , and at the node above b chooses to continue, too. The induced outcome at the node further above is y regardless of the choice of a , and so on up to the root.

Let us make a case disjunction: first case, there exists infinitely many nodes where a chooses y , so b and c always continue by the remark above, so a has an incentive to continue too, to induce outcome x , contradiction. Second case, there exists a node below which a always continues. From then on, one player must stop at some point, otherwise the outcome is x and c has an incentive to stop. The first player to stop cannot be b , otherwise a would stop before b , and it cannot be c , otherwise b would stop before c , contradiction.

Proposition 3 below shows that considering only strict weak orders incurs a loss of generality.

Proposition 3. *Let us define two binary relations by $\gamma \prec_a y \prec_a x$ and $z \prec_a \beta \prec_a \alpha$ and $x \prec_b z \prec_b y$ and $\alpha \prec_b \gamma \prec_b \beta$.*

1. The SPE killer occurs in every strict weak order extensions of \prec_a and \prec_b .

2. \mathcal{D}_{ω_1} -games with players a and b and preferences \prec_a and \prec_b have SPE.

Proof. 1. Let \prec'_a be a strict weak order extension of \prec_a . If $z \prec'_a y$, the SPE killer occurs, with x . If $\neg(z \prec'_a y)$, then $\gamma \prec'_a \beta$ and the SPE killer occurs, with α .

2. It suffices to prove the claim for games where each outcome set is the union of an open set and a closed set. (Then using a transfinite induction as in the proof of Lemma 3 will do.) The techniques from Lemmas 2 and 6 are suitable here, and used without details. If the outcome x does not occur in the game, note that $\prec_a|_{O \setminus \{x\}}$ and $\prec_b|_{O \setminus \{x\}}$ can be extended into the strict weak orders $z \prec'_a \gamma \sim'_a \beta \prec'_a y \sim'_a \alpha$ and $z \sim'_b \alpha \prec'_b \gamma \prec'_b y \sim'_b \beta$, respectively, and that the SPE killer is absent from these. So by Theorem 2 the game has an SPE wrt \prec'_a and \prec'_b and therefore also wrt \prec_a and \prec_b . (And likewise if the outcome α does not occur in the game.)

Now one can reduce the set for x to a closed set by letting a choose the clopen balls with constant outcome x and by letting b ignore them. The plays with outcomes different from x can be seen as belonging to a union of subgames without x , so by the remark above they have SPE. It allows us to replace these subgames with (pseudo)-leaves with outcomes the ones induced by the SPE. Now one can let a choose the pseudo-leaves with outcome α and b ignore them, which yields a game without outcome α . So by the remark above there is an SPE for the game.

Acknowledgements

I thank Vassilios Gregoriades and Arno Pauly for useful discussions. The author is supported by the ERC inVEST (279499) project.

References

1. Samson Abramsky and Viktor Winschel. Coalgebraic analysis of subgame-perfect equilibria in infinite games without discounting, 2012. arXiv preprint.
2. Thomas Brihaye, Véronique Bruyère, Julie De Pril, and Hugo Gimbert. On subgame perfection in quantitative reachability games. *Logical Methods in Computer Science*, 9, 2012.
3. Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-Francois Raskin. Weak subgame perfect equilibria and their application to quantitative reachability, 2015.
4. J. Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
5. Martín Escardó and Paulo Oliva. Selection functions, bar recursion and backward induction. *Mathematical Structures in Computer Science*, 20:127–168, 4 2010.
6. J. Fleisch, J. Kuipers, A. Mashiah-Yaakovi, G. Schoenmakers, E. Solan, and K. Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operations Research*, 35:742755, 2010.

7. János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eran Shmaya, Eilon Solan, and Koos Vrieze. Non-existence of subgame-perfect ε -equilibrium in perfect information games with infinite horizon. *International Journal of Game Theory*, 43(4):945–951, 2014.
8. Drew Fudenberg and David Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31(2):251 – 268, 1983.
9. D. Gale and F. M. Stewart. Infinite games with perfect information. *Annals of Math. Studies*, 28:245–266, 1953.
10. Hugo Gimbert and Wiesław Zielonka. Games where you can play optimally without any memory. In *CONCUR 2005 - Concurrency Theory*, volume 3653 of *Lecture Notes in Computer Science*, pages 428–442. Springer Berlin Heidelberg, 2005.
11. Alexander S. Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer, 1995.
12. Harold W. Kuhn. Extensive games and the problem of information. *Contributions to the Theory of Games II*, 1953.
13. Stéphane Le Roux. Acyclic preferences and existence of sequential Nash equilibria: a formal and constructive equivalence. In *TPHOLs, International Conference on Theorem Proving in Higher Order Logics*, Lecture Notes in Computer Science, pages 293–309. Springer, August 2009.
14. Stéphane Le Roux. Infinite sequential Nash equilibrium. *Logical Methods in Computer Science*, 9, 2013. Special Issue for the Conference "Computability and Complexity in Analysis" (CCA 2011).
15. Stéphane Le Roux. From winning strategy to Nash equilibrium. *Mathematical Logic Quarterly*, 2014. (to appear).
16. Stéphane Le Roux and Arno Pauly. Infinite sequential games with real-valued payoffs. In *Proceedings of LiCS*, 2014.
17. Stéphane Le Roux and Arno Pauly. Weihrauch degrees of finding equilibria in sequential games. In *Proceedings of CiE 2015*, 2015. (to appear).
18. Pierre Lescanne and Matthieu Perrinel. Backward coinduction, Nash equilibrium and the rationality of escalation. *Acta Informatica*, 49(3):117–137, 2012.
19. Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.
20. Donald A. Martin. An extension of Borel determinacy. *Annals of Pure and Applied Logic*, 49:279–293, 1990.
21. J-F Mertens. Repeated games. In *Proceedings of the international congress of mathematicians*, pages 1528–1577. American Mathematical Society, 1987.
22. R. Purves and W. Sudderth. Perfect information games with upper-semicontinuous payoffs. *Mathematics of Operations Research*, 36:468473, 2011.
23. Reinhard Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit. *Zeitschrift für die gesamte Staatswissenschaft*, 121, 1965.
24. J. Maynard Smith and G.R. Price. The logic of animal conflicts. *Nature*, 246:15–18, 1973.
25. Eilon Solan and Nicolas Vieille. Deterministic multi-player dynkin games. *Journal of Mathematical Economics*, 39(8):911 – 929, 2003.
26. Michael Ummels. Rational behaviour and strategy construction in infinite multi-player games. In S. Arun-Kumar and Naveen Garg, editors, *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science*, volume 4337 of *Lecture Notes in Computer Science*, pages 212–223. Springer Berlin Heidelberg, 2006.
27. Philip Wolfe. The strict determinateness of certain infinite games. *Pacific Journal of Mathematics*, 5:841–847, 1955.