

On the existence of weak subgame perfect equilibria

Véronique Bruyère¹, Stéphane Le Roux^{2,*,\dagger}, Arno Pauly^{2,*,\ddagger}, and Jean-François Raskin^{2,*}

¹ Département d'informatique, Université de Mons (UMONS), Belgium

² Département d'informatique, Université Libre de Bruxelles (ULB), Belgium

Abstract. We study multi-player turn-based games played on (potentially infinite) directed graphs. An outcome is assigned to every play of the game. Each player has a preference relation on the set of outcomes which allows him to compare plays. We focus on the recently introduced notion of weak subgame perfect equilibrium (weak SPE). This is a variant of the classical notion of SPE, where players who deviate can only use strategies deviating from their initial strategy in a finite number of histories. Having an SPE in a game implies having a weak SPE but the contrary is generally false.

We propose general conditions on the structure of the game graph and on the preference relations of the players that guarantee the existence of a weak SPE, that additionally is finite-memory. From this general result, we derive two large classes of games for which there always exists a weak SPE: (i) the games with a finite-range outcome function, and (ii) the games with a finite underlying graph and a prefix-independent outcome function. For the second class, we identify conditions on the preference relations that guarantee memoryless strategies for the weak SPE.

1 Introduction

Subgame-perfect equilibria (SPEs) are a core solution concept for sequential games. For infinite duration games, they only exist in restricted cases, though. A weaker variant of SPE, *weak SPE* was recently introduced in [5]. While an SPE must be resistant to any unilateral deviation of one player, a weak SPE must be resistant to such deviations where the deviating strategy differs from the original one on a *finite number* of histories only, or, equivalently, a *single* history. The latter class of deviating strategies is a well-known notion that for instance appears in the proof of Kuhn's theorem [19] with the one-step deviation property.

There are games for which there exists a weak SPE but no SPE [5,26]. The notion of weak SPE is important for several reasons (more details are given

* Author supported by ERC Starting Grant (279499: inVEST).

\dagger Le Roux has since moved to TU Darmstadt, Germany.

\ddagger Pauly has since moved to Swansea University, UK.

in the related work discussed below). First, for the large class of games with upper-semicontinuous payoff functions and for games played on finite trees, the notions of SPE and weak SPE are equivalent. Second, it is a central technical ingredient used to reason on SPEs as shown in [5] and [14]. Third, being immune to strategies that finitely deviate from the initial strategy profile may be sufficient from the perspective of synthesis (see more below).

In this paper, we provide the following contributions. First, we identify a general class of games played on potentially infinite graphs and prove that these games always admit weak SPE (Theorem 1). The proof of our result has an algorithmic flavour, and proceeds via transfinite induction. The weak SPEs we construct require only finite memory to execute, meaning that the prescribed action at any history depends only on the current vertex and on the state of some finite automaton. Second, starting from this general existence result, we prove the existence of a weak SPE:

- for games on infinite trees with a *finite* number of outcomes (Theorem 2, reproving a result from [14]);
- for games with a *finite* underlying graph and a *prefix-independent* outcome function (Theorem 4).

Additionally, in the second result, we identify conditions on the players’ outcome preferences that guarantee the existence of a weak SPE composed of *uniform memoryless* strategies only (Theorem 5).

Related work The concept of SPE has been first introduced and studied by the game theory community. In [19], Kuhn proves the existence of SPEs in games played on finite trees. This result has been generalized in several ways. All games with a continuous real-valued outcome function and a finitely branching tree always have an SPE [22] (the special case with finitely many players is first established in [16]). In [14] (resp. [24]), the authors prove that there always exists an SPE for games with a finite number of players and with a real-valued outcome function that is upper-semicontinuous (resp. lower-semicontinuous) and has finite range. The result of [24] is extended to an infinite number of players in [15]. In [22], it is proved using Borel determinacy that all two-player games with antagonistic preferences over finitely many outcomes and a Borel-measurable outcome function have an SPE. In [21], Le Roux shows that all games where the preferences over finitely many outcomes are free of some “bad pattern” and the outcome function is Δ_2^0 measurable (a low level in the Borel hierarchy) have an SPE.

In part of the aforementioned works, the equivalence between SPEs and weak SPEs is implicitly used as a proof technique: in a finite setting in [19], in a continuous setting in [16], and in a lower-semicontinuous setting in [14]. In the latter reference, it is implicitly proven that all games with a finite range real-valued outcome function always have a weak SPE (which appears to be an SPE when the outcome function is additionally lower-semicontinuous). We obtain this result here as a consequence of a more general theorem, with a proof of a more algorithmic nature.

The concept of SPE and other solution concepts for multi-player non zero-sum games have been considered recently by the theoretical computer community, see [2] for a survey. The existence of SPEs (and thus weak SPEs) is established in [27] for games played on graphs by a finite number of players and with Borel Boolean objectives. In [5], weak SPEs are introduced as a technical tool for showing the existence of SPEs in quantitative reachability games played on finite weighted graphs. An algorithm is also provided for the construction of a (finite-memory) weak SPE that appears to be an SPE for this particular class of games. In this paper, we give several existence results that are orthogonal to the results obtained in [5] as they are concerned with possibly infinite graphs or prefix-independent outcome functions.

Other refinements of Nash equilibria (NE) are studied. Let us mention the secure equilibria for two players first introduced in [9] and then used for reactive synthesis in [12]. These equilibria are generalized to multiple players in [13] or to quantitative objectives in [8], see also a variant called Doomsday equilibrium in [10]. Like NEs, they are subject to possible non-credible threats. Other alternatives to NE are provided by the notion of admissible strategy introduced in [1], with computational aspects studied in [4], and potential for synthesis studied in [3]. Note that these notions are free, like (weak) SPEs, of non-credible threats. Finally, in [20], the authors introduce the notion of cooperative and non-cooperative rational synthesis as a general framework where rationality can be specified by either NE, or SPE, or the notion of dominating strategies. In all cases except [8] and [13], the proposed solution concepts are not guaranteed to exist, hence results concern mostly algorithmic techniques to decide their existence, instead of general conditions for existence as in this paper.

Applications to reactive synthesis Games played on graphs have a large number of applications in theoretical computer science. One particularly important application is *reactive synthesis* [23], i.e. the design of a controller that guarantees a good behavior of a reactive system evolving in a possibly hostile environment. One classical model proposed for the synthesis problem is the notion of *two-player zero-sum game played on a graph*. One player is the reactive system and the other one is the environment; the vertices of the graph model their possible states and the edges model their possible actions. Interactions between the players generate an infinite play in the graph which model behaviors of the system within its environment. As one cannot assume cooperation of the environment, the objectives of the two players are considered to be opposite. Constructing a controller for the system then means devising a *winning strategy* for the player modeling it. Reality is often more subtle and the environment is usually not fully adversarial as it has its own objective, meaning that the game should be non zero-sum. Moreover instead of two players, we could consider the more general situation of several players modeling different interacting systems/environments each of them with its own objective.

This has lead to an exploration of a variety of solution concepts for sequential games from the perspective of theoretical computer science (see survey [6]). Weak SPE have the benefit of allowing less unreasonable threats than Nash equilibria,

but existing in more cases than SPE. We can even imagine ruling out infinite deviations by letting a meta-agent punish every one-shot deviation with a (low) fixed probability. A player using an infinitely-deviating strategy will thus be punished by the meta-agent with probability one. Protocols like BitTorrent use similar ideas: every deviant user is temporarily denied suitable bandwidth (see Chapter *Bandwidth Trading as Incentive* in [25] for details).

Structure of the paper In Section 2, we recall the useful notions of game, strategy and weak SPE. In Section 3, we present our general conditions that guarantee the existence of a weak SPE. From this general existence result, we derive two large classes of games with a weak SPE: games with a finite-range outcome function in Section 4, and games with a finite underlying graph and a prefix-independent outcome function in Section 5. In Section 6 we provide an example of a game without weak SPE demonstrating limitations to possible extensions of our main theorem.

An extended abstract omitting most proofs has appeared as [7].

2 Preliminaries

In this section, we recall the useful notions of game, strategy, and weak subgame perfect equilibrium. We illustrate these notions with examples.

2.1 Games

We consider multi-player turn-based games such that an outcome is assigned to every play. Each player has a preference relation on the set of outcomes which allows him to compare plays.

Definition 1. A game is a tuple $G = (\Pi, V, (V_i)_{i \in \Pi}, E, O, \mu, (\prec_i)_{i \in \Pi})$ where:

- Π is a set of players,
- V is a set of vertices and $E \subseteq V \times V$ is a set of edges, such that w.l.o.g. each vertex has at least one outgoing edge,
- $(V_i)_{i \in \Pi}$ is a partition of V such that V_i is the set of vertices controlled by player $i \in \Pi$,
- O is a set of outcomes and $\mu : V^\omega \rightarrow O$ is an outcome function,
- $\prec_i \subseteq O \times O$ is a preference relation for player $i \in \Pi$.

In this definition the underlying graph (V, E) can be infinite (that is, of arbitrarily cardinality), as well as the set Π of players and the set O of outcomes.

A *play* of G is an infinite (countable) sequence $\rho = \rho_0 \rho_1 \dots \in V^\omega$ of vertices such that $(\rho_i, \rho_{i+1}) \in E$ for all $i \in \mathbb{N}$. *Histories* of G are finite sequences $h = h_0 \dots h_n \in V^+$ defined in the same way. We often use notation hv to mention the last vertex $v \in V$ of the history. Usually histories are non empty, but in specific situations it will be useful to consider the empty history ϵ . The set of plays is denoted by *Plays* and the set of histories (ending with a vertex in V_i)

by $Hist$ (resp. by $Hist_i$).³ A *prefix* (resp. *suffix*) of a play $\rho = \rho_0\rho_1\dots$ is a finite sequence $\rho_{\leq n} = \rho_0\dots\rho_n$ (resp. infinite sequence $\rho_{\geq n} = \rho_n\rho_{n+1}\dots$). We use notation $h < \rho$ when a history h is prefix of a play ρ . When an initial vertex $v_0 \in V$ is fixed, we call (G, v_0) an *initialized* game. In this case, plays and histories are supposed to start in v_0 , and we use notations $Plays(v_0)$ and $Hist(v_0)$. In this article, we often *unravel* the graph of the game (G, v_0) from the initial vertex v_0 , which yields an infinite tree rooted at v_0 .

The outcome function assigns an outcome $\mu(\rho) \in O$ to each play $\rho \in V^\omega$. It is *prefix-independent* if $\mu(h\rho) = \mu(\rho)$ for all histories h and play ρ . A *preference* relation $\prec_i \subseteq O \times O$ is an irreflexive and transitive binary relation. It allows for player i to compare two plays $\rho, \rho' \in V^\omega$ with respect to their outcome: $\mu(\rho) \prec_i \mu(\rho')$ means that player i prefers ρ' to ρ . In this paper we restrict to *linear* preferences. (It is w.l.o.g. since the preference properties that we use are preserved by linear extension). We write $o \preceq_i o'$ when $o \prec_i o'$ or $o = o'$; notice that $o \not\prec_i o'$ if and only if $o' \preceq_i o$. We sometimes use notation \prec_v instead of \prec_i when vertex $v \in V_i$ is controlled by player i .

Example 1. Let us mention some classical classes of games where the set of outcomes O is a subset of $(\mathbb{R} \cup \{+\infty, -\infty\})^\Pi$, and for all player $i \in \Pi$, \prec_i is the usual ordering $<$ on $\mathbb{R} \cup \{+\infty, -\infty\}$ on the outcome i -th components. In other words, each player i has a real-valued payoff function $\mu_i : Plays \rightarrow \mathbb{R} \cup \{+\infty, -\infty\}$. The outcome function of the game is then equal to $\mu = (\mu_i)_{i \in \Pi}$, and for all $i \in \Pi$, $\mu(\rho) \prec_i \mu(\rho')$ whenever $\mu_i(\rho) < \mu_i(\rho')$.

Games with *Boolean* objectives are such that $\mu_i : Plays \rightarrow \{0, 1\}$ where 1 (resp. 0) means that the play is won (resp. lost) by player i . Classical objectives are Borel objectives including ω -regular objectives, like reachability, Büchi, parity, aso [18]. Prefix-independence of μ_i holds in the case of Büchi and parity objectives, but not for reachability objective.

We have *quantitative* objectives when $\mu_i : Plays \rightarrow \mathbb{R} \cup \{+\infty, -\infty\}$ replaces $\mu_i : Plays \rightarrow \{0, 1\}$. Usually, such a μ_i is defined from a weight function $w_i : E \rightarrow \mathbb{R}$ that assigns a weight to each edge. Classical examples of μ_i are *limsup* and *mean-payoff* functions [11], that is⁴,

- *limsup*: $\mu_i(\rho) = \limsup_{k \rightarrow \infty} w_i(\rho_k, \rho_{k+1})$
- *mean-payoff*: $\mu_i(\rho) = \limsup_{n \rightarrow \infty} \sum_{k=0}^n \frac{w_i(\rho_k, \rho_{k+1})}{n}$

2.2 Strategies

Let (G, v_0) be an initialized game. A *strategy* σ for player i in (G, v_0) is a function $\sigma : Hist_i(v_0) \rightarrow V$ assigning to each history $hv \in Hist_i(v_0)$ a vertex $v' = \sigma(hv)$ such that $(v, v') \in E$. A strategy σ of player i is *positional* if it only depends on the last vertex of the history, i.e. $\sigma(hv) = \sigma(v)$ for all $hv \in Hist_i(v_0)$. It is a *finite-memory* strategy if it can be encoded by a deterministic *Moore machine* $\mathcal{M} = (M, m_0, \alpha_U, \alpha_N)$ where M is a finite set of states (the memory of the

³ Indexing $Plays_G$ or $Hist_G$ with G allows to recall the related game G .

⁴ The limit inferior can be used instead of the limit superior.

strategy), $m_0 \in M$ is an initial memory state, $\alpha_U : M \times V \rightarrow M$ is an update function, and $\alpha_N : M \times V_i \rightarrow V$ is a next-move function.⁵ Such a machine defines a strategy σ such that $\sigma(hv) = \alpha_N(\hat{\alpha}_U(m_0, h), v)$ for all histories $hv \in \text{Hist}_i(v_0)$, where $\hat{\alpha}_U$ extends α_U to histories as expected. The *memory size* of σ is then the size $|M|$ of \mathcal{M} . In particular σ is positional when it has memory size one.

The previous definitions of (positional, finite-memory) strategy are given for an initialized game (G, v_0) . We call *uniform* every positional strategy σ of player i defined for all $hv \in \text{Hist}_i$ (instead of $\text{Hist}_i(v_0)$), that is, when σ is a positional strategy in all initialized games (G, v) , $v \in V$.

A play ρ is *consistent* with a strategy σ of player i if $\rho_{n+1} = \sigma(\rho_{\leq n})$ for all n such that $\rho_n \in V_i$. A *strategy profile* is a tuple $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ of strategies, where each σ_i is a strategy of player i . It is called *positional* (resp. *finite-memory with memory size bounded by c , uniform*) if all σ_i , $i \in \Pi$, are positional (resp. finite-memory with memory size bounded by c , uniform). Given an initial vertex v_0 , such a strategy profile determines a unique play of (G, v_0) that is consistent with all the strategies. This play induced by $\bar{\sigma}$ in (G, v_0) is denoted by $\langle \bar{\sigma} \rangle_{v_0}$ and we say that $\bar{\sigma}$ has outcome $\mu(\langle \bar{\sigma} \rangle_{v_0})$.

Let $\bar{\sigma}$ be a strategy profile. When all players stick to their own strategy except player i that shifts from σ_i to σ'_i , we denote by $(\sigma'_i, \bar{\sigma}_{-i})$ the derived strategy profile, and by $\langle \sigma'_i, \bar{\sigma}_{-i} \rangle_{v_0}$ the induced play in (G, v_0) . We say that σ'_i is a *deviating* strategy from σ_i . When σ_i and σ'_i only differ on a finite number of histories (resp. on v_0), we say that σ'_i is a *finitely-deviating* (resp. *one-shot deviating*) strategy from σ_i . One-shot deviating strategies is a well-known notion that for instance appears in the proof of Kuhn's theorem [19] with the one-step deviation property. Finitely-deviating strategies have been introduced in [5].

2.3 Variants of subgame perfect equilibria

In this section we recall the notion of subgame perfect equilibrium (SPE) and its variants. Let us first recall the classical notion of Nash equilibrium (NE). Informally, a strategy profile $\bar{\sigma}$ in an initialized game (G, v_0) is an NE if no player has an incentive to deviate (with respect to his preference relation), if the other players stick to their strategies.

Definition 2. *Given an initialized game (G, v_0) , a strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ of (G, v_0) is a Nash equilibrium if for all players $i \in \Pi$, for all strategies σ'_i of player i , we have $\mu(\langle \bar{\sigma} \rangle_{v_0}) \not\prec_i \mu(\langle \sigma'_i, \bar{\sigma}_{-i} \rangle_{v_0})$.*

When $\mu(\langle \bar{\sigma} \rangle_{v_0}) \prec_i \mu(\langle \sigma'_i, \bar{\sigma}_{-i} \rangle_{v_0})$, we say that σ'_i is a *profitable deviation* for player i w.r.t. $\bar{\sigma}$.

The notion of subgame perfect equilibrium is a refinement of NE. In order to define it, we need to introduce the following concepts. Given a game $G = (\Pi, V, (V_i)_{i \in \Pi}, E, \mu, (\prec_i)_{i \in \Pi})$ and a history $h \in \text{Hist}$, we denote by $G|_h$ the

⁵ Moore machines are usually defined for finite sets V of vertices. We here allow infinite sets V .

game $(\Pi, V, (V_i)_{i \in \Pi}, E, \mu_{|h}, (\prec_i)_{i \in \Pi})$ where $\mu_{|h}(\rho) = \mu(h\rho)$ for all plays of $G_{|h}$ ⁶, and we say that $G_{|h}$ is a *subgame* of G . Given an initialized game (G, v_0) and a history $hv \in \text{Hist}(v_0)$, the initialized game $(G_{|h}, v)$ is called the subgame of (G, v_0) with history hv . In particular (G, v_0) is a subgame of itself with history hv_0 such that $h = \epsilon$. Given a strategy σ of player i in (G, v_0) , the strategy $\sigma_{|h}$ in $(G_{|h}, v)$ is defined as $\sigma_{|h}(h') = \sigma(hh')$ for all histories $h' \in \text{Hist}_i(v)$. Given a strategy profile $\bar{\sigma}$ in (G, v_0) , we use notation $\bar{\sigma}_{|h}$ for $(\sigma_{i|h})_{i \in \Pi}$, and $\langle \bar{\sigma}_{|h} \rangle_v$ is the play induced by $\bar{\sigma}_{|h}$ in the subgame $(G_{|h}, v)$.

We can now recall the classical notion of subgame perfect equilibrium: an SPE is a strategy profile in an initialized game that induces an NE in each of its subgames. Two variants of SPE, called weak SPE and very weak SPE, are proposed in [5] such that no player has an incentive to deviate in any subgame using finitely deviating strategies and one-shot deviating strategies respectively (instead of any deviating strategy).

Definition 3. *Given an initialized game (G, v_0) , a strategy profile $\bar{\sigma}$ of (G, v_0) is a (weak, very weak resp.) subgame perfect equilibrium if for all histories $hv \in \text{Hist}(v_0)$, for all players $i \in \Pi$, for all (finitely, one-shot resp.) deviating strategies σ'_i from $\sigma_{i|h}$ of player i in the subgame $(G_{|h}, v)$, we have $\mu(\langle \bar{\sigma}_{|h} \rangle_v) \not\prec_i \mu(\langle \sigma'_i, \bar{\sigma}_{-i|h} \rangle_v)$.*

Trivially, every SPE is a weak SPE, and every weak SPE is a very weak SPE.

Proposition 1 ([5]). *Let $\bar{\sigma}$ be a strategy profile in (G, v_0) . Then $\bar{\sigma}$ is a weak SPE iff $\bar{\sigma}$ is a very weak SPE. There exists an initialized game (G, v_0) with a weak SPE but no SPE.*



Fig. 1. A initialized game (G, v_0) with a (very) weak SPE and no SPE.

Example 2 ([5]). Consider the two-player game (G, v_0) in Figure 1 such that player 1 (resp. player 2) controls vertices v_0, v_2, v_3 (resp. vertex v_1). The set O of outcomes is equal to $\{o_1, o_2, o_3\}$, and the outcome function is prefix-independent such that $\mu((v_0v_1)^\omega) = o_1$, $\mu(v_2^\omega) = o_2$, and $\mu(v_3^\omega) = o_3$. The preference relation for player 1 (resp. player 2) is $o_1 \prec_1 o_2 \prec_1 o_3$ (resp. $o_2 \prec_2 o_3 \prec_2 o_1$).

It is known that this game has no SPE [26]. Nevertheless the positional strategy profile $\bar{\sigma}$ depicted with thick edges is a very weak SPE, and thus a weak SPE by Proposition 1. Let us give some explanation. Due to the simple form of

⁶ In this article, we will always use notation $\mu(h\rho)$ instead of $\mu_{|h}(\rho)$.

the game, only two cases are to be treated. Consider first the subgame $(G|_h, v_0)$ with $h \in (v_0 v_1)^*$, and the one-shot deviating strategy σ'_1 from $\sigma_{1|h}$ such that $\sigma'_1(v_0) = v_2$. Then $\langle \bar{\sigma}|_h \rangle_{v_0} = v_0 v_1 v_3^\omega$ and $\langle \sigma'_1, \sigma_{2|h} \rangle_{v_0} = v_0 v_2^\omega$ with respective outcomes o_3 and o_2 , showing that σ'_1 is not a profitable deviation for player 1 in $(G|_h, v_0)$. Now in the subgame $(G|_h, v_1)$ with $h \in (v_0 v_1)^* v_0$, the one-shot deviating strategy from $\sigma_{2|h}$ such that $\sigma'_2(v_1) = v_0$ is not profitable for player 2 in $(G|_h, v_1)$ because $\langle \bar{\sigma}|_h \rangle_{v_1} = v_1 v_3^\omega$ and $\langle \sigma_{1|h}, \sigma'_2 \rangle_{v_1} = v_1 v_0 v_1 v_3^\omega$ with the same outcome o_3 .

Notice that $\bar{\sigma}$ is not an SPE. Indeed the strategy σ'_2 such that $\sigma'_2(hv_1) = v_0$ for all h , is infinitely deviating from σ_2 , and is a profitable deviation for player 2 in (G, v_0) since $\langle \sigma_1, \sigma'_2 \rangle_{v_0} = (v_0 v_1)^\omega$ with outcome o_1 .

3 General conditions for the existence of weak SPEs

In this section, we propose general conditions to guarantee the existence of weak SPEs. In the next sections, from this result, we will derive two interesting large families of games always having a weak SPE.

Theorem 1. *Let (G, v_0) be an initialized game with a subset $L \subseteq V$ of vertices called leaves with only one outgoing edge (l, l) for all $l \in L$. Suppose that:*

1. *for all $v \in V$, there exists a play $\rho = hl^\omega$ for some $h \in \text{Hist}(v)$ and $l \in L$,*
2. *for all plays $\rho = hl^\omega$ with $h \in \text{Hist}(v)$ and $l \in L$, $\mu(\rho) = \mu(l^\omega)$,*
3. *the set of outcomes $O_L = \{\mu(l^\omega) \mid l \in L\}$ is finite.*

Then there always exists a weak SPE $\bar{\sigma}$ in (G, v_0) . Moreover, $\bar{\sigma}$ is finite-memory with memory size bounded by $|O_L|$.

Let us comment the hypotheses. The first condition means that from each vertex v of the game there is a leaf reachable from v ; in particular L is not empty. The second condition expresses prefix-independence of the outcome function restricted to plays eventually looping in a leaf $l \in L$. The last condition means that even if there is an infinite number of leaves, the set of outcomes assigned by μ to plays eventually looping in L is finite. The next example describes a family of games satisfying the conditions of Theorem 1.

Example 3. For each natural number $n \geq 3$, we build a game G_n with n players, $2n$ vertices, $3n$ edges, and $n+1$ outcomes. The set of players is $\Pi = \{1, 2, \dots, n\}$ and the set of vertices is $V = \{v_1, \dots, v_n, l_1, \dots, l_n\}$ such that $V_i = \{v_i, l_i\}$ for all $i \in \Pi$. The edges are $(v_1, v_2), (v_2, v_3), \dots, (v_n, v_1)$, and $(v_i, l_i), (l_i, l_i)$ for all $i \in \Pi$. The game G_4 is depicted in Figure 2. The set O of outcomes is equal to $\{o_1, \dots, o_n, \perp\}$, and the outcome function is prefix-independent such that $\mu((v_1 v_2 \dots v_n)^\omega) = \perp$ and $\mu(l_i^\omega) = o_i$ for all $i \in \Pi$. Each player i has a preference relation \prec_i satisfying $\perp \prec_i o_{i-1} \prec_i o_i \prec_i o_j$ for all $j \in \Pi \setminus \{i-1, i\}$ (with the convention that $o_0 = o_n$).

Each game (G_n, v_1) satisfies the hypotheses of Theorem 1 with $L = \{l_1, \dots, l_n\}$ and thus has a finite-memory weak SPE. Such a strategy profile $\bar{\sigma}$ is depicted




Fig. 2. Game G_4

in Figure 3 for $n = 4$ (see the thick edges on the unravelling of G_4 from the initial vertex v_1) and can be easily generalized to every $n \geq 3$. One verifies that this profile is a very weak SPE, and thus a weak SPE by Proposition 1. For all $i \in \Pi$, the strategy σ_i of player i is finite-memory with a memory size equal to $n - 1$. Intuitively, along $(v_1 \dots v_n)^\omega$, player i repeatedly produces one move (v_i, l_i) followed by $n - 2$ moves (v_i, v_{i+1}) . Hence the memory states of the Moore machine for σ_i are counters from 1 to $n - 1$. The Moore machine for σ_1 in the game (G_4, v_1) is depicted in Figure 4 (with $M = \{1, 2, 3\}$, $m_0 = 1$, and the update and next-move functions indicated by the edges).




Fig. 3. Weak SPE in (G_4, v_1)




Fig. 4. The Moore machine for σ_1

Let us now proceed to the proof of Theorem 1. Recall that it is enough to prove the existence of a very weak SPE by Proposition 1. The proof idea is the following one. Initially, for each vertex v , we accept all plays $\rho = hl^\omega$ with $h \in \text{Hist}(v)$ and $l \in L$ as *potential* plays induced by a very weak SPE in the initialized game (G, v) . We thus label each v by the set of outcomes $\mu(l^\omega)$ for such leaves l (recall that $\mu(\rho) = \mu(l^\omega)$ by the second condition of Theorem 1). Notice that this labeling is finite (resp. not empty) by the third (resp. first) condition of the theorem. Step after step, we are going to remove some outcomes from the vertex labelings by a *Remove* operation followed by an *Adjust* operation. The *Remove* operation removes an outcome o from the labeling of a given vertex v when there exists an edge (v, v') for which $o \prec_v o'$ for all outcomes o' that label v' . Indeed o cannot be the outcome of a play induced by a very weak SPE since the player who controls v will choose the move (v, v') to get a preferable outcome o' . Now it may happen that for another vertex u having o in its labeling, all potential plays induced by a very weak SPE from u with outcome o necessarily cross vertex v . As o has been removed from the labeling of v , these potential plays do no longer survive and o will also be removed from the labeling of u by the *Adjust* operation. Repeatedly applying these two operations converge to a fixpoint for which we will prove non-emptiness (this is the difficult part of the proof, non-emptiness will be obtained by maintaining three invariants, see Lemma 1). From this fixpoint, for each vertex v and each outcome o of the resulting labeling of v , there exists a play $\rho_{v,o} = hl^\omega$ with outcome o for some $h \in \text{Hist}(v)$ and $l \in L$. We can thus build a very weak SPE $\bar{\sigma}$ in (G, v_0) as follows. The construction of $\bar{\sigma}$ is done step by step: (i) initially $\bar{\sigma}$ is partially defined such that $\langle \bar{\sigma} \rangle_{v_0} = \rho_{v_0, o_0}$ for some o_0 ; (ii) then in the subgame $(G|_h, v)$ such that $\langle \bar{\sigma}|_h \rangle_v = \rho_{v,o}$, if the player who controls v chooses the move (v, v') in a one-shot deviation, then there exists $\rho_{v',o'}$ such that $o \not\prec_v o'$ by definition of the fixpoint, and we thus extend the construction of $\bar{\sigma}$ such that $\langle \bar{\sigma}|_{hv} \rangle_{v'} = \rho_{v',o'}$.

Let us now go into the details of the proof. For each $l \in L$, we denote by o_l the outcome $\mu(l^\omega)$. Recall that for all $\rho = hl^\omega$ we have $\mu(\rho) = o_l$ by the second hypothesis of the theorem. For each $v \in V$, we denote by $\text{Succ}(v)$ the set of successors of v distinct from v , that is, the vertices $v' \neq v$ such that $(v, v') \in E$. Notice that the leaves l are the vertices with only one outgoing edge (l, l) . Thus, by definition, $\text{Succ}(v) = \emptyset$ for all $v \in L$ and $\text{Succ}(v) \neq \emptyset$ for all $v \in V \setminus L$.

The labeling $\lambda_\alpha(v)$ of the vertices v of G by subsets of O_L is an inductive process on the ordinal α . Initially (step $\alpha = 0$), each $v \in V$ is labeled by:

$$\lambda_0(v) = \{o_l \in O_L \mid \text{there exists a play } hl^\omega \text{ with } h \in \text{Hist}(v) \text{ and } l \in L\}.$$

(In particular $\lambda_0(l) = \{o_l\}$ for all $l \in L$). By the first hypothesis of the theorem, $\lambda_0(v) \neq \emptyset$. Let us introduce some additional terminology. At step α , when there is a path⁷ π from v to v' in G , we say that π is (o, α) -labeled if $o \in \lambda_\alpha(u)$ for all the vertices u of π . Thus initially, we have a $(o_l, 0)$ -labeled path from v to l for each $o_l \in \lambda_0(v)$. For $v \in V$, let

$$m_\alpha(v) = \max_{\prec_v} \{\min_{\prec_v} \lambda_\alpha(v') \mid v' \in \text{Succ}(v)\}$$

⁷ By path, we mean a finite path

with the convention that $m_\alpha(v) = \top$ if $Succ(v) = \emptyset$ or if $\lambda_\alpha(v') = \emptyset$ for all $v' \in Succ(v)$.⁸ When $m_\alpha(v) \neq \top$, we says that $v' \in Succ(v)$ *realizes* $m_\alpha(v)$ if $m_\alpha(v) = \min_{\prec_v} \lambda_\alpha(v')$. Notice that even if $Succ(v)$ could be infinite, there are finitely many sets $\lambda_\alpha(v')$ since O_L is finite. This justifies our use of \max_{\prec_v} and \min_{\prec_v} operators in the definition of $m_\alpha(v)$.

We alternate between applying *Remove* and *Adjust* to the current labeling. More formally, we define the labeling λ_α inductively⁹. In the following, γ is always assumed to be a limit ordinal and n to be a natural number.

– **Defining $\lambda_{\gamma+2n+1}$ via *Remove* operation**

Let $\alpha := \gamma + 2n + 1$. Test if for some $v \in V$, there exist $o \in \lambda_{\alpha-1}(v)$ and $v' \in Succ(v)$ such that

$$o \prec_v o', \text{ for all } o' \in \lambda_{\alpha-1}(v').$$

If such a v exists, then $\lambda_\alpha(v) = \lambda_{\alpha-1}(v) \setminus \{o\}$, and $\lambda_\alpha(u) = \lambda_{\alpha-1}(u)$ for the other vertices $u \neq v$. Otherwise $\lambda_\alpha(u) = \lambda_{\alpha-1}(u)$ for all $u \in V$.

– **Defining $\lambda_{\gamma+2n+2}$ via *Adjust* operation**

Let $\alpha := \gamma + 2n + 2$. Suppose that $\lambda_{\alpha-1}(v) = \lambda_{\alpha-2}(v) \setminus \{o\}$ at the previous step. For all $u \in V$ such that $o \in \lambda_{\alpha-1}(u)$, test if there exists a $(o, \alpha - 1)$ -labeled path from u to some $l \in L$. If yes, then $\lambda_\alpha(u) = \lambda_{\alpha-1}(u)$, otherwise $\lambda_\alpha(u) = \lambda_{\alpha-1}(u) \setminus \{o\}$. For all $u \in V$ such that $o \notin \lambda_{\alpha-1}(u)$, let $\lambda_\alpha(u) = \lambda_{\alpha-1}(u)$.

Suppose that $\lambda_{\alpha-1}(v) = \lambda_{\alpha-2}(v)$ for all $v \in V$ at the previous step, then $\lambda_\alpha(v) = \lambda_{\alpha-1}(v)$ for all $v \in V$.

– **Defining λ_γ via intersection**

Let $\lambda_\gamma(v) = \bigcap_{\beta < \gamma} \lambda_\beta(v)$ for all $v \in V$.

For each v , the sequence $(\lambda_\alpha(v))_\alpha$ is nonincreasing (w.r.t. set inclusion), and thus the sequence $(m_\alpha(v))_\alpha$ is nondecreasing (w.r.t. \prec_v). Moreover, the sequence $(\lambda_\alpha)_\alpha$ is nonincreasing w.r.t. pointwise set inclusion. Thus, there exists some ordinal α^* such that $\lambda_{\alpha^*} = \lambda_\beta$ for all $\beta > \alpha^*$. By inspecting the definition, we see that it suffices to check that $\lambda_{\alpha^*} = \lambda_{\alpha^*+1} = \lambda_{\alpha^*+2}$ in order to see that α^* is a fixed point. If V is finite, such a fixed point is reached after at most $2|O_L| \cdot |V|$ steps. The central challenge is to show that this fixed point is non-empty in each component.

Notice that for all leaves $l \in L$ and all steps α , we have $\lambda_\alpha(l) = \{o_l\}$.

Lemma 1. *There exists an ordinal α^* such that*

$$\lambda_{\alpha^*}(v) = \lambda_{\alpha^*+1}(v) = \lambda_{\alpha^*+2}(v) \text{ for all } v \in V.$$

Moreover, $\lambda_{\alpha^*}(v) \neq \emptyset$ for all $v \in V$.

⁸ We suppose that $o \prec_v \top$ for all $o \in O_L$.

⁹ Note that our definition as written makes non-deterministic choices. This is immaterial for our purposes, but could be determinized by demanding a well-ordering of the vertex set and the outcomes.

To be able to prove that $\lambda_{\alpha^*}(v) \neq \emptyset$, we introduce three invariants for which we will prove that they are initially true (Lemma 2) and remain true after each step α (Lemmata 3,4,5). The non emptiness of $\lambda_{\alpha^*}(v)$ will follow from the second invariant.

INV1 For $v \in V$, we have for all $v' \in Succ(v)$ that

$$\{o \in \lambda_{\alpha}(v') \mid m_{\alpha}(v) \preceq_v o\} \subseteq \lambda_{\alpha}(v).$$

In particular, when $m_{\alpha}(v) \neq \top$, for each v' that realizes $m_{\alpha}(v)$, we have

$$\lambda_{\alpha}(v') \subseteq \lambda_{\alpha}(v). \quad (1)$$

INV2 For $v \in V$, $\lambda_{\alpha}(v) \neq \emptyset$.

INV3 For $v \in V$, there exists a path from v to some $l \in L$ such that for all vertices u in this path, $\lambda_{\alpha}(u) \subseteq \lambda_{\alpha}(v)$.

Lemma 2. *All three invariants are true for λ_0 .*

Proof. Consider $v \in V$ at the initial step $\alpha = 0$. By hypothesis there is a path from v to some $l \in L$. Thus $\lambda_{\alpha}(v) \neq \emptyset$ and INV2 is true. Moreover, for all $v' \in Succ(v)$, we have $\lambda_{\alpha}(v') \subseteq \lambda_{\alpha}(v)$ by the initial labeling, and thus INV1 and INV3 are also true. \square

Lemma 3. *All three invariants are preserved by Remove.*

Proof. Consider some $\alpha = \gamma + 2n$ for limit ordinal γ and $n \in \mathbb{N}$ such that all invariants hold for λ_{α} . If $\lambda_{\alpha+1} = \lambda_{\alpha}$, then trivially, all invariants hold for $\lambda_{\alpha+1}$. Otherwise there exist v and o such that $\lambda_{\alpha+1}(v) = \lambda_{\alpha}(v) \setminus \{o\}$ and $\lambda_{\alpha+1}(u) = \lambda_{\alpha}(u)$ for all $u \neq v$. In particular $v \notin L$. For all $u \in V$, we have $m_{\alpha}(u) \preceq_u m_{\alpha+1}(u)$, with the particular case $m_{\alpha}(v) = m_{\alpha+1}(v)$.

- **Remove cannot violate INV1.** We first consider $u \in V$ such that $u \neq v$. For all $u' \in Succ(u)$, we have

$$\begin{aligned} & \{o' \in \lambda_{\alpha+1}(u') \mid m_{\alpha+1}(u) \preceq_u o'\} \\ & \subseteq \{o' \in \lambda_{\alpha}(u') \mid m_{\alpha}(u) \preceq_u o'\} \quad \text{since } \lambda_{\alpha+1}(u') \subseteq \lambda_{\alpha}(u') \\ & \quad \text{and } m_{\alpha}(u) \preceq_u m_{\alpha+1}(u), \\ & \subseteq \lambda_{\alpha}(u) \quad \text{by INV1 at step } \alpha, \\ & = \lambda_{\alpha+1}(u) \quad \text{as } u \neq v. \end{aligned}$$

Let us turn to vertex v . As $o \prec_v m_{\alpha}(v)$, the previous inclusions can be modified as follows. For all $v' \in Succ(v)$, we now have $\{o' \in \lambda_{\alpha+1}(v') \mid m_{\alpha+1}(v) \preceq_v o'\} \subseteq \{o' \in \lambda_{\alpha}(v') \mid m_{\alpha}(v) \preceq_v o'\} \subseteq \lambda_{\alpha}(v) \setminus \{o\} = \lambda_{\alpha+1}(v)$.

- **Remove cannot violate INV2.** We only have to show that $\lambda_{\alpha+1}(v) \neq \emptyset$. As $Succ(v) \neq \emptyset^{10}$ and by INV2, we have $m_{\alpha}(v) \neq \top$. Hence there exists $v' \in Succ(v)$ that realizes $m_{\alpha}(v) = m_{\alpha+1}(v)$. By INV1 and in particular (1) at step $\alpha + 1$, we thus have $\lambda_{\alpha+1}(v') \subseteq \lambda_{\alpha+1}(v)$. As $\lambda_{\alpha+1}(v') = \lambda_{\alpha}(v') \neq \emptyset$, it follows that $\lambda_{\alpha+1}(v) \neq \emptyset$.

¹⁰ Recall that $v \notin L$, and that $Succ(v) \neq \emptyset$ for all $v \in V \setminus L$.

- **Remove cannot violate INV3.** We first consider $u \neq v$. By INV3, there exists a path π from u to some $l \in L$ such that $\lambda_\alpha(w) \subseteq \lambda_\alpha(u)$ for all vertices w in this path. We can keep the path π at step $\alpha + 1$ since $\lambda_{\alpha+1}(w) \subseteq \lambda_\alpha(w)$ for all w in π and $\lambda_{\alpha+1}(u) = \lambda_\alpha(u)$. We now consider vertex v . Consider again $v' \in Succ(v)$ that realizes $m_{\alpha+1}(v)$. By (1), $\lambda_{\alpha+1}(v') \subseteq \lambda_{\alpha+1}(v)$. We know that there exists a path π from v' to some $l \in L$ such that $\lambda_\alpha(w) \subseteq \lambda_\alpha(v')$ for all w in π . This path π augmented with the edge (v, v') is the required path for INV3 at step $\alpha + 1$ because for all w in π , we have $\lambda_{\alpha+1}(w) \subseteq \lambda_\alpha(w) \subseteq \lambda_\alpha(v') = \lambda_{\alpha+1}(v') \subseteq \lambda_{\alpha+1}(v)$. \square

Lemma 4. *All three invariants are preserved by Adjust.*

Proof. Let all three invariants hold for $\alpha = \gamma + 2n + 1$ for limit ordinal γ and $n \in \mathbb{N}$. Then the preceding step was a *Remove* step. If $\lambda_\alpha = \lambda_{\alpha-1}$, then $\lambda_{\alpha+1} = \lambda_\alpha$. Otherwise, there are $v_0 \in V$ and an outcome o such that $\lambda_\alpha(v_0) = \lambda_{\alpha-1}(v_0) \setminus \{o\}$ and $\lambda_\alpha(u) = \lambda_{\alpha-1}(u)$ for all $u \neq v_0$.

For all $v \in V$, either $\lambda_{\alpha+1}(v) = \lambda_\alpha(v)$ or $\lambda_{\alpha+1}(v) = \lambda_\alpha(v) \setminus \{o\}$, and $m_\alpha(v) \preceq_v m_{\alpha+1}(v)$.

Consider $v \in V$ such that $o \notin \lambda_{\alpha+1}(v)$ and $o \in \lambda_\alpha(v)$. Then

$$\forall v' \in Succ(v), o \notin \lambda_{\alpha+1}(v') \quad (2)$$

Otherwise if $o \in \lambda_{\alpha+1}(v')$ for some $v' \in Succ(v)$, this means that o has not been removed from $\lambda_\alpha(v')$, i.e., there exists a (o, α) -labeled path from v' to some $l \in L$, and thus also from v to l by using the edge (v, v') . This is in contradiction with o being removed from $\lambda_\alpha(v)$.

- **Adjust cannot violate INV1.** We first consider $v \in V$ such that $\lambda_{\alpha+1}(v) = \lambda_\alpha(v)$. As done for INV1 and *Remove*, we have for all $v' \in Succ(v)$ that $\{o' \in \lambda_{\alpha+1}(v') \mid m_{\alpha+1}(v) \preceq_v o'\} \subseteq \{o' \in \lambda_\alpha(v') \mid m_\alpha(v) \preceq_v o'\} \subseteq \lambda_\alpha(v) = \lambda_{\alpha+1}(v)$. We now consider $v \in V$ such that $\lambda_{\alpha+1}(v) \neq \lambda_\alpha(v)$. Let $v' \in Succ(v)$. From (2), we have $\{o' \in \lambda_{\alpha+1}(v') \mid m_{\alpha+1}(v) \preceq_v o'\} \subseteq \{o' \in \lambda_\alpha(v') \mid m_\alpha(v) \preceq_v o'\} \setminus \{o\} \subseteq \lambda_\alpha(v) \setminus \{o\} = \lambda_{\alpha+1}(v)$.
- **Adjust cannot violate INV2.** Assume that for some $v \in V$, $\lambda_{\alpha+1}(v) = \emptyset$, that is, $\lambda_\alpha(v) = \{o\}$. By INV3, there exists a path π from v to some $l \in L$ such that $\lambda_\alpha(u) \subseteq \lambda_\alpha(v)$ for all u in π . From $\lambda_\alpha(v) = \{o\}$ and $\lambda_\alpha(u) \neq \emptyset$ (by INV2), we get $\lambda_\alpha(u) = \{o\}$ for all such u . Therefore, the path π from v to l is (o, α) -labeled and o cannot be removed from $\lambda_\alpha(v)$, showing that $\lambda_{\alpha+1}(v) \neq \emptyset$.
- **Adjust cannot violate INV3.** Let $v \in V$ and by INV3 take a path $u_1 \dots u_n$ from $v = u_1$ to some $l = u_n$ with $l \in L$ such that $\lambda_\alpha(u_i) \subseteq \lambda_\alpha(v)$ for all i . Either this path is still valid at step $\alpha + 1$, or there exists a smallest i such that $o \in \lambda_{\alpha+1}(u_i) = \lambda_\alpha(u_i)$, but $o \in \lambda_\alpha(v)$ and $o \notin \lambda_{\alpha+1}(v)$. By minimality of i , $o \notin \lambda_{\alpha+1}(u_j)$ for all $j \leq i - 1$. By the contraposition of (2) with u_{i-1} and u_i , knowing that $o \notin \lambda_{\alpha+1}(u_{i-1})$, it follows that $o \notin \lambda_\alpha(u_{i-1})$. By INV3 there is a path π from u_{i-1} to some

$l' \in L$ such that for all w in π , $\lambda_\alpha(w) \subseteq \lambda_\alpha(u_{i-1}) (\subseteq \lambda_\alpha(v))$. Notice that $o \notin \lambda_\alpha(w)$ for all these w since $o \notin \lambda_\alpha(u_{i-1})$. The path π' obtained by concatenating $u_1 \dots u_{i-1}$ with π is the required path from v for INV3 at step $\alpha + 1$. Indeed for all w' in π' , we have seen that $\lambda_\alpha(w') \subseteq \lambda_\alpha(v)$ and $o \notin \lambda_{\alpha+1}(w')$. Thus $\lambda_{\alpha+1}(w') \subseteq \lambda_\alpha(v) \setminus \{o\} = \lambda_{\alpha+1}(v)$. \square

Lemma 5. *If all three invariants are true for each λ_β , $\beta < \alpha$, α a limit ordinal, then they are true for λ_α .*

Proof. Let α be a limit ordinal, and suppose that the three invariants are true for each ordinal $\beta < \alpha$. Given $v \in V$, as the set $\lambda_\beta(v)$ is finite¹¹ and the sequence $(\lambda_\beta(v))_{\beta < \alpha}$ is nonincreasing, there exists some $\gamma < \alpha$ such that $\lambda_\beta(v) = \lambda_\gamma(v)$ for all $\beta, \gamma \leq \beta < \alpha$. Therefore

$$\lambda_\alpha(v) = \bigcap_{\beta < \alpha} \lambda_\beta(v) = \lambda_\gamma(v). \quad (3)$$

It immediately follows that INV2 holds at step α . To show that INV3 also holds, consider a path π from v to some $l \in L$ such that $\lambda_\gamma(u) \subseteq \lambda_\gamma(v)$ for all u in π (by INV3 at step γ). We can take this path π for INV3 at step α since for all these u , we have $\lambda_\alpha(u) \subseteq \lambda_\gamma(u) \subseteq \lambda_\gamma(v) = \lambda_\alpha(v)$. Finally, the first invariant remains true at step α because for all $v' \in Succ(v)$, we have

$$\begin{aligned} & \{o \in \lambda_\alpha(v') \mid m_\alpha(v) \preceq_v o\} \\ & \subseteq \{o \in \lambda_\gamma(v') \mid m_\gamma(v) \preceq_v o\} \text{ since } \lambda_\alpha(v') \subseteq \lambda_\gamma(v') \text{ and } m_\gamma(v) \preceq_v m_\alpha(v), \\ & \subseteq \lambda_\gamma(v) \quad \text{by INV1 at step } \gamma, \\ & = \lambda_\alpha(v) \quad \text{by (3)}. \end{aligned}$$

\square

To get Theorem 1, it remains to explain how to build a finite-memory weak SPE $\bar{\sigma}$ from the fixed point provided by Lemma 1.

Proof (of Theorem 1). By Lemma 1, we have a fixed point of *Remove* and *Adjust* such that that $\lambda_{\alpha^*}(v) \neq \emptyset$ for all $v \in V$. Since λ_{α^*} is unchanged by *Adjust*, for all $o \in \lambda_{\alpha^*}(v)$, there is a (o, α^*) -labeled path π from v to some $l \in L$ with $o_l = o$. We denote by $\rho_{v,o}$ the play πl^ω :

$$\rho_{v,o} = \pi l^\omega. \quad (4)$$

(*) Recall that $\mu(\rho_{v,o}) = o_l$, and have in mind that $o_l \in \lambda_{\alpha^*}(u)$ for all vertices u in $\rho_{v,o}$.

The construction of $\bar{\sigma}$ will be done step by step thanks to a progressive labeling of the histories by outcomes in O_L and by using the plays $\rho_{v,o}$. This labeling $\kappa : Hist(v_0) \rightarrow O_L$ will allow to recover from history hv the outcome o of the play $(\bar{\sigma}|_h)_v$ induced by $\bar{\sigma}$ in the subgame $(G|_h, v)$.

¹¹ This is the place in the proof where finiteness of the number of outcomes is used in a crucial way.

We start with history v_0 and any $o_0 \in \lambda_{\alpha^*}(v_0)$. Consider ρ_{v_0, o_0} as in (4). The strategy profile $\bar{\sigma}$ is partially built such that $\langle \bar{\sigma} \rangle_{v_0} = \rho_{v_0, o_0}$. The non empty prefixes g of ρ_{v_0, o_0} are all labeled with $\kappa(g) = o_0$.

At the following steps, we consider a history $h'v'$ that is not yet labeled, but such that $h' = hv$ has already been labeled by $\kappa(hv) = o$. The labeling of hv by o means that $\bar{\sigma}$ has already been built to produce the play $\langle \bar{\sigma} \rangle_v$ with outcome o in the subgame $(G|_h, v)$, such that $\langle \bar{\sigma} \rangle_v$ is suffix of $\rho_{u, o}$ from some u . By (*) we have $o \in \lambda_{\alpha^*}(v)$. As λ_{α^*} is invariant under *Remove* (noting $o \in \lambda_{\alpha^*}(v)$ and $v' \in Succ(v)$), there exists $o' \in \lambda_{\alpha^*}(v')$ such that

$$o \not\prec_v o'. \quad (5)$$

With $\rho_{v', o'}$ as in (4), we then extend the construction of $\bar{\sigma}$ such that $\langle \bar{\sigma} \rangle_{v'} = \rho_{v', o'}$, and for each non empty prefix g of $\rho_{v', o'}$, we label $h'g$ by $\kappa(h'g) = o'$ (notice that the prefixes of h' have already been labeled by choice of h'). This process is iterated to complete the construction of $\bar{\sigma}$.

Let us show that the constructed profile $\bar{\sigma}$ is a very weak SPE in (G, v_0) . Consider a history $h' = hv \in Hist(v_0)$ with $v \in V_i$, and a one-shot deviating strategy σ'_i from $\sigma_i|_h$ in the subgame $(G|_h, v)$. Let v' be such that $\sigma'_i(v) = v'$. By definition of $\bar{\sigma}$, we have $\kappa(hv) = o$ and $\kappa(h'v') = o'$ such that (5) holds. Let $\rho = \langle \bar{\sigma} \rangle_v$ and $\rho' = \langle \bar{\sigma} \rangle_{v'}$. Then $o = \mu(h\rho)$ and $o' = \mu(hv\rho')$ by (*). By (5), σ'_i is not a profitable deviation for player i . Hence $\bar{\sigma}$ is a very weak SPE and thus a weak SPE by Proposition 1.

It remains to prove that $\bar{\sigma}$ is finite-memory by correctly choosing the plays $\rho_{v, o}$ of (4). Fix $o \in O_L$ and consider the set U_o of vertices v such that $o \in \lambda_{\alpha^*}(v)$. Then we choose the plays $\rho_{v, o} = \pi l^\omega$ for all $v \in U_o$, such that the set of associated finite paths πl forms a tree. Therefore having o in memory, the required Moore machine can produce positionally each $\rho_{v, o}$ with $v \in U_o$. Hence its set M of states is equal to O_L . \square

The next corollary is an easy consequence of Theorem 1. Under the same conditions except perhaps the second one, and when the underlying graph of G is a tree, it guarantees the existence of a weak SPE that is positional.

Corollary 1. *Let (G, v_0) be an initialized game with a subset $L \subseteq V$ of leaves¹² such that the underlying graph is a tree rooted at v_0 . If (G, v_0) satisfies the first and third conditions of Theorem 1, then there exists a positional weak SPE in (G, v_0) .*

Proof. If the second condition of Theorem 1 is not satisfied, we replace the outcome function μ by a new function μ' defined as follows. For all plays l^ω , with $l \in L$, there is a unique path π from v_0 to l as the underlying graph is a tree. For all suffixes ρ of πl^ω , we let $\mu'(\rho) = \mu(\pi l^\omega)$. For all the remaining plays ρ , we let $\mu'(\rho) = \mu(\rho)$. With the new function μ' , the game (G, v_0) now satisfies

¹² The existence of leaves l with a unique outgoing edge (l, l) is abusive since the graph is a tree: it should be understood as a unique infinite play from l .

all the conditions of Theorem 1 and has thus a weak SPE $\bar{\sigma}$ with respect to μ' . It is easy to see that $\bar{\sigma}$ is also a weak SPE with respect to μ . Notice that this profile is necessarily positional as the underlying graph is a tree. \square

In the next two sections, we present two large families of games for which there always exists a weak SPE. We will explain how these results are obtained from Theorem 1 and its Corollary 1. Before that, we demonstrate the argument establishing Theorem 1 on the game G_4 as introduced in Example 3.

Example 4. Let us describe the inductive process for the game G_4 of Figure 2 (Page 9). For all $i \in \Pi$ and all steps α , we have $\lambda_\alpha(l_i) = \{o_i\}$. Table 1 indicates the different steps until reaching α^* for the vertices v_i , $i \in \Pi$, with $O_L = \{o_1, o_2, o_3, o_4\}$. For instance, at step 1, *Remove* removes o_4 from $\lambda_\alpha(v_1)$ because $o_4 \prec_1 o'$ for all $o' \in \lambda_\alpha(l_1) = \{o_1\}$. At step 2, *Adjust* removes no outcome. For $v = v_1$ and $o \in \lambda_\alpha(v_1)$, the plays $\rho_{v,o}$ are:

$$\rho_{v_1, o_1} = v_1 l_1^\omega, \quad \rho_{v_1, o_2} = v_1 v_2 l_2^\omega, \quad \rho_{v_1, o_3} = v_1 v_2 v_3 l_3^\omega.$$

The other vertices $v \neq v_1$ have similar plays $\rho_{v,o}$.

α	$\lambda_\alpha(v_1)$	$\lambda_\alpha(v_2)$	$\lambda_\alpha(v_3)$	$\lambda_\alpha(v_4)$
0	O_L	O_L	O_L	O_L
1	$O_L \setminus \{o_4\}$	O_L	O_L	O_L
2	$O_L \setminus \{o_4\}$	O_L	O_L	O_L
3	$O_L \setminus \{o_4\}$	$O_L \setminus \{o_1\}$	O_L	O_L
4	$O_L \setminus \{o_4\}$	$O_L \setminus \{o_1\}$	O_L	O_L
5	$O_L \setminus \{o_4\}$	$O_L \setminus \{o_1\}$	$O_L \setminus \{o_2\}$	O_L
6	$O_L \setminus \{o_4\}$	$O_L \setminus \{o_1\}$	$O_L \setminus \{o_2\}$	O_L
7	$O_L \setminus \{o_4\}$	$O_L \setminus \{o_1\}$	$O_L \setminus \{o_2\}$	$O_L \setminus \{o_3\}$
$\alpha^* = 8$	$O_L \setminus \{o_4\}$	$O_L \setminus \{o_1\}$	$O_L \setminus \{o_2\}$	$O_L \setminus \{o_3\}$

Table 1. The different steps until reaching a fixed point for game G_4

In the case of game (G_4, v_1) , the construction of a weak SPE $\bar{\sigma}$, as described in the previous proof, leads to the strategy profile of Figure 3. Indeed, the construction of $\bar{\sigma}$ begins with history v_1 and $\rho_{v_1, o_1} = v_1 l_1^\omega$. At the next step, we consider history $v_1 v_2$ and $\rho_{v_2, o_4} = v_2 v_3 v_4 l_4^\omega$ such that $o_1 \not\prec_1 o_4$, aso. Notice that the previous proof states a memory size equal to 4 for $\bar{\sigma}$ whereas Figure 4 depicts a Moore machine for $\bar{\sigma}$ with a better memory size equal to 3.

4 First application

In this section, we begin with the first application of the results of the previous section (more particularly Corollary 1): when an initialized game has an outcome function with finite range, then it always has a weak SPE.

Theorem 2. *Let (G, v_0) be an initialized game such that the outcome function has finite range. Then there exists a weak SPE in (G, v_0) .*

Let us comment this theorem. (i) Kuhn's theorem [19] states that there always exist an SPE in initialized games played on a *finite tree* (notice that in this particular case, the existence of a weak SPE is equivalent to the existence of an SPE). Theorem 2 can be seen as a generalization of Kuhn's theorem: if we keep the outcome set finite, all initialized games (regardless of the underlying graph and the player set) have weak SPE. (ii) The next theorem is proved in [14] for outcome functions $\mu = (\mu_i)_{i \in \Pi}$ as presented in Example 1 and has strong relationship with Theorem 2. Recall that a payoff function $\mu_i : Plays \rightarrow \mathbb{R}$ is *lower-semicontinuous* if whenever a sequence of plays $(\rho_n)_{n \in \mathbb{N}}$ converges to a play $\rho = \lim_{n \rightarrow \infty} \rho_n$, then $\liminf_{n \rightarrow \infty} \mu_i(\rho_n) \geq \mu_i(\rho)$.

Theorem 3 ([14]). *Let (G, v_0) be an initialized game with a finite set Π of players and an outcome function $\mu = (\mu_i)_{i \in \Pi}$ such that each $\mu_i : Plays \rightarrow \mathbb{R}$ has finite range and is lower-semicontinuous. Then there exists an SPE in (G, v_0) .*

As every weak SPE is an SPE in the case of lower-semicontinuous payoff functions μ_i [5], we recover the previous result with our Theorem 2. Even if it is not explicitly mentioned in [14], a close look at the details of the proof shows that the authors first show the existence of a weak SPE (without the hypothesis of lower-semicontinuity) and then show that it is indeed an SPE (thanks to this hypothesis). The first part of their proof could be replaced by ours, which is simpler: we remove outcomes from the sets $\lambda_\alpha(v)$ (see the proof of Theorem 1) whereas plays are removed in the inductive process of [14].

4.1 Intermediate results

The proofs of Theorem 2 in this section and Theorem 4 in the next section require several intermediate results that we now describe. We begin with the next lemma where the set $\mu^{-1}(\{o\})$, with $o \in O$, is said to be *dense in (G, v_0)* if for all $h \in Hist(v_0)$, there exists ρ such that $h\rho$ is a play with outcome $\mu(h\rho) = o$.

Lemma 6. *Let (G, v_0) be an initialized game. If for some $o \in O$, the set $\mu^{-1}(\{o\})$ is dense in (G, v_0) , then there exists a weak SPE with outcome o in (G, v_0) .*

Proof. The construction of a very¹³ weak SPE $\bar{\sigma}$ is done step by step thanks to a progressive marking of the histories $hv \in Hist(v_0)$. Let us give the construction of $\bar{\sigma}$. Initially, for history v_0 , we know by density that there exists $\rho_0 \in Plays(v_0)$ with outcome o . We partially construct $\bar{\sigma}$ such that it produces ρ_0 , and we mark each non empty prefix of ρ_0 . Then we consider a shortest unmarked history hv , and we choose some $\rho \in Plays(v)$ such that $\mu(h\rho) = o$ (this is possible by density). We continue the construction of $\bar{\sigma}$ such that it produces the play ρ

¹³ As already done before, we apply Proposition 1. It will be the case in the sequel of the article without mentioning anymore this proposition.

in $(G|_h, v)$, and for each non empty prefix g of ρ , we mark hg (notice that the prefixes of h have already been marked by choice of h), and so on. In this way, we get a strategy profile $\bar{\sigma}$ in (G, v_0) that is a weak SPE because in each subgame $(G|_h, v)$, the play ρ induced by $\bar{\sigma}|_h$ has outcome $\mu(h\rho) = o$ and each one-shot deviating strategy in $(G|_h, v)$ leads to a play with outcome o . \square

Lemma 6 leads to the next two corollaries. The first one states the existence of a uniform weak SPE in each initialized game (G, v) , $v \in V$, when the underlying graph of G is strongly connected and the outcome function is prefix-independent. This corollary will provide a first step towards Theorem 4 presented in Section 5; it is already interesting on its own right.

Corollary 2. *Let G be a game such that the underlying graph is strongly connected and the outcome function μ is prefix-independent.*

- *Then for all realizable outcomes o such that $o = \mu(\rho)$ with $\rho \in \text{Plays}(v_0)$, there exists a weak SPE with outcome o in (G, v_0) .*
- *Moreover, there exists a uniform strategy profile $\bar{\sigma}$ and an outcome o such that for all $v \in V$ taken as initial vertex, $\bar{\sigma}$ is a weak SPE in (G, v) with outcome o .*

Proof. For the first statement, take $\rho \in \text{Plays}(v_0)$ such that $o = \mu(\rho)$. By Lemma 6, it is enough to show that $\mu^{-1}(\{o\})$ is dense in (G, v_0) to get a weak SPE in (G, v_0) . For all $hv \in \text{Hist}(v_0)$, there exists a path πv_0 from v to v_0 as the underlying graph is strongly connected. The play $h\pi\rho$ has outcome equal to $\mu(\rho) = o$ since μ is prefix-independent. Hence $\mu^{-1}(\{o\})$ is dense.

To get the second statement, we need to go further by exhibiting a uniform weak SPE with the same outcome o independently of the initial vertex v . Take any simple cycle $\pi_0 v_0$ from v_0 to v_0 . Such a cycle exists since the underlying graph is strongly connected. Let $\rho = \pi_0^\omega$ and $o = \mu(\rho)$ be its outcome. We partially construct a positional strategy profile $\bar{\sigma}$ that produces π_0^ω (recall that π_0 is simple). Let U be the set of vertices that belong to π_0 . Then extend the construction of $\bar{\sigma}$ to all $v \in V \setminus U$ in a way to reach U (i.e. the cycle π_0) positionally. We then get the required uniform strategy profile $\bar{\sigma}$ with outcome o . \square

The second corollary is a generalization of the previous one. It still guarantees the existence of a uniform weak SPE in all games (G, v) , $v \in V$, for graphs that are not necessarily strongly connected but have bottom strongly connected components all containing a play induced by a simple cycle and with the same outcome. This result will be useful in the proof of Theorem 5 in Section 5.

Corollary 3. *Let G be a game such that the underlying graph is finite and the outcome function μ is prefix-independent. Suppose that there exists an outcome o such that in each bottom strongly connected component C of G , one can find a play $\rho_C \in \text{Plays}(v)$ for some $v \in C$ such that $\mu(\rho_C) = o$ and ρ_C is induced by a simple cycle. Then there exists a uniform weak SPE with outcome o in (G, v) , for all $v \in V$.*

Proof. Let \mathcal{C} be the set of bottom strongly connected components of G . The construction of the strategy profile $\bar{\sigma}$ is very close to the one proposed in the previous proof. We partially construct $\bar{\sigma}$ in a way to produce each ρ_C . This is possible positionally since each ρ_C is induced by a simple cycle. Let U be the set of vertices that belong to $\cup_{C \in \mathcal{C}} \rho_C$. Then extend the construction of $\bar{\sigma}$ to all $v \in V \setminus U$ in a way to reach U positionally. This is possible by definition of \mathcal{C} . The resulting strategy profile $\bar{\sigma}$ is uniform and is a weak SPE in each (G, v) , $v \in V$, such that $\mu(\langle \bar{\sigma} \rangle_v) = o$. Indeed each ρ_C has outcome o and μ is prefix-independent. \square

We end with a last lemma which indicates how to combine different weak SPEs into one weak SPE. It will be used in the proofs of Theorems 2 and 4.

Lemma 7. *Consider an initialized game (G, v_0) and a set of vertices $L \subseteq V$ such that for all $hl \in \text{Hist}(v_0)$ with $l \in L$, the subgame $(G|_h, l)$ has a weak SPE with outcome o_{hl} . Consider another initialized game (G', v_0) obtained from (G, v_0)*

- by replacing all edges $(l, v) \in E$ by one edge (l, l) , for all $l \in L$,
- and with outcome function μ' such that for all $\rho' \in \text{Plays}_{G'}(v_0)$, $\mu'(\rho') = o_{hl}$ if $\rho' = hl^\omega$ with $l \in L$ and $\mu'(\rho') = \mu(\rho')$ otherwise.

If (G', v_0) has a weak SPE, then (G, v_0) has also a weak SPE.

Proof. Denote by $\bar{\sigma}^{hl}$ the weak SPE in each $(G|_h, l)$, and by $\bar{\sigma}'$ the weak SPE in (G', v_0) . We then build a strategy profile $\bar{\tau}$ in (G, v_0) as follows. For player $i \in \Pi$ and history $hv \in \text{Hist}_i(v_0)$:

- if no vertex of L occurs in hv , then $\tau_i(hv) = \sigma'_i(hv)$;
- otherwise, decompose hv as $h_1 h_2 v$ such that the first occurrence of a vertex $l \in L$ is the first vertex of h_2 . Then $\tau_i(hv) = \sigma_i^{h_1 l}(h_2 v)$.

Hence in the first case, τ_i mimics σ'_i in the game (G', v_0) , and in the second case, τ_i mimics $\sigma^{h_1 l}$ in the subgame $(G|_{h_1}, l)$.

Let us show that $\bar{\tau}$ is a weak SPE in (G, v_0) . Consider any subgame $(G|_h, v)$ such that $v \in V_i$, and any one-shot deviation strategy τ'_i of player i from $\bar{\tau}|_h$. Either no vertex of L occurs in hv , and τ'_i is not profitable for player i because $\bar{\sigma}'$ is a weak SPE in (G', v_0) and by definition of μ' . Or $h = h_1 h_2 v$ such that the first occurrence of a vertex $l \in L$ is the first vertex of h_2 , and again τ'_i is not profitable because $\bar{\sigma}^{h_1 l}$ is a weak SPE in the subgame $(G|_{h_1}, l)$. \square

4.2 Proof of Theorem 2

Now that we have established all useful intermediate results for this section and the next one, we can finally proceed to the proof of Theorem 2. W.l.o.g. we can suppose that the underlying graph of G is a tree rooted at v_0 (by unraveling this graph from v_0). We first show how to transform a game played on an infinite

tree to a game satisfying Conditions 1 and 2 from Theorem 1 while reflecting weak SPE.

In the following lemma we write $h \sqsubseteq l$ to denote that h is a prefix of l , and denote by $\text{cl}(A)$ the topological closure of A .

Lemma 8. *Consider a game played on an infinite tree C^ω with countable outcome set O and outcome function $\mu : C^\omega \rightarrow O$. There exists a prefix-free set $L \subseteq C^*$ of leaves and an assignment $\Theta : L \rightarrow O$ such that*

1. *For each $h \in C^*$ there exists some $l \in L$ with $h \sqsubseteq l$ or $l \sqsubseteq h$.*
2. *For each $l \in L$ we find that $\mu^{-1}(\{\Theta(l)\})$ is dense in lC^ω .*

Proof. By iterative use of the Baire Category Theorem. We go through all $h \in C^*$ in some order, add elements to L and extend Θ . Let $h \in C^*$ be the current candidate. If we do not yet have added l to L with $l \sqsubseteq h$ or $h \sqsubseteq l$, then consider that $hC^\omega = hC^\omega \cap \bigcup_{o \in O} \mu^{-1}(\{o\})$. As O is countable, the Baire Category Theorem implies that some $\mu^{-1}(\{o_0\})$ is somewhere dense, i.e. that there exists some $l \sqsupseteq h$ such that $lC^\omega \subseteq \text{cl}(\mu^{-1}(\{o_0\}))$. We add l to L and set $\Theta(l) = o_0$. Then we proceed to the next h . In the limit, we have constructed L and Θ as desired.

To see that L is prefix-free, assume that there are $l_1, l_2 \in L$ with $l_1 \sqsubset l_2$. If l_1 was added first, and l_2 was added when dealing with the history h , then $h \sqsubseteq l_2$. But as prefixes of a given history are linearly ordered, either $h \sqsubseteq l_1$ or $l_1 \sqsubseteq h$ follows. Thus, we would not have added l_2 when dealing with h . If l_2 was added first, and then l_1 when dealing with h , then we would find that $h \sqsubseteq l_1 \sqsubseteq l_2$, thus $h \sqsubseteq l_2$, thus we would not have added l_2 . Hence, L is prefix-free. \square

Proof (of Theorem 2). Instead of reasoning with the underlying graph of G , we work w.l.o.g. with its unraveling from the initial vertex v_0 .

We can apply Lemma 8 to transform the game. For each leaf, we can apply Lemma 6 to obtain a weak SPE in the corresponding subgame. Together, the criteria of Lemma 7 are satisfied. The implication of Lemma 7 is true by Corollary 1, and the conclusion yields the desired statement. \square

5 Second application

In this section, we present a second large family of games with a weak SPE, as another application of the general results of Section 3 (more particularly Theorem 1). This family is constituted with all games with a finite underlying graph and a prefix-independent outcome function.

Theorem 4. *Let (G, v_0) be an initialized game such that the underlying graph is finite and the outcome function is prefix-independent. Then there exists a weak SPE in (G, v_0) .*

Let us comment this theorem. (i) It guarantees the existence of a weak SPE for classical games with *quantitative* objectives as presented in Example 1, such that their outcome function is prefix-independent. This is the case of *limsup*

and *mean-payoff* payoff functions (and their limit inferior counterparts). Recall that Example 2 (see also Figure 1) provides a game with no SPE, where the payoff functions μ_i can be seen as either *limsup* or *mean-payoff* (or their limit inferior counterparts). (ii) Later in this section, we will show that under the hypotheses of Theorem 4, there always exists a weak SPE that is *finite-memory* (Corollary 4), and we will study in which cases it can be *positional* or even *uniform* (Theorem 5). (iii) The families of games of Theorems 2 and 4 are incomparable: Boolean reachability games are in the first family but not in the second one, and mean-payoff games are in the second family but not in the first one.

5.1 Proof of Theorem 4

The proof of Theorem 4 follows the same structure as for Theorem 2. The idea is to apply Lemma 7 where L is equal to the union of the bottom strongly connected components of the graph of G . The weak SPEs required by Lemma 7 exist on the subgames $(G|_h, l)$ with $l \in L$ by Corollary 2, and on the game (G', v_0) thanks to Theorem 1.

Proof (of Theorem 4). Let \mathcal{C} be the set of bottom strongly connected components of the finite graph of G . By Corollary 2, for all $C \in \mathcal{C}$, there exist a uniform strategy profile $\bar{\sigma}_C$ and an outcome o_C such that $\bar{\sigma}_C$ is a weak SPE with outcome o_C in each (G, v) with $v \in C$. Notice that as μ is prefix-independent, $\bar{\sigma}_C$ is also a weak SPE with outcome o_C in all subgames $(G|_h, v)$ with $hv \in \text{Hist}(v_0)$ and $v \in C$.

If the initial vertex v_0 belongs to some $C \in \mathcal{C}$, then $\bar{\sigma}_C$ is the required weak SPE in (G, v_0) (it is clearly finite-memory as it is uniform). From now on we suppose that $v_0 \notin C$ for all $C \in \mathcal{C}$.

We consider the graph (G', v_0) constructed from (G, v_0) as described in Lemma 7 with $L = \cup_{C \in \mathcal{C}} C$. This graph satisfies all the hypotheses of Theorem 1.

The set L of leaves is the one used for Lemma 7. The first hypothesis holds because L is the union of the bottom strongly connected components of G . The second hypothesis holds because μ is prefix-independent. The third hypothesis holds because V is finite. Therefore, (G', v_0) has a weak SPE $\bar{\sigma}'$ by Theorem 1.

By the existence of the previous strategy profiles $\bar{\sigma}'$ and $\bar{\sigma}_C$, $C \in \mathcal{C}$, it follows by Lemma 7 that there exists a weak SPE $\bar{\tau}$ in (G, v_0) . \square

5.2 Finite-memory weak SPE

We here make the statement of Theorem 4 more precise by guaranteeing the existence of a weak SPE with finite-memory.

Corollary 4. *Let (G, v_0) be an initialized game such that the underlying graph is finite and the outcome function is prefix-independent. Then there exists a finite-memory weak SPE in (G, v_0) with memory size bounded by the number of bottom*

strongly connected components of the graph. Moreover, a memory size linear in the number of bottom components is necessary.

Proof. In the proof of Theorem 4, we have constructed a weak SPE $\bar{\tau}$. Let us show that $\bar{\tau}$ is a finite-memory strategy profile with memory size bounded by $|\mathcal{C}|$. Let us first come back to the construction of $\bar{\tau}$ given in the proof of Lemma 7. Consider player $i \in \Pi$ and history $hv \in \text{Hist}_i(v_0)$. If no vertex of L occurs in hv , then $\tau_i(hv) = \sigma'_i(hv)$. Otherwise, decompose hv as h_1h_2v such that the first occurrence of a vertex $l \in C \subseteq L$ is the first vertex of h_2 , then

$$\tau_i(hv) = \sigma_{C,i}(v). \quad (6)$$

Notice that in (6) $\tau_i(hv)$ only depends on C , and not on $l \in C$, since $\bar{\sigma}_C$ is uniform. Now let us recall the construction of $\bar{\sigma}'$ with a memory size $|L|$ given in the proof of Theorem 1, and in particular to equation (4). In (G', v_0) the plays $\rho_{v,o} = \pi l^\omega$ can be produced positionally while keeping $l \in L$ in memory. Therefore by (6) and as $\bar{\sigma}_C$ is uniform, it follows that the memory size of $\bar{\tau}$ can be reduced from $|L|$ to $|\mathcal{C}|$.

Let us now prove that there exist games with a finite set V and a prefix-independent function μ , that require a memory size in $O(|\mathcal{C}|)$ for their weak SPEs. To this end, we come back to the family of games G_n of Example 3 with n bottom strongly connected components. Consider the unravelling of G_n from the initial vertex v_1 as depicted in Figure 3 and let us study the form of any weak SPE $\bar{\sigma}$ in (G_n, v_1) . In all subgames $(G_n|_h, v_i)$, the induced play cannot be $(v_i v_{i+1} \dots v_{i-1})^\omega$ with outcome \perp since each player would have a profitable one-shot deviation. W.l.o.g let us suppose that $\sigma_1(v_1) = l_1$ (player 1 decides to move from v_1 to l_1 at the root of the unravelling, as in Figure 3). Then the outcome of the play ρ induced by $\bar{\sigma}|_{v_1}$ in the subgame $(G_n|_{v_1}, v_2)$ is necessarily o_1 or o_n , otherwise player 1 would have a profitable one-shot deviation in (G_n, v_0) (recall that $o_1 \prec_1 o_j$ for all $j \in \Pi \setminus \{1, n\}$). The first case o_1 cannot occur otherwise player 2 would have a profitable one-shot deviation in $(G_n|_{v_1}, v_2)$ (recall that $o_1 \prec_2 o_2$). With similar arguments one can verify that the induced play ρ is necessarily equal to $v_2 v_3 \dots v_n l_n^\omega$ with outcome o_n (as in Figure 3). We can repeat the same reasoning for the play induced by $\bar{\sigma}|_{v_1 v_2 \dots v_n}$ in the subgame $(G_n|_{v_1 v_2 \dots v_n}, v_1)$ which must be equal to $v_1 v_2 \dots v_{n-1} l_{n-1}^\omega$ with outcome o_{n-1} , aso. Hence all weak SPEs of (G_n, v_1) have the form of the one described in Figure 3 and they have finite memory of size $n - 1$ as explained previously in Example 3 (see also Figure 4). Let us show that such a weak SPE $\bar{\sigma}$ cannot have a memory size $< n - 1$. Assume the contrary: wlog consider the previous weak SPE $\bar{\sigma}$ (as in Figure 3) and in particular a Moore machine $\mathcal{M} = (M, m_0, \alpha_U, \alpha_N)$ encoding σ_1 such that $|M| < n - 1$. Let $h_j v_1$, $j \in \{0, \dots, n - 1\}$ be consecutive histories, with $h_j = (v_1 v_2 \dots v_n)^j$. On one hand, we have $\sigma_1(h_j v_1) = \alpha_N(\hat{\alpha}_U(m_0, h_j), v_1)$ for all j . On the other hand, $\sigma_1(h_0 v_1) = \sigma_1(h_{n-1} v_1) = l_1$ and $\sigma_1(h_j v_1) = v_2$ for all $j \in \{1, \dots, n - 2\}$. Therefore there exists $j_1, j_2 \in \{1, \dots, n - 2\}$, $j_1 \neq j_2$, such that the associated memory state is identical, i.e. $\hat{\alpha}_U(m_0, h_{j_1}) = \hat{\alpha}_U(m_0, h_{j_2})$. Thus \mathcal{M} enters into a cycle while reading the prefixes of $(v_1 v_2 \dots v_n)^\omega$. This

means that \mathcal{M} defines $\sigma_1(hv) = v_2$ for all histories h of which h_1 is prefix, in contradiction with $\sigma_1(h_{n-1}v_1) = l_1$. \square

5.3 Positional weak SPE

In the previous section, Corollary 4 guarantees the existence of a finite-memory weak SPE for games with a finite underlying graph and a prefix-independent outcome function. In this section, we identify conditions on the preference relations of the players, as expressed in the next lemma, that guarantee the existence of a *uniform* weak SPE (see Theorem 5).

Lemma 9 (Lemma 4 of [21]). *Let O be a non empty set of outcomes. Let \prec_i be a preference relation over O , for all $i \in \Pi$. The following assertions are equivalent.*

- For all $i, i' \in \Pi$ and all $o, p, q \in O$, we have $\neg(o \prec_i p \prec_i q \wedge q \prec_{i'} o \prec_{i'} p)$.
- There exist a partition $\{O_k\}_{k \in K}$ of O and a linear order $<$ over K such that
 - $k < k'$ implies $o \prec_i o'$ for all $i \in \Pi$, $o \in O_k$ and $o' \in O_{k'}$,
 - $\prec_{i|O_k} = \prec_{i'|O_k}$ or $\prec_{i|O_k} = (\prec_{i'|O_k})^{-1}$ for all $i, i' \in \Pi$.

In the previous lemma, we call each set O_k a *layer*. The second assertion states that (i) if $k < k'$ then all outcomes in $O_{k'}$ are preferred to all outcomes in O_k by all players, and (ii) inside a layer, any two players have either the same preference relations or the inverse preference relations. When a set of outcomes satisfies the conditions of Lemma 9, we say that it is *layered*. In [21], the author characterizes the preference relations that always yield SPE in games with outcome functions in the Hausdorff difference hierarchy of the open sets. One condition is that the set of outcomes is layered.

Theorem 5. *Let G be a game with a finite underlying graph and such that the outcome function is prefix-independent with a layered set O of outcomes. Then there exists a uniform weak SPE in (G, v) , for all $v \in V$.*

Example 5. Remember the class G_n of games, $n \geq 3$, of Example 3, such that $O = \{o_1, \dots, o_n, \perp\}$ and each player i has a preference relation \prec_i satisfying $\perp \prec_i o_{i-1} \prec_i o_i \prec_i o_j$ for all $j \in \Pi \setminus \{i-1, i\}$. This set of outcomes is not layered because the first assertion of Lemma 9 is not satisfied. Indeed we have

$$o_2 \prec_3 o_3 \prec_3 o_1 \wedge o_1 \prec_2 o_2 \prec_2 o_3.$$

Recall that in the proof of Corollary 4 we have shown that all weak SPEs of G_n require a memory size in $O(n)$. Hence the hypothesis of Theorem 5 about the preference relations is not completely dispensable.

Let us proceed to the proof of Theorem 5. Let \mathcal{C} be the set of the bottom strongly connected components of the finite underlying graph of G . For each $C \in \mathcal{C}$, we fix a play $\rho_C \in \text{Plays}(v)$ for some $v \in C$ induced by a simple cycle. The set $O_C = \{o_C \mid o_C = \mu(\rho_C), C \in \mathcal{C}\}$ is finite. It is layered by hypothesis with a finite partition into layers $\{O_k\}_{k \in K}$. The proof of Theorem 5 is by induction on the number of layers and uses the next lemma dealing with one layer.

Lemma 10. *Suppose that $|K| = 1$, then there exists a uniform strategy profile $\bar{\sigma}$ that is a weak SPE in each (G, v) , $v \in V$, such that $\mu(\langle \bar{\sigma} \rangle_v) = o_C$ for some $C \in \mathcal{C}$.*

The proof of this lemma is by induction on $|O_{\mathcal{C}}|$. The case of only one outcome is solved by Corollary 3. When they are several outcomes in $O_{\mathcal{C}}$, we will show how to decompose G into two subgames G' and G'' such that the bottom strongly connected component of G' (resp. G'') are those components $C \in \mathcal{C}$ of G such that $o_C = o$ for some o (resp. $o_C \in O_{\mathcal{C}} \setminus \{o\}$). By Corollary 3 for G' and by induction hypothesis for G'' , we will get two uniform weak SPEs that can be merged to get a uniform weak SPE for G .

Proof (of Lemma 10). The proof is by induction on $|O_{\mathcal{C}}|$. We solve the basic case $|O_{\mathcal{C}}| = 1$ by Corollary 3. Suppose that $|O_{\mathcal{C}}| = n > 1$. By Lemma 9, we have $\prec_i = \prec_{i'}$ or $\prec_i = \prec_{i'}^{-1}$ for all $i, i' \in \Pi$. We can thus merge the players into two *meta-players* \mathcal{P}_1 and \mathcal{P}_2 with their respective preference relations \prec_1, \prec_2 on $O_{\mathcal{C}}$ satisfying $o_1 \prec_1 o_2 \prec_1 \dots \prec_1 o_n$ and $o_n \prec_2 o_{n-1} \prec_2 \dots \prec_2 o_1$. Notice that \mathcal{P}_2 could not exist.

For the sequel, we need the classical concept of *attractor* of $U \subseteq V$ for \mathcal{P}_1 [17]: it is the set $Attr_1(U)$ composed of all $v \in V$ from which \mathcal{P}_1 can force, against \mathcal{P}_2 , to reach U . More precisely, $Attr_1(U)$ is constructed by induction as follows: $Attr_1(U) = \cup_{k \geq 0} X_k$ such that

$$\begin{aligned} X_0 &= U, \\ X_{k+1} &= X_k \cup \{v \in V \mid v \text{ is controlled by } \mathcal{P}_1 \text{ and } \exists (v, v') \in E, v' \in X_k\} \\ &\quad \cup \{v \in V \mid v \text{ is controlled by } \mathcal{P}_2 \text{ and } \forall (v, v') \in E, v' \in X_k\}. \end{aligned}$$

Let $\mathcal{C}' = \{C \in \mathcal{C} \mid o_C = o_n\}$ and $\mathcal{C}'' = \mathcal{C} \setminus \mathcal{C}'$. We construct a subset V' of V as follows:

1. Initially $V' \leftarrow \cup \{C \mid C \in \mathcal{C}'\}$
2. $V' \leftarrow Attr_1(V')$. Let \mathcal{D} be the set of bottom strongly connected components of $G_{|V \setminus V'}$
3. If \mathcal{D} contains components not in \mathcal{C}'' , then add all of them to V' and goto 2, else stop

At the end of the process, we get two sets V' and $V'' = V \setminus V'$, and the related subgames G' and G'' respectively induced by V' and V'' .

Let us prove by induction on the three steps that (*) for all $v \in V'$, there is a path from v to some $C \in \mathcal{C}'$. To this end, we denote $W = Attr_1(V')$ at step 2 and $T = W \cup \{D \in \mathcal{D} \mid D \notin \mathcal{C}'\}$ at step 3. After step 1, (*) is true (with the empty path from v to v). It is also the case after step 2, since by definition of the attractor, there is a path from $v \in W = Attr_1(V')$ to some $v' \in V'$ for which there is a path to some $C \in \mathcal{C}'$ by induction hypothesis. Consider now $v \in D$ such that $D \in \mathcal{D}$ is added to W in step 3. As D does not belong to \mathcal{C}'' and D is a bottom component of $G_{|V \setminus W}$, then there must exist a path from $v \in D$ to some $C \in \mathcal{C}'$ and (*) holds.

By construction each $C \in \mathcal{C}'$ (resp. $C \in \mathcal{C}''$) is a bottom strongly connected component of G' (resp. G''). Let us prove that neither G' nor G'' contain other

bottom components. Assume the contrary and let v be a vertex belonging to such a bottom component D . By step 3 of the previous process, v cannot belong to V'' . By (*), v cannot belong to V' . Therefore the set of bottom strongly connected components of G' and G'' is equal to \mathcal{C} .

By Corollary 3 for G' and by induction hypothesis for G'' , there exist two uniform strategy profiles $\bar{\sigma}'$ and $\bar{\sigma}''$ respectively on G' and G'' such that $\bar{\sigma}'$ (resp. $\bar{\sigma}''$) is a weak SPE in each (G', v') , $v' \in V'$ (in each (G'', v'') , $v'' \in V''$). Moreover $\mu(\langle \bar{\sigma}' \rangle_{v'}) = o_n$ and $\mu(\langle \bar{\sigma}'' \rangle_{v''}) \in P_{\mathcal{C}} \setminus \{o_n\}$. The required uniform strategy profile $\bar{\sigma}$ on G is built such that $\bar{\sigma}|_{V'} = \bar{\sigma}'$ and $\bar{\sigma}|_{V''} = \bar{\sigma}''$. Let us show that it is a weak SPE in all (G, v) , $v \in V$. Consider first a subgame $(G|_h, v')$ such that $\langle \bar{\sigma}|_h \rangle_{v'}$ is a play in G' and a one-shot deviating strategy using an edge (v', v'') with $v' \in V'$ and $v'' \in V''$. By step 2 (i.e. by definition of the attractor), v' belongs to \mathcal{P}_1 who has no incentive to use (v', v'') since the deviating play goes to G'' for which \mathcal{P}_1 receives an outcome o_m such that $o_m \prec_1 o_n$. Consider next a subgame $(G|_h, v'')$ such that $\langle \bar{\sigma}|_h \rangle_{v''}$ is a play in G'' and a one-shot deviating strategy using an edge (v'', v') with $v' \in V'$ and $v'' \in V''$. By step 2, v'' now belongs to \mathcal{P}_2 who has no incentive to use (v'', v') since he will receive an outcome o_m such that $o_n \prec_2 o_m$. \square

We can now proceed to the proof of Theorem 5, which is by induction on the number of layers of O . The case of one layer is treated in Lemma 10. In case of several layers, we show in the proof how to decompose G into two subgames G' and G'' such that there is only one layer in G' and less layers in G'' than in G . From the two uniform weak SPEs obtained for G' by Lemma 10 and for G'' by induction hypothesis, we construct the required uniform weak SPE for G .

Proof (of Theorem 5). We will prove the theorem by induction on the number of layers and additionally show that for all $v \in V$, $\mu(\langle \bar{\sigma} \rangle_v) = o_C$ for some $C \in \mathcal{C}$. Let $O' \subseteq O_{\mathcal{C}}$ be the highest layer of $O_{\mathcal{C}}$ (with respect to the linear order $<$ over K).

If $O' = O_{\mathcal{C}}$, then there is only one layer and the required uniform strategy profile follows from Lemma 10.

If $O' \subset O_{\mathcal{C}}$, we define $V' \subset V$ composed of all vertices v for which there exists a path from v to some component $C \in \mathcal{C}$ such that $o_C \in O'$ (in particular V' includes all such components), and we let $V'' = V \setminus V'$. We obtain two subgames G' and G'' respectively induced by V' and V'' . By construction of V' , one easily checks that the union of the bottom strongly connected components of G' and G'' is equal to \mathcal{C} . Hence, G' has only one layer (equal to O') and G'' has one layer less than G . It follows (by Lemma 10 and by induction hypothesis) the existence of two strategy profiles $\bar{\sigma}'$ and $\bar{\sigma}''$ respectively on G' and G'' : $\bar{\sigma}'$ is a uniform weak SPE in each (G', v') , $v' \in V'$, such that $\mu(\langle \bar{\sigma}' \rangle_{v'}) \in O'$, and $\bar{\sigma}''$ is a uniform weak SPE in each (G'', v'') , $v'' \in V''$, such that $\mu(\langle \bar{\sigma}'' \rangle_{v''}) \in O \setminus O'$. The required strategy profile $\bar{\sigma}$ on G is built such that $\bar{\sigma}|_{V'} = \bar{\sigma}'$ and $\bar{\sigma}|_{V''} = \bar{\sigma}''$. As in the proof of Lemma 10, we consider crossing edges between G' and G'' . By construction, there is no edge (v'', v') with $v' \in V'$ and $v'' \in V''$ showing that a play starting in G'' remains in G'' . On the contrary, there exist edges

(v', v'') with $v' \in V'$ and $v'' \in V''$, but no player has an incentive to use them in a one-shot deviating strategy since the resulting outcome is in a layer smaller than O' . Therefore, $\bar{\sigma}$ is a weak SPE in each (G, v) . \square

6 A counterexample for countably many players and outcomes

We proceed to give an example of a game without weak SPE. It shows that the requirement of only finitely many leaf-outcomes is not dispensable in Theorem 1 or Theorem 2. In [14, Section 4.3] there is an example of a game in extensive form with countably many players, uncountably many outcomes, preference heights 3, but without weak SPE. Our example is similar, but with only countably many outcomes, one single proper infinite play (i.e. not ending in a leaf), and preferences of height 3.

Example 6. We consider the initialized game (G, v_0) of Figure 5. The set of players is \mathbb{N} . The player i acts at most once, at the vertex v_i , and can either enter the leaf l_i or move onwards to v_{i+1} . The play starts with player 0 at v_0 . The outcome attached to reaching l_i is denoted by $2^i 1^\omega$, the outcome attached to the infinite path $v_0 v_1 v_2 \dots$ is denoted by 0^ω . The preferences of player i are given by $p \prec_i q$ iff $p(i) < q(i)$.

The game (G, v_0) has no SPE. To prove this statement, it is enough to show that there is no very weak SPE by Proposition 1 and since every player only acts one. Assume by contradiction that there exists a very weak SPE $\bar{\sigma}$. In each subgame $(G|_h, v_i)$, the play induced by $\bar{\sigma}$ cannot be the one with outcome 0^ω . Otherwise player i has a profitable one-shot deviating strategy by moving to leaf l_i (by increasing his payoff from 0 to 1). Therefore, for all k , there exists a player $i \geq k$ who moves to leaf l_i . Let i be the first such player. It follows that in $(G|_h, v_i)$, he can increase his payoff from 1 to 2 by moving to v_{i+1} instead to l_i , contradiction.




Fig. 5. A game with no weak SPE

References

1. Dietmar Berwanger. Admissibility in infinite games. In *STACS*, volume 4393 of *Lecture Notes in Computer Science*, pages 188–199. Springer, 2007.
2. Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive synthesis. In *LATA*, volume 9618 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2016.
3. Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. In *CONCUR*, volume 42 of *LIPICs*, pages 100–113. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
4. Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS*, pages 23:1–23:10. ACM, 2014.
5. Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Weak subgame perfect equilibria and their application to quantitative reachability. In *CSL*, volume 41 of *LIPICs*, pages 504–518. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
6. Véronique Bruyère. Computer aided synthesis: A game-theoretic approach. In *DLT 2017*, volume 10396 of *Lecture Notes in Computer Science*, pages 3–35, 2017.
7. Véronique Bruyère, Stéphane Le Roux, Arno Pauly, and Jean-François Raskin. On the existence of weak subgame perfect equilibria. In *FOSSACS*, volume 10203 of *Lecture Notes in Computer Science*, pages 145–161. Springer, 2017.
8. Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Secure equilibria in weighted games. In *CSL-LICS*, pages 26:1–26:26. ACM, 2014.
9. K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365:67–82, 2006.
10. Krishnendu Chatterjee, Laurent Doyen, Emmanuel Filiot, and Jean-François Raskin. Doomsday equilibria for omega-regular games. In *VMCAI*, volume 8318 of *Lecture Notes in Computer Science*, pages 78–97. Springer, 2014.
11. Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. *ACM Trans. Comput. Log.*, 11, 2010.
12. Krishnendu Chatterjee and Thomas A. Henzinger. Assume-guarantee synthesis. In *TACAS*, volume 4424 of *Lecture Notes in Comput. Sci.*, pages 261–275. Springer, 2007.
13. Julie De Pril, János Flesch, Jeroen Kuipers, Gijs Schoenmakers, and Koos Vrieze. Existence of secure equilibrium in multi-player games with perfect information. In *MFCS*, volume 8635 of *Lecture Notes in Comput. Sci.*, pages 213–225. Springer, 2014.
14. János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan, and Koos Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Math. Oper. Res.*, 35:742–755, 2010.
15. János Flesch and Arkadi Predtetchinski. A characterization of subgame perfect equilibrium plays in Borel games of perfect information. *Math. Oper. Res.*, 2017. To appear.
16. Drew Fudenberg and David Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31:251–268, 1983.
17. Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.

18. Erich Grädel and Michael Ummels. Solution Concepts and Algorithms for Infinite Multiplayer Games. In *New Perspectives on Games and Interaction*, volume 4, pages 151–178. Amsterdam University Press, 2008.
19. H.W. Kuhn. Extensive games and the problem of information. *Classics in Game Theory*, pages 46–68, 1953.
20. Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. Synthesis with rational environments. *Ann. Math. Artif. Intell.*, 78(1):3–20, 2016.
21. Stéphane Le Roux. Infinite subgame perfect equilibrium in the Hausdorff difference hierarchy. In *TTCS*, volume 9541 of *Lecture Notes in Computer Science*, pages 147–163. Springer, 2015.
22. Stéphane Le Roux and Arno Pauly. Infinite sequential games with real-valued payoffs. In *CSL-LICS*, pages 62:1–62:10. ACM, 2014.
23. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
24. Roger A. Purves and William D. Sudderth. Perfect information games with upper semicontinuous payoffs. *Math. Oper. Res.*, 36(3):468–473, 2011.
25. X.S. Shen, H. Yu, J. Buford, and M. Akon. *Handbook of Peer-to-Peer Networking*. Springer, 2010.
26. Eilon Solan and Nicolas Vieille. Deterministic multi-player Dynkin games. *Journal of Mathematical Economics*, 39:911–929, 2003.
27. Michael Ummels. Rational behaviour and strategy construction in infinite multi-player games. In *FSTTCS*, volume 4337 of *Lecture Notes in Computer Science*, pages 212–223. Springer, 2006.