On the existence of weak subgame perfect equilibria

Véronique Bruyère[a,*], Stéphane Le Roux[b,1,2], Arno Pauly[b,1,3], Jean-François Raskin[b,1]

[a] *Computer Science Department, Université de Mons (UMONS), place du Parc 20, B-7000-Mons, Belgium*
[b] *Computer Science Department, Université Libre de Bruxelles (ULB), Campus de la Plaine, CP212, B-1050-Bruxelles, 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.

Keywords: Multi-player non zero-sum graph game, subgame perfect equilibrium, synthesis of equilibrium

1. Introduction

Games played on graphs have a large number of applications in theoretical computer science. One particularly important application is *reactive synthesis* [1], i.e. the design of a controller that guarantees a good behavior of a

[*] Corresponding author
Email addresses: Veronique.Bruyere@umons.ac.be (Véronique Bruyère), leroux@mathematik.tu-darmstadt.de (Stéphane Le Roux), a.m.pauly@swansea.ac.uk (Arno Pauly), jraskin@ulb.ac.be (Jean-François Raskin)
[1] Author supported by ERC Starting Grant (279499: inVEST).
[2] Stéphane Le Roux has since moved to LSV, ENS Cachan, France.
[3] Arno Pauly has since moved to Swansea University, United Kingdom.

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.

The concept of *Nash equilibrium* (NE) [2] is central to the study of *multi-player non zero-sum games*. A strategy profile is an NE if no player has an incentive to deviate unilaterally from his strategy, i.e., he cannot strictly improve the outcome of the strategy profile by changing his strategy only. However in the context of games played on graphs, which are sequential by nature, it is well-known that NEs present a serious drawback: they allow for *non-credible threats* that rational players should not carry out [3]. Thus the notion of NE has been strengthened into the notion of *subgame perfect equilibrium* (SPE) [4]: a strategy profile is an SPE if it is an NE in each subgame of the original game. This notion behaves better for sequential games and excludes non-credible threats. See [5] for a survey.

Variants of SPE, *weak SPE* and *very weak SPE*, have been recently introduced in [6]. While an SPE must be resistant to any unilateral deviation of one player, a weak (resp. very weak) SPE must be resistant to such deviations where the deviating strategy differs from the original one on a *finite number* of histories only (resp. on the *initial vertex* only). The latter class of deviating strategies is a well-known notion that for instance appears in [7] with the one-step deviation property. Weak SPEs and very weak SPEs are equivalent, but there are games for which there exists a weak SPE but no SPE [6, 8]. The notion of weak SPE is important for several reasons (more details are given 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 [6] and [9]. Third, being immune to strategies that finitely deviate from the initial strategy profile may be sufficient from a practical point of view. Indeed ruling out infinite deviations can be achieved 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 [10] for details).

In this paper, we provide the following contributions. First, we identify *general conditions* to guarantee the existence of a weak SPE (Theorem 7). The result identifies a large class of multi-player non zero-sum games such that an outcome is assigned to every play of the game and each player has a preference relation on the set of play outcomes which allows him to compare plays. This class covers game graphs that may have infinitely many vertices and infinitely many players. Notice that such models are relevant for systems where the players can join or leave the game dynamically, and the number of players is finite yet unbounded overtime: the users in the Internet are a typical example since there is no (clear) bound on the number of possible users. The proof of our result relies on transfinite induction and additionally provides a weak SPE using finite-memory strategies for all players. Second, starting from this general existence result, we prove the existence of a weak SPE:

- for games with a *finite* number of outcomes (Theorem 16, reproving a result from [9]);
- for games with a *finite* underlying graph and a *prefix-independent* outcome function (Theorem 22).

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 25).

Related work. The concept of SPE has been first introduced and studied by the game theory community. In [7], Kuhn proves the existence of SPEs in games played on finite trees using a backward induction procedure and the one-step deviation property. This existence 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 [11] (the special case with finitely many players is first established in [12]). In [9] (resp. [13]), 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 of finite range. The result of [13] is extended to an infinite number of players in [14]. In [11], 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 [15], 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^0_2 measurable (a low level in the Borel hierarchy) have an SPE. In [16], the authors provide topological assumptions on large extensive form games with perfect information under which the strategy combinations fulfilling the backwards induction criterion used in [7] provides an SPE.

In part of the aforementioned work, the equivalence between SPEs and very weak SPEs is implicitly used as a proof technique: in a finite setting in [7], in a continuous setting in [12], and in a lower-semicontinuous setting in [9]. In the latter reference, it is implicitly proven that all games with a finite range real-valued outcome function 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 science community, see [17] for a survey. The existence of SPEs (and thus weak SPEs) is established in [18] for games played on graphs by a finite number of players and with Borel Boolean objectives. In [6], 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 [6] as they are concerned with possibly infinite graphs or prefix-independent outcome functions.

Other refinements of NE are studied. Let us mention the secure equilibria for two players first introduced in [19] and then used for reactive synthesis in [20]. These equilibria are generalized to multiple players in [21] or to quantitative objectives in [22], see also a variant called Doomsday equilibrium in [23]. Like NEs, they are subject to possible non-credible threats. Other alternatives to NE are provided by the notion of admissible strategy introduced in [24], with computational aspects studied in [25], and potential for synthesis studied in [26]. Note that these notions are free, like (weak) SPEs, of non-credible threats. Finally, in [27], 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 [22] and [21], 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.

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. A conclusion is provided in the last section.

An extended abstract omitting most proofs has appeared as [28].

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* $\mu : V^\omega \to O$ *is an outcome function,*
- $\prec_i \subseteq O \times O$ *is a preference relation for player* $i \in \Pi$ *(that is, an irreflexive and transitive binary relation).*

In this definition, the sets V*,* Π *and* O *can be of arbitrary cardinality (finite or infinite).*

A *play* of G is an infinite (countable) sequence $\rho = \rho_0\rho_1 \ldots \in V^\omega$ of vertices such that $(\rho_i, \rho_{i+1}) \in E$ for all $i \in \mathbb{N}$. *Histories* of G are finite sequences $h = h_0 \ldots h_n \in V^+$ 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 $Plays$, the set of histories by $Hist$, and the set of histories ending with a vertex in V_i by $Hist_i$. When it is necessary to recall the related game G, we index $Plays_G$, $Hist_G$, or $Hist_{i,G}$ with G. A *prefix* of a play $\rho = \rho_0\rho_1 \ldots$ is a finite sequence $\rho_{\leq n} = \rho_0 \ldots \rho_n$, and a *suffix* of ρ is an infinite sequence $\rho_{\geq n} = \rho_n\rho_{n+1} \ldots$. We use notation $h < \rho$ when a history h is prefix of a play ρ. We endow $Plays$ with the topology induced by the cylinder sets $C(h) = \{\rho \in Plays \mid h < \rho\}$, for $h \in Hist$.

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 $Plays(v_0)$, $Hist(v_0)$, and $Hist_i(v_0)$. 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 $\rho \in Plays$. It is *prefix-independent* if $\mu(h\rho) = \mu(\rho)$ for all histories h and plays ρ. A preference relation $\prec_i \subseteq O \times O$ 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 each preference relation which is not linear can be extended into a linear one, and 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 2. 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 $\mu_i : Plays \to \mathbb{R} \cup \{+\infty, -\infty\}$. 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 $\mu_i : Plays \to \{0, 1\}$ 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, and so on [29]. Prefix-independence of μ_i holds in the case of Büchi and parity objectives, but not for reachability objective.

We have *quantitative* objectives when $\mu_i : Plays \to \mathbb{R} \cup \{+\infty, -\infty\}$ replaces $\mu_i : Plays \to \{0, 1\}$. Usually, such a μ_i is defined from a weight function $w_i : E \to \mathbb{R}$ that assigns a weight to each edge. Classical examples of μ_i are *limsup* and *mean-payoff* functions [30], that is,

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

Variants of these functions also exist that use the limit inferior instead of the limit superior.

2.2. Strategies

Let (G, v_0) be an initialized game. A *strategy* σ_i for player i in (G, v_0) is a function $\sigma_i : Hist_i(v_0) \to V$ assigning to each history $hv \in Hist_i(v_0)$ a vertex $v' = \sigma_i(hv)$ such that $(v, v') \in E$. Such a strategy σ_i is *positional* if it only depends on the last vertex of the history, i.e. $\sigma_i(hv) = \sigma_i(h'v)$ for all $hv, h'v \in Hist_i(v_0)$. 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 strategy), $m_0 \in M$ is an initial memory state, $\alpha_U : M \times V \to M$ is an update function, and $\alpha_N : M \times V_i \to V$ is a next-move function.[4] Such a machine defines a strategy σ_i such that $\sigma_i(hv) = \alpha_N(\widehat{\alpha}_U(m_0, h), v)$ for all histories $hv \in Hist_i(v_0)$, where $\widehat{\alpha}_U$ extends α_U to histories as expected. The *memory size* of σ_i is then the size $|M|$ of $\mathcal{M}$. In particular σ_i 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 strategy σ_i of player i defined on $Hist_i$ (instead of $Hist_i(v_0)$) such that $\sigma_i(hv) = \sigma_i(h'v)$ for all $hv, h'v \in Hist_i$. In other words, a strategy $\sigma_i : Hist_i \to V$ is uniform if its restriction to each initialized game (G, v_0), $v_0 \in V$, is a positional strategy.

A play ρ is *consistent* with a strategy σ_i of player i if $\rho_{n+1} = \sigma_i(\rho_{\leq n})$ 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,

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

where each σ_i is a strategy of player i. It is called *positional* if all σ_i, $i \in \Pi$, are positional. Similarly we say that it is *uniform* if all σ_i, $i \in \Pi$, are uniform, and that it is *finite-memory with memory size bounded by* c if all σ_i, $i \in \Pi$, are finite-memory with memory size bounded by c.

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, we say that σ'_i is a *finitely-deviating* strategy from σ_i. And when σ_i and σ'_i only differ on v_0, we say that σ'_i is *one-shot deviating* strategy from σ_i. One-shot deviating strategies is a well-known notion that for instance appears in [7] with the one-step deviation property. Finitely-deviating strategies have been introduced in [6]. Notice that a strategy σ'_i that is a one-shot deviating strategy from a positional strategy σ_i might not be positional.

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 3. *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 an initialized game (G, v_0) with $G = (\Pi, V, (V_i)_{i \in \Pi}, E, \mu, (\prec_i)_{i \in \Pi})$ and a history $hv \in \mathit{Hist}(v_0)$, we denote by $(G_{|h}, v)$ the initialized game such that $G_{|h} = (\Pi, V, (V_i)_{i \in \Pi}, E, \mu_{|h}, (\prec_i)_{i \in \Pi})$ where $\mu_{|h}(\rho) = \mu(h\rho)$ for all plays $\rho \in \mathit{Plays}(v)$[5] . This game $(G_{|h}, v)$ is called the *subgame of* (G, v_0) *with history* hv. Hence in $(G_{|h}, v)$, only the initial vertex and the outcome function changed (compared to (G, v_0)). Moreover the unravelling of $(G_{|h}, v)$ from v is the subtree rooted at v after history h in the unravelling of (G, v_0). Notice that 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 $h' \in \mathit{Hist}_i(v)$. 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 $(G_{|h}, v)$.

---

[5]In this article, we will use both notations $\mu(h\rho)$ and $\mu_{|h}(\rho)$.

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 [6] 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 4. *Given an initialized game (G, v_0), a strategy profile $\bar{\sigma}$ of (G, v_0) is a* subgame perfect equilibrium *if for all histories $hv \in Hist(v_0)$, for all players $i \in \Pi$, for all deviating strategies σ'_i from $\sigma_{i|h}$ of player i in the subgame $(G_{|h}, v)$, we have $\mu(\langle \bar{\sigma}_{|h} \rangle_v) \not\prec_i \mu(\langle \sigma'_i, \bar{\sigma}_{-i|h} \rangle_v)$.*
If the deviating strategies σ'_i are restricted to finitely deviating ones, the strategy profile $\bar{\sigma}$ is a weak subgame perfect equilibrium*, and if they are restricted to one-shot deviating strategies, then $\bar{\sigma}$ is a* very weak subgame perfect equilibrium*.*

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

Proposition 5 ([6]). *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.*

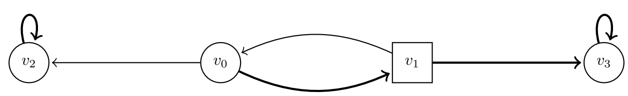

Figure 1: A initialized game (G, v_0) with a (very) weak SPE and no SPE.

Example 6 ([6]). Consider the two-player game (G, v_0) in Figure 1 such that player 1 controls vertices v_0, v_2, v_3 and player 2 controls 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 is $o_1 \prec_1 o_2 \prec_1 o_3$ and the preference relation for player 2 is $o_2 \prec_2 o_3 \prec_2 o_1$.

Let us study the positional strategy profile $\bar{\sigma} = (\sigma_1, \sigma_2)$ such that for all histories $hv_0, hv_2, hv_3 \in Hist(v_0)$, we have $\sigma_1(hv_0) = v_1$, $\sigma_1(hv_2) = v_2$, and $\sigma_1(hv_3) = v_3$, and for all histories $hv_1 \in Hist(v_0)$, we have $\sigma_2(hv_1) = v_3$. This strategy profile is depicted with thick edges on the unravelling T of G from v_0 (see Figure 2). As it is positional, it can also be depicted directly on the graph of the game (see Figure 1). Let us explain why $\bar{\sigma}$ is a very weak SPE (and thus a weak SPE by Proposition 5). Due to the simple form of the game, only two cases are to be treated:

1. Consider first the subgame $(G_{|h}, v_0)$ with $h \in (v_0v_1)^*$, and its unravelling that is the subtree of T reached after hv_0. Consider in this subtree the one-shot deviating strategy σ'_1 from $\sigma_{1|h}$ such that $\sigma'_1(v_0) = v_2$ (σ'_1 is equal

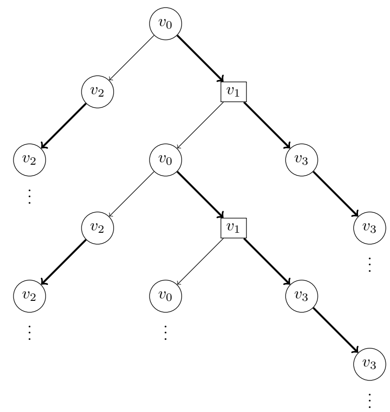

Figure 2: The unravelling of the game of Figure 1 from vertex v_0.

to $\sigma_{1|h}$ elsewhere). Then $\langle\bar{\sigma}_{|h}\rangle_{v_0} = v_0v_1v_3^\omega$ and $\langle\sigma'_1, \sigma_{2|h}\rangle_{v_0} = v_0v_2^\omega$ with respective outcomes o_3 and o_2. As $o_2 \prec_1 o_3$, this shows that σ'_1 is not a profitable deviation for player 1 in $(G_{|h}, v_0)$.

2. Now in the subgame $(G_{|h}, v_1)$ with $h \in (v_0v_1)^*v_0$, the one-shot deviating strategy from $\sigma_{2|h}$ such that $\sigma'_2(v_1) = v_0$ is not profitable for player 2 in $(G_{|h}, v_1)$. Indeed $\langle\bar{\sigma}_{|h}\rangle_{v_1} = v_1v_3^\omega$ and $\langle\sigma_{1|h}, \sigma'_2\rangle_{v_1} = v_1v_0v_1v_3^\omega$ have the same outcome o_3.

Notice that $\bar{\sigma}$ is not an SPE. Indeed in the subgame (G, v_0), the strategy σ'_2 such that $\sigma'_2(hv_1) = v_0$ for all $hv_1 \in Hist(v_0)$, is infinitely deviating from σ_2 since $\sigma'_2(hv_1) = v_0 \neq v_3 = \sigma_2(hv_1)$ for an infinite number of histories. It is a profitable deviation for player 2 because $\langle\bar{\sigma}\rangle_{v_0} = v_0v_1v_3^\omega$ and $\langle\sigma_1, \sigma'_2\rangle_{v_0} = (v_0v_1)^\omega$ with respective outcomes o_3 and o_1 such that $o_3 \prec_2 o_1$.

It is proved in [8] that the game (G, v_0) has no SPE.

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 7. *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 leaf $l \in L$ that is reachable from v,*
2. *for all plays $\rho = hl^\omega$ with $h \in Hist$ 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.*

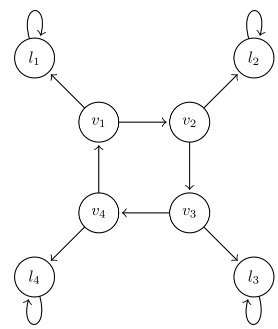

Figure 3: Game G_4

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 implies that for all $v \in V$, there exists a play hl^ω for some history $h \in Hist(v)$ and some leaf $l \in L$, and thus that 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 7.

Example 8. 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 $\Pi = \{1, 2, \ldots, n\}$ and the set of vertices is $V = \{v_1, \ldots, v_n, l_1, \ldots l_n\}$ such that $V_i = \{v_i, l_i\}$ for all $i \in \Pi$. The edges are $(v_1, v_2), (v_2, v_3), \ldots, (v_n, v_1)$, and $(v_i, l_i), (l_i, l_i)$ for all $i \in \Pi$. The game G_4 is depicted in Figure 3. The set O of outcomes is equal to $\{o_1, \ldots, o_n, \bot\}$, and the outcome function is prefix-independent such that $\mu((v_1 v_2 \ldots v_n)^\omega) = \bot$ and $\mu(l_i^\omega) = o_i$ for all $i \in \Pi$. Each player i has a preference relation $\prec_i$ satisfying $\bot \prec_i o_{i-1} \prec_i o_i \prec_i o_j$ for all $j \in \Pi \setminus \{i-1, i\}$ (with the convention that $o_0 = o_n$). This non linear preference relation $\prec_i$ can be arbitrarily extended into a linear one.

Each game (G_n, v_1) satisfies the hypotheses of Theorem 7 with $L = \{l_1, \ldots, l_n\}$ and thus has a finite-memory weak SPE. Such a strategy profile $\bar{\sigma}$ is depicted in Figure 4 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 5. 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 $(v_1 \ldots v_n)^\omega$, 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

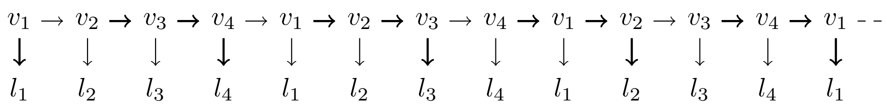

Figure 4: Weak SPE in (G_4, v_1)

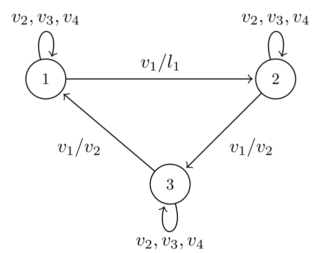

Figure 5: The Moore machine for σ_1

in the game (G_4, v_1) is depicted in Figure 5 (with $M = \{1, 2, 3\}$, $m_0 = 1$, and the update and next-move functions indicated by the edges).

Let us now proceed to the proof of Theorem 7. Recall that it is enough to prove the existence of a very weak SPE by Proposition 5. The proof idea is the following one.

- Initially, for each vertex v, we accept all plays $\rho = hl^\omega$ with $h \in Hist(v)$ 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 7). Notice that this labeling is finite by the third condition of Theorem 7 and it is not empty by the first condition of this 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 converges 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 9). 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 $h \in Hist(v)$ and $l \in L$. Moroever, for all edges (v, v'), there exists $\rho_{v',o'}$ such that $o \not\prec_v o'$ by definition of the fixpoint.

- We can thus build a very weak SPE $\bar{\sigma}$ in (G, v_0) by aggregating some of these plays $\rho_{v,o}$ 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)* Suppose that in the subgame $(G_{|h}, v)$, the partial construction of $\bar{\sigma}$ has been done to induce a play $\langle\bar{\sigma}_{|h}\rangle_v$ with outcome o. If the player who controls v chooses the move (v, v') in a one-shot deviation, then we extend the construction of $\bar{\sigma}$ such that $\langle\bar{\sigma}_{|hv}\rangle_{v'} = \rho_{v',o'}$ such that $o \not\prec_v o'$ (this is possible by the definition of the fixpoint). Such a deviation will not be profitable for the deviating player.

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 $Succ(v)$ 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, $Succ(v) = \emptyset$ for all $v \in L$ and $Succ(v) \neq \emptyset$ 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:

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

(In particular $\lambda_0(l) = \{o_l\}$ for all $l \in L$). By the first hypothesis of the theorem, we have that $\lambda_0(v) \neq \emptyset$. Let us introduce some additional terminology. At step α, when there is a path 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

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

When $Succ(v) = \emptyset$ or when $\lambda_\alpha(v') = \emptyset$ for all $v' \in Succ(v)$, we use the convention that $m_\alpha(v) = \bot$ and we suppose that $\bot \prec_v o$ for all $o \in O_L$. When $m_\alpha(v) \neq \bot$, we say 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 inductively[6]. In the following, γ is always assumed to be a limit ordinal and n to be a natural number.

[6]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.

- **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 $\lambda_\gamma(v) = \cap_{\beta<\gamma}\lambda_\beta(v)$ 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 Knaster-Tarski theorem for monotone functions). 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 fixpoint. If V is finite, such a fixpoint is reached after at most $2 \cdot |O_L| \cdot |V|$ steps. The central challenge is to show that this fixpoint 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 9. *There exists an ordinal α^* such that*

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

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

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 10) and remain true after each step α (Lemmata 11-13). 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 $m_\alpha(v) \neq \bot$, for each v' that realizes $m_\alpha(v)$, we have

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

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 10. *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. □

Lemma 11. *All three invariants are preserved by Remove.*

Roughly speaking, the proof of this lemma is as follows. Consider *odd* step $\alpha + 1$ and the *Remove* operation. *(i) Remove* may remove from $\lambda_\alpha(v)$ only outcomes less than $m_\alpha(v)$, so it preserves INV1. *(ii) Remove* may remove only one outcome at only one vertex, so it preserves INV2 by (1). *(iii) Remove* preserves INV3. Indeed first note that *Remove* might only hurt INV3 at the vertex v subject to outcome removal. Let $v' \in Succ(v)$ that realizes $m_{\alpha+1}(v)$. By INV3 at step α there is a suitable path from v' to a leaf. Prefixing this path with v witnesses INV3 at step $\alpha + 1$, using (1).

PROOF (OF LEMMA 11). 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

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

  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 $Succ(v) \neq \emptyset$[7] and by INV2, we have $m_\alpha(v) \neq \bot$. 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 $\lambda_{\alpha+1}(v) \neq \emptyset$.

- ***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 $\lambda_{\alpha+1}(w) \subseteq \lambda_\alpha(w) \subseteq \lambda_\alpha(v') = \lambda_{\alpha+1}(v') \subseteq \lambda_{\alpha+1}(v)$.

□

Lemma 12. *All three invariants are preserved by Adjust.*

The proof of this lemma works as follows. Consider *even* step $\alpha+1$ and the *Adjust* operation. *(i)* One checks that *Adjust* preserves INV1 by case splitting on whether $\lambda_{\alpha+1}(v) = \lambda_\alpha(v)$. *(ii)* By contradiction assume that $\lambda_\alpha(v) = \{o\}$ from which *Adjust* removes o. By INV3 there would be at prior step one path to a leaf labelled all along with o only. Such labels cannot be removed, leading to a contradiction. *(iii) Adjust* preserves INV3. Indeed from a vertex u_1 let $u_1 \dots u_n$ be a suitable path at step α. If it is no longer suitable at step $\alpha + 1$, some o was removed from some proper prefix $u_1 \dots u_{i-1}$, *i.e.* $o \in \lambda_{\alpha+1}(u_i)$ but $o \notin \lambda_{\alpha+1}(u_{i-1})$, so $o \notin \lambda_\alpha(u_{i-1})$ by definition of *Adjust*. INV3 provides a suitable path (void of o) from u_{i-1} at step α. Concatenating it with $u_1 \dots u_{i-1}$ witnesses INV3 at step $\alpha + 1$.

Proof (of Lemma 12). 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

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

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

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 $u_1 \ldots u_n$ 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 $u_1 \ldots u_{i-1}$ 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 $\lambda_{\alpha+1}(w') \subseteq \lambda_\alpha(v) \setminus \{o\} = \lambda_{\alpha+1}(v)$.

□

Lemma 13. *If all three invariants are true for each λ_β, $\beta < \alpha$, α a limit ordinal, then they are true for λ_α.*

The proof is easier than the previous two. Consider a limit-ordinal step α. Such a step is not explicitly removing outcomes, it is only summarizing what has been removed for lesser ordinals. Indeed for each vertex v, since the sets $\lambda_\beta(v)$ are finite, there is a last outcome removal occurring at some step $\beta < \alpha$. This helps proving that the invariants are indeed preserved at ordinal steps.

Proof (of Lemma 13). 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 finite[8] 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 β, $\gamma \leq \beta < \alpha$. Therefore

$$\lambda_\alpha(v) = \cap_{\beta<\alpha}\lambda_\beta(v) = \lambda_\gamma(v). \tag{3}$$

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

$\{o \in \lambda_\alpha(v') \mid m_\alpha(v) \preceq_v o\}$	
$\subseteq \{o \in \lambda_\gamma(v') \mid m_\gamma(v) \preceq_v o\}$	since $\lambda_\alpha(v') \subseteq \lambda_\gamma(v')$ and $m_\gamma(v) \preceq_v m_\alpha(v)$,
$\subseteq \lambda_\gamma(v)$	by INV1 at step γ,
$= \lambda_\alpha(v)$	by (3).

□

Let us summarize what has been done so far. Each vertex v is initially labelled by the set $\lambda_0(v)$ of the outcomes $\mu(l^\omega)$ for the leaves $l \in L$ that v can reach, meaning that all plays $\rho = hl^\omega$ with $h \in Hist(v)$ are initially accepted as potential plays induced by a very weak SPE in (G, v). Then an iterative process based on a transfinite induction on α

- removes from some $\lambda_\alpha(v)$ an outcome o that can be improved by a one-shot deviation, that is, such that there exists $v' \in Succ(v)$ with $o \prec_v o'$ for all $o' \in \lambda_\alpha(v')$ (by the *Remove* operation),
- and propagates the removal of such an outcome o to the labeling of the other vertices u for which all potential plays induced by a very weak SPE with outcome o in (G, u) necessarily cross vertex v (by the *Adjust* operation).

The limit of this process has been shown non-empty, leading to a not empty fixpoint $\lambda_{\alpha^*}(v) \neq \emptyset$ for all $v \in V$. Therefore for each vertex v and outcome $o \in \lambda_{\alpha^*}(v)$, a potential play $\rho_{v,o} = hl^\omega$ with $o = \mu(l^\omega)$ has survived. Moroever, for all $v' \in Succ(v)$, there exist $o' \in \lambda_{\alpha^*}(v')$ and $\rho_{v',o'}$ such that $o \not\prec_v o'$ by the fixpoint.

To get Theorem 7, it remains to explain how to build a finite-memory weak SPE $\bar{\sigma}$ in (G, v_0). We build $\bar{\sigma}$ by aggregating some of the plays $\rho_{v,o}$ as follows. The construction of $\bar{\sigma}$ is done step by step. Initially $\bar{\sigma}$ is partially defined such that $\langle\bar{\sigma}\rangle_{v_0} = \rho_{v_0,o_0}$ for some $o_0 \in \lambda_{\alpha^*}(v_0)$. Then suppose that in the subgame $(G_{|h}, v)$, the partial construction of $\bar{\sigma}$ has been done to induce a play $\langle\bar{\sigma}_{|h}\rangle_v$ with outcome $o \in \lambda_{\alpha^*}(v)$. If the player who controls v chooses the move

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

(v, v') in a one-shot deviation, then we extend the construction of $\bar{\sigma}$ such that $\langle \bar{\sigma}_{|hv} \rangle_{v'} = \rho_{v',o'}$ such that $o \not<_v o'$ (this is possible by the definition of the fixpoint). Such a deviation will not be profitable for the deviating player.

Proof (of Theorem 7). By Lemma 9, we have a fixpoint of *Remove* and *Adjust* such 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:

$$\rho_{v,o} = \pi l^\omega. \tag{4}$$

(*) 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 $\kappa : Hist(v_0) \to O_L$ will allow to recover from history hv the outcome o of the play $\langle \bar{\sigma}_{|h} \rangle_v$ induced by $\bar{\sigma}$ in the subgame $(G_{|h}, v)$.

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 $\langle \bar{\sigma}_{|h} \rangle_v$ with outcome o in the subgame $(G_{|h}, v)$, such that $\langle \bar{\sigma}_{|h} \rangle_v$ 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

$$o \not<_v o'. \tag{5}$$

With $\rho_{v',o'}$ as in (4), we then extend the construction of $\bar{\sigma}$ such that $\langle \bar{\sigma}_{|h'} \rangle_{v'} = \rho_{v',o'}$, 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 $\rho = \langle \bar{\sigma}_{|h} \rangle_v$ and $\rho' = \langle \bar{\sigma}_{|h'} \rangle_{v'}$. 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 5.

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. □

The next corollary is an easy consequence of Theorem 7. 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 14. *Let (G, v_0) be an initialized game with a subset $L \subseteq V$ of leaves such that the underlying graph is a tree rooted at v_0 when the self-loops (l, l), $l \in L$, are removed. If (G, v_0) satisfies the first and third conditions of Theorem 7, then there exists a positional weak SPE in (G, v_0).*

Proof. If the second condition of Theorem 7 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 all the conditions of Theorem 7 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. □

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 7 and its Corollary 14. Before that, we demonstrate the argument establishing Theorem 7 on the game G_4 as introduced in Example 8.

Example 15. Let us describe the inductive process for the game G_4 of Figure 3 (Page 10). 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 $\rho_{v,o}$.

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 4. 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 $o_1 \not\prec_1 o_4$, and so on. Notice that the previous proof states a memory size equal to 4 for $\bar{\sigma}$ whereas Figure 5 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 14): when an initialized game has an outcome function with finite range, then it always has a weak SPE.

Theorem 16. *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).*

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 fixpoint for game G_4

Let us comment this theorem. *(i)* Kuhn's theorem [7] 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 16 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 [9] for outcome functions $\mu = (\mu_i)_{i \in \Pi}$ as presented in Example 2 and has strong relationship with Theorem 16. Recall that a payoff function $\mu_i : Plays \to \mathbb{R}$ is *lower-semicontinuous* if whenever a sequence of plays $(\rho_n)_{n \in \mathbb{N}}$ converges to a play $\rho = \lim_{n \to \infty} \rho_n$, then $\liminf_{n \to \infty} \mu_i(\rho_n) \geq \mu_i(\rho)$. (Convergence means that for every history $h < \rho$, there exists k such that for all $n \geq k$, we have $h < \rho_n$.)

Theorem 17 ([9]). *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 $\mu_i : Plays \to \mathbb{R}$ 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 [6], we recover the previous result with our Theorem 16. Even if it is not explicitly mentioned in [9], 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 7) whereas plays are removed in the inductive process of [9].

4.1. Intermediate results

The proofs of Theorem 16 in this section and Theorem 22 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 18. *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).*

Such a weak SPE $\bar{\sigma}$ is constructed as follows. By density of $\mu^{-1}(\{o\})$, it is possible to construct $\bar{\sigma}$ such that in each subgame $(G_{|h}, v)$, the play ρ induced by $\bar{\sigma}_{|h}$ has outcome $\mu(h\rho) = o$. Therefore, all one-shot deviations lead to a play with the same outcome o as ρ, and are therefore not profitable for the deviating player.

Proof (of Lemma 18). The construction of a very weak (and thus 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 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. □

Lemma 18 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 22 presented in Section 5; it is already interesting on its own right.

Corollary 19. *Let G be a game such that the underlying graph is strongly connected and the outcome function μ is prefix-independent.*

- *Then for all outcomes o such that $o = \mu(\rho)$ with $\rho \in Plays(v_0)$, 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.*

We obtain the first statement of this corollary by showing that $\mu^{-1}(\{o\})$ is dense in (G, v_0). For the second statement, we construct a weak SPE $\bar{\sigma}$ such that the play induced by $\bar{\sigma}_{|h}$ in $(G_{|h}, v)$ has the same outcome for all $hv \in Hist(v_0)$. Pick a simple cycle from v_0 to v_0 and construct $\bar{\sigma}$ that induces a play iterating this cycle. Extend this construction to the vertices outside this cycle in a way to reach the cycle by a shortest path and then by iterating it.

Proof (of Corollary 19). For the first statement, take $\rho \in Plays(v_0)$ such that $o = \mu(\rho)$. By Lemma 18, 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 $hv \in Hist(v_0)$, 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. □

The second corollary of Lemma 18 is a generalization of Corollary 19. 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 25 in Section 5. The proof of this corollary is very close to the proof of Corollary 19.

Corollary 20. *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 $\rho_C \in Plays(v_C)$ for some $v_C \in C$ 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. □

We end Section 4.1 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 16 and 22.

Lemma 21. *Consider an initialized game (G, v_0) and a set of vertices $L \subseteq V$ such that for all $hl \in Hist(v_0)$ 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 $\rho' \in Plays_{G'}(v_0)$, $\mu'(\rho') = o_{hl}$ if $\rho' = hl^\omega$ with $l \in L$ and $\mu'(\rho') = \mu(\rho')$ otherwise.*

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

In this lemma, the game (G', v_0) can be seen as a simplification of the initial game (G, v_0): each subgame $(G_{|h}, l)$ is summarized as the self-loop (l,l) and the outcome o_{hl} of its SPE. From a weak SPE $\bar{\sigma}'$ of (G', v_0), it is then easy to derive a weak SPE of (G, v_0): begin by mimicking $\bar{\sigma}'$, and after history hl with $l \in L$ continue by mimicking the weak SPE of $(G_{|h}, l)$.

Proof (of Lemma 21). 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 $hv \in Hist_i(v_0)$:

- 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 $(G_{|h_1}, l)$. Clearly $\bar{\tau}$ has the same outcome as $\bar{\sigma}'$.

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 $(G_{|h_1}, l)$. □

4.2. Proof of Theorem 16

Now that we have established all useful intermediate results in Section 4.1, we can finally proceed to the proof of Theorem 16. Recall that this theorem states the existence of a weak SPE in (G, v_0) whenever the outcome function has finite range. 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). The proof idea is to apply previous Lemma 21 the conditions of which will be satisfied thanks to Lemma 18 (to get weak SPEs on some subgames $(G_{|h}, l)$) and Corollary 14 (to get a weak SPE on (G', v_0)). This proof is by induction on the size of the finite set of outcomes.

Proof (of Theorem 16). Instead of reasoning with the underlying graph of G, we work w.l.o.g. with its unraveling from the initial vertex v_0.

By hypothesis, the outcome function μ has finite range. We denote by O the finite set of its outcomes. We are going to show how to get

(*) a weak SPE in each subgame $(G_{|h}, v)$ of (G, v_0)

(and thus in (G, v_0) itself) by induction on the size of O.

The basic case of (*), i.e. when O has size 1, is solved trivially. Indeed for all subgames of (G, v_0), every strategy profile is a weak SPE since all plays have the same outcome.

Suppose that O has size at least two, and that (*) holds for smaller sizes of O. We want to apply Lemma 21 to get a weak SPE in (G, v_0) and thus also in each of its subgames. To this end let us show how to build a set L as required by Lemma 21. Let $o \in O$ and set $L' = \emptyset$. Consider the subgame $(G_{|h}, v)$ with $hv \in Hist(v_0)$. Then either the set $\mu_{|h}^{-1}(o)$ is dense in $(G_{|h}, v)$, or it is not.

- In the first case, there exists a weak SPE in $(G_{|h}, v)$ by Lemma 18. We add v to L'.
- In the second case, as $\mu_{|h}^{-1}(o)$ is not dense, there exists a history $h'v'$ in $Hist(v)$ such that $\mu_{|h}(h'\rho) \neq o$ for all $\rho \in Plays(v')$. Therefore, in the subgame $(G_{|hh'}, v')$, as the range of the outcome function $\mu_{|hh'}$ is smaller, there exists a weak SPE in $(G_{|hh'}, v')$ by induction hypothesis. We add v' to L'.

We repeat this process for all $hv \in Hist(v_0)$. We then define L as a prefix-free subset of L' as follows: we only keep in L those vertices $v \in L'$ such the associated history hv contains no vertex of L' except v. This set $L \subseteq L'$ satisfies the hypotheses of Lemma 21:

- We have a weak SPE in each subgame $(G_{|h}, v)$, $v \in L$, by construction of L' and L;
- The game (G', v_0) as defined in Lemma 21 has also a weak SPE by Corollary 14.

Therefore, it follows by Lemma 21 that there exists a weak SPE in (G, v_0), and thus also in each of its subgames, showing that (*) holds. □

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 7). This family is constituted with all games with a finite underlying graph and a prefix-independent outcome function.

Theorem 22. *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 2, 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 6 (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 22, there always exists a weak SPE that is *finite-memory* (Corollary 23), and we will study in which cases it can be *positional* or even *uniform* (Theorem 25). *(iii)* The families of games of Theorems 16 and 22 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 22

The proof of Theorem 22 follows the same structure as for Theorem 16, by applying again Lemma 21. Here the set L required in this lemma is equal to the union of the bottom strongly connected components of the graph of G. The weak SPEs required by Lemma 21 exist on the subgames $(G_{|h}, l)$ with $l \in L$ by Corollary 19, and on the game (G', v_0) by Theorem 7.

Proof (of Theorem 22). Let $\mathcal{C}$ be the set of bottom strongly connected components of the finite graph of G. By Corollary 19, 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 $hv \in Hist(v_0)$ 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). From now on we suppose that $v_0 \notin C$ for all $C \in \mathcal{C}$.

We want to prove Theorem 22 thanks to Lemma 21. The set L required by this lemma is equal to $L = \cup_{C \in \mathcal{C}} C$. Let us verify that L satisfies the hypotheses of Lemma 21. We already know by the previous arguments that $\bar{\sigma}_C$ is a weak SPE in all subgames $(G_{|h}, v)$ with $hv \in Hist(v_0)$ and $v \in C \in \mathcal{C}$. We now consider the graph (G', v_0) constructed from (G, v_0) as described in Lemma 21 and we prove that it has a weak SPE thanks to Theorem 7. The graph (G', v_0) satisfies all the hypotheses of Theorem 7 with the set of leaves equal to L because:

- 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 the set of outcomes $O_L = \{\mu(l^\omega) \mid l \in L\} = \{o_C \mid C \in \mathcal{C}\}$ is finite since V is finite.

Therefore all the hypotheses of Lemma 21 are satisfied showing that (G, v_0) has a weak SPE. □

5.2. Finite-memory weak SPE

We here make the statement of Theorem 22 more precise by guaranteeing the existence of a weak SPE with finite memory.

Corollary 23. *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. Let us come back to the proof of Theorem 22. Two cases happen: either the initial vertex v_0 belongs to a bottom strongly connected component C or it does not.

In the first case, the required weak SPE is the strategy profile $\bar{\sigma}_C$ that is clearly finite-memory as it is uniform (by Corollary 19).

In the second case, we have constructed a weak SPE $\bar{\tau}$ thanks to Lemma 21. 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 21. Consider player $i \in \Pi$ and history $hv \in Hist_i(v_0)$. 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

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

Notice that in (6) $\tau_i(hv)$ only depends on C, and not on $l \in C$, since $\bar{\sigma}_C$ is uniform (by Corollary 19). Now let us recall the construction of $\bar{\sigma}'$ with a memory size $|L|$ given in the proof of Theorem 7, 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 8 with n bottom strongly connected components. Consider the unravelling of G_n from the initial vertex v_1 as depicted in Figure 4 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 $(v_i v_{i+1} \ldots v_{i-1})^\omega$ 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 4). 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 $v_2 v_3 \ldots v_n l_n^\omega$ with outcome o_n (as in Figure 4). We can repeat the same reasoning for the play induced by $\bar{\sigma}_{|v_1 v_2 \cdots v_n}$ in the subgame $(G_{n|v_1 v_2 \cdots v_n}, v_1)$ which must be equal to $v_1 v_2 \ldots v_{n-1} l_{n-1}^\omega$ with outcome o_{n-1}, and so on.

Hence all weak SPEs of (G_n, v_1) have the form of the one described in Figure 4 and they have finite memory of size $n-1$ as explained previously in

Example 8 (see also Figure 5). Let us show that such a weak SPE $\bar{\sigma}$ cannot have a memory size $< n-1$. Assume the contrary: w.l.o.g. consider the previous weak SPE $\bar{\sigma}$ (as in Figure 4) and in particular a Moore machine $\mathcal{M} = (M, m_0, \alpha_U, \alpha_N)$ encoding σ_1 such that $|M| < n-1$. Let $h_j v_1$, $j \in \{0, \ldots, n-1\}$ be consecutive histories, with $h_j = (v_1 v_2 \cdots v_n)^j$. On one hand, we have $\sigma_1(h_j v_1) = \alpha_N(\widehat{\alpha}_U(m_0, h_j), v_1)$ 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 $j \in \{1, \ldots, n-2\}$. Therefore there exists $j_1, j_2 \in \{1, \ldots, n-2\}, j_1 \neq j_2$, such that the associated memory state is identical, i.e, $\widehat{\alpha}_U(m_0, h_{j_1}) = \widehat{\alpha}_U(m_0, h_{j_2})$. Thus $\mathcal{M}$ enters into a cycle while reading the prefixes of $(v_1 v_2 \cdots v_n)^\omega$. This means that $\mathcal{M}$ defines $\sigma_1(hv) = v_2$ for all histories h of which h_1 is prefix, in contradiction with $\sigma_1(h_{n-1} v_1) = l_1$. □

5.3. Positional weak SPE

In the previous section, Corollary 23 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 25).

Lemma 24 (Lemma 4 of [15]). *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 24, we say that it is *layered*. In [15], 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 25. *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 26. Remember the class G_n of games, $n \geq 3$, of Example 8, such that $O = \{o_1, \ldots, o_n, \bot\}$ and each player i has a preference relation $\prec_i$ satisfying

$\bot \prec_i o_{i-1} \prec_i o_i \prec_i o_j$ for all $j \in \Pi \setminus \{i-1, i\}$. This set of outcomes is not layered because the first assertion of Lemma 24 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 23 we have shown that all weak SPEs of G_n require a memory size in $O(n)$. Hence the hypothesis of Theorem 25 about the preference relations is not completely dispensable.

Let us proceed to the proof of Theorem 25. 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 $\rho_C \in \mathit{Plays}(v)$ for some $v \in C$ induced by a simple cycle. The set $O_{\mathcal{C}} = \{o_C \mid o_C = \mu(\rho_C), C \in \mathcal{C}\}$ is finite. It is layered by hypothesis with a finite partition into layers $\{O_k\}_{k\in K}$. The proof of Theorem 25 is by induction on the number of layers and uses the next lemma dealing with one layer.

Lemma 27. *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 20. When there 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 20 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 27). The proof is by induction on $|O_{\mathcal{C}}|$. We solve the basic case $|O_{\mathcal{C}}| = 1$ by Corollary 20. Suppose that $|O_{\mathcal{C}}| = n > 1$. By Lemma 24, 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 $o_1 \prec_1 o_2 \prec_1 \ldots \prec_1 o_n$ and $o_n \prec_2 o_{n-1} \prec_2 \ldots \prec_2 o_1$. Notice that $\mathcal{P}_2$ could not exist.

For the sequel, we need the classical concept of *attractor* of $U \subseteq V$ for $\mathcal{P}_1$ [31]: it is the set $\mathit{Attr}_1(U)$ composed of all $v \in V$ from which $\mathcal{P}_1$ can force, against $\mathcal{P}_2$, to reach U. More precisely, $\mathit{Attr}_1(U)$ is constructed by induction as follows: $\mathit{Attr}_1(U) = \cup_{k\geq 0} X_k$ such that

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

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 $T = W \bigcup \cup \{D \in \mathcal{D} \mid D \notin \mathcal{C}\}$ at step 3. After step 1, ($\star$) 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 $C \in \mathcal{C}'$ is a bottom strongly connected component of G', and each $C \in \mathcal{C}''$ is a bottom strongly connected component of 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 20 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 $\bar{\sigma}'$ is a weak SPE in each (G', v'), $v' \in V'$, and $\bar{\sigma}''$ is a weak SPE 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 $o_n \prec_2 o_m$. □

We can now proceed to the proof of Theorem 25, which is by induction on the number of layers of O. The case of one layer is treated in Lemma 27. 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 27 and for G'' by induction hypothesis, we construct the required uniform weak SPE for G.

Proof (of Theorem 25). 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 27.

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 27 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 27, 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). □

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 7 or Theorem 16. In [9, 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 28. We consider the initialized game (G, v_0) of Figure 6. 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 $v_0 v_1 v_2 \ldots$ is denoted by 0^ω. The preferences of player i are given by $o \prec_i o'$ iff $o_i < o'_i$ (where o_i is the outcome i-th component). 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 5 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

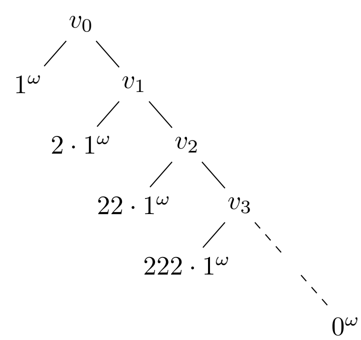

Figure 6: A game with no weak SPE

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.

7. Conclusion

In this paper, we study weak SPEs in multiplayer games (G, v_0). We provide a general result (Theorem 7) that identifies three conditions for the existence of a (finite-memory) weak SPE in (G, v_0). The proof of this result uses a labeling of the vertices with some of the outcomes. An iterative process removes those outcomes that can be impoved by a one-shot deviation, and propagates each such removal. The limit is a fixpoint that is shown non-empty and can be used to define a finite-memory weak SPE.

From Theorem 7, we derive two large classes of games that always admit a weak SPE:

- games where the outcome function has finite range (Theorem 16),
- games played on a finite graph where the outcome function is prefix-independent (Theorem 22).

For the latter family, the constructed SPE has finite memory, with a memory bounded by the number of strongly connected components of the graph. When the set of outcomes is layered, we can go further by constructing a uniform weak SPE in that case.

References

[1] A. Pnueli, R. Rosner, On the synthesis of a reactive module, in: POPL, ACM Press, 1989, pp. 179–190.

[2] J. F. Nash, Equilibrium points in n-person games, in: PNAS, Vol. 36, National Academy of Sciences, 1950, pp. 48–49.

[3] A. Rubinstein, Comments on the interpretation of game theory, Econometrica 59 (1991) 909–924.

[4] R. Selten, Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, Zeitschrift für die gesamte Staatswissenschaft 121 (1965) 301–324 and 667–689.

[5] V. Bruyère, Computer aided synthesis: A game-theoretic approach, in: DLT 2017, Vol. 10396 of Lecture Notes in Computer Science, 2017, pp. 3–35.

[6] T. Brihaye, V. Bruyère, N. Meunier, J. Raskin, Weak subgame perfect equilibria and their application to quantitative reachability, in: CSL, Vol. 41 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015, pp. 504–518.

[7] H. Kuhn, Extensive games and the problem of information, Classics in Game Theory (1953) 46–68.

[8] E. Solan, N. Vieille, Deterministic multi-player Dynkin games, Journal of Mathematical Economics 39 (2003) 911–929.

[9] J. Flesch, J. Kuipers, A. Mashiah-Yaakovi, G. Schoenmakers, E. Solan, K. Vrieze, Perfect-information games with lower-semicontinuous payoffs, Math. Oper. Res. 35 (2010) 742–755.

[10] X. Shen, H. Yu, J. Buford, M. Akon, Handbook of Peer-to-Peer Networking, Springer, 2010.

[11] S. Le Roux, A. Pauly, Infinite sequential games with real-valued payoffs, in: CSL-LICS, ACM, 2014, pp. 62:1–62:10.

[12] D. Fudenberg, D. Levine, Subgame-perfect equilibria of finite- and infinite-horizon games, Journal of Economic Theory 31 (1983) 251–268.

[13] R. A. Purves, W. D. Sudderth, Perfect information games with upper semicontinuous payoffs, Math. Oper. Res. 36 (3) (2011) 468–473.

[14] J. Flesch, A. Predtetchinski, A characterization of subgame perfect equilibrium plays in Borel games of perfect information, Math. Oper. Res. 42 (4) (2017) 1162–1179.

[15] S. Le Roux, Infinite subgame perfect equilibrium in the Hausdorff difference hierarchy, in: TTCS, Vol. 9541 of Lecture Notes in Computer Science, Springer, 2015, pp. 147–163.

[16] C. Alós-Ferrer, K. Ritzberger, Does backwards induction imply subgame perfection?, Games and Economic Behavior 103 (2017) 19–29.

[17] R. Brenguier, L. Clemente, P. Hunter, G. A. Pérez, M. Randour, J. Raskin, O. Sankur, M. Sassolas, Non-zero sum games for reactive synthesis, in: LATA, Vol. 9618 of Lecture Notes in Computer Science, Springer, 2016, pp. 3–23.

[18] M. Ummels, Rational behaviour and strategy construction in infinite multiplayer games, in: FSTTCS, Vol. 4337 of Lecture Notes in Computer Science, Springer, 2006, pp. 212–223.

[19] K. Chatterjee, T. A. Henzinger, M. Jurdzinski, Games with secure equilibria, Theor. Comput. Sci. 365 (2006) 67–82.

[20] K. Chatterjee, T. A. Henzinger, Assume-guarantee synthesis, in: TACAS, Vol. 4424 of Lecture Notes in Computer Science, Springer, 2007, pp. 261–275.

[21] J. De Pril, J. Flesch, J. Kuipers, G. Schoenmakers, K. Vrieze, Existence of secure equilibrium in multi-player games with perfect information, in: MFCS, Vol. 8635 of Lecture Notes in Computer Science, Springer, 2014, pp. 213–225.

[22] V. Bruyère, N. Meunier, J. Raskin, Secure equilibria in weighted games, in: CSL-LICS, ACM, 2014, pp. 26:1–26:26.

[23] K. Chatterjee, L. Doyen, E. Filiot, J. Raskin, Doomsday equilibria for omega-regular games, in: VMCAI, Vol. 8318 of Lecture Notes in Computer Science, Springer, 2014, pp. 78–97.

[24] D. Berwanger, Admissibility in infinite games, in: STACS, Vol. 4393 of Lecture Notes in Computer Science, Springer, 2007, pp. 188–199.

[25] R. Brenguier, J. Raskin, M. Sassolas, The complexity of admissibility in omega-regular games, in: CSL-LICS, ACM, 2014, pp. 23:1–23:10.

[26] R. Brenguier, J. Raskin, O. Sankur, Assume-admissible synthesis, in: CONCUR, Vol. 42 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015, pp. 100–113.

[27] O. Kupferman, G. Perelli, M. Y. Vardi, Synthesis with rational environments, Ann. Math. Artif. Intell. 78 (1) (2016) 3–20.

[28] V. Bruyère, S. Le Roux, A. Pauly, J.-F. Raskin, On the existence of weak subgame perfect equilibria, in: FOSSACS, Vol. 10203 of Lecture Notes in Computer Science, Springer, 2017, pp. 145–161.

[29] E. Grädel, M. Ummels, Solution Concepts and Algorithms for Infinite Multiplayer Games, in: New Perspectives on Games and Interaction, Vol. 4, Amsterdam University Press, 2008, pp. 151–178.

[30] K. Chatterjee, L. Doyen, T. A. Henzinger, Quantitative languages, ACM Trans. Comput. Log. 11.

[31] E. Grädel, W. Thomas, T. Wilke (Eds.), Automata, Logics, and Infinite Games: A Guide to Current Research, Vol. 2500 of Lecture Notes in Computer Science, Springer, 2002.