On the Existence of Weak Subgame Perfect Equilibria

Véronique Bruyère[1(✉)], Stéphane Le Roux[2], Arno Pauly[2], and Jean-François Raskin[2]

[1] Département d'informatique, Université de Mons (UMONS), Mons, Belgium
veronique.bruyere@umons.ac.be
[2] Département d'informatique, Université Libre de Bruxelles (ULB), Brussels, Belgium

Abstract. We study multi-player turn-based games played on a directed graph, where the number of players and vertices can be infinite. 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), 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. We give general conditions on the structure of the game graph and the preference relations of the players that guarantee the existence of a weak SPE, which moreover is finite-memory.

1 Introduction

Games played on graphs have a large number of applications in theoretical computer science. One particularly important application is *reactive synthesis* [21], i.e. the design of a controller that guarantees a good behavior of a reactive system evolving in a possibly hostile environment. One classical model proposed for the synthesis problem is the notion of *two-player zero-sum game played on a graph*. One player is the reactive system and the other one is the environment; the vertices of the graph model their possible states and the edges model their possible actions. Interactions between the players generate an infinite play in the graph which model behaviors of the system within its environment. As one cannot assume cooperation of the environment, the objectives of the two players are considered to be opposite. Constructing a controller for the system then means devising a *winning strategy* for the player modeling it. Reality is often more subtle and the environment is usually not fully adversarial as it has its own objective, meaning that the game should be non zero-sum. Moreover instead of two players, we could consider the more general situation of several players modeling different interacting systems/environments each of them with its own objective.

S. Le Roux, A. Pauly, J.-F. Raskin—Supported by ERC Starting Grant (279499: inVEST).

© Springer-Verlag GmbH Germany 2017
J. Esparza and A.S. Murawski (Eds.): FOSSACS 2017, LNCS 10203, pp. 145–161, 2017.
DOI: 10.1007/978-3-662-54458-7_9

The concept of *Nash equilibrium* (NE) [20] 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 [23]. Thus the notion of NE has been strengthened into the notion of *subgame perfect equilibrium* (SPE) [24]: 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.

Variants of SPE, *weak SPE* and *very weak SPE*, have been recently introduced in [5]. 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 the proof of Kuhn's theorem [16] 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 [5,26]. 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 [5,12]. 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 [25] for details).

In this paper, we provide the following contributions. First, we identify *general conditions* to guarantee the existence of a weak SPE (Theorem 1). 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 2);
– for games with a *finite* underlying graph and a *prefix-independent* outcome function (Theorem 4).

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

Related work. The concept of SPE has been first introduced and studied by the game theory community. Kuhn proves in [16] the existence of SPEs in games played on finite trees. This result has been generalized in several ways. Games with a continuous real-valued outcome function and a finitely branching tree always have an SPE [19] (the case with finitely many players is first established in [14]). In [12] (resp. [22]), 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 [22] is extended to an infinite number of players in [13]. In [19], 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 [18], 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 part of the former work, the equivalence between SPEs and very weak SPEs is implicitly used as a proof technique: in a finite setting in [16], continuous setting in [14], and lower-semicontinuous setting in [12]. In the latter reference, the authors implicitly prove that all games with a finite range real-valued outcome function have a weak SPE (which is an SPE when the outcome function is additionally lower-semicontinuous). Inspired by this result and its proof, we here generalize it to an infinite number of players using a simpler proof technique: our algorithm discards outcomes instead of discarding plays.

The concept of SPE and other solution concepts for multi-player non zero-sum games have been recently studied by the theoretical computer community, see [2] for a survey. In [27], the existence of SPEs (and thus weak SPEs) is proved for games with a finite number of players and Borel Boolean objectives. We here generalize the existence of weak SPEs to games with infinitely many players. In [5], weak SPEs are introduced as a technical tool for showing the existence of SPEs in quantitative reachability games played on finite weighted graphs. An algorithm is also provided for the construction of a (finite-memory) weak SPE that appears to be an SPE for this particular class of games. We here give several existence results that are orthogonal to the results of [5] 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 [7] and then used for reactive synthesis in [10]. These equilibria are generalized to multiple players in [11] or to quantitative objectives in [6], see also a variant called Doomsday equilibrium in [8]. Like NEs, they are subject to possible non-credible threats. Other refinements of NE are provided by the notion of admissible strategy introduced in [1], with computational aspects studied in [4], and potential for synthesis studied in [3]. Note that these notions are immune, as (weak) SPEs, of non-credible threats. Finally, in [17], 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 [6] and [11], the proposed solution concepts are not guaranteed to exist, hence results concern mostly algorithmic techniques to decide their existence and not general conditions for existence as in this paper.

2 Preliminaries

In this section, 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 play outcomes which allows him to compare plays.

Games. A *game* is a tuple $G = (\Pi, V, (V_i)_{i\in\Pi}, E, O, \mu, (\prec_i)_{i\in\Pi})$ where *(i)* Π is a set of players, *(ii)* 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, *(iii)* $(V_i)_{i\in\Pi}$ is a partition of V such that V_i is the set of vertices controlled by player $i \in \Pi$, *(iv)* O is a set of outcomes and $\mu : V^\omega \to O$ is an outcome function, and *(v)* $\prec_i \subseteq O\times O$ is a preference relation for player $i \in \Pi$. In this definition the underlying graph (V, E) can be infinite (that is, of arbitrarily cardinality), as well as the set Π of players and the set O of outcomes.

A *play* of G is an infinite (countable) sequence $\rho = \rho_0\rho_1 \ldots \in V^\omega$ 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 $Plays$ and the set of histories (ending with a vertex in V_i) by $Hist$ (resp. by $Hist_i$).[1] A *prefix* (resp. *suffix*) of a play $\rho = \rho_0\rho_1 \ldots$ is a finite sequence $\rho_{\leq n} = \rho_0 \ldots \rho_n$ (resp. 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. When an initial vertex $v_0 \in V$ is fixed, we call (G, v_0) an *initialized* game. In this case, plays and histories are supposed to start in v_0, and we use notations $Plays(v_0)$ and $Hist(v_0)$. In this article, we often *unravel* the graph of the game (G, v_0) from the initial vertex v_0, which yields an infinite tree rooted at v_0.

The outcome function μ assigns an outcome $\mu(\rho) \in O$ to each play $\rho \in V^\omega$. It is *prefix-independent* if $\mu(h\rho) = \mu(\rho)$ for all histories h and play ρ. A *preference* relation $\prec_i \subseteq O \times O$ is an irreflexive and transitive binary relation. It allows for player i to compare two plays $\rho, \rho' \in V^\omega$ with respect to their outcome: $\mu(\rho) \prec_i \mu(\rho')$ means that player i prefers ρ' to ρ. In this paper we restrict to *linear* preferences. (It is w.l.o.g. since the preference properties that we use are preserved by linear extension). We write $o \preceq_i o'$ when $o \prec_i o'$ or $o = o'$; notice

[1] Indexing $Plays_G$ or $Hist_G$ with G allows to recall the related game G.

that $o \not\prec_i o'$ if and only if $o' \preceq_i o$. We sometimes use notation $\prec_v$ instead of $\prec_i$ when vertex $v \in V_i$ is controlled by player i.

Example 1. Let us mention some classical classes of games where the set of outcomes O is a subset of $(\mathbb{R} \cup \{+\infty, -\infty\})^{\Pi}$, and for all player $i \in \Pi$, $\prec_i$ is the usual ordering $<$ on $\mathbb{R} \cup \{+\infty, -\infty\}$ on the outcome i-th components. In other words, each player i has a real-valued payoff function $\mu_i : Plays \to \mathbb{R} \cup \{+\infty, -\infty\}$. 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, also [15]. 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 [9][2]: *(i) limsup*: $\mu_i(\rho) = \limsup_{k\to\infty} w_i(\rho_k, \rho_{k+1})$, *(ii) mean-payoff*: $\mu_i(\rho) = \limsup_{n\to\infty} \sum_{k=0}^{n} \frac{w_i(\rho_k, \rho_{k+1})}{n}$.

Strategies. Let (G, v_0) be an initialized game. A *strategy* σ for player i in (G, v_0) is a function $\sigma : Hist_i(v_0) \to V$ assigning to each history $hv \in Hist_i(v_0)$ a vertex $v' = \sigma(hv)$ such that $(v, v') \in E$. A strategy σ of player i is *positional* if it only depends on the last vertex of the history, i.e. $\sigma(hv) = \sigma(v)$ for all $hv \in Hist_i(v_0)$. It is a *finite-memory* strategy if $\sigma(hv)$ only needs finite memory of the history hv (recorded by a Moore machine[3] with a finite number of memory states). These definitions of (positional, finite-memory) strategy are given for an initialized game (G, v_0). We call *uniform* every positional strategy σ of player i defined for all $hv \in Hist_i$ (instead of $Hist_i(v_0)$), that is, when σ is a positional strategy in all initialized games (G, v), $v \in V$.

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

Let $\bar{\sigma}$ be a strategy profile. When all players stick to their own strategy except player i that shifts from σ_i to σ'_i, we denote by $(\sigma'_i, \bar{\sigma}_{-i})$ the derived strategy profile, and by $\langle\sigma'_i, \bar{\sigma}_{-i}\rangle_{v_0}$ the induced play in (G, v_0). We say that σ'_i is a *deviating* strategy from σ_i. When σ_i and σ'_i only differ on a finite number of histories (resp. on v_0), we say that σ'_i is a *finitely-deviating* (resp. *one-shot deviating*) strategy from σ_i. One-shot deviating strategies is a well-known notion

[2] The limit inferior can be used instead of the limit superior.
[3] Moore machines are usually defined for finite sets V. We here allow infinite sets V.

that for instance appears in the proof of Kuhn's theorem [16] with the one-step deviation property. Finitely-deviating strategies have been introduced in [5].

Variants of subgame perfect equilibria. 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 1. *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 (SPE) is a refinement of NE. To define it, we introduce the following concepts. Given a game $G = (\Pi, V, (V_i)_{i\in\Pi}, E, \mu, (\prec_i)_{i\in\Pi})$ and a history $h \in Hist$, we denote by $G_{|h}$ the game $(\Pi, V, (V_i)_{i\in\Pi}, E, \mu_{|h}, (\prec_i)_{i\in\Pi})$ where $\mu_{|h}(\rho) = \mu(h\rho)$ for all plays of $G_{|h}$[4], and we say that $G_{|h}$ is a *subgame* of G. Given an initialized game (G, v_0) and a history $hv \in Hist(v_0)$, the initialized game $(G_{|h}, v)$ is called the subgame of (G, v_0) with history hv. In particular (G, v_0) is a subgame of itself with history hv_0 such that $h = \epsilon$. Given a strategy σ of player i in (G, v_0), the strategy $\sigma_{|h}$ in $(G_{|h}, v)$ is defined as $\sigma_{|h}(h') = \sigma(hh')$ for all $h' \in Hist_i(v)$. 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 induced play in the subgame $(G_{|h}, v)$.

Now a strategy profile is an SPE in an initialized game if it induces an NE in each of its subgames. Two variants of SPE, called weak SPE and very weak SPE, are proposed in [5] such that no player has an incentive to deviate in any subgame using finitely deviating strategies and one-shot deviating strategies respectively (instead of any deviating strategy).

Definition 2. *Given an initialized game* (G, v_0)*, a strategy profile* $\bar{\sigma}$ *of* (G, v_0) *is a* (weak, very weak resp.) subgame perfect equilibrium *if for all histories* $hv \in Hist(v_0)$*, for all players* $i \in \Pi$*, for all (finitely, one-shot resp.) deviating strategies* σ'_i *from* $\sigma_{i|h}$ *of player* i *in the subgame* $(G_{|h}, v)$*, we have* $\mu(\langle\bar{\sigma}_{|h}\rangle_v) \not\prec_i \mu(\langle\sigma'_i, \bar{\sigma}_{-i|h}\rangle_v)$.

Every SPE is a weak SPE, and every weak SPE is a very weak SPE. The next proposition states that weak SPE and very weak SPE are equivalent notions, but this is not true for SPE and weak SPE (see also Example 2 below).

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

---

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

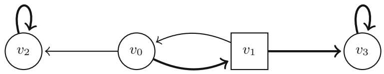

Fig. 1. An initialized game (G, v_0) with a (very) weak SPE and no SPE.

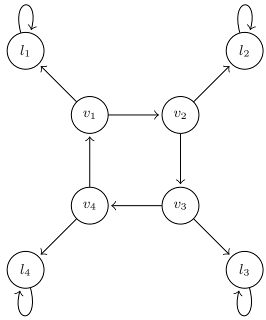

Fig. 2. Game G_4

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

It is known that this game has no SPE [26]. Nevertheless the strategy profile $\bar{\sigma}$ depicted with thick edges is a very weak SPE, and thus a weak SPE by Proposition 1. Let us give some explanation. Due to the simple form of the game, only two cases are to be treated. Consider first the subgame $(G_{|h}, v_0)$ with $h \in (v_0v_1)^*$, and the one-shot deviating strategy σ'_1 from $\sigma_{1|h}$ such that $\sigma'_1(v_0) = v_2$. Then $\langle \bar{\sigma}_{|h} \rangle_{v_0} = v_0v_1v_3^\omega$ and $\langle \sigma'_1, \sigma_{2|h} \rangle_{v_0} = v_0v_2^\omega$ with respective outcomes o_3 and o_2, showing that σ'_1 is not a profitable deviation for player 1 in $(G_{|h}, v_0)$. Now in the subgame $(G_{|h}, v_1)$ with $h \in (v_0v_1)^*v_0$, the one-shot deviating strategy from $\sigma_{2|h}$ such that $\sigma'_2(v_1) = v_0$ is not profitable for player 2 in $(G_{|h}, v_1)$ because $\langle \bar{\sigma}_{|h} \rangle_{v_1} = v_1v_3^\omega$ and $\langle \sigma_{1|h}, \sigma'_2 \rangle_{v_1} = v_1v_0v_1v_3^\omega$ with the same outcome o_3.

Notice that $\bar{\sigma}$ is not an SPE. Indeed the strategy σ'_2 such that $\sigma'_2(hv_1) = v_0$ for all h, is infinitely deviating from σ_2, and is a profitable deviation for player 2 in (G, v_0) since $\langle \sigma_1, \sigma'_2 \rangle_{v_0} = (v_0v_1)^\omega$ with outcome o_1.

3 General Conditions for the Existence of Weak SPEs

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

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

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

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

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

Example 3. For each natural number $n \geq 3$, we build a game G_n with n players, $2n$ vertices, $3n$ edges, and $n+1$ outcomes. The set of players is $\Pi = \{1, 2, \ldots, n\}$ 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 Fig. 2. 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 $j \in \Pi \backslash \{i-1, i\}$ (with the convention that $o_0 = o_n$).

Each game (G_n, v_1) satisfies the hypotheses of Theorem 1 with $L = \{l_1, \ldots, l_n\}$ and thus has a finite-memory weak SPE. Such a strategy profile $\bar{\sigma}$ is depicted in Fig. 3 for $n = 4$ (see the thick edges on the unravelling of G_4 from the initial vertex v_1) and can be easily generalized to every $n \geq 3$. One verifies that this profile is a very weak SPE, and thus a weak SPE by Proposition 1. For all $i \in \Pi$, the strategy σ_i of player i is finite-memory with a memory size equal to $n - 1$. Intuitively, along $(v_1 \ldots v_n)^\omega$, 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 $n-1$.

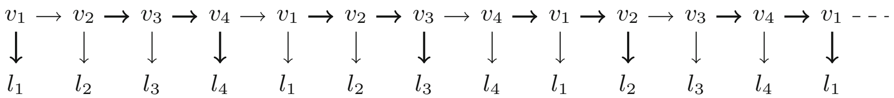

Fig. 3. Weak SPE in (G_4, v_1)

Let us now proceed to the proof of Theorem 1. Recall that it is enough to prove the existence of a very weak SPE by Proposition 1. The proof idea is the following one. Initially, for each vertex v, we accept all plays $\rho = hl^\omega$ with $h \in Hist(v)$ and $l \in L$ as *potential* plays induced by a very weak SPE in the initialized game (G, v). We thus label each v by the set of outcomes $\mu(l^\omega)$ for such leaves l (recall that $\mu(\rho) = \mu(l^\omega)$ by the second condition of Theorem 1). Notice that this labeling is finite (resp. not empty) by the third (resp. first) condition of the theorem. Step after step, we are going to remove some outcomes from the vertex labelings by a *Remove* operation followed by an *Adjust*operation.

The *Remove* operation removes an outcome o from the labeling of a given vertex v when there exists an edge (v, v') for which $o \prec_v o'$ for all outcomes o' that label v'. Indeed o cannot be the outcome of a play induced by a very weak SPE since the player who controls v will choose the move (v, v') to get a preferable outcome o'. Now it may happen that for another vertex u having o in its labeling, all potential plays induced by a very weak SPE from u with outcome o necessarily cross vertex v. As o has been removed from the labeling of v, these potential plays do no longer survive and o will also be removed from the labeling of u by the *Adjust* operation. Repeatedly applying these two operations converge to a fixpoint for which we will prove non-emptiness (this is the difficult part of the proof, non-emptiness will be obtained by maintaining three invariants, see Lemma 1). From this fixpoint, for each vertex v and each outcome o of the resulting labeling of v, there exists a play $\rho_{v,o} = hl^\omega$ with outcome o for some $h \in Hist(v)$ and $l \in L$. We can thus build a very weak SPE $\bar{\sigma}$ in (G, v_0) as follows. The construction of $\bar{\sigma}$ is done step by step: *(i)* initially $\bar{\sigma}$ is partially defined such that $\langle \bar{\sigma} \rangle_{v_0} = \rho_{v_0,o_0}$ for some o_0; *(ii)* then in the subgame $(G_{|h}, v)$ such that $\langle \bar{\sigma}_{|h} \rangle_v = \rho_{v,o}$, if the player who controls v chooses the move (v, v') in a one-shot deviation, then there exists $\rho_{v',o'}$ such that $o \not\prec_v o'$ by definition of the fixpoint, and we thus extend the construction of $\bar{\sigma}$ such that $\langle \bar{\sigma}_{|hv} \rangle_{v'} = \rho_{v',o'}$.

Let us now go into the details of the proof. For each $l \in L$, we denote by o_l the outcome $\mu(l^\omega)$. Recall that for all $\rho = hl^\omega$ we have $\mu(\rho) = o_l$ by the second hypothesis of the theorem. For each $v \in V$, we denote by $Succ(v)$ 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 $v \in V \backslash L$.

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, $\lambda_0(v) \neq \emptyset$. 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

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

with the convention that $m_\alpha(v) = \top$ if $Succ(v) = \emptyset$ or if $\lambda_\alpha(v') = \emptyset$ for all $v' \in Succ(v)$.[5] When $m_\alpha(v) \neq \top$, we says that $v' \in Succ(v)$ *realizes* $m_\alpha(v)$ if $m_\alpha(v) = \min_{\prec_v} \lambda_\alpha(v')$. Notice that even if $Succ(v)$ could be infinite, there are finitely many sets $\lambda_\alpha(v')$ since O_L is finite. This justifies our use of $\max_{\prec_v}$ and $\min_{\prec_v}$ operators in the definition of $m_\alpha(v)$.

[5] We suppose that $o \prec_v \top$ for all $o \in O_L$.

We alternate between *Remove* and *Adjust* that remove outcomes from labeling $\lambda_\alpha(v)$ in the following way:

- For an even[6] successor ordinal $\alpha + 2$,
  ***Remove* operation.** Test if for some $v \in V$, there exist $o \in \lambda_\alpha(v)$ and $v' \in Succ(v)$ such that
  $$o \prec_v o', \text{for all} o' \in \lambda_\alpha(v').$$
  If such a v exists, then $\lambda_{\alpha+1}(v) = \lambda_\alpha(v) \backslash \{o\}$, and $\lambda_{\alpha+1}(u) = \lambda_\alpha(u)$ for the other vertices $u \neq v$. Otherwise $\lambda_{\alpha+1}(u) = \lambda_\alpha(u)$ for all $u \in V$.
  ***Adjust* operation.** Suppose that $\lambda_{\alpha+1}(v) = \lambda_\alpha(v) \backslash \{o\}$ 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 $\lambda_{\alpha+2}(u) = \lambda_{\alpha+1}(u)$, otherwise $\lambda_{\alpha+2}(u) = \lambda_{\alpha+1}(u) \backslash \{o\}$. For all $u \in V$ such that $o \notin \lambda_{\alpha+1}(u)$, let $\lambda_{\alpha+2}(u) = \lambda_{\alpha+1}(u)$.

  Suppose that $\lambda_{\alpha+1}(v) = \lambda_\alpha(v)$ for all $v \in V$ at the previous step, then $\lambda_{\alpha+2}(v) = \lambda_{\alpha+1}(v)$ for all $v \in V$.
  (Thus *Remove* is performed at odd step $\alpha + 1$, whereas *Adjust* is performed at even step $\alpha + 2$.)
- For a limit ordinal α, let $\lambda_\alpha(v) = \cap_{\beta<\alpha} \lambda_\beta(v)$ for all $v \in V$.

For each v, the sequence $(\lambda_\alpha(v))_\alpha$ is nonincreasing (for the set inclusion), and thus the sequence $(m_\alpha(v))_\alpha$ is nondecreasing (for the $\prec_v$ relation). Notice that for all leaves $l \in L$ and all steps α, we have $\lambda_\alpha(l) = \{o_l\}$. The next lemma states that we get a non empty fixpoint in the following sense:

Lemma 1. *There exists an ordinal* α^* *such that* $\lambda_{\alpha^*}(v) = \lambda_{\alpha^*+1}(v) = \lambda_{\alpha^*+2}(v)$ *for all* $v \in V$. *Moreover,* $\lambda_{\alpha^*}(v) \neq \emptyset$ *for all* $v \in V$.

Proof. Each set $\lambda_\alpha(v)$ has size bounded by $|O_L|$. During the inductive process, from step α (with α even) to step $\alpha + 1$, *Remove* removes one outcome from one of these sets, and from step $\alpha + 1$ to step $\alpha + 2$, *Adjust* can remove outcomes from several such sets (it can remove no outcome at all). Therefore there exists an ordinal α^* such that $\lambda_{\alpha^*}(v) = \lambda_{\alpha^*+1}(v) = \lambda_{\alpha^*+2}(v)$ for all $v \in V$, and a fixpoint is then reached.[7] To prove that $\lambda_{\alpha^*}(v) \neq \emptyset$, we consider the next three invariants for which we will briefly explain that they are initially true and remain true after each step α. The non emptiness of $\lambda_{\alpha^*}(v)$ will follow from the second invariant.

INV1. For $v \in V$, we have for all $v' \in Succ(v)$ that
$$\{o \in \lambda_\alpha(v') \mid m_\alpha(v) \preceq_v o\} \subseteq \lambda_\alpha(v).$$

In particular, when $m_\alpha(v) \neq \top$, for each v' that realizes $m_\alpha(v)$, we have
$$\lambda_\alpha(v') \subseteq \lambda_\alpha(v). \tag{1}$$

[6] Ordinal 0 and each limit ordinal are even, and each successor ordinal $\alpha + 1$ is even (resp. odd) if α is odd (resp. even).
[7] When V is finite, it is reached after at most $2|O_L| \cdot |V|$ steps.

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, $\lambda_\alpha(u) \subseteq \lambda_\alpha(v)$.

The three invariants are initially true. 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. The successor-ordinal steps are the difficult ones. The detailed proof invokes many times that the $\lambda_\alpha(v)$ and $m_\alpha(v)$ are monotone with respect to α.

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).

Consider even step $\alpha + 2$ and the *Adjust* operation. *(i)* One checks that *Adjust* preserves INV1 by case splitting on whether $\lambda_{\alpha+2}(v) = \lambda_{\alpha+1}(v)$. *(ii)* By contradiction assume that $\lambda_{\alpha+1}(v) = \{o\}$ 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 $\alpha + 1$. If it is no longer suitable at step $\alpha + 2$, some o was removed from some proper prefix $u_1 \dots u_{i-1}$, *i.e.* $o \in \lambda_{\alpha+2}(u_i)$ but $o \notin \lambda_{\alpha+2}(u_{i-1})$, so $o \notin \lambda_{\alpha+1}(u_{i-1})$ by definition of *Adjust*. INV3 provides a suitable path (void of o) from u_{i-1} at step $\alpha + 1$. Concatenating it with $u_1 \dots u_{i-1}$ witnesses INV3 at step $\alpha + 2$. □

By the previous lemma, we have a fixpoint such that that $\lambda_{\alpha^*}(v) \neq \emptyset$ for all $v \in V$. Moreover 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 $\pi l^\omega \in Plays(v)$:

$$\rho_{v,o} = \pi l^\omega. \tag{2}$$

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

To get Theorem 1, it remains to explain how to build a weak SPE $\bar{\sigma}$ from this fixpoint that is finite-memory.

Proof (of Theorem 1). 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 (2). 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 $o \in \lambda_{\alpha^*}(v)$. By the fixpoint and in particular by *Remove* (with $o \in \lambda_{\alpha^*}(v)$ and $v' \in Succ(v)$), there exists $o' \in \lambda_{\alpha^*}(v')$ such that

$$o \not\prec_v o'. \tag{3}$$

With $\rho_{v',o'}$ as in (2), 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 $\bar{\sigma}$ 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 (3) 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 (3), σ'_i is not a profitable deviation for player i. Hence $\bar{\sigma}$ is a very weak SPE and thus a weak SPE by Proposition 1.

It remains to prove that $\bar{\sigma}$ is finite-memory by correctly choosing the plays $\rho_{v,o}$ of (2). Fix $o \in O_L$ and consider the set U_o of vertices v such that $o \in \lambda_{\alpha^*}(v)$. 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. Hence having o in memory, one can produce positionally each $\rho_{v,o}$ with $v \in U_o$. Thus the memory-size of $\bar{\sigma}$ is equal to O_L. □

The next corollary is an easy consequence of Theorem 1.

Corollary 1. *Let (G, v_0) be an initialized game with a subset $L \subseteq V$ of leaves*[8] *such that the underlying graph is a tree rooted at v_0. If (G, v_0) satisfies the first and third conditions of Theorem 1, then there is a positional weak SPE in (G, v_0).*

In the next sections, we present two large families of games for which there always exists a weak SPE, as a consequence of Theorem 1 and its Corollary 1.

4 First Application

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

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

Let us comment this theorem. *(i)* Kuhn's theorem [16] states that there exists an SPE in all initialized games played on *finite trees* (note that in this particular case, the existence of a weak SPE is equivalent to the existence of an SPE).

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

Theorem 2 can be seen as a generalization of Kuhn's theorem: if we keep the outcome set finite, all initialized games (regardless of the underlying graph and the player set) have a weak SPE. *(ii)* Theorem 2 guarantees the existence of a weak SPE for games with *Boolean* objectives as presented in Example 1, since their outcome function μ has finite range. It is proved in [27] that each initialized game with a finite number of players and Borel objectives has an SPE and thus a weak SPE. We thus here extend the existence of a weak SPE to an infinite number of players. *(iii)* The next theorem is proved in [12] for outcome functions $\mu = (\mu_i)_{i \in \Pi}$ as presented in Example 1 and has strong relationship with Theorem 2. Recall that $\mu_i : Plays \to \mathbb{R}$ is *lower-semicontinuous* if whenever a sequence of plays $(\rho_n)_{n \in \mathbb{N}}$ converges to play $\rho = \lim_{n\to\infty} \rho_n$, then $\liminf_{n\to\infty} \mu_i(\rho_n) \geq \mu_i(\rho)$.

Theorem 3 ([12]). *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 [5], we recover the previous result with our Theorem 2, however with a set of players of any cardinality and general outcome functions $\mu : Plays \to O$. Even if it is not explicitly mentioned in [12], 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 (indeed we remove outcomes from the sets $\lambda_\alpha(v)$ (see the proof of Theorem 1) whereas plays are removed in the inductive process of [12]).

Intermediate results. The proofs of Theorem 2 in this section and Theorem 4 in the next section require intermediate results that we now describe. We begin with the next lemma where the set $\mu^{-1}(o)$, with $o \in O$, is called *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 o.

Lemma 2. *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).*

Lemma 2 leads to the next corollary. This corollary will provide a first step towards Theorem 4; it is already interesting on its own right.

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

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

Lemma 3. *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 the 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, then (G, v_0) has also a weak SPE.

Proof of Theorem 2. We now proceed to the proof of Theorem 2. W.l.o.g. we can suppose that the underlying graph of G is a tree rooted at v_0 (by unraveling this graph from v_0). The proof idea is to apply Lemma 3 the conditions of which will be satisfied thanks to Lemma 2 (to get weak SPEs on some subgames) and Corollary 1 (to get a weak SPE on (G', v_0)).

Proof (of Theorem 2). Let us reason on the unraveling of G from 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 (*) is trivial since for all subgames of (G, v_0), each strategy profile is a weak SPE when μ has range one. Suppose that O has size at least two, and that (*) holds for smaller sizes. We are going to build a set L as required by Lemma 3 to get a weak SPE in (G, v_0) and thus also in each of its subgames.

Let $o \in O$ and set $L' = \emptyset$. Consider the subgame $(G_{|h}, v)$ with $hv \in Hist_G(v_0)$. 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 2. 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. As in the first case, we add v' to L'.

We repeat this process for all $hv \in Hist(v_0)$. We then get the set $L \subseteq L'$ as required by Lemma 3 by only keeping[9] the vertices $v \in L'$ such the associated history hv contains no vertex of L' except v. For each subgame $(G_{|h}, v)$ with $v \in L$, we thus have a weak SPE. The game (G', v_0) as defined in Lemma 3 has also a weak SPE by Corollary 1. It thus follows by Lemma 3 that there exists a weak SPE in (G, v_0), and thus also in each of its subgames. □

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 Sect. 3 (more particularly Theorem 1). This family is constituted with all games with a finite underlying graph and a prefix-independent outcome function.

[9] L is the prefix-free subset of L'.

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

Let us comment this theorem. *(i)* It guarantees the existence of a weak SPE for classical games with *quantitative* objectives as presented in Example 1, such that their outcome function is prefix-independent. This is the case of *limsup* and *mean-payoff* functions (and their limit inferior counterparts). Recall that Example 2 (see also Fig. 1) provides a game with no SPE, where the payoff functions μ_i can be seen as either *limsup* or *mean-payoff* (or their limit inferior counterparts). *(ii)* Later in this section, we will show that under the hypotheses of Theorem 4, there always exists a weak SPE that is *finite-memory* (Corollary 3), and we will study in which cases it can be *positional* or even *uniform* (Theorem 5). *(iii)* The families of games of Theorems 2 and 4 are incomparable: Boolean reachability games are in the first family but not in the second one, and mean-payoff games are in the second family but not in the first one.

The proof of Theorem 4 follows the same structure as for Theorem 2. The idea is to apply Lemma 3 where L is the union of the bottom strongly connected components of the graph of G. The weak SPEs required by Lemma 3 exist on the subgames $(G_{|h}, l)$, $l \in L$, by Corollary 2, and on the game (G', v_0) by Theorem 1.

Discussion on the memory. First we make the statement of Theorem 4 more precise by guaranteeing the existence of a weak SPE with finite-memory. The necessity of memory is illustrated by the family of games G_n of Example 3.

Corollary 3. *Let* (G, v_0) *be an initialized game with a finite underlying graph and a prefix-independent outcome function. Then there is a finite-memory weak SPE in* (G, v_0) *with* $O(m)$ *memory size where* m *is the number of bottom strongly connected components of the graph. A memory size linear in* m *is necessary.*

Second we identify conditions on the preference relations of the players, as expressed in the next lemma, that guarantee the existence of a *uniform* (instead of finite-memory) weak SPE (see Theorem 5).

Lemma 4 (Lemma 4 of [18]). *Let* O *be a set of outcomes. Let* $\prec_i \subseteq O \times O$ *be a preference relation 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 this 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 ones. A set of outcomes satisfying the conditions of Lemma 4 is called *layered*. In [18], the author characterizes the preference relations that always yield SPE in games with outcome functions in the Hausdorff difference hierarchy of the open sets. One condition is that the set of outcomes is layered.

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

Example 4. Remember the class G_n of games, $n \geq 3$, of Example 3, 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 4 is not satisfied. Indeed we have $o_2 \prec_3 o_3 \prec_3 o_1$ and $o_1 \prec_2 o_2 \prec_2 o_3$. Recall that all weak SPEs of the games G_n require a memory size in $O(n)$ (by Corollary 3). Hence the hypothesis of Theorem 5 about the preference relations is not completely dispensable.

References

1. Berwanger, D.: Admissibility in infinite games. In: Thomas, W., Weil, P. (eds.) STACS 2007. LNCS, vol. 4393. Springer, Heidelberg (2007)
2. Brenguier, R., Clemente, L., Hunter, P., Pérez, G.A., Randour, M., Raskin, J.-F., Sankur, O., Sassolas, M.: Non-zero sum games for reactive synthesis. In: Dediu, A.-H., Janoušek, J., Martín-Vide, C., Truthe, B. (eds.) LATA 2016. LNCS, vol. 9618, pp. 3–23. Springer, Cham (2016). doi:10.1007/978-3-319-30000-9_1
3. Brenguier, R., Raskin, J.-F., Sankur, O.: Assume-admissible synthesis. In: CONCUR, LIPIcs, vol. 42, pp. 100–113. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2015)
4. Brenguier, R., Raskin, J.-F., Sassolas, M.: The complexity of admissibility in omega-regular games. In: CSL-LICS, pp. 23:1–23:10. ACM (2014)
5. Brihaye, T., Bruyère, V., Meunier, N., Raskin, J.-F.: Weak subgame perfect equilibria and their application to quantitative reachability. In: CSL, LIPIcs, vol. 41, pp. 504–518. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2015)
6. Bruyère, V., Meunier, N., Raskin, J.-F.: Secure equilibria in weighted games. In: CSL-LICS, pp. 26:1–26:26. ACM (2014)
7. Chatterjee, K., Henzinger, T.A., Jurdzinski, M.: Games with secure equilibria. Theor. Comput. Sci. **365**, 67–82 (2006)
8. Chatterjee, K., Doyen, L., Filiot, E., Raskin, J.-F.: Doomsday equilibria for omega-regular games. In: McMillan, K.L., Rival, X. (eds.) VMCAI 2014. LNCS, vol. 8318, pp. 78–97. Springer, Heidelberg (2014). doi:10.1007/978-3-642-54013-4_5
9. Chatterjee, K., Doyen, L., Henzinger, T.A.: Quantitative languages. ACM Trans. Comput. Log. **11**, 23:1–23:38 (2010)
10. Chatterjee, K., Henzinger, T.A.: Assume-guarantee synthesis. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 261–275. Springer, Heidelberg (2007). doi:10.1007/978-3-540-71209-1_21
11. Pril, J., Flesch, J., Kuipers, J., Schoenmakers, G., Vrieze, K.: Existence of secure equilibrium in multi-player games with perfect information. In: Csuhaj-Varjú, E., Dietzfelbinger, M., Ésik, Z. (eds.) MFCS 2014. LNCS, vol. 8635, pp. 213–225. Springer, Heidelberg (2014). doi:10.1007/978-3-662-44465-8_19
12. Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Solan, E., Vrieze, K.: Perfect-information games with lower-semicontinuous payoffs. Math. Oper. Res. **35**, 742–755 (2010)
13. Flesch, J., Predtetchinski, A.: A characterization of subgame perfect equilibrium plays in borel games of perfect information. Math. Oper. Res. (2017, to appear)

14. Fudenberg, D., Levine, D.: Subgame-perfect equilibria of finite- and infinite-horizon games. J. Econ. Theor. **31**, 251–268 (1983)
15. Grädel, E., Ummels, M.: Solution concepts and algorithms for infinite multiplayer games. In: New Perspectives on Games and Interaction, vol. 4, pp. 151–178. University Press, Amsterdam (2008)
16. Kuhn, H.W.: Extensive games and the problem of information, pp. 46–68. Classics in Game Theory (1953)
17. Kupferman, O., Perelli, G., Vardi, M.Y.: Synthesis with rational environments. Ann. Math. Artif. Intell. **78**(1), 3–20 (2016)
18. Le Roux, S.: Infinite subgame perfect equilibrium in the Hausdorff difference. In: Hajiaghayi, M.T., Mousavi, M.R. (eds.) TTCS 2015. LNCS, vol. 9541, pp. 147–163. Springer, Cham (2016)
19. Le Roux, S., Pauly, A.: Infinite sequential games with real-valued payoffs. In: CSL-LICS, pp. 62:1–62:10. ACM (2014)
20. Nash, J.F.: Equilibrium points in n-person games. In: PNAS, vol. 36, pp. 48–49. National Academy of Sciences (1950)
21. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: POPL, pp. 179–190. ACM Press (1989)
22. Purves, R.A., Sudderth, W.D.: Perfect information games with upper semicontinuous payoffs. Math. Oper. Res. **36**(3), 468–473 (2011)
23. Rubinstein, A.: Comments on the interpretation of game theory. Econometrica **59**, 909–924 (1991)
24. Selten, R.: Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. Zeitschrift für die gesamte Staatswissenschaft **121**, 301–324, 667–689 (1965)
25. Shen, X.S., Yu, H., Buford, J., Akon, M.: Handbook of Peer-to-Peer Networking. Springer, Heidelberg (2010)
26. Solan, E., Vieille, N.: Deterministic multi-player Dynkin games. J. Math. Econ. **39**, 911–929 (2003)
27. Ummels, M.: Rational behaviour and strategy construction in infinite multiplayer games. In: Arun-Kumar, S., Garg, N. (eds.) FSTTCS 2006. LNCS, vol. 4337, pp. 212–223. Springer, Heidelberg (2006). doi:10.1007/11944836_21