Conference contribution :
Le Roux, S. & Pauly, A. (2014). *Infinite Sequential Games with Real-valued Payoffs.* CSL-LICS '14, ACM.
http://dx.doi.org/10.1145/2603088.2603120

Infinite sequential games with real-valued payoffs

Stéphane Le Roux
Department of Mathematics
Technische Universität Darmstadt, Germany
leroux@mathematik.tu-darmstadt.de

Arno Pauly
Computer Laboratory
University of Cambridge, United Kingdom
Arno.Pauly@cl.cam.ac.uk

Abstract

We investigate the existence of certain types of equilibria (Nash, ε-Nash, subgame perfect, ε-subgame perfect) in infinite sequential games with real-valued payoff functions depending on the class of payoff functions (continuous, upper semi-continuous, Borel) and whether the game is zero-sum. Our results hold for games with two or up to countably many players.

Several of these results are corollaries of stronger results that we establish about equilibria in infinite sequential games with some weak conditions on the occurring preference relations. We also formulate an abstract equilibrium transfer result about games with compact strategy spaces and open preferences. Finally, we consider a dynamical improvement rule for infinite sequential games with continuous payoff functions.

Categories and Subject Descriptors F.4.1 [*Mathematical logic*]

General Terms Theory

Keywords infinite games, solution concepts, determinacy, Borel measurability

1. Introduction

The present article continues the research programme to investigate sequential games in a very general setting, which was initiated by the first author in [15, 16]. This programme reunites two mostly separate developments in the study of games. On the one hand, the first development is the investigation of variations on solution concepts for games, and of different formalizations of the preferences of the players, which primarily happened inside game theory proper. Related to that, game theory has also seen an interest in relaxing the continuity and convexity assumptions of NASH's original existence theorem [20], both for an example and references to further work see [21]. So typically, the countably infinite is absent from game theory: sets are either finite, or in cases such as randomized strategies, have the structure of the continuum.

On the other hand, the study of infinite sequential games has a long history in logic. Many variations on the rules of games have been studied, albeit mostly restricted to zero-sum games with two players and two outcomes. Thus, here the continuum is entirely absent (discounting its internal occurrence in the set of potential plays), and the countably infinite is only used for time, not for e.g. the number of agents or outcomes.

In our work we study infinite sequential games with perfect information (i.e. generalized Gale-Stewart games [9]) in a setting as general as possible. We can have any countable number of players, and we investigate various ways to represent preferences. Some results do put restrictions on the number of distinguished outcomes (as being countable), and some require a (generalized) zero-sum condition. A similar synthesis of the approaches is found in [7, 10, 26, 31].

A classical result by MARTIN [17] established that such a game played by two players who only care about whether or not the play falls into some fixed Borel set is determined, i.e. admits a winning strategy for one of the players. While determinacy can be understood as a special case of existence of Nash equilibrium, Nash equilibrium is often regarded as an unsatisfactory solution concept for sequential games in game theory. Subgame perfect equilibria are a more convincing solution concept from a rationality perspective. This article studies these two concepts, and as optimal strategies are not always available, we also investigate the existence of ε-Nash equilibria and ε-subgame perfect equilibria.

The proofs that we provide fall into two broad categories: some of our results are obtained by lifting Borel determinacy to more complicated settings, similar to [15, 16] or to a sketched observation by MERTENS and NEYMAN in [19]; other results are based on topological arguments to show strong existence results, albeit at the cost of continuity requirements in the game characterizations. Both proof techniques provide general results in rather abstract settings. In this sense, our main results are Theorem 6 on the one hand, and Theorems 18, 19, 22 (which share parts of their proofs) on the other hand.

A common theme of the results is that they are transfer principles: they tell us how to take a pre-existing result on a more restricted class of games and obtain from it a result for a more general class of games. Theorem 6 is then used to transfer the existence of subgame perfect equilibria from finite sequential games to certain infinite ones. With Theorems 18, 19, 22, we can in particular extend Borel determinacy to yield equilibria in multi-player multi-outcome settings.

In an attempt to make the rather abstract results somewhat more accessible, we shall consider in addition the corollaries obtained in the situation where the goals of the players are to maximize real-valued payoff functions. Distinguishing properties here are continuity, upper or lower semicontinuity and Borel measurability. These settings have been studied before [7, 8, 19, 26], and usually our corollaries improve upon known results by extending them from the case of finite players to the case of countably many players. In particular, prior proofs sometimes use induction on the number of players, and thus do not generalize to the countably infinite. An overview of past and new results is given in the table on Page 3.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.

http://dx.doi.org/10.1145/2603088.2603120

Since our existence result for Nash equilibria in games with continuous payoff functions in Corollary 8 does not, prima facie, come with a means to construct these, we do consider a dynamical updating procedure in Section 6 which has precisely the Nash equilibria of such a game as its accumulation points[1].

Our results showcase which requirements are actually needed for which aspects of determinacy, and as such may contribute to the understanding of strategic behaviour in general.

Additionally, the emergence of quantitative objectives in addition to qualitative structure in traditional verification/synthesis games [2], such as mean-payoff parity games [4, 5], provides an area of applications for abstract theorems about the existence of equilibria. Existence results in such settings are usually not trivial, but are proven together with the introduction of the setting – thus we do not answer open questions, but are hopeful that our results may be useful in the future.

Having results for countably many players is important for applications if multi-agent interactions in open systems are studied. In order to employ an equilibrium existence result for finitely many players, a bound on the number of agents involved in the interaction might need to be common knowledge from the beginning on. Our results on the other hand easily enable a setting where additional agents may join the interaction later on, and only the number of agents who have acted in the past is finite.

2. Background

In our most abstract definition, a game is a tuple $\langle A, (S_a)_{a\in A}, (\prec_a)_{a\in A}\rangle$ consisting of a non-empty set A of *agents* or *players*, for each agent $a \in A$ a non-empty set S_a of *strategies*, and for each agent $a \in A$ a *preference* relation $\prec_a \subseteq \left(\prod_{a\in A} S_a\right) \times \left(\prod_{a\in A} S_a\right)$. The generic setting suffices to introduce the notion of a Nash equilibrium: A *strategy profile* $\sigma \in \left(\prod_{a\in A} S_a\right)$ is called a Nash equilibrium, if for any agent $a \in A$ and any strategy $s_a \in S_a$ we find $\neg\left(\sigma \prec_a \sigma_{a\mapsto s_a}\right)$, where $\sigma_{a\mapsto s_a}$ is defined by $\sigma_{a\mapsto s_a}(b) = \sigma(b)$ for $b \in A \setminus \{a\}$ and $\sigma_{a\mapsto s_a}(a) = s_a$. In words, no agent prefers over a Nash equilibrium some other situation that only differs in her choice of strategy.

We will give additional structure to games in two primary ways: In Section 3 we add topologies to the strategy spaces, and then impose some topological constraints on both strategy spaces and preferences. Beyond that, we will consider games where strategy spaces and preferences are derived objects from more structured variants of games. One such variant is the infinite sequential game:

Definition 1 (Infinite sequential game, cf. [15, Definition 1.1])**.** An *infinite sequential game* is an object
$\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$ complying with the following.

1. A is a non-empty set (of agents).
2. C is a non-empty set (of choices).
3. $d : C^* \to A$ (assigns a decision maker to each stage of the game).
4. O is a non-empty set (of possible outcomes of the game).
5. $v : C^\omega \to O$ (uses outcomes to value the infinite sequences of choices).
6. Each $\prec_a$ is a binary relation over O (modelling the preference of agent a).

The intuition behind the definition is that agents take turns to make a choice. Whose turn it is depends on the past choices via the function d. Over time, the agents thus jointly generate some infinite sequence, which is mapped by v to the outcome of the game. Note that using a single set of actions C for each step just simplifies the notation, a generalization to varying action sets is straightforward.

The infinite sequential games are linked to abstract games as follows: the agents remain the agents and the strategies of agent a are the functions $s_a : d^{-1}(\{a\}) \to C$. We can then safely regard a strategy profile as a function $\sigma : C^* \to C$ whose *induced play* is defined below, where for an infinite sequence $p \in C^\omega$ we let p_n be its n-th value, and $p_{\leq n} = p_{<n+1} \in C^*$ be its finite prefix of length n.

Definition 2 (Induced play and outcome, cf. [15, Definition 1.3])**.** Let $s : C^* \to C$ be a strategy profile. The *play* $p = p^\lambda(s) \in C^\omega$ *induced by* s *starting at* $\lambda \in C^*$ is defined inductively through its prefixes: $p_n = \lambda_n$ for $n \leq |\lambda|$ and $p_n := s(p_{<n})$ for $n > |\lambda|$. Also, $v \circ p_\lambda(s)$ is the *outcome induced by* s *starting at* λ. The play (resp. outcome) induced by s is the play (resp. outcome) induced by s starting at ε.

In the usual way to regard an infinite sequential game as a special abstract game, an agent prefers a strategy profile σ to σ', iff he prefers the outcome induced by σ to the outcome induced by σ'. And indeed we shall call a strategy profile of an infinite sequential game a Nash equilibrium, iff it is a Nash equilibrium with these preferences. In a certain notation overload, we will in particular use the same symbols for the preferences over strategy profiles and the preferences over outcomes.

However, there is a certain criticism of this choice as being not rational: Essentially, the resulting concept of a Nash equilibrium means that players can use empty threats – declarations they would play in a certain way from a position onwards, even if that would be against their own interests once the position is reached, as long as this threat keeps other players from moving to that position. This can be fixed by considering subgame perfect equilibria [28]. We understand these as the Nash equilibria derived from a different translation of preferences from infinite sequential games to abstract games. (Similar remarks were made in [12, Lemma 144 in Section 7.2.3, Section 7.3.2].)

Definition 3. Given an infinite sequential game $\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$, let the subgame perfect preferences[2] $\prec_a^{sgp} \subseteq C^{C^*} \times C^{C^*}$ be defined by $\sigma \prec_a^{sgp} \sigma'$ iff $\exists\lambda \in C^*$ such that $p^\lambda(\sigma) \prec_a p^\lambda(\sigma)$.

The subgame perfect equilibria of $\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$ are the Nash equilibria of $\langle A, (C^{d^{-1}(\{a\})})_{a\in A}, (\prec_a^{sgp})_{a\in A}\rangle$.

We consider a further variant, namely the *infinite sequential games with real-valued payoffs*, which can (but do not have to) be understood as a special case of infinite sequential games.

Definition 4. An infinite sequential game with real-valued payoffs is a tuple $\langle A, C, d, (f_a)_{a\in\mathbb{N}}\rangle$ where A, C, d are as above, and $f_a : C^\omega \to [0, 1]$ is the payoff function of player a.

Such a game can be identified with the infinite sequential game $\langle A, C, d, [0, 1]^A, v, (\prec_a)_{a\in A}\rangle$ where $v(p) = \prod_{a\in\mathbb{A}} f_a(p)$ and $x \prec_a y$ iff $x_a < y_a$.

As with the introduction of subgame perfect equilibria, we can consider infinite sequential games with real-valued payoffs as infinite sequential games in a different way, which then gives rise to another commonly studied equilibrium concept, namely ε-Nash equilibria. Depending on how we then translate from infinite sequential games to abstract games, we obtain also ε-subgame perfect equilib-

[1] It should be pointed out that this does not yield a constructive proof of the existence of Nash equilibria, as finding accumulation points of a sequence is not a computable operation [3]. In fact, the existence of Nash equilibria generally cannot be proven constructively, cf. also [23].

[2] Note that the translation of preferences in the following definition does not preserve acyclicity. Preservation could be ensured, *e.g.*, by giving the nodes a linear "priority" order, in a lexicographic fashion. This, however, would complicate the definition against little benefit for the point that we want to make.

Overview

Payoff functions	Type	ε-Nash	Nash	ε-subgame perfect	subgame perfect
continuous	finitely many players	yes -	yes -	yes -	yes [8, Corollary 4.2]
continuous	countably many players	yes -	yes -	yes -	yes Corollary 8
upper semi-continuous	finitely many players	yes -	yes Corollary 23	yes [26, Theorem 2.1]	no Example 28
upper semi-continuous	countably many players	yes -	yes Corollary 23	?	no Example 28
lower semi-continuous	finitely many players	yes -	no Example 11	yes [7, Theorem 2.3]	no -
lower semi-continuous	countably many players	yes -	no Example 11	no [7, Subsection 4.3]	no -
Borel	zero-sum	yes -	no Example 11	yes Corollary 21	no -
Borel	finitely many players	yes Corollary 10	no Example 11	no Example 28	no -
Borel	countably many players	yes Corollary 10	no Example 11	no Example 28	no -

Note: The table shows for any relevant combination of properties of the payoff functions, type of the game and type of equilibria whether such equilibria exist always or whether there is a counterexample. The results as listed here pertain to having finitely many choices at each stage of the game. There is no difference between two player games and games with finitely many players in any situation we investigate. As we consider perfect information games only, any zero-sum game is understood to be a two player game. The combination of semi-continuity and zero-sum would imply continuity, and is thus left out. For continuous payoff functions, we already see complete positive results without the zero-sum condition, and thus do not mention it explicitly either. Both Corollary 10 and Example 11 seem to be folklore results.

ria. Given some $\varepsilon > 0$, we define the relation $\prec_a^\varepsilon \subseteq [0,1]^A \times [0,1]^A$ by $x \prec_a^\varepsilon y$ iff $y_a - x_a \geq \varepsilon$. Using $\prec_a^\varepsilon$ in place of $\prec_a$ in Definition 4 then provides the above-mentioned equilibrium notions.

For infinite sequential games with real-valued payoffs, every Nash equilibrium (w.r.t. the standard preferences) is an ε-Nash equilibrium; and every subgame perfect equilibrium is an ε-subgame perfect equilibrium. For infinite sequential games, every subgame perfect equilibrium is a Nash equilibrium, in particular, any ε-subgame perfect equilibrium is an ε-Nash equilibrium.

If the preferences of the players satisfy $\prec_a = \prec_b^{-1}$, where $x \prec^{-1} y \Leftrightarrow y \prec x$, we speak of an abstract zero-sum game. As we consider perfect information games only, any abstract zero-sum game is understood to be a two player game. (In a zero-sum game, there are only two distinct preferences. In a perfect information game, we can identify all players with identical preferences, as there will be no coordination problems.)

We proceed to recall a few more notions that are only tangentially related to the formulation of our results, but that do show up in the proofs.

Definition 5. A two-player win-loose game is a tuple $\langle C, D, W\rangle$ with $D \subseteq C^*$ and $W \subseteq C^\omega$. It corresponds to the infinite sequential game $\langle \{a,b\}, C, d, \{0,1\}, v, \{<, <^{-1}\}\rangle$ where d is defined via $d^{-1}(\{a\}) = D$ and v is defined via $v^{-1}(\{1\}) = W$.

Finally, we will extend the notion of the induced play. Given some partial function $s :\subseteq C^* \to C$, we define the consistency set $P(s) \subseteq C^\omega$ by:

$$P(s) = \{p(\sigma) \mid \sigma : C^* \to C \wedge \sigma|_{\mathrm{dom}(s)} = s\}$$

3. The continuity argument

A strong transfer result can be obtained using topological arguments alone, with the reasoning being particularly well-adapted to a formulation in synthetic topology (originally [6], [24] for a short introduction). Consider games in normal form, with potentially countably many agents with strategy spaces $\mathbf{S}_1, \mathbf{S}_2, \ldots$. Our first condition is that each $\mathbf{S}_i$ be compact (subsequently, by Tychonoff's Theorem, also $\Pi_{i\in\mathbb{N}}\mathbf{S}_i$). This restriction is very common and usually combined with continuity of the outcome function, as it avoids pathological games such as *pick-the-largest-natural-number*. Our second condition is that each preference relation $\prec_i$ is open (as a subset of $(\Pi_{i\in\mathbb{N}}\mathbf{S}_i) \times (\Pi_{i\in\mathbb{N}}\mathbf{S}_i)$).[3] In the reading of synthetic topology, this means that any agent will be able to eventually confirm that he prefers a given strategy profile to another, provided he does indeed do so. We shall call a class of games $\mathcal{G}$ satisfying these conditions (in a uniform way) to be compact-strategies, open-preferences. Uniformity here means that we assume a topology on $\mathcal{G}$ such that the function mapping a game to the preferences is continuous itself.

We will write $\mathcal{O}(\mathbf{X})$ for the hyperspace of open subsets of $\mathbf{X}$, and $\mathcal{K}(\mathbf{X})$ for the hyperspace of compact sets. By $\mathcal{C}(\mathbf{X}, \mathbf{Y})$ we denote the space of continuous functions from $\mathbf{X}$ to $\mathbf{Y}$, in particular $\mathcal{C}(\mathbb{N}, \mathbf{X})$ we denote the space of sequences in $\mathbf{X}$. For precise definitions, see [24]. There we also find that the following operations are continuous:

1. [4] $\exists : \mathcal{O}(\mathbf{X} \times \mathbf{Y}) \to \mathcal{O}(\mathbf{X})$, defined by $\exists(U) = \{x \in \mathbf{X} \mid \exists y \in \mathbf{Y}\ (x, y) \in U\}$
2. $\bigcup : \mathcal{C}(\mathbb{N}, \mathcal{O}(\mathbf{X})) \to \mathcal{O}(\mathbf{X})$
3. ${}^C : \mathcal{O}(\mathbf{X}) \to \mathcal{K}(\mathbf{X})$, provided that $\mathbf{X}$ is compact.
4. $\mathrm{NonEmptyValue} : \mathcal{C}(\mathbf{X}, \mathcal{K}(\mathbf{Y})) \to \mathcal{O}(\mathbf{X})$ defined by $\mathrm{NonEmptyValue}(f) = \{x \in \mathbf{X} \mid f(x) \neq \emptyset\}$.

[3] There actually is a third condition, that any strategy space is overt. A space $\mathbf{X}$ is overt, if $\{\emptyset\} \subseteq \mathcal{O}(\mathbf{X})$ is a closed set, i.e. if there is a way to detect non-emptiness of open subsets. This condition could only ever fail in a constructive reading, but is always valid for topological spaces in classical logic. Synthetic topology however would also allow us to read *continuous map* to mean *computable map*, in which case overtness becomes non-trivial. In this reading, though, we actually obtain an algorithmic result.

[4] In general, the continuity of this map would require $\mathbf{Y}$ to be overt. As explained above, in a classical reading, this condition is always satisfied.

Theorem 6. Let $\mathcal{G}$ be compact-strategies, open-preferences, and let $\mathcal{G}' \subseteq \mathcal{G}$ be a dense subclass. If every $G \in \mathcal{G}'$ has a Nash equilibrium, then every $G \in \mathcal{G}$ has a Nash equilibrium.

Proof. By combining our continuous operations, we may obtain the set of all games in $\mathcal{G}$ with a Nash equilibrium as an open set in the following way:

$$\mathcal{NE} := \text{NonEmptyValue}\left(G \mapsto \left(\bigcup_{i\in\mathbb{N}} \exists(\prec_i^G)\right)^C\right) \in \mathcal{O}(\mathcal{G})$$

Formulating the individual steps in words: With $\prec_i$ being open[5], we immediately obtain that the set of all strategy profiles such that player i has a better response is uniformly open in the game. Taking the union over all players again yields an open set, which complement now is the closed set $\text{NE}(G)$ of all Nash equilibria of the respective game G. As this is a subset of the compact space $(\Pi_{i\in\mathbb{N}}\mathbf{S}_i)$, we can even treat $\text{NE}(G)$ as a compact set uniformly in G. By the synthetic definition of compactness, we obtain that $\{G \in \mathcal{G} \mid \text{NE}(G) = \emptyset\} \subseteq \mathcal{G}$ is an open set.

Because we have assumed $\mathcal{G}'$ to be dense in $\mathcal{G}$, we see that if any game in $\mathcal{G}$ would fail to have a Nash equilibrium, this would imply that some game in $\mathcal{G}'$ would fail, too, contrary to the assumption. □

Consider sequential games with continuous payoff-functions and finite choices sets. We find:

Lemma 7.
1. The subgame-perfect preferences produce a compact-strategies, open-preference class $\mathcal{S}$.
2. The games with payoffs fully determined after finitely many moves are a dense subset $\mathcal{S}^f \subseteq \mathcal{S}$.
3. All games in $\mathcal{S}^f$ have a Nash equilibrium.

Proof.
1. First, note that the mapping $(\lambda, \sigma) \mapsto p_\lambda(\sigma)$ is continuous, so for any fixed $\lambda \in C^*$, $\{(\sigma, \sigma') \mid f_a(p_\lambda(\sigma)) < f_a(p_\lambda(\sigma'))\}$ is an open set. By taking countable union, we learn that $\prec_a^{sgp}$ is open. Compactness and overtness of the strategy spaces are straight-forward.
2. As the argument in (1) is uniform in the continuous functions f_a, it suffices to argue that the payoff functions $f : C^\omega \to [0,1]$ depending only on some finite prefix of the input are dense in $\mathcal{C}(C^\omega, [0,1])$. A countable base for the applicable topology is found in all
$$\{f \mid \forall\lambda \in C^k, p \in C^\omega \,.\, p_{\leq k} = \lambda \Rightarrow f(p) \in (x_\lambda, y_\lambda)\}$$
for
$$k \in \mathbb{N}, x_{(\cdot)}, y_{(\cdot)} : C^k \to \mathbb{Q} \cap [0,1]$$
A base element is non-empty, iff $\forall\lambda \in C^k\ x_\lambda < y_\lambda$; and then it will contain the function $f_0 : C^\omega \to [0,1]$ defined via $f_0(p) = \frac{1}{2}\left(x_{p_{\leq k}} + y_{p_{\leq k}}\right)$, which clearly depends only on the prefix of length k of its argument.
3. As the actions of the players beyond the finite prefix determining the outputs is irrelevant, and taking into consideration the definition of the subgame-perfect preferences, the claim is that any finite game in extensive form has a subgame perfect equilibrium. This well-known result is easily proven by backwards induction: Let the players who move last pick an optimal (for them) choice. Then the players who move second-but-last have guaranteed outcomes associated with their moves, so they can optimize, and so on.

□

[5] And $\mathbf{S}_i$ being overt, see above.

Corollary 8 ([6]). Any sequential game with continuous payoff-functions and finitely many choices has a subgame perfect Nash equilibrium.

In a similar fashion we could consider multi-player multi-outcome Blackwell games [1, 18] with continuous payoff functions (into $[0,1]$). Here each vertex of a finitely branching tree is labeled with a finite game in normal form, together with a bijection from the children of the vertex to the (pure) strategy profiles in the finite game. The players choose a randomized strategy in each finite game, this then induces a probability distribution over the paths through the tree. As the payoff for a given player is then $\int f_i d_{\mu_{s_1,\ldots}}$, we again see that the continuity argument applies (and even w.r.t. subgame perfect preferences):

Corollary 9 ([7]). Any multi-player multi-outcome Blackwell game with continuous payoff functions has a subgame-perfect Nash equilibrium.

4. From discrete to real-valued payoffs

A rather simple argument allows us to transfer existence theorems for equilibria in games with finitely many players, from Borel-measurable valuations over finitely many outcomes to Borel-measurable real-valued payoff functions, if one is willing to replace the original notions by their ε-counterparts. If $v : \mathbf{S} \to [0,1]^m$ is the Borel-measurable payoff function, then for any $k \in \mathbb{N}$ we define $v_{(k)} : \mathbf{S} \to \{1, \ldots, (k+1)\}^m$ by $v_{(k)}^{-1}(i_1, \ldots, i_m) := v^{-1}\left([\frac{i_1-1}{k}, \frac{i_1}{k}) \times \ldots \times [\frac{i_m-1}{k}, \frac{i_m}{k})\right)$. Then any $v_{(k)}$ is again a Borel measurable valuation. Furthermore, we define the preferences $\prec_n$ for the n-th player by $(i_1, \ldots, i_m) \prec_n (j_1, \ldots, j_m)$ iff $i_n < j_n$. Now any Nash equilibrium of the resulting game is a $\frac{1}{k}$-Nash equilibrium of the original game, and every subgame perfect equilibrium of the resulting game is a $\frac{1}{k}$-subgame perfect equilibrium of the original game. Note that the same argument can be used for games with countably many players and countably many outcomes, without any restriction on the branching of the game-tree.

Corollary 10. [8] Sequential games with countably many players and Borel-measurable payoff functions with upper-bounds admit ε-Nash equilibria.

Proof. By combining the statement of [15, Theorem 3.2] with the argument above. □

5. Infinite sequential games with infinitely ascending preferences

As soon as the continuity requirement for the payoff function (or, more generally, the openness of the preferences) is dropped, Nash equilibria may fail to exist. We provide a generic folklore counterexample, and will proceed to demonstrate that the underlying feature is essential for the failure of existence of Nash equilibria. The counterexample only requires a single player, and its payoff function is in a sense *the least discontinuous* payoff function, and in particular is Δ_2^0-measurable.

Example 11. Let the payoff function $P : \{0,1\}^\mathbb{N} \to [0,1]$ for the single player be defined by $P(1^n0p) = \frac{n}{n+1}$ for all $p \in \{0,1\}^\mathbb{N}$

[6] This extends [8, Corollary 4.2] from finitely many players to countably many players.

[7] This extends [8, Theorem 6.1] from finitely many players to countably many players.

[8] In his survey [19], MERTENS sketches an observation by himself and NEYMAN that one may use Borel determinacy to directly obtain this result.

and $P(1^{\mathbb{N}}) = 0$. As P does not attain its supremum, the resulting game cannot have a Nash equilibrium.

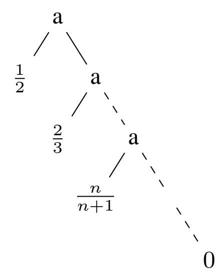

We proceed to show that the presence of a converging sequence of plays $(p^n)_{n\in\mathbb{N}}$ such that a player prefers p^{n+1} to p^n, but prefers any p^n to $\lim_{i\to\infty} p^i$, is a crucial feature of the example above to have no Nash equilibrium. The proof will be an adaption of the main result of [15] by the first author. Under the additional assumption of antagonistic preferences in a two-player game, we can even obtain subgame perfect equilibria.

In this section the preferences of the players are restricted to strict weak orders, so we recall their definition below.

Definition 12 (Strict weak order)**.** A relation $\prec$ is called a *strict weak order* if it satisfies:

$$\forall x, \quad \neg(x \prec x)$$
$$\forall x, y, z, \quad x \prec y \wedge y \prec z \Rightarrow x \prec z$$
$$\forall x, y, z, \quad \neg(x \prec y) \wedge \neg(y \prec z) \Rightarrow \neg(x \prec z)$$

Definition 13 below slightly rephrases Definitions 2.3 and 2.5 from [15]: the guarantee of a player is the smallest set of outcomes that is upper-closed w.r.t. the strict-weak-order preference of the player and includes every incomparability classe (of the preference) that contains any outcome compatible with a given strategy of the player in the subgame at a given node of a given infinite sequential game. The best guarantee of a player consists of the intersection of all her guarantees over the set of strategies.

Definition 13 (Agent (best) guarantee)**.** Let $\langle A, C, d, O, v, (\prec_a)_{a\in A}\rangle$ be a game where the $\prec_a$ are strict weak orders.

$$\forall a \in A, \forall \gamma \in C^*, \forall s : d^{-1}(a) \to C, \quad g_a(\gamma, s) :=$$
$$\{o \in O \mid \exists p \in P(s|_{\gamma C^\omega}) \cap \gamma C^\omega, \neg(o \prec_a v(p))\}$$
$$G_a(\gamma) := \bigcap_s g_a(\gamma, s)$$

We write $g_a(s)$ and G_a instead of $g_a(\gamma, s)$ and $G_a(\gamma)$ when γ is the empty word.

Lemma 2.4. from [15] still holds without major changes in the proofs, so we do not display it, but note that when speaking about $\prec_a$-terminal intervals (which are upper-closed sets), we now actually refer to the terminal intervals of the lift of $\prec_a$ from outcomes to the equivalence classes of outcomes induced by the strict weak order. Also, we collect some more useful facts in Observation 14 below.

Observation 14. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$, let $a \in A$, assume that $\prec_a$ is a strict weak order, and let $\gamma \in C^*$.

1. $d(\gamma) \neq a \Rightarrow G_a(\gamma) = \cup_{c\in C} G_a(\gamma \cdot c)$
2. $d(\gamma) = a \Rightarrow G_a(\gamma) = \cap_{c\in C} G_a(\gamma \cdot c)$
3. $d(\gamma) = a \wedge |C| < \infty \Rightarrow \exists c \in C, G_a(\gamma) = G_a(\gamma \cdot c)$

Proof. For example, for 2. note that $G_a(\gamma) = \cap_s g_a(\gamma, s) = \cap_{c\in C} \cap_{s(\gamma)=c} g_a(\gamma, s) = \cap_{c\in C} \cap_s g_a(\gamma \cdot c, s)$ $= \cap_{c\in C} G_a(\gamma \cdot c)$. □

Lemma 15. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ be a game where C is finite, let $a \in A$, and let us assume the following.

1. $\prec_a$ is a strict weak order.
2. For every play $p \in C^\omega$ and every sequence of plays $(p^n)_{n\in\mathbb{N}}$ converging towards p, if $v(p^n) \prec_a v(p^{n+1})$ for all n, then $v(p^n) \prec_a v(p)$ for all n.

Then there exists $s \in S_a$ such that $g_a(s) = G_a$.

Proof. Let $s_0 : d^{-1}(a) \to C$ be a strategy of Player a and let us build inductively a sequence $(s_n)_{n\in\mathbb{N}}$ of strategies of Player a, as follows, where case 3. implicitly invokes Observation 14.

- Let $s_{n+1}|_{C^{<n}} := s_n|_{C^{<n}}$.
- For all $\gamma \in C^n \backslash d^{-1}(a)$, let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$.
- For all $\gamma \in C^n \cap d^{-1}(a)$,
  1. if $g_a(\gamma, s_n) \subseteq G_a$ then let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$,
  2. if $G_a \subsetneq g_a(\gamma, s_n)$ and there exists $\mu : d^{-1}(a) \cap \gamma C^* \to C$ such that $g_a(\gamma, \mu) \subseteq G_a$, let $s_{n+1}|_{\gamma C^*} := \mu$,
  3. otherwise[9] let $s_{n+1}(\gamma) := c$ such that $G_a(\gamma \cdot c) = G_a(\gamma)$, and let $s_{n+1}|_{\gamma C C^*} := s_n|_{\gamma C C^*}$.

Let s be the limit strategy of the sequence $(s_n)_{n\in\mathbb{N}}$ and first note that, using Observation 14, one can prove by induction on γ that $G_a(\gamma) \subseteq G_a$ for every $\gamma \in C^*$ that is compatible with s. Next, let $p \in P(s)$ be a path compatible with s. If p has a prefix γ that fell into Cases 1. or 2. during the recursive construction above, then $v(p) \in G_a$, so let us now assume that case 3. applies at every node $p_{<n} \in d^{-1}(a)$. If such nodes are finitely many, let $p_{<n}$ be the deepest one, so that $G_a = G_a(p_{<n+1}) = g_a(p_{n+1}, s)$ holds for all strategies s, which implies $v(p) \in G_a$. Let us now assume that such nodes are infinitely many. If $G_a(p_{<n}) \subsetneq G_a$ for some $p_{<n} \in d^{-1}(a)$, there exists $\mu : d^{-1}(a) \cap p_{<n}C^* \to C$ such that $G_a(p_{<n}) \subseteq g_a(p_{<n}, \mu) \subsetneq G_a$ since $G_a(p_{<n}) := \bigcap_s g_a(p_{<n}, s)$ by definition, which would mean that Case 1. or 2. applies; so $G_a(p_{<n}) = G_a$ for all $p_{<n} \in d^{-1}(a)$, and subsequently for all n. Also, since the guarantee is never witnessed (through Case 2.) at any node $p_{<n} \in d^{-1}(a)$, for every $o \notin G_a$ and every $n \in \mathbb{N}$ there exists $q \in C^\omega$ such that $q_{<n} = p_{<n}$ and $o \prec_a v(q) \notin G_a(p_{<n}) = G_a$. Let us assume for a contradiction that $v(p) \notin G_a$, let $q^0 \in C^\omega$ be such that $v(p) \prec_a v(q_0) \notin G_a$, and let us define a sequence $(q^n)_{n\in\mathbb{N}}$ inductively such that $v(q^n) \prec_a v(q^{n+1}) \notin G_a$ and $q^n_{<n} = p_{<n}$ for all n. It implies that $v(p) \prec_a v(p)$, contradiction. □

Conditions 1 and 2 of Lemma 16 below are the same as in Lemma 15, Condition 3 is the conclusion of Lemma 15, which is similar to the conclusion of Lemma 16, after a key quantifier inversion, though. We do not merge these two lemmas into one, nonetheless, since no cardinality assumption on C is needed for Lemma 16.

Lemma 16. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ be a game, let $a \in A$, and let us assume the following.

1. $\prec_a$ is a strict weak order.
2. For every play $p \in C^\omega$ and every sequence of plays $(p^n)_{n\in\mathbb{N}}$ converging towards p, if $v(p^n) \prec_a v(p^{n+1})$ for all n, then $v(p^n) \prec_a v(p)$ for all n.
3. For all $\gamma \in C^\omega$, there exists s such that $g_a(\gamma, s) = G_a(\gamma)$.

Then there exists s such that $g_a(\gamma, s) = G_a(\gamma)$ for all $\gamma \in C^\omega$.

Proof. We proceed similarly as in the proof of Lemma 15. Let s_0 be a strategy of Player a and let us build inductively a sequence $(s_n)_{n\in\mathbb{N}}$ of strategies of Player a. The recursive definition below is

[9] Note that due to the properties of a strict weak order, the sets of the form $g_a(\lambda, s)$ and $G_a(\lambda)$ are linearly ordered by inclusion $\subseteq$. Thus, $G_a \subsetneq g_a(\gamma, s_n)$ holds in this case, too.

different from the one in the proof of Lemma 15 in three respects: the three occurrences of G_a in Cases 1. and 2. are replaced with $G_a(\gamma)$; Case 3. is deleted since we now know that it never applies; and two inclusions are replaced with equalities.

- Let $s_{n+1}|_{C^{<n}} := s_n|_{C^{<n}}$
- For all $\gamma \in C^n \backslash d^{-1}(a)$, let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$.
- For all $\gamma \in C^n \cap d^{-1}(a)$,
  1. if $g_a(\gamma, s_n) = G_a(\gamma)$ then let $s_{n+1}|_{\gamma C^*} := s_n|_{\gamma C^*}$,
  2. if $G_a(\gamma) \subsetneq g_a(\gamma, s_n)$, let $s_{n+1}|_{\gamma C^*} := \mu$ where $\mu : d^{-1}(a) \cap \gamma C^* \to C$ is such that $g_a(\gamma, \mu) = G_a(\gamma)$.

Let s be the limit strategy of the sequence $(s_n)_{n\in\mathbb{N}}$ and first note that, using Observation 14, one can prove by induction on γ that $G_a(\gamma) \subseteq G_a$ for every $\gamma \in C^*$ that is compatible with s. Next, let $p \in P(s)$ be a path compatible with s. Due to the uniformity of the recursive definition, it suffices to show that $v(p) \in G_a$ to prove the full statement.

Let us make a case distinction. If Case 2. applies only finitely many times in the construction of s, the sequence $(s_n|_{\{\gamma\in C^* \mid p\in\gamma C^\omega\}})_{n\in\mathbb{N}}$ is eventually constant, so $p \in g_a(p_{<n}, s_n) = G_a(p_{<n}) \subseteq G_a$ for some n. Otherwise, there exists an increasing function $\varphi : \mathbb{N} \to \mathbb{N}$ such that $G_a(p_{\varphi(n+1)}) \subsetneq G_a(p_{\varphi(n)})$. Let us assume for a contradiction that $v(p) \notin G_a$ and define inductively a sequence $(q^n)_{n\in\mathbb{N}}$ such that $q^n_{<\varphi(n)} = p_{<\varphi(n)}$ and $v(q^n) \in G_a(p_{\varphi(n)}) \backslash G_a(p_{\varphi(n+1)})$ for all n, which implies[10] that $p \prec_a v(q^n) \prec_a v(q^{n+1})$ for all n, and subsequently that $v(p) \prec_a v(p)$, contradiction, so $g_a(s) = G_a$, and actually $g_a(\gamma, s) = G_a(\gamma)$ for all $\gamma \in C^*$, by uniformity of the recursive definition of s. □

Theorems 18 and 19 below both prove the existence of subgame perfect equilibria for "abstract-zero-sum" games, either when the choice set C is finite, or when the outcome set O is finite. Since their proofs are similar, most is factored out in Lemma 17 below.

Lemma 17. Let $\langle\{a, b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ be a two-player game. Let $\Gamma \subseteq \mathcal{P}(C^\omega)$ and assume the following.

1. $\prec$ is a strict weak order.
2. For every play $p \in C^\omega$ and every sequence of plays $(p^n)_{n\in\mathbb{N}}$ converging towards p, if $v(p^n) \prec v(p^{n+1})$ (resp. $v(p^n) \succ v(p^{n+1})$) for all n, then $v(p^n) \prec v(p)$ (resp. $v(p^n) \succ v(p)$ for all n.
3. For all $\gamma \in C^\omega$, there exists s such that $g_a(\gamma, s) = G_a(\gamma)$ (resp. $g_b(\gamma, s) = G_b(\gamma)$).
4. For all non-empty closed $E \subseteq C^\omega$, there are $\prec$-extremal elements in $v[E]$.
5. For every $\prec$-extremal interval I and $\gamma \in C^*$, we have $\left(v^{-1}[I] \cap \gamma C^\omega\right) \in \Gamma$.
6. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$, $D \subseteq C^*$.

Then the game $\langle\{a, b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ has a subgame perfect equilibrium.

Proof. By invoking Lemma 16 once for Player a and once for Player b, let us build a strategy profile $s : C^* \to C$, such that $g_X(\gamma, s_X) = G_X(\gamma)$ for all $\gamma \in C^*$ and $X \in \{a, b\}$. Let $\gamma \in C^*$ and let us prove that $G_a(\gamma) \cap G_b(\gamma) = \{\min_<(G_a(\gamma))\} = \{\max_<(G_b(\gamma))\}$. Consider the game $\langle C, D, W\rangle$ (as in Definition 5) where the winning set is defined by $W := \{\alpha \in \gamma C^\omega \mid v(\alpha) \in G_a(\gamma) \backslash \{\min_<(G_a(\gamma))\}\}$ and where Player a owns exactly the nodes in $D := (C^* \setminus \gamma C^*) \cup (d^{-1}(\{a\}) \cap \gamma C^*)$. By Assumption 5 the set W is in Γ, so by Assumption 6 the game $\langle C, D, W\rangle$ is determined. By definition of the best guarantee, Player a has no winning strategy for this game, so Player b has a winning strategy, which means that $G_b(\gamma) \subseteq \{\min_<(G_a(\gamma))\} \cup O\backslash G_a(\gamma)$. Since $G_a(\gamma) \cap G_b(\gamma)$ must be non-empty, otherwise the two guarantees are contradictory, $G_a(\gamma) \cap G_b(\gamma) = \{\min_<(G_a(\gamma))\}$. This means that the subprofile of s rooted at γ induces the outcome $\min_<(G_a(\gamma))$ (which equals $\max_<(G_b(\gamma))$ by symmetry), and it is optimal for both players. □

Theorem 18. Let $\langle\{a, b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ be a two-player "abstract-zero-sum" game, where C is finite. Let $\Gamma \subseteq \mathcal{P}(C^\omega)$ and assume the following.

1. $\prec$ is a strict weak order.
2. For every play $p \in C^\omega$ and every sequence of plays $(p^n)_{n\in\mathbb{N}}$ converging towards p, if $v(p^n) \prec v(p^{n+1})$ (resp. $v(p^n) \succ v(p^{n+1})$) for all n, then $v(p^n) \prec v(p)$ (resp. $v(p^n) \succ v(p)$ for all n.
3. For every $\prec$-extremal interval I and $\gamma \in C^*$, we have $\left(v^{-1}[I] \cap \gamma C^\omega\right) \in \Gamma$.
4. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$, $D \subseteq C^*$.

Then the game $\langle\{a, b\}, C, O, d, v, \{\prec, \prec^{-1}\}\rangle$ has a subgame perfect equilibrium.

Proof. By application of Lemma 17. Condition 3 is proved by Lemma 15. For Condition 4, let E be a non-empty closed subset of C^ω, and let T be the tree such that $[T] = E$. Consider the game where Player a plays alone on T. Since Player a can maximise her best guarantee by Lemma 15, and since all her guarantees are singletons, $v[E]$ has a $\prec$-maximum. Likewise, it has a $\prec$-minimum, by considering Player b. □

Theorem 19. Let $\langle\{a, b\}, C, d, O, v, \{<, <^{-1}\}\rangle$ be an infinite sequential game where O is finite and $<$ is a strict linear order. Let $\Gamma \subseteq \mathcal{P}(C^\omega)$ and assume the following.

1. $\forall O' \subseteq O, \forall\gamma \in C^*, \quad \{\alpha \in C^\omega \mid v(\gamma\alpha) \in O'\} \in \Gamma$
2. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$ and $D \subseteq C^*$.

Then the game $\langle\{a, b\}, C, d, O, v, \{<, <^{-1}\}\rangle$ has a subgame perfect equililbrium.

Proof. by Lemma 17 where Conditions 2, 3, and 4 hold by finiteness of O. □

Corollary 20. Let $\langle\{a, b\}, C, d, O, v, \{<, <^{-1}\}\rangle$ be an infinite sequential game where O is finite and $<$ is a strict linear order. If $v^{-1}(o)$ is quasi-Borel for all $o \in O$, the game has a subgame perfect equilibrium.

Proof. From Theorem 19, quasi-Borel determinacy [18], and Lemma 3.1. in [15]. □

Corollary 21. A sequential two-player zero-sum game with Borel measurable payoffs has ε-subgame perfect equilibria.

Proof. By Corollary 20 and the argument in Section 4. □

Theorem 22. Let $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ be a game, $\Gamma \subseteq \mathcal{P}(C^\omega)$ and assume the following.

1. The $\prec_a$ are strict weak orders.
2. C is finite.
3. For every play $p \in C^\omega$, every sequence of plays $(p^n)_{n\in\mathbb{N}}$ converging towards p, and every player $a \in A$, if $v(p^n) \prec_a v(p^{n+1})$ for all n, then $v(p^n) \prec_a v(p)$ for all n.
4. For every player $a \in A$ and $\prec_a$-terminal interval I and $\gamma \in C^*$, we have $\left(v^{-1}[I] \cap \gamma C^\omega\right) \in \Gamma$.

[10] Here we are using again the properties of a strict weak order.

5. The game $\langle C, D, W\rangle$ is determined for all $W \in \Gamma$, $D \subseteq C^*$.

Then the game $\langle A, C, O, d, v, (\prec_a)_{a\in A}\rangle$ has a Nash equilibrium.

Proof. Since the proof is similar to that of Theorem 2.9 in [15], we rephrase and give it a more intuitive flavour. Let σ be a strategy profile where every player is using a witness to Lemma 16 (which is applicable by Lemma 15). Let p be the induced play. We now turn σ into a Nash equilibrium with p as induced play by use of threats. More specifically, at each node $p_{<n}$ we let the players other than $a := d(p_{<n})$ threaten Player a that if she deviates from p exactly at $p_{<n}$, they will team up against her at every subsequent position γ after $p_{<n}$ other than those extending the prescribed $p_{<n+1}$.

We claim that if they team up, they can prevent Player a from getting better outcome than $v(p)$ by deviating to γ, which will suffice. Let us build a win-lose game $\langle C, D, W\rangle$, with Player a against her threatening opponents gathered as a meta-player, and where the winning set for Player a is defined by $W = v^{-1}[I] \cap \gamma C^\omega$, where $I := \{o \in O \mid v(p) \prec_a o\}$, and D is defined by $D = d^{-1}(\{a\}) \cup (C^* \setminus \gamma C^*)$. This game is determined by Assumptions 4 and 5, and Player a looses it, otherwise her winning strategy would guarantee that $v(p) \notin G_a(p_{<n})$ and thus contradict the choice of p. Therefore the threat of the opponents of Player a is effective. □

For comparison, the preparatory work before [15, Theorem 2.9] considers strict well-orders only; then [15, Theorem 2.9] considers strict well-founded orders, since linear extensions of these make it possible to invoke the special, linear case, knowing that any Nash equilibrium for these extensions is still a Nash equilibrium for the original preferences. However, let us explain why Theorem 22 assumes that preferences are strict weak orders, instead of more general strict partial orders. In the preparatory work before both results, the algorithm that builds a play step by step makes decisions based on the guarantees that the subgames offer. If the guarantees of one player were not ordered by a strict weak order, the player might eventually regret a previous decision, in the same way that backward induction on partially ordered preferences may not yield a Nash equilibrium (see *e.g.*, [11] for a concrete example or page 3 of [13] for a generic one). So the algorithm has to run on strict weak orders. (In [15, Theorem 2.9] it even runs on strict linear orders.)

If we wanted to consider strict partial orders and extend them linearly for the algorithm to work, we would potentially run into two problems: First, there may not exist any linear extension preserving Condition 3. Second, assumptions 4 and 5 of Theorem 22 make sure that the win-lose games associated with the $\prec_a$-terminal intervals are determined, which is a requirement for the proof to work. If the preferences were not strict weak orders, we might think of replacing the condition on terminal intervals by a condition on the upper-closed sets and then extend the preferences linearly for the algorithm to work, but in the special case where the preference of one player were the empty relation, every subset would be an upper-closed set and its preimage by v would be in the pointclass with nice closure property, by assumption. If, in addition, each outcome is mapped to at most one play, it implies that each subset of C^ω is in the pointclass, so Theorem 22 could be used with the axiom of determinacy only, but not with, *e.g.*, Borel determinacy. On the contrary, [15, Theorem 2.9, Assumption 3] is not an issue since there are only countably many outcomes in that setting.

Theorem 22 has a corollary pertaining to sequential games with real-valued payoffs. Rather than the usual Euclidean topology, we consider the lower topology generated by $\{(-\infty, a) \mid a \in \mathbb{Q}\}$. This space will be denoted by $\mathbb{R}_>$. Note that continuous functions with codomain $\mathbb{R}_>$ are often called upper semi-continuous. As $\mathrm{id} : \mathbb{R}_> \to \mathbb{R}$ is complete for the Σ^0_2-measurable functions [30, 32], we see that the Borel sets[11] on $\mathbb{R}_>$ are the same as the Borel sets on $\mathbb{R}$. Moreover, if $(p^n)_{n\in\mathbb{N}}$ is a converging sequence of plays, and $P : \{0,1\}^{\mathbb{N}} \to \mathbb{R}_>$ is a continuous payoff function, then $P(\lim_{i\in\mathbb{N}} p^i) \geq \limsup_{i\in\mathbb{N}} P(p^i)$. In particular, Condition 3 in Theorem 22 is always satisfied for the preferences obtained from upper semi-continuous payoff functions.

Corollary 23. Sequential games with countably many players, finitely many choices and upper semi-continuous payoff functions have Nash equilibria.

6. Lazy improvement

The *lazy improvement* was introduced in [14] for finite sequential games. Intuitively, it consists in repeating the game and letting, at each repetition, an arbitrary player improve upon the current outcome by first identifying a reachable and improving leaf of the game tree, and second by changing her strategy only as much as necessary to reach this leaf. In particular a player never changes her choices that are rendered irrelevant by prior choices of other players. For instance below, consider the four strategy profiles on the same game with players a and b, where strategical choices are represented by double lines and where only the payoff of Player a is displayed at the two relevant leaves. The top-left profile can be lazily improved into the top-right one, but not into the bottom-left or bottom-right ones.

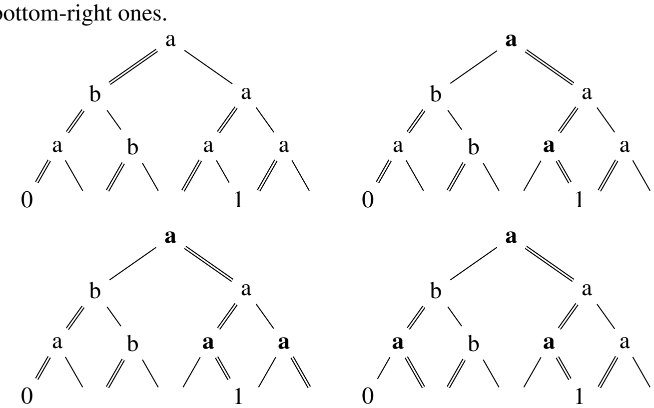

For finite sequential games, lazy improvement terminates quickly and settles at Nash equilibria, as proved in [14]. The idea of lazy improvement can be extended to infinite sequential games with continuous payoff functions. Since we consider continuous payoff functions, for each improvement step there is a similar improvement step with finitely many node changes, so we rule out improvement steps with changes at infinitely many nodes. In the infinite case, the dynamics are somewhat less well-behaved than one could wish for; but still provide a characterization of Nash equilibria as the accumulation points of certain sequences.

As a slight oddity, we find that the lazy improvement dynamics are not fully determined by the game, but there is a degree of consistent choice needed to ensure the intended behaviour. We can use the continuous payoff function to label any game vertex with some rational interval in a way[12] that the label of any vertex is a subset of its predecessor, and that the intersection of all labels along an infinite path is the singleton set containing the payoff for this path. Now lazy improvement decisions are made based only by

[11] As $\mathbb{R}_>$ is not metric (but still countably based), the definition of the Borel hierarchy has to be modified as demonstrated by SELIVANOV [27]. A move towards definitions of Borel measurability on even more general spaces can be found in [25].

[12] The idea behind this corresponds to the representation of real numbers in computable analysis [32].

inspecting the strategy profile and game tree (including labels) up to some finite depth d, using the definitions from [14].

Once all players are stable at the current inspection depth, the inspection depth is incremented by one. The incrementing shall be counted as an updating step, thus always some infinite sequence of strategy profile arises. We shall call the subsequence corresponding to steps where the inspection depth is incremented the *stable subsequence*.

Note that the choice of labeling system is not uniquely determined by the payoff function, and that the labeling in turn influences the lazy improvement dynamics. Moreover, note that while we are dealing with linear preferences only in the case of infinite games, we do make use of finite approximations that lack linear preferences – yet we are guaranteed that any preferences occurring in our finite approximations are acyclic, which is sufficient for [14]. Finally, as the investigation depth is never reset, the dynamics do depend on the history – however, only on its length, not on any details.

Observation 24. The lazy improvement dynamics are computable, i.e. given a game and an initial strategy profile, we can compute a sequence of strategy profiles arising from lazy improvement, as well as the indices of the stable subsequence.

Theorem 25. The following properties are equivalent for a strategy profile s:

1. s is a Nash equilibrium.
2. s is a fixed point[13] for lazy improvement.
3. s is an accumulation point of the stable subsequence of a sequence obtained from iterated lazy improvement.

Proof. 1. $\Leftrightarrow$ 2. By continuity of the preferences, a player prefers a strategy profile s to another profile s', if and only if there is an inspection depth d such that he prefers the restriction of s to the restriction of s' in the corresponding finite approximation. This in turn implies that a strategy profile is a Nash equilibrium of the infinite game, if and only if all its finite prefixes are Nash equilibria in the corresponding finite games. The same holds for fixed points by construction of the lazy improvement steps for infinite games. Thus, the claim for infinite games follows from the result for finite games, i.e. [14, Proposition 10].

2. $\Rightarrow$ 3. If s is a fixed point, then the lazy improvement sequence with starting point s is constant, hence has s as accumulation point.

3. $\Rightarrow$ 2. Let the strategy profile s arise as an accumulation point of the stable subsequence of a sequence $(s_n)_{n\in\mathbb{N}}$ obtained by iterated lazy improvement, and assume that s is not a fixed point. Then there is some minimal inspection depth d necessary to find a lazy improvement step in s, which is executed by some player p. The detection at inspection depth d means that any strategy profile s' sharing a finite prefix of depth d with s will admit exactly the same lazy improvement step.
The assumption that s is an accumulation point of the stable subsequence in particular implies that infinitely many strategy profiles occur that share a prefix of length d with s. In particular, there would have to be a strategy profile that shares a prefix of length d with s, and that is stable at inspection depth $d' > d$. But, as explained above, player p would then wish to change his strategy, i.e. we have arrived at a contradiction. Hence, s has to be a fixed point.

$\square$

[13] Given the history dependence of the lazy improvement dynamics, a *fixed point* is understood to be any starting point of a lazy improvement sequence resulting in a constant sequence, i.e. no improvement step is found at any inspection depth.

Next, we shall provide an example exhibiting the necessity of the various restrictions of the preceding theorem.

Example 26. Let us consider games with four players a, b, c, and d. Given four real-valued sequences $\mathcal{A} = (\alpha_n)_{n\in\mathbb{N}}, \mathcal{B} = (\beta_n)_{n\in\mathbb{N}}$, $\mathcal{C} = (\gamma_n)_{n\in\mathbb{N}}$, and $\mathcal{D} = (\delta_n)_{n\in\mathbb{N}}$ converging towards α, β, γ, and δ, let $T(\mathcal{A},\mathcal{B},\mathcal{C},\mathcal{D})$ be the following game and strategy profile. Note that apart from the payoffs, the underlying game effectively involves Players c and d only. If $\mathcal{C}$ and $\mathcal{D}$ are decreasing, the lazy improvement dynamics see players c and d alternating in switching their top left-move to a right-move.

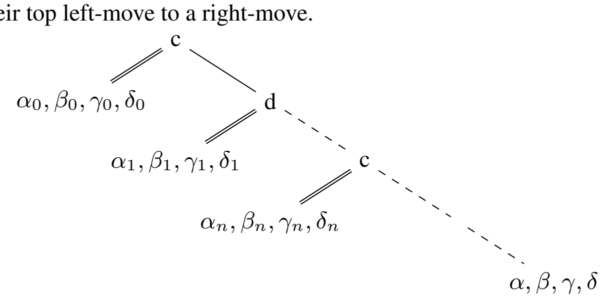

Let $\mathcal{A} := \mathcal{B} := (1 + \frac{1}{n+1})_{n\in\mathbb{N}}$ and let $\mathcal{C} := \mathcal{D} := (1 - \frac{1}{n+1})_{n\in\mathbb{N}}$. Starting from the profile below, players c and d will continue to unravel the subgame currently chosen jointly by a and b. Player b will keep alternating her choices to pick the least-unraveled subgame available to her. Player a will prefer to chose a subgame where player b currently chooses right, and also prefers less-unraveled subgames.

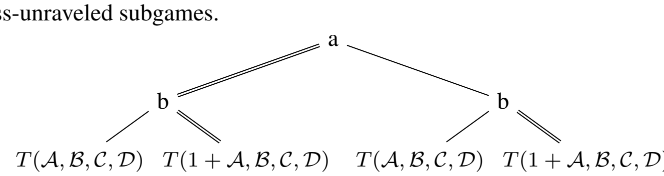

First of all, already the subgame where b moves first demonstrates that the lazy improvement dynamics will not always converge, hence we have to consider accumulation points rather than limit points. For the next feature, note that there is an infinite sequence of lazy improvement where Players a and b (at both node that she owns) switch infinitely often, and where Player a switches only when Player b chooses the right subgame (on the induced play). Then the following strategy profile is an accumulation point (however, not an accumulation point of the stable subsequence). It is clearly not a Nash equilibrium, thus justifying the restriction to accumulation points of the stable subsequence in Theorem 25.

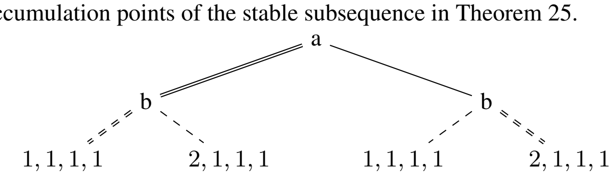

7. Absence of subgame perfect equilibria

In this section we will show that in the simultaneous absence of continuity and the zero-sum property, even a two-player game with three distinct outcomes may fail to have subgame perfect equilibria. It is a straightforward consequence that moving to ε-subgame perfect equilibria cannot help, either. As our (counter-) Example 11, the valuation function here is Δ^0_2-measurable, hence, in a sense, not very discontinuous. A similar counterexample is also exhibited in [29, Example 3].

Example 27. The following game where $z <_a y <_a x$ and $x <_b z <_b y$ has no subgame perfect equilibrium.

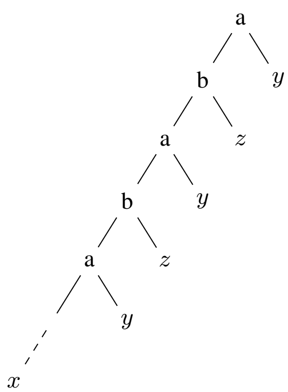

The game above is formally defined as $\langle\{a, b\}, \{0, 1\}, d, O, v, \{<_a, <_b\}\rangle$, where $d^{-1}(a) := 0^{2*}$ and $v(0^\omega) := x$ and $v[0^{2*}1\{0, 1\}^\omega] := \{y\}$ and $v[0^{2*+1}1\{0, 1\}^\omega] := \{z\}$.

Proof. Assume for a contradiction that there is a subgame perfect equilibrium for this game. Then no subprofile (starting at some node in 0^*) induces the outcome x, because Player b could then switch to the right and obtain z. So for infinitely many nodes in 0^*, Players a or b chooses 1. Also, if Player b chooses 1 at some node 0^{2n+1}, Player a chooses 1 at the node 0^{2n} right above it. This implies that every subprofile rooted at nodes in 0^{2*} induces the outcome y, and subsequently, Player b always chooses 0 at nodes in 0^{2*+1}. But then Player a could always choose 0 and obtain x, contradiction. □

A further example shows us that we can rule out subgame perfect equilibria even with stronger conditions on the functions by using countably many distinct payoffs. This example no longer extends to ε-subgame perfect equilibria.

Example 28. The following game where $y_n := (2^{-n}, 2^{-n})$ and $z_n := (0, 2^{-n-2})$ for all $n \in \mathbb{N}$ has no subgame perfect equilibrium, although the payoff functions are upper-semicontinuous.

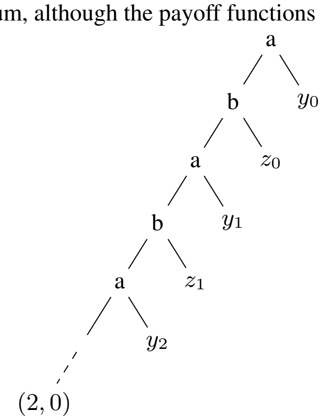

Proof. If the payoffs are $(2, 0)$, Player b can improve her payoff as late as required, so there are infinitely many "right" choices in a subgame perfect equilibrium. If the payoffs are not $(2, 0)$ then Player a chooses "right" at every node that she owns, so that Player b chooses "left", but then Player a chooses "left" too. □

8. Outlook

There is one open question regarding infinite sequential games with real-valued payoff functions (see the question mark in the Overview table on page 3), namely:

Open Question 29. Do games with countably many players and upper-semicontinuous payoff functions have ε-subgame perfect equilibria?

It seems surprising to have both Nash equilibria and ε-subgame perfect equilibria, but no subgame perfect equilibria – but this is precisely the situation for finitely many players. On the other hand, given the split between finitely many players and countably many players for lower-semicontinuous payoff functions, one should be cautious about assuming that this result would extend. Thus, we do not present a conjecture regarding the answer to the open question.

A more abstract question is whether there is a connection between existence results for infinite games and complexity / decidability results for finite games with related winning conditions[14]. In an intermediate step, one could investigate the connections to finitary infinite[15] games such as generalized Muller games studied in [22].

Finally, the condition on the payoff functions used in Section 5 seems to merit further investigation. This was that for any sequence $(p^i)_{i\in\mathbb{N}}$ converging to p in C^ω, we find that $\forall i \in \mathbb{N}\ v(p^i) \prec v(p^{i+1})$ implies $\forall i \in \mathbb{N}\ v(p^i) \prec v(p)$. This is a weaker condition than continuity of the function where the upper order topology is used on the codomain, which still seems to be strong enough to formulate some results. In a sense, it is a weakening of continuity that is orthogonal to Borel-measurability.

References

[1] D. Blackwell. Infinite G_δ games with imperfect information. *Zastowania Matematyki Applicationes Mathematicae*, 1969.

[2] R. Bloem, K. Chatterjee, T. A. Henzinger, and B. Jobstmann. Better quality in synthesis through quantitative objectives. In A. Bouajjani and O. Maler, editors, *Computer Aided Verification*, volume 5643 of *Lecture Notes in Computer Science*, pages 140–156. Springer Berlin Heidelberg, 2009. .

[3] V. Brattka, G. Gherardi, and A. Marcone. The Bolzano-Weierstrass Theorem is the jump of Weak König's Lemma. *Annals of Pure and Applied Logic*, 163(6):623–625, 2012. . also arXiv:1101.0792.

[4] K. Chatterjee and L. Doyen. Energy parity games. *Theor. Comput. Sci.*, 458:49–60, 2012. .

[5] K. Chatterjee, T. A. Henzinger, and M. Jurdziński. Mean-payoff parity games. In *Proceedings, Twentieth Annual IEEE Symposium on Logic in Computer Science, LICS 2005*, pages 178–187. IEEE Computer Society Press, 2005.

[6] M. Escardó. Synthetic topology of datatypes and classical spaces. *Electronic Notes in Theoretical Computer Science*, 87, 2004.

[7] J. Flesch, J. Kuipers, A. Mashiah-Yaakovi, G. Schoenmakers, E. Solan, and K. Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operations Research*, 35(4): 742–755, 2010.

[8] D. Fudenberg and D. Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31(2):251 – 268, 1983. .

[9] D. Gale and F. Stewart. Infinite games with perfect information. In *Contributions to the theory of games*, volume 28 of *Annals of Mathematical Studies*, pages 245–266. Princeton University Press, 1953.

[10] E. Grädel and M. Ummels. Solution concepts and algorithms for infinite multiplayer game. In K. Apt and R. van Rooij, editors, *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 151–178. Amsterdam University Press, 2008.

[11] T. Krieger. On Pareto equilibria in vector-valued extensive form games. *Mathematical Methods of Operations Research*, 58:449–458, 2003.

[14] This question was raised by one of the referees.

[15] The game duration is infinite, but the game is described in a finite way.

[12] S. Le Roux. *Generalisation and formalisation in game theory*. Ph.D. thesis, Ecole Normale Supérieure de Lyon, January 2008.

[13] S. Le Roux. Acyclic preferences and existence of sequential Nash equilibria: a formal and constructive equivalence. In *TPHOLs, International Conference on Theorem Proving in Higher Order Logics*, Lecture Notes in Computer Science, pages 293–309. Springer, August 2009.

[14] S. Le Roux. Lazy beats crazy: a spurious yet mathematical justification for laissez-faire. arXiv:1309.2798, 2013.

[15] S. Le Roux. Infinite sequential Nash equilibria. *Logical Methods in Computer Science*, 9(2), 2013.

[16] S. Le Roux. From winning strategy to Nash equilibrium. *Mathematical Logic Quarterly*, 2014. (to appear, cf. arXiv 1203.1866).

[17] D. A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):pp. 363–371, 1975. URL `http://www.jstor.org/stable/1971035`.

[18] D. A. Martin. The determinacy of blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998. .

[19] J. F. Mertens. Repeated games. In *Proc. Internat. Congress Mathematicians*, pages 1528–1577. American Mathematical Society, 1987.

[20] J. F. Nash. Equilibrium Points in N-Person Games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950. URL `http://www.pnas.org/content/36/1/48.short`.

[21] R. Nessah and K. Kerstens. Characterizations of the existence of Nash equilibria with non-convex strategy sets. Working Papers 2008-ECO-13, IESEG School of Management, 2008. URL `http://ideas.repec.org/p/ies/wpaper/e200813.html`.

[22] S. Paul and S. Simon. Nash Equilibrium in Generalised Muller Games. In R. Kannan and K. N. Kumar, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 4 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 335–346, Dagstuhl, Germany, 2009. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. . URL `http://drops.dagstuhl.de/opus/volltexte/2009/2330`.

[23] A. Pauly. How incomputable is finding Nash equilibria? *Journal of Universal Computer Science*, 16(18):2686–2710, 2010. .

[24] A. Pauly. A new introduction to the theory of represented spaces. http://arxiv.org/abs/1204.3763, 2012.

[25] A. Pauly and M. de Brecht. Towards synthetic descriptive set theory: An instantiation with represented spaces. http://arxiv.org/abs/1307.1850, 2013.

[26] R. Purves and W. Sudderth. Perfect information games with upper semicontinuous payoffs. *Mathematics of Operations Research*, 36(3): 468–473, 2011.

[27] V. L. Selivanov. Difference hierarchy in φ-spaces. *Algebra and Logic*, 43(4):238–248, 2004.

[28] R. Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Z. Gesamte Staatswissenschaft*, 12:301–324, 1965.

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

[30] T. v. Stein. Vergleich nicht konstruktiv lösbarer probleme in der analysis. Master's thesis, Fachbereich Informatik, FernUniversität Hagen, 1989.

[31] M. Ummels. *Rational behaviour and strategy construction in infinite multiplayer games*. Diploma thesis, RWTH Aachen, 2005.

[32] K. Weihrauch. *Computable Analysis*. Springer-Verlag, 2000.

Acknowledgements

We are very grateful to Vassilios Gregoriades and the anonymous referees for their helpful comments. This work benefited from the Royal Society International Exchange Grant IE111233. The authors participate in Marie Curie International Research Staff Exchange Scheme *Computable Analysis*, PIRSES-GA-2011- 294962.