Determinacy for infinite games with more than two players with preferences

Benedikt Löwe

Institute for Logic, Language and Computation, Universiteit van Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands,
`bloewe@science.uva.nl`

Abstract

We discuss infinite zero-sum perfect-information games with more than two players. They are not determined in the traditional sense, but as soon as you fix a preference function for the players and assume common knowledge of rationality and this preference function among the players, you get determinacy for open and closed payoff sets.

2000 AMS Mathematics Subject Classification. **91A06 91A10 03B99 03E99**.

Keywords. Perfect Information Games, Coalitions, Preferences, Determinacy Axioms, Common Knowledge, Rationality, Gale-Stewart Theorem, Labellings

1 Introduction

The theory of infinite two-player games is connected to the foundations of mathematics. Statements like

> "Every infinite two-player game with payoff set in the complexity class Γ is determined, *i.e.*, one of the two players has a winning strategy"

[1] The author was partially financed by NWO Grant B 62-584 in the project *Verzamelingstheoretische Modelvorming voor Oneindige Spelen van Imperfecte Informatie* and by DFG Travel Grant KON 192/2003 LO 834/4-1. He would like to extend his gratitude towards Johan van Benthem (Amsterdam), Boudewijn de Bruin (Amsterdam), Balder ten Cate (Amsterdam), Derrick DuBose (Las Vegas NV), Philipp Rohde (Aachen), and Merlijn Sevenster (Amsterdam) for discussions on many-player games.

are called **determinacy axioms (for two-player games)**. When we supplement the standard (Zermelo-Fraenkel) axiom system for mathematics with determinacy axioms, we get a surprisingly fine calibration of very interesting logical systems. Research in set theory from the 1960es to this day has shown that many if not most of the interesting features of Higher Set Theory are connected to determinacy axioms for two-player games: they have connections to infinitary combinatorics of large sets, to topology, to the theory of the real numbers, and to many other areas of interest in foundations of mathematics. In general, determinacy axioms provide a rich theory of interesting features. (*Cf.* [Ka94, §§ 27–32].)

A natural but naïve approach would be to assume that if you increase the number of players and look at "determinacy axioms for three-player games", the number of interesting features for set theory should also increase. It is well-known that the opposite is true: determinacy axioms of the above form for n-player games are only interesting if $n = 2$. In all other cases, they are trivial: If $n = 1$ they are all true, regardless of Γ, if $n > 2$ they are all false for almost all Γ (*cf.* Proposition 1). It turns out that if we want to give solution concepts for infinite many-player games, determinacy in the classical sense is not the right concept. Solutions that have been offered in the literature include giving up the notion of a pure strategy and moving to mixed strategies (*cf.* [Ga53] and [Br00]), and understanding many-player games as coalitional games. In this paper, we want to work with pure strategies and stay within the realm of non-cooperative perfect information games

In addition to common knowledge of rationality of all players involved, we need to add fixed preferences of the players and common knowledge of these preferences. Then we are able to give solution concepts for many-player infinite games (for payoffs of very restricted complexity, though).

The results of this paper are (for definitions, *cf.* Sections 2 and 3):

Let I be an arbitrary set of players, X an arbitrary set of moves, μ an arbitrary moving function, Π an arbitrary total preference.

- If P is an open or closed payoff, then there is a rational labelling ℓ^* such that one of the players has an $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy. (Theorems 8 and 9)
- If $\langle P, \Pi \rangle$ is a exceptional least evil situation, then there is a rational labelling ℓ^* such that one of the players has an $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy. (Theorem 13)
- If P is a payoff such that a finite set of players has a closed payoff set and the rest of the players have an open payoff set, then there is a rational labelling ℓ^* such that one of the players has an $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy. (Theorem 19) [1]

[1] Theorem 13 is a special case of Theorem 19, but uses a different technique, and is thus interesting in its own right.

2 Definitions & Motivation

In this paper, we will look at a very general form of infinite zero-sum non-cooperative perfect information games. We have an arbitrary set I of players for this game. Our games will be games of length ω with a set X of possible moves.[2] We denote the set of finite sequences of elements of X by $X^{<\omega}$. The length of a finite sequence $p \in X^{<\omega}$ will be denoted by $\mathrm{lh}(p)$, and we write $\langle x_0 x_1 x_2 \ldots x_n \rangle$ for a finite sequence of length $n+1$. Note that $X^{<\omega}$ is a tree ordered by inclusion. Runs of the game are infinite sequences of elements from X, and we denote that set by X^ω. If $x = \langle x_k \,;\, k \in \omega \rangle \in X^\omega$, we denote by $x \restriction m := \langle x_0 x_1 \ldots x_{m-1} \rangle$ its **restriction to** m.

We have a function $\mu : X^{<\omega} \to I$ called the **moving function** determining which player has to move in some position $p \in X^{<\omega}$. The set $M_i := \{p \in X^{<\omega} \,;\, \mu(p) = i\}$ is the set of positions where player i has to move.

The branches through $X^{<\omega}$, *i.e.*, the elements of X^ω will be partitioned by a payoff function

$$P : X^\omega \to I$$

into the payoff sets for the players. For each $i \in I$, we interpret $P_i := \{x \,;\, p(x) = i\}$ to be the set of all those ω-sequences of moves that result in a win for player i.

Since we are dealing with perfect information games, all functions $\sigma_i := M_i \to X$ will be called i**-strategies** (that means a player has access to the entire game played so far when making his decision about what to play in a position p). We call a sequence of strategies $\vec{\sigma} = \langle \sigma_i \,;\, i \in I \rangle$ a μ**-frame** (or just **frame** if μ is implicitly determined) if for each $i \in I$, σ_i is an i-strategy. Clearly, a μ-frame determines a unique element of X^ω by letting the strategies in $\vec{\sigma}$ play against each other, and we call it ✠$\vec{\sigma}$.

Our choice of payoff functions of the type $P : X^\omega \to I$ affects our language: if a μ-frame $\vec{\sigma}$ is a Nash equilibrium in the usual sense, and $P(✠\vec{\sigma}) = i$, this means that while all players locally optimize their outcome, for all players j with $j \neq i$ this optimal outcome is still a loss. Consequently, the game-theoretically central notion of a Nash equilibrium becomes a bit stale, and we choose the notion of a **winning strategy** over the notion of an **equilibrium** to be central for this paper. We shall briefly discuss the translation of our determinacy theorems into the language of equilibria after Theorem 8.

[2] We are basically working in set theory without the Axiom of Choice, so we would like X to be well-orderable (the standard example would be either a finite set X or $X = \omega$), or at least $\mathsf{AC}_X(X)$, *i.e.*, the existence of choice functions for X-indexed families of subsets of X. The choice of the exact set theory is not of any importance for this paper.

A proper analysis of a class $\mathcal{C}$ of games would be something like the following theorem scheme:

> *For each payoff $P \in \mathcal{C}$ there is some player $i \in I$ and an i-strategy σ such that for all **relevant** μ-frames $\vec{\tau}$ with $\tau_i = \sigma$, we have $P(\maltese\vec{\tau}) = i$.*

The meaning of "relevant" has to be determined by the game analysis. In the classical two-player case of set theory mentioned above, we have the degenerate case where all strategies are relevant. This yields **winning strategies**: σ is **winning** in the game with payoff P if for every μ-frame $\vec{\tau}$ with $\tau_i = \sigma$, we have $P(\maltese\vec{\tau}) = i$. As usual, we call a payoff P **determined** if there is a winning i-strategy for some $i \in I$. For games with more than two players, winning strategies might not always exist (Proposition 1).

For our classes Γ of payoffs we use the usual complexity classes of descriptive set theory: X^ω is endowed with a natural topology (the product topology of the discrete topology on X), and if X is finite or countable, this is a Polish space. We have the usual pointclasses of open, closed, Borel, projective sets on X^ω, and for a pointclass Γ, we say that the payoff $P : X^\omega \to I$ is a Γ-payoff if for all i, $P_i \in \Gamma$.[3]

Proposition 1 *There is an open three-player game that is not determined.*

Proof. Let $X := \{0,1\}$, $I = 3 = \{0,1,2\}$, $\mu(p) := \mathrm{lh}(p) \mod 3$ and

$$P_0 := \{x\,;\, \langle 010\rangle \subseteq x \text{ or } \langle 011\rangle \subseteq x\},$$

$$P_1 := \{x\,;\, \langle 101\rangle \subseteq x \text{ or } \langle 100\rangle \subseteq x \text{ or } \langle 110\rangle \subseteq x \text{ or } \langle 111\rangle \subseteq x\}, \text{ and}$$
$$P_2 := \{x\,;\, \langle 000\rangle \subseteq x \text{ or } \langle 001\rangle \subseteq x\},$$

as depicted in Figure 1 or (simplified) in Figure 2.

Clearly, the payoff sets are all open. Figure 2 shows that none of the three players can have a winning strategy. q.e.d.

Consequently, for games with more than two players, we have to specify a smaller class of strategies as "relevant" in the sense of the above theorem scheme. This is well-known in game theory and led to studying many-player games in terms of coalitions.

As mentioned, we shall stay within the non-cooperative paradigm, but add information about the preferences of the players and their rationality to the

[3] Note that this usage is slightly different from the usage in the two-player case. A two-player game is normally called closed if the payoff for player I is a closed subset of X^ω (with no conditions on the payoff for player II).

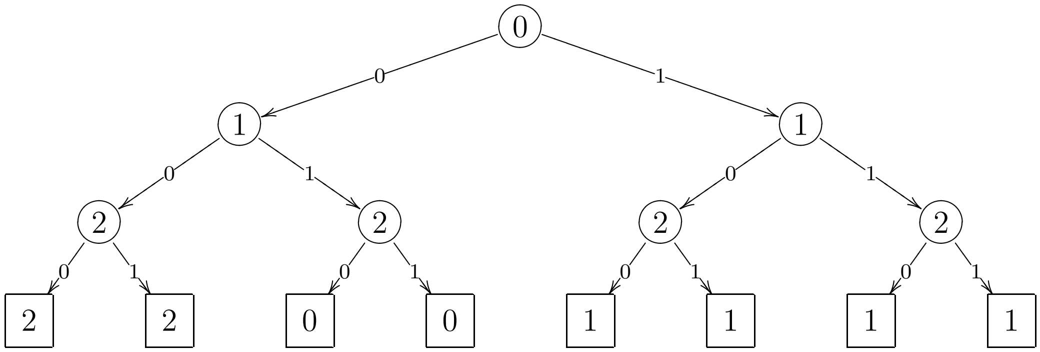

Fig. 1. The game graph of the game from Proposition 1. Round nodes with an i represent opportunities for player i to move, square nodes denote the winner of a given game path.

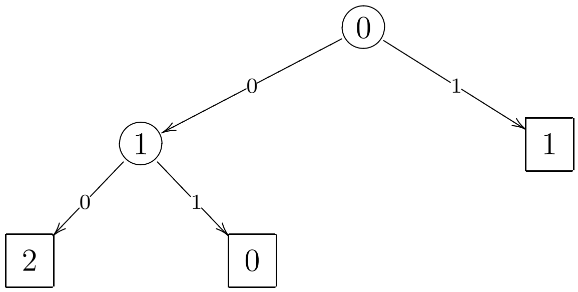

Fig. 2. Simplified version of Figure 1 with all irrelevant moves removed.

analysis. We shall presuppose full rationality of all players and common knowledge of that fact. This allows us to use the usual **backwards induction techniques** known from perfect information game analysis. [4] Assuming common knowledge of rationality, we can restrict the set of relevant strategies; this will be done when we define the notion of a $\langle \Pi, \ell \rangle$-strategy in Section 3.

As an example, consider the game defined in Figure 1 and imagine that we have the three players Jill (0), her husband Jeff (1), and an invited friend John (2). If Jill knows for sure that Jeff prefers her over John, she can exclude his strategy "play 0" from her considerations as irrelevant, and win by playing 0 herself.

Note that this is not a coalition: there is no agreement between Jill and Jeff, there isn't even a benefit for Jeff. It is just calculation of Jill, taking into account predictions about the behaviour of her husband. Also note that playing 0 can only guarantee Jill to win if she can be sure of both Jeff's preference for

[4] *Cf.* [Au95], [Bi88], [St96]. For a discussion of the rôle of common knowledge of rationality for the backwards induction technique in terms of modal logic, *cf.* [dB∞]. Note that the usual Gale-Stewart analysis for two players doesn't need any rationality assumptions: the winning strategy constructed in the proof of the Gale-Stewart Theorem will win against all players, including irrational ones. This is not true for our games anymore; we briefly discuss this at the end of this section.

her and Jeff's rationality. If Jeff prefers Jill over John, but doesn't properly understand what's going on in the game, he might play 0 and thus let John win. (This type of the game will show up later in Section 7 as Evening at a Couple.)

3 Games with a preference

We call a function Π a **partial preference** if its domain is the set I, and $\Pi(i)$ is a well-founded partial order on the set I — we also write $<^i_\Pi$ for this partial order; the intended interpretation of $j_0 <^i_\Pi j_1$ is: "player i prefers a win of j_0 over a win of j_1". We call a partial preference Π a **(total) preference** if for each $i \in I$, $\Pi(i)$ is a wellordering. In the following, we shall always assume that players want to win, *i.e.*, that

$$\min_{<^i_\Pi} I = i$$

for all $i \in I$.

Let us define Π_0 as follows:

$$i_0 <^i_{\Pi_0} i_1 \iff i_0 = i \;\&\; i_1 \neq i.$$

Then Π_0 is a partial preference that corresponds to the classical zero-sum situation: all players prefer to win, but if they don't, they have no preferences about the actual winner. Note that Proposition 1 shows that there is no solution for games with the partial preference Π_0.

In the following, we will offer a backwards induction solution concept for infinite many-player games based on a (total) preference that is commonly known. As in the usual situation for infinite perfect information games, we assume that all players are fully rational, and their rationality is common knowledge.

Let I be an arbitrary set of players, μ a move function, and Π a total preference. We call any partial function $\ell : X^{<\omega} \to I$ an **(partial) labelling**.

For each partial labelling ℓ and a total preference Π, we can define its **Gale-Stewart procedure** in a transfinite recursion as follows:

Start of the recursion and the successor step. We let $\mathrm{GS}^\Pi_0(\ell) := \ell$. Suppose that $\mathrm{GS}^\Pi_\alpha(\ell)$ is already defined, we then define $\mathrm{GS}^\Pi_{\alpha+1}(\ell)$ as follows: for each $s \in X^{<\omega}$, we check whether one of the cases $(s+)$ or $(s-)$ holds. If so, we follow the instructions described below.

- ($s+$) If $\mu(s) = i$ and there is an immediate successor t of s such that

$$\mathrm{GS}^{\Pi}_{\alpha}(\ell)(t) = i,$$

  then we let $\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s) := i$.
- ($s-$) If $\mu(s) = i$ and all immediate successors of s are already labelled, let

$$\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s) := \min_{<^i_\Pi}\{j \in I\,;\, \exists\xi \in X\,(\mathrm{GS}^{\Pi}_{\alpha}(\ell)(s^\frown\langle\xi\rangle) = j)\}.$$

If none of the conditions ($s+$) or ($s-$) hold, we let $\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s) := \mathrm{GS}^{\Pi}_{\alpha}(\ell)(s)$ (which might be undefined).

In words: We interpret a label $\ell(s) = j$ as "at s it is determined that player j will win". If player i has to play at s, the case ($s+$) means that there is an immediate successor node t of s such that "at t it is determined that player i will win". Assuming rationality of player i, and assuming that player i wants to win (note that we demanded that of our preferences), this means that "at s it is determined that player i will win", since player i will play into such a successor node. The case ($s-$) means that in all of the successor nodes, the game is determined. In that case, player i can look at the possible labels of the successor nodes, pick the one that he likes best (according to his preference relation $<^i_\Pi$), and play a successor with such a label. Again, assuming the rationality of player i and knowledge of the preference Π, the outcome of the game is determined at s.

Note that in general, the Gale-Stewart procedure is non-monotonic: both ($s+$) and ($s-$) are able to change labels of $\mathrm{GS}^{\Pi}_{\alpha}(\ell)$, and so there needn't be a fixed point. The key idea of Gale and Stewart was that for nice labellings ℓ, the Gale-Stewart procedure is monotonic, and we have a fixed point. In Section 5, we shall see an example where we can do a game analysis even if the procedure isn't fully monotonic.

The limit step of the recursion. The possibility of non-monotonicity also causes potential trouble with the limit case. In the spirit of Herzberger's limit rule for the Revision Theory of Truth[5] we define

$$\mathrm{GS}^{\Pi}_{\lambda}(\ell)(s) := \begin{cases} \mathrm{GS}^{\Pi}_{\alpha}(\ell)(s) & \text{if } \forall\beta(\alpha \leq \beta < \lambda \;\rightarrow\; \mathrm{GS}^{\Pi}_{\alpha}(\ell)(s) = \mathrm{GS}^{\Pi}_{\beta}(\ell)(s)), \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Of course, if the procedure is monotonic below λ, this is the same as the usual definition $\mathrm{GS}^{\Pi}_{\lambda}(\ell) := \bigcup_{\alpha<\lambda} \mathrm{GS}^{\Pi}_{\alpha}(\ell)$.

[5] *Cf.* [GuBe93].

If η is a fixed point of the Gale-Stewart construction, *i.e.*, $\mathrm{GS}^{\Pi}_{\eta}(\ell) = \mathrm{GS}^{\Pi}_{\eta+1}(\ell)$, we call $\mathrm{GS}^{\Pi}_{\eta}(\ell)$ the **Gale-Stewart closure of ℓ relative to** Π and denote it by $\mathrm{GSC}^{\Pi}(\ell)$. Regardless of whether the Gale-Stewart procedure has a fixed point or not, we call the least α such that $\mathrm{GS}^{\Pi}_{\alpha}(\ell)(s)$ is defined the **index of** s.

Special properties of labellings. Let ℓ be a partial labelling, Π a total preference, and

$$\langle \mathrm{GS}^{\Pi}_{\alpha}(\ell)\,;\,\alpha \in \mathrm{Ord}\rangle$$

be the Gale-Stewart procedure starting from ℓ.

We say that ℓ has the **antichain property** if for each infinite sequence $\langle s_n\,;\,n \in \mathbb{N}\rangle$ with $s_{n+1} \supsetneqq s_n$ there is some n such that $s_n \in \mathrm{dom}(\ell)$. Note that this implies that $\mathrm{dom}(\ell)$ is a maximal antichain in $X^{<\omega}$ (and for the labellings ℓ discussed in this paper, it is actually equivalent).

We say that ℓ has the **Π-fixed point property** if the Gale-Stewart procedure starting from ℓ relative to Π has a fixed point, *i.e.*, for some α, $\mathrm{GS}^{\Pi}_{\alpha}(\ell) = \mathrm{GS}^{\Pi}_{\alpha+1}(\ell)$.

We say that ℓ has the **Π-monotonicity property** if for all $\alpha \in \mathrm{Ord}$, if $s \in \mathrm{dom}(\mathrm{GS}^{\Pi}_{\alpha}(\ell))$, then

$$\mathrm{GS}^{\Pi}_{\alpha}(\ell)(s) = \mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s).$$

Obviously, this implies that the Gale-Stewart procedure is monotonic in the usual sense[6], and we get a fixed point by the standard fixed point theorem for monotonic operators:

Lemma 2 *If ℓ has the Π-monotonicity property, it also has the Π-fixed point property.*

For some ℓ that has the Π-fixed point property, we say that it has the **Π-totality property**, if $\mathrm{GSC}^{\Pi}(\ell)$ is a total function. We say that it has the **Π-root property** if $\varnothing \in \mathrm{dom}(\mathrm{GSC}^{\Pi}(\ell))$.

Lemma 3 (Totality Lemma) *If ℓ has the antichain property and the Π-monotonicity property, then it has the Π-totality property.*

Proof. Suppose that s is not labelled by $\mathrm{GSC}^{\Pi}(\ell)$. By $(s-)$, all positions that have only labelled immediate successors are labelled in the Gale-Stewart

[6] *I.e.*, for all $\alpha < \beta$, we have $\mathrm{dom}(\mathrm{GS}^{\Pi}_{\alpha}(\ell)) \subseteq \mathrm{dom}(\mathrm{GS}^{\Pi}_{\beta}(\ell))$ and for $s \in \mathrm{dom}(\mathrm{GS}^{\Pi}_{\alpha}(\ell))$, we have $\mathrm{GS}^{\Pi}_{\alpha}(\ell)(s) = \mathrm{GS}^{\Pi}_{\beta}(\ell)(s)$.

closure. Consequently, we can define an infinite sequence $\langle s_n\,;\, n \in \omega\rangle$ with $s_n \subsetneq s_{n+1}$ starting from s such that for all $n \in \omega$, $s_n \notin \mathrm{dom}(\mathrm{GSC}^{\Pi}(\ell))$.

By the antichain property, there is some m such that $s_m \in \mathrm{dom}(\ell)$. But the monotonicity property of ℓ implies that $\mathrm{dom}(\ell) \subseteq \mathrm{dom}(\mathrm{GSC}^{\Pi}(\ell))$ which is a contradiction. q.e.d.

Special properties of strategies. Let $\mathcal{S}$ be some class of strategies, and ℓ be some labelling. We call a frame $\vec{\tau} = \langle \tau_j\,;\, j \in I\rangle$ an **$\mathcal{S}$-frame** if for all $j \in I$, we have $\tau_j \in \mathcal{S}$.

We call an i-strategy σ **ℓ-$\mathcal{S}$-consequent** if for every $\mathcal{S}$-frame $\vec{\tau}$ such that $\tau_i = \sigma$, we have that

$$\ell(\maltese\,\vec{\tau}\restriction n) = i$$

for all $n \in \omega$. We call an i-strategy $\mathcal{S}$-winning if for all $\mathcal{S}$-frames $\vec{\tau}$ and $x := \maltese\,\vec{\tau}$, we have $P(x) = i$.

We call a strategy σ a **$\langle \Pi, \ell\rangle$-strategy** if for all p with $\mu(p) = i$, if $\ell(p^\frown\langle\sigma(p)\rangle) = k$, then k is the $<^i_\Pi$-least element of

$$\{j \in I\,;\, \exists x\,(\ell(p^\frown\langle x\rangle) = j)\}.$$

We denote the set of $\langle \Pi, \ell\rangle$-strategies by $\mathcal{S}^{\Pi}_{\ell}$.

Heuristics: We again assume that $\ell(s) = j$ is interpreted as "at s it is determined that player j wins". Suppose $\mu(p) = i$; using his knowledge about ℓ, player i will be able to check his options at p by looking at the set

$$\{\ell(p^\frown\langle\xi\rangle)\,;\, \xi \in X\}.$$

Given the above interpretation of ℓ, the only rational choice for player i is to choose one of the successors that rank highest in his preference. A $\langle \Pi, \ell\rangle$-strategy is one in which the player is required to behave rational in this way.

Lemma 4 *Let ℓ be a labelling with the Π-totality property. Then there is a $\mathrm{GSC}^{\Pi}(\ell)$-$\mathcal{S}^{\Pi}_{\mathrm{GSC}^{\Pi}(\ell)}$-consequent strategy.*

Proof. We write $\ell^* := \mathrm{GSC}^{\Pi}(\ell)$. By the Π-totality property, we have $\ell^*(\varnothing) = i$ for some $i \in I$. We define the following i-strategy σ: if $\mu(p) = i$, play some ξ such that $p^\frown\langle\xi\rangle$ has the $<^i_\Pi$-least label among the immediate successors of p. The labelling ℓ^* is a total function and a fixed point of the Gale-Stewart procedure, so for each node s, $\ell^*(s)$ must also be the label of at least one immediate successor of s.

Fix any $\mathcal{S}^{\Pi}_{\ell^*}$-frame $\vec{\tau}$ such that $\tau_i = \sigma$, and let $x := \maltese\,\vec{\tau}$. We shall show that $\ell^*(x\restriction n) = i$ for every $n \in \omega$: Suppose it's not, and let $n + 1$ be the least counterexample (thus $\ell^*(x\restriction n) = i$ and $\ell^*(x\restriction n+1) \neq i$). By the above remark, we know that there is some immediate successor p of $x\restriction n$ such that $\ell^*(p) = i$.

Thus, by the choice of σ, we know that $\mu(x\restriction n) = j \neq i$ because player i would have chosen to play p instead of $x\restriction n+1$. But now by $(x\restriction n-)$, we have

$$i = \min_{<^j_\Pi}\{k \in I\,;\, \exists\xi \in X\,(\ell^*(s^\frown\langle\xi\rangle) = k)\},$$

yet $\tau_j(x\restriction n) = x(n)$ with $i <^j_\Pi \ell^*(x\restriction n+1)$. Thus τ_j is not a $\langle\Pi, \ell^*\rangle$-strategy. Contradiction. q.e.d.

4 Open and Closed Preference Games

4.1 Open Payoffs

Now let I be an arbitrary set of players, μ a move function, Π a total preference, and P an open payoff.

We define

$$S_i := \{p \in X^{<\omega}\,;\, \forall x\,(p \subseteq x \text{ implies } P(x) = i)\}$$

to be the set of positions at which player i has won, and set

$$\ell(p) := i :\Longleftrightarrow\ p \in S_i.$$

We call a labelling ℓ an **open labelling** if there is an open payoff P as above.

Lemma 5 *Every open labelling has the antichain property.*

Proof. Obvious. q.e.d.

Lemma 6 *Every open labelling has the* Π*-monotonicity property (for arbitrary preferences* Π*).*

Proof. Let $\langle \mathrm{GS}^\Pi_\ell(\alpha)\,;\, \alpha \in \mathrm{Ord}\rangle$ be the Gale-Stewart procedure derived from ℓ. We prove the claim by induction on the index of s. Suppose s is a counterexample of minimal index, so $\mathrm{GS}^\Pi_\alpha(\ell)(s) = i \neq j = \mathrm{GS}^\Pi_{\alpha+1}(\ell)(s)$. Take α to be minimal among those as well.

Case 1. The index is $\beta + 1$, and s was labelled by $(s+)$. Then there is some immediate successor t with strictly lower index such that $\mathrm{GS}^\Pi_\beta(\ell)(t) = i$. By minimality, $\mathrm{GS}^\Pi_\alpha(\ell)(t) = i$, so $(s+)$ is applied in the construction of $\mathrm{GS}^\Pi_{\alpha+1}(\ell)$, hence $\mathrm{GS}^\Pi_{\alpha+1}(\ell)(s) = i$. Contradiction.

Case 2. The index is $\beta+1$, and s was labelled by $(s-)$. Then all immediate successors t have strictly lower index, and thus by minimality, their labels are fixed. So $(s-)$ is applied in the construction of $\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)$, hence $\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s) = i$. Contradiction.

Case 3. Suppose that the index of s is 0. In that case, all successors (not only immediate) are i-labelled with index 0, so by induction $\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s) = i$. Contradiction. q.e.d.

Corollary 7 *Every open labelling has the* Π*-totality property (for arbitrary preferences* Π*).*

Proof. From Lemmas 5 and 6 *via* the Totality Lemma 3. q.e.d.

Theorem 8 *Let* I *be an arbitrary set of players,* μ *an arbitrary move function,* P *an open payoff,* ℓ *the open labelling derived from* P*,* ℓ^* *its Gale-Stewart closure, and* Π *a (total) preference. Then there is an* $i \in I$ *and a* $\mathcal{S}^{\Pi}_{\ell^*}$*-winning strategy* σ *for player* i*.*

Proof. By Corollary 7, $\mathrm{GSC}^{\Pi}(\ell)$ is total, so let $\mathrm{GSC}^{\Pi}(\ell)(\varnothing) = i$. Using (the proof of) Lemma 4, we get a $\mathrm{GSC}^{\Pi}(\ell)$-$\mathcal{S}^{\Pi}_{\ell^*}$-consequent strategy σ for player i. Let $\vec{\tau}$ be an arbitrary $\mathcal{S}^{\Pi}_{\ell^*}$-frame with $\tau_i = \sigma$, and let $x := \maltese\vec{\tau}$. We have that for every $n \in \omega$, $\mathrm{GSC}^{\Pi}(\ell)(x{\restriction}n) = i$.

Lemma 5 gives us some m such that $x{\restriction}m \in \mathrm{dom}(\ell)$, and Lemma 6 yields that $\ell(x{\restriction}m) = \mathrm{GSC}^{\Pi}(\ell)(x{\restriction}m) = i$, *i.e.*, $x \in S_i$, and thus $P(x) = i$. q.e.d.

In terms of determinacy axioms, the analysis of this section yields a many-player determinacy statement:

> For open payoffs and arbitrary sets of players, if we fix a (total) preference Π, then there is a Π-winning strategy for one of the players.

Let us rephrase this in terms of equilibria: If we restrict the class of relevant strategies to the strategies in $\mathcal{S}^{\Pi}_{\ell^*}$, then Theorem 8 gives us a large set of very strong equilibria: Each frame $\vec{\tau}$ such that $\tau_i = \sigma$ is an equilibrium since none of the players can increase his payoff while playing a strategy in $\mathcal{S}^{\Pi}_{\ell^*}$; the players $j \neq i$ necessarily lose against σ.

4.2 Closed Payoffs

If I is finite and P is closed, then all sets P_i are also open, and Theorem 8 can be applied. If I is infinite, it could be that the payoff sets are closed but not clopen. However, a slight modification of the above argument (using closed labellings instead of open labellings) yields the same result:

Let I be an arbitrary set of players, μ a move function, Π a total preference, and P a closed payoff.

We define

$$C_i := \{p \in X^{<\omega} \,;\, \forall x\,(p \subseteq x \text{ implies } P(x) \neq i)\}$$

to be the set of positions where player i has irrevocably lost, and define the labelling

$$\ell(s) = i :\Longleftrightarrow \forall j \in I\ (j \neq i \rightarrow s \in C_j).$$

We call such a labelling an **closed labelling**.

Now it's easy to see closed labellings also have the antichain and the monotonicity property. So by the Totality Lemma 3, $\ell^* = \mathrm{GSC}^{\Pi}(\ell)$ is a total function. Lemma 4 gives an $\mathcal{S}^{\Pi}_{\ell^*}$-consequent strategy, which then by definition of ℓ yields a $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy:

Theorem 9 *Let I be an arbitrary set of players, μ an arbitrary move function, P a closed payoff, ℓ the derived closed labelling, ℓ^* its Gale-Stewart closure, and Π a (total) preference. Then there is an $i \in I$ and a $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy σ for player i.*

4.3 Exceptional Least Evil

We shall now briefly discuss a special situation which occurs rather frequently and can be solved with (a variant of) the methods of this section.

If Π is a preference, we say that $i_0 \in I$ is a **least evil** in Π, if for all $i \neq i_0$, we have

$$i <^i_\Pi i_0 <^i_\Pi j$$

for all $j \notin \{i, i_0\}$.

We call a pair $\langle P, \Pi \rangle$ an **exceptional least evil situation** if i_0 is a least evil, and P is a payoff such that P_i is open for all $i \neq i_0$. Clearly, P_{i_0} is then closed.

For $i \neq i_0$, we define

$$S_i := \{p \in X^{<\omega} \,;\, \forall x\,(p \subseteq x \text{ implies } P(x) = i)\}$$

and

$$\ell(p) := i :\Longleftrightarrow p \in S_i$$

as before, and call these labellings **exceptional least evil labellings**.

We get analogues of Lemmas 5 and 6:

Lemma 10 *If $\langle P, \Pi\rangle$ is an exceptional least evil situation and ℓ the derived labelling, then for every $x \in X^\omega$, we have*

- *either there is some n such that $x{\upharpoonright}n \in \mathrm{dom}(\ell)$,*
- *or $P(x) = i_0$.*

Lemma 11 *If $\langle P, \Pi\rangle$ is an exceptional least evil situation, then its derived labelling has the Π-monotonicity property.*

Consequently, exceptional least evil labellings have the Π-fixed point property. However, the Gale-Stewart closure needn't be total this time as they do not have the antichain property. We define

$$\ell^*(s) = \begin{cases} \mathrm{GSC}^\Pi_\ell(s) \text{ if } s \in \mathrm{dom}(\mathrm{GSC}^\Pi(\ell)), \text{ and} \\ \quad i_0 \qquad \text{otherwise.} \end{cases}$$

If $s \notin \mathrm{dom}(\mathrm{GSC}^\Pi(\ell))$, we say that the index of s is 0.

Lemma 12 *If $\langle P, \Pi\rangle$ is an exceptional least evil situation, ℓ the derived labelling and ℓ^* defined as above, then there is an ℓ^*-$\mathcal{S}^\Pi_{\ell^*}$-consequent strategy.*

Proof. Suppose that $\ell^*(\varnothing) = i$. As in the proof of Lemma 4, we define a strategy σ for player i: if $\mu(p) = i$, let $i^p_{\min} := \min_{<^i_\Pi}\{\ell^*(p^\frown\langle\xi\rangle)\,;\, \xi \in X\}$. Look at the set $\Xi^p := \{\xi \in X\,;\, \ell^*(p^\frown\langle\xi\rangle) = i^p_{\min}\}$. Pick ξ_0 from this set such that the index of $p^\frown\langle\xi_0\rangle$ is minimal. Note that if the index of p is not zero, there is some $\xi_0 \in \Xi^p$ such that the index of $p^\frown\langle\xi_0\rangle$ is strictly lower than the index of p.

We claim that this strategy is ℓ^*-$\mathcal{S}^\Pi_{\ell^*}$-consequent.

If $\ell^*(\varnothing) = i \neq i_0$, this proof is essentially the same as the proof of Lemma 4. Suppose that $\ell^*(\varnothing) = i_0$. Observe that the only way s can be labelled i_0 is if s has an immediate successor that is labelled i_0 (otherwise, all immediate successors have are labelled by $\mathrm{GSC}^\Pi(\ell)$, and hence also s itself). Let $\vec{\tau}$ be an arbitrary $\mathcal{S}^\Pi_{\ell^*}$-frame with $\tau_{i_0} = \sigma$ and $x := \maltese\vec{\tau}$. Take n such that $\ell^*(x{\upharpoonright}n) = i_0$ and $\ell^*(x{\upharpoonright}n+1) = j \neq i_0$. Obviously, $\mu(x{\upharpoonright}n) = k \neq i_0$. If $j = k$, then $\ell^*(x{\upharpoonright}n) = \mathrm{GSC}^\Pi_\ell(x{\upharpoonright}n) = j$ by $(x{\upharpoonright}n+)$. Consequently, $j \neq k$. But now we can use that i_0 is the least evil, so in particular, $k <^k_\Pi i_0 <^k_\Pi j$. But

$$\ell^*(x{\upharpoonright}n)^\frown\tau_k(x{\upharpoonright}n)) = j,$$

so τ_k is not a $\langle \Pi, \ell^*\rangle$-strategy. Contradiction. q.e.d.

Theorem 13 *Let I be an arbitrary set of players, μ an arbitrary move func-*

tion, $\langle P, \Pi \rangle$ an exceptional least evil situation, ℓ the derived labelling, and ℓ^* as defined above. There is an $i \in I$ and an $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy σ for player i.

Proof. The strategy σ as defined in the proof of Lemma 12 will be the $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy. Remember that σ always plays into positions of strictly lower index (unless the index is 0). Let $\vec{\tau}$ be a $\mathcal{S}^{\Pi}_{\ell^*}$-frame with $\tau_i = \sigma$ and $x := \maltese\vec{\tau}$. Then by Lemma 12, $\ell^*(x{\upharpoonright}n) = i$ for all $n \in \omega$.

We are left to show that $P(x) = i$. By Lemma 10, we know that this is the case if $i = i_0$. So let $i \neq i_0$.

Note that if $\mu(x{\upharpoonright}n) = j \neq i$, then it was either already i-labelled in ℓ or it was i-labelled by $(x{\upharpoonright}n-)$, so all immediate successors have strictly lower index than $x{\upharpoonright}n$. This together with the choice of σ says that the sequence of indices of $x{\upharpoonright}n$ is a decreasing sequence of ordinals. Hence it must eventually reach 0, which means that $x{\upharpoonright}n \in S_i$ for some $n \in \omega$. q.e.d.

The fact that the index function needs to be used in order to define a winning strategy in this case is no coincidence: as opposed to games in which all payoffs are open or all payoffs are closed, the games with an exceptional least evil contain for example all so-called **combinatorial games** (games with counters on graphs; the last player to move the counter wins, if the counter is moved infinitely many times, it's a draw). That labelling functions for combinatorial games need transfinite ordinal values to describe winning strategies has been discussed in [FrRa01].

5 Mixed Labels

The analysis of Section 4, in particular Section 4.3 covers quite a lot: The usual open and closed two-player games (*i.e.*, P_0 open and P_1 closed, or *vice versa*) are an instance of an exceptional least evil (the closed player is the exception); also all games with open payoffs and a draw are instances of an exceptional least evil (as mentioned, the combinatorial games on graphs are examples of these).

But there are other situations that the analysis of Section 4.3 cannot deal with, for example a closed player who is not a least evil, or two closed and one open players.

For these situation we need to mix open and closed labellings, and give up monotonicity.

Let I be a set of players, μ a move function and P be a payoff such that P_i

is either open or closed for all $i \in I$. We shall call such a payoff function a **mixed payoff**. Let $I_{\mathrm{open}} \cup I_{\mathrm{closed}}$ be a disjoint partition of I such that for all $i \in I_{\mathrm{open}}$, P_i is open and for all $i \in I_{\mathrm{closed}}$, P_i is closed. Assume in addition that I_{closed} is finite (this is used in the proof of Lemma 14).

We are now joining the ideas of open and closed labellings. For each $i \in I_{\mathrm{closed}}$, let

$$C_i := \{p \in X^{<\omega}\,;\, \forall x\,(p \subseteq x \text{ implies } P(x) \neq i)\}$$

be the set of positions where player i has irrevocably lost, and for each $i \in I_{\mathrm{open}}$, let

$$S_i := \{p \in X^{<\omega}\,;\, \forall x\,(p \subseteq x \text{ implies } P(x) = i)\}$$

be the set of positions at which player i has won.

We define the following **mixed labelling**:

$$\ell(s) := \begin{cases} i \in I_{\mathrm{open}} & \text{if } s \in S_i, \\ j \in I_{\mathrm{closed}} & \text{if } s \notin \bigcup\{S_i\,;\, i \in I_{\mathrm{open}}\} \cup \bigcup\{C_k\,;\, k \in I_{\mathrm{closed}}, k \neq j\}. \end{cases}$$

Labellings derived from a mixed payoff function in this way are called **mixed labellings**.

Lemma 14 *Every mixed labelling has the antichain property.*

Proof. Let $x \in X^\omega$, and let $P(x) = i$. If $i \in I_{\mathrm{open}}$, then there is some n such that $x{\restriction}n \in S_i$.

If $i \in I_{\mathrm{closed}}$, then for each $j \in I_{\mathrm{closed}}$ such that $j \neq i$ there is some natural number n_j with $x{\restriction}n_j \in C_j$. Let $n := \max\{n_j\,;\, j \in I_{\mathrm{closed}}, j \neq i\}$ (note that this exists because I_{closed} is finite). Then $\ell(x{\restriction}n) = i$. q.e.d.

If now $\langle \mathrm{GS}^\Pi_\alpha(\ell)\,;\, \alpha \in \mathrm{Ord}\rangle$ is the Gale-Stewart procedure starting from a mixed labelling ℓ, it is not necessarily monotonic anymore:

For example, if $I = X = \{0,1\}$, $\mu(\varnothing) = 0$, $P_0 := \{x\,;\, x(0) = 0\}$, then $\ell(\varnothing) = 1$ and $\ell(\langle 0\rangle) = 0$. Then $(\varnothing+)$ gives $\mathrm{GS}^\Pi_1(\ell)(\varnothing) = 0 \neq 1 = \mathrm{GS}^\Pi_0(\ell)(\varnothing)$. We call such a situation, *i.e.*, a pair $\langle s, \alpha\rangle$ such that

$$\mathrm{GS}^\Pi_\alpha(\ell)(s) = i \neq j = \mathrm{GS}^\Pi_{\alpha+1}(\ell)(s)$$

an **overwriting instance** (for ℓ).

Lemma 15 *If ℓ is a mixed labelling and $\langle s, \alpha\rangle$ is an overwriting instance with*

$$\mathrm{GS}^\Pi_\alpha(\ell)(s) = i \neq j = \mathrm{GS}^\Pi_{\alpha+1}(\ell)(s),$$

then $i \in I_{\mathrm{closed}}$, $j \in I_{\mathrm{open}}$, and $\mathrm{GS}^\Pi_\alpha(\ell)(s) = \ell(s)$.

Proof. We prove this by induction on the index of s. Suppose s is a counterexample of minimal index. Take α to be minimal among those as well.

Case 1., **Case 2.** and **Case 3.** from the proof of Lemma 6, show that the index of s must be 0, and that i cannot be an open label, so $i \in I_{\text{closed}}$. By definition of ℓ, we know that

$$I_{\text{closed}} \cap \{\ell(t)\,;\, t \supseteq s\} = \{i\}.$$

So by induction, $\mathrm{GS}^{\Pi}_{\alpha}(\ell)(t) \in I_{\text{open}} \cup \{i\}$ for all successors $t \supsetneq s$. Consequently, if $\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s) = j \neq i$, then $j \in I_{\text{open}}$. But now the minimal choice of α gives $\ell(s) = \mathrm{GS}^{\Pi}_{\alpha}(\ell)(s)$ by induction, so s was no counterexample. q.e.d.

Corollary 16 *If ℓ is a mixed labelling and $s \in X^{<\omega}$, then there is at most one α such that $\mathrm{GS}^{\Pi}_{\alpha}(\ell)(s) \neq \mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s)$.*

Proof. By Lemma 15, if $\mathrm{GS}^{\Pi}_{\alpha}(\ell)(s) \neq \mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s)$, then $\mathrm{GS}^{\Pi}_{\alpha}(\ell)(s) \in I_{\text{closed}}$ and $\mathrm{GS}^{\Pi}_{\alpha+1}(\ell)(s) \in I_{\text{open}}$. But now again by Lemma 15 this means that for no $\beta > \alpha$, we can have $\mathrm{GS}^{\Pi}_{\beta}(\ell)(s) \neq \mathrm{GS}^{\Pi}_{\beta+1}(\ell)(s)$. q.e.d.

Lemma 17 *Every mixed labelling has the Π-fixed point property (for arbitrary Π).*

Proof. By Corollary 16, there is only a set of overwriting instances (in fact, the cardinality of the set is at most $\kappa := \mathrm{Card}(X^{<\omega})$. Let

$$\Sigma := \sup\{\alpha + 1\,;\, \langle s, \alpha\rangle \text{ is an overwriting instance for some } s \in X^{<\omega}\}.$$

Then after Σ, the Gale-Stewart procedure is fully monotonic, and hence has a fixed point by the usual fixed point theorem. q.e.d.

Lemma 18 *Every mixed labelling has the Π-totality property (for arbitrary Π).*

Proof. We know by Lemma 14 that ℓ has the antichain property. Note that the proof of the Totality Lemma 3 doesn't really need the full Π-monotonicity property but only

$$s \in \mathrm{dom}(\ell) \;\to\; s \in \mathrm{dom}(\mathrm{GSC}^{\Pi}(\ell)) \tag{$\dagger$}$$

which is a consequence of Corollary 16. q.e.d.

Theorem 19 *Let $I = I_{\text{open}} \cup I_{\text{closed}}$ be an arbitrary set of players where I_{closed} is finite, μ an arbitrary move function, P a payoff such that P_i is open for $i \in I_{\text{open}}$ and P_i is closed for $i \in I_{\text{closed}}$, ℓ the derived mixed labelling, and ℓ^* its Gale-Stewart closure (which exists by Lemma 17). Then there is an $i \in I$ and an $\mathcal{S}^{\Pi}_{\ell^*}$-winning strategy σ for player i.*

Proof. By Lemma 18, ℓ^* is total, so by (the proof of) Lemma 4, we have a ℓ^*-$\mathcal{S}^{\Pi}_{\ell^*}$-consequent strategy σ for player i such that $\ell^*(\varnothing) = i$.

To prove the theorem, it is enough to show the following:

We show that if $x \in X^\omega$ such that for all n, $\ell^*(x{\restriction}n) = i$, then $P(x) = i$.

Case 1. Let $i \in I_{\text{open}}$. By Lemma 14, we get some n such that $\ell(x{\restriction}n) = i$, so $x \in S_i$, so $P(x) = i$.

Case 2. Let $i \in I_{\text{closed}}$. If $x \in \bigcup\{P_j\,;\, j \in I_{\text{open}}\}$, then there is some n such that $\ell(x{\restriction}n) \in I_{\text{open}}$. But then by Lemma 15, $\ell^*(x{\restriction}n) = \ell(x{\restriction}n) \in I_{\text{open}}$. Contradiction.

So, we have that $x \in \bigcup\{P_j\,;\, j \in I_{\text{closed}}\}$ and by Lemma 14 that for some n, $\ell(x{\restriction}n) = i$. By definition of the closed labels in ℓ, this means that $x \notin \bigcup\{P_j\,;\, j \in I_{\text{closed}}, j \neq i\}$, so $P(x) = i$. q.e.d.

6 Temporal Game Logic

In his [vB03, § 8.2], van Benthem discusses the Gale-Stewart theorem in terms of a branching time logic with additional game operators. We think of X^ω as the set of branches in a model for a Prior-style branching time logic with the usual operators $\mathbf{G}$ ("for all future times") and $\mathbf{H}$ ("for all past times"), and the derived operator $\mathbf{A}\varphi \equiv \mathbf{G}\varphi \wedge \mathbf{H}\varphi$.[7] Motivated by looking at finite games [vB03, § 5.3], van Benthem adds game modalities $\mathbf{W}_i$ for each of the players $i \in I$ with the intended meaning "player i can force".

If $x \in X^\omega$ and $s \in X^{<\omega}$, then the semantics for $\mathbf{W}_i$ is $\langle x, s\rangle \models \mathbf{W}_i\,\varphi$ if and only if there is a strategy for player i such that every run of the game $y \supseteq s$ consistent with y (beyond s), we get $\langle y, t\rangle \models \varphi$ for all $s \subseteq t \subseteq y$.

Considering a two-player situation with an open payoff P_0 and a closed payoff P_1 where membership in P_0 is described by φ, the key step of the Gale-Stewart argument transforms into the modal formula

$$\mathbf{W}_0\,\varphi \vee \mathbf{W}_1\,\mathbf{A}\,\neg\mathbf{W}_0\,\varphi:$$

[7] For a thorough account of Prior's tense logics, *cf.* [Mü02]; for the original development of the modern technicalities, *cf.* [Th70] and [BAMaPn83]. Branching Time Logic has been connected to games and Backward Induction phenomena by Bonanno [Bo01] who proves in the framework of temporal logic that each *internally consistent* prediction or recommendation for the players must come from a backwards induction solution.

either player 0 can force the outcome into P_0 or player 1 can make sure that player 0 can never force the outcome into P_0.[8]

This is exactly the claim of Lemma 12, and if we are in the case of the second alternative ($\mathbf{W}_1\,\mathbf{A}\,\neg\mathbf{W}_0\,\varphi$), we use the fact that P_1 is closed to show that this is enough to prove that the outcome is in P_1.

Of course, when we move to more general formulas φ of temporal logic, the Weak Determinacy formula might not be enough to prove determinacy anymore, and also the provability of the Weak Determinacy formula might depend on the system we're working in.

In the cases of open payoffs, closed payoffs and mixed payoffs with the non-monotone analysis of Section 5, the explication in a temporal game logic term becomes even simpler and is identical with the statement of determinacy:

$$\bigvee_{i\in I} \mathbf{W}_i\,\varphi_i$$

where φ_i is a formula describing membership in P_i.[9] The case of the Exceptional Least Evil reverts to the Gale-Stewart situation and gives

$$\bigvee_{i\neq i_0} \mathbf{W}_i\,\varphi_i \;\vee\; \mathbf{W}_{i_0}\,\mathbf{A} \bigwedge_{i\neq i_0} \neg\mathbf{W}_i\,\varphi_i.$$

7 Games with three players

We now give some examples of games solved by the solution concepts by open labellings and mixed labellings given in Sections 4 and 5.

For this let's look a bit more closely at three-player games with preferences. Up to renaming of the players, there are only two different three-player games with total preferences, we call them Evening at a Married Ex and Love Triangle (depicted in Figure 3). In our pictures of the preferences, an arrow from i to j means "i prefers j over k" (where $\{i, j, k\} = \{0, 1, 2\}$).

Of course, the solution concept presented in Theorem 19 gives us optimal strategies for some player in these games if the payoffs are open and/or closed.

[8] This is called **Weak Determinacy** by van Benthem [vB03].

[9] If I is infinite, we either need to interpret the disjunction $\bigvee$ as a quantifier with the obvious semantics or move to an infinitary logic.

Fig. 3. Evening at a Married Ex and Love Triangle

For three players, if $\Pi(i)$ is not a wellordering, it is necessarily of the form $i <^i_\Pi j, k$ where j and k are in no particular ordering. So if Π is not a total preference, then either one, two, or all three of the relations $<^i_\Pi$ are of this form. There are (up to renaming) three cases with two wellorderings, and one case each with one and zero wellorderings. They are depicted in Figures 4 and 5.

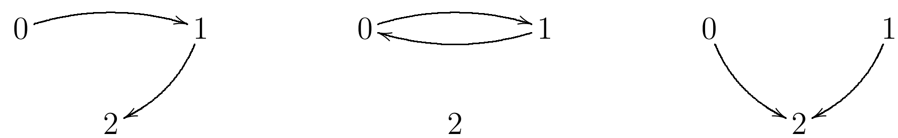

Fig. 4. Beatrice's Revenge, Evening at a Couple, and Least Evil

In general, there are no solution concepts for partial preferences. Of course, Hobbes is just the case of full non-cooperation and as mentioned in Proposition 1, there is no solution for this situation. A notable exception is Least Evil: In the case where the least evil itself doesn't move (*i.e.*, $M_2 = \varnothing$), the "Exceptional Least Evil" analysis of Section 4.3 is a solution.

Let us briefly describe the different situations by examples:

Love Triangle. This situation is truly pseudo-Shakespearean: Beatrice (0) is the fiancée of the poor but handsome Captain Antonio (1). She recently started to question her fiancé's character. Faking a family emergency, she pretends to go to the countryside while dressing as a rich wine merchant to check on Antonio's behaviour. Alas, her suspicions prove to be correct: as soon as she apparently leaves town, Antonio starts a liaison with the beautiful Cressida (2). Beatrice in the rôle the wine merchant is intent on confronting Antonio with his deeds and invites Antonio and Cressida to play an infinite three-player game. Now the treacherous heart of Cressida abandons poor Antonio for the rich wine merchant ($0 <^2_\Pi 1$) while the good-hearted Beatrice is full of pity for Antonio seeing him so deceived by Cressida ($1 <^0_\Pi 2$). Antonio, of course, doesn't evaluate the situation correctly: he doesn't recognize Beatrice, and although he realizes that Cressida has abandoned him, he tries to win her back by preferring her in the game ($2 <^1_\Pi 0$).

Beatrice's Revenge. In the situation of Love Triangle, Beatrice realizes that Antonio is an idiot. Even though Cressida shamelessly flirts with the rich wine merchant, Antonio gazes in benumbed infatuation at her. Suddenly, Beatrice realizes that this buffoon of a Captain doesn't deserve her compassion, and she gives up favouring Antonio over Cressida.

Evening at a Couple. John (2) visits his friends, the married couple Jeff (1) and Jill (0) at home. They decide to kill some time by playing an infinite three-player game. Although Jeff and Jill are good sports and try to play the game without prejudice, subconsciously, they prefer their marital partner over John ($0 <^1_\Pi 2$, $1 <^0_\Pi 2$). John knows that pretty well, and realizes that it makes no difference whether Jeff or Jill wins.

Evening at a Married Ex. We are in the situation of Evening at a Couple but with an added twist: Jill was the girlfriend of John for a long time while they were in graduate school, and (not unbeknownst to Jill) John is still in love with her ($0 <^2_\Pi 1$).

Least Evil. We are in a Mathematics Department with Professors Smith (0), Johnson (1) and Williams (2) eligible to become the new Chair.[10] Smith and Miller are ambitious administrators and know that the only way to become Dean of the Faculty of Sciences is to become Department Chair. Of course, they also know that if the other one becomes Chair, he will probably use his chance to become Dean, and the position of the Dean will be blocked for at least five if not ten years. Williams is not ambitious at all – both Smith and Miller realize that if Williams were to become Chair, he would never aim at the office of Dean ($2 <^0_\Pi 1$, $2 <^1_\Pi 0$). As laid out in the bylaws of the Mathematics Department, Smith, Johnson and Williams have to engage in an infinite three-player game to determine the next Chair.

Of course, Least Evil is a very common situation. As mentioned, if $M_2 = \varnothing$, we end up with a two-player game with draw.

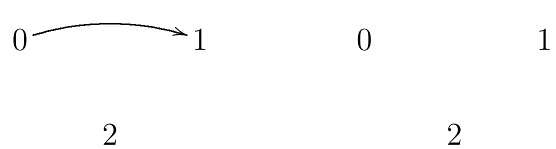

Fig. 5. Sidekick and Hobbes

Sidekick. Luigi (1) and Paolo (2), two mafia bosses meet to deal with each other, and Luigi brought his faithful follower Giacomo (0). They play an infinite game, and whoever wins the game will become Overlord of Crime. Each of them is given a chance of winning that title, but what about the consequences of losing? Clearly, if Paolo wins, he will most probably kill both Luigi and

[10] To stifle any discussions about the choice of these surnames: According to at least one poll, the most common surnames in the United States in the year 2000 were (in that order): Smith, Johnson, Williams, Jones, Brown, Davis, Miller, Wilson, Moore, Taylor.

Giacomo to thwart any opposition forming around them; similarly, he can be sure to be killed for the same reason by either of the others if they win. Luigi has humiliated Giacomo over many years, so Luigi can't be sure of his future if Giacomo wins. The only one who has a preference is Giacomo: if Luigi wins, he will stay alive –and continue to be humiliated by Luigi–; if Paolo wins, he'll die ($1 <^0_\Pi 2$).

Hobbes. In this three-player game there are no preferences except that all players wish to win, and if they don't, they don't care who does. This type of game is played in the Hobbesian *Natural Condition of Mankind*:

> *If any two men desire the same thing, which nevertheless they cannot both enjoy, they become enemies; and in the way to their end (which is principally their own conservation, and sometimes their delectation only) endeavour to destroy or subdue one another. ... Men have no pleasure (but on the contrary a great deal of grief) in keeping company where there is no power able to overawe them all. ... It is manifest that during the time men live without a common power to keep them all in awe, they are in that condition which is called war; and such a war as is of every man against every man* (Hobbes, Leviathan XIII).

As pointed out, this situation is the fully non-cooperative situation described in Proposition 1, and thus Π-preference determinacy cannot offer a solution concept for these games.

8 Conclusion

We gave solution concepts for infinite multi-player games with the strongest common knowledge assumptions and for very simple payoffs. Of course, there are many directions to discuss variations of these themes:

More complicated payoffs. As can be made explicit, $\mathbf{\Delta}^0_2$ payoffs still allow combinatorial labellings, thus it is possible to give a version of backwards induction. That suggests that the arguments from this paper could possibly be extended to $\mathbf{\Delta}^0_2$ payoffs.

Furthermore, almost all determinacy proofs in set theory go back to open determinacy. The key notion here is the tree representation of more complicated sets by κ-homogeneous trees (*cf.* [Ka94, § 32]). Winning a game on the closed payoff set given by the tree can be transformed into winning the original game. Can we give homogeneous tree version for many-players games with more complicated payoffs (*e.g.*, assuming that the players have agreed to be playing according to one fixed homogeneous tree before the game)?

Other knowledge assumptions. Consider the following version of Love Triangle where Antonio is even less observant: He doesn't recognize Beatrice, but moreover he also doesn't realize that Cressida has abandoned him, so in evaluating the game tree, he will evaluate the positions as if $1 <^2_{\mathrm{II}} 0$ instead of (the correct) $0 <^2_{\mathrm{II}} 1$. Assuming that Beatrice and Cressida know about this error in judgement of Antonio's, can we give a solution of the game? What if we allow Antonio to change his mind (and thus the labelling of the game tree in his mind) as soon as he sees that Cressida has betrayed him? [11] This necessarily leads to dynamic models of these games similar to dynamic models for epistemic solutions of finite games (*cf.* [vB01]).

Proof-theoretic analysis & more non-monotonicity. The existence of winning strategies has been analyzed proof-theoretically in terms of weak systems of second-order arithmetic. [12]

Deleting and overwriting of labels for games as in the analysis of mixed labellings occurs in some proof-theoretic analyses of subsystems of second-order number theory [Mö03], and in full generality might be connected to games corresponding to the proof-theoretic strength of Gupta-Belnap revision theory. [13] Thus extending the ideas of Section 5 to labellings with the antichain property but with truly non-monotonic Gale-Stewart procedures that need to be analysed as revision sequences is an interesting approach to game-theoretic analyses of Revision Theory.

References

[Au95] Robert J. **Aumann**, Backward Induction and Common Knowledge of Rationality, **Games and Economic Behavior** 8 (1995), p. 6–19

[BAMaPn83] Mordechai **Ben-Ari**, Zohar **Manna**, Amir **Pnueli**, The temporal logic of branching time, **Acta Informatica** 20 (1983), p. 207–226

[Bi88] Ken **Binmore**, Modelling Rational Players. Part II, **Economics and Philosophy** 4 (1988), p. 9-55

[Bo01] Giacomo **Bonanno**, Branching Time Logic, Perfect Information Games and Backward Induction, **Games and Economic Behavior** 36 (2001), p. 57–73

[11] For finite games, this has been considered by Kreps *et al.* in [Kr+82].
[12] *Cf.* [Ta90], [Ta91], [Mö03], and in particular the textbook [Si99].
[13] *Cf.* [KüLöMöWe∞] for a discussion of the strength of Revision Theory and of the connections between Revision Theory and games.

[Br00] Rodica **Brânzei**, On the Determinateness of n-person games with information energy, **Revue Roumaine de Mathématiques Pures et Appliquées** 45 (2000), p. 67–76

[dB∞] Boudewijn **de Bruin**, An Application of Epistemic Logic to Some Questions in Game Theory, *typoscript*

[FrRa01] Aviezri S. **Fraenkel**, Ofer **Rahat**, Infinite cyclic impartial games, Theoretical Computer Science 252 (2001), p. 13–22

[Ga53] David **Gale**, A theory of n-person games with perfect information, **Proceedings of the National Academy of Sciences U.S.A.** 39 (1953), p. 496–501

[GaSt53] David **Gale**, Frank M. **Stewart**, Infinite Games with Perfect Information, *in:* Harold W. Kuhn, Albert W. Tucker (*eds.*), Contributions to the Theory of Games II, Princeton 1953 [Annals of Mathematical Studies 28], p. 245–266

[GuBe93] Anil **Gupta**, Nuel **Belnap**, The Revision Theory of Truth, Cambridge MA 1993

[Ka94] Akihiro **Kanamori**, The Higher Infinite, Large Cardinals in Set Theory from Their Beginnings, Berlin 1994 [Perspectives in Mathematical Logic]

[Kr+82] David M. **Kreps**, Paul **Milgrom**, John **Roberts**, Robert **Wilson**, Rational cooperation in the finitely repeated prisoners' dilemma, **Journal of Economic Theory** 27 (1982), p. 245–252

[KüLöMöWe∞] Kai-Uwe **Kühnberger**, Benedikt **Löwe**, Michael **Möllerfeld**, Philip D. **Welch**, Comparing inductive and circular definitions: parameters, complexities and games, *in preparation*

[Mö03] Michael **Möllerfeld**, Systems of Inductive Definitions, Ph.D. Thesis, Westfälische Wilhelms-Universität Münster, 2003

[Mü02] Thomas **Müller**, Arthur Priors Zeitlogik, Mentis Verlag, Paderborn 2002

[Si99] Stephen G. **Simpson**, Subsystems of second order arithmetic, Berlin 1999 [Perspectives in Mathematical Logic]

[St96] Robert **Stalnaker**, Knowledge, Belief, and Counterfactual Reasoning in Games, **Economics and Philosophy** 12 (1996), p. 133–163

[Ta90] Kazuyuki **Tanaka**, Weak axioms of determinacy and subsystems of analysis I: Δ^0_2 games, **Zeitschrift für Mathematische Logik und Grundlagen der Mathematik** 36 (1990), p. 481–491

[Ta91] Kazuyuki **Tanaka**, Weak axioms of determinacy and subsystems of analysis II: Σ^0_2 games **Annals of Pure and Applied Logic** 52 (1991), p. 181–193

[Th70] Richmond H. **Thomason**, Indeterminist time and truth-value gaps, **Theoria** 36 (1970), p. 264–281

[vB01] Johan **van Benthem**, Games in dynamic-epistemic logic, **Bulletin of Economic Research** 53 (2001), p. 219–248

[vB03] Johan **van Benthem**, What Logic Games are Trying to Tell Us, **ILLC Publications** PP-2003-05, *preprint*